Amazon-Web-Services

Exam Questions SAA-C03

AWS Certified Solutions Architect - Associate (SAA-C03)

NEW QUESTION 1
- (Exam 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 2
- (Exam 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 3
- (Exam Topic 1)
An application development team is designing a microservice that will convert large images to smaller, compressed images. When a user uploads an image through the web interface, the microservice should store the image in an Amazon S3 bucket, process and compress the image with an AWS Lambda function, and store the image in its compressed form in a different S3 bucket.
A solutions architect needs to design a solution that uses durable, stateless components to process the images automatically.
Which combination of actions will meet these requirements? (Choose two.)

A. Create an Amazon Simple Queue Service (Amazon SQS) queue Configure the S3 bucket to send a notification to the SQS queue when an image is uploaded to the S3 bucket
B. Configure the Lambda function to use the Amazon Simple Queue Service (Amazon SQS) queue as the invocation source When the SQS message is successfully processed, delete the message in the queue
C. Configure the Lambda function to monitor the S3 bucket for new uploads When an uploaded image is detected write the file name to a text file in memory and use the text file to keep track of the images that were processed
D. Launch an Amazon EC2 instance to monitor an Amazon Simple Queue Service (Amazon SQS) queue When items are added to the queue log the file name in a text file on the EC2 instance and invoke the Lambda function
E. Configure an Amazon EventBridge (Amazon CloudWatch Events) event to monitor the S3 bucket When an image is uploade
F. send an alert to an Amazon Simple Notification Service (Amazon SNS) topic with the application owner's email address for further processing

Answer: AB

Explanation:
➤ Creating an Amazon Simple Queue Service (SQS) queue and configuring the S3 bucket to send a notification to the SQS queue when an image is uploaded to the S3 bucket will ensure that the Lambda function is triggered in a stateless and durable manner.
➤ Configuring the Lambda function to use the SQS queue as the invocation source, and deleting the message in the queue after it is successfully processed will ensure that the Lambda function processes the image in a stateless and durable manner.
Amazon SQS is a fully managed message queuing service that enables you to decouple and scale microservices, distributed systems, and serverless applications. SQS eliminates the complexity and overhead associated with managing and operating-message oriented middleware, and empowers developers to focus on differentiating work. When new images are uploaded to the S3 bucket, SQS will trigger the Lambda function to process the image and compress it. Once the image is processed, the SQS message is deleted, ensuring that the Lambda function is stateless and durable.

NEW QUESTION 4
- (Exam 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)%20redu "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 5
- (Exam 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 instances.Connect the database by using native Java Database Connectivity (JDBC) drivers.
B. Change the platform from Aurora to Amazon DynamoD
C. Provision a DynamoDB Accelerator (DAX) cluste
D. Use the DAX client SDK to point the existing DynamoDB API calls at the DAX cluster.
E. Set up two Lambda function
F. Configure one function to receive the informatio
G. Configure the other function to load the information into the databas
H. Integrate the Lambda functions by using Amazon Simple Notification Service (Amazon SNS).
I. Set up two Lambda function
J. Configure one function to receive the informatio
K. Configure the other function to load the information into the databas
L. 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 6
- (Exam Topic 1)
A company is preparing to store confidential data in Amazon S3 For compliance reasons the data must be encrypted at rest Encryption key usage must be logged tor auditing purposes. Keys must be rotated every year.
Which solution meets these requirements and «the MOST operationally efferent?

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

Answer: D

Explanation:
https://docs.aws.amazon.com/kms/latest/developerguide/rotate-keys.html
When you enable automatic key rotation for a customer managed key, AWS KMS generates new cryptographic material for the KMS key every year. AWS KMS also saves the KMS key's older cryptographic material in perpetuity so it can be used to decrypt data that the KMS key encrypted.
Key rotation in AWS KMS is a cryptographic best practice that is designed to be transparent and easy to use. AWS KMS supports optional automatic key rotation only for customer managed CMKs. Enable and disable key rotation. Automatic key rotation is disabled by default on customer managed CMKs. When you enable (or re-enable) key rotation, AWS KMS automatically rotates the CMK 365 days after the enable date and every 365 days thereafter.

NEW QUESTION 7

- (Exam Topic 1)
A solutions architect is designing a new hybrid architecture to extend a company s on-premises infrastructure to AWS The company requires a highly available connection with consistent low latency to an AWS Region. The company needs to minimize costs and is willing to accept slower traffic if the primary connection fails.
What should the solutions architect do to meet these requirements?

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

Answer: A

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

NEW QUESTION 8
- (Exam Topic 1)
A company runs a shopping application that uses Amazon DynamoDB to store customer information. In case of data corruption, a solutions architect needs to design a solution that meets a recovery point objective (RPO) of 15 minutes and a recovery time objective (RTO) of 1 hour.
What should the solutions architect recommend to meet these requirements?

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

Answer: B

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

NEW QUESTION 9
- (Exam Topic 1)
A company is planning to use an Amazon DynamoDB table for data storage. The company is concerned about cost optimization. The table will not be used on most mornings. In the evenings, the read and write traffic will often be unpredictable. When traffic spikes occur, they will happen very quickly.
What should a solutions architect recommend?

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

Answer: A

NEW QUESTION 10
- (Exam 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 server.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 10
- (Exam Topic 1)
A company hosts its multi-tier applications on AWS. For compliance, governance, auditing, and security, the company must track configuration changes on its AWS resources and record a history of API calls made to these resources.

What should a solutions architect do to meet these requirements?

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

Answer: B

NEW QUESTION 14
- (Exam Topic 1)
A company hosts a containerized web application on a fleet of on-premises servers that process incoming requests. The number of requests is growing quickly. The on-premises servers cannot handle the increased number of requests. The company wants to move the application to AWS with minimum code changes and minimum development effort.
Which solution will meet these requirements with the LEAST operational overhead?

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

Answer: A

NEW QUESTION 16
- (Exam Topic 1)
A company needs to store data in Amazon S3 and must prevent the data from being changed. The company wants new objects that are uploaded to Amazon S3 to remain unchangeable for a nonspecific amount of time until the company decides to modify the objects. Only specific users in the company's AWS account can have the ability to delete the objects. What should a solutions architect do to meet these requirements?

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

Answer: D

Explanation:
"The Object Lock legal hold operation enables you to place a legal hold on an object version. Like setting a retention period, a legal hold prevents an object version from being overwritten or deleted. However, a legal hold doesn't have an associated retention period and remains in effect until removed."
https://docs.aws.amazon.com/AmazonS3/latest/userguide/batch-ops-legal-hold.html

NEW QUESTION 20
- (Exam 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

NEW QUESTION 24
- (Exam 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 to one of the marketing firm's S3 buckets

Answer: A

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

NEW QUESTION 29
- (Exam Topic 1)
A company wants to improve its ability to clone large amounts of production data into a test environment in the same AWS Region. The data is stored in Amazon EC2 instances on Amazon Elastic Block Store (Amazon EBS) volumes. Modifications to the cloned data must not affect the production environment. The software that accesses this data requires consistently high I/O performance.
A solutions architect needs to minimize the time that is required to clone the production data into the test environment.
Which solution will meet these requirements?

