Number: AWS SAA-C03
Passing Score: 800
Time Limit: 120
File Version: 15.0

Exam Code: SAA-C03

Exam Name: AWS Certified Solutions Architect – Associate

Exam A

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

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

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

QUESTION 2
A solutions architect is designing the cloud architecture for a new application being deployed on AWS. The process should run in parallel while adding and removing application nodes as needed based on the number of jobs to be processed. The processor application is stateless. The solutions architect must ensure that the application is loosely coupled and the job items are durably stored. Which design should the solutions architect use?

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

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

QUESTION 3
A company hosts its web applications in the AWS Cloud. The company configures Elastic Load Balancers to use certificate that are imported into AWS Certificate Manager (ACM). The company's security team must be notified 30 days before the expiration of each certificate.
What should a solutions architect recommend to meet the requirement?

A. Add a rule m ACM to publish a custom message to an Amazon Simple Notification Service (Amazon SNS) topic every day beginning 30 days before any certificate will expire.

B. Create an AWS Config rule that checks for certificates that will expire within 30 days. Configure Amazon EventBridge (Amazon CloudWatch Events) to invoke a custom alert by way of Amazon Simple Notification Service (Amazon SNS) when AWS Config reports a noncompliant resource
C. Use AWS trusted Advisor to check for certificates that will expire within to days. Create an Amazon CloudWatch alarm that is based on Trusted Advisor metrics for check status changes Configure the alarm to send a custom alert by way of Amazon Simple rectification Service (Amazon SNS)
D. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to detect any certificates that will expire within 30 days. Configure the rule to invoke an AWS Lambda function. Configure the Lambda function to send a custom alert by way of Amazon Simple Notification Service (Amazon SNS).

Correct Answer: B
Section:
Explanation:
https://aws.amazon.com/premiumsupport/knowledge-center/acm-certificate-expiration/

QUESTION 4
A company's dynamic website is hosted using on-premises servers in the United States. The company is launching its product in Europe, and it wants to optimize site loading times for new European users. The site's backend must remain in the United States. The product is being launched in a few days, and an immediate solution is needed.
What should the solutions architect recommend?

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

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

QUESTION 5
A company wants to reduce the cost of its existing three-tier web architecture. The web, application, and database servers are running on Amazon EC2 instances for the development, test, and production environments. The EC2 instances average 30% CPU utilization during peak hours and 10% CPU utilization during non-peak hours.
The production EC2 instances run 24 hours a day. The development and test EC2 instances run for at least 8 hours each day. The company plans to implement automation to stop the development and test EC2 instances when they are not in use.
Which EC2 instance purchasing solution will meet the company's requirements MOST costeffectively?

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

Correct Answer: B
Section:

QUESTION 6
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 bucket. Configure an S3 Lifecycle policy to archive the documents periodically.

C. Store the uploaded documents in an Amazon S3 bucket with S3 Versioning enabled Configure an ACL to restrict all access to read-only.
D. Store the uploaded documents on an Amazon Elastic File System (Amazon EFS) volume. Access the data by mounting the volume in read-only mode.

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

QUESTION 7
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 system. The web server should be able to decrypt the files and access the database

Correct Answer: A
Section:
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.

QUESTION 8
A company is hosting a web application on AWS using a single Amazon EC2 instance that stores useruploaded documents in an Amazon EBS volume. For better scalability and availability, the company duplicated the architecture and created a second EC2 instance and EBS volume in another Availability Zone placing both behind an Application Load Balancer After completing this change, users reported that, each time they refreshed the website, they could see one subset of their documents or the other, but never all of the documents at the same time.
What should a solutions architect propose to ensure users see all of their documents at once?

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

Correct Answer: C
Section:
Explanation:

QUESTION 9
A company uses NFS to store large video files in on-premises network attached storage. Each video file ranges in size from 1MB to 500 GB. The total storage is 70 TB and is no longer growing. The company decides to migrate the video files to Amazon S3. The company must migrate the video files as soon as possible while using the least possible network bandwidth. Which solution will meet these requirements?

A. Create an S3 bucket Create an IAM role that has permissions to write to the S3 bucket. Use the AWS CLI to copy all files locally to the S3 bucket.

B. Create an AWS Snowball Edge job. Receive a Snowball Edge device on premises. Use the Snowball Edge client to transfer data to the device. Return the device so that AWS can import the data into Amazon S3.
C. Deploy an S3 File Gateway on premises. Create a public service endpoint to connect to the S3 File Gateway Create an S3 bucket Create a new NFS file share on the S3 File Gateway Point the new file share to the S3 bucket. Transfer the data from the existing NFS file share to the S3 File Gateway.
D. Set up an AWS Direct Connect connection between the on-premises network and AWS. Deploy an S3 File Gateway on premises. Create a public virtual interlace (VIF) to connect to the S3 File Gateway. Create an S3 bucket. Create a new NFS file share on the S3 File Gateway. Point the new file share to the S3 bucket. Transfer the data from the existing NFS file share to the S3 File Gateway.

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

QUESTION 10
A company has an application that ingests incoming messages. These messages are then quickly consumed by dozens of other applications and microservices. The number of messages varies drastically and sometimes spikes as high as 100,000 each second.
The company wants to decouple the solution and increase scalability.
Which solution meets these requirements?

A. Persist the messages to Amazon Kinesis Data Analytics. All the applications will read and process the messages.
B. Deploy the application on Amazon EC2 instances in an Auto Scaling group, which scales the number of EC2 instances based on CPU metrics.
C. Write the messages to Amazon Kinesis Data Streams with a single shard. All applications will read from the stream and process the messages.
D. Publish the messages to an Amazon Simple Notification Service (Amazon SNS) topic with one or more Amazon Simple Queue Service (Amazon SQS) subscriptions. All applications then process the messages from the queues.

Correct Answer: D
Section:
Explanation:
https://aws.amazon.com/sqs/features/
By routing incoming requests to Amazon SQS, the company can decouple the job requests from the processing instances. This allows them to scale the number of instances based on the size of the queue, providing more resources when needed. Additionally, using an Auto Scaling group based on the queue size will automatically scale the number of instances up or down depending on the workload. Updating the software to read from the queue will allow it to process the job requests in a more efficient manner, improving the performance of the system.

QUESTION 11
A company is migrating a distributed application to AWS The application serves variable workloads The legacy platform consists of a primary server trial coordinates jobs across multiple compute nodes The company wants to modernize the application with a solution that maximizes resiliency and scalability How should a solutions architect design the architecture to meet these requirements?

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

Correct Answer: B
Section:

QUESTION 12
A company is running an SMB file server in its data center. The file server stores large files that are accessed frequently for the first few days after the files are created. After 7 days the files are rarely accessed. The total data size is increasing and is close to the company's total storage capacity. A solutions architect must increase the company's available storage space without losing low-latency access to the most recently accessed files. The solutions architect must also provide file lifecycle management to avoid future storage issues.
Which solution will meet these requirements?

A. Use AWS DataSync to copy data that is older than 7 days from the SMB file server to AWS.
B. Create an Amazon S3 File Gateway to extend the company's storage space. Create an S3 Lifecycle policy to transition the data to S3 Glacier Deep Archive after 7 days.
C. Create an Amazon FSx for Windows File Server file system to extend the company's storage space.
D. 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.

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

QUESTION 13
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 order. Subscribe an AWS Lambda function to the topic to perform processing.
B. Use an API Gateway integration to send a message to an Amazon Simple Queue Service (Amazon SQS) FIFO queue when the application receives an order. Configure the SQS FIFO queue to invoke an AWS Lambda function for processing.
C. Use an API Gateway authorizer to block any requests while the application processes an order.
D. Use an API Gateway integration to send a message to an Amazon Simple Queue Service (Amazon SQS) standard queue when the application receives an order. Configure the SQS standard queue to invoke an AWS Lambda function for processing.

Correct Answer: B
Section:

QUESTION 14
A company has an application that runs on Amazon EC2 instances and uses an Amazon Aurora database. The EC2 instances connect to the database by using user names and passwords that are stored locally in a file. The company wants to minimize the operational overhead of credential management.
What should a solutions architect do to accomplish this goal?

A. Use AWS Secrets Manager. Turn on automatic rotation.
B. Use AWS Systems Manager Parameter Store. Turn on automatic rotation.
C. Create an Amazon S3 bucket lo store objects that are encrypted with an AWS Key C. Management Service (AWS KMS) encryption key. Migrate the credential file to the S3 bucket. Point the application to the S3 bucket.
D. Create an encrypted Amazon Elastic Block Store (Amazon EBS) volume (or each EC2 instance.
Attach the new EBS volume to each EC2 instance. Migrate the credential file to the new EBS volume.
Point the application to the new EBS volume.

Correct Answer: A

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

QUESTION 15
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 dat a. The company stores its static data in an Amazon S3 bucket. The company wants to improve performance and reduce latency for the static data and dynamic data. The company is using its own domain name registered with Amazon Route 53. What should a solutions architect do to meet these requirements?

A. Create an Amazon CloudFront distribution that has the S3 bucket and the ALB as origins Configure Route 53 to route traffic to the CloudFront distribution.
B. Create an Amazon CloudFront distribution that has the ALB as an origin Create an AWS Global Accelerator standard accelerator that has the S3 bucket as an endpoint. Configure Route 53 to route traffic to the CloudFront distribution.
C. 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.
D. Create an Amazon CloudFront distribution that has the ALB as an origin C. Create an AWS Global Accelerator standard accelerator that has the S3 bucket as an endpoint Create two domain names. 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.

Correct Answer: A
Section:
Explanation:

QUESTION 16
A company performs monthly maintenance on its AWS infrastructure. During these maintenance activities, the company needs to rotate the credentials tor its Amazon ROS tor MySQL databases across multiple AWS Regions Which solution will meet these requirements with the LEAST operational overhead?

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

Correct Answer: A
Section:
Explanation:
https://aws.amazon.com/blogs/security/how-to-replicate-secrets-aws-secrets-manager-multiple- regions/

QUESTION 17
A company runs an ecommerce application on Amazon EC2 instances behind an Application Load Balancer. The instances run in an Amazon EC2 Auto Scaling group across multiple Availability Zones. The Auto Scaling group scales based on CPU utilization metrics. The ecommerce application stores the transaction data in a MySQL 8.0 database that is hosted on a large EC2 instance. The database's performance degrades quickly as application load increases. The application handles more read requests than write transactions. The company wants a solution that will automatically scale the database to meet the demand of unpredictable read workloads while maintaining high availability.
Which solution will meet these requirements?

A. Use Amazon Redshift with a single node for leader and compute functionality.
B. Use Amazon RDS with a Single-AZ deployment Configure Amazon RDS to add reader instances in a different Availability Zone.
C. Use Amazon Aurora with a Multi-AZ deployment. Configure Aurora Auto Scaling with Aurora Replicas.
D. Use Amazon ElastiCache for Memcached with EC2 Spot Instances.

Correct Answer: C
Section:
Explanation:
AURORA is 5x performance improvement over MySQL on RDS and handles more read requests than write,; maintaining high availability = Multi-AZ deployment

QUESTION 18
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 volumes. Restore the snapshots onto EC2 instance store volumes in the test environment.
B. Configure the production EBS volumes to use the EBS Multi-Attach feature. Take EBS snapshots of the production EBS volumes. Attach the production EBS volumes to the EC2 instances in the test environment.
C. Take EBS snapshots of the production EBS volumes. Create and initialize new EBS volumes. Attach the new EBS volumes to EC2 instances in the test environment before restoring the volumes from the production EBS snapshots.
D. Take EBS snapshots of the production EBS volumes. Turn on the EBS fast snapshot restore feature on the EBS snapshots. Restore the snapshots into new EBS volumes. Attach the new EBS volumes to EC2 instances in the test environment.

Correct Answer: D
Section:
Explanation:

QUESTION 19
An ecommerce company wants to launch a one-deal-a-day website on AWS. Each day will feature exactly one product on sale for a period of 24 hours. The company wants to be able to handle millions of requests each hour with millisecond latency during peak hours.
Which solution will meet these requirements with the LEAST operational overhead?

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

Correct Answer: D
Section:
Explanation:
To launch a one-deal-a-day website on AWS with millisecond latency during peak hours and with the least operational overhead, the best option is to use an Amazon S3 bucket to host the website's static content, deploy an Amazon CloudFront distribution, set the S3 bucket as the origin, use Amazon API

QUESTION 20
A solutions architect is using Amazon S3 to design the storage architecture of a new digital media application. The media files must be resilient to the loss of an Availability Zone Some files are accessed frequently while other files are rarely accessed in an unpredictable pattern. The solutions architect must minimize the costs of storing and retrieving the media files. Which storage option meets these requirements?

A. S3 Standard
B. S3 Intelligent-Tiering
C. S3 Standard-Infrequent Access {S3 Standard-IA)

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

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

QUESTION 21
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.

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

QUESTION 22
A company observes an increase in Amazon EC2 costs in its most recent bill The billing team notices unwanted vertical scaling of instance types for a couple of EC2 instances A solutions architect needs to create a graph comparing the last 2 months of EC2 costs and perform an in-depth analysis to identify the root cause of the vertical scaling How should the solutions architect generate the information with the LEAST operational overhead?

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

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

QUESTION 23
A company is designing an application. The application uses an AWS Lambda function to receive information through Amazon API Gateway and to store the information in an Amazon Aurora PostgreSQL database. During the proof-of-concept stage, the company has to increase the Lambda quotas significantly to handle the high volumes of data that the company needs to load into the database. A solutions architect must recommend a new design

to improve scalability and minimize the configuration effort.
Which solution will meet these requirements?

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

Correct Answer: D
Section:
Explanation:

QUESTION 24
A company needs to review its AWS Cloud deployment to ensure that its Amazon S3 buckets do not have unauthorized configuration changes. What should a solutions architect do to accomplish this goal?

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

Correct Answer: A
Section:
Explanation:

QUESTION 25
A company is launching a new application and will display application metrics on an Amazon CloudWatch dashboard. The company's product manager needs to access this dashboard periodically. The product manager does not have an AWS account. A solution architect must provide access to the product manager by following the principle of least privilege. Which solution will meet these requirements?

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

Correct Answer: A
Section:
Explanation:

QUESTION 26
A company is migrating applications to AWS. The applications are deployed in different accounts. The company manages the accounts centrally by using AWS Organizations. The company's security team needs a single sign-on (SSO) solution across all the company's accounts. The company must continue managing the users and groups in its on-premises self-managed Microsoft Active Directory. Which solution will meet these requirements?

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

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

QUESTION 27
A company provides a Voice over Internet Protocol (VoIP) service that uses UDP connections. The service consists of Amazon EC2 instances that run in an Auto Scaling group. The company has deployments across multiple AWS Regions. The company needs to route users to the Region with the lowest latency. The company also needs automated failover between Regions. Which solution will meet these requirements?

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

Correct Answer: A
Section:
Explanation:
https://aws.amazon.com/global-accelerator/faqs/

QUESTION 28
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 completed. Restart the DB instance when required.
B. Use an Auto Scaling policy with the DB instance to automatically scale when tests are completed.
C. Create a snapshot when tests are completed. Terminate the DB instance and restore the snapshot when required.
D. Modify the DB instance to a low-capacity instance when tests are completed. Modify the DB instance again when required.

Correct Answer: C
Section:
Explanation:
To reduce the cost of running the tests without reducing the compute and memory attributes of the Amazon RDS for MySQL DB instance, the development team can stop the instance when tests are completed and restart it when required. Stopping the DB instance when not in use can help save costs because customers are only charged for storage while the DB instance is stopped. During this time, automated backups and automated DB instance maintenance are suspended. When the instance is restarted, it retains the same configurations, security groups, and DB parameter groups as when it was stopped.Reference: Amazon RDS Documentation: Stopping and Starting a DB instance(https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_StopInstance.html)

QUESTION 29
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

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

QUESTION 30
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 creation. Delete the fees 4 years after object creation
C. Create an S3 bucket lifecycle policy to move files from S3 Standard-infrequent Access (S3 Standard -lA) 30 from object creation. Delete the ties 4 years after object creation
D. 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.

Correct Answer: C
Section:

QUESTION 31
A company wants to manage Amazon Machine Images (AMIs). The company currently copies AMIs to the same AWS Region where the AMIs were created. The company needs to design an application that captures AWS API calls and sends alerts whenever the Amazon EC2 Createlmage API operation is called within the company's account. Which solution will meet these requirements with the LEAST operational overhead?

A. Create an AWS Lambda function to query AWS CloudTrail logs and to send an alert when a Createlmage API call is detected.
B. Configure AWS CloudTrail with an Amazon Simple Notification Service {Amazon SNS) notification that occurs when updated logs are sent to Amazon S3. Use Amazon Athena to create a new table and to query on Createlmage when an API call is detected.
C. Create an Amazon EventBridge (Amazon CloudWatch Events) rule for the Createlmage API call.
Configure the target as an Amazon Simple Notification Service (Amazon SNS) topic to send an alert when a Createlmage API call is detected.
D. Configure an Amazon Simple Queue Service (Amazon SQS) FIFO queue as a target for AWS CloudTrail logs. Create an AWS Lambda function to send an alert to an Amazon Simple Notification Service (Amazon SNS) topic when a Createlmage API call is detected.

Correct Answer: C
Section:
Explanation:

QUESTION 32
A company has a financial application that produces reports. The reports average 50 KB in size and are stored in Amazon S3. The reports are frequently accessed during the first week after production and must be stored for several years. The reports must be retrievable within 6 hours.
Which solution meets these requirements MOST cost-effectively?

A. Use S3 Standard. Use an S3 Lifecycle rule to transition the reports to S3 Glacier after 7 days.
B. Use S3 Standard. Use an S3 Lifecycle rule to transition the reports to S3 Standard-Infrequent Access (S3 Standard-IA) after 7 days.
C. Use S3 Intelligent-Tiering. Configure S3 Intelligent-Tiering to transition the reports to S3 Standard-Infrequent Access (S3 Standard-IA) and S3 Glacier.
D. Use S3 Standard. Use an S3 Lifecycle rule to transition the reports to S3 Glacier Deep Archive after 7 days.

