

# Amazon-Web-Services

## Exam Questions SAA-C03

AWS Certified Solutions Architect - Associate (SAA-C03)



### NEW QUESTION 1

- (Topic 1)

A company performs monthly maintenance on its AWS infrastructure. During these maintenance activities, the company needs to rotate the credentials for its Amazon ROS for MySQL databases across multiple AWS Regions  
Which solution will meet these requirements with the LEAST operational overhead?

- A. Store the credentials as secrets in AWS Secrets Manager
- B. Use multi-Region secret replication for the required Regions Configure Secrets Manager to rotate the secrets on a schedule
- C. Store the credentials as secrets in AWS Systems Manager by creating a secure string parameter Use multi-Region secret replication for the required Regions Configure Systems Manager to rotate the secrets on a schedule
- D. Store the credentials in an Amazon S3 bucket that has server-side encryption (SSE) enabled Use Amazon EventBridge (Amazon CloudWatch Events) to invoke an AWS Lambda function to rotate the credentials
- E. Encrypt the credentials as secrets by using AWS Key Management Service (AWS KMS) multi-Region customer managed keys Store the secrets in an Amazon DynamoDB global table Use an AWS Lambda function to retrieve the secrets from DynamoDB Use the RDS API to rotate the secrets.

**Answer:** A

#### Explanation:

<https://aws.amazon.com/blogs/security/how-to-replicate-secrets-aws-secrets-manager-multiple-regions/>

### NEW QUESTION 2

- (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=AL!4422!3!536452716950!p!!g!!aws%20s3%20pricing&ef\\_id=Cj0KCQjwnbmaBhD- ARIsAGTPcfVHUZN5\\_BMrzI5zBcaC8KnqpnNZvbZzqPkH6k7q4JcYO5KFLx0YYgaAm6nEALw\\_wcB:s&s\\_kwid=AL!4422!3!536452716950!p!!g!!aws%20s3%20pricing](https://aws.amazon.com/s3/storage-classes/?trk=66264cd8-3b73-416c-9693-ea7cf4fe846a&sc_channel=ps&s_kwid=AL!4422!3!536452716950!p!!g!!aws%20s3%20pricing&ef_id=Cj0KCQjwnbmaBhD- ARIsAGTPcfVHUZN5_BMrzI5zBcaC8KnqpnNZvbZzqPkH6k7q4JcYO5KFLx0YYgaAm6nEALw_wcB:s&s_kwid=AL!4422!3!536452716950!p!!g!!aws%20s3%20pricing)

### NEW QUESTION 3

- (Topic 1)

A company observes an increase in Amazon EC2 costs in its most recent bill The billing team notices unwanted vertical scaling of instance types for a couple of EC2 instances A solutions architect needs to create a graph comparing the last 2 months of EC2 costs and perform an in-depth analysis to identify the root cause of the vertical scaling  
How should the solutions architect generate the information with the LEAST operational overhead?

- A. Use AWS Budgets to create a budget report and compare EC2 costs based on instance types
- B. Use Cost Explorer's granular filtering feature to perform an in-depth analysis of EC2 costs based on instance types
- C. Use graphs from the AWS Billing and Cost Management dashboard to compare EC2 costs based on instance types for the last 2 months
- D. Use AWS Cost and Usage Reports to create a report and send it to an Amazon S3 bucket Use Amazon QuickSight with Amazon S3 as a source to generate an interactive graph based on instance types.

**Answer:** B

#### Explanation:

AWS Cost Explorer is a tool that enables you to view and analyze your costs and usage. You can explore your usage and costs using the main graph, the Cost Explorer cost and usage reports, or the Cost Explorer RI reports. You can view data for up to the last 12 months, forecast how much you're likely to spend for the next 12 months, and get recommendations for what Reserved Instances to purchase. You can use Cost Explorer to identify areas that need further inquiry and see trends that you can use to understand your costs. <https://docs.aws.amazon.com/cost-management/latest/userguide/ce-what-is.html>

### NEW QUESTION 4

- (Topic 1)

A company has an application that provides marketing services to stores. The services are based on previous purchases by store customers. The stores upload transaction data to the company through SFTP, and the data is processed and analyzed to generate new marketing offers. Some of the files can exceed 200 GB in size.

Recently, the company discovered that some of the stores have uploaded files that contain personally identifiable information (PII) that should not have been included. The company wants administrators to be alerted if PII is shared again. The company also wants to automate remediation.

What should a solutions architect do to meet these requirements with the LEAST development effort?

- A. Use an Amazon S3 bucket as a secure transfer point
- B. Use Amazon Inspector to scan the objects in the bucket
- C. If objects contain PII
- D. trigger an S3 Lifecycle policy to remove the objects that contain PII.

- E. Use an Amazon S3 bucket as a secure transfer point
- F. Use Amazon Macie to scan the objects in the bucket
- G. If objects contain PII
- H. Use Amazon Simple Notification Service (Amazon SNS) to trigger a notification to the administrators to remove the objects that contain PII.
- I. Implement custom scanning algorithms in an AWS Lambda function
- J. Trigger the function when objects are loaded into the bucket
- K. If objects contain PII
- L. Use Amazon Simple Notification Service (Amazon SNS) to trigger a notification to the administrators to remove the objects that contain PII.
- M. Implement custom scanning algorithms in an AWS Lambda function
- N. Trigger the function when objects are loaded into the bucket
- O. If objects contain PII
- P. Use Amazon Simple Email Service (Amazon SES) to trigger a notification to the administrators and trigger on S3 Lifecycle policy to remove the objects that contain PII.

**Answer: B**

**Explanation:**

To meet the requirements of detecting and alerting the administrators when PII is shared and automating remediation with the least development effort, the best approach would be to use Amazon S3 bucket as a secure transfer point and scan the objects in the bucket with Amazon Macie. Amazon Macie is a fully managed data security and data privacy service that uses machine learning and pattern matching to discover and protect sensitive data stored in Amazon S3. It can be used to classify sensitive data, monitor access to sensitive data, and automate remediation actions.

In this scenario, after uploading the files to the Amazon S3 bucket, the objects can be scanned for PII by Amazon Macie, and if it detects any PII, it can trigger an Amazon Simple Notification Service (SNS) notification to alert the administrators to remove the objects containing PII. This approach requires the least development effort, as Amazon Macie already has pre-built data classification rules that can detect PII in various formats. Hence, option B is the correct answer.

References:

? Amazon Macie User Guide: <https://docs.aws.amazon.com/macie/latest/userguide/what-is-macie.html>

? AWS Well-Architected Framework - Security Pillar: <https://docs.aws.amazon.com/wellarchitected/latest/security-pillar/welcome.html>

**NEW QUESTION 5**

- (Topic 1)

A company runs a highly available image-processing application on Amazon EC2 instances in a single VPC. The EC2 instances run inside several subnets across multiple Availability Zones. The EC2 instances do not communicate with each other. However, the EC2 instances download images from Amazon S3 and upload images to Amazon S3 through a single NAT gateway. The company is concerned about data transfer charges. What is the MOST cost-effective way for the company to avoid Regional data transfer charges?

- A. Launch the NAT gateway in each Availability Zone
- B. Replace the NAT gateway with a NAT instance
- C. Deploy a gateway VPC endpoint for Amazon S3
- D. Provision an EC2 Dedicated Host to run the EC2 instances

**Answer: A**

**Explanation:**

In this scenario, the company wants to avoid regional data transfer charges while downloading and uploading images from Amazon S3. To accomplish this at the lowest cost, the NAT gateway should be launched in each availability zone that the EC2 instances are running in. This allows the EC2 instances to route traffic through the local NAT gateway instead of sending traffic across an availability zone boundary and incurring regional data transfer fees. This method will help reduce the data transfer costs since inter-availability zone data transfers in a single region are free of charge.

Reference:

AWS NAT Gateway documentation: <https://docs.aws.amazon.com/vpc/latest/userguide/vpc-nat-gateway.html>

**NEW QUESTION 6**

- (Topic 1)

A company has an on-premises application that generates a large amount of time-sensitive data that is backed up to Amazon S3. The application has grown and there are user complaints about internet bandwidth limitations. A solutions architect needs to design a long-term solution that allows for both timely backups to Amazon S3 and with minimal impact on internet connectivity for internal users.

Which solution meets these requirements?

- A. Establish AWS VPN connections and proxy all traffic through a VPC gateway endpoint
- B. Establish a new AWS Direct Connect connection and direct backup traffic through this new connection.
- C. Order daily AWS Snowball devices. Load the data onto the Snowball devices and return the devices to AWS each day.
- D. Submit a support ticket through the AWS Management Console. Request the removal of S3 service limits from the account.

