Amazon-Web-Services

Exam Questions SAA-C03

AWS Certified Solutions Architect - Associate (SAA-C03)

NEW QUESTION 1
- (Topic 1)
A company is running a popular social media website. The website gives users the ability to upload images to share with other users. The company wants to make sure that the images do not contain inappropriate content. The company needs a solution that minimizes development effort.
What should a solutions architect do to meet these requirements?

A. Use Amazon Comprehend to detect inappropriate conten
B. Use human review for low- confidence predictions.
C. Use Amazon Rekognition to detect inappropriate conten
D. Use human review for low- confidence predictions.
E. Use Amazon SageMaker to detect inappropriate conten
F. Use ground truth to label low- confidence predictions.
G. Use AWS Fargate to deploy a custom machine learning model to detect inappropriate conten
H. Use ground truth to label low-confidence predictions.

Answer: B

Explanation:
https://docs.aws.amazon.com/rekognition/latest/dg/moderation.html?pg=ln&sec=ft https://docs.aws.amazon.com/rekognition/latest/dg/a2i-rekognition.html

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 ou
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 bucke
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 dat
F. Configure the S3 bucket as the rule's targe
G. Create a second EventBridge (CloudWatch Events) rule to send events when the upload to the S3 bucket is complet
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 instanc
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 observes an increase in Amazon EC2 costs in its most recent bill The billing team notices unwanted vertical scaling of instance types for a couple of EC2 instances A solutions architect needs to create a graph comparing the last 2 months of EC2 costs and perform an in-depth analysis to identify the root cause of the vertical scaling
How should the solutions architect generate the information with the LEAST operational overhead?

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

Answer: B

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

NEW QUESTION 4
- (Topic 1)
A company has an application that provides marketing services to stores. The services are based on previous purchases by store customers. The stores upload transaction data to the company through SFTP, and the data is processed and analyzed to generate new marketing offers. Some of the files can exceed 200 GB in size.
Recently, the company discovered that some of the stores have uploaded files that contain personally identifiable information (PII) that should not have been included. The company wants administrators to be alerted if PII is shared again. The company also wants to automate remediation.
What should a solutions architect do to meet these requirements with the LEAST development effort?

A. Use an Amazon S3 bucket as a secure transfer poin

B. Use Amazon Inspector to scan me objects in the bucke
C. If objects contain PI
D. trigger an S3 Lifecycle policy to remove the objects that contain Pll.
E. Use an Amazon S3 bucket as a secure transfer poin
F. Use Amazon Macie to scan the objects in the bucke
G. If objects contain PI
H. Use Amazon Simple Notification Service (Amazon SNS) to trigger a notification to the administrators to remove the objects mat contain Pll.
I. Implement custom scanning algorithms in an AWS Lambda functio
J. Trigger the function when objects are loaded into the bucke
K. It objects contain RI
L. use Amazon Simple Notification Service (Amazon SNS) to trigger a notification to the administrators to remove the objects that contain Pll.
M. Implement custom scanning algorithms in an AWS Lambda functio
N. Trigger the function when objects are loaded into the bucke
O. If objects contain PI
P. use Amazon Simple Email Service (Amazon STS) to trigger a notification to the administrators and trigger on S3 Lifecycle policy to remove the objects mot contain PII.

Answer: B

Explanation:
To meet the requirements of detecting and alerting the administrators when PII is shared and automating remediation with the least development effort, the best approach would be to use Amazon S3 bucket as a secure transfer point and scan the objects in the bucket with Amazon Macie. Amazon Macie is a fully managed data security and data privacy service that uses machine learning and pattern matching to discover and protect sensitive data stored in Amazon S3. It can be used to classify sensitive data, monitor access to sensitive data, and automate remediation actions.
In this scenario, after uploading the files to the Amazon S3 bucket, the objects can be scanned for PII by Amazon Macie, and if it detects any PII, it can trigger an Amazon Simple Notification Service (SNS) notification to alert the administrators to remove the objects containing PII. This approach requires the least development effort, as Amazon Macie already has pre-built data classification rules that can detect PII in various formats. Hence, option B is the correct answer.
References:
? Amazon Macie User Guide: https://docs.aws.amazon.com/macie/latest/userguide/what-is-macie.html
? AWS Well-Architected Framework - Security Pillar: https://docs.aws.amazon.com/wellarchitected/latest/security-pillar/welcome.html

NEW QUESTION 5
- (Topic 1)
A company has a three-tier web application that is deployed on AWS. The web servers are
deployed in a public subnet in a VPC. The application servers and database servers are deployed in private subnets in the same VPC. The company has deployed a third-party virtual firewall appliance from AWS Marketplace in an inspection VPC. The appliance is configured with an IP interface that can accept IP packets.
A solutions architect needs to Integrate the web application with the appliance to inspect all traffic to the application before the traffic teaches the web server.
Which solution will moot these requirements with the LEAST operational overhead?

A. Create a Network Load Balancer the public subnet of the application's VPC to route the traffic lo the appliance for packet inspection
B. Create an Application Load Balancer in the public subnet of the application's VPC to route the traffic to the appliance for packet inspection
C. Deploy a transit gateway m the inspection VPC Configure route tables to route the incoming pockets through the transit gateway
D. Deploy a Gateway Load Balancer in the inspection VPC Create a Gateway Load Balancer endpoint to receive the incoming packets and forward the packets to the appliance

Answer: D

Explanation:
https://aws.amazon.com/blogs/networking-and-content-delivery/scaling-network-traffic-inspection-using-aws-gateway-load-balancer/

NEW QUESTION 6
- (Topic 1)
A global company hosts its web application on Amazon EC2 instances behind an Application Load Balancer (ALB). The web application has static data and dynamic data. The company stores its static data in an Amazon S3 bucket. The company wants to improve performance and reduce latency for the static data and dynamic data. The company is using its own domain name registered with Amazon Route 53.
What should a solutions architect do to meet these requirements?

A. Create an Amazon CloudFront distribution that has the S3 bucket and the ALB as origins Configure Route 53 to route traffic to the CloudFront distribution.
B. Create an Amazon CloudFront distribution that has the ALB as an origin Create an AWS Global Accelerator standard accelerator that has the S3 bucket as an endpoin
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 7
- (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 bucke
B. Attach the role to the EC2 instances.
C. Create an IAM policy that grants access to the S3 bucke
D. Attach the policy to the EC2 instances.
E. Create an IAM group that grants access to the S3 bucke
F. Attach the group to the EC2 instances.
G. Create an IAM user that grants access to the S3 bucke
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 8
- (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 9
- (Topic 1)
A company's containerized application runs on an Amazon EC2 instance. The application needs to download security certificates before it can communicate with other business applications. The company wants a highly secure solution to encrypt and decrypt the certificates in near real time. The solution also needs to store data in highly available storage after the data is encrypted.
Which solution will meet these requirements with the LEAST operational overhead?

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

Answer: D

NEW QUESTION 10
- (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 10
- (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 tile 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 11
- (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 queu
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 12
- (Topic 1)
A company needs to keep user transaction data in an Amazon DynamoDB table. The company must retain the data for 7 years.
What is the MOST operationally efficient solution that meets these requirements?

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

Answer: C

NEW QUESTION 14
- (Topic 1)
An ecommerce company wants to launch a one-deal-a-day website on AWS. Each day will feature exactly one product on sale for a period of 24 hours. The company wants to be able to handle millions of requests each hour with millisecond latency during peak hours.
Which solution will meet these requirements with the LEAST operational overhead?

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

Answer: D

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

NEW QUESTION 18
- (Topic 1)
A company hosts a data lake on AWS. The data lake consists of data in Amazon S3 and Amazon RDS for PostgreSQL. The company needs a reporting solution that provides data visualization and includes all the data sources within the data lake. Only the company's management team should have full access to all the visualizations. The rest of the company should have only limited access.
Which solution will meet these requirements?

