

Amazon-Web-Services

Exam Questions SAA-C03

AWS Certified Solutions Architect - Associate (SAA-C03)



NEW QUESTION 1

- (Topic 1)

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

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

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

Answer: A

Explanation:

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

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

NEW QUESTION 2

- (Topic 1)

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

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

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

Answer: B

Explanation:

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

<https://aws.amazon.com/appflow/>

NEW QUESTION 3

- (Topic 1)

A company hosts its web applications in the AWS Cloud. The company configures Elastic Load Balancers to use certificates that are imported into AWS Certificate Manager (ACM). The company's security team must be notified 30 days before the expiration of each certificate.

What should a solutions architect recommend to meet the requirement?

- A. Add a rule in ACM to publish a custom message to an Amazon Simple Notification Service (Amazon SNS) topic every day beginning 30 days before any certificate will expire.
- B. Create an AWS Config rule that checks for certificates that will expire within 30 days
- C. Configure Amazon EventBridge (Amazon CloudWatch Events) to invoke a custom alert by way of Amazon Simple Notification Service (Amazon SNS) when AWS Config reports a noncompliant resource
- D. Use AWS Trusted Advisor to check for certificates that will expire within 30 days
- E. Create an Amazon CloudWatch alarm that is based on Trusted Advisor metrics for check status changes. Configure the alarm to send a custom alert by way of Amazon Simple Notification Service (Amazon SNS)
- F. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to detect any certificates that will expire within 30 days
- G. Configure the rule to invoke an AWS Lambda function
- H. Configure the Lambda function to send a custom alert by way of Amazon Simple Notification Service (Amazon SNS).

Answer: B

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/acm-certificate-expiration/>

NEW QUESTION 4

- (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 5

- (Topic 1)

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

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

Answer: D

Explanation:

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

NEW QUESTION 6

- (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 7

- (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 8

- (Topic 1)

A company's website uses an Amazon EC2 instance store for its catalog of items. The company wants to make sure that the catalog is highly available and that the catalog is stored in a durable location.

What should a solutions architect do to meet these requirements?

- A. Move the catalog to Amazon ElastiCache for Redis.
- B. Deploy a larger EC2 instance with a larger instance store.
- C. Move the catalog from the instance store to Amazon S3 Glacier Deep Archive.
- D. Move the catalog to an Amazon Elastic File System (Amazon EFS) file system.

Answer: D

Explanation:

Moving the catalog to an Amazon Elastic File System (Amazon EFS) file system provides both high availability and durability. Amazon EFS is a fully-managed, highly-available, and durable file system that is built to scale on demand. With Amazon EFS, the catalog data can be stored and accessed from multiple EC2 instances in different availability zones, ensuring high availability. Also, Amazon EFS automatically stores files redundantly within and across multiple availability zones, making it a durable storage option.

NEW QUESTION 9

- (Topic 1)

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

What should a solutions architect do to meet these requirements?

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

Answer: C

Explanation:

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

NEW QUESTION 10

- (Topic 1)

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

What should the solutions architect do to meet this requirement?

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

Answer: A

Explanation:

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

NEW QUESTION 10

- (Topic 1)

A company uses NFS to store large video files in on-premises network attached storage. Each video file ranges in size from 1MB to 500 GB. The total storage is 70 TB and is no longer growing. The company decides to migrate the video files to Amazon S3. The company must migrate the video files as soon as possible while using the least possible network bandwidth.

Which solution will meet these requirements?

- A. Create an S3 bucket Create an IAM role that has permissions to write to the S3 bucket

- B. Use the AWS CLI to copy all files locally to the S3 bucket.
- C. Create an AWS Snowball Edge job
- D. Receive a Snowball Edge device on premise
- E. Use the Snowball Edge client to transfer data to the device
- F. Return the device so that AWS can import the data into Amazon S3.
- G. Deploy an S3 File Gateway on premise
- H. Create a public service endpoint to connect to the S3 File Gateway Create an S3 bucket Create a new NFS file share on the S3 File Gateway Point the new file share to the S3 bucket
- I. Transfer the data from the existing NFS file share to the S3 File Gateway.
- J. Set up an AWS Direct Connect connection between the on-premises network and AWS
- K. Deploy an S3 File Gateway on premise
- L. Create a public virtual interface (VIF) to connect to the S3 File Gateway
- M. Create an S3 bucket
- N. Create a new NFS file share on the S3 File Gateway
- O. Point the new file share to the S3 bucket
- P. Transfer the data from the existing NFS file share to the S3 File Gateway.

Answer: B

Explanation:

The basic difference between Snowball and Snowball Edge is the capacity they provide. Snowball provides a total of 50 TB or 80 TB, out of which 42 TB or 72 TB is available, while Amazon Snowball Edge provides 100 TB, out of which 83 TB is available.