**Answer: B**

**Explanation:**

To address the issue of bandwidth limitations on the company's on-premises application, and to minimize the impact on internal user connectivity, a new AWS Direct Connect connection should be established to direct backup traffic through this new connection. This solution will offer a secure, high-speed connection between the company's data center and AWS, which will allow the company to transfer data quickly without consuming internet bandwidth.

Reference:

AWS Direct Connect documentation: <https://aws.amazon.com/directconnect/>

**NEW QUESTION 7**

- (Topic 1)

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

Which solution meets these requirements MOST cost-effectively?

- A. Stop the DB instance when tests are complete
- B. Restart the DB instance when required.

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

**Answer:** A

**Explanation:**

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

Reference:

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

**NEW QUESTION 8**

- (Topic 1)

A company runs an on-premises application that is powered by a MySQL database. The company is migrating the application to AWS to increase the application's elasticity and availability.

The current architecture shows heavy read activity on the database during times of normal operation. Every 4 hours, the company's development team pulls a full export of the production database to populate a database in the staging environment. During this period, users experience unacceptable application latency. The development team is unable to use the staging environment until the procedure completes.

A solutions architect must recommend replacement architecture that alleviates the application latency issue. The replacement architecture also must give the development team the ability to continue using the staging environment without delay.

Which solution meets these requirements?

- A. Use Amazon Aurora MySQL with Multi-AZ Aurora Replicas for production.
- B. Populate the staging database by implementing a backup and restore process that uses the mysqldump utility.
- C. Use Amazon Aurora MySQL with Multi-AZ Aurora Replicas for production. Use database cloning to create the staging database on-demand.
- D. Use Amazon RDS for MySQL with a Multi-AZ deployment and read replicas for production. Use the standby instance for the staging database.
- E. Use Amazon RDS for MySQL with a Multi-AZ deployment and read replicas for production.
- F. Populate the staging database by implementing a backup and restore process that uses the mysqldump utility.

**Answer:** B

**Explanation:**

<https://aws.amazon.com/blogs/aws/amazon-aurora-fast-database-cloning/>

**NEW QUESTION 9**

- (Topic 1)

A company uses 50 TB of data for reporting. The company wants to move this data from on-premises to AWS. A custom application in the company's data center runs a weekly data transformation job. The company plans to pause the application until the data transfer is complete and needs to begin the transfer process as soon as possible.

The data center does not have any available network bandwidth for additional workloads. A solutions architect must transfer the data and must configure the transformation job to continue to run in the AWS Cloud.

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

- A. Use AWS DataSync to move the data. Create a custom transformation job by using AWS Glue.
- B. Order an AWS Snowcone device to move the data. Deploy the transformation application to the device.
- C. Order an AWS Snowball Edge Storage Optimized device.
- D. Copy the data to the device.
- E. Create a custom transformation job by using AWS Glue.
- F. Order an AWS Snowball Edge Storage Optimized device that includes Amazon EC2 compute. Copy the data to the device. Create a new EC2 instance on AWS to run the transformation application.

**Answer:** D

**Explanation:**

AWS Snowball Edge is a type of Snowball device with on-board storage and compute power for select AWS capabilities. Snowball Edge can do local processing and edge-computing workloads in addition to transferring data between your local environment and the AWS Cloud<sup>1</sup>. Users can order an AWS Snowball Edge Storage Optimized device that includes Amazon EC2 compute to move 50 TB of data from on-premises to AWS. The Storage Optimized device has 80 TB of usable storage and 40 vCPUs of compute power<sup>2</sup>. Users can copy the data to the device using the AWS OpsHub graphical user interface or the Snowball client command line tool<sup>3</sup>. Users can also create and run Amazon EC2 instances on the device using Amazon Machine Images (AMIs) that are compatible with the sbec1 instance type. Users can use the Snowball Edge device to transfer the data and run the transformation job locally without using any network bandwidth. Users can also create a new EC2 instance on AWS to run the transformation application after the data transfer is complete. Amazon EC2 is a web service that provides secure, resizable compute capacity in the cloud. Users can launch an EC2 instance in the same AWS Region where they send their Snowball Edge device and choose an AMI that matches their application requirements. Users can use the EC2 instance to continue running the transformation job in the AWS Cloud.

**NEW QUESTION 10**

- (Topic 1)

A solutions architect is using Amazon S3 to design the storage architecture of a new digital media application. The media files must be resilient to the loss of an Availability Zone. Some files are accessed frequently while other files are rarely accessed in an unpredictable pattern. The solutions architect must minimize the costs of storing and retrieving the media files.

Which storage option meets these requirements?

- A. S3 Standard
- B. S3 Intelligent-Tiering

- C. S3 Standard-Infrequent Access (S3 Standard-IA)
- D. S3 One Zone-Infrequent Access (S3 One Zone-IA)

**Answer:** B

**Explanation:**

S3 Intelligent-Tiering - Perfect use case when you don't know the frequency of access or irregular patterns of usage. Amazon S3 offers a range of storage classes designed for different use cases. These include S3 Standard for general-purpose storage of frequently accessed data; S3 Intelligent-Tiering for data with unknown or changing access patterns; S3 Standard- Infrequent Access (S3 Standard-IA) and S3 One Zone-Infrequent Access (S3 One Zone- IA) for long-lived, but less frequently accessed data; and Amazon S3 Glacier (S3 Glacier) and Amazon S3 Glacier Deep Archive (S3 Glacier Deep Archive) for long-term archive and digital preservation. If you have data residency requirements that can't be met by an existing AWS Region, you can use the S3 Outposts storage class to store your S3 data on- premises. Amazon S3 also offers capabilities to manage your data throughout its lifecycle. Once an S3 Lifecycle policy is set, your data will automatically transfer to a different storage class without any changes to your application.  
[https://aws.amazon.com/getting-started/hands-on/getting-started-using-amazon-s3-intelligent-tiering/?nc1=h\\_ls](https://aws.amazon.com/getting-started/hands-on/getting-started-using-amazon-s3-intelligent-tiering/?nc1=h_ls)

**NEW QUESTION 10**

- (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 Zones.
- 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 Zones.
- 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 web<sup>1</sup>. 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 object<sup>2</sup>. 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 resources<sup>3</sup>. 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 14**

- (Topic 1)

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

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

**Answer:** B

**Explanation:**

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

**NEW QUESTION 19**

- (Topic 1)

An image-processing company has a web application that users use to upload images. The application uploads the images into an Amazon S3 bucket. The company has set up S3 event notifications to publish the object creation events to an Amazon Simple Queue Service (Amazon SQS) standard queue. The SQS

queue serves as the event source for an AWS Lambda function that processes the images and sends the results to users through email. Users report that they are receiving multiple email messages for every uploaded image. A solutions architect determines that SQS messages are invoking the Lambda function more than once, resulting in multiple email messages. What should the solutions architect do to resolve this issue with the LEAST operational overhead?

- A. Set up long polling in the SQS queue by increasing the ReceiveMessage wait time to 30 seconds.
- B. Change the SQS standard queue to an SQS FIFO queue
- C. Use the message deduplication ID to discard duplicate messages.
- D. Increase the visibility timeout in the SQS queue to a value that is greater than the total of the function timeout and the batch window timeout.
- E. Modify the Lambda function to delete each message from the SQS queue immediately after the message is read before processing.

**Answer: C**

#### NEW QUESTION 22

- (Topic 1)

A company hosts a data lake on AWS. The data lake consists of data in Amazon S3 and Amazon RDS for PostgreSQL. The company needs a reporting solution that provides data visualization and includes all the data sources within the data lake. Only the company's management team should have full access to all the visualizations. The rest of the company should have only limited access. Which solution will meet these requirements?

- A. Create an analysis in Amazon QuickSight
- B. Connect all the data sources and create new dataset
- C. Publish dashboards to visualize the data
- D. Share the dashboards with the appropriate IAM roles.
- E. Create an analysis in Amazon QuickSight
- F. Connect all the data sources and create new dataset
- G. Publish dashboards to visualize the data
- H. Share the dashboards with the appropriate users and groups.
- I. Create an AWS Glue table and crawler for the data in Amazon S3. Create an AWS Glue extract, transform, and load (ETL) job to produce report
- J. Publish the reports to Amazon S3. Use S3 bucket policies to limit access to the reports.
- K. Create an AWS Glue table and crawler for the data in Amazon S3. Use Amazon Athena Federated Query to access data within Amazon RDS for PostgreSQL
- L. Generate reports by using Amazon Athena
- M. Publish the reports to Amazon S3. Use S3 bucket policies to limit access to the reports.