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

Answer: C

NEW QUESTION 32
- (Exam Topic 1)
A company has an on-premises application that generates a large amount of time-sensitive data that is backed up to Amazon S3. The application has grown and there are user complaints about internet bandwidth limitations. A solutions architect needs to design a long-term solution that allows for both timely backups to Amazon S3 and with minimal impact on internet connectivity for internal users.
Which solution meets these requirements?

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

Answer: B

NEW QUESTION 37
- (Exam Topic 1)
A company uses NFS to store large video files in on-premises network attached storage. Each video file ranges in size from 1MB to 500 GB. The total storage is 70 TB and is no longer growing. The company decides to migrate the video files to Amazon S3. The company must migrate the video files as soon as possible while using the least possible network bandwidth.
Which solution will meet these requirements?

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

Answer: B

NEW QUESTION 38
- (Exam Topic 1)
A company wants to reduce the cost of its existing three-tier web architecture. The web, application, and database servers are running on Amazon EC2 instances for the development, test, and production environments. The EC2 instances average 30% CPU utilization during peak hours and 10% CPU utilization during non-peak hours.
The production EC2 instances run 24 hours a day. The development and test EC2 instances run for at least 8 hours each day. The company plans to implement automation to stop the development and test EC2 instances when they are not in use.
Which EC2 instance purchasing solution will meet the company's requirements MOST cost-effectively?

A. Use Spot Instances for the production EC2 instance

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

Answer: B

NEW QUESTION 42
- (Exam Topic 1)
A development team needs to host a website that will be accessed by other teams. The website contents consist of HTML, CSS, client-side JavaScript, and images Which method is the MOST cost-effective for hosting the website?

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

Answer: B

Explanation:
In Static Websites, Web pages are returned by the server which are prebuilt. They use simple languages such as HTML, CSS, or JavaScript.
There is no processing of content on the server (according to the user) in Static Websites. Web pages are returned by the server with no change therefore, static Websites are fast.
There is no interaction with databases.
Also, they are less costly as the host does not need to support server-side processing with different languages.
============
In Dynamic Websites, Web pages are returned by the server which are processed during runtime means they are not prebuilt web pages but they are built during runtime according to the user's demand.
These use server-side scripting languages such as PHP, Node.js, ASP.NET and many more supported by the server.
So, they are slower than static websites but updates and interaction with databases are possible.

NEW QUESTION 47
- (Exam 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 52
- (Exam Topic 1)
A company hosts an application on multiple Amazon EC2 instances The application processes messages from an Amazon SQS queue writes to an Amazon RDS table and deletes the message from the queue Occasional duplicate records are found in the RDS table. The SQS queue does not contain any duplicate messages.
What should a solutions architect do to ensure messages are being processed once only?

A. Use the CreateQueue API call to create a new queue
B. Use the Add Permission API call to add appropriate permissions
C. Use the ReceiveMessage API call to set an appropriate wail time
D. Use the ChangeMessageVisibility APi call to increase the visibility timeout

Answer: D

Explanation:
The visibility timeout begins when Amazon SQS returns a message. During this time, the consumer processes and deletes the message. However, if the consumer fails before deleting the message and your system doesn't call the DeleteMessage action for that message before the visibility timeout expires, the message becomes visible to other consumers and the message is received again. If a message must be received only once, your consumer should delete it within the duration of the visibility timeout. https://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/sqs-visibility-timeout.html
Keyword: SQS queue writes to an Amazon RDS From this, Option D best suite & other Options ruled out [Option A - You can't intruduce one more Queue in the existing one; Option B - only Permission & Option C - Only Retrieves Messages] FIF O queues are designed to never introduce duplicate messages. However, your message producer might introduce duplicates in certain scenarios: for example, if the producer sends a message, does not receive a response, and then resends the same message. Amazon SQS APIs provide deduplication functionality that prevents your message producer from sending duplicates. Any duplicates introduced by the message producer are removed within a 5-minute deduplication interval. For standard queues, you might occasionally receive a duplicate copy of a message (at-least- once delivery). If you use a standard queue, you must design your applications to be idempotent (that is, they must not be affected adversely when processing the same message more than once).

NEW QUESTION 54
- (Exam Topic 1)
A company provides a Voice over Internet Protocol (VoIP) service that uses UDP connections. The service consists of Amazon EC2 instances that run in an Auto Scaling group. The company has deployments across multiple AWS Regions.

The company needs to route users to the Region with the lowest latency. The company also needs automated failover between Regions.
Which solution will meet these requirements?

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

Answer: D

Explanation:
https://aws.amazon.com/global-accelerator/faqs/
HTTP /HTTPS - ALB ; TCP and UDP - NLB; Lowest latency routing and more throughput. Also supports failover, uses Anycast Ip addressing - Global Accelerator Caching at Egde Locations – Cloutfront
WS Global Accelerator automatically checks the health of your applications and routes user traffic only to healthy application endpoints. If the health status changes or you make configuration updates, AWS Global Accelerator reacts instantaneously to route your users to the next available endpoint..

NEW QUESTION 57
- (Exam Topic 1)
A solutions architect is designing a two-tier web application The application consists of a public-facing web tier hosted on Amazon EC2 in public subnets The database tier consists of Microsoft SQL Server running on Amazon EC2 in a private subnet Security is a high priority for the company
How should security groups be configured in this situation? (Select TWO)

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

Answer: AC

Explanation:
"Security groups create an outbound rule for every inbound rule." Not completely right. Statefull does NOT mean that if you create an inbound (or outbound) rule, it will create an outbound (or inbound) rule. What it does mean is: suppose you create an inbound rule on port 443 for the X ip. When a request enters on port 443 from X ip, it will allow traffic out for that request in the port 443. However, if you look at the outbound rules, there will not be any outbound rule on port 443 unless explicitly create it. In ACLs, which are stateless, you would have to create an inbound rule to allow incoming requests and an outbound rule to allow your application responds to those incoming requests.
https://docs.aws.amazon.com/vpc/latest/userguide/VPC_SecurityGroups.html#SecurityGroupRules

NEW QUESTION 62
- (Exam Topic 1)
A hospital recently deployed a RESTful API with Amazon API Gateway and AWS Lambda The hospital uses API Gateway and Lambda to upload reports that are in PDF format and JPEG format The hospital needs to modify the Lambda code to identify protected health information (PHI) in the reports
Which solution will meet these requirements with the LEAST operational overhead?

A. Use existing Python libraries to extract the text from the reports and to identify the PHI from the extracted text.
B. Use Amazon Textract to extract the text from the reports Use Amazon SageMaker to identify the PHI from the extracted text.
C. Use Amazon Textract to extract the text from the reports Use Amazon Comprehend Medical to identify the PHI from the extracted text
D. Use Amazon Rekognition to extract the text from the reports Use Amazon Comprehend Medical to identify the PHI from the extracted text

Answer: C

NEW QUESTION 65
- (Exam 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 68
- (Exam 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 database subnets.