Correct Answer: A
Section:
Explanation:
To store and retrieve reports that are frequently accessed during the first week and must be stored for several years, S3 Standard and S3 Glacier are suitable solutions. S3 Standard offers high durability, availability, and performance for frequently accessed data. S3 Glacier offers secure and durable storage for long-term data archiving at a low cost. S3 Lifecycle rules can be used to transition the reports from S3 Standard to S3 Glacier after 7 days, which can reduce storage costs. S3 Glacier also supports retrieval within 6 hours.
Reference:
Storage Classes
Object Lifecycle Management
Retrieving Archived Objects from Amazon S3 Glacier

QUESTION 33
A company has data collection sensors at different locations. The data collection sensors stream a high volume of data to the company. The company wants to design a platform on AWS to ingest and process high-volume streaming dat a. The solution must be scalable and support data collection in near real time. The company must store the data in Amazon S3 for future reporting.
Which solution will meet these requirements with the LEAST operational overhead?

A. Use Amazon Kinesis Data Firehose to deliver streaming data to Amazon S3.
B. Use AWS Glue to deliver streaming data to Amazon S3.
C. Use AWS Lambda to deliver streaming data and store the data to Amazon S3.
D. Use AWS Database Migration Service (AWS DMS) to deliver streaming data to Amazon S3.

Correct Answer: A
Section:
Explanation:
To ingest and process high-volume streaming data with the least operational overhead, Amazon Kinesis Data Firehose is a suitable solution. Amazon Kinesis Data Firehose can capture, transform, and deliver streaming data to Amazon S3 or other destinations. Amazon Kinesis Data Firehose can scale automatically to match the throughput of the data and handle any amount of data. Amazon Kinesis Data Firehose is also a fully managed service that does not require any servers to provision or manage.
Reference:
What Is Amazon Kinesis Data Firehose?
Amazon Kinesis Data Firehose Pricing

QUESTION 34
A solutions architect is designing the storage architecture for a new web application used for storing and viewing engineering drawings. All application components will be deployed on the AWS infrastructure.
The application design must support caching to minimize the amount of time that users wait for the engineering drawings to load. The application must be able to store petabytes of dat a. Which combination of storage and caching should the solutions architect use?

A. Amazon S3 with Amazon CloudFront
B. Amazon S3 Glacier with Amazon ElastiCache

C. Amazon Elastic Block Store (Amazon EBS) volumes with Amazon CloudFront
D. AWS Storage Gateway with Amazon ElastiCache

Correct Answer: A
Section:
Explanation:
To store and view engineering drawings with caching support, Amazon S3 and Amazon CloudFront are suitable solutions. Amazon S3 can store any amount of data with high durability, availability, and performance. Amazon CloudFront can distribute the engineering drawings to edge locations closer to the users, which can reduce the latency and improve the user experience. Amazon CloudFront can also cache the engineering drawings at the edge locations, which can minimize the amount of time that users wait for the drawings to load.
Reference:
What Is Amazon S3?
What Is Amazon CloudFront?

QUESTION 35
A company runs its applications on Amazon EC2 instances that are backed by Amazon Elastic Block Store (Amazon EBS). The EC2 instances run the most recent Amazon Linux release. The applications are experiencing availability issues when the company's employees store and retrieve files that are 25 GB or larger. The company needs a solution that does not require the company to transfer files between EC2 instances. The files must be available across many EC2 instances and across multiple Availability Zones.
Which solution will meet these requirements?

A. Migrate all the files to an Amazon S3 bucket. Instruct the employees to access the files from the S3 bucket.
B. Take a snapshot of the existing EBS volume. Mount the snapshot as an EBS volume across the EC2 instances. Instruct the employees to access the files from the EC2 instances.
C. Mount an Amazon Elastic File System (Amazon EFS) file system across all the EC2 instances.
Instruct the employees to access the files from the EC2 instances.
D. Create an Amazon Machine Image (AMI) from the EC2 instances. Configure new EC2 instances from the AMI that use an instance store volume. Instruct the employees to access the files from the EC2 instances

Correct Answer: C
Section:
Explanation:
To store and access files that are 25 GB or larger across many EC2 instances and across multiple Availability Zones, Amazon Elastic File System (Amazon EFS) is a suitable solution. Amazon EFS provides a simple, scalable, elastic file system that can be mounted on multiple EC2 instances concurrently. Amazon EFS supports high availability and durability by storing data across multiple Availability Zones within a Region.
Reference:
What Is Amazon Elastic File System?
Using EFS with EC2

QUESTION 36
A solutions architect is using an AWS CloudFormation template to deploy a three-tier web application. The web application consists of a web tier and an application tier that stores and retrieves user data in Amazon DynamoDB tables. The web and application tiers are hosted on
Amazon EC2 instances, and the database tier is not publicly accessible. The application EC2 instances need to access the DynamoDB tables without exposing API credentials in the template.
What should the solutions architect do to meet these requirements?

A. Create an IAM role to read the DynamoDB tables. Associate the role with the application instances by referencing an instance profile.
B. Create an IAM role that has the required permissions to read and write from the DynamoDB tables. Add the role to the EC2 instance profile, and associate the instance profile with the application instances.
C. Use the parameter section in the AWS CloudFormation template to have the user input access and secret keys from an already-created IAM user that has the required permissions to read and write from the DynamoDB tables.
D. Create an IAM user in the AWS CloudFormation template that has the required permissions to read and write from the DynamoDB tables. Use the GetAtt function to retrieve the access and secret keys, and pass them to the application instances through the user data.

Correct Answer: B

Section:
Explanation:
it allows the application EC2 instances to access the DynamoDB tables without exposing API credentials in the template. By creating an IAM role that has the required permissions to read and write from the DynamoDB tables and adding it to the EC2 instance profile, the application instances can use temporary security credentials that are automatically rotated by AWS. This is a secure and best practice way to grant access to AWS resources from EC2 instances. Reference:
IAM Roles for Amazon EC2 Using Instance Profiles

QUESTION 37
A company hosts an application on multiple Amazon EC2 instances The application processes messages from an Amazon SQS queue writes to an Amazon RDS table and deletes the message from the queue Occasional duplicate records are found in the RDS table. The SQS queue does not contain any duplicate messages.
What should a solutions architect do to ensure messages are being processed once only?

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

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

QUESTION 38
A solutions architect is designing a new hybrid architecture to extend a company s on-premises infrastructure to AWS The company requires a highly available connection with consistent low latency to an AWS Region. The company needs to minimize costs and is willing to accept slower traffic if the primary connection fails.
What should the solutions architect do to meet these requirements?

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

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

QUESTION 39
A company is running a business-critical web application on Amazon EC2 instances behind an Application Load Balancer. The EC2 instances are in an Auto Scaling group. The application uses an Amazon Aurora PostgreSQL

database that is deployed in a single Availability Zone. The company wants the application to be highly available with minimum downtime and minimum loss of data. Which solution will meet these requirements with the LEAST operational effort?

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

Correct Answer: B
Section:
Explanation:

QUESTION 40
A company's HTTP application is behind a Network Load Balancer (NLB). The NLB's target group isconfigured to use an Amazon EC2 Auto Scaling group with multiple EC2 instances that run the webservice. The company notices that the NLB is not detecting HTTP errors for the application. These errorsrequire a manual restart of the EC2 instances that run the web service. The company needs toimprove the application's availability without writing custom scripts or code.
What should a solutions architect do to meet these requirements?

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

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

QUESTION 41
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 tables. For RPO recovery, point the application to a different AWS Region.
B. Configure DynamoDB point-in-time recovery. For RPO recovery, restore to the desired point in time.
C. Export the DynamoDB data to Amazon S3 Glacier on a daily basis. For RPO recovery, import the data from S3 Glacier to DynamoDB.
D. Schedule Amazon Elastic Block Store (Amazon EBS) snapshots for the DynamoDB table every 15 minutes. For RPO recovery, restore the DynamoDB table by using the EBS snapshot.

Correct Answer: B
Section:
Explanation:
https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/PointInTimeRecovery.html

QUESTION 42

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.

Correct Answer: D
Section:
Explanation:

QUESTION 43
A company recently launched Linux-based application instances on Amazon EC2 in a private subnet and launched a Linux-based bastion host on an Amazon EC2 instance in a public subnet of a VPC A solutions architect needs to connect from the on-premises network, through the company's internet connection to the bastion host and to the application servers The solutions architect must make sure that the security groups of all the EC2 instances will allow that access Which combination of steps should the solutions architect take to meet these requirements? (Select TWO)

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

Correct Answer: C, D
Section:
Explanation:
https://digitalcloud.training/ssh-into-ec2-in-private-subnet/

QUESTION 44
A solutions architect is designing a two-tier web application The application consists of a publicfacing web tier hosted on Amazon EC2 in public subnets The database tier consists of Microsoft SQL Server running on Amazon EC2 in a private subnet Security is a high priority for the company How should security groups be configured in this situation? (Select TWO)

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

Correct Answer: A, C
Section:
Explanation:
https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/security-group-rules-reference.html

QUESTION 45
A company wants to move a multi-tiered application from on premises to the AWS Cloud to improve the application's performance. The application consists of application tiers that communicate with each other by way of RESTful services. Transactions are dropped when one tier becomes overloaded.
A solutions architect must design a solution that resolves these issues and modernizes the application. Which solution meets these requirements and is the MOST operationally efficient?

A. Use Amazon API Gateway and direct transactions to the AWS Lambda functions as the application layer. Use Amazon Simple Queue Service (Amazon SQS) as the communication layer between application services. Most Voted
B. Use Amazon CloudWatch metrics to analyze the application performance history to determine the server's peak utilization during the performance failures. Increase the size of the application server's Amazon EC2 instances to meet the peak requirements.
C. Use Amazon Simple Notification Service (Amazon SNS) to handle the messaging between application servers running on Amazon EC2 in an Auto Scaling group. Use Amazon CloudWatch to monitor the SNS queue length and scale up and down as required.
D. Use Amazon Simple Queue Service (Amazon SQS) to handle the messaging between application servers running on Amazon EC2 in an Auto Scaling group. Use Amazon CloudWatch to monitor the SQS queue length and scale up when communication failures are detected.

Correct Answer: A
Section:
Explanation:
https://aws.amazon.com/getting-started/hands-on/build-serverless-web-app-lambda-apigateways3- dynamodb-cognito/module-4/ Build a Serverless Web Application with AWS Lambda, Amazon API Gateway, AWS Amplify, Amazon DynamoDB, and Amazon Cognito. This example showed similar setup as question: Build a Serverless Web Application with AWS Lambda, Amazon API Gateway, AWS Amplify, Amazon DynamoDB, and Amazon Cognito

QUESTION 46
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 onpremises 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

Correct Answer: B
Section:
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/

QUESTION 47
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 source. Use AWS Lambda functions to transform the data. Use the Kinesis Data Firehose delivery stream to send the data to Amazon S3.
B. Deploy an Amazon EC2 instance to host an API that sends data to AWS Glue. Stop source/destination checking on the EC2 instance. Use AWS Glue to transform the data and to send the data to Amazon S3.
C. Configure an Amazon API Gateway API to send data to an Amazon Kinesis data stream. Create an Amazon Kinesis Data Firehose delivery stream that uses the Kinesis data stream as a data source. Use AWS Lambda functions to transform the data. Use the Kinesis Data Firehose delivery stream to send the data to Amazon S3.
D. Configure an Amazon API Gateway API to send data to AWS Glue. Use AWS Lambda functions to transform the data. Use AWS Glue to send the data to Amazon S3.

Correct Answer: C
Section:

QUESTION 48
A company needs to keep user transaction data in an Amazon DynamoDB table.
The company must retain the data for 7 years.
What is the MOST operationally efficient solution that meets these requirements?

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

Correct Answer: B
Section:

QUESTION 49
A company is planning to use an Amazon DynamoDB table for data storage. The company is concerned about cost optimization. The table will not be used on most mornings. In the evenings, the read and write traffic will often be unpredictable. When traffic spikes occur, they will happen very quickly.
What should a solutions architect recommend?

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

Correct Answer: A
Section:

QUESTION 50
A company hosts an application on AWS Lambda functions mat are invoked by an Amazon API Gateway API The Lambda functions save customer data to an Amazon Aurora MySQL database Whenever the company upgrades the database, the Lambda functions fail to establish database connections until the upgrade is complete The result is that customer data Is not recorded for some of the event A solutions architect needs to design a solution that stores customer data that is created during database upgrades Which solution will meet these requirements?

A. Provision an Amazon RDS proxy to sit between the Lambda functions and the database Configure the Lambda functions to connect to the RDS proxy
B. Increase the run time of me Lambda functions to the maximum Create a retry mechanism in the code that stores the customer data in the database
C. Persist the customer data to Lambda local storage. Configure new Lambda functions to scan the local storage to save the customer data to the database.
D. Store the customer data m an Amazon Simple Queue Service (Amazon SOS) FIFO queue Create a new Lambda function that polls the queue and stores the customer data in the database

Correct Answer: D
Section:
Explanation:
https://www.learnaws.org/2020/12/13/aws-rds-proxy-deep-dive/
RDS proxy can improve application availability in such a situation by waiting for the new database instance to be functional and maintaining any requests received from the application during this time. The end result is that the application is more resilient to issues with the underlying database.
This will enable solution to hold data till the time DB comes back to normal. RDS proxy is to optimally utilize the connection between Lambda and DB. Lambda can open multiple connection concurrently which can be taxing on DB compute resources, hence RDS proxy was introduced to manage and leverage these connections efficiently.

QUESTION 51
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

Correct Answer: A
Section:
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

QUESTION 52
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.

Correct Answer: A
Section:

QUESTION 53
A company is using a SQL database to store movie data that is publicly accessible. The database runs on an Amazon RDS Single-AZ DB instance A script runs queries at random intervals each day to record the number of new movies that have been added to the database. The script must report a final total during business hours The company's development team notices that the database performance is inadequate for development tasks when the script is running. A solutions architect must recommend a solution to resolve this issue. Which solution will meet this requirement with the LEAST operational overhead?

A. Modify the DB instance to be a Multi-AZ deployment
B. Create a read replica of the database Configure the script to query only the read replica
C. Instruct the development team to manually export the entries in the database at the end of each day
D. Use Amazon ElastiCache to cache the common queries that the script runs against the database

Correct Answer: B
Section:

QUESTION 54
A company has applications that run on Amazon EC2 instances in a VPC. One of the applications needs to call the Amazon S3 API to store and read objects. According to the company's security regulations, no traffic from the applications is allowed to travel across the internet.
Which solution will meet these requirements?

A. Configure an S3 interface endpoint.
B. Configure an S3 gateway endpoint.
C. Create an S3 bucket in a private subnet.
D. Create an S3 bucket in the same Region as the EC2 instance.

Correct Answer: A
Section:
Explanation:
https://docs.aws.amazon.com/vpc/latest/privatelink/gateway-endpoints.html

QUESTION 55
A company is storing sensitive user information in an Amazon S3 bucket The company wants to provide secure access to this bucket from the application tier running on Ama2on EC2 instances inside a VPC Which combination of steps should a solutions architect take to accomplish this? (Select TWO.)

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

Correct Answer: A, C
Section:
Explanation:
https://aws.amazon.com/premiumsupport/knowledge-center/s3-private-connection-noauthentication/

QUESTION 56
A company runs an on-premises application that is powered by a MySQL database The company is migrating the application to AWS to Increase the application's elasticity and availability The current architecture shows heavy read activity on the database during times of normal operation Every 4 hours the company's development team pulls a full export of the production database to populate a database in the staging environment During this period, users experience unacceptable application latency The development team is unable to use the staging environment until the procedure completes A solutions architect must recommend replacement architecture that alleviates the application latency issue
The replacement architecture also must give the development team the ability to continue using the staging environment without delay Which solution meets these requirements?

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

Correct Answer: B
Section:

QUESTION 57
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 data. Store the resulting JSON file in an Amazon Aurora DB cluster.
B. Configure Amazon S3 to send an event notification to an Amazon Simple Queue Service (Amazon SQS) queue. Use Amazon EC2 instances to read from the queue and process the data. Store the resulting JSON file in Amazon DynamoDB.
C. Configure Amazon S3 to send an event notification to an Amazon Simple Queue Service (Amazon SQS) queue. Use an AWS Lambda function to read from the queue and process the data. Store the resulting JSON file in Amazon DynamoDB. Most Voted
D. Configure Amazon EventBridge (Amazon CloudWatch Events) to send an event to Amazon Kinesis Data Streams when a new file is uploaded. Use an AWS Lambda function to consume the event from the stream and process the data. Store the resulting JSON file in Amazon Aurora DB cluster.

Correct Answer: C
Section:
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-functionto- event-notifications-from-s3-buckets-in-different-aws-regions.html

QUESTION 58
An application allows users at a company's headquarters to access product dat a. The product data is stored in an Amazon RDS MySQL DB instance. The operations team has isolated an application performance slowdown and wants to separate read traffic from write traffic.
A solutions architect needs to optimize the application's performance quickly.
What should the solutions architect recommend?

A. Change the existing database to a Multi-AZ deployment. Serve the read requests from the primary Availability Zone.
B. Change the existing database to a Multi-AZ deployment. Serve the read requests from the secondary Availability Zone.
C. Create read replicas for the database. Configure the read replicas with half of the compute and storage resources as the source database.
D. Create read replicas for the database. Configure the read replicas with the same compute and storage resources as the source database.

Correct Answer: D
Section:
Explanation:
https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_MySQL.Replication.ReadReplica s.html

QUESTION 59
An Amazon EC2 administrator created the following policy associated with an IAM group containing several users

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

What is the effect of this policy?