NEW QUESTION 14

- (Topic 1)

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

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

Answer: B

Explanation:

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

NEW QUESTION 18

- (Topic 1)

A company is developing an application that provides order shipping statistics for retrieval by a REST API. The company wants to extract the shipping statistics, organize the data into an easy-to-read HTML format, and send the report to several email addresses at the same time every morning. Which combination of steps should a solutions architect take to meet these requirements? (Choose two.)

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

Answer: BD

Explanation:

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

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

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

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

NEW QUESTION 19

- (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 20

- (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
- G. 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 Cloud¹. 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 power². Users can copy the data to the device using the AWS OpsHub graphical user interface or the Snowball client command line tool³. Users can also create and run Amazon EC2 instances on the device using Amazon Machine Images (AMIs) that are compatible with the sbe1 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 24

- (Topic 1)

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

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

Which solution will meet these requirements?

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

Answer: D

Explanation:

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

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

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

NEW QUESTION 26

- (Topic 1)

A company is preparing to launch a public-facing web application in the AWS Cloud. The architecture consists of Amazon EC2 instances within a VPC behind an Elastic Load Balancer (ELB). A third-party service is used for the DNS. The company's solutions architect must recommend a solution to detect and protect against large-scale DDoS attacks.

Which solution meets these requirements?

- A. Enable Amazon GuardDuty on the account.
- B. Enable Amazon Inspector on the EC2 instances.
- C. Enable AWS Shield and assign Amazon Route 53 to it.

D. Enable AWS Shield Advanced and assign the ELB to it.

Answer: D

Explanation:

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

NEW QUESTION 27

- (Topic 1)

A company's HTTP application is behind a Network Load Balancer (NLB). The NLB's target group is configured to use an Amazon EC2 Auto Scaling group with multiple EC2 instances that run the web service.

The company notices that the NLB is not detecting HTTP errors for the application. These errors require a manual restart of the EC2 instances that run the web service. The company needs to improve the application's availability without writing custom scripts or code.

What should a solutions architect do to meet these requirements?

- A. Enable HTTP health checks on the NL
- B. supplying the URL of the company's application.
- C. Add a cron job to the EC2 instances to check the local application's logs once each minut
- D. If HTTP errors are detected, the application will restart.
- E. Replace the NLB with an Application Load Balance
- F. Enable HTTP health checks by supplying the URL of the company's applicatio
- G. Configure an Auto Scaling action to replace unhealthy instances.
- H. Create an Amazon Cloud Watch alarm that monitors the UnhealthyHostCount metric for the NL
- I. Configure an Auto Scaling action to replace unhealthy instances when the alarm is in the ALARM state.

Answer: C

Explanation:

Application availability: NLB cannot assure the availability of the application. This is because it bases its decisions solely on network and TCP-layer variables and has no awareness of the application at all. Generally, NLB determines availability based on the ability of a server to respond to ICMP ping or to correctly complete the three-way TCP handshake. ALB goes much deeper and is capable of determining availability based on not only a successful HTTP GET of a particular page but also the verification that the content is as was expected based on the input parameters.

NEW QUESTION 30

- (Topic 1)

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

The company has noticed that some insert operations are taking 10 seconds or longer The company has determined that the database storage performance is the problem

Which solution addresses this performance issue?

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

Answer: A

Explanation:

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

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

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

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

NEW QUESTION 33

- (Topic 1)

A company has a website hosted on AWS. The website is behind an Application Load Balancer (ALB) that is configured to handle HTTP and HTTPS separately.

The company wants to forward all requests to the website so that the requests will use HTTPS.

What should a solutions architect do to meet this requirement?

- A. Update the ALB's network ACL to accept only HTTPS traffic
- B. Create a rule that replaces the HTTP in the URL with HTTPS.
- C. Create a listener rule on the ALB to redirect HTTP traffic to HTTPS.
- D. Replace the ALB with a Network Load Balancer configured to use Server Name Indication (SNI).

Answer: C

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/elb-redirect-http-to-https- using-alb/>

How can I redirect HTTP requests to HTTPS using an Application Load Balancer? Last updated: 2020-10-30 I want to redirect HTTP requests to HTTPS using Application Load Balancer listener rules. How can I do this? Resolution Reference: <https://aws.amazon.com/premiumsupport/knowledge-center/elb-redirect-http-to-https- using-alb/>

NEW QUESTION 34

- (Topic 1)

A company needs to store data in Amazon S3 and must prevent the data from being changed. The company wants new objects that are uploaded to Amazon S3 to remain unchangeable for a nonspecific amount of time until the company decides to modify the objects. Only specific users in the company's AWS account can have the ability to delete the objects. What should a solutions architect do to meet these requirements?