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#VPCSecurityGroups

NEW QUESTION 71
- (Exam Topic 1)
A company is running a business-critical web application on Amazon EC2 instances behind an Application Load Balancer. The EC2 instances are in an Auto Scaling group. The application uses an Amazon Aurora PostgreSQL database that is deployed in a single Availability Zone. The company wants the application to be highly available with minimum downtime and minimum loss of data.
Which solution will meet these requirements with the LEAST operational effort?

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

Answer: B

NEW QUESTION 75
- (Exam Topic 1)
A company is developing a two-tier web application on AWS. The company's developers have deployed the application on an Amazon EC2 instance that connects directly to a backend Amazon RDS database. The company must not hardcode database credentials in the application. The company must also implement a solution to automatically rotate the database credentials on a regular basis.
Which solution will meet these requirements with the LEAST operational overhead?

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

Answer: C

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

NEW QUESTION 79
- (Exam Topic 1)
A company is hosting a static website on Amazon S3 and is using Amazon Route 53 for DNS. The website is experiencing increased demand from around the world. The company must decrease latency for users who access the website.
Which solution meets these requirements MOST cost-effectively?

A. Replicate the S3 bucket that contains the website to all AWS Region
B. Add Route 53 geolocation routing entries.

C. Provision accelerators in AWS Global Accelerato
D. Associate the supplied IP addresses with the S3 bucke
E. Edit the Route 53 entries to point to the IP addresses of the accelerators.
F. Add an Amazon CloudFront distribution in front of the S3 bucke
G. Edit the Route 53 entries to point to the CloudFront distribution.
H. Enable S3 Transfer Acceleration on the bucke
I. Edit the Route 53 entries to point to the new endpoint.

Answer: C

NEW QUESTION 83
- (Exam Topic 1)
An Amazon EC2 administrator created the following policy associated with an IAM group containing several users

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

What is the effect of this policy?

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

Answer: C

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

NEW QUESTION 85
- (Exam 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: A

NEW QUESTION 86
- (Exam Topic 1)
An image-processing company has a web application that users use to upload images. The application uploads the images into an Amazon S3 bucket. The company has set up S3 event notifications to publish the object creation events to an Amazon Simple Queue Service (Amazon SQS) standard queue. The SQS queue serves as the event source for an AWS Lambda function that processes the images and sends the results to users through email.
Users report that they are receiving multiple email messages for every uploaded image. A solutions architect determines that SQS messages are invoking the Lambda function more than once, resulting in multiple email messages.
What should the solutions architect do to resolve this issue with the LEAST operational overhead?

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

Answer: C

NEW QUESTION 87
- (Exam 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: A

NEW QUESTION 92
- (Exam Topic 1)
A company recently migrated a message processing system to AWS. The system receives messages into an ActiveMQ queue running on an Amazon EC2 instance. Messages are processed by a consumer application running on Amazon EC2. The consumer application processes the messages and writes results to a MySQL database funning on Amazon EC2. The company wants this application to be highly available with tow operational complexity
Which architecture otters the HGHEST availability?

A. Add a second ActiveMQ server to another Availably Zone Add an additional consumer EC2 instance in another Availability Zon
B. Replicate the MySQL database to another Availability Zone.
C. Use Amazon MO with active/standby brokers configured across two Availability Zones Add an additional consumer EC2 instance in another Availability Zon
D. Replicate the MySQL database to another Availability Zone.
E. Use Amazon MO with active/standby blotters configured across two Availability Zone
F. Add an additional consumer EC2 instance in another Availability Zon
G. Use Amazon ROS tor MySQL with Multi-AZ enabled.
H. Use Amazon MQ with active/standby brokers configured across two Availability Zones Add an Auto Scaling group for the consumer EC2 instances across two Availability Zone
I. Use Amazon RDS for MySQL with Multi-AZ enabled.

Answer: D

NEW QUESTION 97
- (Exam 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: D

NEW QUESTION 101
- (Exam Topic 1)
A company is migrating a distributed application to AWS The application serves variable workloads The legacy platform consists of a primary server trial coordinates jobs across multiple compute nodes The company wants to modernize the application with a solution that maximizes resiliency and scalability.
How should a solutions architect design the architecture to meet these requirements?

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

Answer: B

NEW QUESTION 106
- (Exam 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 107
- (Exam 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 112
- (Exam 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-priv https://aws.amazon.com/blogs/security/rotate-amazon-rds-database-credentials-automatically-with-aws-secrets

NEW QUESTION 117
- (Exam Topic 1)
A company has a data ingestion workflow that consists the following:

> An Amazon Simple Notification Service (Amazon SNS) topic for notifications about new data deliveries

> An AWS Lambda function to process the data and record metadata

The company observes that the ingestion workflow fails occasionally because of network connectivity issues. When such a failure occurs, the Lambda function does not ingest the corresponding data unless the company manually reruns the job.
Which combination of actions should a solutions architect take to ensure that the Lambda function ingests all data in the future? (Select TWO.)

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

Answer: BE

NEW QUESTION 121
- (Exam Topic 1)
A company's website uses an Amazon EC2 instance store for its catalog of items. The company wants to make sure that the catalog is highly available and that the catalog is stored in a durable location.
What should a solutions architect do to meet these requirements?

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

Answer: A

NEW QUESTION 122
- (Exam 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-p 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 127
- (Exam Topic 1)
A company has an application that generates a large number of files, each approximately 5 MB in size. The files are stored in Amazon S3. Company policy requires the files to be stored for 4 years before they can be
deleted Immediate accessibility is always required as the files contain critical business data that is not easy to reproduce. The files are frequently accessed in the first 30 days of the object creation but are rarely accessed after the first 30 days
Which storage solution is MOST cost-effective?

A. Create an S3 bucket lifecycle policy to move Mm from S3 Standard to S3 Glacier 30 days from object creation Delete the Tiles 4 years after object creation
B. Create an S3 bucket lifecycle policy to move tiles from S3 Standard to S3 One Zone-infrequent Access (S3 One Zone-IA] 30 days from object creatio
C. Delete the fees 4 years after object creation
D. Create an S3 bucket lifecycle policy to move files from S3 Standard-infrequent Access (S3 Standard-IA) 30 from object creatio
E. Delete the ties 4 years after object creation
F. Create an S3 bucket Lifecycle policy to move files from S3 Standard to S3 Standard-Infrequent Access (S3 Standard-IA) 30 days from object creation Move the files to S3 Glacier 4 years after object carton.

Answer: B

Explanation:
https://aws.amazon.com/s3/storage-classes/?trk=66264cd8-3b73-416c-9693-ea7cf4fe846a&sc_channel=ps&s_k

NEW QUESTION 128
- (Exam Topic 1)
A company wants to migrate its on-premises application to AWS. The application produces output files that vary in size from tens of gigabytes to hundreds of terabytes The application data must be stored in a standard file system structure The company wants a solution that scales automatically, is highly available, and requires minimum operational overhead.
Which solution will meet these requirements?