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

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

QUESTION 60
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

Correct Answer: D
Section:

QUESTION 61
An image-processing company has a web application that users use to upload images. The application uploads the images into an Amazon S3 bucket. The company has set up S3 event notifications to publish the object creation events to an Amazon Simple Queue Service (Amazon SQS) standard queue. The SQS queue serves as the event source for an AWS Lambda function that processes the images and sends the results to users through email. Users report that they are receiving multiple email messages for every uploaded image. A solutions architect determines that SQS messages are invoking the Lambda function more than once, resulting in multiple email messages. What should the solutions architect do to resolve this issue with the LEAST operational overhead?

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

Correct Answer: C
Section:

QUESTION 62
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 gateway. Create a file share that uses the required client protocol Connect the application server to the tile share.
B. Create an AWS Storage Gateway tape gateway Configure (apes to use Amazon S3 Connect the application server lo the tape gateway
C. Create an Amazon EC2 Windows instance Install and configure a Windows file share role on the instance. Connect the application server to the file share.
D. Create an Amazon FSx for Windows File Server tile system Attach the fie system to the origin server. Connect the application server to the tile system

Correct Answer: D

Section:
Explanation:
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

QUESTION 63
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 certificates. Manually update the certificates as needed. Control access to the data by using fine-grained IAM access.
B. Create an AWS Lambda function that uses the Python cryptography library to receive and perform encryption operations. Store the function in an Amazon S3 bucket.
C. Create an AWS Key Management Service (AWS KMS) customer managed key. Allow the EC2 role to use the KMS key for encryption operations. Store the encrypted data on Amazon S3.
D. Create an AWS Key Management Service (AWS KMS) customer managed key. Allow the EC2 role to use the KMS key for encryption operations. Store the encrypted data on Amazon Elastic Block Store (Amazon EBS) volumes.

Correct Answer: C
Section:

QUESTION 64
A solutions architect is designing a VPC with public and private subnets. The VPC and subnets use IPv4 CIDR blocks. There is one public subnet and one private subnet in each of three Availability Zones (AZs) for high availability. An internet gateway is used to provide internet access for the public subnets. The private subnets require access to the internet to allow Amazon EC2 instances to download software updates. What should the solutions architect do to enable Internet access for the private subnets?

A. Create three NAT gateways, one for each public subnet in each AZ. Create a private route table for each AZ that forwards non-VPC traffic to the NAT gateway in its AZ.
B. Create three NAT instances, one for each private subnet in each AZ. Create a private route table for each AZ that forwards non-VPC traffic to the NAT instance in its AZ.
C. Create a second internet gateway on one of the private subnets. Update the route table for the private subnets that forward non-VPC traffic to the private internet gateway.
D. Create an egress-only internet gateway on one of the public subnets. Update the route table for the private subnets that forward non-VPC traffic to the egress- only internet gateway.

Correct Answer: A
Section:
Explanation:
https://aws.amazon.com/about-aws/whats-new/2018/03/introducing-amazon-vpc-nat-gateway-inthe- aws-govcloud-usregion/#:~: text=NAT%20Gateway%20is%20a%20highly,instances%20in%20a%20private%20subnet. https://docs.aws.amazon.com/vpc/latest/userguide/vpc-nat-comparison.html

QUESTION 65
A company wants to migrate an on-premises data center to AWS. The data canter hosts an SFTP server that stores its data on an NFS-based file system. The server holds 200 GB of data that needs to be transferred. The server must be hosted on an Amazon EC2 instance that uses an Amazon Elastic File System (Amazon EFS) file system When combination of steps should a solutions architect take to automate this task? (Select TWO)

A. Launch the EC2 instance into the same Avalability Zone as the EFS fie system
B. install an AWS DataSync agent m the on-premises data center
C. Create a secondary Amazon Elastic Block Store (Amazon EBS) volume on the EC2 instance tor the data
D. Manually use an operating system copy command to push the data to the EC2 instance
E. Use AWS DataSync to create a suitable location configuration for the onprermises SFTP server

Correct Answer: B, E
Section:
Explanation:

AWS DataSync is an online data movement and discovery service that simplifies data migration and helps users quickly, easily, and securely move their file or object data to, from, and between AWS

QUESTION 66
A company has an AWS Glue extract. transform, and load (ETL) job that runs every day at the same time. The job processes XML data that is in an Amazon S3 bucket. New data is added to the S3 bucket every day. A solutions architect notices that AWS Glue is processing all the data during each run. What should the solutions architect do to prevent AWS Glue from reprocessing old data?

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

Correct Answer: A
Section:
Explanation:

QUESTION 67
A solutions architect must design a highly available infrastructure for a website. The website is powered by Windows web servers that run on Amazon EC2 instances. The solutions architect must implement a solution that can mitigate a large-scale DDoS attack that originates from thousands of IP addresses. Downtime is not acceptable for the website. Which actions should the solutions architect take to protect the website from such an attack? (Select TWO.)

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

Correct Answer: A, C
Section:
Explanation:
https://aws.amazon.com/cloudfron

QUESTION 68
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.

Correct Answer: D
Section:
Explanation:
https://docs.aws.amazon.com/eventbridge/latest/userguide/resource-based-policieseventbridge. html#lambda-permissions

QUESTION 69

A company is preparing to store confidential data in Amazon S3 For compliance reasons the data must be encrypted at rest Encryption key usage must be logged tor auditing purposes. Keys must be rotated every year. Which solution meets these requirements and «the MOST operationally efferent?

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

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

QUESTION 70
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 subnets. Deploy an RDS Multi-AZ DB instance in private subnets.
B. Configure a VPC with two private subnets and two NAT gateways across two Availability Zones. Deploy an Application Load Balancer in the private subnets.
C. 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.
D. Configure a VPC with one public subnet, one private subnet, and two NAT gateways across two Availability Zones. Deploy an Application Load Balancer in the public subnet.
E. Configure a VPC with two public subnets, two private subnets, and two NAT gateways across two Availability Zones. Deploy an Application Load Balancer in the public subnets.

Correct Answer: A, E
Section:
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.

QUESTION 71
A solutions architect needs to implement a solution to reduce a company's storage costs. All the company's data is in the Amazon S3 Standard storage class. The company must keep all data for at least 25 years. Data from the most recent 2 years must be highly available and immediately retrievable.
Which solution will meet these requirements?

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

Correct Answer: B
Section:

QUESTION 72
A media company is evaluating the possibility ot moving rts systems to the AWS Cloud The company needs at least 10 TB of storage with the maximum possible I/O performance for video processing. 300 TB of very durable storage for storing media content, and 900 TB of storage to meet requirements for archival media that is not in use anymore Which set of services should a solutions architect recommend to meet these requirements?

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

Correct Answer: A
Section:
Explanation:
https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/InstanceStorage.html

QUESTION 73
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.

Correct Answer: B
Section:
Explanation:

QUESTION 74
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).

Correct Answer: A, E
Section:

QUESTION 75
An application runs on Amazon EC2 instances across multiple Availability Zones The instances run in an Amazon EC2 Auto Scaling group behind an Application Load Balancer The application performs best when the CPU utilization of the EC2 instances is at or near 40%.
What should a solutions architect do to maintain the desired performance across all instances in the group?

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

Correct Answer: B
Section:
Explanation:
https://docs.aws.amazon.com/autoscaling/application/userguide/application-auto-scaling-targettracking. html

QUESTION 76
A company is developing a file-sharing application that will use an Amazon S3 bucket for storage. The company wants to serve all the files through an Amazon CloudFront distribution. The company does not want the files to be accessible through direct navigation to the S3 URL.
What should a solutions architect do to meet these requirements?

A. Write individual policies for each S3 bucket to grant read permission for only CloudFront access.
B. Create an IAM user. Grant the user read permission to objects in the S3 bucket. Assign the user to CloudFront.
C. 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).
D. Create an origin access identity (OAI). Assign the OAI to the CloudFront distribution. Configure the S3 bucket permissions so that only the OAI has read permission.

Correct Answer: D
Section:
Explanation:
https://aws.amazon.com/premiumsupport/knowledge-center/cloudfront-access-to-amazon-s3/ https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/private-contentrestricting- access-to-s3.html#private-content-restricting- access-to-s3-overview

QUESTION 77
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

Correct Answer: A
Section:
Explanation:
Cloudfront for rapid response and s3 to minimize infrastructure.

QUESTION 78
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 instance. Set up database replication to a different AWS Region.
B. Migrate the Oracle database to Amazon RDS for Oracle. Activate Cross-Region automated backups to replicate the snapshots to another AWS Region.

C. Migrate the Oracle database to Amazon RDS Custom for Oracle. Create a read replica for the database in another AWS Region.
D. Migrate the Oracle database to Amazon RDS for Oracle. Create a standby database in another Availability Zone.

Correct Answer: C
Section:
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

QUESTION 79
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 bucket. Load the data into the new S3 bucket. Use S3 Cross-Region Replication (CRR) to replicate encrypted objects to an S3 bucket in another Region. Use server-side encryption with AWS KMS multi-Region kays (SSE-KMS). Use Amazon Athena to query the data.
B. Create a new S3 bucket. Load the data into the new S3 bucket. Use S3 Cross-Region Replication (CRR) to replicate encrypted objects to an S3 bucket in another Region. Use server-side encryption with AWS KMS multi-Region keys (SSE-KMS). Use Amazon RDS to query the data.
C. Load the data into the existing S3 bucket. Use S3 Cross-Region Replication (CRR) to replicate encrypted objects to an S3 bucket in another Region. Use server-side encryption with Amazon S3 managed encryption keys (SSE-S3). Use Amazon Athena to query the data.
D. Load the data into the existing S3 bucket. Use S3 Cross-Region Replication (CRR) to replicate encrypted objects to an S3 bucket in another Region. Use server-side encryption with Amazon S3 managed encryption keys (SSE-S3). Use Amazon RDS to query the data.

Correct Answer: A
Section:

QUESTION 80
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 VPC. Update the route table to connect to the target service.
B. Ask the provider to create a virtual private gateway in its VPC. Use AWS PrivateLink to connect to the target service.
C. Create a NAT gateway in a public subnet of the company's VPC. Update the route table to connect to the target service.
D. Ask the provider to create a VPC endpoint for the target service. Use AWS PrivateLink to connect to the target service.

Correct Answer: D
Section:

QUESTION 81
A company is migrating its on-premises PostgreSQL database to Amazon Aurora PostgreSQL. The onpremises 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

Correct Answer: A, C
Section:

QUESTION 82
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 alerts. Configure AWS account alternate contacts in the AWS Organizations console or programmatically.
C. 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.
D. Configure all existing AWS accounts and all newly created accounts to use the same root user email address. Configure AWS account alternate contacts in the AWS Organizations console or programmatically.

Correct Answer: B
Section:
Explanation:

QUESTION 83
A company runs its ecommerce application on AWS. Every new order is published as a message in a RabbitMQ queue that runs on an Amazon EC2 instance in a single Availability Zone. These messages are processed by a different application that runs on a separate EC2 instance. This application stores the details in a PostgreSQL database on another EC2 instance. All the EC2 instances are in the same Availability Zone. The company needs to redesign its architecture to provide the highest availability with the least operational overhead. What should a solutions architect do to meet these requirements?

A. Migrate the queue to a redundant pair (active/standby) of RabbitMQ instances on Amazon MQ.
Create a Multi-AZ Auto Scaling group (or EC2 instances that host the application. Create another Multi-AZ Auto Scaling group for EC2 instances that host the PostgreSQL database.
B. Migrate the queue to a redundant pair (active/standby) of RabbitMQ instances on Amazon MQ.
Create a Multi-AZ Auto Scaling group for EC2 instances that host the application. Migrate the database to run on a Multi-AZ deployment of Amazon RDS for PostgreSQL.
C. Create a Multi-AZ Auto Scaling group for EC2 instances that host the RabbitMQ queue. Create another Multi-AZ Auto Scaling group for EC2 instances that host the application. Migrate the database to run on a Multi-AZ deployment of Amazon RDS fqjPostgreSQL.
D. Create a Multi-AZ Auto Scaling group for EC2 instances that host the RabbitMQ queue. Create another Multi-AZ Auto Scaling group for EC2 instances that host the application. Create a third Multi- AZ Auto Scaling group for EC2 instances that host the PostgreSQL database.

Correct Answer: B
Section:
Explanation:
Migrating to Amazon MQ reduces the overhead on the queue management. C and D are dismissed. Deciding between A and B means deciding to go for an AutoScaling group for EC2 or an RDS for Postgress (both multi- AZ). The RDS option has less operational impact, as provide as a service the tools and software required. Consider for instance, the effort to add an additional node like a read replica, to the DB. https://docs.aws.amazon.com/amazon-mq/latest/developer-guide/active-standby-broker- deployment.html https://aws.amazon.com/rds/postgresql/

QUESTION 84
A reporting team receives files each day in an Amazon S3 bucket. The reporting team manually reviews and copies the files from this initial S3 bucket to an analysis S3 bucket each day at the same time to use with Amazon QuickSight. Additional teams are starting to send more files in larger sizes to the initial S3 bucket.
The reporting team wants to move the files automatically analysis S3 bucket as the files enter the initial S3 bucket. The reporting team also wants to use AWS Lambda functions to run patternmatching code on the copied dat a. In addition, the reporting team wants to send the data files to a pipeline in Amazon SageMaker Pipelines.
What should a solutions architect do to meet these requirements with the LEAST operational overhead?

A. Create a Lambda function to copy the files to the analysis S3 bucket. Create an S3 event notification for the analysis S3 bucket. Configure Lambda and SageMaker Pipelines as destinations of the event notification. Configure s30bjectCreated:Put as the event type.
B. Create a Lambda function to copy the files to the analysis S3 bucket. Configure the analysis S3 bucket to send event notifications to Amazon EventBridge (Amazon CloudWatch Events). Configure an ObjectCreated rule in EventBridge (CloudWatch Events). Configure Lambda and SageMaker Pipelines as targets for the rule.
C. Configure S3 replication between the S3 buckets. Create an S3 event notification for the analysis S3 bucket. Configure Lambda and SageMaker Pipelines as destinations of the event notification. Configure s30bjectCreated:Put as the event type.
D. Configure S3 replication between the S3 buckets. Configure the analysis S3 bucket to send event notifications to Amazon EventBridge (Amazon CloudWatch Events). Configure an ObjectCreated rule in EventBridge (CloudWatch Events). Configure Lambda and SageMaker Pipelines as targets for the rule.

Correct Answer: D
Section:
Explanation:
This solution meets the requirements of moving the files automatically, running Lambda functions on the copied data, and sending the data files to SageMaker Pipelines with the least operational overhead. S3 replication can copy the files from the initial S3 bucket to the analysis S3 bucket as they arrive. The analysis S3 bucket can send event notifications to Amazon EventBridge (Amazon CloudWatch Events) when an object is created. EventBridge can trigger Lambda and SageMaker Pipelines as targets for the ObjectCreated rule. Lambda can run pattern-matching code on the copied data, and SageMaker Pipelines can execute a pipeline with the data files.Option A is incorrect because creating a Lambda function to copy the files to the analysis S3 bucket is not necessary when S3 replication can do that automatically. It also adds operational overhead to manage the Lambda function. Option B is incorrect because creating a Lambda function to copy the files to the analysis S3 bucket is not necessary when S3 replication can do that automatically. It also adds operational overhead to manage the Lambda function. Option C is incorrect because using S3 event notification with multiple destinations can result in throttling or delivery failures if there are too many events.Reference:https://aws.amazon.com/blogs/machine-learning/automate-feature-engineering-pipelines-with- amazon-sagemaker/ https://docs.aws.amazon.com/sagemaker/latest/dg/automating-sagemaker-with-eventbridge.html

QUESTION 85
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.

Correct Answer: A, C
Section:

QUESTION 86
A company runs a web-based portal that provides users with global breaking news, local alerts, and weather updates. The portal delivers each user a personalized view by using mixture of static and dynamic content. Content is served over HTTPS through an API server running on an Amazon EC2 instance behind an Application Load Balancer (ALB). The company wants the portal to provide this content to its users across the world as quickly as possible. How should a solutions architect design the application to ensure the LEAST amount of latency for all users?

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

Correct Answer: A
Section:
Explanation:

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

QUESTION 87
A gaming company is designing a highly available architecture. The application runs on a modified Linux kernel and supports only UDP-based traffic. The company needs the front-end tier to provide the best possible user experience. That tier must have low latency, route traffic to the nearest edge location, and provide static IP addresses for entry into the application endpoints. What should a solutions architect do to meet these requirements?

A. Configure Amazon Route 53 to forward requests to an Application Load Balancer. Use AWS Lambda for the application in AWS Application Auto Scaling.
B. Configure Amazon CloudFront to forward requests to a Network Load Balancer. Use AWS Lambda for the application in an AWS Application Auto Scaling group.
C. Configure AWS Global Accelerator to forward requests to a Network Load Balancer. Use Amazon EC2 instances for the application in an EC2 Auto Scaling group.
D. Configure Amazon API Gateway to forward requests to an Application Load Balancer. Use Amazon EC2 instances for the application in an EC2 Auto Scaling group.

Correct Answer: C
Section:

QUESTION 88
A company wants to migrate its existing on-premises monolithic application to AWS.
The company wants to keep as much of the front- end code and the backend code as possible.
However, the company wants to break the application into smaller applications. A different team will manage each application. The company needs a highly scalable solution that minimizes operational overhead. Which solution will meet these requirements?

A. Host the application on AWS Lambda Integrate the application with Amazon API Gateway.
B. Host the application with AWS Amplify. Connect the application to an Amazon API Gateway API that is integrated with AWS Lambda.
C. Host the application on Amazon EC2 instances. Set up an Application Load Balancer with EC2 instances in an Auto Scaling group as targets.
D. Host the application on Amazon Elastic Container Service (Amazon ECS) Set up an Application Load Balancer with Amazon ECS as the target.