A. Create an analysis in Amazon QuickSigh
B. Connect all the data sources and create new dataset
C. Publish dashboards to visualize the dat
D. Share the dashboards with the appropriate IAM roles.
E. Create an analysis in Amazon OuickSigh
F. Connect all the data sources and create new dataset
G. Publish dashboards to visualize the dat
H. Share the dashboards with the appropriate users and groups.

I. Create an AWS Glue table and crawler for the data in Amazon S3. Create an AWS Glue extract, transform, and load (ETL) job to produce report
J. Publish the reports to Amazon S3. Use S3 bucket policies to limit access to the reports.
K. Create an AWS Glue table and crawler for the data in Amazon S3. Use Amazon Athena Federated Query to access data within Amazon RDS for PoslgreSQ
L. Generate reports by using Amazon Athen
M. Publish the reports to Amazon S3. Use S3 bucket policies to limit access to the reports.

Answer: B

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

NEW QUESTION 22
- (Topic 1)
A company is implementing a shared storage solution for a media application that is hosted m the AWS Cloud The company needs the ability to use SMB clients to access data The solution must he fully managed.
Which AWS solution meets these requirements?

A. Create an AWS Storage Gateway volume gatewa
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 (apes to use Amazon S3 Connect the application server lo the tape gateway
D. Create an Amazon EC2 Windows instance Install and configure a Windows file share role on the instanc
E. Connect the application server to the file share.
F. Create an Amazon FSx for Windows File Server tile system Attach the fie system to the origin serve
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 23
- (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 Gatewa
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 configuratio
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 configuratio
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 25
- (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 30
- (Topic 1)
A company is running an SMB file server in its data center. The file server stores large files that are accessed frequently for the first few days after the files are created. After 7 days the files are rarely accessed.

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

Which solution will meet these requirements?

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

Answer: B

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

NEW QUESTION 34
- (Topic 1)
A solutions architect is designing the cloud architecture for a new application being deployed on AWS. The process should run in parallel while adding and removing application nodes as needed based on the number of jobs to be processed. The processor application is stateless. The solutions architect must ensure that the application is loosely coupled and the job items are durably stored.
Which design should the solutions architect use?

A. Create an Amazon SNS topic to send the jobs that need to be processed Create an Amazon Machine Image (AMI) that consists of the processor application Create a launch configuration that uses the AMI Create an Auto Scaling group using the launch configuration Set the scaling policy for the Auto Scaling group to add and remove nodes based on CPU usage
B. Create an Amazon SQS queue to hold the jobs that need to be processed Create an Amazon Machine image (AMI) that consists of the processor application Create a launch configuration that uses the AM' Create an Auto Scaling group using the launch configuration Set the scaling policy for the Auto Scaling group to add and remove nodes based on network usage
C. Create an Amazon SQS queue to hold the jobs that needs to be processed Create an Amazon Machine image (AMI) that consists of the processor application Create a launch template that uses the AMI Create an Auto Scaling group using the launch template Set the scaling policy for the Auto Scaling group to add and remove nodes based on the number of items in the SQS queue
D. Create an Amazon SNS topic to send the jobs that need to be processed Create an Amazon Machine Image (AMI) that consists of the processor application Create a launch template that uses the AMI Create an Auto Scaling group using the launch template Set the scaling policy for the Auto Scaling group to add and remove nodes based on the number of messages published to the SNS topic

Answer: C

Explanation:
"Create an Amazon SQS queue to hold the jobs that needs to be processed. Create an Amazon EC2 Auto Scaling group for the compute application. Set the scaling policy for the Auto Scaling group to add and remove nodes based on the number of items in the SQS queue"
In this case we need to find a durable and loosely coupled solution for storing jobs. Amazon SQS is ideal for this use case and can be configured to use dynamic scaling based on the number of jobs waiting in the queue.To configure this scaling you can use the backlog per instance metric with the target value being the acceptable backlog per instance to maintain. You can calculate these numbers as follows: Backlog per instance: To calculate your backlog per instance, start with the ApproximateNumberOfMessages queue attribute to determine the length of the SQS queue

NEW QUESTION 35
- (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 loggin
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 39
- (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 Ama2on 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 to 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 43
- (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 bucke
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 Regio
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 Regio
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 Regio
H. Store the data in an Amazon Elastic Block Store (Amazon EBS) volum
I. Once a day take an EBS snapshot and copy it to the centralized Regio
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:
"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 48
- (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 ke
C. Create an S3 bucket in each Regio
D. Configure replication between the S3 bucket
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 m 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 53
- (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 strea
B. Create an Amazon Kinesis Data Firehose delivery stream that uses the Kinesis data stream as a data sourc
C. Use AWS Lambda functions to transform the dat
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 Glu
F. Stop source/destination checking on the EC2 instanc

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 strea
I. Create an Amazon Kinesis Data Firehose delivery stream that uses the Kinesis data stream as a data sourc
J. Use AWS Lambda functions to transform the dat
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 Glu
M. Use AWS Lambda functions to transform the dat
N. Use AWS Glue to send the data to Amazon S3.

Answer: C

NEW QUESTION 56
- (Topic 1)
A company has an Amazon S3 bucket that contains critical data. The company must protect the data from accidental deletion.
Which combination of steps should a solutions architect take to meet these requirements?
(Choose two.)

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

Answer: AB

Explanation:
To protect data in an S3 bucket from accidental deletion, versioning should be enabled, which enables you to preserve, retrieve, and restore every version of every object in an S3 bucket. Additionally, enabling MFA (multi-factor authentication) Delete on the S3 bucket adds an extra layer of protection by requiring an authentication token in addition to the user's access keys to delete objects in the bucket.
Reference:
AWS S3 Versioning documentation: https://docs.aws.amazon.com/AmazonS3/latest/dev/Versioning.html
AWS S3 MFA Delete documentation: https://docs.aws.amazon.com/AmazonS3/latest/dev/UsingMFADelete.html

NEW QUESTION 59
- (Topic 1)
A company has more than 5 TB of file data on Windows file servers that run on premises Users and applications interact with the data each day
The company is moving its Windows workloads to AWS. As the company continues this process, the company requires access to AWS and on-premises file storage with minimum latency The company needs a solution that minimizes operational overhead and requires no significant changes to the existing file access patterns. The company uses an AWS Site-to-Site VPN connection for connectivity to AWS
What should a solutions architect do to meet these requirements?

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

Answer: D

Explanation:
https://docs.aws.amazon.com/filegateway/latest/filefsxw/what-is-file-fsxw.html
To meet the requirements of the company to have access to both AWS and on-premises file storage with minimum latency, a hybrid cloud architecture can be used. One solution is to deploy and configure Amazon FSx for Windows File Server on AWS, which provides fully managed Windows file servers. The on-premises file data can be moved to the FSx File Gateway, which can act as a bridge between on-premises and AWS file storage. The cloud workloads can be configured to use FSx for Windows File Server on AWS, while the on-premises workloads can be configured to use the FSx File Gateway. This solution minimizes operational overhead and requires no significant changes to the existing file access patterns. The connectivity between on-premises and AWS can be established using an AWS Site-to-Site VPN connection.
Reference:
AWS FSx for Windows File Server: https://aws.amazon.com/fsx/windows/ AWS FSx File Gateway: https://aws.amazon.com/fsx/file-gateway/
AWS Site-to-Site VPN: https://aws.amazon.com/vpn/site-to-site-vpn/

NEW QUESTION 62
- (Topic 1)
A company's dynamic website is hosted using on-premises servers in the United States. The company is launching its product in Europe, and it wants to optimize site loading times for new European users. The site's backend must remain in the United States. The product is being launched in a few days, and an immediate solution is needed.
What should the solutions architect recommend?