- A. Create an S3 Glacier vault Apply a write-once, read-many (WORM) vault lock policy to the objects
- B. Create an S3 bucket with S3 Object Lock enabled Enable versioning Set a retention period of 100 years Use governance mode as the S3 bucket's default retention mode for new objects
- C. Create an S3 bucket Use AWS CloudTrail to track any S3 API events that modify the objects Upon notification, restore the modified objects from any backup versions that the company has
- D. Create an S3 bucket with S3 Object Lock enabled Enable versioning Add a legal hold to the objects Add the s3 PutObjectLegalHold permission to the IAM policies of users who need to delete the objects

Answer: D

Explanation:

"The Object Lock legal hold operation enables you to place a legal hold on an object version. Like setting a retention period, a legal hold prevents an object version from being overwritten or deleted. However, a legal hold doesn't have an associated retention period and remains in effect until removed."

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/batch-ops-legal-hold.html>

NEW QUESTION 39

- (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 43

- (Topic 1)

A company is hosting a web application on AWS using a single Amazon EC2 instance that stores user-uploaded documents in an Amazon EBS volume. For better scalability and availability, the company duplicated the architecture and created a second EC2 instance and EBS volume in another Availability Zone placing both behind an Application Load Balancer. After completing this change, users reported that, each time they refreshed the website, they could see one subset of their documents or the other, but never all of the documents at the same time.

What should a solutions architect propose to ensure users see all of their documents at once?

- A. Copy the data so both EBS volumes contain all the documents.
- B. Configure the Application Load Balancer to direct a user to the server with the documents
- C. Copy the data from both EBS volumes to Amazon EFS Modify the application to save new documents to Amazon EFS
- D. Configure the Application Load Balancer to send the request to both servers Return each document from the correct server.

Answer: C

Explanation:

<https://docs.aws.amazon.com/efs/latest/ug/how-it-works.html#how-it-works-ec2>

NEW QUESTION 45

- (Topic 1)

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

Which solution will satisfy these requirements?

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

Answer: D

NEW QUESTION 49

- (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 53

- (Topic 1)

A company is implementing a shared storage solution for a media application that is hosted in the AWS Cloud. The company needs the ability to use SMB clients to access data. The solution must be fully managed. Which AWS solution meets these requirements?

- A. Create an AWS Storage Gateway volume gateway.
- B. Create a file share that uses the required client protocol. Connect the application server to the file share.
- C. Create an AWS Storage Gateway tape gateway. Configure tapes to use Amazon S3. Connect the application server to the tape gateway.
- D. Create an Amazon EC2 Windows instance. Install and configure a Windows file share role on the instance.
- E. Connect the application server to the file share.
- F. Create an Amazon FSx for Windows File Server file system. Attach the file system to the origin server.
- G. Connect the application server to the file system.

Answer: D

Explanation:

<https://aws.amazon.com/fsx/lustre/>

Amazon FSx has native support for Windows file system features and for the industry-standard Server Message Block (SMB) protocol to access file storage over a network. <https://docs.aws.amazon.com/fsx/latest/WindowsGuide/what-is.html>

NEW QUESTION 55

- (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 57

- (Topic 1)

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

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

Answer: CD

Explanation:

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

NEW QUESTION 60

- (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 61

- (Topic 1)

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

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

- A. Use the CreateQueue API call to create a new queue.
- B. Use the AddPermission API call to add appropriate permissions.
- C. Use the ReceiveMessage API call to set an appropriate wait time.
- D. Use the ChangeMessageVisibility API call to increase the visibility timeout.

Answer: D

Explanation:

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

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

NEW QUESTION 64

- (Topic 1)

A company needs to review its AWS Cloud deployment to ensure that its Amazon S3 buckets do not have unauthorized configuration changes.

What should a solutions architect do to accomplish this goal?

- A. Turn on AWS Config with the appropriate rules.
- B. Turn on AWS Trusted Advisor with the appropriate checks.
- C. Turn on Amazon Inspector with the appropriate assessment template.
- D. Turn on Amazon S3 server access logging.
- E. Configure Amazon EventBridge (Amazon Cloud Watch Events).

Answer: A

Explanation:

To ensure that Amazon S3 buckets do not have unauthorized configuration changes, a solutions architect should turn on AWS Config with the appropriate rules. AWS Config is a service that allows users to audit and assess their AWS resource configurations for compliance with industry standards and internal policies. It provides a detailed view of the resources and their configurations, including information on how the resources are related to each other. By turning on AWS Config with the appropriate rules, users can identify and remediate unauthorized configuration changes to their Amazon S3 buckets.

NEW QUESTION 69

- (Topic 1)

A company is storing sensitive user information in an Amazon S3 bucket. The company wants to provide secure access to this bucket from the application tier running on Amazon EC2 instances inside a VPC.

Which combination of steps should a solutions architect take to accomplish this? (Select TWO.)