**Answer: B**

#### Explanation:

Amazon QuickSight is a data visualization service that allows you to create interactive dashboards and reports from various data sources, including Amazon S3 and Amazon RDS for PostgreSQL. You can connect all the data sources and create new datasets in QuickSight, and then publish dashboards to visualize the data. You can also share the dashboards with the appropriate users and groups, and control their access levels using IAM roles and permissions. Reference: <https://docs.aws.amazon.com/quicksight/latest/user/working-with-data-sources.html>

#### NEW QUESTION 27

- (Topic 1)

A company runs multiple Windows workloads on AWS. The company's employees use Windows file shares that are hosted on two Amazon EC2 instances. The file shares synchronize data between themselves and maintain duplicate copies. The company wants a highly available and durable storage solution that preserves how users currently access the files. What should a solutions architect do to meet these requirements?

- A. Migrate all the data to Amazon S3 Set up IAM authentication for users to access files
- B. Set up an Amazon S3 File Gateway
- C. Mount the S3 File Gateway on the existing EC2 Instances.
- D. Extend the file share environment to Amazon FSx for Windows File Server with a Multi-AZ configuration
- E. Migrate all the data to FSx for Windows File Server.
- F. Extend the file share environment to Amazon Elastic File System (Amazon EFS) with a Multi-AZ configuration
- G. Migrate all the data to Amazon EFS.

**Answer: C**

#### Explanation:

<https://docs.aws.amazon.com/AWSEC2/latest/WindowsGuide/AmazonEFS.html> Amazon FSx for Windows File Server provides fully managed Microsoft Windows file servers, backed by a fully native Windows file system. <https://docs.aws.amazon.com/fsx/latest/WindowsGuide/what-is.html>

#### NEW QUESTION 32

- (Topic 1)

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

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

Which solution will meet these requirements?

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

- J. Turn on the EBS fast snapshot restore feature on the EBS snapshot
- K. Restore the snapshots into new EBS volume
- L. Attach the new EBS volumes to EC2 instances in the test environment.

**Answer: C**

**Explanation:**

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

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

**NEW QUESTION 33**

- (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 37**

- (Topic 1)

A solutions architect must design a highly available infrastructure for a website. The website is powered by Windows web servers that run on Amazon EC2 instances. The solutions architect must implement a solution that can mitigate a large-scale DDoS attack that originates from thousands of IP addresses.

Downtime is not acceptable for the website.

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

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

**Answer: AC**

**Explanation:**

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

**NEW QUESTION 40**

- (Topic 1)

An Amazon EC2 administrator created the following policy associated with an IAM group containing several users