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

Answer: C

Explanation:
 https://aws.amazon.com/pt/blogs/aws/amazon-cloudfront-support-for- custom-origins/
You can now create a CloudFront distribution using a custom origin. Each distribution will can point to an S3 or to a custom origin. This could be another storage service, or it could be something more interesting and more dynamic, such as an EC2 instance or even an Elastic Load Balancer

NEW QUESTION 64
- (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 tagge
C. Tag those resources manually.
D. Write API calls to check all resources for proper tag allocatio
E. Periodically run the code on an EC2 instance.
F. Write API calls to check all resources for proper tag allocatio
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 66
- (Topic 1)
A company is deploying a new public web application to AWS. The application will run behind an Application Load Balancer (ALB). The application needs to be encrypted at the edge with an SSL/TLS certificate that is issued by an external certificate authority (CA). The certificate must be rotated each year before the certificate expires.
What should a solutions architect do to meet these requirements?

A. Use AWS Certificate Manager (ACM) to issue an SSL/TLS certificat
B. Apply the certificate to the AL
C. Use the managed renewal feature to automatically rotate the certificate.
D. Use AWS Certificate Manager (ACM) to issue an SSL/TLS certificat
E. Import the key material from the certificat
F. Apply the certificate to the AL
G. Use the managed renewal feature to automatically rotate the certificate.
H. Use AWS Certificate Manager (ACM) Private Certificate Authority to issue an SSL/TLS certificate from the root C
I. Apply the certificate to the AL
J. Use the managed renewal feature to automatically rotate the certificate.
K. Use AWS Certificate Manager (ACM) to import an SSL/TLS certificat
L. Apply the certificate to the AL
M. Use Amazon EventBridge (Amazon CloudWatch Events) to send a notification when the certificate is nearing expiratio
N. Rotate the certificate manually.

Answer: D

Explanation:
 https://www.amazonaws.cn/en/certificate-manager/faqs/#Managed_renewal_and_deployment

NEW QUESTION 70
- (Topic 1)
A company is preparing to deploy a new serverless workload. A solutions architect must use the principle of least privilege to configure permissions that will be used to run an AWS Lambda function. An Amazon EventBridge (Amazon CloudWatch Events) rule will invoke the function.
Which solution meets these requirements?

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

Answer: D

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

NEW QUESTION 75
- (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 80
- (Topic 1)
A company recently signed a contract with an AWS Managed Service Provider (MSP) Partner for help with an application migration initiative. A solutions architect needs to share an Amazon Machine Image (AMI) from existing AWS account with the MSP Partner's AWS account. The AMI is backed by Amazon Elastic Block Store (Amazon EBS) and uses a customer managed customer master key (CMK) to encrypt EBS volume snapshots.
What is the MOST secure way for the solutions architect to share the AMI with the MSP Partner's AWS account?

A. Make the encrypted AMI and snapshots publicly availabl
B. Modify the CMK's key policy to allow the MSP Partner's AWS account to use the key
C. Modify the launchPermission property of the AM
D. Share the AMI with the MSP Partner's AWS account onl
E. Modify the CMK's key policy to allow the MSP Partner's AWS account to use the key.
F. Modify the launchPermission property of the AMI Share the AMI with the MSP Partner's AWS account onl
G. Modify the CMK's key policy to trust a new CMK that is owned by the MSP Partner for encryption.
H. Export the AMI from the source account to an Amazon S3 bucket in the MSP Partner's AWS accoun
I. Encrypt the S3 bucket with a CMK that is owned by the MSP Partner Copy and launch the AMI in the MSP Partner's AWS account.

Answer: B

Explanation:
Share the existing KMS key with the MSP external account because it has already been used to encrypt the AMI snapshot.
https://docs.aws.amazon.com/kms/latest/developerguide/key-policy-modifying-external- accounts.html

NEW QUESTION 84
- (Topic 1)
A company runs a photo processing application that needs to frequently upload and download pictures from Amazon S3 buckets that are located in the same AWS Region. A solutions architect has noticed an increased cost in data transfer fees and needs to implement a solution to reduce these costs.
How can the solutions architect meet this requirement?

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

Answer: D

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

NEW QUESTION 87
- (Topic 1)
A solutions architect is developing a multiple-subnet VPC architecture. The solution will consist of six subnets in two Availability Zones. The subnets are defined as public, private and dedicated for databases. Only the Amazon EC2 instances running in the private subnets should be able to access a database.
Which solution meets these requirements?

A. Create a now route table that excludes the route to the public subnets' CIDR block
B. Associate the route table to the database subnets.
C. Create a security group that denies ingress from the security group used by instances in the public subnet
D. Attach the security group to an Amazon RDS DB instance.
E. Create a security group that allows ingress from the security group used by instances in the private subnet
F. Attach the security group to an Amazon RDS DB instance.
G. Create a new peering connection between the public subnets and the private subnet
H. Create a different peering connection between the private subnets and the databasesubnets.

Answer: C

Explanation:
Security groups are stateful. All inbound traffic is blocked by default. If you create an inbound rule allowing traffic in, that traffic is automatically allowed back out again. You cannot block specific IP address using Security groups (instead use Network Access Control Lists).
"You can specify allow rules, but not deny rules." "When you first create a security group, it has no inbound rules. Therefore, no inbound traffic originating from another host to your instance is allowed until you add inbound rules to the security group." Source:
https://docs.aws.amazon.com/vpc/latest/userguide/VPC_SecurityGroups.html#VPCSecurit yGroups

NEW QUESTION 89
- (Topic 1)
A survey company has gathered data for several years from areas m\ the United States. The company hosts the data in an Amazon S3 bucket that is 3 TB m 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 toone 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 91
- (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 consol
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 consol
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 Servic
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 (AWS 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 93
- (Topic 1)
A company has an automobile sales website that stores its listings in a database on Amazon RDS When an automobile is sold the listing needs to be removed from the website and the data must be sent to multiple target systems.
Which design should a solutions architect recommend?

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

Answer: D

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

NEW QUESTION 94
- (Topic 1)
A company has an application that runs on Amazon EC2 instances and uses an Amazon
Aurora database. The EC2 instances connect to the database by using user names and passwords that are stored locally in a file. The company wants to minimize the operational overhead of credential management.
What should a solutions architect do to accomplish this goal?

A. Use AWS Secrets Manage
B. Turn on automatic rotation.
C. Use AWS Systems Manager Parameter Stor
D. Turn on automatic rotation.
E. Create an Amazon S3 bucket lo store objects that are encrypted with an AWS Key
F. Management Service (AWS KMS) encryption ke
G. Migrate the credential file to the S3 bucke
H. Point the application to the S3 bucket.
I. Create an encrypted Amazon Elastic Block Store (Amazon EBS) volume (or each EC2 instanc
J. Attach the new EBS volume to each EC2 instanc
K. Migrate the credential file to the new EBS volum
L. Point the application to the new EBS volume.

Answer: A

Explanation:
 https://aws.amazon.com/cn/blogs/security/how-to-connect-to-aws-secrets-manager-service-within-a-virtual-private-cloud/
https://aws.amazon.com/blogs/security/rotate-amazon-rds-database-credentials-automatically-with-aws-secrets-manager/

NEW QUESTION 96
- (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 97
- (Topic 1)
A company uses AWS Organizations to manage multiple AWS accounts for different departments. The management account has an Amazon S3 bucket that contains project reports. The company wants to limit access to this S3 bucket to only users of accounts within the organization in AWS Organizations.
Which solution meets these requirements with the LEAST amount of operational overhead?

