SAA-C03 Dumps

AWS Certified Solutions Architect - Associate (SAA-C03)

https://www.certleader.com/SAA-C03-dumps.html
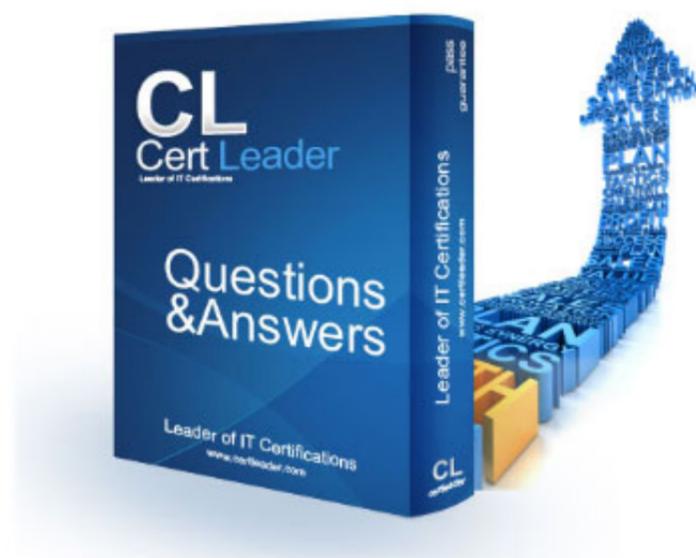
NEW QUESTION 1
- (Exam Topic 1)
A company is preparing to launch a public-facing web application in the AWS Cloud. The architecture consists of Amazon EC2 instances within a VPC behind an Elastic Load Balancer (ELB). A third-party service is used for the DNS. The company's solutions architect must recommend 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 2
- (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 3
- (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 4
- (Exam Topic 1)
A company is building an ecommerce web application on AWS. The application sends information about new orders to an Amazon API Gateway REST API to process. The company wants to ensure that orders are processed in the order that they are received.
Which solution will meet these requirements?

A. Use an API Gateway integration to publish a message to an Amazon Simple Notification Service (Amazon SNS) topic when the application receives an orde
B. Subscribe an AWS Lambda function to the topic to perform processing.
C. Use an API Gateway integration to send a message to an Amazon Simple Queue Service (Amazon SQS) FIFO queue when the application receives an orde
D. Configure the SQS FIFO queue to invoke an AWS Lambda function for processing.
E. Use an API Gateway authorizer to block any requests while the application processes an order.
F. Use an API Gateway integration to send a message to an Amazon Simple Queue Service (Amazon SQS) standard queue when the application receives an orde
G. Configure the SQS standard queue to invoke an AWS Lambda function for processing.

Answer: B

NEW QUESTION 5
- (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 6
- (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 7
- (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 8
- (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 9
- (Exam Topic 1)
A company is storing backup files by using Amazon S3 Standard storage. The files are accessed frequently for 1 month. However, the files are not accessed after 1 month. The company must keep the files indefinitely.
Which storage solution will meet these requirements MOST cost-effectively?

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

Answer: B

NEW QUESTION 10
- (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 10
- (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 11
- (Exam Topic 1)
A company is designing an application where users upload small files into Amazon S3. After a user uploads a file, the file requires one-time simple processing to transform the data and save the data in JSON format for later analysis.

Each file must be processed as quickly as possible after it is uploaded. Demand will vary. On some days, users will upload a high number of files. On other days, users will upload a few files or no files.
Which solution meets these requirements with the LEAST operational overhead?

A. Configure Amazon EMR to read text files from Amazon S3. Run processing scripts to transform the dat
B. Store the resulting JSON file in an Amazon Aurora DB cluster.
C. Configure Amazon S3 to send an event notification to an Amazon Simple Queue Service (Amazon SQS) queu
D. Use Amazon EC2 instances to read from the queue and process the dat
E. Store the resulting JSON file in Amazon DynamoDB.
F. Configure Amazon S3 to send an event notification to an Amazon Simple Queue Service (Amazon SQS) queu
G. Use an AWS Lambda function to read from the queue and process the dat
H. Store the resulting JSON file in Amazon DynamoD
I. Most Voted
J. Configure Amazon EventBridge (Amazon CloudWatch Events) to send an event to Amazon Kinesis Data Streams when a new file is uploade
K. Use an AWS Lambda function to consume the event from the stream and process the dat
L. Store the resulting JSON file in Amazon Aurora DB cluster.

Answer: C

Explanation:
Amazon S3 sends event notifications about S3 buckets (for example, object created, object removed, or object restored) to an SNS topic in the same Region.
The SNS topic publishes the event to an SQS queue in the central Region.
The SQS queue is configured as the event source for your Lambda function and buffers the event messages for the Lambda function.
The Lambda function polls the SQS queue for messages and processes the Amazon S3 event notifications according to your application's requirements.
https://docs.aws.amazon.com/prescriptive-guidance/latest/patterns/subscribe-a-lambda-function-to-event-notific

NEW QUESTION 16
- (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 18
- (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 21
- (Exam Topic 1)
A company's containerized application runs on an Amazon EC2 instance. The application needs to download security certificates before it can communicate with other business applications. The company wants a highly secure solution to encrypt and decrypt the certificates in near real time. The solution also needs to store data in highly available storage after the data is encrypted.
Which solution will meet these requirements with the LEAST operational overhead?

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

J. Allow the EC2 role to use the KMS key for encryption operation
K. Store the encrypted data on Amazon Elastic Block Store (Amazon EBS) volumes.

Answer: D

NEW QUESTION 24
- (Exam Topic 1)
A company runs its Infrastructure on AWS and has a registered base of 700.000 users for res document management application The company intends to create a product that converts large pdf files to jpg Imago files. The .pdf files average 5 MB in size. The company needs to store the original files and the converted files. A solutions architect must design a scalable solution to accommodate demand that will grow rapidly over lime.
Which solution meets these requirements MOST cost-effectively?

A. Save the pdf files to Amazon S3 Configure an S3 PUT event to invoke an AWS Lambda function to convert the files to jpg format and store them back in Amazon S3
B. Save the pdf files to Amazon DynamoD
C. Use the DynamoDB Streams feature to invoke an AWS Lambda function to convert the files to jpg format and store them hack in DynamoDB
D. Upload the pdf files to an AWS Elastic Beanstalk application that includes Amazon EC2 instances.Amazon Elastic Block Store (Amazon EBS) storage and an Auto Scaling grou
E. Use a program In the EC2 instances to convert the files to jpg format Save the .pdf files and the .jpg files In the EBS store.
F. Upload the .pdf files to an AWS Elastic Beanstalk application that includes Amazon EC2 instances, Amazon Elastic File System (Amazon EPS) storage, and an Auto Scaling grou
G. Use a program in the EC2 instances to convert the file to jpg format Save the pdf files and the jpg files in the EBS store.