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

Answer: C

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

NEW QUESTION 133
- (Exam Topic 1)
A company has thousands of edge devices that collectively generate 1 TB of status alerts each day. Each alert is approximately 2 KB in size. A solutions architect needs to implement a solution to ingest and store the alerts for future analysis.
The company wants a highly available solution. However, the company needs to minimize costs and does not want to manage additional infrastructure. Ad ditionally, the company wants to keep 14 days of data available for immediate analysis and archive any data older than 14 days.
What is the MOST operationally efficient solution that meets these requirements?

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

Answer: A

Explanation:
https://aws.amazon.com/kinesis/data-firehose/features/?nc=sn&loc=2#:~:text=into%20Amazon%20S3%2C%20

NEW QUESTION 137
- (Exam Topic 2)
A company runs a global web application on Amazon EC2 instances behind an Application Load Balancer The application stores data in Amazon Aurora. The company needs to create a disaster recovery solution and can tolerate up to 30 minutes of downtime and potential data loss. The solution does not need to handle the load when the primary infrastructure is healthy
What should a solutions architect do to meet these requirements?

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

Answer: A

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

NEW QUESTION 142
- (Exam 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: C

NEW QUESTION 146
- (Exam Topic 2)
A company stores its application logs in an Amazon CloudWatch Logs log group. A new policy requires the company to store all application logs in Amazon OpenSearch Service (Amazon Elasticsearch Service) in near-real time.
Which solution will meet this requirement with the LEAST operational overhead?

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

Answer: B

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

NEW QUESTION 151
- (Exam Topic 2)
An online retail company has more than 50 million active customers and receives more than 25,000 orders each day. The company collects purchase data for customers and stores this data in Amazon S3. Additional customer data is stored in Amazon RDS.
The company wants to make all the data available to various teams so that the teams can perform analytics.
The solution must provide the ability to manage fine-grained permissions for the data and must minimize operational overhead.
Which solution will meet these requirements?

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

Answer: D

NEW QUESTION 153
- (Exam Topic 2)
A company is migrating an application from on-premises servers to Amazon EC2 instances. As part of the migration design requirements, a solutions architect must implement infrastructure metric alarms. The company does not need to take action if CPU utilization increases to more than 50% for a short burst of time. However, if the CPU utilization increases to more than 50% and read IOPS on the disk are high at the same time, the company needs to act as soon as possible. The solutions architect also must reduce false alarms.
What should the solutions architect do to meet these requirements?

A. Create Amazon CloudWatch composite alarms where possible.
B. Create Amazon CloudWatch dashboards to visualize the metrics and react to issues quickly.
C. Create Amazon CloudWatch Synthetics canaries to monitor the application and raise an alarm.
D. Create single Amazon CloudWatch metric alarms with multiple metric thresholds where possible.

Answer: A

NEW QUESTION 155
- (Exam Topic 2)
A company wants to direct its users to a backup static error page if the company's primary website is unavailable. The primary website's DNS records are hosted in Amazon Route 53. The domain is pointing to an Application Load Balancer (ALB). The company needs a solution that minimizes changes and infrastructure overhead.
Which solution will meet these requirements?

A. Update the Route 53 records to use a latency routing polic
B. Add a static error page that is hosted in an Amazon S3 bucket to the records so that the traffic is sent to the most responsive endpoints.
C. Set up a Route 53 active-passive failover configuratio
D. Direct traffic to a static error page that is hosted in an Amazon S3 bucket when Route 53 health checks determine that the ALB endpoint is unhealthy.

E. Set up a Route 53 active-active configuration with the ALB and an Amazon EC2 instance that hosts a static error page as endpoint
F. Configure Route 53 to send requests to the instance only if the health checks fail for the ALB.
G. Update the Route 53 records to use a multivalue answer routing polic
H. Create a health chec
I. Direct traffic to the website if the health check passe
J. Direct traffic to a static error page that is hosted in Amazon S3 if the health check does not pass.

Answer: B

NEW QUESTION 157
- (Exam 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

NEW QUESTION 159
- (Exam Topic 2)
A company wants to run applications in containers in the AWS Cloud. These applications are stateless and can tolerate disruptions within the underlying infrastructure. The company needs a solution that minimizes cost and operational overhead.
What should a solutions architect do to meet these requirements?

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

Answer: A

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

NEW QUESTION 160
- (Exam Topic 2)
A solutions architect is designing a customer-facing application for a company. The application's database will have a clearly defined access pattern throughout the year and will have a variable number of reads and writes that depend on the time of year. The company must retain audit records for the database for 7 days. The recovery point objective (RPO) must be less than 5 hours.
Which solution meets these requirements?

A. Use Amazon DynamoDB with auto scaling Use on-demand backups and Amazon DynamoDB Streams
B. Use Amazon Redshif
C. Configure concurrency scalin
D. Activate audit loggin
E. Perform database snapshots every 4 hours.
F. Use Amazon RDS with Provisioned IOPS Activate the database auditing parameter Perform database snapshots every 5 hours
G. Use Amazon Aurora MySQL with auto scalin
H. Activate the database auditing parameter

Answer: B

NEW QUESTION 165
- (Exam Topic 2)
A company has a Windows-based application that must be migrated to AWS. The application requires the use of a shared Windows file system attached to multiple Amazon EC2 Windows instances that are deployed across multiple Availability Zones.
What should a solutions architect do to meet this requirement?

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

Answer: B

NEW QUESTION 169
- (Exam Topic 2)
A company is concerned about the security of its public web application due to recent web attacks. The application uses an Application Load Balancer (ALB). A

solutions architect must reduce the risk of DDoS attacks against the application.
What should the solutions architect do to meet this requirement?

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

Answer: C

NEW QUESTION 171
- (Exam Topic 2)
A company has two applications: a sender application that sends messages with payloads to be processed and a processing application intended to receive the messages with payloads. The company wants to implement an AWS service to handle messages between the two applications. The sender application can send about 1.000 messages each hour. The messages may take up to 2 days to be processed. If the messages fail to process, they must be retained so that they do not impact the processing of any remaining messages.
Which solution meets these requirements and is the MOST operationally efficient?

A. Set up an Amazon EC2 instance running a Redis databas
B. Configure both applications to use the instanc
C. Store, process, and delete the messages, respectively.
D. Use an Amazon Kinesis data stream to receive the messages from the sender applicatio
E. Integrate the processing application with the Kinesis Client Library (KCL).
F. Integrate the sender and processor applications with an Amazon Simple Queue Service (Amazon SQS) queu
G. Configure a dead-letter queue to collect the messages that failed to process.
H. Subscribe the processing application to an Amazon Simple Notification Service (Amazon SNS) topic to receive notifications to proces
I. Integrate the sender application to write to the SNS topic.

Answer: C

Explanation:
https://aws.amazon.com/blogs/compute/building-loosely-coupled-scalable-c-applications-with-amazon-sqs-and-
https://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/sqs-dead-letter-queues.htm

NEW QUESTION 172
- (Exam Topic 2)
A solutions architect needs to help a company optimize the cost of running an application on AWS. The application will use Amazon EC2 instances, AWS Fargate, and AWS Lambda for compute within the architecture.
The EC2 instances will run the data ingestion layer of the application. EC2 usage will be sporadic and unpredictable. Workloads that run on EC2 instances can be interrupted at any time. The application front end will run on Fargate, and Lambda will serve the API layer. The front-end utilization and API layer utilization will be predictable over the course of the next year.
Which combination of purchasing options will provide the MOST cost-effective solution for hosting this application? (Choose two.)