- A. Configure a VPC gateway endpoint for Amazon S3 within the VPC.
- B. Create a bucket policy to make the objects in the S3 bucket public.
- C. Create a bucket policy that limits access to only the application tier running in the VPC.
- D. Create an IAM user with an S3 access policy and copy the IAM credentials to the EC2 instance.
- E. Create a NAT instance and have the EC2 instances use the NAT instance to access the S3 bucket.

Answer: AC

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/s3-private-connection-no-authentication/>

NEW QUESTION 70

- (Topic 1)

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

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

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

Answer: A

Explanation:

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

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

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

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

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

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

Customers who have either web or mobile applications with widespread users or applications hosted far away from their S3 bucket can experience long and

variable upload and download speeds over the Internet"

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

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

NEW QUESTION 74

- (Topic 1)

A company is building an application in the AWS Cloud. The application will store data in Amazon S3 buckets in two AWS Regions. The company must use an AWS Key Management Service (AWS KMS) customer managed key to encrypt all data that is stored in the S3 buckets. The data in both S3 buckets must be encrypted and decrypted with the same KMS key. The data and the key must be stored in each of the two Regions.

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

- A. Create an S3 bucket in each Region Configure the S3 buckets to use server-side encryption with Amazon S3 managed encryption keys (SSE-S3) Configure replication between the S3 buckets.
- B. Create a customer managed multi-Region KMS key
- C. Create an S3 bucket in each Region
- D. Configure replication between the S3 buckets
- E. Configure the application to use the KMS key with client-side encryption.
- F. Create a customer managed KMS key and an S3 bucket in each Region Configure the S3 buckets to use server-side encryption with Amazon S3 managed encryption keys (SSE-S3) Configure replication between the S3 buckets.
- G. Create a customer managed KMS key and an S3 bucket in each Region Configure the S3 buckets to use server-side encryption with AWS KMS keys (SSE-KMS) Configure replication between the S3 buckets.

Answer: B

Explanation:

From <https://docs.aws.amazon.com/kms/latest/developerguide/custom-key-store-overview.html>

For most users, the default AWS KMS key store, which is protected by FIPS 140-2 validated cryptographic modules, fulfills their security requirements. There is no need to add an extra layer of maintenance responsibility or a dependency on an additional service. However, you might consider creating a custom key store if your organization has any of the following requirements: Key material cannot be stored in a shared environment. Key material must be subject to a secondary, independent audit path. The HSMs that generate and store key material must be certified at FIPS 140-2 Level 3.

<https://docs.aws.amazon.com/kms/latest/developerguide/custom-key-store-overview.html>

<https://docs.aws.amazon.com/kms/latest/developerguide/multi-region-keys-overview.html>

NEW QUESTION 77

- (Topic 1)

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

What should a solutions architect recommend?

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

Answer: A

NEW QUESTION 82

- (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 83

- (Topic 1)

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

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

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

Answer: C

NEW QUESTION 88

- (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)

NEW QUESTION 89

- (Topic 1)

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

Which solution will meet these requirements?

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

F. Use Amazon Elastic Block Store (Amazon EBS) for storage.

Answer: C

Explanation:

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

NEW QUESTION 94

- (Topic 1)

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

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

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

Answer: C

Explanation:

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

NEW QUESTION 97

- (Topic 1)

A company needs guaranteed Amazon EC2 capacity in three specific Availability Zones in a specific AWS Region for an upcoming event that will last 1 week.

What should the company do to guarantee the EC2 capacity?

- A. Purchase Reserved instances that specify the Region needed
- B. Create an On Demand Capacity Reservation that specifies the Region needed
- C. Purchase Reserved instances that specify the Region and three Availability Zones needed
- D. Create an On-Demand Capacity Reservation that specifies the Region and three Availability Zones needed

Answer: D

Explanation:

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-capacity-reservations.html>

Reserve instances: You will have to pay for the whole term (1 year or 3 years) which is not cost effective

NEW QUESTION 100

- (Topic 1)

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

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

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

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

Answer: CD

Explanation:

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

NEW QUESTION 104

- (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 106

- (Topic 1)

A bicycle sharing company is developing a multi-tier architecture to track the location of its bicycles during peak operating hours. The company wants to use these data points in its existing analytics platform. A solutions architect must determine the most viable multi-tier option to support this architecture. The data points must be accessible from the REST API.

Which action meets these requirements for storing and retrieving location data?

- A. Use Amazon Athena with Amazon S3
- B. Use Amazon API Gateway with AWS Lambda
- C. Use Amazon QuickSight with Amazon Redshift.
- D. Use Amazon API Gateway with Amazon Kinesis Data Analytics

Answer: D

Explanation:

<https://aws.amazon.com/solutions/implementations/aws-streaming-data-solution-for-amazon-kinesis/>

NEW QUESTION 109

- (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 114

- (Topic 1)

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

Which solution will meet these requirements?