Answer: A

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

NEW QUESTION 28
- (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 33
- (Exam Topic 1)
A company needs to store its accounting records in Amazon S3. The records must be immediately accessible for 1 year and then must be archived for an additional 9 years. No one at the company, including administrative users and root users, can be able to delete the records during the entire 10-year period. The records must be stored with maximum resiliency.
Which solution will meet these requirements?

A. Store the records in S3 Glacier for the entire 10-year perio
B. Use an access control policy to deny deletion of the records for a period of 10 years.
C. Store the records by using S3 Intelligent-Tierin
D. Use an IAM policy to deny deletion of the records.After 10 years, change the IAM policy to allow deletion.
E. Use an S3 Lifecycle policy to transition the records from S3 Standard to S3 Glacier Deep Archive after 1 yea
F. Use S3 Object Lock in compliance mode for a period of 10 years.
G. Use an S3 Lifecycle policy to transition the records from S3 Standard to S3 One Zone-Infrequent Access (S3 One Zone-IA) after 1 yea
H. Use S3 Object Lock in governance mode for a period of 10 years.

Answer: C

NEW QUESTION 38
- (Exam Topic 1)
A company needs the ability to analyze the log files of its proprietary application. The logs are stored in JSON format in an Amazon S3 bucket Queries will be simple and will run on-demand A solutions architect needs to perform the analysis with minimal changes to the existing architecture
What should the solutions architect do to meet these requirements with the LEAST amount of operational overhead?

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

Answer: C

Explanation:
Amazon Athena can be used to query JSON in S3

NEW QUESTION 40
- (Exam Topic 1)
A company is running an SMB file server in its data center. The file server stores large files that are accessed frequently for the first few days after the files are

created. After 7 days the files are rarely accessed.
The total data size is increasing and is close to the company's total storage capacity. A solutions architect must increase the company's available storage space without losing low-latency access to the most recently accessed files. The solutions architect must also provide file lifecycle management to avoid future storage issues.
Which solution will meet these requirements?

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

Answer: A

NEW QUESTION 44
- (Exam Topic 1)
A development team runs monthly resource-intensive tests on its general purpose Amazon RDS for MySQL DB instance with Performance Insights enabled. The testing lasts for 48 hours once a month and is the only process that uses the database. The team wants to reduce the cost of running the tests without reducing the compute and memory attributes of the DB instance.
Which solution meets these requirements MOST cost-effectively?

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

Answer: A

NEW QUESTION 45
- (Exam Topic 1)
A company needs to configure a real-time data ingestion architecture for its application. The company needs an API, a process that transforms data as the data is streamed, and a storage solution for the data.
Which solution will meet these requirements with the LEAST operational overhead?

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

Answer: C

NEW QUESTION 48
- (Exam Topic 1)
A company uses 50 TB of data for reporting. The company wants to move this data from on premises to AWS A custom application in the company's data center runs a weekly data transformation job. The company plans to pause the application until the data transfer is complete and needs to begin the transfer process as soon as possible.
The data center does not have any available network bandwidth for additional workloads A solutions architect must transfer the data and must configure the transformation job to continue to run in the AWS Cloud
Which solution will meet these requirements with the LEAST operational overhead?

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

Answer: C

NEW QUESTION 53
- (Exam Topic 1)
A company is preparing to deploy a new serverless workload. A solutions architect must use the principle of least privilege to configure permissions that will be used to run an AWS Lambda function. An Amazon EventBridge (Amazon CloudWatch Events) rule will invoke the function.
Which solution meets these requirements?

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

Answer: D

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

NEW QUESTION 58
- (Exam Topic 1)
A social media company allows users to upload images to its website. The website runs on Amazon EC2 instances. During upload requests, the website resizes the images to a standard size and stores the resized images in Amazon S3. Users are experiencing slow upload requests to the website.
The company needs to reduce coupling within the application and improve website performance. A solutions architect must design the most operationally efficient process for image uploads.
Which combination of actions should the solutions architect take to meet these requirements? (Choose two.)

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

Answer: BD

NEW QUESTION 63
- (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 68
- (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 72
- (Exam Topic 1)
A company runs a photo processing application that needs to frequently upload and download pictures from Amazon S3 buckets that are located in the same AWS Region. A solutions architect has noticed an increased cost in data transfer fees and needs to implement a solution to reduce these costs.
How can the solutions architect meet this requirement?

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

Answer: D

NEW QUESTION 76
- (Exam Topic 1)
A company has a production web application in which users upload documents through a web interlace or a mobile app. According to a new regulatory requirement, new documents cannot be modified or deleted after they are stored.
What should a solutions architect do to meet this requirement?

A. Store the uploaded documents in an Amazon S3 bucket with S3 Versioning and S3 Object Lock enabled
B. Store the uploaded documents in an Amazon S3 bucke
C. Configure an S3 Lifecycle policy to archive the documents periodically.
D. Store the uploaded documents in an Amazon S3 bucket with S3 Versioning enabled Configure an ACL to restrict all access to read-only.
E. Store the uploaded documents on an Amazon Elastic File System (Amazon EFS) volum
F. Access the data by mounting the volume in read-only mode.

Answer: A

Explanation:
https://docs.aws.amazon.com/AmazonS3/latest/userguide/object-lock-overview.html

NEW QUESTION 79
- (Exam Topic 1)
A company has registered its domain name with Amazon Route 53. The company uses Amazon API Gateway in the ca-central-1 Region as a public interface for its backend microservice APIs. Third-party services consume the APIs securely. The company wants to design its API Gateway URL with the company's domain name and corresponding certificate so that the third-party services can use HTTPS.
Which solution will meet these requirements?

A. Create stage variables in API Gateway with Name="Endpoint-URL" and Value="Company Domain Name" to overwrite the default UR
B. Import the public certificate associated with the company's domain name into AWS Certificate Manager (ACM).
C. Create Route 53 DNS records with the company's domain nam
D. Point the alias record to the Regional API Gateway stage endpoin
E. Import the public certificate associated with the company's domain name into AWS Certificate Manager (ACM) in the us-east-1 Region.
F. Create a Regional API Gateway endpoin
G. Associate the API Gateway endpoint with the company's domain nam
H. Import the public certificate associated with the company's domain name into AWS Certificate Manager (ACM) in the same Regio
I. Attach the certificate to the API Gateway endpoin
J. Configure Route 53 to route traffic to the API Gateway endpoint.
K. Create a Regional API Gateway endpoin
L. Associate the API Gateway endpoint with the company's domain nam
M. Import the public certificate associated with the company's domain name into AWS Certificate Manager (ACM) in the us-east-1 Regio
N. Attach the certificate to the API Gateway APIs.Create Route 53 DNS records with the company's domain nam
O. Point an A record to the company's domain name.