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "ec2:TerminateInstances",
      "Resource": "*",
      "Condition": {
        "IpAddress": {
          "aws:SourceIp": "10.100.100.0/24"
        }
      }
    },
    {
      "Effect": "Deny",
      "Action": "ec2:*",
      "Resource": "*",
      "Condition": {
        "StringNotEquals": {
          "ec2:Region": "us-east-1"
        }
      }
    }
  ]
}
```

What is the effect of this policy?

- A. Users can terminate an EC2 instance in any AWS Region except us-east-1.
- B. Users can terminate an EC2 instance with the IP address 10 100 100 1 in the us-east-1 Region
- C. Users can terminate an EC2 instance in the us-east-1 Region when the user's source IP is 10.100.100.254.
- D. Users cannot terminate an EC2 instance in the us-east-1 Region when the user's source IP is 10.100 100 254

**Answer: C**

**Explanation:**

as the policy prevents anyone from doing any EC2 action on any region except us-east-1 and allows only users with source ip 10.100.100.0/24 to terminate instances. So user with source ip 10.100.100.254 can terminate instances in us-east-1 region.

**NEW QUESTION 43**

- (Topic 1)

A solutions architect is designing a new hybrid architecture to extend a company's on-premises infrastructure to AWS. The company requires a highly available connection with consistent low latency to an AWS Region. The company needs to minimize costs and is willing to accept slower traffic if the primary connection fails.

What should the solutions architect do to meet these requirements?

- A. Provision an AWS Direct Connect connection to a Region. Provision a VPN connection as a backup if the primary Direct Connect connection fails.
- B. Provision a VPN tunnel connection to a Region for private connectivity.
- C. Provision a second VPN tunnel for private connectivity and as a backup if the primary VPN connection fails.
- D. Provision an AWS Direct Connect connection to a Region. Provision a second Direct Connect connection to the same Region as a backup if the primary Direct Connect connection fails.
- E. Provision an AWS Direct Connect connection to a Region. Use the Direct Connect failover attribute from the AWS CLI to automatically create a backup connection if the primary Direct Connect connection fails.

**Answer: A**

**Explanation:**

"In some cases, this connection alone is not enough. It is always better to guarantee a fallback connection as the backup of DX. There are several options, but implementing it with an AWS Site-To-Site VPN is a real cost-effective solution that can be exploited to reduce costs or, in the meantime, wait for the setup of a second DX." <https://www.proud2becloud.com/hybrid-cloud-networking-backup-aws-direct-connect-network-connection-with-aws-site-to-site-vpn/>

**NEW QUESTION 45**

- (Topic 1)

A company has an Amazon S3 bucket that contains critical data. The company must protect the data from accidental deletion. Which combination of steps should a solutions architect take to meet these requirements?

(Choose two.)

- A. Enable versioning on the S3 bucket.
- B. Enable MFA Delete on the S3 bucket.
- C. Create a bucket policy on the S3 bucket.
- D. Enable default encryption on the S3 bucket.
- E. Create a lifecycle policy for the objects in the S3 bucket.

**Answer:** AB

**Explanation:**

To protect data in an S3 bucket from accidental deletion, versioning should be enabled, which enables you to preserve, retrieve, and restore every version of every object in an S3 bucket. Additionally, enabling MFA (multi-factor authentication) Delete on the S3 bucket adds an extra layer of protection by requiring an authentication token in addition to the user's access keys to delete objects in the bucket.

Reference:

AWS S3 Versioning documentation: <https://docs.aws.amazon.com/AmazonS3/latest/dev/Versioning.html>

AWS S3 MFA Delete documentation: <https://docs.aws.amazon.com/AmazonS3/latest/dev/UsingMFADelete.html>

**NEW QUESTION 49**

- (Topic 1)

A company has more than 5 TB of file data on Windows file servers that run on premises. Users and applications interact with the data each day.

The company is moving its Windows workloads to AWS. As the company continues this process, the company requires access to AWS and on-premises file storage with minimum latency. The company needs a solution that minimizes operational overhead and requires no significant changes to the existing file access patterns. The company uses an AWS Site-to-Site VPN connection for connectivity to AWS.

What should a solutions architect do to meet these requirements?

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

**Answer:** D

**Explanation:**

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

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

Reference:

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

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

**NEW QUESTION 51**

- (Topic 1)

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

What should the solutions architect recommend?

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

**Answer:** C

**Explanation:**

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

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

**NEW QUESTION 52**

- (Topic 1)

A company that hosts its web application on AWS wants to ensure all Amazon EC2 instances, Amazon RDS DB instances, and Amazon Redshift clusters are configured with tags. The company wants to minimize the effort of configuring and operating this check.

What should a solutions architect do to accomplish this?

- A. Use AWS Config rules to define and detect resources that are not properly tagged.
- B. Use Cost Explorer to display resources that are not properly tagged.
- C. Tag those resources manually.
- D. Write API calls to check all resources for proper tag allocation.
- E. Periodically run the code on an EC2 instance.
- F. Write API calls to check all resources for proper tag allocation.
- G. Schedule an AWS Lambda function through Amazon CloudWatch to periodically run the code.

**Answer:** A

**Explanation:**

To ensure all Amazon EC2 instances, Amazon RDS DB instances, and Amazon Redshift clusters are configured with tags, a solutions architect should use AWS

Config rules to define and detect resources that are not properly tagged. AWS Config rules are a set of customizable rules that AWS Config uses to evaluate AWS resource configurations for compliance with best practices and company policies. Using AWS Config rules can minimize the effort of configuring and operating this check because it automates the process of identifying non-compliant resources and notifying the responsible teams. Reference: AWS Config Developer Guide: AWS Config Rules ([https://docs.aws.amazon.com/config/latest/developerguide/evaluate-config\\_use-managed-rules.html](https://docs.aws.amazon.com/config/latest/developerguide/evaluate-config_use-managed-rules.html))

#### NEW QUESTION 57

- (Topic 1)

A social media company allows users to upload images to its website. The website runs on Amazon EC2 instances. During upload requests, the website resizes the images to a standard size and stores the resized images in Amazon S3. Users are experiencing slow upload requests to the website. The company needs to reduce coupling within the application and improve website performance. A solutions architect must design the most operationally efficient process for image uploads.

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

- A. Configure the application to upload images to S3 Glacier.
- B. Configure the web server to upload the original images to Amazon S3.
- C. Configure the application to upload images directly from each user's browser to Amazon S3 through the use of a presigned URL.
- D. Configure S3 Event Notifications to invoke an AWS Lambda function when an image is uploaded.
- E. Use the function to resize the image.
- F. Create an Amazon EventBridge (Amazon CloudWatch Events) rule that invokes an AWS Lambda function on a schedule to resize uploaded images.

**Answer: CD**

#### Explanation:

Amazon S3 is a highly scalable and durable object storage service that can store and retrieve any amount of data from anywhere on the web<sup>1</sup>. 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 object<sup>2</sup>. 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 resources<sup>3</sup>. 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 61

- (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 63

- (Topic 1)

A company is storing backup files by using Amazon S3 Standard storage. The files are accessed frequently for 1 month. However, the files are not accessed after 1 month. The company must keep the files indefinitely.

Which storage solution will meet these requirements MOST cost-effectively?

- A. Configure S3 Intelligent-Tiering to automatically migrate objects.
- B. Create an S3 Lifecycle configuration to transition objects from S3 Standard to S3 Glacier Deep Archive after 1 month.
- C. Create an S3 Lifecycle configuration to transition objects from S3 Standard to S3 Standard-Infrequent Access (S3 Standard-IA) after 1 month.
- D. Create an S3 Lifecycle configuration to transition objects from S3 Standard to S3 One Zone-Infrequent Access (S3 One Zone-IA) after 1 month.

**Answer: B**

#### Explanation:

The storage solution that will meet these requirements most cost-effectively is B: Create an S3 Lifecycle configuration to transition objects from S3 Standard to S3 Glacier Deep Archive after 1 month. Amazon S3 Glacier Deep Archive is a secure, durable, and extremely low-cost Amazon S3 storage class for long-term retention of data that is rarely accessed and for which retrieval times of several hours are acceptable. It is the lowest-cost storage option in Amazon S3, making it a cost-effective choice for storing backup files that are not accessed after 1 month. You can use an S3 Lifecycle configuration to automatically transition objects from S3 Standard to S3 Glacier Deep Archive after 1 month. This will minimize the storage costs for the backup files that are not accessed frequently.

#### NEW QUESTION 67

- (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 databases
- F. Configure the read replicas with half of the compute and storage resources as the source database.
- G. Create read replicas for the databases
- 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](https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_MySQL.Replication.ReadReplicas.html)

#### NEW QUESTION 71

- (Topic 1)

A company has thousands of edge devices that collectively generate 1 TB of status alerts each day. Each alert is approximately 2 KB in size. A solutions architect needs to implement a solution to ingest and store the alerts for future analysis.

The company wants a highly available solution. However, the company needs to minimize costs and does not want to manage additional infrastructure. Additionally, the company wants to keep 14 days of data available for immediate analysis and archive any data older than 14 days.

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

- A. Create an Amazon Kinesis Data Firehose delivery stream to ingest the alerts. Configure the Kinesis Data Firehose stream to deliver the alerts to an Amazon S3 bucket. Set up an S3 Lifecycle configuration to transition data to Amazon S3 Glacier after 14 days.
- B. Launch Amazon EC2 instances across two Availability Zones and place them behind an Elastic Load Balancer to ingest the alerts. Create a script on the EC2 instances that will store the alerts in an Amazon S3 bucket. Set up an S3 Lifecycle configuration to transition data to Amazon S3 Glacier after 14 days.