Correct Answer: D
Section:
Explanation:
https://aws.amazon.com/blogs/compute/microservice-delivery-with-amazon-ecs-and-applicationload- balancers/

QUESTION 89
A company recently started using Amazon Aurora as the data store for its global ecommerce application When large reports are run developers report that the ecommerce application is performing poorly After reviewing metrics in Amazon CloudWatch, a solutions architect finds that the ReadlOPS and CPUUtilization metrics are spiking when monthly reports run. What is the MOST cost-effective solution?

A. Migrate the monthly reporting to Amazon Redshift.
B. Migrate the monthly reporting to an Aurora Replica
C. Migrate the Aurora database to a larger instance class
D. Increase the Provisioned IOPS on the Aurora instance

Correct Answer: B
Section:
Explanation:
https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Aurora.Replication.html #Aurora.Replication.Replicas Aurora Replicas have two main purposes. You can issue queries to them to scale the read operations for your application. You typically do so by connecting to the reader endpoint of the cluster. That way, Aurora can spread the load for read-only connections across as many Aurora Replicas as you have in the cluster. Aurora Replicas also help to increase availability. If the writer instance in a cluster becomes unavailable, Aurora automatically promotes one of the reader instances to take its place as the new writer. https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Aurora.Overview.html

QUESTION 90

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

A. Migrate the database to an Amazon RDS for MySQL DB instance. Create an AMI of the web application. Use the AMI to launch a second EC2 On-Demand Instance. Use an Application Load Balancer to distribute the load to each EC2 instance.
B. Migrate the database to an Amazon RDS for MySQL DB instance. Create an AMI of the web application. Use the AMI to launch a second EC2 On-Demand Instance. Use Amazon Route 53 weighted routing to distribute the load across the two EC2 instances.
C. Migrate the database to an Amazon Aurora MySQL DB instance. Create an AWS Lambda function to stop the EC2 instance and change the instance type. Create an Amazon CloudWatch alarm to invoke the Lambda function when CPU utilization surpasses 75%.
D. Migrate the database to an Amazon Aurora MySQL DB instance. Create an AMI of the web application. Apply the AMI to a launch template. Create an Auto Scaling group with the launch template Configure the launch template to use a Spot Fleet. Attach an Application Load Balancer to the Auto Scaling group.

Correct Answer: D
Section:

QUESTION 91

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 usage. Use Spot Instances for any additional capacity that the application needs
D. Use Dedicated Instances for the baseline level of usage. Use On-Demand Instances for any additional capacity that the application needs

Correct Answer: B
Section:

QUESTION 92
A company hosts a two-tier application on Amazon EC2 instances and Amazon RDS. The application's demand varies based on the time of day. The load is minimal after work hours and on weekends. The EC2 instances run in an EC2 Auto Scaling group that is configured with a minimum of two instances and a maximum of five instances. The application must be available at all times, but the company is concerned about overall cost. Which solution meets the availability requirement MOST cost-effectively?

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

Correct Answer: C
Section:
Explanation:
This solution meets the requirements of a two-tier application that has a variable demand based on the time of day and must be available at all times, while minimizing the overall cost. EC2 Reserved Instances can provide significant savings compared to On-Demand Instances for the baseline level of usage, and they can guarantee capacity reservation when needed. EC2 Spot Instances can provide up to 90% savings compared to On-Demand Instances for any additional capacity that the application needs during peak hours. Spot Instances are suitable for stateless applications that can tolerate interruptions and can be replaced by other instances. Stopping the RDS database when it is not in use can reduce the cost of running the database tier.Option A is incorrect because using all EC2 Spot Instances can affect the availability of the application if there are not enough spare capacity or if the Spot price exceeds the maximum price. Stopping the RDS database when it is not in use can reduce the cost of running the database tier, but it can also affect the availability of the application. Option B is incorrect because purchasing EC2 Instance Savings Plans to cover five EC2 instances can lock in a fixed amount of compute usage per hour, which may not match the actual usage pattern of the application. Purchasing an RDS Reserved DB Instance can provide savings for the database tier, but it does not allow stopping the database when it is not in use. Option D is incorrect because purchasing EC2 Instance Savings Plans to cover two EC2 instances can lock in a fixed amount of compute usage per hour, which may not match the actual usage pattern of the application. Using up to three additional EC2 On-Demand Instances as needed can incur higher costs than using Spot Instances.

QUESTION 93
A company has an ecommerce checkout workflow that writes an order to a database and calls a service to process the payment. Users are experiencing timeouts during the checkout process. When users resubmit the checkout form, multiple unique orders are created for the same desired transaction.
How should a solutions architect refactor this workflow to prevent the creation of multiple orders?

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

Correct Answer: D
Section:

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

QUESTION 94
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 highperforming storage.
Which solution will meet these requirements?

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

Correct Answer: A
Section:

QUESTION 95
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 connections. Configure retry logic in the Lambda functions for attempts to connect to the database
B. Set up Amazon ElastiCache tor Redls to cache commonly read items from the database. Configure the Lambda functions to connect to ElastiCache for reads.
C. Add an Aurora Replica as a reader node. Configure the Lambda functions to connect to the reader endpoint of the OB cluster rather than lo the writer endpoint.
D. Use Amazon ROS Proxy to create a proxy. Set the DB cluster as the target database Configure the Lambda functions lo connect to the proxy rather than to the DB cluster.

Correct Answer: D
Section:

QUESTION 96
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) repository. Use an Amazon Elastic Container Service (Amazon ECS) cluster with the AWS Fargate launch type to run the containers. Use target tracking to scale automatically based on demand.
B. Store container images in an Amazon Elastic Container Registry (Amazon ECR) repository. Use an Amazon Elastic Container Service (Amazon ECS) cluster with the Amazon EC2 launch type to run the containers. Use target tracking to scale automatically based on demand.
C. Store container images in a repository that runs on an Amazon EC2 instance. Run the containers on EC2 instances that are spread across multiple Availability Zones. Monitor the average CPU utilization in Amazon CloudWatch. Launch new EC2 instances as needed
D. Create an Amazon EC2 Amazon Machine Image (AMI) that contains the container image Launch EC2 Instances in an Auto Scaling group across multiple Availability Zones. Use an Amazon CloudWatch alarm to scale out EC2 instances when the average CPU utilization threshold is breached.

Correct Answer: A
Section:

QUESTION 97
A company's application Is having performance issues The application staleful and needs to complete m-memory tasks on Amazon EC2 instances. The company used AWS CloudFormation to deploy infrastructure and used the

M5 EC2 Instance family As traffic increased, the application performance degraded Users are reporting delays when the users attempt to access the application. Which solution will resolve these issues in the MOST operationally efficient way?

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

Correct Answer: C
Section:
Explanation:

QUESTION 98
An ecommerce company has an order-processing application that uses Amazon API Gateway and an AWS Lambda function. The application stores data in an Amazon Aurora PostgreSQL database. During a recent sales event, a sudden surge in customer orders occurred. Some customers experienced timeouts and the application did not process the orders of those customers A solutions architect determined that the CPU utilization and memory utilization were high on the database because of a large number of open connections The solutions architect needs to prevent the timeout errors while making the least possible changes to the application. Which solution will meet these requirements?

A. Configure provisioned concurrency for the Lambda function Modify the database to be a global database in multiple AWS Regions
B. Use Amazon RDS Proxy to create a proxy for the database Modify the Lambda function to use the RDS Proxy endpoint instead of the database endpoint
C. Create a read replica for the database in a different AWS Region Use query string parameters in API Gateway to route traffic to the read replica
D. Migrate the data from Aurora PostgreSQL to Amazon DynamoDB by using AWS Database Migration Service (AWS DMS| Modify the Lambda function to use the OynamoDB table

Correct Answer: B
Section:
Explanation:

QUESTION 99
A company runs a global web application on Amazon EC2 instances behind an Application Load Balancer The application stores data in Amazon Auror a. The company needs to create a disaster recovery solution and can tolerate up to 30 minutes of downtime and potential data loss. The solution does not need to handle the load when the primary infrastructure is healthy What should a solutions architect do to meet these requirements?

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

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

QUESTION 100
A company wants to measure the effectiveness of its recent marketing campaigns. The company performs batch processing on csv files of sales data and stores the results «i an Amazon S3 bucket once every hour. The S3 bi petabytes of objects. The company runs one-time queries in Amazon Athena to determine which products are most popular on a particular date for a particular region Queries sometimes fail or take longer than expected to finish. Which actions should a solutions architect take to improve the query performance and reliability?

(Select TWO.)

A. Reduce the S3 object sizes to less than 126 MB
B. Partition the data by date and region n Amazon S3
C. Store the files as large, single objects in Amazon S3.
D. Use Amazon Kinosis Data Analytics to run the Queries as pan of the batch processing operation
E. Use an AWS duo extract, transform, and load (ETL) process to convert the csv files into Apache Parquet format.

Correct Answer: B, E
Section:
Explanation:

QUESTION 101
A company needs a solution to prevent AWS CloudFormation stacks from deploying AWS Identity and Access Management (1AM) resources that include an inline policy or '*' in the statement The solution must also prohibit deployment ot Amazon EC2 instances with public IP addresses The company has AWS Control Tower enabled in its organization in AWS Organizations.
Which solution will meet these requirements?

A. Use AWS Control Tower proactive controls to block deployment of EC2 instances with public IP addresses and inline policies with elevated access or '*'
B. Use AWS Control Tower detective controls to block deployment of EC2 instances with public IP addresses and inline policies with elevated access or ''
C. Use AWS Config to create rules for EC2 and 1AM compliance Configure the rules to run an AWS Systems Manager Session Manager automation to delete a resource when it is not compliant
D. Use a service control policy (SCP) to block actions for the EC2 instances and 1AM resources if the actions lead to noncompliance

Correct Answer: D
Section:
Explanation:
A service control policy (SCP) is a type of policy that you can use to manage permissions in your organization. SCPs offer central control over the maximum available permissions for all accounts in your organization, allowing you to ensure your accounts stay within your organization's access control guidelines. SCPs are available only in an organization that has all features enabled. SCPs do not grant permissions; instead, they specify the maximum permissions for an organization or organizational unit (OU). SCPs limit permissions that identity-based policies or resource-based policies grant to entities (users or roles) within the account, but do not grant permissions to entities. You can use SCPs to restrict the actions that the root user in an account can perform. You can also use SCPs to prevent users or roles in any account from creating or modifying certain AWS resources, such as EC2 instances with public IP addresses or IAM resources with inline policies or '''. For example, you can create an SCP that denies the ec2:RunInstances action if the request includes the AssociatePublicIpAddress parameter set to true. You can also create an SCP that denies the iam:PutUserPolicy and iam:PutRolePolicy actions if the request includes a policy document that contains '''. By attaching these SCPs to your organization or OUs, you can prevent the deployment of AWS CloudFormation stacks that violate these rules.
AWS Control Tower proactive controls are guardrails that enforce preventive policies on your accounts and resources. Proactive guardrails are implemented as AWS Organizations service control policies (SCPs) and AWS Config rules. However, AWS Control Tower does not provide a built-in proactive guardrail to block EC2 instances with public IP addresses or IAM resources with inline policies or ''*''. You would have to create your own custom guardrails using AWS CloudFormation templates and SCPs, which is essentially the same as option D. Therefore, option A is not correct.
AWS Control Tower detective controls are guardrails that detect and alert on policy violations in your accounts and resources. Detective guardrails are implemented as AWS Config rules and Amazon CloudWatch alarms. Detective guardrails do not block or remediate noncompliant resources; they only notify you of the issues. Therefore, option B is not correct.
AWS Config is a service that enables you to assess, audit, and evaluate the configurations of your AWS resources. AWS Config continuously monitors and records your AWS resource configurations and allows you to automate the evaluation of recorded configurations against desired configurations. AWS Config rules are customizable, AWS Lambda functions that AWS Config invokes to evaluate your AWS resource configurations. You can use AWS Config rules to check for compliance with your policies, such as ensuring that EC2 instances have public IP addresses disabled or IAM resources do not have inline policies or ''*''. However, AWS Config rules alone cannot prevent the deployment of AWS CloudFormation stacks that violate these policies; they can only report the compliance status. You would need to use another service, such as AWS Systems Manager Session Manager, to run automation scripts to delete or modify the noncompliant resources. This would require additional configuration and permissions, and may not be the most efficient or secure way to enforce your policies. Therefore, option C is not correct.
Service Control Policies
AWS Control Tower Guardrails
AWS Config

QUESTION 102

A company has a mobile app for customers The app's data is sensitive and must be encrypted at rest The company uses AWS Key Management Service (AWS KMS)
The company needs a solution that prevents the accidental deletion of KMS keys The solution must use Amazon Simple Notification Service (Amazon SNS) to send an email notification to administrators when a user attempts to delete a KMS key
Which solution will meet these requirements with the LEAST operational overhead''

A. Create an Amazon EventBndge rule that reacts when a user tries to delete a KMS key Configure an AWS Config rule that cancels any deletion of a KMS key Add the AWS Config rule as a target of the EventBridge rule Create an SNS topic that notifies the administrators
B. Create an AWS Lambda function that has custom logic to prevent KMS key deletion Create an Amazon CloudWatch alarm that is activated when a user tries to delete a KMS key Create an Amazon EventBridge rule that invokes the Lambda function when the DeleteKey operation is performed Create an SNS topic Configure the EventBndge rule to publish an SNS message that notifies the administrators
C. Create an Amazon EventBndge rule that reacts when the KMS DeleteKey operation is performed Configure the rule to initiate an AWS Systems Manager Automation runbook Configure the runbook to cancel the deletion of the KMS key Create an SNS topic Configure the EventBndge rule to publish an SNS message that notifies the administrators.
D. Create an AWS CloudTrail trail Configure the trail to delrver logs to a new Amazon CloudWatch log group Create a CloudWatch alarm based on the metric filter for the CloudWatch log group Configure the alarm to use Amazon SNS to notify the administrators when the KMS DeleteKey operation is performed

Correct Answer: C
Section:
Explanation:
This solution meets the requirements with the least operational overhead because it uses AWS services that are fully managed and scalable. The EventBridge rule can detect the DeleteKey operation from the AWS KMS API and trigger the Systems Manager Automation runbook, which can execute a predefined workflow to cancel the key deletion. The EventBridge rule can also publish an SNS message to the topic that sends an email notification to the administrators. This way, the company can prevent the accidental deletion of KMS keys and notify the administrators of any attempts to delete them.
Option A is not a valid solution because AWS Config rules are used to evaluate the configuration of AWS resources, not to cancel the deletion of KMS keys. Option B is not a valid solution because it requires creating and maintaining a custom Lambda function that has logic to prevent KMS key deletion, which adds operational overhead. Option D is not a valid solution because it only notifies the administrators of the DeleteKey operation, but does not cancel it.
Using Amazon EventBridge rules to trigger Systems Manager Automation workflows - AWS Systems Manager
Using Amazon SNS for system-to-administrator communications - Amazon Simple Notification Service
Deleting AWS KMS keys - AWS Key Management Service

QUESTION 103
A solutions architect is designing a user authentication solution for a company The solution must invoke two-factor authentication for users that log in from inconsistent geographical locations. IP addresses, or devices. The solution must also be able to scale up to accommodate millions of users.
Which solution will meet these requirements'?

A. Configure Amazon Cognito user pools for user authentication Enable the nsk-based adaptive authentication feature with multi-factor authentication (MFA)
B. Configure Amazon Cognito identity pools for user authentication Enable multi-factor authentication (MFA).
C. Configure AWS Identity and Access Management (1AM) users for user authentication Attach an 1AM policy that allows the AllowManageOwnUserMFA action
D. Configure AWS 1AM Identity Center (AWS Single Sign-On) authentication for user authentication Configure the permission sets to require multi-factor authentication (MFA)

Correct Answer: A
Section:
Explanation:
Amazon Cognito user pools provide a secure and scalable user directory for user authentication and management. User pools support various authentication methods, such as username and password, email and password, phone number and password, and social identity providers. User pools also support multi-factor authentication (MFA), which adds an extra layer of security by requiring users to provide a verification code or a biometric factor in addition to their credentials. User pools can also enable risk-based adaptive authentication, which dynamically adjusts the authentication challenge based on the risk level of the sign-in attempt. For example, if a user tries to sign in from an unfamiliar device or location, the user pool can require a stronger authentication factor, such as SMS or email verification code. This feature helps to protect user accounts from unauthorized access and reduce the friction for legitimate users. User pools can scale up to millions of users and integrate with other AWS services, such as Amazon SNS, Amazon SES, AWS Lambda, and AWS KMS.
Amazon Cognito identity pools provide a way to federate identities from multiple identity providers, such as user pools, social identity providers, and corporate identity providers. Identity pools allow users to access AWS resources with temporary, limited-privilege credentials. Identity pools do not provide user authentication or management features, such as MFA or adaptive authentication. Therefore, option B is not correct.
AWS Identity and Access Management (IAM) is a service that helps to manage access to AWS resources. IAM users are entities that represent people or applications that need to interact with AWS. IAM users can be authenticated with a password or an access key. IAM users can also enable MFA for their own accounts, by using the AllowManageOwnUserMFA action in an IAM policy. However, IAM users are not suitable for user

authentication for web or mobile applications, as they are intended for administrative purposes. IAM users also do not support adaptive authentication based on risk factors. Therefore, option C is not correct.
AWS IAM Identity Center (AWS Single Sign-On) is a service that enables users to sign in to multiple AWS accounts and applications with a single set of credentials. AWS SSO supports various identity sources, such as AWS SSO directory, AWS Managed Microsoft AD, and external identity providers. AWS SSO also supports MFA for user authentication, which can be configured in the permission sets that define the level of access for each user. However, AWS SSO does not support adaptive authentication based on risk factors. Therefore, option D is not correct.
Amazon Cognito User Pools
Adding Multi-Factor Authentication (MFA) to a User Pool
Risk-Based Adaptive Authentication
Amazon Cognito Identity Pools
IAM Users
Enabling MFA Devices
AWS Single Sign-On
How AWS SSO Works

QUESTION 104
A company has an Amazon S3 data lake The company needs a solution that transforms the data from the data lake and loads the data into a data warehouse every day The data warehouse must have massively parallel processing (MPP) capabilities.
Data analysts then need to create and train machine learning (ML) models by using SQL commands on the data The solution must use serverless AWS services wherever possible
Which solution will meet these requirements?