Answer: D

NEW QUESTION 83
- (Exam Topic 1)
A global company hosts its web application on Amazon EC2 instances behind an Application Load Balancer (ALB). The web application has static data and dynamic data. The company stores its static data in an Amazon S3 bucket. The company wants to improve performance and reduce latency for the static data and dynamic data. The company is using its own domain name registered with Amazon Route 53.
What should a solutions architect do to meet these requirements?

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

Answer: C

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

NEW QUESTION 86
- (Exam Topic 1)

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

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

Answer: A

NEW QUESTION 89
- (Exam Topic 1)
A company receives 10 TB of instrumentation data each day from several machines located at a single factory. The data consists of JSON files stored on a storage area network (SAN) in an on-premises data center located within the factory. The company wants to send this data to Amazon S3 where it can be accessed by several additional systems that provide critical near-real-lime analytics. A secure transfer is important because the data is considered sensitive.
Which solution offers the MOST reliable data transfer?

A. AWS DataSync over public internet
B. AWS DataSync over AWS Direct Connect
C. AWS Database Migration Service (AWS DMS) over public internet
D. AWS Database Migration Service (AWS DMS) over AWS Direct Connect

Answer: B

Explanation:
These are some of the main use cases for AWS DataSync: • Data migration – Move active datasets rapidly
over the network into Amazon S3, Amazon EFS, or FSx for Windows File Server. DataSync includes automatic encryption and data integrity validation to help make sure that your data arrives securely, intact, and ready to use.
"DataSync includes encryption and integrity validation to help make sure your data arrives securely, intact, and ready to use."
https://aws.amazon.com/datasync/faqs/

NEW QUESTION 90
- (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 94
- (Exam Topic 1)
An application runs on an Amazon EC2 instance in a VPC. The application processes logs that are stored in an Amazon S3 bucket. The EC2 instance needs to access the S3 bucket without connectivity to the internet.
Which solution will provide private network connectivity to Amazon S3?

A. Create a gateway VPC endpoint to the S3 bucket.
B. Stream the logs to Amazon CloudWatch Log
C. Export the logs to the S3 bucket.
D. Create an instance profile on Amazon EC2 to allow S3 access.
E. Create an Amazon API Gateway API with a private link to access the S3 endpoint.

Answer: A

Explanation:
VPC endpoint allows you to connect to AWS services using a private network instead of using the public Internet

NEW QUESTION 96
- (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 98
- (Exam Topic 1)
A company has several web servers that need to frequently access a common Amazon RDS MySQL Multi-AZ DB instance The company wants a secure method for the web servers to connect to the database while meeting a security requirement to rotate user credentials frequently.
Which solution meets these requirements?

A. Store the database user credentials in AWS Secrets Manager Grant the necessary IAM permissions to allow the web servers to access AWS Secrets Manager
B. Store the database user credentials in AWS Systems Manager OpsCenter Grant the necessary IAM permissions to allow the web servers to access OpsCenter
C. Store the database user credentials in a secure Amazon S3 bucket Grant the necessary IAM permissions to allow the web servers to retrieve credentials and access the database.
D. Store the database user credentials in files encrypted with AWS Key Management Service (AWS KMS) on the web server file syste
E. The web server should be able to decrypt the files and access the database

Answer: A

Explanation:
AWS Secrets Manager helps you protect secrets needed to access your applications, services, and IT resources. The service enables you to easily rotate, manage, and retrieve database credentials, API keys, and other secrets throughout their lifecycle.
https://docs.aws.amazon.com/secretsmanager/latest/userguide/intro.html
Secrets Manager enables you to replace hardcoded credentials in your code, including passwords, with an API call to Secrets Manager to retrieve the secret programmatically. This helps ensure the secret can't be compromised by someone examining your code, because the secret no longer exists in the code. Also, you can configure Secrets Manager to automatically rotate the secret for you according to a specified schedule. This enables you to replace long-term secrets with short-term ones, significantly reducing the risk of compromise.

NEW QUESTION 101
- (Exam Topic 1)
A company maintains a searchable repository of items on its website. The data is stored in an Amazon RDS for MySQL database table that contains more than 10 million rows The database has 2 TB of General Purpose SSD storage There are millions of updates against this data every day through the company's website The company has noticed that some insert operations are taking 10 seconds or longer The company has determined that the database storage performance is the problem
Which solution addresses this performance issue?

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

Answer: A

Explanation:
https://aws.amazon.com/ebs/features/
"Provisioned IOPS volumes are backed by solid-state drives (SSDs) and are the highest performance EBS volumes designed for your critical, I/O intensive database applications.
These volumes are ideal for both IOPS-intensive and throughput-intensive workloads that require extremely low latency."
https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/CHAP_Storage.html

NEW QUESTION 104
- (Exam Topic 2)
A solutions architect needs to securely store a database user name and password that an application uses to access an Amazon RDS DB instance. The application that accesses the database runs on an Amazon EC2 instance. The solutions architect wants to create a secure parameter in AWS Systems Manager Parameter Store.
What should the solutions architect do to meet this requirement?

A. Create an IAM role that has read access to the Parameter Store paramete
B. Allow Decrypt access to an AWS Key Management Service (AWS KMS) key that is used to encrypt the paramete
C. Assign this IAM role to the EC2 instance.
D. Create an IAM policy that allows read access to the Parameter Store paramete
E. Allow Decrypt access to an AWS Key Management Service (AWS KMS) key that is used to encrypt the paramete
F. Assign this IAM policy to the EC2 instance.
G. Create an IAM trust relationship between the Parameter Store parameter and the EC2 instanc
H. Specify Amazon RDS as a principal in the trust policy.
I. Create an IAM trust relationship between the DB instance and the EC2 instanc
J. Specify Systems Manager as a principal in the trust policy.