- C. Create an Amazon Kinesis Data Firehose delivery stream to ingest the alerts. Configure the Kinesis Data Firehose stream to deliver the alerts to an Amazon Elasticsearch Service (Amazon ES) cluster. Set up the Amazon ES cluster to take manual snapshots every day and delete data from the cluster that is older than 14 days.
- D. Create an Amazon Simple Queue Service (Amazon SQS) standard queue to ingest the alerts and set the message retention period to 14 days. Configure consumers to poll the SQS queue, check the age of the message, and analyze the message data as needed. If the message is 14 days old, the consumer should copy the message to an Amazon S3 bucket and delete the message from the SQS queue.

**Answer: A**

**Explanation:**

<https://aws.amazon.com/kinesis/data-firehose/features/?nc=sn&loc=2#:~:text=into%20Amazon%20S3%2C%20Amazon%20Redshift%2C%20Amazon%20OpenSearch%20Service%2C%20Kinesis,Delivery%20streams>

#### NEW QUESTION 73

- (Topic 1)

A company uses AWS Organizations to manage multiple AWS accounts for different departments. The management account has an Amazon S3 bucket that contains project reports. The company wants to limit access to this S3 bucket to only users of accounts within the organization in AWS Organizations.

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

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

**Answer: A**

**Explanation:**

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

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

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

[https://docs.aws.amazon.com/IAM/latest/UserGuide/reference\\_policies\\_condition-keys.html](https://docs.aws.amazon.com/IAM/latest/UserGuide/reference_policies_condition-keys.html)

#### NEW QUESTION 74

- (Topic 1)

A company runs its Infrastructure on AWS and has a registered base of 700,000 users for a document management application. The company intends to create a product that converts large PDF files to JPG/IMAGO files. The PDF files average 5 MB in size. The company needs to store the original files and the converted files. A solutions architect must design a scalable solution to accommodate demand that will grow rapidly over time.

Which solution meets these requirements MOST cost-effectively?

- A. Save the PDF files to Amazon S3. Configure an S3 PUT event to invoke an AWS Lambda function to convert the files to JPG format and store them back in Amazon S3.

- B. Save the pdf files to Amazon DynamoD
- C. Use the DynamoDB Streams feature to invoke an AWS Lambda function to convert the files to jpg format and store them hack in DynamoDB
- D. Upload the pdf files to an AWS Elastic Beanstalk application that includes Amazon EC2 instance
- E. Amazon Elastic Block Store (Amazon EBS) storage and an Auto Scaling grou
- F. Use a program In the EC2 instances to convert the files to jpg format Save the .pdf files and the .jpg files In the EBS store.
- G. Upload the .pdf files to an AWS Elastic Beanstalk application that includes Amazon EC2 instances, Amazon Elastic File System (Amazon EPS) storage, and an Auto Scaling grou
- H. Use a program in the EC2 instances to convert the file to jpg format Save the pdf files and the jpg files in the EBS store.

**Answer:** A

**Explanation:**

Elastic BeanStalk is expensive, and DocumentDB has a 400KB max to upload files. So Lambda and S3 should be the one.

**NEW QUESTION 76**

- (Topic 1)

A company is running a business-critical web application on Amazon EC2 instances behind an Application Load Balancer. The EC2 instances are in an Auto Scaling group. The application uses an Amazon Aurora PostgreSQL database that is deployed in a single Availability Zone. The company wants the application to be highly available with minimum downtime and minimum loss of data.

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

- A. Place the EC2 instances in different AWS Region
- B. Use Amazon Route 53 health checks to redirect traffi
- C. Use Aurora PostgreSQL Cross-Region Replication.
- D. Configure the Auto Scaling group to use multiple Availability Zone
- E. Configure the database as Multi-A
- F. Configure an Amazon RDS Proxy instance for the database.
- G. Configure the Auto Scaling group to use one Availability Zon
- H. Generate hourly snapshots of the databas
- I. Recover the database from the snapshots in the event of a failure.
- J. Configure the Auto Scaling group to use multiple AWS Region
- K. Write the data from the application to Amazon S3. Use S3 Event Notifications to launch an AWS Lambda function to write the data to the database.

**Answer:** B

**Explanation:**

To achieve high availability with minimum downtime and minimum loss of data, the Auto Scaling group should be configured to use multiple Availability Zones to ensure that there is no single point of failure. The database should be configured as Multi- AZ to enable automatic failover in case of an outage in the primary Availability Zone. Additionally, an Amazon RDS Proxy instance can be used to improve the scalability and availability of the database by reducing connection failures and improving failover times.

**NEW QUESTION 80**

- (Topic 1)

A company recently launched a variety of new workloads on Amazon EC2 instances in its AWS account. The company needs to create a strategy to access and administer the instances remotely and securely. The company needs to implement a repeatable process that works with native AWS services and follows the AWS Well-Architected Framework.

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

- A. Use the EC2 serial console to directly access the terminal interface of each instance foradministration.
- B. Attach the appropriate IAM role to each existing instance and new instanc
- C. Use AWS Systems Manager Session Manager to establish a remote SSH session.
- D. Create an administrative SSH key pai
- E. Load the public key into each EC2 instanc
- F. Deploy a bastion host in a public subnet to provide a tunnel for administration of each instance.
- G. Establish an AWS Site-to-Site VPN connectio
- H. Instruct administrators to use their local on-premises machines to connect directly to the instances by using SSH keys across the VPN tunnel.

**Answer:** B

**Explanation:**

<https://docs.aws.amazon.com/systems-manager/latest/userguide/setup-launch-managed-instance.html>

**NEW QUESTION 84**

- (Topic 1)

A company wants to reduce the cost of its existing three-tier web architecture. The web, application, and database servers are running on Amazon EC2 instances for the development, test, and production environments. The EC2 instances average 30% CPU utilization during peak hours and 10% CPU utilization during non-peak hours.

The production EC2 instances run 24 hours a day. The development and test EC2 instances run for at least 8 hours each day. The company plans to implement automation to stop the development and test EC2 instances when they are not in use.

Which EC2 instance purchasing solution will meet the company's requirements MOST cost-effectively?

- A. Use Spot Instances for the production EC2 instance
- B. Use Reserved Instances for the development and test EC2 instances.
- C. Use Reserved Instances for the production EC2 instance
- D. Use On-Demand Instances for the development and test EC2 instances.
- E. Use Spot blocks for the production EC2 instance
- F. Use Reserved Instances for the development and test EC2 instances.
- G. Use On-Demand Instances for the production EC2 instance
- H. Use Spot blocks for the development and test EC2 instances.

Answer: B

#### NEW QUESTION 85

- (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](https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_StopInstance.html)

#### NEW QUESTION 87

- (Topic 2)

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

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

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

Answer: D

#### Explanation:

This solution meets the requirements of saving money on storage while keeping the most accessed files readily available for the users. S3 Lifecycle policy can automatically move objects from one storage class to another based on predefined rules. S3 Standard-IA is a lower-cost storage class for data that is accessed less frequently, but requires rapid access when needed. It is suitable for ringtones older than 90 days that are downloaded infrequently.

Option A is incorrect because configuring S3 Standard-IA for the initial storage tier of the objects can incur higher costs for frequent access and retrieval fees.

Option B is incorrect

because moving the files to S3 Intelligent-Tiering can incur additional monitoring and automation fees that may not be necessary for ringtones older than 90 days.

Option C is incorrect because using S3 inventory to manage objects and move them to S3 Standard-IA can be complex and time-consuming, and it does not provide automatic cost savings. References:

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

? <https://aws.amazon.com/s3/cloud-storage-cost-optimization-ebook/>

#### NEW QUESTION 88

- (Topic 2)

A company has a Windows-based application that must be migrated to AWS. The application requires the use of a shared Windows file system attached to multiple Amazon EC2 Windows instances that are deployed across multiple Availability Zones.

What should a solutions architect do to meet this requirement?

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

**Answer:** B

**Explanation:**

This solution meets the requirement of migrating a Windows-based application that requires the use of a shared Windows file system attached to multiple Amazon EC2 Windows instances that are deployed across multiple Availability Zones. Amazon FSx for Windows File Server provides fully managed shared storage built on Windows Server, and delivers a wide range of data access, data management, and administrative capabilities. It supports the Server Message Block (SMB) protocol and can be mounted to EC2 Windows instances across multiple Availability Zones.

Option A is incorrect because AWS Storage Gateway in volume gateway mode provides cloud-backed storage volumes that can be mounted as iSCSI devices from on-premises application servers, but it does not support SMB protocol or EC2 Windows instances. Option C is incorrect because Amazon Elastic File System (Amazon EFS) provides a scalable and elastic NFS file system for Linux-based workloads, but it does not support SMB protocol or EC2 Windows instances. Option D is incorrect because Amazon Elastic Block Store (Amazon EBS) provides persistent block storage volumes for use with EC2 instances, but it does not support SMB protocol or attaching multiple instances to the same volume.

References:

? <https://aws.amazon.com/fsx/windows/>

? <https://docs.aws.amazon.com/fsx/latest/WindowsGuide/using-file-shares.html>

**NEW QUESTION 89**

- (Topic 2)

A company wants to migrate its on-premises data center to AWS. According to the company's compliance requirements, the company can use only the ap-northeast-3 Region. Company administrators are not permitted to connect VPCs to the internet.

Which solutions will meet these requirements? (Choose two.)

- A. Use AWS Control Tower to implement data residency guardrails to deny internet access and deny access to all AWS Regions except ap-northeast-3.
- B. Use rules in AWS WAF to prevent internet access
- C. Deny access to all AWS Regions except ap-northeast-3 in the AWS account settings.
- D. Use AWS Organizations to configure service control policies (SCPs) that prevent VPCs from gaining internet access
- E. Deny access to all AWS Regions except ap-northeast-3.
- F. Create an outbound rule for the network ACL in each VPC to deny all traffic from 0.0.0.0/0. Create an IAM policy for each user to prevent the use of any AWS Region other than ap-northeast-3.
- G. Use AWS Config to activate managed rules to detect and alert for internet gateways and to detect and alert for new resources deployed outside of ap-northeast-3.