A. Run a daily Amazon EMR job to transform the data and load the data into Amazon Redshift Use Amazon Redshift ML to create and train the ML models
B. Run a daily Amazon EMR job to transform the data and load the data into Amazon Aurora Serverless Use Amazon Aurora ML to create and train the ML models
C. Run a daily AWS Glue job to transform the data and load the data into Amazon Redshift Serverless Use Amazon Redshift ML to create and tram the ML models
D. Run a daily AWS Glue job to transform the data and load the data into Amazon Athena tables Use Amazon Athena ML to create and train the ML models

Correct Answer: C
Section:
Explanation:
AWS Glue is a fully managed extract, transform, and load (ETL) service that makes it easy to prepare and load your data for analytics. AWS Glue can automatically discover your data in Amazon S3 and catalog it, so you can query and search the data using SQL. AWS Glue can also run serverless ETL jobs using Apache Spark and Python to transform and load your data into various destinations, such as Amazon Redshift, Amazon Athena, or Amazon Aurora. AWS Glue is a serverless service, so you only pay for the resources consumed by the jobs, and you don't need to provision or manage any infrastructure.
Amazon Redshift is a fully managed, petabyte-scale data warehouse service that enables you to use standard SQL and your existing business intelligence (BI) tools to analyze your data. Amazon Redshift also supports massively parallel processing (MPP), which means it can distribute and execute queries across multiple nodes in parallel, delivering fast performance and scalability. Amazon Redshift Serverless is a new option that automatically scales query compute capacity based on the queries being run, so you don't need to manage clusters or capacity. You only pay for the query processing time and the storage consumed by your data.
Amazon Redshift ML is a feature that enables you to create, train, and deploy machine learning (ML) models using familiar SQL commands. Amazon Redshift ML can automatically discover the best model and hyperparameters for your data, and store the model in Amazon SageMaker, a fully managed service that provides a comprehensive set of tools for building, training, and deploying ML models. You can then use SQL functions to apply the model to your data in Amazon Redshift and generate predictions.
The combination of AWS Glue, Amazon Redshift Serverless, and Amazon Redshift ML meets the requirements of the question, as it provides a serverless, scalable, and SQL-based solution to transform, load, and analyze the data from the Amazon S3 data lake, and to create and train ML models on the data.
Option A is not correct, because Amazon EMR is not a serverless service. Amazon EMR is a managed service that simplifies running Apache Spark, Apache Hadoop, and other big data frameworks on AWS. Amazon EMR requires you to launch and configure clusters of EC2 instances to run your ETL jobs, which adds complexity and cost compared to AWS Glue.
Option B is not correct, because Amazon Aurora Serverless is not a data warehouse service, and it does not support MPP. Amazon Aurora Serverless is an on-demand, auto-scaling configuration for Amazon Aurora, a relational database service that is compatible with MySQL and PostgreSQL. Amazon Aurora Serverless can automatically adjust the database capacity based on the traffic, but it does not distribute the data and queries across multiple nodes like Amazon Redshift does. Amazon Aurora Serverless is more suitable for transactional workloads than analytical workloads.
Option D is not correct, because Amazon Athena is not a data warehouse service, and it does not support MPP. Amazon Athena is an interactive query service that enables you to analyze data in Amazon S3 using standard SQL. Amazon Athena is serverless, so you only pay for the queries you run, and you don't need to load the data into a database. However, Amazon Athena does not store the data in a columnar format, compress the data, or optimize the query execution plan like Amazon Redshift does. Amazon Athena is more suitable for ad-hoc queries than complex analytics and ML.
AWS Glue
Amazon Redshift
Amazon Redshift Serverless

Amazon Redshift ML
Amazon EMR
Amazon Aurora Serverless
Amazon Athena

QUESTION 105
A company runs containers in a Kubernetes environment in the company's local data center. The company wants to use Amazon Elastic Kubernetes Service (Amazon EKS) and other AWS managed services Data must remain locally in the company's data center and cannot be stored in any remote site or cloud to maintain compliance
Which solution will meet these requirements?

A. Deploy AWS Local Zones in the company's data center
B. Use an AWS Snowmobile in the company's data center
C. Install an AWS Outposts rack in the company's data center
D. Install an AWS Snowball Edge Storage Optimized node in the data center

Correct Answer: C
Section:
Explanation:
AWS Outposts is a fully managed service that delivers AWS infrastructure and services to virtually any on-premises or edge location for a consistent hybrid experience. AWS Outposts supports Amazon EKS, which is a managed service that makes it easy to run Kubernetes on AWS and on-premises. By installing an AWS Outposts rack in the company's data center, the company can run containers in a Kubernetes environment using Amazon EKS and other AWS managed services, while keeping the data locally in the company's data center and meeting the compliance requirements. AWS Outposts also provides a seamless connection to the local AWS Region for access to a broad range of AWS services.
Option A is not a valid solution because AWS Local Zones are not deployed in the company's data center, but in large metropolitan areas closer to end users. AWS Local Zones are owned, managed, and operated by AWS, and they provide low-latency access to the public internet and the local AWS Region. Option B is not a valid solution because AWS Snowmobile is a service that transports exabytes of data to AWS using a 45-foot long ruggedized shipping container pulled by a semi-trailer truck. AWS Snowmobile is not designed for running containers or AWS managed services on-premises, but for large-scale data migration. Option D is not a valid solution because AWS Snowball Edge Storage Optimized is a device that provides 80 TB of HDD or 210 TB of NVMe storage capacity for data transfer and edge computing. AWS Snowball Edge Storage Optimized does not support Amazon EKS or other AWS managed services, and it is not suitable for running containers in a Kubernetes environment.
AWS Outposts - Amazon Web Services
Amazon EKS on AWS Outposts - Amazon EKS
AWS Local Zones - Amazon Web Services
AWS Snowmobile - Amazon Web Services
[AWS Snowball Edge Storage Optimized - Amazon Web Services]

QUESTION 106
A company is running several business applications in three separate VPCs within me us-east-1 Region. The applications must be able to communicate between VPCs. The applications also must be able to consistently send hundreds to gigabytes of data each day to a latency-sensitive application that runs in a single on-premises data center. A solutions architect needs to design a network connectivity solution that maximizes costeffectiveness
Which solution moots those requirements?

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

Correct Answer: D
Section:
Explanation:
https://docs.aws.amazon.com/whitepapers/latest/aws-vpc-connectivity-options/aws-directconnect-aws-transit-gateway.html

QUESTION 107
A company's web application is running on Amazon EC2 instances behind an Application Load Balancer. The company recently changed its policy, which now requires the application to be accessed from one specific country only. Which configuration will meet this requirement?

A. Configure the security group for the EC2 instances.
B. Configure the security group on the Application Load Balancer.
C. Configure AWS WAF on the Application Load Balancer in a VPC.
D. Configure the network ACL for the subnet that contains the EC2 instances.

Correct Answer: C
Section:
Explanation:
https://aws.amazon.com/about-aws/whats-new/2017/10/aws-waf-now-supports-geographicmatch/

QUESTION 108
A company has a mulli-tier application that runs six front-end web servers in an Amazon EC2 Auto Scaling group in a single Availability Zone behind an Application Load Balancer (ALB). A solutions architect needs lo modify the infrastructure to be highly available without modifying the application.
Which architecture should the solutions architect choose that provides high availability?

A. Create an Auto Scaling group that uses three Instances across each of tv/o Regions.
B. Modify the Auto Scaling group to use three instances across each of two Availability Zones.
C. Create an Auto Scaling template that can be used to quickly create more instances in another Region.
D. Change the ALB in front of the Amazon EC2 instances in a round-robin configuration to balance traffic to the web tier.

Correct Answer: B
Section:
Explanation:
High availability can be enabled for this architecture quite simply by modifying the existing Auto Scaling group to use multiple availability zones. The ASG will automatically balance the load so you don't actually need to specify the instances per AZ.

QUESTION 109
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.

Correct Answer: D
Section:

QUESTION 110
A company runs an application using Amazon ECS. The application creates esi/ed versions of an original image and then makes Amazon S3 API calls to store the resized images in Amazon S3. How can a solutions architect ensure that the application has permission to access Amazon S3?

A. Update the S3 role in AWS IAM to allow read/write access from Amazon ECS, and then relaunch the container.

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

Correct Answer: B
Section:

QUESTION 111
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 parameter. Allow Decrypt access to an AWS Key Management Service (AWS KMS) key that is used to encrypt the parameter. Assign this IAM role to the EC2 instance.
B. Create an IAM policy that allows read access to the Parameter Store parameter. Allow Decrypt access to an AWS Key Management Service (AWS KMS) key that is used to encrypt the parameter. Assign this IAM policy to the EC2 instance.
C. Create an IAM trust relationship between the Parameter Store parameter and the EC2 instance.
Specify Amazon RDS as a principal in the trust policy.
D. Create an IAM trust relationship between the DB instance and the EC2 instance. Specify Systems Manager as a principal in the trust policy.

Correct Answer: A
Section:
Explanation:

QUESTION 112
An entertainment company is using Amazon DynamoDB to store media metadat a. 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.

Correct Answer: B
Section:
Explanation:
https://aws.amazon.com/dynamodb/dax/

QUESTION 113
A security team wants to limit access to specific services or actions in all of the team's AWS accounts.
All accounts belong to a large organization in AWS Organizations. The solution must be scalable and there must be a single point where permissions can be maintained.
What should a solutions architect do to accomplish this?

A. Create an ACL to provide access to the services or actions.
B. Create a security group to allow accounts and attach it to user groups.
C. Create cross-account roles in each account to deny access to the services or actions.

D. Create a service control policy in the root organizational unit to deny access to the services or actions.

Correct Answer: D
Section:
Explanation:
Service control policies (SCPs) are one type of policy that you can use to manage your organization. SCPs offer central control over the maximum available permissions for all accounts in your organization, allowing you to ensure your accounts stay within your organization's access control guidelines. See https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scp.html.

QUESTION 114
A company is concerned about the security of its public web application due to recent web attacks.
The application uses an Application Load Balancer (ALB). A solutions architect must reduce the risk of DDoS attacks against the application. What should the solutions architect do to meet this requirement?

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

Correct Answer: C
Section:

QUESTION 115
A company runs a production application on a fleet of Amazon EC2 instances. The application reads the data from an Amazon SQS queue and processes the messages in parallel. The message volume is unpredictable and often has intermittent traffic. This application should continually process messages without any downtime.
Which solution meets these requirements MOST cost-effectively?

A. Use Spot Instances exclusively to handle the maximum capacity required.
B. Use Reserved Instances exclusively to handle the maximum capacity required.
C. Use Reserved Instances for the baseline capacity and use Spot Instances to handle additional capacity.
D. Use Reserved Instances for the baseline capacity and use On-Demand Instances to handle additional capacity.

Correct Answer: D
Section:
Explanation:
We recommend that you use On-Demand Instances for applications with short-term, irregular workloads that cannot be interrupted. https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-on-demand-instances.html

QUESTION 116
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 bucket. Enable AWS WAF on the distribution.

Correct Answer: D
Section:
Explanation:

https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/private-contentrestricting-access-to-s3.html
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/distribution-webawswaf.html

QUESTION 117
A company sends AWS CloudTrail logs from multiple AWS accounts to an Amazon S3 bucket in a centralized account. The company must keep the CloudTrail logs. The company must also be able to query the CloudTrail logs at any time Which solution will meet these requirements?

A. Use the CloudTrail event history in the centralized account to create an Amazon Athena table.
Query the CloudTrail logs from Athena.
B. Configure an Amazon Neptune instance to manage the CloudTrail logs. Query the CloudTrail logs from Neptune.
C. Configure CloudTrail to send the logs to an Amazon DynamoDB table. Create a dashboard in Amazon QulCkSight to query the logs in the table.
D. use Amazon Athena to create an Athena notebook. Configure CloudTrail to send the logs to the notebook. Run queries from Athena.

Correct Answer: A
Section:
Explanation:
it allows the company to keep the CloudTrail logs and query them at any time. By using the CloudTrail event history in the centralized account, the company can view, filter, and download recent API activity across multiple AWS accounts. By creating an Amazon Athena table from the CloudTrail event history, the company can use a serverless interactive query service that makes it easy to analyze data in S3 using standard SQL. By querying the CloudTrail logs from Athena, the company can gain insights into user activity and resource changes. Reference:
Viewing Events with CloudTrail Event History
Querying AWS CloudTrail Logs
Amazon Athena

QUESTION 118
A company has five organizational units (OUs) as part of its organization in AWS Organizations. Each OU correlates to the five businesses that the company owns. The company's research and development (R&D) business is separating from the company and will need its own organization. A solutions architect creates a separate new management account for this purpose.
What should the solutions architect do next in the new management account?

A. Have the R&D AWS account be part of both organizations during the transition.
B. Invite the R&D AWS account to be part of the new organization after the R&D AWS account has left the prior organization.
C. Create a new R&D AWS account in the new organization. Migrate resources from the prior R&D AWS account to the new R&D AWS account.
D. Have the R&D AWS account join the new organization. Make the new management account a member of the prior organization.

Correct Answer: B
Section:
Explanation:
it allows the solutions architect to create a separate organization for the research and development (R&D) business and move its AWS account to the new organization. By inviting the R&D AWS account to be part of the new organization after it has left the prior organization, the solutions architect can ensure that there is no overlap or conflict between the two organizations. The R&D AWS account can accept or decline the invitation to join the new organization. Once accepted, it will be subject to any policies and controls applied by the new organization. Reference:
Inviting an AWS Account to Join Your Organization
Leaving an Organization as a Member Account

QUESTION 119
A company is deploying an application that processes large quantities of data in parallel. The company plans to use Amazon EC2 instances for the workload. The network architecture must be configurable to prevent groups of nodes from sharing the same underlying hardware.
Which networking solution meets these requirements?

A. Run the EC2 instances in a spread placement group.

B. Group the EC2 instances in separate accounts.
C. Configure the EC2 instances with dedicated tenancy.
D. Configure the EC2 instances with shared tenancy.

Correct Answer: A
Section:
Explanation:
it allows the company to deploy an application that processes large quantities of data in parallel and prevent groups of nodes from sharing the same underlying hardware. By running the EC2 instances in a spread placement group, the company can launch a small number of instances across distinct underlying hardware to reduce correlated failures. A spread placement group ensures that each instance is isolated from each other at the rack level.
Reference:
Placement Groups
Spread Placement Groups

QUESTION 120
A solutions architect is creating a new Amazon CloudFront distribution for an application. Some of the information submitted by users is sensitive. The application uses HTTPS but needs another layer of security. The sensitive information should.be protected throughout the entire application stack, and access to the information should be restricted to certain applications.
Which action should the solutions architect take?

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

Correct Answer: C
Section:
Explanation:
it allows the company to protect sensitive information submitted by users throughout the entire application stack and restrict access to certain applications. By configuring a CloudFront field-level encryption profile, the company can encrypt specific fields of user data at the edge locations before sending it to the origin servers. By using public-private key pairs, the company can ensure that only authorized applications can decrypt and access the sensitive information. Reference:
Field-Level Encryption
Encrypting and Decrypting Data

QUESTION 121
A company runs its applications on Amazon EC2 instances. The company performs periodic financial assessments of itsAWS costs. The company recently identified unusual spending.
The company needs a solution to prevent unusual spending. The solution must monitor costs and notify responsible stakeholders in the event of unusual spending.
Which solution will meet these requirements?

A. Use an AWS Budgets template to create a zero spend budget
B. Create an AWS Cost Anomaly Detection monitor in the AWS Billing and Cost Management console.
C. CreateAWS Pricing Calculator estimates for the current running workload pricing details_
D. Use Amazon CloudWatch to monitor costs and to identify unusual spending

Correct Answer: B
Section:
Explanation:
it allows the company to monitor costs and notify responsible stakeholders in the event of unusual spending. By creating an AWS Cost Anomaly Detection monitor in the AWS Billing and Cost Management console, the company can use a machine learning service that automatically detects and alerts on anomalous spend. By configuring alert thresholds, notification preferences, and root cause analysis, the company can prevent unusual spending and identify its source. Reference:

AWS Cost Anomaly Detection
Creating a Cost Anomaly Monitor

QUESTION 122
A recent analysis of a company's IT expenses highlights the need to reduce backup costs. The company's chief information officer wants to simplify the on- premises backup infrastructure and reduce costs by eliminating the use of physical backup tapes. The company must preserve the existing investment in the on- premises backup applications and workflows.
What should a solutions architect recommend?

A. Set up AWS Storage Gateway to connect with the backup applications using the NFS interface.
B. Set up an Amazon EFS file system that connects with the backup applications using the NFS interface.
C. Set up an Amazon EFS file system that connects with the backup applications using the iSCSI interface.
D. Set up AWS Storage Gateway to connect with the backup applications using the iSCSI-virtual tape library (VTL) interface.

Correct Answer: D
Section:
Explanation:
it allows the company to simplify the on-premises backup infrastructure and reduce costs by eliminating the use of physical backup tapes. By setting up AWS Storage Gateway to connect with the backup applications using the iSCSI-virtual tape library (VTL) interface, the company can store backup data on virtual tapes in S3 or Glacier. This preserves the existing investment in the on-premises backup applications and workflows while leveraging AWS storage services. Reference:
AWS Storage Gateway
Tape Gateway

QUESTION 123
A company is moving its data and applications to AWS during a multiyear migration project. The company wants to securely access data on Amazon S3 from the company's AWS Region and from the company's on-premises location. The data must not traverse the internet. The company has established an AWS Direct Connect connection between its Region and its on-premises location Which solution will meet these requirements?