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

Answer: A

Explanation:
https://aws.amazon.com/blogs/security/control-access-to-aws-resources-by-using-the-aws-organization-of-iam-principals/
The aws:PrincipalOrgID global key provides an alternative to listing all the account IDs for all AWS accounts in an organization. For example, the following Amazon S3 bucket policy allows members of any account in the XXX organization to add an object into the
examtopics bucket.
{"Version": "2020-09-10",
"Statement": {
"Sid": "AllowPutObject", "Effect": "Allow",
"Principal": "*",
"Action": "s3:PutObject",
"Resource": "arn:aws:s3:::examtopics/*", "Condition": {"StringEquals":
{"aws:PrincipalOrgID":["XXX"]}}}}
https://docs.aws.amazon.com/IAM/latest/UserGuide/reference_policies_condition- keys.html

NEW QUESTION 101
- (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 102
- (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 gatewa
B. Configure the route table for the public subnets to send traffic to Amazon S3 through the NAT gateway.
C. Configure the security group for the EC2 instances to restrict outbound traffic so that only traffic to the S3 prefix list is permitted.
D. Move the EC2 instances to private subnet
E. Create a VPC endpoint for Amazon S3, and link the endpoint to the route table for the private subnets
F. Remove the internet gateway from the VP
G. Set up an AWS Direct Connect connection, and route traffic to Amazon S3 over the Direct Connect connection.

Answer: C

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

NEW QUESTION 106
- (Topic 2)
A company wants to migrate its on-premises data center to AWS. According to the company's compliance requirements, the company can use only the ap-northeast-3 Region. Company administrators are not permitted to connect VPCs to the internet.
Which solutions will meet these requirements? (Choose two.)

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

Answer: AC

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

NEW QUESTION 108
- (Topic 2)
A company uses a popular content management system (CMS) for its corporate website. However, the required patching and maintenance are burdensome. The company is redesigning its website and wants anew solution. The website will be updated four times a year and does not need to have any dynamic content available. The solution must provide high scalability and enhanced security.
Which combination of changes will meet these requirements with the LEAST operational overhead? (Choose two.)

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

Answer: AD

Explanation:
 A -> We can configure CloudFront to require HTTPS from clients (enhanced security)
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/using- https-viewers-to-cloudfront.html D -> storing static website on S3 provides scalability and less operational overhead, then configuration of Application LB and EC2 instances (hence E is out)

NEW QUESTION 111
- (Topic 2)
A company is building a web-based application running on Amazon EC2 instances in multiple Availability Zones. The web application will provide access to a repository of text documents totaling about 900 TB in size. The company anticipates that the web application will experience periods of high demand. A solutions architect must ensure that the storage component for the text documents can scale to meet the demand of the application at all times. The company is concerned

about the overall cost of the solution.
Which storage solution meets these requirements MOST cost-effectively?

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

Answer: D

Explanation:
Amazon S3 is cheapest and can be accessed from anywhere.

NEW QUESTION 115
- (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 117
- (Topic 2)
A company wants to move its application to a serverless solution. The serverless solution needs to analyze existing and new data by using SL. The company stores the data in an Amazon S3 bucket. The data requires encryption and must be replicated to a different AWS Region.
Which solution will meet these requirements with the LEAST operational overhead?

A. Create a new S3 bucke
B. Load the data into the new S3 bucke
C. Use S3 Cross-Region Replication (CRR) to replicate encrypted objects to an S3 bucket in another Regio
D. Use server-side encryption with AWS KMS multi-Region kays (SSE-KMS). Use Amazon Athena to query the data.
E. Create a new S3 bucke
F. Load the data into the new S3 bucke
G. Use S3 Cross-Region Replication (CRR) to replicate encrypted objects to an S3 bucket in another Regio
H. Use server-side encryption with AWS KMS multi-Region keys (SSE-KMS). Use Amazon RDS to query the data.
I. Load the data into the existing S3 bucke
J. Use S3 Cross-Region Replication (CRR) to replicate encrypted objects to an S3 bucket in another Regio
K. Use server-side encryption with Amazon S3managed encryption keys (SSE-S3). Use Amazon Athena to query the data.
L. Load the data into the existing S3 bucke
M. Use S3 Cross-Region Replication (CRR) to replicate encrypted objects to an S3 bucket in another Regio
N. Use server-side encryption with Amazon S3 managed encryption keys (SSE-S3). Use Amazon RDS to query the data.

Answer: A

Explanation:
This solution meets the requirements of a serverless solution, encryption, replication, and SQL analysis with the least operational overhead. Amazon Athena is a serverless interactive query service that can analyze data in S3 using standard SQL. S3 Cross-Region Replication (CRR) can replicate encrypted objects to an S3 bucket in another Region automatically. Server-side encryption with AWS KMS multi-Region keys (SSE-KMS) can encrypt the data at rest using keys that are replicated across multiple Regions. Creating a new S3 bucket can avoid potential conflicts with existing data or configurations.
Option B is incorrect because Amazon RDS is not a serverless solution and it cannot query data in S3 directly. Option C is incorrect because server-side encryption with Amazon S3 managed encryption keys (SSE-S3) does not use KMS keys and it does not support multi- Region replication. Option D is incorrect because Amazon RDS is not a serverless solution and it cannot query data in S3 directly. It is also incorrect for the same reason as option C. References:
? https://docs.aws.amazon.com/AmazonS3/latest/userguide/replication-walkthrough-4.html
?https://aws.amazon.com/blogs/storage/considering-four-different-replication-options-for-data-in-amazon-s3/
? https://docs.aws.amazon.com/AmazonS3/latest/userguide/UsingEncryption.html
? https://aws.amazon.com/athena/

NEW QUESTION 119
- (Topic 2)
A media company is evaluating the possibility ot moving rts 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 performanc
D. Amazon EFS for durable data storage and Amazon S3 for archival storage
E. Amazon EC2 Instance store for maximum performanc
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 122
- (Topic 2)
A solutions architect is creating a new Amazon CloudFront distribution for an application. Some of the information submitted by users is sensitive. The application uses HTTPS but needs another layer of security. The sensitive information should be protected throughout the entire application stack, and access to the information should be restricted to certain applications.
Which action should the solutions architect take?

A. Configure a CloudFront signed URL.
B. Configure a CloudFront signed cookie.
C. Configure a CloudFront field-level encryption profile.
D. Configure CloudFront and set the Origin Protocol Policy setting to HTTPS Only for the Viewer Protocol Policy.

Answer: C

Explanation:
 https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/field-level- encryption.html
"With Amazon CloudFront, you can enforce secure end-to-end connections to origin servers by using HTTPS. Field-level encryption adds an additional layer of security that lets you protect specific data throughout system processing so that only certain applications can see it."

NEW QUESTION 125
- (Topic 2)
A company has an event-driven application that invokes AWS Lambda functions up to 800 times each minute with varying runtimes. The Lambda functions access data that is stored in an Amazon Aurora MySQL OB cluster. The company is noticing connection timeouts as user activity increases The database shows no signs of being overloaded. CPU. memory, and disk access metrics are all low.
Which solution will resolve this issue with the LEAST operational overhead?

A. Adjust the size of the Aurora MySQL nodes to handle more connection
B. Configure retry logic in the Lambda functions for attempts to connect to the database
C. Set up Amazon ElastiCache tor Redls to cache commonly read items from the databas
D. Configure the Lambda functions to connect to ElastiCache for reads.
E. Add an Aurora Replica as a reader nod
F. Configure the Lambda functions to connect to the reader endpoint of the OB cluster rather than lo the writer endpoint.
G. Use Amazon ROS Proxy to create a prox
H. Set the DB cluster as the target database Configure the Lambda functions lo connect to the proxy rather than to the DB cluster.