**Answer:** AC

**Explanation:**

[https://docs.aws.amazon.com/organizations/latest/userguide/orgs\\_manage\\_policies\\_scps\\_examples\\_vpc.html#example\\_vpc\\_2](https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps_examples_vpc.html#example_vpc_2)

**NEW QUESTION 92**

- (Topic 2)

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

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

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

**Answer:** AE

**Explanation:**

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

**NEW QUESTION 97**

- (Topic 2)

A solutions architect needs to help a company optimize the cost of running an application on AWS. The application will use Amazon EC2 instances, AWS Fargate, and AWS Lambda for compute within the architecture.

The EC2 instances will run the data ingestion layer of the application. EC2 usage will be sporadic and unpredictable. Workloads that run on EC2 instances can be interrupted at any time. The application front end will run on Fargate, and Lambda will serve the API layer. The front-end utilization and API layer utilization will be predictable over the course of the next year.

Which combination of purchasing options will provide the MOST cost-effective solution for hosting this application? (Choose two.)

- A. Use Spot Instances for the data ingestion layer
- B. Use On-Demand Instances for the data ingestion layer
- C. Purchase a 1-year Compute Savings Plan for the front end and API layer.
- D. Purchase 1-year All Upfront Reserved instances for the data ingestion layer.
- E. Purchase a 1-year EC2 instance Savings Plan for the front end and API layer.

**Answer:** AC

**Explanation:**

EC2 instance Savings Plan saves 72% while Compute Savings Plans saves 66%. But according to link, it says "Compute Savings Plans provide the most flexibility and help to reduce your costs by up to 66%. These plans automatically apply to EC2 instance usage regardless of instance family, size, AZ, region, OS or tenancy, and also apply to Fargate and Lambda usage." EC2 instance Savings Plans are not applied to Fargate or Lambda

**NEW QUESTION 100**

- (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](https://aws.amazon.com/s3/faqs/#Durability_and_data_protection)

**NEW QUESTION 102**

- (Topic 2)

A company is migrating its on-premises PostgreSQL database to Amazon Aurora PostgreSQL. The on-premises database must remain online and accessible during the migration. The Aurora database must remain synchronized with the on-premises database.

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

- A. Create an ongoing replication task.
- B. Create a database backup of the on-premises database
- C. Create an AWS Database Migration Service (AWS DMS) replication server
- D. Convert the database schema by using the AWS Schema Conversion Tool (AWS SCT).
- E. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to monitor the database synchronization

**Answer: AC**

**Explanation:**

AWS Database Migration Service supports homogeneous migrations such as Oracle to Oracle, as well as heterogeneous migrations between different database platforms, such as Oracle or Microsoft SQL Server to Amazon Aurora. With AWS Database Migration Service, you can also continuously replicate data with low latency from any supported source to any supported target. For example, you can replicate from multiple sources to Amazon Simple Storage Service (Amazon S3) to build a highly available and scalable data lake solution. You can also consolidate databases into a petabyte-scale data warehouse by streaming data to Amazon Redshift. Learn more about the supported source and target databases. <https://aws.amazon.com/dms/>

**NEW QUESTION 104**

- (Topic 2)

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

Which storage solution meets these requirements MOST cost-effectively?

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

**Answer: D**

**Explanation:**

Amazon S3 is cheapest and can be accessed from anywhere.

**NEW QUESTION 109**

- (Topic 2)

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

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

Which solution will meet these requirements?

- A. Migrate the purchase data to write directly to Amazon RD

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

**Answer: C**

**Explanation:**

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

**NEW QUESTION 110**

- (Topic 2)

A reporting team receives files each day in an Amazon S3 bucket. The reporting team manually reviews and copies the files from this initial S3 bucket to an analysis S3 bucket each day at the same time to use with Amazon QuickSight. Additional teams are starting to send more files in larger sizes to the initial S3 bucket.

The reporting team wants to move the files automatically to the analysis S3 bucket as the files enter the initial S3 bucket. The reporting team also wants to use AWS Lambda functions to run pattern-matching code on the copied data. In addition, the reporting team wants to send the data files to a pipeline in Amazon SageMaker Pipelines.

What should a solutions architect do to meet these requirements with the LEAST operational overhead?

- A. Create a Lambda function to copy the files to the analysis S3 bucket
- B. Create an S3 event notification for the analysis S3 bucket
- C. Configure Lambda and SageMaker Pipelines as destinations of the event notification
- D. Configure s3:ObjectCreated:Put as the event type.
- E. Create a Lambda function to copy the files to the analysis S3 bucket
- F. Configure the analysis S3 bucket to send event notifications to Amazon EventBridge (Amazon CloudWatch Events). Configure an ObjectCreated rule in EventBridge (CloudWatch Events). Configure Lambda and SageMaker Pipelines as targets for the rule.
- G. Configure S3 replication between the S3 buckets
- H. Create an S3 event notification for the analysis S3 bucket
- I. Configure Lambda and SageMaker Pipelines as destinations of the event notification
- J. Configure s3:ObjectCreated:Put as the event type.
- K. Configure S3 replication between the S3 buckets
- L. Configure the analysis S3 bucket to send event notifications to Amazon EventBridge (Amazon CloudWatch Events). Configure an ObjectCreated rule in EventBridge (CloudWatch Events). Configure Lambda and SageMaker Pipelines as targets for the rule.

**Answer: D**

**Explanation:**

This solution meets the requirements of moving the files automatically, running Lambda functions on the copied data, and sending the data files to SageMaker Pipelines with the least operational overhead. S3 replication can copy the files from the initial S3 bucket to the analysis S3 bucket as they arrive. The analysis S3 bucket can send event notifications to Amazon EventBridge (Amazon CloudWatch Events) when an object is created. EventBridge can trigger Lambda and SageMaker Pipelines as targets for the ObjectCreated rule. Lambda can run pattern-matching code on the copied data, and SageMaker Pipelines can execute a pipeline with the data files.

Option A is incorrect because creating a Lambda function to copy the files to the analysis S3 bucket is not necessary when S3 replication can do that automatically. It also adds operational overhead to manage the Lambda function. Option B is incorrect because creating a Lambda function to copy the files to the analysis S3 bucket is not necessary when S3 replication can do that automatically. It also adds operational overhead to manage the Lambda function. Option C is incorrect because using S3 event notification with multiple destinations can result in throttling or delivery failures if there are too many events. References:

? <https://aws.amazon.com/blogs/machine-learning/automate-feature-engineering-pipelines-with-amazon-sagemaker/>

? <https://docs.aws.amazon.com/sagemaker/latest/dg/automating-sagemaker-with-eventbridge.html>

? <https://aws.amazon.com/about-aws/whats-new/2021/04/new-options-trigger-amazon-sagemaker-pipeline-executions/>

**NEW QUESTION 112**

- (Topic 2)

A company has an ecommerce checkout workflow that writes an order to a database and calls a service to process the payment. Users are experiencing timeouts during the checkout process. When users resubmit the checkout form, multiple unique orders are created for the same desired transaction.

How should a solutions architect refactor this workflow to prevent the creation of multiple orders?

- A. Configure the web application to send an order message to Amazon Kinesis Data Firehose
- B. Set the payment service to retrieve the message from Kinesis Data Firehose and process the order.
- C. Create a rule in AWS CloudTrail to invoke an AWS Lambda function based on the logged application path request. Use Lambda to query the database, call the payment service, and pass in the order information.
- D. Store the order in the database
- E. Send a message that includes the order number to Amazon Simple Notification Service (Amazon SNS). Set the payment service to poll Amazon SNS
- F. retrieve the message, and process the order.
- G. Store the order in the database
- H. Send a message that includes the order number to an Amazon Simple Queue Service (Amazon SQS) FIFO queue
- I. Set the payment service to retrieve the message and process the order
- J. Delete the message from the queue.

**Answer: D**

**Explanation:**

This approach ensures that the order creation and payment processing steps are separate and atomic. By sending the order information to an SQS FIFO queue, the payment service can process the order one at a time and in the order they were received. If the payment service is unable to process an order, it can be retried later, preventing the creation of multiple orders. The deletion of the message from the queue after it is processed will prevent the same message from being processed multiple times.

**NEW QUESTION 114**

- (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-performing storage.

Which solution will meet these requirements?

- A. Use a duster 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 instance==>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-efs-mult-attach>

**NEW QUESTION 116**

- (Topic 2)

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

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

**Answer: A**

**Explanation:**

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

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

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

**NEW QUESTION 119**

- (Topic 2)

A company runs workloads on AWS. The company needs to connect to a service from an external provider. The service is hosted in the provider's VPC. According to the company's security team, the connectivity must be private and must be restricted to the target service. The connection must be initiated only from the company's VPC. Which solution will meet these requirements?

- A. Create a VPC peering connection between the company's VPC and the provider's VPC
- B. Update the route table to connect to the target service.
- C. Ask the provider to create a virtual private gateway in its VPC
- D. Use AWS PrivateLink to connect to the target service.
- E. Create a NAT gateway in a public subnet of the company's VPC

- F. Update the route table to connect to the target service.
- G. Ask the provider to create a VPC endpoint for the target service.
- H. Use AWS PrivateLink to connect to the target service.

Answer: D

**Explanation:**

**\*\*AWS PrivateLink** provides private connectivity between VPCs, AWS services, and your on-premises networks, without exposing your traffic to the public internet. AWS PrivateLink makes it easy to connect services across different accounts and VPCs to significantly simplify your network architecture. Interface **\*\*VPC endpoints\*\***, powered by AWS PrivateLink, connect you to services hosted by AWS Partners and supported solutions available in AWS Marketplace.  
<https://aws.amazon.com/privatelink/>