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

Answer: AC

NEW QUESTION 176
- (Exam 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 178
- (Exam Topic 2)
A company is building a containerized application on premises and decides to move the application to AWS. The application will have thousands of users soon after li is deployed. The company Is unsure how to manage the deployment of containers at scale. The company needs to deploy the containerized application in a highly available architecture that minimizes operational overhead.
Which solution will meet these requirements?

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

Answer: A

NEW QUESTION 183
- (Exam Topic 2)
A company is running several business applications in three separate VPCs within me us-east-1 Region. The applications must be able to communicate between VPCs. The applications also must be able to consistently send hundreds to gigabytes of data each day to a latency-sensitive application that runs in a single on-premises data center.
A solutions architect needs to design a network connectivity solution that maximizes cost-effectiveness Which solution moots those requirements?

A. Configure three AWS Site-to-Site VPN connections from the data center to AWS Establish connectivity by configuring one VPN connection for each VPC
B. Launch a third-party virtual network appliance in each VPC Establish an iPsec VPN tunnel between the Data center and each virtual appliance
C. Set up three AWS Direct Connect connections from the data center to a Direct Connect gateway inus-east-1 Establish connectivity by configuring each VPC to use one of the Direct Connect connections
D. Set up one AWS Direct Connect connection from the data center to AW
E. Create a transit gateway, and attach each VPC to the transit gatewa
F. Establish connectivity between the Direct Connect connection and the transit gateway.

Answer: D

Explanation:
https://docs.aws.amazon.com/whitepapers/latest/aws-vpc-connectivity-options/aws-direct-connect-aws-transit-g

NEW QUESTION 186
- (Exam Topic 2)
A company hosts a website analytics application on a single Amazon EC2 On-Demand Instance. The analytics software is written in PHP and uses a MySQL database. The analytics software, the web server that provides PHP, and the database server are all hosted on the EC2 instance. The application is showing signs of performance degradation during busy times and is presenting 5xx errors. The company needs to make the application scale seamlessly.
Which solution will meet these requirements MOST cost-effectively?

A. Migrate the database to an Amazon RDS for MySQL DB instanc
B. Create an AMI of the web applicatio
C. Use the AMI to launch a second EC2 On-Demand Instanc
D. Use an Application Load Balancer to distribute the load to each EC2 instance.
E. Migrate the database to an Amazon RDS for MySQL DB instanc
F. Create an AMI of the web applicatio
G. Use the AMI to launch a second EC2 On-Demand Instanc
H. Use Amazon Route 53 weighted routing to distribute the load across the two EC2 instances.
I. Migrate the database to an Amazon Aurora MySQL DB instanc
J. Create an AWS Lambda function to stop the EC2 instance and change the instance typ
K. Create an Amazon CloudWatch alarm to invoke the Lambda function when CPU utilization surpasses 75%.
L. Migrate the database to an Amazon Aurora MySQL DB instanc
M. Create an AMI of the web application.Apply the AMI to a launch templat
N. Create an Auto Scaling group with the launch template Configure the launch template to use a Spot Flee
O. Attach an Application Load Balancer to the Auto Scaling group.

Answer: D

NEW QUESTION 187
- (Exam Topic 2)
A company's website provides users with downloadable historical performance reports. The website needs a solution that will scale to meet the company's website demands globally. The solution should be
cost-effective, limit the provisioning of infrastructure resources, and provide the fastest possible response time.
Which combination should a solutions architect recommend to meet these requirements?

A. Amazon CloudFront and Amazon S3
B. AWS Lambda and Amazon DynamoDB
C. Application Load Balancer with Amazon EC2 Auto Scaling
D. Amazon Route 53 with internal Application Load Balancers

Answer: A

Explanation:
Cloudfront for rapid response and s3 to minimize infrastructure.

NEW QUESTION 192
- (Exam Topic 2)
A gaming company has a web application that displays scores. The application runs on Amazon EC2 instances behind an Application Load Balancer. The application stores data in an Amazon RDS for MySQL database. Users are starting to experience long delays and interruptions that are caused by database read performance. The company wants to improve the user experience while minimizing changes to the application's architecture.
What should a solutions architect do to meet these requirements?

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

Answer: A

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

NEW QUESTION 197
- (Exam Topic 2)
A global company is using Amazon API Gateway to design REST APIs for its loyalty club users in the
us-east-1 Region and the ap-southeast-2 Region. A solutions architect must design a solution to protect these API Gateway managed REST APIs across multiple accounts from SQL injection and cross-site scripting attacks.
Which solution will meet these requirements with the LEAST amount of administrative effort?

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

Answer: A

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

NEW QUESTION 199
- (Exam Topic 2)
An entertainment company is using Amazon DynamoDB to store media metadata. The application is read intensive and experiencing delays. The company does not have staff to handle additional operational overhead and needs to improve the performance efficiency of DynamoDB without reconfiguring the application.
What should a solutions architect recommend to meet this requirement?

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

Answer: B

Explanation:
https://aws.amazon.com/dynamodb/dax/

NEW QUESTION 200
- (Exam Topic 2)
A reporting team receives files each day in an Amazon S3 bucket. The reporting team manually reviews and copies the files from this initial S3 bucket to an analysis S3 bucket each day at the same time to use with Amazon QuickSight. Additional teams are starting to send more files in larger sizes to the initial S3 bucket.
The reporting team wants to move the files automatically analysis S3 bucket as the files enter the initial S3 bucket. The reporting team also wants to use AWS Lambda functions to run pattern-matching code on the copied data. In addition, the reporting team wants to send the data files to a pipeline in Amazon SageMaker Pipelines.
What should a solutions architect do to meet these requirements with the LEAST operational overhead?

A. Create a Lambda function to copy the files to the analysis S3 bucke
B. Create an S3 event notification for the analysis S3 bucke
C. Configure Lambda and SageMaker Pipelines as destinations of the event notificatio
D. Configure s3ObjectCreated:Put as the event type.
E. Create a Lambda function to copy the files to the analysis S3 bucke

F. Configure the analysis S3 bucket to send event notifications to Amazon EventBridge (Amazon CloudWatch Events). Configure an ObjectCreated rule in EventBridge (CloudWatch Events). Configure Lambda and SageMaker Pipelines as targets for the rule.

G. Configure S3 replication between the S3 bucket

H. Create an S3 event notification for the analysis S3 bucke

I. Configure Lambda and SageMaker Pipelines as destinations of the event notificatio

J. Configure s3ObjectCreated:Put as the event type.

K. Configure S3 replication between the S3 bucket

L. Configure the analysis S3 bucket to send event notifications to Amazon EventBridge (Amazon CloudWatch Events). Configure an ObjectCreated rule in EventBridge (CloudWatch Events). Configure Lambda and SageMaker Pipelines as targets for the rule.