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

Answer: A

Explanation:

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

NEW QUESTION 117

- (Topic 1)

A company is migrating applications to AWS. The applications are deployed in different accounts. The company manages the accounts centrally by using AWS Organizations. The company's security team needs a single sign-on (SSO) solution across all the company's accounts. The company must continue managing the

users and groups in its on-premises self-managed Microsoft Active Directory.
Which solution will meet these requirements?

- A. Enable AWS Single Sign-On (AWS SSO) from the AWS SSO console
- B. Create a one-way forest trust or a one-way domain trust to connect the company's self-managed Microsoft Active Directory with AWS SSO by using AWS Directory Service for Microsoft Active Directory.
- C. Enable AWS Single Sign-On (AWS SSO) from the AWS SSO console
- D. Create a two-way forest trust to connect the company's self-managed Microsoft Active Directory with AWS SSO by using AWS Directory Service for Microsoft Active Directory.
- E. Use AWS Directory Service
- F. Create a two-way trust relationship with the company's self-managed Microsoft Active Directory.
- G. Deploy an identity provider (IdP) on-premise
- H. Enable AWS Single Sign-On (AWS SSO) from the AWS SSO console.

Answer: A

Explanation:

To provide single sign-on (SSO) across all the company's accounts while continuing to manage users and groups in its on-premises self-managed Microsoft Active Directory, the solution is to enable AWS Single Sign-On (SSO) from the AWS SSO console and create a one-way forest trust or a one-way domain trust to connect the company's self-managed Microsoft Active Directory with AWS SSO by using AWS Directory Service for Microsoft Active Directory. This solution is described in the AWS documentation

NEW QUESTION 120

- (Topic 1)

A company has an AWS Glue extract, transform, and load (ETL) job that runs every day at the same time. The job processes XML data that is in an Amazon S3 bucket.

New data is added to the S3 bucket every day. A solutions architect notices that AWS Glue is processing all the data during each run.

What should the solutions architect do to prevent AWS Glue from reprocessing old data?

- A. Edit the job to use job bookmarks.
- B. Edit the job to delete data after the data is processed
- C. Edit the job by setting the NumberOfWorkers field to 1.
- D. Use a FindMatches machine learning (ML) transform.

Answer: C

Explanation:

This is the purpose of bookmarks: "AWS Glue tracks data that has already been processed during a previous run of an ETL job by persisting state information from the job run. This persisted state information is called a job bookmark. Job bookmarks help AWS Glue maintain state information and prevent the reprocessing of old data." <https://docs.aws.amazon.com/glue/latest/dg/monitor-continuations.html>

NEW QUESTION 122

- (Topic 1)

A company uses Amazon S3 to store its confidential audit documents. The S3 bucket uses bucket policies to restrict access to audit team IAM user credentials according to the principle of least privilege. Company managers are worried about accidental deletion of documents in the S3 bucket and want a more secure solution.

What should a solutions architect do to secure the audit documents?

- A. Enable the versioning and MFA Delete features on the S3 bucket.
- B. Enable multi-factor authentication (MFA) on the IAM user credentials for each audit team IAM user account.
- C. Add an S3 Lifecycle policy to the audit team's IAM user accounts to deny the s3:DeleteObject action during audit dates.
- D. Use AWS Key Management Service (AWS KMS) to encrypt the S3 bucket and restrict audit team IAM user accounts from accessing the KMS key.

Answer: A

NEW QUESTION 124

- (Topic 1)

A solutions architect is designing a two-tier web application. The application consists of a public-facing web tier hosted on Amazon EC2 in public subnets. The database tier consists of Microsoft SQL Server running on Amazon EC2 in a private subnet. Security is a high priority for the company.

How should security groups be configured in this situation? (Select TWO)

- A. Configure the security group for the web tier to allow inbound traffic on port 443 from 0.0.0.0/0.
- B. Configure the security group for the web tier to allow outbound traffic on port 443 from 0.0.0.0/0.
- C. Configure the security group for the database tier to allow inbound traffic on port 1433 from the security group for the web tier.
- D. Configure the security group for the database tier to allow outbound traffic on ports 443 and 1433 to the security group for the web tier.
- E. Configure the security group for the database tier to allow inbound traffic on ports 443 and 1433 from the security group for the web tier.

Answer: AC

Explanation:

"Security groups create an outbound rule for every inbound rule." Not completely right. Stateful does NOT mean that if you create an inbound (or outbound) rule, it will create an outbound (or inbound) rule. What it does mean is: suppose you create an inbound rule on port 443 for the X IP. When a request enters on port 443 from X IP, it will allow traffic out for that request in the port 443. However, if you look at the outbound rules, there will not be any outbound rule on port 443 unless explicitly create it. In ACLs, which are stateless, you would have to create an inbound rule to allow incoming requests and an outbound rule to allow your application responds to those incoming requests.

https://docs.aws.amazon.com/vpc/latest/userguide/VPC_SecurityGroups.html#SecurityGroupRules

NEW QUESTION 126

- (Topic 1)

A company is designing an application. The application uses an AWS Lambda function to receive information through Amazon API Gateway and to store the information in an Amazon Aurora PostgreSQL database. During the proof-of-concept stage, the company has to increase the Lambda quotas significantly to handle the high volumes of data that the company needs to load into the database. A solutions architect must recommend a new design to improve scalability and minimize the configuration effort. Which solution will meet these requirements?