**NEW QUESTION 120**

- (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](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 121**

- (Topic 2)

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

How should a solutions architect design the application to ensure the LEAST amount of latency for all users?

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

- F. Use Amazon CloudFront to serve the static content
- G. Serve the dynamic content directly from the ALB.
- H. Deploy the application stack in two AWS Region
- I. Use an Amazon Route 53 geolocation routing policy to serve all content from the ALB in the closest Region.

**Answer:** A

**Explanation:**

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

**NEW QUESTION 124**

- (Topic 2)

A hospital wants to create digital copies for its large collection of historical written records. The hospital will continue to add hundreds of new documents each day. The hospital's data team will scan the documents and will upload the documents to the AWS Cloud.

A solutions architect must implement a solution to analyze the documents, extract the medical information, and store the documents so that an application can run SQL queries on the data. The solution must maximize scalability and operational efficiency.

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

- A. Write the document information to an Amazon EC2 instance that runs a MySQL database.
- B. Write the document information to an Amazon S3 bucket
- C. Use Amazon Athena to query the data.
- D. Create an Auto Scaling group of Amazon EC2 instances to run a custom application that processes the scanned files and extracts the medical information.
- E. Create an AWS Lambda function that runs when new documents are uploaded
- F. Use Amazon Rekognition to convert the documents to raw text
- G. Use Amazon Transcribe Medical to detect and extract relevant medical information from the text.
- H. Create an AWS Lambda function that runs when new documents are uploaded
- I. Use Amazon Textract to convert the documents to raw text
- J. Use Amazon Comprehend Medical to detect and extract relevant medical information from the text.

**Answer:** BE

**Explanation:**

This solution meets the requirements of creating digital copies for a large collection of historical written records, analyzing the documents, extracting the medical information, and storing the documents so that an application can run SQL queries on the data. Writing the document information to an Amazon S3 bucket can provide scalable and durable storage for the scanned files. Using Amazon Athena to query the data can provide serverless and interactive SQL analysis on data stored in S3. Creating an AWS Lambda function that runs when new documents are uploaded can provide event-driven and serverless processing of the scanned files. Using Amazon Textract to convert the documents to raw text can provide

accurate optical character recognition (OCR) and extraction of structured data such as tables and forms from documents using artificial intelligence (AI). Using Amazon Comprehend Medical to detect and extract relevant medical information from the text can provide natural language processing (NLP) service that uses machine learning that has been pre-trained to understand and extract health data from medical text.

Option A is incorrect because writing the document information to an Amazon EC2 instance that runs a MySQL database can increase the infrastructure overhead and complexity, and it may not be able to handle large volumes of data. Option C is incorrect because creating an Auto Scaling group of Amazon EC2 instances to run a custom application that processes the scanned files and extracts the medical information can increase the infrastructure overhead and complexity, and it may not be able to leverage existing AI and NLP services such as Textract and Comprehend Medical. Option D is incorrect because using Amazon Rekognition to convert the documents to raw text can provide image and video analysis, but it does not support OCR or extraction of structured data from documents. Using Amazon Transcribe Medical to detect and extract relevant medical information from the text can provide speech-to-text transcription service for medical conversations, but it does not support text analysis or extraction of health data from medical text.

References:

- ? <https://aws.amazon.com/s3/>
- ? <https://aws.amazon.com/athena/>
- ? <https://aws.amazon.com/lambda/>
- ? <https://aws.amazon.com/textract/>
- ? <https://aws.amazon.com/comprehend/medical/>

**NEW QUESTION 127**

- (Topic 2)

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

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

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

**Answer:** D

**Explanation:**

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

**NEW QUESTION 129**

- (Topic 2)

A company uses AWS Organizations to create dedicated AWS accounts for each business unit to manage each business unit's account independently upon request. The root email recipient missed a notification that was sent to the root user email address of one account. The company wants to ensure that all future

notifications are not missed. Future notifications must be limited to account administrators.  
Which solution will meet these requirements?

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

**Answer: B**

**Explanation:**

Use a group email address for the management account's root user [https://docs.aws.amazon.com/organizations/latest/userguide/orgs\\_best-practices\\_mgmt-acct.html#best-practices\\_mgmt-acct\\_email-address](https://docs.aws.amazon.com/organizations/latest/userguide/orgs_best-practices_mgmt-acct.html#best-practices_mgmt-acct_email-address)

**NEW QUESTION 131**

- (Topic 2)

A gaming company has a web application that displays scores. The application runs on Amazon EC2 instances behind an Application Load Balancer. The application stores data in an Amazon RDS for MySQL database. Users are starting to experience long delays and interruptions that are caused by database read performance. The company wants to improve the user experience while minimizing changes to the application's architecture.  
What should a solutions architect do to meet these requirements?

- A. Use Amazon ElastiCache in front of the database.
- B. Use RDS Proxy between the application and the database.
- C. Migrate the application from EC2 instances to AWS Lambda.
- D. Migrate the database from Amazon RDS for MySQL to Amazon DynamoDB.

**Answer: A**

**Explanation:**

ElastiCache can help speed up the read performance of the database by caching frequently accessed data, reducing latency and allowing the application to access the data more quickly. This solution requires minimal modifications to the current architecture, as ElastiCache can be used in conjunction with the existing Amazon RDS for MySQL database.

**NEW QUESTION 132**

- (Topic 2)

A company has a service that produces event data. The company wants to use AWS to process the event data as it is received. The data is written in a specific order that must be maintained throughout processing. The company wants to implement a solution that minimizes operational overhead.  
How should a solutions architect accomplish this?

- A. Create an Amazon Simple Queue Service (Amazon SQS) FIFO queue to hold messages. Set up an AWS Lambda function to process messages from the queue.
- B. Create an Amazon Simple Notification Service (Amazon SNS) topic to deliver notifications containing payloads to process. Configure an AWS Lambda function as a subscriber.
- C. Create an Amazon Simple Queue Service (Amazon SQS) standard queue to hold messages.
- D. Set up an AWS Lambda function to process messages from the queue independently.
- E. Create an Amazon Simple Notification Service (Amazon SNS) topic to deliver notifications containing payloads to process.
- F. Configure an Amazon Simple Queue Service (Amazon SQS) queue as a subscriber.

**Answer: A**

**Explanation:**

The details are revealed in below url: <https://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/FIFO-queues.html>  
FIFO (First-In-First-Out) queues are designed to enhance messaging between applications when the order of operations and events is critical, or where duplicates can't be tolerated. Examples of situations where you might use FIFO queues include the following: To make sure that user-entered commands are run in the right order. To display the correct product price by sending price modifications in the right order. To prevent a student from enrolling in a course before registering for an account.

**NEW QUESTION 136**

- (Topic 2)

A company is developing a file-sharing application that will use an Amazon S3 bucket for storage. The company wants to serve all the files through an Amazon CloudFront distribution. The company does not want the files to be accessible through direct navigation to the S3 URL.  
What should a solutions architect do to meet these requirements?

- A. Write individual policies for each S3 bucket to grant read permission for only CloudFront access.
- B. Create an IAM user.
- C. Grant the user read permission to objects in the S3 bucket.
- D. Assign the user to CloudFront.
- E. Write an S3 bucket policy that assigns the CloudFront distribution ID as the Principal and assigns the target S3 bucket as the Amazon Resource Name (ARN).
- F. Create an origin access identity (OAI). Assign the OAI to the CloudFront distribution.
- G. Configure the S3 bucket permissions so that only the OAI has read permission.

**Answer: D**

**Explanation:**

<https://aws.amazon.com/premiumsupport/knowledge-center/cloudfront-access-to-amazon-s3/>  
<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/private-content-restricting-access-to-s3.html#private-content-restricting-access-to-s3-overview>

### NEW QUESTION 139

- (Topic 2)

A company has a dynamic web application hosted on two Amazon EC2 instances. The company has its own SSL certificate, which is on each instance to perform SSL termination.

There has been an increase in traffic recently, and the operations team determined that SSL encryption and decryption is causing the compute capacity of the web servers to reach their maximum limit.

What should a solutions architect do to increase the application's performance?

- A. Create a new SSL certificate using AWS Certificate Manager (ACM) install the ACM certificate on each instance
- B. Create an Amazon S3 bucket Migrate the SSL certificate to the S3 bucket Configure the EC2 instances to reference the bucket for SSL termination
- C. Create another EC2 instance as a proxy server Migrate the SSL certificate to the new instance and configure it to direct connections to the existing EC2 instances
- D. Import the SSL certificate into AWS Certificate Manager (ACM) Create an Application Load Balancer with an HTTPS listener that uses the SSL certificate from ACM

**Answer:** D

#### Explanation:

<https://aws.amazon.com/certificate-manager/>:

"With AWS Certificate Manager, you can quickly request a certificate, deploy it on ACM- integrated AWS resources, such as Elastic Load Balancers, Amazon CloudFront distributions, and APIs on API Gateway, and let AWS Certificate Manager handle certificate renewals. It also enables you to create private certificates for your internal resources and manage the certificate lifecycle centrally."

### NEW QUESTION 140

- (Topic 2)

A company's application is having performance issues. The application is stateful and needs to complete in-memory tasks on Amazon EC2 instances. The company used AWS CloudFormation to deploy infrastructure and used the M5 EC2 Instance family. As traffic increased, the application performance degraded. Users are reporting delays when they attempt to access the application.