Answer: B

Explanation:
https://docs.aws.amazon.com/IAM/latest/UserGuide/reference_aws-services-that-work-with-iam.html

NEW QUESTION 105
- (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 107
- (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 108
- (Exam Topic 2)
A gaming company hosts a browser-based application on AWS. The users of the application consume a large number of videos and images that are stored in Amazon S3. This content is the same for all users.
The application has increased in popularity, and millions of users worldwide are accessing these media files. The company wants to provide the files to the users while reducing the load on the origin.
Which solution meets these requirements MOST cost-effectively?

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

Answer: B

NEW QUESTION 109
- (Exam Topic 2)
A company needs to move data from an Amazon EC2 instance to an Amazon S3 bucket. The company must ensure that no API calls and no data are routed through public internet routes. Only the EC2 instance can have access to upload data to the S3 bucket.
Which solution will meet these requirements?

A. Create an interface VPC endpoint for Amazon S3 in the subnet where the EC2 instance is located.Attach a resource policy to the S3 bucket to only allow the EC2 instance's IAM role for access.
B. Create a gateway VPC endpoint for Amazon S3 in the Availability Zone where the EC2 instance is locate
C. Attach appropriate security groups to the endpoin
D. Attach a resource policy lo the S3 bucket to only allow the EC2 instance's IAM role for access.
E. Run the nslookup tool from inside the EC2 instance to obtain the private IP address of the S3 bucket's service API endpoin
F. Create a route in the VPC route table to provide the EC2 instance with access to the S3 bucke
G. Attach a resource policy to the S3 bucket to only allow the EC2 instance's IAM role for access.
H. Use the AWS provided, publicly available ip-ranges.json tile to obtain the private IP address of the S3 bucket's service API endpoin
I. Create a route in the VPC route table to provide the EC2 instance with access to the S3 bucke
J. Attach a resource policy to the S3 bucket to only allow the EC2 instance's IAM role for access.

Answer: A

Explanation:
(https://aws.amazon.com/blogs/security/how-to-restrict-amazon-s3-bucket-access-to-a-specific-iam-role/)

NEW QUESTION 114
- (Exam Topic 2)
A company is migrating its on-premises PostgreSQL database to Amazon Aurora PostgreSQL. The
on-premises database must remain online and accessible during the migration. The Aurora database must remain synchronized with the on-premises database.
Which combination of actions must a solutions architect take to meet these requirements? (Choose two.)

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

Answer: CD

NEW QUESTION 119
- (Exam Topic 2)
A company wants to use the AWS Cloud to make an existing application highly available and resilient. The current version of the application resides in the company's data center. The application recently experienced data loss after a database server crashed because of an unexpected power outage.
The company needs a solution that avoids any single points of failure. The solution must give the application the ability to scale to meet user demand.
Which solution will meet these requirements?

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

Answer: A

NEW QUESTION 122
- (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 126
- (Exam Topic 2)
A company is building a web-based application running on Amazon EC2 instances in multiple Availability Zones. The web application will provide access to a repository of text documents totaling about 900 TB in size. The company anticipates that the web application will experience periods of high demand. A solutions architect must ensure that the storage component for the text documents can scale to meet the demand of the application at all times. The company is concerned about the overall cost of the solution.
Which storage solution meets these requirements MOST cost-effectively?

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

Answer: D

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

NEW QUESTION 130
- (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 132
- (Exam Topic 2)
A solutions architect must design a solution that uses Amazon CloudFront with an Amazon S3 origin to store a static website. The company's security policy requires that all website traffic be inspected by AWS WAR
How should the solutions architect comply with these requirements?

A. Configure an S3 bucket policy lo accept requests coming from the AWS WAF Amazon Resource Name (ARN) only.
B. Configure Amazon CloudFront to forward all incoming requests to AWS WAF before requesting content from the S3 origin.
C. Configure a security group that allows Amazon CloudFront IP addresses to access Amazon S3 only.Associate AWS WAF to CloudFront.
D. Configure Amazon CloudFront and Amazon S3 to use an origin access identity (OAI) to restrict access to the S3 bucke
E. Enable AWS WAF on the distribution.

Answer: D

Explanation:
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/private-content-restricting-access-to-s3
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/distribution-web-awswaf.html

NEW QUESTION 137
- (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 138
- (Exam Topic 2)
A company runs workloads on AWS. The company needs to connect to a service from an external provider. The service is hosted in the provider's VPC. According to the company's security team, the connectivity must be private and must be restricted to the target service. The connection must be initiated only from the company's VPC.
Which solution will mast these requirements?

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

Answer: D

NEW QUESTION 143
- (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 145
- (Exam Topic 2)
A company is running an online transaction processing (OLTP) workload on AWS. This workload uses an unencrypted Amazon RDS DB instance in a Multi-AZ deployment. Daily database snapshots are taken from this instance.
What should a solutions architect do to ensure the database and snapshots are always encrypted moving forward?

A. Encrypt a copy of the latest DB snapsho
B. Replace existing DB instance by restoring the encrypted snapshot
C. Create a new encrypted Amazon Elastic Block Store (Amazon EBS) volume and copy the snapshots to it Enable encryption on the DB instance
D. Copy the snapshots and enable encryption using AWS Key Management Service (AWS KMS) Restoreencrypted snapshot to an existing DB instance
E. Copy the snapshots to an Amazon S3 bucket that is encrypted using server-side encryption with AWS Key Management Service (AWS KMS) managed keys (SSE-KMS)

Answer: A

Explanation:
https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_RestoreFromSnapshot.html#USER_Restor Under "Encrypt unencrypted resources"
https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSEncryption.html

NEW QUESTION 146
- (Exam Topic 2)
A company needs to save the results from a medical trial to an Amazon S3 repository. The repository must allow a few scientists to add new files and must restrict all other users to read-only access. No users can have the ability to modify or delete any files in the repository. The company must keep every file in the repository for a minimum of 1 year after its creation date.
Which solution will meet these requirements?

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

Answer: C

NEW QUESTION 150
- (Exam Topic 2)
A company runs its two-tier ecommerce website on AWS. The web tier consists of a load balancer that sends traffic to Amazon EC2 instances. The database tier uses an Amazon RDS DB instance. The EC2 instances and the RDS DB instance should not be exposed to the public internet. The EC2 instances require internet access to complete payment processing of orders through a third-party web service. The application must be highly available.
Which combination of configuration options will meet these requirements? (Choose two.)

A. Use an Auto Scaling group to launch the EC2 instances in private subnet
B. Deploy an RDS Multi-AZ DB instance in private subnets.
C. Configure a VPC with two private subnets and two NAT gateways across two Availability Zones.Deploy an Application Load Balancer in the private subnets.
D. Use an Auto Scaling group to launch the EC2 instances in public subnets across two Availability Zones.Deploy an RDS Multi-AZ DB instance in private subnets.
E. Configure a VPC with one public subnet, one private subnet, and two NAT gateways across two Availability Zone
F. Deploy an Application Load Balancer in the public subnet.
G. Configure a VPC with two public subnets, two private subnets, and two NAT gateways across two Availability Zone
H. Deploy an Application Load Balancer in the public subnets.