A. Create gateway endpoints for Amazon S3. Use the gateway endpoints to securely access the data from the Region and the on-premises location.
B. Create a gateway in AWS Transit Gateway to access Amazon S3 securely from the Region and the on-premises location.
C. Create interface endpoints for Amazon S3_ Use the interface endpoints to securely access the data from the Region and the on-premises location.
D. Use an AWS Key Management Service (AWS KMS) key to access the data securely from the Region and the on-premises location.

Correct Answer: B
Section:
Explanation:
A gateway endpoint is a gateway that is a target for a specified route in your route table, used for traffic destined to a supported AWS service1. Amazon S3 does not support gateway endpoints, only interface endpoints2. Therefore, option A is incorrect.
An interface endpoint is an elastic network interface with a private IP address that serves as an entry point for traffic destined to a supported service1. An interface endpoint can provide secure access to Amazon S3 from within the Region, but not from the on-premises location. Therefore, option C is incorrect.
AWS Key Management Service (AWS KMS) is a service that allows you to create and manage encryption keys to protect your data3. AWS KMS does not provide a way to access data on Amazon S3 without traversing the internet. Therefore, option D is incorrect.
AWS Transit Gateway is a service that enables you to connect your Amazon Virtual Private Clouds (VPCs) and your on-premises networks to a single gateway. You can create a gateway in AWS Transit Gateway to access Amazon S3 securely from both the Region and the on-premises location using AWS Direct Connect. Therefore, option B is correct.

QUESTION 124
A company's website handles millions of requests each day, and the number of requests continues to increase. A solutions architect needs to improve the response time of the web application. The solutions architect determines that the application needs to decrease latency when retrieving product details from the Amazon DynamoDB table.
Which solution will meet these requirements with the LEAST amount of operational overhead?

A. Set up a DynamoDB Accelerator (DAX) cluster. Route all read requests through DAX.
B. Set up Amazon ElastiCache for Redis between the DynamoDB table and the web application. Route all read requests through Redis.
C. Set up Amazon ElastiCache for Memcached between the DynamoDB table and the web application. Route all read requests through Memcached.
D. Set up Amazon DynamoDB Streams on the table, and have AWS Lambda read from the table and populate Amazon ElastiCache. Route all read requests through ElastiCache.

Correct Answer: A
Section:
Explanation:
it allows the company to improve the response time of the web application and decrease latency when retrieving product details from the Amazon DynamoDB table. By setting up a DynamoDB Accelerator (DAX) cluster, the company can use a fully managed, highly available, in-memory cache for DynamoDB that delivers up to a 10x performance improvement. By routing all read requests through DAX, the company can reduce the number of read operations on the DynamoDB table and improve the user experience. Reference:
Amazon DynamoDB Accelerator (DAX)
Using DAX with DynamoDB

QUESTION 125
A company runs a three-tier application in two AWS Regions. The web tier, the application tier, and the database tier run on Amazon EC2 instances. The company uses Amazon RDS for Microsoft SQL Server Enterprise for the database tier The database tier is experiencing high load when weekly and monthly reports are run. The company wants to reduce the load on the database tier.
Which solution will meet these requirements with the LEAST administrative effort?

A. Create read replicas. Configure the reports to use the new read replicas.
B. Convert the RDS database to Amazon DynamoDB_ Configure the reports to use DynamoDB
C. Modify the existing RDS DB instances by selecting a larger instance size.
D. Modify the existing ROS DB instances and put the instances into an Auto Scaling group.

Correct Answer: A
Section:
Explanation:
it allows the company to create read replicas of its RDS database and reduce the load on the database tier. By creating read replicas, the company can offload read traffic from the primary database instance to one or more replicas. By configuring the reports to use the new read replicas, the company can improve performance and availability of its database tier. Reference:
Working with Read Replicas
Read Replicas for Amazon RDS for SQL Server

QUESTION 126
A company runs a website that stores images of historical events. Website users need the ability to search and view images based on the year that the event in the image occurred. On average, users request each image only once or twice a year The company wants a highly available solution to store and deliver the images to users.
Which solution will meet these requirements MOST cost-effectively?

A. Store images in Amazon Elastic Block Store (Amazon EBS). Use a web server that runs on Amazon EC2.
B. Store images in Amazon Elastic File System (Amazon EFS). Use a web server that runs on Amazon EC2.
C. Store images in Amazon S3 Standard. use S3 Standard to directly deliver images by using a static website.
D. Store images in Amazon S3 Standard-InfrequentAccess (S3 Standard-IA). use S3 Standard-IA to directly deliver images by using a static website.

Correct Answer: C
Section:
Explanation:
it allows the company to store and deliver images to users in a highly available and cost-effective way. By storing images in Amazon S3 Standard, the company can use a durable, scalable, and secure object storage service that offers high availability and performance. By using S3 Standard to directly deliver images by using a static website, the company can avoid running web servers and reduce operational overhead. S3 Standard also offers low

storage pricing and free data transfer within AWS Regions. Reference:
Amazon S3 Storage Classes
Hosting a Static Website on Amazon S3

QUESTION 127
A solutions architect needs to review a company's Amazon S3 buckets to discover personally identifiable information (Pll). The company stores the Pll data in the us-east-I Region and us-west-2 Region.
Which solution will meet these requirements with the LEAST operational overhead?

A. Configure Amazon Macie in each Region. Create a job to analyze the data that is in Amazon S3_
B. Configure AWS Security Hub for all Regions. Create an AWS Config rule to analyze the data that is in Amazon S3_
C. Configure Amazon Inspector to analyze the data that IS in Amazon S3.
D. Configure Amazon GuardDuty to analyze the data that is in Amazon S3.

Correct Answer: A
Section:
Explanation:
it allows the solutions architect to review the S3 buckets to discover personally identifiable information (Pll) with the least operational overhead. Amazon Macie is a fully managed data security and data privacy service that uses machine learning and pattern matching to discover and protect sensitive data in AWS. Amazon Macie can analyze data in S3 buckets across multiple regions and provide insights into the type, location, and level of sensitivity of the data. Reference:
Amazon Macie
Analyzing data with Amazon Macie

QUESTION 128
A company is building an ecommerce application and needs to store sensitive customer information.
The company needs to give customers the ability to complete purchase transactions on the website.
The company also needs to ensure that sensitive customer data is protected, even from database administrators.
Which solution meets these requirements?

A. Store sensitive data in an Amazon Elastic Block Store (Amazon EBS) volume. Use EBS encryption to encrypt the data. Use an IAM instance role to restrict access.
B. Store sensitive data in Amazon RDS for MySQL. Use AWS Key Management Service (AWS KMS) client-side encryption to encrypt the data.
C. Store sensitive data in Amazon S3. Use AWS Key Management Service (AWS KMS) server-side encryption to encrypt the data. Use S3 bucket policies to restrict access.
D. Store sensitive data in Amazon FSx for Windows Server. Mount the file share on application servers. Use Windows file permissions to restrict access.

Correct Answer: B
Section:
Explanation:
it allows the company to store sensitive customer information in a managed AWS service and give customers the ability to complete purchase transactions on the website. By using AWS Key Management Service (AWS KMS) client-side encryption, the company can encrypt the data before sending it to Amazon RDS for MySQL. This ensures that sensitive customer data is protected, even from database administrators, as only the application has access to the encryption keys. Reference:
Using Encryption with Amazon RDS for MySQL
Encrypting Amazon RDS Resources

QUESTION 129
A solutions architect needs to ensure that API calls to Amazon DynamoDB from Amazon EC2 instances in a VPC do not travel across the internet.
Which combination of steps should the solutions architect take to meet this requirement? (Choose two.)

A. Create a route table entry for the endpoint.

B. Create a gateway endpoint for DynamoDB.
C. Create an interface endpoint for Amazon EC2.
D. Create an elastic network interface for the endpoint in each of the subnets of the VPC.
E. Create a security group entry in the endpoint's security group to provide access.

Correct Answer: B, E
Section:
Explanation:
B and E are the correct answers because they allow the solutions architect to ensure that API calls to Amazon DynamoDB from Amazon EC2 instances in a VPC do not travel across the internet. By creating a gateway endpoint for DynamoDB, the solutions architect can enable private connectivity between the VPC and DynamoDB. By creating a security group entry in the endpoint's security group to provide access, the solutions architect can control which EC2 instances can communicate with DynamoDB through the endpoint. Reference:
Gateway Endpoints Controlling Access to Services with VPC Endpoints

QUESTION 130
A company has a service that reads and writes large amounts of data from an Amazon S3 bucket in the same AWS Region. The service is deployed on Amazon EC2 instances within the private subnet of a VPC. The service communicates with Amazon S3 over a NAT gateway in the public subnet.
However, the company wants a solution that will reduce the data output costs.
Which solution will meet these requirements MOST cost-effectively?

A. Provision a dedicated EC2 NAT instance in the public subnet. Configure the route table for the private subnet to use the elastic network interface of this instance as the destination for all S3 traffic.
B. Provision a dedicated EC2 NAT instance in the private subnet. Configure the route table for the public subnet to use the elastic network interface of this instance as the destination for all S3 traffic.
C. Provision a VPC gateway endpoint. Configure the route table for the private subnet to use the gateway endpoint as the route for all S3 traffic.
D. Provision a second NAT gateway. Configure the route table for the private subnet to use this NAT gateway as the destination for all S3 traffic.

Correct Answer: C
Section:
Explanation:
it allows the company to reduce the data output costs for accessing Amazon S3 from Amazon EC2 instances in a VPC. By provisioning a VPC gateway endpoint, the company can enable private connectivity between the VPC and S3. By configuring the route table for the private subnet to use the gateway endpoint as the route for all S3 traffic, the company can avoid using a NAT gateway, which charges for data processing and data transfer.
Reference:
VPC Endpoints for Amazon S3
VPC Endpoints Pricing

QUESTION 131
A company runs multiple Amazon EC2 Linux instances in a VPC across two Availability Zones. The instances host applications that use a hierarchical directory structure. The applications need to read and write rapidly and concurrently to shared storage.
What should a solutions architect do to meet these requirements?

A. Create an Amazon S3 bucket. Allow access from all the EC2 instances in the VPC.
B. Create an Amazon Elastic File System (Amazon EFS) file system. Mount the EFS file system from each EC2 instance.
C. Create a file system on a Provisioned IOPS SSD (102) Amazon Elastic Block Store (Amazon EBS) volume. Attach the EBS volume to all the EC2 instances.
D. Create file systems on Amazon Elastic Block Store (Amazon EBS) volumes that are attached to each EC2 instance. Synchromze the EBS volumes across the different EC2 instances.

Correct Answer: B
Section:
Explanation:
it allows the EC2 instances to read and write rapidly and concurrently to shared storage across two
Availability Zones. Amazon EFS provides a scalable, elastic, and highly available file system that can be mounted from multiple EC2 instances. Amazon EFS supports high levels of throughput and IOPS, and consistent low

latencies. Amazon EFS also supports NFSv4 lock upgrading and downgrading, which enables high levels of concurrency. Reference:
Amazon EFS Features
Using Amazon EFS with Amazon EC2

QUESTION 132
A company is using AWS Key Management Service (AWS KMS) keys to encrypt AWS Lambda environment variables. A solutions architect needs to ensure that the required permissions are in place to decrypt and use the environment variables.
Which steps must the solutions architect take to implement the correct permissions? (Choose two.)

A. Add AWS KMS permissions in the Lambda resource policy.
B. Add AWS KMS permissions in the Lambda execution role.
C. Add AWS KMS permissions in the Lambda function policy.
D. Allow the Lambda execution role in the AWS KMS key policy.
E. Allow the Lambda resource policy in the AWS KMS key policy.

Correct Answer: B, D
Section:
Explanation:
B and D are the correct answers because they ensure that the Lambda execution role has the permissions to decrypt and use the environment variables, and that the AWS KMS key policy allows the Lambda execution role to use the key. The Lambda execution role is an IAM role that grants the Lambda function permission to access AWS resources, such as AWS KMS. The AWS KMS key policy is a resource-based policy that controls access to the key. By adding AWS KMS permissions in the Lambda execution role and allowing the Lambda execution role in the AWS KMS key policy, the solutions architect can implement the correct permissions for encrypting and decrypting environment variables. Reference:
AWS Lambda Execution Role
Using AWS KMS keys in AWS Lambda

QUESTION 133
A company wants to use an AWS CloudFormatlon stack for its application in a test environment. The company stores the CloudFormation template in an Amazon S3 bucket that blocks public access. The company wants to grant CloudFormation access to the template in the S3 bucket based on specific user requests to create the test environment The solution must follow security best practices.
Which solution will meet these requirements?

A. Create a gateway VPC endpoint for Amazon S3. Configure the CloudFormation stack to use the S3 object URL
B. Create an Amazon API Gateway REST API that has the S3 bucket as the target. Configure the CloudFormat10n stack to use the API Gateway URL _
C. Create a presigned URL for the template object_ Configure the CloudFormation stack to use the presigned URL.
D. Allow public access to the template object in the S3 bucket. Block the public access after the test environment is created

Correct Answer: C
Section:
Explanation:
it allows CloudFormation to access the template in the S3 bucket without granting public access or creating additional resources. A presigned URL is a URL that is signed with the access key of an IAM user or role that has permission to access the object. The presigned URL can be used by anyone who receives it, but it expires after a specified time. By creating a presigned URL for the template object and configuring the CloudFormation stack to use it, the company can grant CloudFormation access to the template based on specific user requests and follow security best practices. Reference:
Using Amazon S3 Presigned URLs
Using Amazon S3 Buckets

QUESTION 134
A company runs an application on AWS. The application receives inconsistent amounts of usage. The application uses AWS Direct Connect to connect to an on-premises MySQL-compatible database. The on-premises database consistently uses a minimum of 2 GiB of memory.
The company wants to migrate the on-premises database to a managed AWS service. The company wants to use auto scaling capabilities to manage unexpected workload increases.

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

A. Provision an Amazon DynamoDB database with default read and write capacity settings.
B. Provision an Amazon Aurora database with a minimum capacity of 1 Aurora capacity unit (ACU).
C. Provision an Amazon Aurora Serverless v2 database with a minimum capacity of 1 Aurora capacity unit (ACU).
D. Provision an Amazon RDS for MySQL database with 2 GiB of memory.

Correct Answer: C
Section:
Explanation:
it allows the company to migrate the on-premises database to a managed AWS service that supports auto scaling capabilities and has the least administrative overhead. Amazon Aurora Serverless v2 is a configuration of Amazon Aurora that automatically scales compute capacity based on workload demand. It can scale from hundreds to hundreds of thousands of transactions in a fraction of a second. Amazon Aurora Serverless v2 also supports MySQL-compatible databases and AWS Direct Connect connectivity. Reference:
Amazon Aurora Serverless v2 Connecting to an Amazon Aurora DB Cluster

QUESTION 135
A company uses Amazon Elastic Kubernetes Service (Amazon EKS) to run a container application. The EKS cluster stores sensitive information in the Kubernetes secrets object. The company wants to ensure that the information is encrypted Which solution will meet these requirements with the LEAST operational overhead?

A. Use the container application to encrypt the information by using AWS Key Management Service (AWS KMS).
B. Enable secrets encryption in the EKS cluster by using AWS Key Management Service (AWS KMS)_
C. Implement an AWS Lambda tuncüon to encrypt the information by using AWS Key Management Service (AWS KMS).
D. use AWS Systems Manager Parameter Store to encrypt the information by using AWS Key Management Service (AWS KMS).

Correct Answer: B
Section:
Explanation:
it allows the company to encrypt the Kubernetes secrets object in the EKS cluster with the least operational overhead. By enabling secrets encryption in the EKS cluster, the company can use AWS Key Management Service (AWS KMS) to generate and manage encryption keys for encrypting and decrypting secrets at rest. This is a simple and secure way to protect sensitive information in EKS clusters. Reference:
Encrypting Kubernetes secrets with AWS KMS
Kubernetes Secrets

QUESTION 136
A company runs a web application on Amazon EC2 instances in an Auto Scaling group that has a target group. The company desgned the application to work with session affinity (sticky sessions) for a better user experience. The application must be available publicly over the internet as an endpoint_ A WAF must be applied to the endpoint for additional security. Session affinity (sticky sessions) must be configured on the endpoint Which combination of steps will meet these requirements? (Select TWO)

A. Create a public Network Load Balancer Specify the application target group.
B. Create a Gateway Load Balancer Specify the application target group.
C. Create a public Application Load Balancer Specify the application target group.
D. Create a second target group. Add Elastic IP addresses to the EC2 instances
E. Create a web ACL in AWS WAF Associate the web ACL with the endpoint

Correct Answer: C, E
Section:
Explanation:
C and E are the correct answers because they allow the company to create a public endpoint for its web application that supports session affinity (sticky sessions) and has a WAF applied for additional security. By creating a public Application Load Balancer, the company can distribute incoming traffic across multiple EC2 instances in an Auto Scaling group and specify the application target group. By creating a web ACL in AWS WAF and associating

it with the Application Load Balancer, the company can protect its web application from common web exploits. By enabling session stickiness on the Application Load Balancer, the company can ensure that subsequent requests from a user during a session are routed to the same target. Reference:
Application Load Balancers
AWS WAF
Target Groups for Your Application Load Balancers
How Application Load Balancer Works with Sticky Sessions

QUESTION 137
A company runs an application on Amazon EC2 instances. The company needs to implement a disaster recovery (DR) solution for the application. The DR solution needs to have a recovery time objective (RTO) of less than 4 hours. The DR solution also needs to use the fewest possible AWS resources during normal operations.
Which solution will meet these requirements in the MOST operationally efficient way?

A. Create Amazon Machine Images (AMIs) to back up the EC2 instances. Copy the AMIs to a secondary AWS Region. Automate infrastructure deployment in the secondary Region by using AWS Lambda and custom scripts.
B. Create Amazon Machine Images (AMIs) to back up the EC2 instances. Copy the AMIs to a secondary AWS Region. Automate infrastructure deployment in the secondary Region by using AWS CloudFormation.
C. Launch EC2 instances in a secondary AWS Region. Keep the EC2 instances in the secondary Region active at all times.
D. Launch EC2 instances in a secondary Availability Zone. Keep the EC2 instances in the secondary Availability Zone active at all times.