Answer: A

NEW QUESTION 201

- (Exam Topic 2)

A company wants to measure the effectiveness of its recent marketing campaigns. The company performs batch processing on csv files of sales data and stores the results «i an Amazon S3 bucket once every hour. The S3 bi petabytes of objects. The company runs one-time queries in Amazon Athena to determine which products are most popular on a particular date for a particular region Queries sometimes fail or take longer than expected to finish.

Which actions should a solutions architect take to improve the query performance and reliability? (Select TWO.)

A. Reduce the S3 object sizes to less than 126 MB

B. Partition the data by date and region n Amazon S3

C. Store the files as large, single objects in Amazon S3.

D. Use Amazon Kinosis Data Analytics to run the Queries as pan of the batch processing operation

E. Use an AWS duo extract, transform, and load (ETL) process to convert the csv files into Apache Parquet format.

Answer: CE

NEW QUESTION 205

- (Exam 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 AmazonEC2 instances for the application in an EC2 Auto Scaling group.

Answer: C

NEW QUESTION 209

- (Exam Topic 2)

A company needs to retain application logs files for a critical application for 10 years. The application team regularly accesses logs from the past month for troubleshooting, but logs older than 1 month are rarely accessed. The application generates more than 10 TB of logs per month.

Which storage option meets these requirements MOST cost-effectively?

A. Store the logs in Amazon S3 Use AWS Backup lo move logs more than 1 month old to S3 Glacier Deep Archive

B. Store the logs in Amazon S3 Use S3 Lifecycle policies to move logs more than 1 month old to S3 Glacier Deep Archive

C. Store the logs in Amazon CloudWatch Logs Use AWS Backup to move logs more then 1 month old toS3 Glacier Deep Archive

D. Store the logs in Amazon CloudWatch Logs Use Amazon S3 Lifecycle policies to move logs more than 1 month old to S3 Glacier Deep Archive

Answer: B

Explanation:

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

NEW QUESTION 213

- (Exam 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 poll Amazon 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 queue

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 214
- (Exam Topic 2)
A company runs an application using Amazon ECS. The application creates esi/ed versions of an original image and then makes Amazon S3 API calls to store the resized images in Amazon S3.
How can a solutions architect ensure that the application has permission to access Amazon S3?

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

Answer: B

NEW QUESTION 216
- (Exam Topic 2)
A company hosts a two-tier application on Amazon EC2 instances and Amazon RDS. The application's demand varies based on the time of day. The load is minimal after work hours and on weekends. The EC2 instances run in an EC2 Auto Scaling group that is configured with a minimum of two instances and a maximum of five instances. The application must be available at all times, but the company is concerned about overall cost.
Which solution meets the availability requirement MOST cost-effectively?

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

Answer: D

NEW QUESTION 218
- (Exam Topic 2)
A company produces batch data that comes from different databases. The company also produces live stream data from network sensors and application APIs. The company needs to consolidate all the data into one place for business analytics. The company needs to process the incoming data and then stage the data in different Amazon S3 buckets. Teams will later run one-time queries and import the data into a business intelligence tool to show key performance indicators (KPIs).
Which combination of steps will meet these requirements with the LEAST operational overhead? (Choose two.)

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

Answer: CE

NEW QUESTION 221
- (Exam 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

NEW QUESTION 226
- (Exam Topic 3)
An ecommerce company needs to run a scheduled daily job to aggregate and filler sales records for analytics. The company stores the sales records in an Amazon S3 bucket. Each object can be up to 10 G6 in size Based on the number of sales events, the job can take up to an hour to complete. The CPU and memory usage of the fob are constant and are known in advance.
A solutions architect needs to minimize the amount of operational effort that is needed for the job to run. Which solution meets these requirements?

A. Create an AWS Lambda function that has an Amazon EventBridge notification Schedule the EventBridge event to run once a day
B. Create an AWS Lambda function Create an Amazon API Gateway HTTP API, and integrate the API with the function Create an Amazon EventBridge scheduled avert that calls the API and invokes the function.
C. Create an Amazon Elastic Container Service (Amazon ECS) duster with an AWS Fargate launch type.Create an Amazon EventBridge scheduled event that launches an ECS task on the cluster to run the job.
D. Create an Amazon Elastic Container Service (Amazon ECS) duster with an Amazon EC2 launch type and an Auto Scaling group with at least one EC2 instanc
E. Create an Amazon EventBridge scheduled event that launches an ECS task on the duster to run the job.

Answer: C

NEW QUESTION 229
- (Exam Topic 3)
A company needs to export its database once a day to Amazon S3 for other teams to access. The exported object size vanes between 2 GB and 5 GB. The S3 access pattern for the data is variable and changes rapidly. The data must be immediately available and must remain accessible for up to 3 months. The company needs the most cost-effective solution that will not increase retrieval time
Which S3 storage class should the company use to meet these requirements?

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

Answer: D

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

NEW QUESTION 232
- (Exam Topic 3)
A company has a Microsoft NET application that runs on an on-premises Windows Server Trie application stores data by using an Oracle Database Standard Edition server The company is planning a migration to AWS and wants to minimize development changes while moving the application The AWS application environment should be highly available
Which combination of actions should the company take to meet these requirements? (Select TWO)

A. Refactor the application as serverless with AWS Lambda functions running NET Cote
B. Rehost the application in AWS Elastic Beanstalk with the NET platform in a Multi-AZ deployment
C. Replatform the application to run on Amazon EC2 with the Amazon Linux Amazon Machine Image (AMI)
D. Use AWS Database Migration Service (AWS DMS) to migrate from the Oracle database to Amazon DynamoDB in a Multi-AZ deployment
E. Use AWS Database Migration Service (AWS DMS) to migrate from the Oracle database to Oracle on Amazon RDS in a Multi-AZ deployment

Answer: BE

NEW QUESTION 235
- (Exam Topic 3)
A company is implementing new data retention policies for all databases that run on Amazon RDS DB instances. The company must retain daily backups for a minimum period of 2 years. The backups must be consistent and restorable.
Which solution should a solutions architect recommend to meet these requirements?

A. Create a backup vault in AWS Backup to retain RDS backup
B. Create a new backup plan with a daily schedule and an expiration period of 2 years after creatio
C. Assign the RDS DB instances to the backup plan.
D. Configure a backup window for the RDS DB instances for daily snapshot
E. Assign a snapshot retention policy of 2 years to each RDS DB instanc
F. Use Amazon Data Lifecycle Manager (Amazon DLM) to schedule snapshot deletions.
G. Configure database transaction logs to be automatically backed up to Amazon CloudWatch Logs with an expiration period of 2 years.
H. Configure an AWS Database Migration Service (AWS DMS) replication tas
I. Deploy a replication instance, and configure a change data capture (CDC) task to stream database changes to Amazon S3 as the targe
J. Configure S3 Lifecycle policies to delete the snapshots after 2 years.

Answer: A

NEW QUESTION 239
- (Exam Topic 3)
A company hosts a multi-tier web application that uses an Amazon Aurora MySQL DB cluster for storage. The application tier is hosted on Amazon EC2 instances. The company's IT security guidelines mandate that the database credentials be encrypted and rotated every 14 days
What should a solutions architect do to meet this requirement with the LEAST operational effort?