Answer: AE

Explanation:
Before you begin: Decide which two Availability Zones you will use for your EC2 instances. Configure your virtual private cloud (VPC) with at least one public subnet in each of these Availability Zones. These public subnets are used to configure the load balancer. You can launch your EC2 instances in other subnets of these Availability Zones instead.

NEW QUESTION 154
- (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) queue
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 156
- (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 161
- (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 162
- (Exam Topic 2)
A company is running a multi-tier web application on premises. The web application is containerized and runs on a number of Linux hosts connected to a PostgreSQL database that contains user records. The operational overhead of maintaining the infrastructure and capacity planning is limiting the company's growth. A solutions architect must improve the application's infrastructure.
Which combination of actions should the solutions architect take to accomplish this? (Choose two.)

A. Migrate the PostgreSQL database to Amazon Aurora
B. Migrate the web application to be hosted on Amazon EC2 instances.
C. Set up an Amazon CloudFront distribution for the web application content.
D. Set up Amazon ElastiCache between the web application and the PostgreSQL database.
E. Migrate the web application to be hosted on AWS Fargate with Amazon Elastic Container Service (Amazon ECS).

Answer: AE

NEW QUESTION 163
- (Exam Topic 2)
A company is developing a file-sharing application that will use an Amazon S3 bucket for storage. The company wants to serve all the files through an Amazon CloudFront distribution. The company does not want the files to be accessible through direct navigation to the S3 URL.
What should a solutions architect do to meet these requirements?

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

Answer: D

Explanation:
https://aws.amazon.com/premiumsupport/knowledge-center/cloudfront-access-to-amazon-s3/
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/private-content-restricting-access-to-s3

NEW QUESTION 168
- (Exam Topic 2)
An ecommerce company hosts its analytics application in the AWS Cloud. The application generates about 300 MB of data each month. The data is stored in

JSON format. The company is evaluating a disaster recovery solution to back up the data. The data must be accessible in milliseconds if it is needed, and the data must be kept for 30 days.
Which solution meets these requirements MOST cost-effectively?

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

Answer: C

NEW QUESTION 170
- (Exam Topic 2)
A company is planning to build a high performance computing (HPC) workload as a service solution that Is hosted on AWS A group of 16 AmazonEC2Ltnux
Instances requires the lowest possible latency for
node-to-node communication. The instances also need a shared block device volume for high-performing
storage.
Which solution will meet these requirements?

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

Answer: A

NEW QUESTION 171
- (Exam Topic 2)
A company wants to migrate its on-premises data center to AWS. According to the company's compliance requirements, the company can use only the ap-
northeast-3 Region. Company administrators are not permitted to connect VPCs to the internet.
Which solutions will meet these requirements? (Choose two.)

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

Answer: AC

NEW QUESTION 172
- (Exam Topic 2)
Organizers for a global event want to put daily reports online as static HTML pages. The pages are expected to generate millions of views from users around the
world. The files are stored In an Amazon S3 bucket. A solutions architect has been asked to design an efficient and effective solution.
Which action should the solutions architect take to accomplish this?

A. Generate presigned URLs for the files.
B. Use cross-Region replication to all Regions.
C. Use the geoproximtty feature of Amazon Route 53.
D. Use Amazon CloudFront with the S3 bucket as its origin.

Answer: D

NEW QUESTION 174
- (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 175
- (Exam Topic 2)
A company uses AWS Organizations to create dedicated AWS accounts for each business unit to manage each business unit's account independently upon request. The root email recipient missed a notification that was sent to the root user email address of one account. The company wants to ensure that all future notifications are not missed. Future notifications must be limited to account administrators.
Which solution will meet these requirements?

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

Answer: D

NEW QUESTION 180
- (Exam Topic 2)
A company has a data ingestion workflow that includes the following components:
• An Amazon Simple Notation Service (Amazon SNS) topic that receives notifications about new data deliveries
• An AWS Lambda function that processes and stores the data
The ingestion workflow occasionally fails because of network connectivity issues. When tenure occurs the corresponding data is not ingested unless the company manually reruns the job. What should a solutions architect do to ensure that all notifications are eventually processed?

A. Configure the Lambda function (or deployment across multiple Availability Zones
B. Modify me Lambda functions configuration to increase the CPU and memory allocations tor the (unction
C. Configure the SNS topic's retry strategy to increase both the number of retries and the wait time between retries
D. Configure an Amazon Simple Queue Service (Amazon SQS) queue as the on failure destination Modify the Lambda function to process messages in the queue

Answer: A

NEW QUESTION 181
- (Exam Topic 2)
A company has implemented a self-managed DNS solution on three Amazon EC2 instances behind a Network Load Balancer (NLB) in the us-west-2 Region. Most of the company's users are located in the United States and Europe. The company wants to improve the performance and availability of the solution. The company launches and configures three EC2 instances in the eu-west-1 Region and adds the EC2 instances as targets for a new NLB.
Which solution can the company use to route traffic to all the EC2 instances?

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

Answer: B

Explanation:
For standard accelerators, Global Accelerator uses the AWS global network to route traffic to the optimal regional endpoint based on health, client location, and policies that you configure, which increases the availability of your applications. Endpoints for standard accelerators can be Network Load Balancers, Application Load Balancers, Amazon EC2 instances, or Elastic IP addresses that are located in one AWS Region or multiple Regions.
https://docs.aws.amazon.com/global-accelerator/latest/dg/what-is-global-accelerator.html

NEW QUESTION 185
- (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 186
- (Exam Topic 2)
A company has an event-driven application that invokes AWS Lambda functions up to 800 times each minute with varying runtimes. The Lambda functions access data that is stored in an Amazon Aurora MySQL OB cluster. The company is noticing connection timeouts as user activity increases The database shows no signs of being overloaded. CPU. memory, and disk access metrics are all low.
Which solution will resolve this issue with the LEAST operational overhead?

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

Answer: D

NEW QUESTION 189
- (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 193
- (Exam Topic 2)
A company wants to move its application to a serverless solution. The serverless solution needs to analyze existing and new data by using SL. The company stores the data in an Amazon S3 bucket. The data requires encryption and must be replicated to a different AWS Region.
Which solution will meet these requirements with the LEAST operational overhead?