- A. Refactor the Lambda function code to Apache Tomcat code that runs on Amazon EC2 instance
- B. Connect the database by using native Java Database Connectivity (JDBC) drivers.
- C. Change the platform from Aurora to Amazon DynamoD
- D. Provision a DynamoDB Accelerator (DAX) cluste
- E. Use the DAX client SDK to point the existing DynamoDB API calls at the DAX cluster.
- F. Set up two Lambda function
- G. Configure one function to receive the informatio
- H. Configure the other function to load the information into the databas
- I. Integrate the Lambda functions by using Amazon Simple Notification Service (Amazon SNS).
- J. Set up two Lambda function
- K. Configure one function to receive the informatio
- L. Configure the other function to load the information into the databas
- M. Integrate the Lambda functions by using an Amazon Simple Queue Service (Amazon SQS) queue.

Answer: B

Explanation:

bottlenecks can be avoided with queues (SQS).

NEW QUESTION 128

- (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 for administration.
- 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 133

- (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 135

- (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 neede
- F. Stop the RDS database when it is not in use.
- G. Purchase EC2 Instance Savings Plans to cover two EC2 instance

- H. Use up to three additional EC2 On-Demand Instances as needed
- I. Purchase an RDS Reserved DB Instance.

Answer: C

Explanation:

This solution meets the requirements of a two-tier application that has a variable demand based on the time of day and must be available at all times, while minimizing the overall cost. EC2 Reserved Instances can provide significant savings compared to On-Demand Instances for the baseline level of usage, and they can guarantee capacity reservation when needed. EC2 Spot Instances can provide up to 90% savings compared to On-Demand Instances for any additional capacity that the application needs during peak hours. Spot Instances are suitable for stateless applications that can tolerate interruptions and can be replaced by other instances. Stopping the RDS database when it is not in use can reduce the cost of running the database tier.

Option A is incorrect because using all EC2 Spot Instances can affect the availability of the application if there are not enough spare capacity or if the Spot price exceeds the maximum price. Stopping the RDS database when it is not in use can reduce the cost of running the database tier, but it can also affect the availability of the application. Option B is incorrect because purchasing EC2 Instance Savings Plans to cover five EC2 instances can lock in a fixed amount of compute usage per hour, which may not match the actual usage pattern of the application. Purchasing an RDS Reserved DB Instance can provide savings for the database tier, but it does not allow stopping the database when it is not in use. Option D is incorrect because purchasing EC2 Instance Savings Plans to cover two EC2 instances can lock in a fixed amount of compute usage per hour, which may not match the actual usage pattern of the application. Using up to three additional EC2 On-Demand Instances as needed can incur higher costs than using Spot Instances.

References:

? <https://aws.amazon.com/ec2/pricing/reserved-instances/>

? <https://aws.amazon.com/ec2/spot/>

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

NEW QUESTION 138

- (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 141

- (Topic 2)

A medical records company is hosting an application on Amazon EC2 instances. The application processes customer data files that are stored on Amazon S3. The EC2 instances are hosted in public subnets. The EC2 instances access Amazon S3 over the internet, but they do not require any other network access.

A new requirement mandates that the network traffic for file transfers take a private route and not be sent over the internet.

Which change to the network architecture should a solutions architect recommend to meet this requirement?

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

Answer: C

Explanation:

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

NEW QUESTION 143

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

Answer: B

Explanation:

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

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

Option D is incorrect because Amazon Elastic Block Store (Amazon EBS) provides persistent block storage volumes for use with EC2 instances, but it does not support SMB protocol or attaching multiple instances to the same volume.

References:

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

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

NEW QUESTION 147

- (Topic 2)

A company wants to direct its users to a backup static error page if the company's primary website is unavailable. The primary website's DNS records are hosted in Amazon Route 53. The domain is pointing to an Application Load Balancer (ALB). The company needs a solution that minimizes changes and infrastructure overhead.

Which solution will meet these requirements?

- A. Update the Route 53 records to use a latency routing polic
- B. Add a static error page that is hosted in an Amazon S3 bucket to the records so that the traffic is sent to the most responsive endpoints.
- C. Set up a Route 53 active-passive failover configuratio
- D. Direct traffic to a static error page that is hosted in an Amazon S3 bucket when Route 53 health checks determine that the ALB endpoint is unhealthy.
- E. Set up a Route 53 active-active configuration with the ALB and an Amazon EC2 instance that hosts a static error page as endpoint
- F. Configure Route 53 to send requests to the instance only if the health checks fail for the ALB.
- G. Update the Route 53 records to use a multivalue answer routing polic
- H. Create a health chec
- I. Direct traffic to the website if the health check passe
- J. Direct traffic to a static error page that is hosted in Amazon S3 if the health check does not pass.