Answer: D

Explanation:
 1. database shows no signs of being overloaded. CPU, memory, and disk access metrics are all low==>A and C out. We cannot only add nodes instance or add read replica, because database workload is totally fine, very low. 2. "least operational overhead"==>B out, because b need to configure lambda. 3. ROS proxy: Shares infrequently used connections; High availability with failover; Drives increased efficiency==>proxy can leverage failover to redirect traffic from timeout rds instance to
healthy rds instance. So D is right.

NEW QUESTION 127
- (Topic 2)
A gaming company hosts a browser-based application on AWS. The users of the application consume a large number of videos and images that are stored in Amazon S3. This content is the same for all users.
The application has increased in popularity, and millions of users worldwide are accessing these media files. The company wants to provide the files to the users while reducing the load on the origin.
Which solution meets these requirements MOST cost-effectively?

A. Deploy an AWS Global Accelerator accelerator in front of the web servers.
B. Deploy an Amazon CloudFront web distribution in front of the S3 bucket.
C. Deploy an Amazon ElastiCache for Redis instance in front of the web servers.
D. Deploy an Amazon ElastiCache for Memcached instance in front of the web servers.

Answer: B

Explanation:
ElastiCache, enhances the performance of web applications by quickly retrieving information from fully-managed in-memory data stores. It utilizes Memcached and Redis, and manages to considerably reduce the time your applications would, otherwise, take to read data from disk-based databases. Amazon CloudFront supports dynamic content from HTTP and WebSocket protocols, which are based on the Transmission Control Protocol (TCP) protocol. Common use cases include dynamic API calls, web pages and web applications, as well as an application's static files such as audio and images. It also supports on-demand media streaming over HTTP. AWS Global Accelerator supports both User Datagram Protocol (UDP) and TCP-based protocols. It is commonly used for non- HTTP use cases, such as gaming, IoT and voice over IP. It is also good for HTTP use cases that need static IP addresses or fast regional failover

NEW QUESTION 129
- (Topic 2)
A company runs a stateless web application in production on a group of Amazon EC2 On- Demand Instances behind an Application Load Balancer. The application experiences heavy usage during an 8-hour period each business day. Application usage is moderate and steady overnight Application usage is low during weekends.
The company wants to minimize its EC2 costs without affecting the availability of the application.
Which solution will meet these requirements?

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

Answer: B

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

NEW QUESTION 132
- (Topic 2)
A solutions architect needs to implement a solution to reduce a company's storage costs. All the company's data is in the Amazon S3 Standard storage class. The company must keep all data for at least 25 years. Data from the most recent 2 years must be highly available and immediately retrievable.
Which solution will meet these requirements?

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

Answer: B

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

NEW QUESTION 136
- (Topic 2)
A company has a legacy data processing application that runs on Amazon EC2 instances. Data is processed sequentially, but the order of results does not matter. The application uses a monolithic architecture. The only way that the company can scale the application to meet increased demand is to increase the size of the instances.
The company's developers have decided to rewrite the application to use a microservices architecture on Amazon Elastic Container Service (Amazon ECS).
What should a solutions architect recommend for communication between the microservices?

A. Create an Amazon Simple Queue Service (Amazon SQS) queu
B. Add code to the data producers, and send data to the queu
C. Add code to the data consumers to process data from the queue.
D. Create an Amazon Simple Notification Service (Amazon SNS) topi
E. Add code to the data producers, and publish notifications to the topi
F. Add code to the data consumers to subscribe to the topic.
G. Create an AWS Lambda function to pass message
H. Add code to the data producers to call the Lambda function with a data objec
I. Add code to the data consumers to receive a data object that is passed from the Lambda function.
J. Create an Amazon DynamoDB tabl
K. Enable DynamoDB Stream
L. Add code to the data producers to insert data into the tabl
M. Add code to the data consumers to use the DynamoDB Streams API to detect new table entries and retrieve the data.

Answer: A

Explanation:
Queue has Limited throughput (300 msg/s without batching, 3000 msg/s with batching whereby up-to 10 msg per batch operation; Msg duplicates not allowed in the queue (exactly-once delivery); Msg order is preserved (FIFO); Queue name must end with .fifo

NEW QUESTION 138
- (Topic 2)
A gaming company is designing a highly available architecture. The application runs on a modified Linux kernel and supports only UDP-based traffic. The company needs the front- end tier to provide the best possible user experience. That tier must have low latency, route traffic to the nearest edge location, and provide static IP addresses for entry into the application endpoints.
What should a solutions architect do to meet these requirements?

A. Configure Amazon Route 53 to forward requests to an Application Load Balance
B. Use AWS Lambda for the application in AWS Application Auto Scaling.
C. Configure Amazon CloudFront to forward requests to a Network Load Balance
D. Use AWS Lambda for the application in an AWS Application Auto Scaling group.

E. Configure AWS Global Accelerator to forward requests to a Network Load Balance
F. Use Amazon EC2 instances for the application in an EC2 Auto Scaling group.
G. Configure Amazon API Gateway to forward requests to an Application Load Balance
H. Use Amazon EC2 instances for the application in an EC2 Auto Scaling group.

Answer: C

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

NEW QUESTION 139
- (Topic 2)
A company uses a three-tier web application to provide training to new employees. The application is accessed for only 12 hours every day. The company is using an Amazon RDS for MySQL DB instance to store information and wants to minimize costs.
What should a solutions architect do to meet these requirements?

A. Configure an IAM policy for AWS Systems Manager Session Manage
B. Create an IAM role for the polic
C. Update the trust relationship of the rol
D. Set up automatic start and stop for the DB instance.
E. Create an Amazon ElastiCache for Redis cache cluster that gives users the ability to access the data from the cache when the DB instance is stoppe
F. Invalidate the cache after the DB instance is started.
G. Launch an Amazon EC2 instanc
H. Create an IAM role that grants access to Amazon RD
I. Attach the role to the EC2 instanc
J. Configure a cron job to start and stop the EC2 instance on the desired schedule.
K. Create AWS Lambda functions to start and stop the DB instanc
L. Create Amazon EventBridge (Amazon CloudWatch Events) scheduled rules to invoke the Lambda function
M. Configure the Lambda functions as event targets for the rules

Answer: D

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

NEW QUESTION 143
- (Topic 2)
A company wants to migrate its MySQL database from on premises to AWS. The company recently experienced a database outage that significantly impacted the business. To ensure this does not happen again, the company wants a reliable database solution on AWS that minimizes data loss and stores every transaction on at least two nodes.
Which solution meets these requirements?

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

Answer: B

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

NEW QUESTION 146
- (Topic 3)
A company has migrated an application to Amazon EC2 Linux instances. One of these EC2 instances runs several 1-hour tasks on a schedule. These tasks were written by different teams and have no common programming language. The company is concerned about performance and scalability while these tasks run on a single instance. A solutions architect needs to implement a solution to resolve these concerns.
Which solution will meet these requirements with the LEAST operational overhead?