Which solution will resolve these issues in the MOST operationally efficient way?

- A. Replace the EC2 instances with T3 EC2 instances that run in an Auto Scaling group
- B. Make the changes by using the AWS Management Console.
- C. Modify the CloudFormation templates to run the EC2 instances in an Auto Scaling group
- D. Increase the desired capacity and the maximum capacity of the Auto Scaling group manually when an increase is necessary
- E. Modify the CloudFormation template
- F. Replace the EC2 instances with R5 EC2 instances
- G. Use Amazon CloudWatch built-in EC2 memory metrics to track the application performance for future capacity planning.
- H. Modify the CloudFormation template
- I. Replace the EC2 instances with R5 EC2 instances
- J. Deploy the Amazon CloudWatch agent on the EC2 instances to generate custom application latency metrics for future capacity planning.

**Answer:** D

#### Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/cloudwatch-memory-metrics-ec2/>

### NEW QUESTION 145

- (Topic 2)

A company runs a high performance computing (HPC) workload on AWS. The workload requires low-latency network performance and high network throughput with tightly coupled node-to-node communication. The Amazon EC2 instances are properly sized for compute and storage capacity, and are launched using default options.

What should a solutions architect propose to improve the performance of the workload?

- A. Choose a cluster placement group while launching Amazon EC2 instances.
- B. Choose dedicated instance tenancy while launching Amazon EC2 instances.
- C. Choose an Elastic Inference accelerator while launching Amazon EC2 instances.
- D. Choose the required capacity reservation while launching Amazon EC2 instances.

**Answer:** A

#### Explanation:

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-resource-ec2-placementgroup.html>

"A cluster placement group is a logical grouping of instances within a single Availability Zone that benefit from low network latency, high network throughput"

### NEW QUESTION 149

- (Topic 2)

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

What should a solutions architect do to meet these requirements?

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

**Answer:** A

**Explanation:**

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

**NEW QUESTION 151**

- (Topic 2)

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

What should a solutions architect do to meet these requirements?

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

**Answer:** A

**Explanation:**

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

**NEW QUESTION 153**

- (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 155**

- (Topic 2)

A company has implemented a self-managed DNS solution on three Amazon EC2 instances behind a Network Load Balancer (NLB) in the us-west-2 Region. Most of the company's users are located in the United States and Europe. The company wants to improve the performance and availability of the solution. The company launches and configures three EC2 instances in the eu-west-1 Region and adds the EC2 instances as targets for a new NLB.

Which solution can the company use to route traffic to all the EC2 instances?

- A. Create an Amazon Route 53 geolocation routing policy to route requests to one of the two NLBs
- B. Create an Amazon CloudFront distribution
- C. Use the Route 53 record as the distribution's origin.
- D. Create a standard accelerator in AWS Global Accelerator
- E. Create endpoint groups in us-west-2 and eu-west-1. Add the two NLBs as endpoints for the endpoint groups.
- F. Attach Elastic IP addresses to the six EC2 instances
- G. Create an Amazon Route 53 geolocation routing policy to route requests to one of the six EC2 instances
- H. Create an Amazon CloudFront distribution
- I. Use the Route 53 record as the distribution's origin.
- J. Replace the two NLBs with two Application Load Balancers (ALBs). Create an Amazon Route 53 latency routing policy to route requests to one of the two ALBs
- K. Create an Amazon CloudFront distribution
- L. Use the Route 53 record as the distribution's origin.

**Answer:** B

**Explanation:**

For standard accelerators, Global Accelerator uses the AWS global network to route traffic to the optimal regional endpoint based on health, client location, and policies that you configure, which increases the availability of your applications. Endpoints for standard accelerators can be Network Load Balancers, Application Load Balancers, Amazon EC2 instances, or Elastic IP addresses that are located in one AWS Region or multiple Regions.

<https://docs.aws.amazon.com/global-accelerator/latest/dg/what-is-global-accelerator.html>

#### NEW QUESTION 159

- (Topic 3)

An ecommerce 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 email 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 164

- (Topic 3)

A company wants to migrate a Windows-based application from on premises to the AWS Cloud. The application has three tiers, a business tier, and a database tier with Microsoft SQL Server. The company wants to use specific features of SQL Server such as native backups and Data Quality Services. The company also needs to share files for process between the tiers.

How should a solution architect design the architecture to meet these requirements?

- A. Host all three on Amazon instance
- B. Use Amazon FSx File Gateway for file sharing between tiers.
- C. Host all three on Amazon EC2 instance
- D. Use Amazon FSx for Windows file sharing between the tiers.
- E. Host the application tier and the business tier on Amazon EC2 instance
- F. Host the database tier on Amazon RD
- G. Use Amazon Elastic File system (Amazon EFS) for file sharing between the tiers.
- H. Host the application tier and the business tier on Amazon EC2 instance
- I. Host the database tier on Amazon RD
- J. Use a Provisioned IOPS SSD (io2) Amazon Elastic Block Store (Amazon EBS) volume for file sharing between the tiers.

**Answer: B**

#### Explanation:

This solution will allow the company to host all three tiers on Amazon EC2 instances while using Amazon FSx for Windows File Server to provide Windows-based file sharing between the tiers. This will allow the company to use specific features of SQL Server, such as native backups and Data Quality Services, while sharing files for processing between the tiers.

#### NEW QUESTION 167

- (Topic 3)

A company has hundreds of Amazon EC2 Linux-based instances in the AWS Cloud. Systems administrators have used shared SSH keys to manage the instances. After a recent audit, the company's security team is mandating the removal of all shared keys. A solutions architect must design a solution that provides secure access to the EC2 instances.

Which solution will meet this requirement with the LEAST amount of administrative overhead?

- A. Use AWS Systems Manager Session Manager to connect to the EC2 instances.
- B. Use AWS Security Token Service (AWS STS) to generate one-time SSH keys on demand.
- C. Allow shared SSH access to a set of bastion instance
- D. Configure all other instances to allow only SSH access from the bastion instances
- E. Use an Amazon Cognito custom authorizer to authenticate user
- F. Invoke an AWS Lambda function to generate a temporary SSH key.

**Answer: A**

#### Explanation:

Session Manager is a fully managed AWS Systems Manager capability. With Session Manager, you can manage your Amazon Elastic Compute Cloud (Amazon EC2) instances, edge devices, on-premises servers, and virtual machines (VMs). You can use either an interactive one-click browser-based shell or the AWS Command Line Interface (AWS CLI). Session Manager provides secure and auditable node management without the need to open inbound ports, maintain bastion hosts, or manage SSH keys. Session Manager also allows you to comply with corporate policies that require controlled access to managed nodes, strict security practices, and fully auditable logs with node access details, while providing end users with simple one-click cross-platform access to your managed nodes.  
<https://docs.aws.amazon.com/systems-manager/latest/userguide/session-manager.html>

#### NEW QUESTION 168

- (Topic 3)

A company's facility has badge readers at every entrance throughout the building. When badges are scanned, the readers send a message over HTTPS to indicate who attempted to access that particular entrance.

A solutions architect must design a system to process these messages from the sensors. The solution must be highly available, and the results must be made available for the company's security team to analyze.

Which system architecture should the solutions architect recommend?

- A. Launch an Amazon EC2 instance to serve as the HTTPS endpoint and to process the messages. Configure the EC2 instance to save the results to an Amazon S3 bucket.

- B. Create an HTTPS endpoint in Amazon API Gateway
- C. Configure the API Gateway endpoint to invoke an AWS Lambda function to process the messages and save the results to an Amazon DynamoDB table.
- D. Use Amazon Route 53 to direct incoming sensor messages to an AWS Lambda function
- E. Configure the Lambda function to process the messages and save the results to an Amazon DynamoDB table.
- F. Create a gateway VPC endpoint for Amazon S3. Configure a Site-to-Site VPN connection from the facility network to the VPC so that sensor data can be written directly to an S3 bucket by way of the VPC endpoint.

**Answer: B**

**Explanation:**

Deploy Amazon API Gateway as an HTTPS endpoint and AWS Lambda to process and save the messages to an Amazon DynamoDB table. This option provides a highly available and scalable solution that can easily handle large amounts of data. It also integrates with other AWS services, making it easier to analyze and visualize the data for the security team.

**NEW QUESTION 169**

- (Topic 3)

A rapidly growing global ecommerce company is hosting its web application on AWS. The web application includes static content and dynamic content. The website stores online transaction processing (OLTP) data in an Amazon RDS database. The website's users are experiencing slow page loads. Which combination of actions should a solutions architect take to resolve this issue? (Select TWO.)

- A. Configure an Amazon Redshift cluster.
- B. Set up an Amazon CloudFront distribution
- C. Host the dynamic web content in Amazon S3
- D. Create a read replica for the RDS DB instance.
- E. Configure a Multi-AZ deployment for the RDS DB instance

**Answer: BD**

**Explanation:**

To resolve the issue of slow page loads for a rapidly growing e-commerce website hosted on AWS, a solutions architect can take the following two actions:

- \* 1. Set up an Amazon CloudFront distribution
- \* 2. Create a read replica for the RDS DB instance

Configuring an Amazon Redshift cluster is not relevant to this issue since Redshift is a data warehousing service and is typically used for the analytical processing of large amounts of data.

Hosting the dynamic web content in Amazon S3 may not necessarily improve performance since S3 is an object storage service, not a web application server. While S3 can be used to host static web content, it may not be suitable for hosting dynamic web content since S3 doesn't support server-side scripting or processing.

Configuring a Multi-AZ deployment for the RDS DB instance will improve high availability but may not necessarily improve performance.

**NEW QUESTION 172**

- (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 173**

- (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 174**

.....

## **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](#)**