Answer: B

Explanation:

This solution meets the requirements of directing users to a backup static error page if the primary website is unavailable, minimizing changes and infrastructure overhead. Route 53 active-passive failover configuration can route traffic to a primary resource when it is healthy or to a secondary resource when the primary resource is unhealthy. Route 53 health checks can monitor the health of the ALB endpoint and trigger the failover when needed. The static error page can be hosted in an S3 bucket that is configured as a website, which is a simple and cost-effective way to serve static content.

Option A is incorrect because using a latency routing policy can route traffic based on the lowest network latency for users, but it does not provide failover functionality. Option C is incorrect because using an active-active configuration with the ALB and an EC2 instance can increase the infrastructure overhead and complexity, and it does not guarantee that the EC2 instance will always be healthy. Option D is incorrect because using a multivalue answer routing policy can return multiple values for a query, but it does not provide failover functionality.

References:

? <https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-policy-failover.html>

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

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

NEW QUESTION 150

- (Topic 2)

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

Which solution will meet these requirements?

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

Answer: B

Explanation:

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

Compliance:

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

Governance:

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

NEW QUESTION 152

- (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 157

- (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 save 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 161

- (Topic 2)

A company wants to use the AWS Cloud to make an existing application highly available and resilient. The current version of the application resides in the company's data center. The application recently experienced data loss after a database server crashed because of an unexpected power outage.

The company needs a solution that avoids any single points of failure. The solution must give the application the ability to scale to meet user demand.

Which solution will meet these requirements?

- A. Deploy the application servers by using Amazon EC2 instances in an Auto Scaling group across multiple Availability Zones
- B. Use an Amazon RDS DB instance in a Multi-AZ configuration.
- C. Deploy the application servers by using Amazon EC2 instances in an Auto Scaling group in a single Availability Zone
- D. Deploy the database on an EC2 instance
- E. Enable EC2 Auto Recovery.
- F. Deploy the application servers by using Amazon EC2 instances in an Auto Scaling group across multiple Availability Zones
- G. Use an Amazon RDS DB instance with a read replica in a single Availability Zone
- H. Promote the read replica to replace the primary DB instance if the primary DB instance fails.
- I. Deploy the application servers by using Amazon EC2 instances in an Auto Scaling group across multiple Availability Zones Deploy the primary and secondary database servers on EC2 instances across multiple Availability Zones Use Amazon Elastic Block Store (Amazon EBS) Multi-Attach to create shared storage between the instances.

Answer: A

Explanation:

Deploy the application servers by using Amazon EC2 instances in an Auto Scaling group across multiple Availability Zones. Use an Amazon RDS DB instance in a Multi-AZ configuration. To make an existing application highly available and resilient while avoiding any single points of failure and giving the application the ability to scale to meet user demand, the best solution would be to deploy the application servers using Amazon EC2 instances in an Auto Scaling group across multiple Availability Zones and use an Amazon RDS DB instance in a Multi-AZ configuration. By using an Amazon RDS DB instance in a Multi-AZ configuration, the

database is automatically replicated across multiple Availability Zones, ensuring that the database is highly available and can withstand the failure of a single Availability Zone. This provides fault tolerance and avoids any single points of failure.

NEW QUESTION 163

- (Topic 2)

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

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

Answer: B

Explanation:

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

NEW QUESTION 166

- (Topic 2)

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

Which solution meets these requirements MOST cost-effectively?

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

Answer: C

Explanation:

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

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

References:

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

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

NEW QUESTION 168

- (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 172

- (Topic 2)

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

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

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

Answer: B

Explanation:

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

NEW QUESTION 173

- (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 crawle
- D. Use Amazon Athena to query the dat
- E. Use S3 policies to limit access.
- F. Create a data lake by using AWS Lake Formatio
- G. Create an AWS Glue JDBC connection to Amazon RD
- H. Register (he S3 bucket in Lake Formatio
- I. Use Lake Formation access controls to limit access.
- J. Create an Amazon Redshift cluste
- K. Schedule an AWS Lambda function to periodically copy data from Amazon S3 and Amazon RDS to Amazon Redshif
- 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 175

- (Topic 2)

A company wants to build a scalable key management Infrastructure to support developers who need to encrypt data in their applications.

What should a solutions architect do to reduce the operational burden?