A. Use AWS Batch to run the tasks as job
B. Schedule the jobs by using Amazon EventBridge (Amazon CloudWatch Events).
C. Convert the EC2 instance to a containe
D. Use AWS App Runner to create the containeron demand to run the tasks as jobs.
E. Copy the tasks into AWS Lambda function
F. Schedule the Lambda functions by using Amazon EventBridge (Amazon CloudWatch Events).
G. Create an Amazon Machine Image (AMI) of the EC2 instance that runs the task

H. Create an Auto Scaling group with the AMI to run multiple copies of the instance.

Answer: A

Explanation:
AWS Batch is a fully managed service that enables users to run batch jobs on AWS. It can handle different types of tasks written in different languages and run them on EC2 instances. It also integrates with Amazon EventBridge (Amazon CloudWatch Events) to schedule jobs based on time or event triggers. This solution will meet the requirements of performance, scalability and low operational overhead12.
* B. Convert the EC2 instance to a container. Use AWS App Runner to create the container on demand to run the tasks as jobs. This solution will not meet the requirement of low operational overhead, as it involves converting the EC2 instance to a container and using AWS App Runner, which is a service that automatically builds and deploys web applications and load balances traffic2. This is not necessary for running batch jobs.
* C. Copy the tasks into AWS Lambda functions. Schedule the Lambda functions by using Amazon EventBridge (Amazon CloudWatch Events). This solution will not meet the requirement of performance, as AWS Lambda has a limit of 15 minutes for execution time and 10 GB for memory allocation3. These limits may not be sufficient for running 1-hour tasks.
* D. Create an Amazon Machine Image (AMI) of the EC2 instance that runs the tasks. Create an Auto Scaling group with the AMI to run multiple copies of the instance. This solution will not meet the requirement of low operational overhead, as it involves creating and maintaining AMIs and Auto Scaling groups, which are additional resources that need to be configured and managed2.
Reference URL: https://docs.aws.amazon.com/whitepapers/latest/aws-overview/compute- services.html

NEW QUESTION 150
- (Topic 3)
A company has hundreds of Amazon EC2 Linux-based instances in the AWS Cloud. Systems administrators have used shared SSH keys to manage the instances After a recent audit, the company's security team is mandating the removal of all shared keys. A solutions architect must design a solution that provides secure access to the EC2 instances.
Which solution will meet this requirement with the LEAST amount of administrative overhead?

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

Answer: A

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

NEW QUESTION 154
- (Topic 3)
A telemarketing company is designing its customer call center functionality on AWS. The company needs a solution that provides multiples speaker recognition and generates transcript files The company wants to query the transcript files to analyze the business patterns The transcript files must be stored for 7 years for auditing piloses.
Which solution will meet these requirements?

A. Use Amazon Recognition for multiple speaker recognitio
B. Store the transcript files in Amazon S3 Use machine teaming models for transcript file analysis
C. Use Amazon Transcribe for multiple speaker recognitio
D. Use Amazon Athena for transcript file analysts
E. Use Amazon Translate lor multiple speaker recognitio
F. Store the transcript files in Amazon Redshift Use SQL queues lor transcript file analysis
G. Use Amazon Recognition for multiple speaker recognitio
H. Store the transcript files in Amazon S3 Use Amazon Textract for transcript file analysis

Answer: B

Explanation:
Amazon Transcribe now supports speaker labeling for streaming transcription. Amazon Transcribe is an automatic speech recognition (ASR) service that makes it easy for you to convert speech-to-text. In live audio transcription, each stream of audio may contain multiple speakers. Now you can conveniently turn on the ability to label speakers, thus helping to identify who is saying what in the output transcript. https://aws.amazon.com/about-aws/whats-new/2020/08/amazon-transcribe-supports- speaker-labeling-streaming-transcription/

NEW QUESTION 158
- (Topic 3)
A solutions architect is designing the architecture for a software demonstration environment The environment will run on Amazon EC2 instances in an Auto Scaling group behind an Application Load Balancer (ALB) The system will experience significant increases in traffic during working hours but Is not required to operate on weekends.
Which combination of actions should the solutions architect take to ensure that the system can scale to meet demand? (Select TWO)

A. Use AWS Auto Scaling to adjust the ALB capacity based on request rate
B. Use AWS Auto Scaling to scale the capacity of the VPC internet gateway
C. Launch the EC2 instances in multiple AWS Regions to distribute the load across Regions
D. Use a target tracking scaling policy to scale the Auto Scaling group based on instance CPU utilization
E. Use scheduled scaling to change the Auto Scaling group minimum, maximum, and desired capacity to zero for weekends Revert to the default values at the start of the week

Answer: DE

Explanation:
https://docs.aws.amazon.com/autoscaling/ec2/userguide/as-scaling-target- tracking.html#target-tracking-choose-metrics
A target tracking scaling policy is a type of dynamic scaling policy that adjusts the capacity of an Auto Scaling group based on a specified metric and a target value1. A target tracking scaling policy can automatically scale out or scale in the Auto Scaling group to keep the actual metric value at or near the target value1. A target tracking scaling policy is suitable for scenarios where the load on the application changes frequently and unpredictably, such as during working hours2.
To meet the requirements of the scenario, the solutions architect should use a target tracking scaling policy to scale the Auto Scaling group based on instance CPU utilization. Instance CPU utilization is a common metric that reflects the demand on the
application1. The solutions architect should specify a target value that represents the ideal average CPU utilization level for the application, such as 50 percent1. Then, the Auto Scaling group will scale out or scale in to maintain that level of CPU utilization.
Scheduled scaling is a type of scaling policy that performs scaling actions based on a date and time3. Scheduled scaling is suitable for scenarios where the load on the application changes periodically and predictably, such as on weekends2.
To meet the requirements of the scenario, the solutions architect should also use scheduled scaling to change the Auto Scaling group minimum, maximum, and desired capacity to zero for weekends. This way, the Auto Scaling group will terminate all instances on weekends when they are not required to operate. The solutions architect should also revert to the default values at the start of the week, so that the Auto Scaling group can resume normal operation.

NEW QUESTION 159
- (Topic 3)
A company needs to migrate a legacy application from an on-premises data center to the AWS Cloud because of hardware capacity constraints. The application runs 24 hours a day. & days a week,. The application database storage continues to grow over time.
What should a solution architect do to meet these requirements MOST cost-affectivity?

A. Migrate the application layer to Amazon FC2 Spot Instances Migrate the data storage layer to Amazon S3.
B. Migrate the application layer to Amazon EC2 Reserved Instances Migrate the data storage layer to Amazon RDS On-Demand Instances.
C. Migrate the application layer to Amazon EC2 Reserved instances Migrate the data storage layer to Amazon Aurora Reserved Instances.
D. Migrate the application layer to Amazon EC2 On Demand Amazon Migrate the data storage layer to Amazon RDS Reserved instances.

Answer: C

Explanation:
https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Aurora.AuroraMySQL. html

NEW QUESTION 161
- (Topic 3)
A company wants to run an in-memory database for a latency-sensitive application that runs on Amazon EC2 instances. The application processes more than 100,000 transactions each minute and requires high network throughput. A solutions architect needs to provide a cost-effective network design that minimizes data transfer charges.
Which solution meets these requirements?

A. Launch all EC2 instances in the same Availability Zone within the same AWS Regio
B. Specify a placement group with cluster strategy when launching EC2 instances.
C. Launch all EC2 instances in different Availability Zones within the same AWS Regio
D. Specify a placement group with partition strategy when launching EC2 instances.
E. Deploy an Auto Scaling group to launch EC2 instances in different Availability Zones based on a network utilization target.
F. Deploy an Auto Scaling group with a step scaling policy to launch EC2 instances in different Availability Zones.

Answer: A

Explanation:
• Launching instances within a single AZ and using a cluster placement group provides the lowest network latency and highest bandwidth between instances. This maximizes performance for an in-memory database and high-throughput application.
• Communications between instances in the same AZ and placement group are free, minimizing data transfer charges. Inter-AZ and public IP traffic can incur charges.
• A cluster placement group enables the instances to be placed close together within the
AZ, allowing the high network throughput required. Partition groups span AZs, reducing bandwidth.
• Auto Scaling across zones could launch instances in AZs that increase data transfer charges. It may reduce network throughput, impacting performance.

NEW QUESTION 165
- (Topic 3)
A company is using Amazon Route 53 latency-based routing to route requests to its UDP- based application for users around the world. The application is hosted on redundant servers in the company's on-premises data centers in the United States. Asia, and Europe. The company's compliance requirements state that the application must be hosted on premises The company wants to improve the performance and availability of the application
What should a solutions architect do to meet these requirements?