Correct Answer: B
Section:
Explanation:
it allows the company to implement a disaster recovery (DR) solution for the application that has a recovery time objective (RTO) of less than 4 hours and uses the fewest possible AWS resources during normal operations. By creating Amazon Machine Images (AMIs) to back up the EC2 instances and copying the AMIs to a secondary AWS Region, the company can create point-in-time snapshots of the application and store them in a different geographical location. By automating infrastructure deployment in the secondary Region by using AWS CloudFormation, the company can quickly launch a stack of resources from a template in case of a disaster. This is a cost-effective and operationally efficient way to implement a DR solution for EC2 instances. Reference:
Amazon Machine Images (AMI)
Copying an AMI
AWS CloudFormation
Working with Stacks

QUESTION 138
A company stores its data on premises. The amount of data is growing beyond the company's available capacity.
The company wants to migrate its data from the on-premises location to an Amazon S3 bucket The company needs a solution that will automatically validate the integrity of the data after the transfer Which solution will meet these requirements?

A. Order an AWS Snowball Edge device Configure the Snowball Edge device to perform the online data transfer to an S3 bucket.
B. Deploy an AWS DataSync agent on premises. Configure the DataSync agent to perform the online data transfer to an S3 bucket.
C. Create an Amazon S3 File Gateway on premises. Configure the S3 File Gateway to perform the online data transfer to an S3 bucket
D. Configure an accelerator in Amazon S3 Transfer Acceleration on premises. Configure the accelerator to perform the online data transfer to an S3 bucket.

Correct Answer: B
Section:
Explanation:
it allows the company to migrate its data from the on-premises location to an Amazon S3 bucket and automatically validate the integrity of the data after the transfer. By deploying an AWS DataSync agent on premises, the company can use a fully managed data transfer service that makes it easy to move large amounts of data to and from AWS. By configuring the DataSync agent to perform the online data transfer to an S3 bucket, the company can take advantage of DataSync's features, such as encryption, compression, bandwidth throttling, and data validation. DataSync automatically verifies data integrity at both source and destination after each transfer task.
Reference:
AWS DataSync
Deploying an Agent for AWS DataSync
How AWS DataSync Works

QUESTION 139
A company has a large workload that runs every Friday evening. The workload runs on Amazon EC2 instances that are in two Availability Zones in the us-east-1 Region. Normally, the company must run no more than two instances at all times. However, the company wants to scale up to six instances each Friday to handle a regularly repeating increased workload.
Which solution will meet these requirements with the LEAST operational overhead?

A. Create a reminder in Amazon EventBridge to scale the instances.
B. Create an Auto Scaling group that has a scheduled action.
C. Create an Auto Scaling group that uses manual scaling.
D. Create an Auto Scaling group that uses automatic scaling.

Correct Answer: B
Section:
Explanation:
An Auto Scaling group is a collection of EC2 instances that share similar characteristics and can be scaled in or out automatically based on demand. An Auto Scaling group can have a scheduled action, which is a configuration that tells the group to scale to a specific size at a specific time. This way, the company can scale up to six instances each Friday evening to handle the increased workload, and scale down to two instances at other times to save costs. This solution meets the requirements with the least operational overhead, as it does not require manual intervention or custom scripts.
1 explains how to create a scheduled action for an Auto Scaling group.
2 describes the concept and benefits of an Auto Scaling group.

QUESTION 140
A company uses an on-premises network-attached storage (NAS) system to provide file shares to its high performance computing (HPC) workloads. The company wants to migrate its latency-sensitive HPC workloads and its storage to the AWS Cloud. The company must be able to provide NFS and SMB multi-protocol access from the file system.
Which solution will meet these requirements with the LEAST latency? (Select TWO.)

A. Deploy compute optimized EC2 instances into a cluster placement group.
B. Deploy compute optimized EC2 instances into a partition placement group.
C. Attach the EC2 instances to an Amazon FSx for Lustre file system.
D. Attach the EC2 instances to an Amazon FSx for OpenZFS file system.
E. Attach the EC2 instances to an Amazon FSx for NetApp ONTAP file system.

Correct Answer: A, E
Section:
Explanation:
A cluster placement group is a logical grouping of EC2 instances within a single Availability Zone that are placed close together to minimize network latency. This is suitable for latency-sensitive HPC workloads that require high network performance. A compute optimized EC2 instance is an instance type that has a high ratio of vCPUs to memory, which is ideal for compute-intensive applications. Amazon FSx for NetApp ONTAP is a fully managed service that provides NFS and SMB multi-protocol access from the file system, as well as features such as data deduplication, compression, thin provisioning, and snapshots. This solution will meet the requirements with the least latency, as it leverages the low-latency network and storage performance of AWS.
1 explains how cluster placement groups work and their benefits.
2 describes the characteristics and use cases of compute optimized EC2 instances.
3 provides an overview of Amazon FSx for NetApp ONTAP and its features.

QUESTION 141
A solutions architect needs to copy files from an Amazon S3 bucket to an Amazon Elastic File System (Amazon EFS) file system and another S3 bucket. The files must be copied continuously. New files are added to the original S3 bucket consistently. The copied files should be overwritten only if the source file changes.
Which solution will meet these requirements with the LEAST operational overhead?

A. Create an AWS DataSync location for both the destination S3 bucket and the EFS file system. Create a task for the destination S3 bucket and the EFS file system. Set the transfer mode to transfer only data that has changed.

B. Create an AWS Lambda function. Mount the file system to the function. Set up an S3 event notification to invoke the function when files are created and changed in Amazon S3. Configure the function to copy files to the file system and the destination S3 bucket.
C. Create an AWS DataSync location for both the destination S3 bucket and the EFS file system. Create a task for the destination S3 bucket and the EFS file system. Set the transfer mode to transfer all data.
D. Launch an Amazon EC2 instance in the same VPC as the file system. Mount the file system. Create a script to routinely synchronize all objects that changed in the origin S3 bucket to the destination S3 bucket and the mounted file system.

Correct Answer: A
Section:
Explanation:
AWS DataSync is a service that makes it easy to move large amounts of data between AWS storage services and on-premises storage systems. AWS DataSync can copy files from an S3 bucket to an EFS file system and another S3 bucket continuously, as well as overwrite only the files that have changed in the source. This solution will meet the requirements with the least operational overhead, as it does not require any code development or manual intervention.
4 explains how to create AWS DataSync locations for different storage services.
5 describes how to create and configure AWS DataSync tasks for data transfer.
6 discusses the different transfer modes that AWS DataSync supports.

QUESTION 142
A company deployed a serverless application that uses Amazon DynamoDB as a database layer The application has experienced a large increase in users. The company wants to improve database response time from milliseconds to microseconds and to cache requests to the database.
Which solution will meet these requirements with the LEAST operational overhead?

A. Use DynamoDB Accelerator (DAX).
B. Migrate the database to Amazon Redshift.
C. Migrate the database to Amazon RDS.
D. Use Amazon ElastiCache for Redis.

Correct Answer: A
Section:
Explanation:
DynamoDB Accelerator (DAX) is a fully managed, highly available caching service built for Amazon DynamoDB. DAX delivers up to a 10 times performance improvement---from milliseconds to microseconds---even at millions of requests per second. DAX does all the heavy lifting required to add in-memory acceleration to your DynamoDB tables, without requiring developers to manage cache invalidation, data population, or cluster management. Now you can focus on building great applications for your customers without worrying about performance at scale. You do not need to modify application logic because DAX is compatible with existing DynamoDB API calls. This solution will meet the requirements with the least operational overhead, as it does not require any code development or manual intervention.
1provides an overview of Amazon DynamoDB Accelerator (DAX) and its benefits.
2explains how to use DAX with DynamoDB for in-memory acceleration.

QUESTION 143
A company hosts multiple applications on AWS for different product lines. The applications use different compute resources, including Amazon EC2 instances and Application Load Balancers. The applications run in different AWS accounts under the same organization in AWS Organizations across multiple AWS Regions. Teams for each product line have tagged each compute resource in the individual accounts.
The company wants more details about the cost for each product line from the consolidated billing feature in Organizations.
Which combination of steps will meet these requirements? (Select TWO.)

A. Select a specific AWS generated tag in the AWS Billing console.
B. Select a specific user-defined tag in the AWS Billing console.
C. Select a specific user-defined tag in the AWS Resource Groups console.
D. Activate the selected tag from each AWS account.
E. Activate the selected tag from the Organizations management account.

Correct Answer: B, E
Section:
Explanation:
User-defined tags are key-value pairs that can be applied to AWS resources to categorize and track them. User-defined tags can also be used to allocate costs and create detailed billing reports in the AWS Billing console. To use user-defined tags for cost allocation, the tags must be activated from the Organizations management account, which is the root account that has full control over all the member accounts in the organization. Once activated, the user-defined tags will appear as columns in the cost allocation report, and can be used to filter and group costs by product line. This solution will meet the requirements with the least operational overhead, as it leverages the existing tagging strategy and does not require any code development or manual intervention.
1 explains how to use user-defined tags for cost allocation.
2 describes how to access and manage member accounts from the Organizations management account.
3 discusses how to create and view cost allocation reports in the AWS Billing console.

QUESTION 144
A company uses Amazon EC2 instances and Amazon Elastic Block Store (Amazon EBS) volumes to run an application. The company creates one snapshot of each EBS volume every day to meet compliance requirements. The company wants to implement an architecture that prevents the accidental deletion of EBS volume snapshots. The solution must not change the administrative rights of the storage administrator user.
Which solution will meet these requirements with the LEAST administrative effort?

A. Create an IAM role that has permission to delete snapshots. Attach the role to a new EC2 instance. Use the AWS CLI from the new EC2 instance to delete snapshots.
B. Create an IAM policy that denies snapshot deletion. Attach the policy to the storage administrator user.
C. Add tags to the snapshots. Create retention rules in Recycle Bin for EBS snapshots that have the tags.
D. Lock the EBS snapshots to prevent deletion.

Correct Answer: D
Section:
Explanation:
EBS snapshots are point-in-time backups of EBS volumes that can be used to restore data or create new volumes. EBS snapshots can be locked to prevent accidental deletion using a feature called EBS Snapshot Lock. When a snapshot is locked, it cannot be deleted by any user, including the root user, until it is unlocked. The lock policy can also specify a retention period, after which the snapshot can be deleted. This solution will meet the requirements with the least administrative effort, as it does not require any code development or policy changes.
1 explains how to lock and unlock EBS snapshots using EBS Snapshot Lock.
2 describes the concept and benefits of EBS snapshots.

QUESTION 145
A company maintains an Amazon RDS database that maps users to cost centers. The company has accounts in an organization in AWS Organizations. The company needs a solution that will tag all resources that are created in a specific AWS account in the organization. The solution must tag each resource with the cost center ID of the user who created the resource.
Which solution will meet these requirements?

A. Move the specific AWS account to a new organizational unit (OU) in Organizations from the management account. Create a service control policy (SCP) that requires all existing resources to have the correct cost center tag before the resources are created. Apply the SCP to the new OU.
B. Create an AWS Lambda function to tag the resources after the Lambda function looks up the appropriate cost center from the RDS database. Configure an Amazon EventBridge rule that reacts to AWS CloudTrail events to invoke the Lambda function.
C. Create an AWS CloudFormation stack to deploy an AWS Lambda function. Configure the Lambda function to look up the appropriate cost center from the RDS database and to tag resources. Create an Amazon EventBridge scheduled rule to invoke the CloudFormation stack.
D. Create an AWS Lambda function to tag the resources with a default value. Configure an Amazon EventBridge rule that reacts to AWS CloudTrail events to invoke the Lambda function when a resource is missing the cost center tag.

Correct Answer: B
Section:
Explanation:
AWS Lambda is a serverless compute service that lets you run code without provisioning or managing servers. Lambda can be used to tag resources with the cost center ID of the user who created the resource, by querying the RDS database that maps users to cost centers. Amazon EventBridge is a serverless event bus service that enables event-driven architectures. EventBridge can be configured to react to AWS CloudTrail events, which are

recorded API calls made by or on behalf of the AWS account. EventBridge can invoke the Lambda function when a resource is created in the specific AWS account, passing the user identity and resource information as parameters. This solution will meet the requirements, as it enables automatic tagging of resources based on the user and cost center mapping.
1 provides an overview of AWS Lambda and its benefits.
2 provides an overview of Amazon EventBridge and its benefits.
3 explains the concept and benefits of AWS CloudTrail events.

QUESTION 146
A company is migrating its multi-tier on-premises application to AWS. The application consists of a single-node MySQL database and a multi-node web tier. The company must minimize changes to the application during the migration. The company wants to improve application resiliency after the migration.
Which combination of steps will meet these requirements? (Select TWO.)

A. Migrate the web tier to Amazon EC2 instances in an Auto Scaling group behind an Application Load Balancer.
B. Migrate the database to Amazon EC2 instances in an Auto Scaling group behind a Network Load Balancer.
C. Migrate the database to an Amazon RDS Multi-AZ deployment.
D. Migrate the web tier to an AWS Lambda function.
E. Migrate the database to an Amazon DynamoDB table.

Correct Answer: A, C
Section:
Explanation:
An Auto Scaling group is a collection of EC2 instances that share similar characteristics and can be scaled in or out automatically based on demand. An Auto Scaling group can be placed behind an Application Load Balancer, which is a type of Elastic Load Balancing load balancer that distributes incoming traffic across multiple targets in multiple Availability Zones. This solution will improve the resiliency of the web tier by providing high availability, scalability, and fault tolerance. An Amazon RDS Multi-AZ deployment is a configuration that automatically creates a primary database instance and synchronously replicates the data to a standby instance in a different Availability Zone. When a failure occurs, Amazon RDS automatically fails over to the standby instance without manual intervention. This solution will improve the resiliency of the database tier by providing data redundancy, backup support, and availability. This combination of steps will meet the requirements with minimal changes to the application during the migration.
1 describes the concept and benefits of an Auto Scaling group.
2 provides an overview of Application Load Balancers and their benefits.
3explains how Amazon RDS Multi-AZ deployments work and their benefits.

QUESTION 147
A company runs an SMB file server in its data center. The file server stores large files that the company frequently accesses for up to 7 days after the file creation date. After 7 days, the company needs to be able to access the files with a maximum retrieval time of 24 hours.
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 increase the company's storage space. Create an S3 Lifecycle policy to transition the data to S3 Glacier Deep Archive after 7 days.
C. Create an Amazon FSx File Gateway to increase the company's storage space. Create an Amazon S3 Lifecycle policy to transition the data after 7 days.
D. Configure access to Amazon S3 for each user. Create an S3 Lifecycle policy to transition the data to S3 Glacier Flexible Retrieval after 7 days.

Correct Answer: B
Section:
Explanation:
Amazon S3 File Gateway is a service that provides a file-based interface to Amazon S3, which appears as a network file share. It enables you to store and retrieve Amazon S3 objects through standard file storage protocols such as SMB. S3 File Gateway can also cache frequently accessed data locally for low-latency access. S3 Lifecycle policy is a feature that allows you to define rules that automate the management of your objects throughout their lifecycle. You can use S3 Lifecycle policy to transition objects to different storage classes based on their age and access patterns. S3 Glacier Deep Archive is a storage class that offers the lowest cost for long-term data archiving, with a retrieval time of 12 hours or 48 hours. This solution will meet the requirements, as it allows the company to store large files in S3 with SMB file access, and to move the files to S3 Glacier Deep Archive after 7 days for cost savings and compliance.
1provides an overview of Amazon S3 File Gateway and its benefits.

2explains how to use S3 Lifecycle policy to manage object storage lifecycle.
3 describes the features and use cases of S3 Glacier Deep Archive storage class.

QUESTION 148
A company has a three-tier environment on AWS that ingests sensor data from its users' devices The traffic flows through a Network Load Balancer (NIB) then to Amazon EC2 instances for the web tier and finally to EC2 instances for the application tier that makes database calls
What should a solutions architect do to improve the security of data in transit to the web tier?

A. Configure a TLS listener and add the server certificate on the NLB
B. Configure AWS Shield Advanced and enable AWS WAF on the NLB
C. Change the load balancer to an Application Load Balancer and attach AWS WAF to it
D. Encrypt the Amazon Elastic Block Store (Amazon EBS) volume on the EC2 instances using AWS Key Management Service (AWS KMS)

Correct Answer: A
Section:
Explanation:
A) How do you protect your data in transit?
Best Practices:
Implement secure key and certificate management: Store encryption keys and certificates securely and rotate them at appropriate time intervals while applying strict access control; for example, by using a certificate management service, such as AWS Certificate Manager (ACM).
Enforce encryption in transit: Enforce your defined encryption requirements based on appropriate standards and recommendations to help you meet your organizational, legal, and compliance requirements.
Automate detection of unintended data access: Use tools such as GuardDuty to automatically detect attempts to move data outside of defined boundaries based on data classification level, for example, to detect a trojan that is copying data to an unknown or untrusted network using the DNS protocol.
Authenticate network communications: Verify the identity of communications by using protocols that support authentication, such as Transport Layer Security (TLS) or IPsec.
https://wa.aws.amazon.com/wat.question.SEC_9.en.html

QUESTION 149
A company wants to migrate its on-premises Microsoft SQL Server Enterprise edition database to AWS. The company's online application uses the database to process transactions. The data analysis team uses the same production database to run reports for analytical processing. The company wants to reduce operational overhead by moving to managed services wherever possible.
Which solution will meet these requirements with the LEAST operational overhead?

A. Migrate to Amazon RDS for Microsoft SQL Server. Use read replicas for reporting purposes.
B. Migrate to Microsoft SQL Server on Amazon EC2. Use Always On read replicas for reporting purposes.
C. Migrate to Amazon DynamoDB. Use DynamoDB on-demand replicas for reporting purposes.
D. Migrate to Amazon Aurora MySQL. Use Aurora read replicas for reporting purposes.