- A. Use multifactor authentication (MFA) to protect the encryption keys.
- B. Use AWS Key Management Service (AWS KMS) to protect the encryption keys
- C. Use AWS Certificate Manager (ACM) to create, store, and assign the encryption keys
- D. Use an IAM policy to limit the scope of users who have access permissions to protect the encryption keys

Answer: B

Explanation:

<https://aws.amazon.com/kms/faqs/#:~:text=If%20you%20are%20a%20developer%20who%20needs%20to%20digitally,a%20broad%20set%20of%20industry%20and%20regional%20compliance%20regimes.>

NEW QUESTION 178

- (Topic 2)

A company is concerned about the security of its public web application due to recent web attacks. The application uses an Application Load Balancer (ALB). A solutions architect must reduce the risk of DDoS attacks against the application.

What should the solutions architect do to meet this requirement?

- A. Add an Amazon Inspector agent to the ALB.
- B. Configure Amazon Macie to prevent attacks.
- C. Enable AWS Shield Advanced to prevent attacks.
- D. Configure Amazon GuardDuty to monitor the ALB.

Answer: C

Explanation:

AWS Shield Advanced provides expanded DDoS attack protection for your Amazon EC2 instances, Elastic Load Balancing load balancers, CloudFront distributions, Route 53 hosted zones, and AWS Global Accelerator standard accelerators. <https://docs.aws.amazon.com/waf/latest/developerguide/what-is-aws-waf.html>

NEW QUESTION 179

- (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 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 bucke
- B. Create an S3 event notification for the analysis S3 bucke

- C. Configure Lambda and SageMaker Pipelines as destinations of the event notificatio
- D. Configure s3objectCreated: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 bucket
- H. Create an S3 event notification for the analysis S3 bucket
- I. Configure Lambda and SageMaker Pipelines as destinations of the event notificatio
- J. Configure s3objectCreated:Put as the event type.
- K. Configure S3 replication between the S3 bucket
- 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 183

- (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 Firehos
- 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 databas
- E. Send a message that includes the order number to Amazon Simple Notification Service (Amazon SNS). Set the payment service to pollAmazon SN
- F. retrieve the message, and process the order.
- G. Store the order in the databas
- H. Send a message that includes the order number to an Amazon Simple Queue Service (Amazon SQS) FIFO queu
- I. Set the payment service to retrieve the message and process the orde
- 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 187

- (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 AmazonEC2Linux 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 grou
- 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 grou
- D. Create shared file systems across the instances by using Amazon Elastic File System (Amazon EFS)
- E. Use a partition placement grou
- F. Create shared tile systems across the instances by using Amazon Elastic File System (Amazon EFS).
- G. Use a spread placement grou
- 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-ebs-mult-attach>

NEW QUESTION 191

- (Topic 2)

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

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

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

Answer: A

Explanation:

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

NEW QUESTION 196

- (Topic 2)

An entertainment company is using Amazon DynamoDB to store media metadata. The application is read intensive and experiencing delays. The company does not have staff to handle additional operational overhead and needs to improve the performance efficiency of DynamoDB without reconfiguring the application.

What should a solutions architect recommend to meet this requirement?

- A. Use Amazon ElastiCache for Redis.
- B. Use Amazon DynamoDB Accelerator (DAX).
- C. Replicate data by using DynamoDB global tables.
- D. Use Amazon ElastiCache for Memcached with Auto Discovery enabled.

Answer: B

Explanation:

<https://aws.amazon.com/dynamodb/dax/>

NEW QUESTION 198

- (Topic 2)

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

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

- A. Deploy the application stack in a single AWS Region
- B. Use Amazon CloudFront to serve all static and dynamic content by specifying the ALB as an origin.
- C. Deploy the application stack in two AWS Regions
- 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 Regions
- 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 201

- (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 206

- (Topic 2)

A company stores its application logs in an Amazon CloudWatch Logs log group. A new policy requires the company to store all application logs in Amazon OpenSearch Service (Amazon Elasticsearch Service) in near-real time.

Which solution will meet this requirement with the LEAST operational overhead?

- A. Configure a CloudWatch Logs subscription to stream the logs to Amazon OpenSearch Service (Amazon Elasticsearch Service).
- B. Create an AWS Lambda function
- C. Use the log group to invoke the function to write the logs to Amazon OpenSearch Service (Amazon Elasticsearch Service).
- D. Create an Amazon Kinesis Data Firehose delivery stream
- E. Configure the log group as the delivery stream's source
- F. Configure Amazon OpenSearch Service (Amazon Elasticsearch Service) as the delivery stream's destination.
- G. Install and configure Amazon Kinesis Agent on each application server to deliver the logs to Amazon Kinesis Data Stream
- H. Configure Kinesis Data Streams to deliver the logs to Amazon OpenSearch Service (Amazon Elasticsearch Service)

Answer: B

Explanation:

<https://computingforgeeks.com/stream-logs-in-aws-from-cloudwatch-to-elasticsearch/>

NEW QUESTION 207

.....

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](#)