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

Answer: A

NEW QUESTION 196
- (Exam Topic 2)
A company runs an Oracle database on premises. As part of the company's migration to AWS, the company wants to upgrade the database to the most recent available version. The company also wants to set up disaster recovery (DR) for the database. The company needs to minimize the operational overhead for normal operations and DR setup. The company also needs to maintain access to the database's underlying operating system.
Which solution will meet these requirements?

A. Migrate the Oracle database to an Amazon EC2 instanc
B. Set up database replication to a different AWS Region.
C. Migrate the Oracle database to Amazon RDS for Oracl
D. Activate Cross-Region automated backups to replicate the snapshots to another AWS Region.
E. Migrate the Oracle database to Amazon RDS Custom for Oracl
F. Create a read replica for the database in another AWS Region.
G. Migrate the Oracle database to Amazon RDS for Oracl
H. Create a standby database in another Availability Zone.

Answer: C

Explanation:
https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/rds-custom.html and https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/working-with-custom-oracle.html

NEW QUESTION 201
- (Exam Topic 2)
A company has a service that produces event data. The company wants to use AWS to process the event data as it is received. The data is written in a specific order that must be maintained throughout processing The company wants to implement a solution that minimizes operational overhead. How should a solutions architect accomplish this?

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

Answer: A

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

NEW QUESTION 205
- (Exam Topic 2)
A company runs a high performance computing (HPC) workload on AWS. The workload required low-latency network performance and high network throughput with tightly coupled node-to-node communication. The Amazon EC2 instances are properly sized for compute and storage capacity, and are launched using default options.
What should a solutions architect propose to improve the performance of the workload?

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

Answer: A

Explanation:
https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-resource-ec2-placementgroup.html "A cluster placement group is a logical grouping of instances within a single Availability Zone that benefit from low network latency, high network throughput"

NEW QUESTION 208
- (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 211
- (Exam Topic 2)
A new employee has joined a company as a deployment engineer. The deployment engineer will be using AWS CloudFormation templates to create multiple AWS resources. A solutions architect wants the deployment engineer to perform job activities while following the principle of least privilege.
Which steps should the solutions architect do in conjunction to reach this goal? (Select two.)

A. Have the deployment engineer use AWS account roof user credentials for performing AWS CloudFormation stack operations.
B. Create a new IAM user for the deployment engineer and add the IAM user to a group that has the PowerUsers IAM policy attached.
C. Create a new IAM user for the deployment engineer and add the IAM user to a group that has the Administrate/Access IAM policy attached.
D. Create a new IAM User for the deployment engineer and add the IAM user to a group that has an IAM policy that allows AWS CloudFormation actions only.
E. Create an IAM role for the deployment engineer to explicitly define the permissions specific to the AWS CloudFormation stack and launch stacks using Dial IAM role.

Answer: DE

Explanation:
https://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles.html https://docs.aws.amazon.com/IAM/latest/UserGuide/id_users.html

NEW QUESTION 216
- (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 218
- (Exam Topic 2)
A company runs a stateless web application in production on a group of Amazon EC2 On-Demand Instances behind an Application Load Balancer. The application experiences heavy usage during an 8-hour period each business day. Application usage is moderate and steady overnight Application usage is low during weekends.
The company wants to minimize its EC2 costs without affecting the availability of the application. Which solution will meet these requirements?

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

Answer: B

NEW QUESTION 222
- (Exam Topic 3)
A company is migrating a Linux-based web server group to AWS. The web servers must access files in a shared file store for some content. The company must not make any changes to the application.
What should a solutions architect do to meet these requirements?

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

Answer: C

NEW QUESTION 224
- (Exam Topic 3)
A company wants to deploy a new public web application on AWS The application includes a web server tier that uses Amazon EC2 instances The application also includes a database tier that uses an Amazon RDS for MySQL DB instance
The application must be secure and accessible for global customers that have dynamic IP addresses How should a solutions architect configure the security groups to meet these requirements'?

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

Answer: A

NEW QUESTION 229
- (Exam Topic 3)
A company uses a payment processing system that requires messages for a particular payment ID to be received in the same order that they were sent Otherwise, the payments might be processed incorrectly.
Which actions should a solutions architect take to meet this requirement? (Select TWO.)

A. Write the messages to an Amazon DynamoDB table with the payment ID as the partition key
B. Write the messages to an Amazon Kinesis data stream with the payment ID as the partition key.
C. Write the messages to an Amazon ElastiCache for Memcached cluster with the payment ID as the key
D. Write the messages to an Amazon Simple Queue Service (Amazon SQS) queue Set the message attribute to use the payment ID
E. Write the messages to an Amazon Simple Queue Service (Amazon SQS) FIFO queu
F. Set the message group to use the payment ID.

Answer: AE

NEW QUESTION 231
- (Exam Topic 3)
A company s order system sends requests from clients to Amazon EC2 instances The EC2 instances process the orders and men store the orders in a database on Amazon RDS Users report that they must reprocess orders when the system fails. The company wants a resilient solution that can process orders automatically

it a system outage occurs.
What should a solutions architect do to meet these requirements?

A. Move (he EC2 Instances into an Auto Scaling group Create an Amazon EventBridge (Amazon CloudWatch Events) rule to target an Amazon Elastic Container Service (Amazon ECS) task
B. Move the EC2 instances into an Auto Scaling group behind an Application Load Balancer (ALB) Update the order system to send messages to the ALB endpoint.
C. Move the EC2 instances into an Auto Scaling group Configure the order system to send messages to an Amazon Simple Queue Service (Amazon SQS) queue Configure the EC2 instances to consume messages from the queue
D. Create an Amazon Simple Notification Service (Amazon SNS) topic Create an AWS Lambda function, and subscribe the function to the SNS topic Configure the order system to send messages to the SNS topic Send a command to the EC2 instances to process the messages by using AWS Systems Manager Run Command

Answer: C

NEW QUESTION 233
- (Exam Topic 3)
A company is planning to migrate a commercial off-the-shelf application from is on-premises data center to AWS. The software has a software licensing model using sockets and cores with predictable capacity and uptime requirements. The company wants to use its existing licenses, which were purchased earlier this year.
Which Amazon EC2 pricing option is the MOST cost-effective?

A. Dedicated Reserved Hosts
B. Dedicated On-Demand Hosts
C. Dedicated Reserved Instances
D. Dedicated On-Oemand Instances

Answer: A

NEW QUESTION 236
- (Exam Topic 3)
A company needs to create an Amazon Elastic Kubernetes Service (Amazon EKS) cluster to host a digital
media streaming application. The EKS cluster will use a managed node group that is backed by Amazon Elastic Block Store (Amazon EBS) volumes for storage.
The company must encrypt all data at rest by using a customer managed key that is stored in AWS Key Management Service (AWS KMS)
Which combination of actions will meet this requirement with the LEAST operational overhead? (Select TWO.)