A. A Configure three Network Load Balancers (NLBs) in the three AWS Regions to address the on-premises endpoints Create an accelerator by using AWS Global Accelerator, and register the NLBs as its endpoint
B. Provide access to the application by using a CNAME that points to the accelerator DNS
C. Configure three Application Load Balancers (ALBs) in the three AWS Regions to address the on-premises endpoint
D. Create an accelerator by using AWS Global Accelerator and register the ALBs as its endpoints Provide access to the application by using a CNAME that points to the accelerator DNS
E. Configure three Network Load Balancers (NLBs) in the three AWS Regions to address the on-premises endpoints In Route 53. create a latency-based record that points to the three NLB
F. and use it as an origin for an Amazon CloudFront distribution Provide access to the application by using a CNAME that points to the CloudFront DNS
G. Configure three Application Load Balancers (ALBs) in the three AWS Regions to address the on-premises endpoints In Route 53 create a latency-based record that points to the three ALBs and use it as an origin for an Amazon CloudFront distribution- Provide access to the application by using a CNAME that points to the CloudFront DNS

Answer: A

Explanation:
 https://aws.amazon.com/step-functions/#:~:text=AWS%20Step%20Functions%20is%20a,machine%20learning%20(ML)%20pipelines.
"A common use case for AWS Step Functions is a task that requires human intervention (for example, an approval process). Step Functions makes it easy to coordinate the components of distributed applications as a series of steps in a visual workflow called a state machine. You can quickly build and run state machines to execute the steps of your application in a reliable and scalable fashion. (https://aws.amazon.com/pt/blogs/compute/implementing-serverless-manual-approval- steps-in-aws-step-functions-and-amazon-api-gateway/)"

NEW QUESTION 170
- (Topic 3)
A rapidly growing global ecommerce company is hosting its web application on AWS. The web application includes static content and dynamic content. The website stores online transaction processing (OLTP) data in an Amazon RDS database. The website's users are experiencing slow page loads.
Which combination of actions should a solutions architect take to resolve this issue? (Select TWO.)

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

Answer: BD

Explanation:
To resolve the issue of slow page loads for a rapidly growing e-commerce website hosted on AWS, a solutions architect can take the following two actions:
* 1. Set up an Amazon CloudFront distribution
* 2. Create a read replica for the RDS DB instance
Configuring an Amazon Redshift cluster is not relevant to this issue since Redshift is a data warehousing service and is typically used for the analytical processing of large amounts of data.
Hosting the dynamic web content in Amazon S3 may not necessarily improve performance since S3 is an object storage service, not a web application server. While S3 can be used to host static web content, it may not be suitable for hosting dynamic web content since S3 doesn't support server-side scripting or processing.
Configuring a Multi-AZ deployment for the RDS DB instance will improve high availability but may not necessarily improve performance.

NEW QUESTION 171
- (Topic 3)
A developer has an application that uses an AWS Lambda function to upload files to Amazon S3 and needs the required permissions to perform the task The developer already has an IAM user with valid IAM credentials required for Amazon S3
What should a solutions architect do to grant the permissions?

A. Add required IAM permissions in the resource policy of the Lambda function
B. Create a signed request using the existing IAM credentials n the Lambda function
C. Create a new IAM user and use the existing IAM credentials in the Lambda function.
D. Create an IAM execution role with the required permissions and attach the IAM rote to the Lambda function

Answer: D

Explanation:
 To grant the necessary permissions to an AWS Lambda function to upload files to Amazon S3, a solutions architect should create an IAM execution role with the required permissions and attach the IAM role to the Lambda function. This approach follows the principle of least privilege and ensures that the Lambda function can only access the resources it needs to perform its specific task.

NEW QUESTION 176
- (Topic 3)
A company has an application thai runs on several Amazon EC2 instances Each EC2 instance has multiple Amazon Elastic Block Store (Amazon EBS) data volumes attached to it The application's EC2 instance configuration and data need to be backed up nightly The application also needs to be recoverable in a different AWS Region
Which solution will meet these requirements in the MOST operationally efficient way?

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

Answer: B

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

NEW QUESTION 178
- (Topic 3)
A company is experiencing sudden increases in demand. The company needs to provision large Amazon EC2 instances from an Amazon Machine image (AMI) The instances will run m an Auto Scaling group. The company needs a solution that provides minimum initialization latency to meet the demand.
Which solution meets these requirements?

A. Use the aws ec2 register-image command to create an AMI from a snapshot Use AWS Step Functions to replace the AMI in the Auto Scaling group
B. Enable Amazon Elastic Block Store (Amazon EBS) fast snapshot restore on a snapshot Provision an AMI by using the snapshot Replace the AMI m the Auto Scaling group with the new AMI

C. Enable AMI creation and define lifecycle rules in Amazon Data Lifecycle Manager (Amazon DLM) Create an AWS Lambda function that modifies the AMI in the Auto Scaling group
D. Use Amazon EventBridge (Amazon CloudWatch Events) to invoke AWS Backup lifecycle policies that provision AMIs Configure Auto Scaling group capacity limits as an event source in EventBridge

Answer: B

Explanation:
Enabling Amazon Elastic Block Store (Amazon EBS) fast snapshot restore on a snapshot allows you to quickly create a new Amazon Machine Image (AMI) from a snapshot, which can help reduce the initialization latency when provisioning new instances. Once the AMI is provisioned, you can replace the AMI in the Auto Scaling group with the new AMI. This will ensure that new instances are launched from the updated AMI and are able to meet the increased demand quickly.

NEW QUESTION 181
- (Topic 3)
A company wants to deploy a new public web application on AWS The application includes a web server tier that uses Amazon EC2 instances The application also includes a database tier that uses an Amazon RDS for MySQL DB instance
The application must be secure and accessible for global customers that have dynamic IP addresses
How should a solutions architect configure the security groups to meet these requirements'?

A. Configure the security group tor the web servers lo allow inbound traffic on port 443 from 0.0.0. 0/0) Configure the security group for the DB instance to allow inbound traffic on port 3306 from the security group of the web servers
B. Configure the security group for the web servers to allow inbound traffic on port 443 from the IP addresses of the customers Configure the security group for the DB instance lo allow inbound traffic on port 3306 from the security group of the web servers
C. Configure the security group for the web servers to allow inbound traffic on port 443 from the IP addresses of the customers Configure the security group for the DB instance to allow inbound traffic on port 3306 from the IP addresses of the customers
D. Configure the security group for the web servers to allow inbound traffic on port 443 from 0.0.0.0.0 Configure the security group for the DB instance to allow inbound traffic on port 3306 from 0.0.0.0/0)

Answer: A

Explanation:
? Restricting inbound access to the web servers to only port 443, which is used for HTTPS traffic, and allowing access from any IP address (0.0.0.0/0), since the application is public and accessible for global customers.
? Restricting inbound access to the DB instance to only port 3306, which is used for MySQL traffic, and allowing access only from the security group of the web servers, which creates a secure connection between the two tiers and prevents unauthorized access to the database.
? Restricting outbound access to the minimum required for both tiers, which is not specified in the question but can be assumed to be similar to the inbound rules.
References:
? Security groups - Amazon Virtual Private Cloud
? 5 Best Practices for AWS Security Groups - DZone

NEW QUESTION 182
- (Topic 3)
A company selves a dynamic website from a flee! of Amazon EC2 instances behind an Application Load Balancer (ALB) The website needs to support multiple languages to serve customers around the world The website's architecture is running in the us-west-1 Region and is exhibiting high request latency tor users that are located in other parts of the world
The website needs to serve requests quickly and efficiently regardless of a user's location However the company does not want to recreate the existing architecture across multiple Regions
What should a solutions architect do to meet these requirements?