A. Create a new AWS Key Management Service (AWS KMS) encryption key Use AWS Secrets Manager to create a new secret that uses the KMS key with the appropriate credentials Associate the secret with the Aurora DB cluster Configure a custom rotation period of 14 days
B. Create two parameters in AWS Systems Manager Parameter Store one for the user name as a string parameter and one that uses the SecureStnng type for the password Select AWS Key Management Service (AWS KMS) encryption for the password parameter, and load these parameters in the application tier Implement an AWS Lambda function that rotates the password every 14 days.
C. Store a file that contains the credentials in an AWS Key Management Service (AWS KMS) encryptedAmazon Elastic File System (Amazon EFS) file system Mount the EFS file system in all EC2 instances of the application tie
D. Restrict the access to the file on the file system so that the application can read the file and that only super users can modify the file Implement an AWS

Lambda function that rotates the key in Aurora every 14 days and writes new credentials into the file
E. Store a file that contains the credentials in an AWS Key Management Service (AWS KMS) encrypted Amazon S3 bucket that the application uses to load the credentials Download the file to the application regularly to ensure that the correct credentials are used Implement an AWS Lambda function that rotates the Aurora credentials every 14 days and uploads these credentials to the file in the S3 bucket

Answer: A

NEW QUESTION 243
- (Exam Topic 3)
At part of budget planning. management wants a report of AWS billed dams listed by user. The data will be used to create department budgets. A solution architect needs to determine the most efficient way to obtain this report Information
Which solution meets these requirement?

A. Run a query with Amazon Athena to generate the report.
B. Create a report in Cost Explorer and download the report
C. Access the bill details from the runing dashboard and download Via bill.
D. Modify a cost budget in AWS Budgets to alert with Amazon Simple Email Service (Amazon SES).

Answer: B

NEW QUESTION 248
- (Exam Topic 3)
A company is designing a cloud communications platform that is driven by APIs. The application is hosted on Amazon EC2 instances behind a Network Load Balancer (NLB). The company uses Amazon API Gateway to provide external users with access to the application through APIs. The company wants to protect the platform against web exploits like SQL injection and also wants to detect and mitigate large, sophisticated DDoS attacks.
Which combination of solutions provides the MOST protection? (Select TWO.)

A. Use AWS WAF to protect the NLB.
B. Use AWS Shield Advanced with the NLB.
C. Use AWS WAF to protect Amazon API Gateway.
D. Use Amazon GuardDuty with AWS Shield Standard.
E. Use AWS Shield Standard with Amazon API Gateway.

Answer: BC

Explanation:
AWS Shield Advanced provides expanded DDoS attack protection for your Amazon EC2 instances, Elastic Load Balancing load balancers, CloudFront distributions, Route 53 hosted zones, and AWS Global Accelerator standard accelerators.
AWS WAF is a web application firewall that lets you monitor the HTTP and HTTPS requests that are forwarded to your protected web application resources. You can protect the following resource types:
Amazon CloudFront distribution Amazon API Gateway REST API Application Load Balancer
AWS AppSync GraphQL API Amazon Cognito user pool
https://docs.aws.amazon.com/waf/latest/developerguide/what-is-aws-waf.html

NEW QUESTION 250
- (Exam Topic 3)
A company has an On-premises volume backup solution that has reached its end of life. The company wants to use AWS as part of a new backup solution and wants to maintain local access to all the data while it is backed up on AWS. The company wants to ensure that the data backed up on AWS is automatically and securely transferred.
Which solution meets these requirements?

A. Use AWS Snowball to migrate data out of the on-premises solution to Amazon S3. Configure on-premises systems to mount the Snowball S3 endpoint to provide local access to the data.
B. Use AWS Snowball Edge to migrate data out of the on-premises solution to Amazon S3.Use the Snowball Edge file interface to provide on-premises systems with local access to the data.
C. Use AWS Storage Gateway and configure a cached volume gatewa
D. Run the Storage Gateway software application on premises and configure a percentage of data to cache locall
E. Mount the gateway storage volumes to provide local access to the data.
F. Use AWS Storage Gateway and configure a stored volume gatewa
G. Run the Storage software application on premises and map the gateway storage volumes to on-premises storag
H. Mount the gateway storage volumes to provide local access to the data.

Answer: C

NEW QUESTION 252
- (Exam Topic 3)
A company wants to migrate an Oracle database to AWS. The database consists of a single table that contains millions of geographic information systems (GIS) images that are high resolution and are identified by a geographic code.
When a natural disaster occurs tens of thousands of images get updated every few minutes. Each geographic code has a single image or row that is associated with it. The company wants a solution that is highly available and scalable during such events
Which solution meets these requirements MOST cost-effectively?

A. Store the images and geographic codes in a database table Use Oracle running on an Amazon RDS Multi-AZ DB instance
B. Store the images in Amazon S3 buckets Use Amazon DynamoDB with the geographic code as the key and the image S3 URL as the value
C. Store the images and geographic codes in an Amazon DynamoDB table Configure DynamoDB Accelerator (DAX) during times of high load
D. Store the images in Amazon S3 buckets Store geographic codes and image S3 URLs in a database table Use Oracle running on an Amazon RDS Multi-AZ DB instance.

Answer: A

NEW QUESTION 255
- (Exam Topic 3)
A company has a web application that is based on Java and PHP The company plans to move the application from on premises to AWS The company needs the ability to test new site features frequently. The company also needs a highly available and managed solution that requires minimum operational overhead
Which solution will meet these requirements?

A. Create an Amazon S3 bucket Enable static web hosting on the S3 bucket Upload the static content to the S3 bucket Use AWS Lambda to process all dynamic content
B. Deploy the web application to an AWS Elastic Beanstalk environment Use URL swapping to switch between multiple Elastic Beanstalk environments for feature testing
C. Deploy the web application lo Amazon EC2 instances that are configured with Java and PHP Use Auto Scaling groups and an Application Load Balancer to manage the website's availability
D. Containerize the web application Deploy the web application to Amazon EC2 instances Use the AWS Load Balancer Controller to dynamically route traffic between containers thai contain the new site features for testing

Answer: B

NEW QUESTION 259
- (Exam Topic 3)
An application running on an Amazon EC2 instance in VPC-A needs to access files in another EC2 instance in VPC-B. Both VPCs are in separate AWS accounts. The network administrator needs to design a solution to configure secure access to EC2 instance in VPC-B from VPC-A. The connectivity should not have a single point of failure or bandwidth concerns.
Which solution will meet these requirements?

A. Set up a VPC peering connection between VPC-A and VPC-B.
B. Set up VPC gateway endpoints for the EC2 instance running in VPC-B.
C. Attach a virtual private gateway to VPC-B and set up routing from VPC-A.
D. Create a private virtual interface (VIF) for the EC2 instance running in VPC-B and add appropriate routes from VPC-A.