A. Use a Kubernetes plugin that uses the customer managed key to perform data encryption.
B. After creation of the EKS cluster, locate the EBS volume
C. Enable encryption by using the customer managed key.
D. Enable EBS encryption by default in the AWS Region where the EKS cluster will be create
E. Select the customer managed key as the default key.
F. Create the EKS cluster Create an IAM role that has cuwlicy that grants permission to the customer managed ke
G. Associate the role with the EKS cluster.
H. Store the customer managed key as a Kubernetes secret in the EKS cluste
I. Use the customer managed key to encrypt the EBS volumes.

Answer: AD

NEW QUESTION 241
- (Exam Topic 3)
An image-hosting company stores its objects in Amazon S3 buckets. The company wants to avoid accidental exposure of the objects in the S3 buckets to the public. All S3 objects in the entire AWS account need to remain private
Which solution will meal these requirements?

A. Use Amazon GuardDuty to monitor S3 bucket policies Create an automatic remediation action rule that uses an AWS Lambda function to remediate any change that makes the objects public
B. Use AWS Trusted Advisor to find publicly accessible S3 Dockets Configure email notifications In Trusted Advisor when a change is detected manually change the S3 bucket policy if it allows public access
C. Use AWS Resource Access Manager to find publicly accessible S3 buckets Use Amazon Simple Notification Service (Amazon SNS) to invoke an AWS Lambda function when a change it detected.Deploy a Lambda function that programmatically remediates the change.
D. Use the S3 Block Public Access feature on the account leve
E. Use AWS Organizations to create a service control policy (SCP) that prevents IAM users from changing the setting Apply tie SCP to tie account

Answer: D

NEW QUESTION 244
- (Exam Topic 3)
An ecommerce company is building a distributed application that involves several serverless functions and AWS services to complete order-processing tasks.
These tasks require manual approvals as part of the workflow A solutions architect needs to design an architecture for the order-processing application The solution must be able to combine multiple AWS Lambda functions into responsive serverless applications The solution also must orchestrate data and services that run on Amazon EC2 instances, containers, or on-premises servers
Which solution will meet these requirements with the LEAST operational overhead?

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

Answer: D

NEW QUESTION 247
- (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 252
- (Exam Topic 3)
A research laboratory needs to process approximately 8 TB of data The laboratory requires sub-millisecond latencies and a minimum throughput of 6 GBps for the storage subsystem Hundreds of Amazon EC2 instances that run Amazon Linux will distribute and process the data
Which solution will meet the performance requirements?

A. Create an Amazon FSx for NetApp ONTAP file system Set each volume's tiering policy to ALL Import the raw data into the file system Mount the file system on the EC2 instances
B. Create an Amazon S3 bucket to stofe the raw data Create an Amazon FSx for Lustre file system that uses persistent SSD storage Select the option to import data from and export data to Amazon S3 Mount the file system on the EC2 instances
C. Create an Amazon S3 bucket to store the raw data Create an Amazon FSx for Lustre file system that uses persistent HDD storage Select the option to import data from and export data to Amazon S3 Mount the file system on the EC2 instances
D. Create an Amazon FSx for NetApp ONTAP file system Set each volume's tienng policy to NON
E. Import the raw data into the file system Mount the file system on the EC2 instances

Answer: B

Explanation:
Create an Amazon S3 bucket to store the raw data Create an Amazon FSx for Lustre file system that uses persistent SSD storage Select the option to import data from and export data to Amazon S3 Mount the file system on the EC2 instances. Amazon FSx for Lustre uses SSD storage for sub-millisecond latencies and up to 6 GBps throughput, and can import data from and export data to Amazon S3. Additionally, the option to select persistent SSD storage will ensure that the data is stored on the disk and not lost if the file system is stopped.

NEW QUESTION 254
- (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 255
- (Exam Topic 3)
A company runs a public three-Tier web application in a VPC The application runs on Amazon EC2 instances across multiple Availability Zones. The EC2 instances that run in private subnets need to communicate with a license server over the internet The company needs a managed solution that minimizes operational maintenance
Which solution meets these requirements"

A. Provision a NAT instance in a public subnet Modify each private subnets route table with a default route that points to the NAT instance
B. Provision a NAT instance in a private subnet Modify each private subnet's route table with a default route that points to the NAT instance
C. Provision a NAT gateway in a public subnet Modify each private subnet's route table with a default route that points to the NAT gateway
D. Provision a NAT gateway in a private subnet Modify each private subnet's route table with a default route that points to the NAT gateway .

Answer: C

NEW QUESTION 260
- (Exam Topic 3)
A company has a three-tier application for image sharing. The application uses an Amazon EC2 instance for the front-end layer, another EC2 instance for the application layer, and a third EC2 instance for a MySQL database. A solutions architect must design a scalable and highly available solution that requires the least amount of change to the application.
Which solution meets these requirements?

A. Use Amazon S3 to host the front-end laye
B. Use AWS Lambda functions for the application laye
C. Move the database to an Amazon DynamoDB tabl
D. Use Amazon S3 to store and serve users' images.
E. Use load-balanced Multi-AZ AWS Elastic Beanstalk environments for the front-end layer and the application laye
F. Move the database to an Amazon RDS DB instance with multiple read replicas to serve users' images.
G. Use Amazon S3 to host the front-end laye
H. Use a fleet of EC2 instances in an Auto Scaling group for the application laye
I. Move the database to a memory optimized instance type to store and serve users' images.
J. Use load-balanced Multi-AZ AWS Elastic Beanstalk environments for the front-end layer and the application laye
K. Move the database to an Amazon RDS Multi-AZ DB instanc
L. Use Amazon S3 to store and serve users' images.

Answer: D

Explanation:
for "Highly available": Multi-AZ & for "least amount of changes to the application": Elastic Beanstalk automatically handles the deployment, from capacity provisioning, load balancing, auto-scaling to application health monitoring

NEW QUESTION 263
- (Exam Topic 3)
A company has an API that receives real-time data from a fleet of monitoring devices. The API stores this data in an Amazon RDS DB instance for later analysis. The amount of data that the monitoring devices send to the API fluctuates. During periods of heavy traffic, the API often returns timeout errors.
After an inspection of the logs, the company determines that the database is not capable of processing the volume of write traffic that comes from the API. A solutions architect must minimize the number of
connections to the database and must ensure that data is not lost during periods of heavy traffic. Which solution will meet these requirements?