A. Replace the existing architecture with a website that is served from an Amazon S3 bucket Configure an Amazon CloudFront distribution with the S3 bucket as the origin Set the cache behavior settings to cache based on the Accept-Language request header
B. Configure an Amazon CloudFront distribution with the ALB as the origin Set the cache behavior settings to cache based on the Accept-Language request header
C. Create an Amazon API Gateway API that is integrated with the ALB Configure the API to use the HTTP integration type Set up an API Gateway stage to enable the API cache based on the Accept-Language request header
D. Launch an EC2 instance in each additional Region and configure NGINX to act as a cache server for that Region Put all the EC2 instances and the ALB behind an Amazon Route 53 record set with a geolocation routing policy

Answer: B

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

NEW QUESTION 183
- (Topic 3)
A solutions architect observes that a nightly batch processing job is automatically scaled up for 1 hour before the desired Amazon EC2 capacity is reached. The peak capacity is the 'same every night and the batch jobs always start at 1 AM. The solutions architect needs to find a cost-effective solution that will allow for the desired EC2 capacity to be reached quickly and allow the Auto Scaling group to scale down after the batch jobs are complete.
What should the solutions architect do to meet these requirements?

A. Increase the minimum capacity for the Auto Scaling group.
B. Increase the maximum capacity for the Auto Scaling group.
C. Configure scheduled scaling to scale up to the desired compute level.
D. Change the scaling policy to add more EC2 instances during each scaling operation.

Answer: C

Explanation:

By configuring scheduled scaling, the solutions architect can set the Auto Scaling group to automatically scale up to the desired compute level at a specific time (IAM) when the batch job starts and then automatically scale down after the job is complete. This will allow the desired EC2 capacity to be reached quickly and also help in reducing the cost.

NEW QUESTION 186
- (Topic 3)
A company has a regional subscription-based streaming service that runs in a single AWS Region. The architecture consists of web servers and application servers on Amazon EC2 instances. The EC2 instances are in Auto Scaling groups behind Elastic Load Balancers. The architecture includes an Amazon Aurora database cluster that extends across multiple Availability Zones.
The company wants to expand globally and to ensure that its application has minimal downtime.

A. Extend the Auto Scaling groups for the web tier and the application tier to deploy instances in Availability Zones in a second Regio
B. Use an Aurora global database to deploy the database in the primary Region and the second Regio
C. Use Amazon Route 53 health checks with a failover routing policy to the second Region.
D. Deploy the web tier and the application tier to a second Regio
E. Add an Aurora PostgreSQL cross-Region Aurara Replica in the second Regio
F. Use Amazon Route 53 health checks with a failovers routing policy to the second Region, Promote the secondary to primary as needed.
G. Deploy the web tier and the applicatin tier to a second Regio
H. Create an Aurora PostSQL database in the second Regio
I. Use AWS Database Migration Service (AWS DMS) to replicate the primary database to the second Regio
J. Use Amazon Route 53 health checks with a failover routing policy to the second Region.
K. Deploy the web tier and the application tier to a second Regio
L. Use an Amazon Aurora global database to deploy the database in the primary Region and the second Regio
M. Use Amazon Route 53 health checks with a failover routing policy to the second Regio
N. Promote the secondary to primary as needed.

Answer: D

Explanation:
This option is the most efficient because it deploys the web tier and the application tier to a second Region, which provides high availability and redundancy for the application. It also uses an Amazon Aurora global database, which is a feature that allows a single Aurora database to span multiple AWS Regions1. It also deploys the database in the primary Region and the second Region, which provides low latency global reads and fast recovery from a Regional outage. It also uses Amazon Route 53 health checks with a failover routing policy to the second Region, which provides data protection by routing traffic to healthy endpoints in different Regions2. It also promotes the secondary to primary as needed, which provides data consistency by allowing write operations in one of the Regions at a time3. This solution meets the requirement of expanding globally and ensuring that its application has minimal downtime. Option A is less efficient because it extends the Auto Scaling groups for the web tier and the application tier to deploy instances in Availability Zones in a second Region, which could incur higher costs and complexity than deploying them separately. It also uses an Aurora global database to deploy the database in the primary Region and the second Region, which is correct. However, it does not use Amazon Route 53 health checks with a failover routing policy to the second Region, which could result in traffic being routed to unhealthy endpoints. Option B is less efficient because it deploys the web tier and the application tier to a second Region, which is correct. It also adds an Aurora PostgreSQL cross-Region Aurora Replica in the second Region, which provides read scalability across Regions. However, it does not use an Aurora global database, which provides faster replication and recovery than cross-Region replicas. It also uses Amazon Route 53 health checks with a failover routing policy to the second Region, which is correct. However, it does not promote the secondary to primary as needed, which could result in data inconsistency or loss. Option C is less efficient because it deploys the web tier and the application tier to a second Region, which is correct. It also creates an Aurora PostgreSQL database in the second Region, which provides data redundancy across Regions. However, it does not use an Aurora global database or cross-Region replicas, which provide faster replication and recovery than creating separate databases. It also uses AWS Database Migration Service (AWS DMS) to replicate the primary database to the second Region, which provides data migration between different sources and targets. However, it does not use an Aurora global database or cross-Region replicas, which provide faster replication and recovery than using AWS DMS. It also uses Amazon Route 53 health checks with a failover routing policy to the second Region, which is correct.

NEW QUESTION 190
- (Topic 3)
A rapidly growing ecommerce company is running its workloads in a single AWS Region. A solutions architect must create a disaster recovery (DR) strategy that includes a different AWS Region The company wants its database to be up to date in the DR Region with the least possible latency The remaining infrastructure in the DR Region needs to run at reduced capacity and must be able to scale up it necessary
Which solution will meet these requirements with the LOWEST recovery time objective (RTO)?

A. Use an Amazon Aurora global database with a pilot light deployment
B. Use an Amazon Aurora global database with a warm standby deployment
C. Use an Amazon RDS Multi-AZ DB instance with a pilot light deployment
D. Use an Amazon RDS Multi-AZ DB instance with a warm standby deployment

Answer: B

Explanation:
https://docs.aws.amazon.com/whitepapers/latest/disaster-recovery-workloads-on-aws/disaster-recovery-options-in-the-cloud.html

NEW QUESTION 192
- (Topic 3)
A solutions architect needs to design a system to store client case files. The files are core company assets and are important. The number of files will grow over time.
The files must be simultaneously accessible from multiple application servers that run on Amazon EC2 instances. The solution must have built-in redundancy.
Which solution meets these requirements?

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

Answer: A

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

NEW QUESTION 196
- (Topic 3)
A company is launching a new application deployed on an Amazon Elastic Container Service (Amazon ECS) cluster and is using the Fargate launch type tor ECS tasks The company is monitoring CPU and memory usage because it is expecting high traffic to the application upon its launch However the company wants to reduce costs when utilization decreases
What should a solutions architect recommend?

A. Use Amazon EC2 Auto Scaling to scale at certain periods based on previous traffic patterns
B. Use an AWS Lambda function to scale Amazon ECS based on metric breaches that trigger an Amazon CloudWatch alarm
C. Use Amazon EC2 Auto Scaling with simple scaling policies to scale when ECS metric breaches trigger an Amazon CloudWatch alarm
D. Use AWS Application Auto Scaling with target tracking policies to scale when ECS metric breaches trigger an Amazon CloudWatch alarm

Answer: D

Explanation:
https://docs.aws.amazon.com/autoscaling/application/userguide/what-is-application-auto- scaling.html

NEW QUESTION 198
......

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questons 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 FirstTry

* SAA-C03 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

100% Actual & Verified — Instant Download, Please Click
Order The SAA-C03 Practice Test Here