Answer: A

Explanation:
AWS uses the existing infrastructure of a VPC to create a VPC peering connection; it is neither a gateway nor a VPN connection, and does not rely on a separate piece of physical hardware. There is no single point of failure for communication or a bandwidth bottleneck.
https://docs.aws.amazon.com/vpc/latest/peering/what-is-vpc-peering.html

NEW QUESTION 260
- (Exam Topic 3)
A company needs to provide its employee with secure access to confidential and sensitive files. The company wants to ensure that the files can be accessed only by authorized users. The files must be downloaded security to the employees devices.
The files are stored in an on-premises Windows files server. However, due to an increase in remote usage, the file server out of capacity.
Which solution will meet these requirement?

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

Answer: C

NEW QUESTION 262
- (Exam Topic 3)
A company wants to configure its Amazon CloudFront distribution to use SSL/TLS certificates. The company does not want to use the default domain name for the distribution. Instead, the company wants to use a different domain name for the distribution.
Which solution will deploy the certificate with icurring any additional costs?

A. Request an Amazon issued private certificate from AWS Certificate Manager (ACM) in the us-east-1 Region
B. Request an Amazon issued private certificate from AWS Certificate Manager (ACM) in the us-west-1 Region.
C. Request an Amazon issued public certificate from AWS Certificate Manager (ACU) in the us-east-1 Region
D. Request an Amazon issued public certificate from AWS Certificate Manager (ACU) in the us-west-1 Regon.

Answer: B

NEW QUESTION 266
- (Exam Topic 3)
An online retail company has more than 50 million active customers and receives more than 25,000 orders each day. The company collects purchase data for customers and stores this data in Amazon S3. Additional customer data is stored in Amazon RDS.
The company wants to make all the data available to various teams so that the teams can perform analytics. The solution must provide the ability to manage fine-grained permissions for the data and must minimize operational overhead.
Which solution will meet these requirements?

A. Migrate the purchase data to write directly to Amazon RD
B. Use RDS access controls to limit access.
C. Schedule an AWS Lambda function to periodically copy data from Amazon RDS to Amazon S3. Create an AWS Glue crawle
D. Use Amazon Athena to query the dat

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

Answer: C

Explanation:
https://aws.amazon.com/blogs/big-data/manage-fine-grained-access-control-using-aws-lake-formation/

NEW QUESTION 267
- (Exam Topic 3)
A company is using AWS to design a web application that will process insurance quotes Users will request quotes from the application Quotes must be separated by quote type, must be responded to within 24 hours, and must not get lost The solution must maximize operational efficiency and must minimize maintenance. Which solution meets these requirements?

A. Create multiple Amazon Kinesis data streams based on the quote type Configure the web application to send messages to the proper data stream Configure each backend group of application servers to use the Kinesis Client Library (KCL) to pool messages from its own data stream
B. Create an AWS Lambda function and an Amazon Simple Notification Service (Amazon SNS) topic for each quote type Subscribe the Lambda function to its associated SNS topic Configure the application to publish requests tot quotes to the appropriate SNS topic
C. Create a single Amazon Simple Notification Service (Amazon SNS) topic Subscribe Amazon Simple Queue Service (Amazon SQS) queues to the SNS topic Configure SNS message filtering to publish messages to the proper SQS queue based on the quote type Configure each backend application server to use its own SQS queue
D. Create multiple Amazon Kinesis Data Firehose delivery streams based on the quote type to deliver data streams to an Amazon Elasucsearch Service (Amazon ES) cluster Configure the application to send messages to the proper delivery stream Configure each backend group of application servers to search for the messages from Amazon ES and process them accordingly

Answer: C

NEW QUESTION 269
- (Exam Topic 3)
A media company hosts its website on AWS. The website application's architecture includes a fleet of Amazon EC2 instances behind an Application Load Balancer (ALB) and a database that is hosted on Amazon Aurora The company's cyber security teem reports that the application is vulnerable to SOL injection.
How should the company resolve this issue?

A. Use AWS WAF in front of the ALB Associate the appropriate web ACLs with AWS WAF.
B. Create an ALB listener rule to reply to SQL injection with a fixed response
C. Subscribe to AWS Shield Advanced to block all SQL injection attempts automatically.
D. Set up Amazon Inspector to block all SOL injection attempts automatically

Answer: A

NEW QUESTION 273
- (Exam Topic 3)
An application that is hosted on Amazon EC2 instances needs to access an Amazon S3 bucket Traffic must not traverse the internet How should a solutions architect configure access to meet these requirements?

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

Answer: B

NEW QUESTION 276
- (Exam Topic 3)
A gaming company is moving its public scoreboard from a data center to the AWS Cloud. The company uses Amazon EC2 Windows Server instances behind an Application Load Balancer to host its dynamic application. The company needs a highly available storage solution for the application. The application consists of static files and dynamic server-side code.
Which combination of steps should a solutions architect take to meet these requirements? (Select TWO.) ' A. Store the static files on Amazon S3. Use Amazon A. CloudFront to cache objects at the edge.

A. Store the static files on Amazon S3. Use Amazon ElastiCache to cache objects at the edge.
B. Store the server-side code on Amazon Elastic File System (Amazon EFS). Mount the EFS volume on each EC2 instance to share the files.
C. Store the server-side code on Amazon FSx for Windows File Serve
D. Mount the FSx for Windows File Server volume on each EC2 instance to share the files.
E. Store the server-side code on a General Purpose SSD (gp2) Amazon Elastic Block Store (Amazon EBS) volum
F. Mount the EBS volume on each EC2 instance to share the files.

Answer: AE

NEW QUESTION 281
- (Exam Topic 3)
A company has deployed a database in Amazon RDS for MySQL. Due to increased transactions, the database support team is reporting slow reads against the

DB instance and recommends adding a read replica.
Which combination of actions should a solutions architect take before implementing this change? (Choose two.)

A. Enable binlog replication on the RDS primary node.
B. Choose a failover priority for the source DB instance.
C. Allow long-running transactions to complete on the source DB instance.
D. Create a global table and specify the AWS Regions where the table will be available.
E. Enable automatic backups on the source instance by setting the backup retention period to a value other than 0.

Answer: CE

Explanation:
"An active, long-running transaction can slow the process of creating the read replica. We recommend that you wait for long-running transactions to complete before creating a read replica. If you create multiple read replicas in parallel from the same source DB instance, Amazon RDS takes only one snapshot at the start of the first create action. When creating a read replica, there are a few things to consider. First, you must enable automatic backups on the source DB instance by setting the backup retention period to a value other than 0. This requirement also applies to a read replica that is the source DB instance for another read replica"
https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_ReadRepl.html

NEW QUESTION 284
- (Exam Topic 3)
A company has deployed a server less application that invokes an AWS Lambda function when new documents are uploaded to an Amazon S3 bucket The application uses the Lambda function to process the documents After a recent marketing campaign the company noticed that the application did not process many of The documents
What should a solutions architect do to improve the architecture of this application?

A. Set the Lambda function's runtime timeout value to 15 minutes
B. Configure an S3 bucket replication policy Stage the documents m the S3 bucket for later processing
C. Deploy an additional Lambda function Load balance the processing of the documents across the two Lambda functions
D. Create an Amazon Simple Queue Service (Amazon SOS) queue Send the requests to the queue Configure the queue as an event source for Lambda.

Answer: B

NEW QUESTION 287
......

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