A. Increase the size of the DB instance to an instance type that has more available memory.
B. Modify the DB instance to be a Multi-AZ DB instanc
C. Configure the application to write to all active RDS DB instances.
D. Modify the API to write incoming data to an Amazon Simple Queue Service (Amazon SQS) queu
E. Use an AWS Lambda function that Amazon SQS invokes to write data from the queue to the database.
F. Modify the API to write incoming data to an Amazon Simple Notification Service (Amazon SNS) topic.Use an AWS Lambda function that Amazon SNS invokes to write data from the topic to the database.

Answer: C

Explanation:
Using Amazon SQS will help minimize the number of connections to the database, as the API will write data to a queue instead of directly to the database. Additionally, using an AWS Lambda function that Amazon SQS invokes to write data from the queue to the database will help ensure that data is not lost during periods of heavy traffic, as the queue will serve as a buffer between the API and the database.

NEW QUESTION 267
- (Exam Topic 3)
A company is using Amazon CloudFront with this website. The company has enabled logging on the CloudFront distribution, and logs are saved in one of the company's Amazon S3 buckets The company needs to perform advanced analyses on the logs and build visualizations
What should a solutions architect do to meet these requirements'?

A. Use standard SQL queries in Amazon Athena to analyze the CloudFront togs in the S3 bucket Visualize the results with AWS Glue
B. Use standard SQL queries in Amazon Athena to analyze the CloudFront togs in the S3 bucket Visualize the results with Amazon QuickSight
C. Use standard SQL queries in Amazon DynamoDB to analyze the CloudFront logs m the S3 bucket Visualize the results with AWS Glue
D. Use standard SQL queries in Amazon DynamoDB to analyze the CtoudFront logs m the S3 bucket Visualize the results with Amazon QuickSight

Answer: D

NEW QUESTION 269
- (Exam Topic 3)
A company is building a new web-based customer relationship management application. The application will use several Amazon EC2 instances that are backed by Amazon Elastic Block Store (Amazon EBS) volumes behind an Application Load Balancer (ALB). The application will also use an Amazon Aurora database. All data for the application must be encrypted at rest and in transit.
Which solution will meet these requirements?

A. Use AWS Key Management Service (AWS KMS) certificates on the ALB to encrypt data in transi
B. Use AWS Certificate Manager (ACM) to encrypt the EBS volumes and Aurora database storage at rest.
C. Use the AWS root account to log in to the AWS Management Consol
D. Upload the company's encryption certificate
E. While in the root account, select the option to turn on encryption for all data at rest and in transit for the account.
F. Use a AWS Key Management Service (AWS KMS) to encrypt the EBS volumes and Aurora database storage at res
G. Attach an AWS Certificate Manager (ACM) certificate to the ALB to encrypt data in transit.
H. Use BitLocker to encrypt all data at res
I. Import the company's TLS certificate keys to AWS key Management Service (AWS KMS). Attach the KMS keys to the ALB to encrypt data in transit.

Answer: C

NEW QUESTION 272
- (Exam Topic 3)
A financial company hosts a web application on AWS. The application uses an Amazon API Gateway Regional API endpoint to give users the ability to retrieve current stock prices. The company's security team has noticed an increase in the number of API requests. The security team is concerned that HTTP flood attacks might take the application offline.

A solutions architect must design a solution to protect the application from this type of attack. Which solution meats these requirements with the LEAST operational overhead?

A. Create an Amazon CloudFront distribution in front of the API Gateway Regional API endpoint with a maximum TTL of 24 hours
B. Create a Regional AWS WAF web ACL with a rate-based rul
C. Associate the web ACL with the API Gateway stage.
D. Use Amazon CloudWatch metrics to monitor the Count metric and alert the security team when the predefined rate is reached
E. Create an Amazon CloudFront distribution with Lambda@Edge in front of the API Gateway Regional API endpoint Create an AWS Lambda function to block requests from IP addresses that exceed the predefined rate.

Answer: B

NEW QUESTION 277
- (Exam Topic 3)
A company hosts multiple production applications. One of the applications consists of resources from Amazon EC2, AWS Lambda, Amazon RDS, Amazon Simple Notification Service (Amazon SNS), and Amazon Simple Queue Service (Amazon SQS) across multiple AWS Regions. All company resources are tagged with a tag name of "application" and a value that corresponds to each application. A solutions architect must provide the quickest solution for identifying all of the tagged components.
Which solution meets these requirements?

A. Use AWS CloudTrail to generate a list of resources with the application tag.
B. Use the AWS CLI to query each service across all Regions to report the tagged components.
C. Run a query in Amazon CloudWatch Logs Insights to report on the components with the application tag.
D. Run a query with the AWS Resource Groups Tag Editor to report on the resources globally with the application tag.

Answer: D

Explanation:
https://docs.aws.amazon.com/tag-editor/latest/userguide/tagging.html

NEW QUESTION 278
- (Exam Topic 3)
A solutions architect is designing the architecture for a software demonstration environment The environment will run on Amazon EC2 instances in an Auto Scaling group behind an Application Load Balancer (ALB) The system will experience significant increases in traffic during working hours but Is not required to operate on weekends.
Which combination of actions should the solutions architect take to ensure that the system can scale to meet demand? (Select TWO)

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

Answer: DE

NEW QUESTION 283
- (Exam Topic 3)
A solutions architect is designing a new API using Amazon API Gateway that will receive requests from users. The volume of requests is highly variable; several hours can pass without receiving a single request. The data processing will take place asynchronously, but should be completed within a few seconds after a request is made.
Which compute service should the solutions architect have the API invoke to deliver the requirements at the lowest cost?

A. An AWS Glue job
B. An AWS Lambda function
C. A containerized service hosted in Amazon Elastic Kubernetes Service (Amazon EKS)
D. A containerized service hosted in Amazon ECS with Amazon EC2

Answer: B

Explanation:
API Gateway + Lambda is the perfect solution for modern applications with serverless architecture.

NEW QUESTION 285
- (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 286
- (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 290
......

Thank You for Trying Our Product

* 100% Pass or Money Back

    All our products come with a 90-day Money Back Guarantee.

* One year free update

    You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

    We currently serve more than 30,000,000 customers.

* Shop Securely

    All transactions are protected by VeriSign!

100% Pass Your SAA-C03 Exam with Our Prep Materials Via below:

https://www.certleader.com/SAA-C03-dumps.html