Correct Answer: A
Section:
Explanation:
Amazon RDS for Microsoft SQL Server is a fully managed service that offers SQL Server 2014, 2016, 2017, and 2019 editions while offloading database administration tasks such as backups, patching, and scaling. Amazon RDS supports read replicas, which are read-only copies of the primary database that can be used for reporting purposes without affecting the performance of the online application. This solution will meet the requirements with the least operational overhead, as it does not require any code changes or manual intervention.
1provides an overview of Amazon RDS for Microsoft SQL Server and its benefits.
2 explains how to create and use read replicas with Amazon RDS.

QUESTION 150
A company has deployed its newest product on AWS. The product runs in an Auto Scaling group behind a Network Load Balancer. The company stores the product's objects in an Amazon S3 bucket.

The company recently experienced malicious attacks against its systems. The company needs a solution that continuously monitors for malicious activity in the AWS account, workloads, and access patterns to the S3 bucket. The solution must also report suspicious activity and display the information on a dashboard.
Which solution will meet these requirements?

A. Configure Amazon Made to monitor and report findings to AWS Config.
B. Configure Amazon Inspector to monitor and report findings to AWS CloudTrail.
C. Configure Amazon GuardDuty to monitor and report findings to AWS Security Hub.
D. Configure AWS Config to monitor and report findings to Amazon EventBridge.

Correct Answer: C
Section:
Explanation:
Amazon GuardDuty is a threat detection service that continuously monitors for malicious activity and unauthorized behavior across the AWS account and workloads. GuardDuty analyzes data sources such as AWS CloudTrail event logs, Amazon VPC Flow Logs, and DNS logs to identify potential threats such as compromised instances, reconnaissance, port scanning, and data exfiltration. GuardDuty can report its findings to AWS Security Hub, which is a service that provides a comprehensive view of the security posture of the AWS account and workloads. Security Hub aggregates, organizes, and prioritizes security alerts from multiple AWS services and partner solutions, and displays them on a dashboard. This solution will meet the requirements, as it enables continuous monitoring, reporting, and visualization of malicious activity in the AWS account, workloads, and access patterns to the S3 bucket.
1 provides an overview of Amazon GuardDuty and its benefits.
2 explains how GuardDuty generates and reports findings based on threat detection.
3 provides an overview of AWS Security Hub and its benefits.
4 describes how Security Hub collects and displays findings from multiple sources on a dashboard

QUESTION 151
A company is developing an application that will run on a production Amazon Elastic Kubernetes Service (Amazon EKS) cluster The EKS cluster has managed node groups that are provisioned with On-Demand Instances.
The company needs a dedicated EKS cluster for development work. The company will use the development cluster infrequently to test the resiliency of the application. The EKS cluster must manage all the nodes.
Which solution will meet these requirements MOST cost-effectively?

A. Create a managed node group that contains only Spot Instances.
B. Create two managed node groups. Provision one node group with On-Demand Instances. Provision the second node group with Spot Instances.
C. Create an Auto Scaling group that has a launch configuration that uses Spot Instances. Configure the user data to add the nodes to the EKS cluster.
D. Create a managed node group that contains only On-Demand Instances.

Correct Answer: A
Section:
Explanation:
Spot Instances are EC2 instances that are available at up to a 90% discount compared to On-Demand prices. Spot Instances are suitable for stateless, fault-tolerant, and flexible workloads that can tolerate interruptions. Spot Instances can be reclaimed by EC2 when the demand for On-Demand capacity increases, but they provide a two-minute warning before termination. EKS managed node groups automate the provisioning and lifecycle management of nodes for EKS clusters. Managed node groups can use Spot Instances to reduce costs and scale the cluster based on demand. Managed node groups also support features such as Capacity Rebalancing and Capacity Optimized allocation strategy to improve the availability and resilience of Spot Instances. This solution will meet the requirements most cost-effectively, as it leverages the lowest-priced EC2 capacity and does not require any manual intervention.
1 explains how to create and use managed node groups with EKS.
2 describes how to use Spot Instances with managed node groups.
3 provides an overview of Spot Instances and their benefits.

QUESTION 152
A global company runs its applications in multiple AWS accounts in AWS Organizations. The company's applications use multipart uploads to upload data to multiple Amazon S3 buckets across AWS Regions. The company wants to report on incomplete multipart uploads for cost compliance purposes.
Which solution will meet these requirements with the LEAST operational overhead?

A. Configure AWS Config with a rule to report the incomplete multipart upload object count.
B. Create a service control policy (SCP) to report the incomplete multipart upload object count.
C. Configure S3 Storage Lens to report the incomplete multipart upload object count.
D. Create an S3 Multi-Region Access Point to report the incomplete multipart upload object count.

Correct Answer: C
Section:
Explanation:
S3 Storage Lens is a cloud storage analytics feature that provides organization-wide visibility into object storage usage and activity across multiple AWS accounts in AWS Organizations. S3 Storage Lens can report the incomplete multipart upload object count as one of the metrics that it collects and displays on an interactive dashboard in the S3 console. S3 Storage Lens can also export metrics in CSV or Parquet format to an S3 bucket for further analysis. This solution will meet the requirements with the least operational overhead, as it does not require any code development or policy changes.
1explains how to use S3 Storage Lens to gain insights into S3 storage usage and activity.
2 describes the concept and benefits of multipart uploads.

QUESTION 153
A company has deployed its application on Amazon EC2 instances with an Amazon RDS database. The company used the principle of least privilege to configure the database access credentials. The company's security team wants to protect the application and the database from SQL injection and other web-based attacks.
Which solution will meet these requirements with the LEAST operational overhead?

A. Use security groups and network ACLs to secure the database and application servers.
B. Use AWS WAF to protect the application. Use RDS parameter groups to configure the security settings.
C. Use AWS Network Firewall to protect the application and the database.
D. Use different database accounts in the application code for different functions. Avoid granting excessive privileges to the database users.

Correct Answer: B
Section:
Explanation:
AWS WAF is a web application firewall that helps protect web applications from common web exploits that could affect application availability, compromise security, or consume excessive resources. AWS WAF allows users to create rules that block, allow, or count web requests based on customizable web security rules. One of the types of rules that can be created is an SQL injection rule, which allows users to specify a list of IP addresses or IP address ranges that they want to allow or block. By using AWS WAF to protect the application, the company can prevent SQL injection and other web-based attacks from reaching the application and the database.
RDS parameter groups are collections of parameters that define how a database instance operates. Users can modify the parameters in a parameter group to change the behavior and performance of the database. By using RDS parameter groups to configure the security settings, the company can enforce best practices such as disabling remote root login, requiring SSL connections, and limiting the maximum number of connections.
The other options are not correct because they do not effectively protect the application and the database from SQL injection and other web-based attacks. Using security groups and network ACLs to secure the database and application servers is not sufficient because they only filter traffic at the network layer, not at the application layer. Using AWS Network Firewall to protect the application and the database is not necessary because it is a stateful firewall service that provides network protection for VPCs, not for individual applications or databases. Using different database accounts in the application code for different functions is a good practice, but it does not prevent SQL injection attacks from exploiting vulnerabilities in the application code.
AWS WAF
How AWS WAF works
Working with IP match conditions
Working with DB parameter groups
Amazon RDS security best practices

QUESTION 154
A research company uses on-premises devices to generate data for analysis. The company wants to use the AWS Cloud to analyze the dat
a. The devices generate .csv files and support writing the data to SMB file share. Company analysts must be able to use SQL commands to query the data. The analysts will run queries periodically throughout the day.
Which combination of steps will meet these requirements MOST cost-effectively? (Select THREE.)

A. Deploy an AWS Storage Gateway on premises in Amazon S3 File Gateway mode.

B. Deploy an AWS Storage Gateway on premises in Amazon FSx File Gateway mode.
C. Set up an AWS Glue crawler to create a table based on the data that is in Amazon S3.
D. Set up an Amazon EMR cluster with EMR Fife System (EMRFS) to query the data that is in Amazon S3. Provide access to analysts.
E. Set up an Amazon Redshift cluster to query the data that is in Amazon S3. Provide access to analysts.
F. Set up Amazon Athena to query the data that is in Amazon S3. Provide access to analysts.

Correct Answer: A, C, F
Section:
Explanation:
To meet the requirements of the use case in a cost-effective way, the following steps are recommended:
Deploy an AWS Storage Gateway on premises in Amazon S3 File Gateway mode. This will allow the company to write the .csv files generated by the devices to an SMB file share, which will be stored as objects in Amazon S3 buckets. AWS Storage Gateway is a hybrid cloud storage service that integrates on-premises environments with AWS storage.Amazon S3 File Gateway mode provides a seamless way to connect to Amazon S3 and access a virtually unlimited amount of cloud storage1.
Set up an AWS Glue crawler to create a table based on the data that is in Amazon S3. This will enable the company to use standard SQL to query the data stored in Amazon S3 buckets. AWS Glue is a serverless data integration service that simplifies data preparation and analysis.AWS Glue crawlers can automatically discover and classify data from various sources, and create metadata tables in the AWS Glue Data Catalog2.The Data Catalog is a central repository that stores information about data sources and how to access them3.
Set up Amazon Athena to query the data that is in Amazon S3. This will provide the company analysts with a serverless and interactive query service that can analyze data directly in Amazon S3 using standard SQL. Amazon Athena is integrated with the AWS Glue Data Catalog, so users can easily point Athena at the data source tables defined by the crawlers.Amazon Athena charges only for the queries that are run, and offers a pay-per-query pricing model, which makes it a cost-effective option for periodic queries4.
The other options are not correct because they are either not cost-effective or not suitable for the use case. Deploying an AWS Storage Gateway on premises in Amazon FSx File Gateway mode is not correct because this mode provides low-latency access to fully managed Windows file shares in AWS, which is not required for the use case. Setting up an Amazon EMR cluster with EMR File System (EMRFS) to query the data that is in Amazon S3 is not correct because this option involves setting up and managing a cluster of EC2 instances, which adds complexity and cost to the solution. Setting up an Amazon Redshift cluster to query the data that is in Amazon S3 is not correct because this option also involves provisioning and managing a cluster of nodes, which adds overhead and cost to the solution.
What is AWS Storage Gateway?
What is AWS Glue?
AWS Glue Data Catalog
What is Amazon Athena?

QUESTION 155
An ecommerce company is running a seasonal online sale. The company hosts its website on Amazon EC2 instances spanning multiple Availability Zones. The company wants its website to manage sudden traffic increases during the sale.
Which solution will meet these requirements MOST cost-effectively?

A. Create an Auto Scaling group that is large enough to handle peak traffic load. Stop half of the Amazon EC2 instances. Configure the Auto Scaling group to use the stopped instances to scale out when traffic increases.
B. Create an Auto Scaling group for the website. Set the minimum size of the Auto Scaling group so that it can handle high traffic volumes without the need to scale out.
C. Use Amazon CloudFront and Amazon ElastiCache to cache dynamic content with an Auto Scaling group set as the origin. Configure the Auto Scaling group with the instances necessary to populate CloudFront and ElastiCache. Scale in after the cache is fully populated.
D. Configure an Auto Scaling group to scale out as traffic increases. Create a launch template to start new instances from a preconfigured Amazon Machine Image (AMI).

Correct Answer:
Section:
Explanation:
Question no: 564 Verified Answer: D Comprehensive and Detailed The
AWS Transfer for SFTP
Amazon S3
Amazon EC2 Auto Scaling

QUESTION 156

A company has a nightly batch processing routine that analyzes report files that an on-premises file system receives daily through SFTP. The company wants to move the solution to the AWS Cloud. The solution must be highly available and resilient. The solution also must minimize operational effort.
Which solution meets these requirements?

A. Deploy AWS Transfer for SFTP and an Amazon Elastic File System (Amazon EFS) file system for storage. Use an Amazon EC2 instance in an Auto Scaling group with a scheduled scaling policy to run the batch operation.
B. Deploy an Amazon EC2 instance that runs Linux and an SFTP service. Use an Amazon Elastic Block Store {Amazon EBS) volume for storage. Use an Auto Scaling group with the minimum number of instances and desired number of instances set to 1.
C. Deploy an Amazon EC2 instance that runs Linux and an SFTP service. Use an Amazon Elastic File System (Amazon EFS) file system for storage. Use an Auto Scaling group with the minimum number of instances and desired number of instances set to 1.
D. Deploy AWS Transfer for SFTP and an Amazon S3 bucket for storage. Modify the application to pull the batch files from Amazon S3 to an Amazon EC2 instance for processing. Use an EC2 instance in an Auto Scaling group with a scheduled scaling policy to run the batch operation.

Correct Answer: D
Section:
Explanation:
The solution that meets the requirements of high availability, performance, security, and static IP addresses is to use Amazon CloudFront, Application Load Balancers (ALBs), Amazon Route 53, and AWS WAF. This solution allows the company to distribute its HTTP-based application globally using CloudFront, which is a content delivery network (CDN) service that caches content at edge locations and provides static IP addresses for each edge location. The company can also use Route 53 latency-based routing to route requests to the closest ALB in each Region, which balances the load across the EC2 instances. The company can also deploy AWS WAF on the CloudFront distribution to protect the application against common web exploits by creating rules that allow, block, or count web requests based on conditions that are defined. The other solutions do not meet all the requirements because they either use Network Load Balancers (NLBs), which do not support HTTP-based applications, or they do not use CloudFront, which provides better performance and security than AWS Global Accelerator.Reference:=
Amazon CloudFront
Application Load Balancer
Amazon Route 53
AWS WAF

QUESTION 157
A company has users all around the world accessing its HTTP-based application deployed on Amazon EC2 instances in multiple AWS Regions. The company wants to improve the availability and performance of the application. The company also wants to protect the application against common web exploits that may affect availability, compromise security, or consume excessive resources. Static IP addresses are required.
What should a solutions architect recommend to accomplish this?

A. Put the EC2 instances behind Network Load Balancers (NLBs) in each Region. Deploy AWS WAF on the NLBs. Create an accelerator using AWS Global Accelerator and register the NLBs as endpoints.
B. Put the EC2 instances behind Application Load Balancers (ALBs) in each Region. Deploy AWS WAF on the ALBs. Create an accelerator using AWS Global Accelerator and register the ALBs as endpoints.
C. Put the EC2 instances behind Network Load Balancers (NLBs) in each Region. Deploy AWS WAF on the NLBs. Create an Amazon CloudFront distribution with an origin that uses Amazon Route 53 latency-based routing to route requests to the NLBs.
D. Put the EC2 instances behind Application Load Balancers (ALBs) in each Region. Create an Amazon CloudFront distribution with an origin that uses Amazon Route 53 latency-based routing to route requests to the ALBs. Deploy AWS WAF on the CloudFront distribution.

Correct Answer: A
Section:
Explanation:
The company wants to improve the availability and performance of the application, as well as protect it against common web exploits. The company also needs static IP addresses for the application. To meet these requirements, a solutions architect should recommend the following solution:
Put the EC2 instances behind Network Load Balancers (NLBs) in each Region. NLBs are designed to handle millions of requests per second while maintaining high throughput at ultra-low latency. NLBs also support static IP addresses for each Availability Zone, which can be useful for whitelisting or firewalling purposes.
Deploy AWS WAF on the NLBs. AWS WAF is a web application firewall that helps protect web applications from common web exploits that could affect availability, security, or performance. AWS WAF lets you define customizable web security rules that control which traffic to allow or block to your web applications.
Create an accelerator using AWS Global Accelerator and register the NLBs as endpoints. AWS Global Accelerator is a service that improves the availability and performance of your applications with local or global users. It provides static IP addresses that act as a fixed entry point to your application endpoints in any AWS Region. It uses the AWS global network to optimize the path from your users to your applications, improving the

performance of your TCP and UDP traffic.
This solution will provide high availability across Availability Zones and Regions, improve performance by routing traffic over the AWS global network, protect the application from common web attacks, and provide static IP addresses for the application.
Network Load Balancer
AWS WAF
AWS Global Accelerator

QUESTION 158
A company wants to deploy its containerized application workloads to a VPC across three Availability Zones. The company needs a solution that is highly available across Availability Zones. The solution must require minimal changes to the application.
Which solution will meet these requirements with the LEAST operational overhead?

A. Use Amazon Elastic Container Service (Amazon ECS). Configure Amazon ECS Service Auto Scaling to use target tracking scaling. Set the minimum capacity to 3. Set the task placement strategy type to spread with an Availability Zone attribute.
B. Use Amazon Elastic Kubernetes Service (Amazon EKS) self-managed nodes. Configure Application Auto Scaling to use target tracking scaling. Set the minimum capacity to 3.
C. Use Amazon EC2 Reserved Instances. Launch three EC2 instances in a spread placement group. Configure an Auto Scaling group to use target tracking scaling. Set the minimum capacity to 3.
D. Use an AWS Lambda function. Configure the Lambda function to connect to a VPC. Configure Application Auto Scaling to use Lambda as a scalable target. Set the minimum capacity to 3.

Correct Answer: A
Section:
Explanation:
The company wants to deploy its containerized application workloads to a VPC across three Availability Zones, with high availability and minimal changes to the application. The solution that will meet these requirements with the least operational overhead is:
Use Amazon Elastic Container Service (Amazon ECS). Amazon ECS is a fully managed container orchestration service that allows you to run and scale containerized applications on AWS. Amazon ECS eliminates the need for you to install, operate, and scale your own cluster management infrastructure. Amazon ECS also integrates with other AWS services, such as VPC, ELB, CloudFormation, CloudWatch, IAM, and more.
Configure Amazon ECS Service Auto Scaling to use target tracking scaling. Amazon ECS Service Auto Scaling allows you to automatically adjust the number of tasks in your service based on the demand or custom metrics. Target tracking scaling is a policy type that adjusts the number of tasks in your service to keep a specified metric at a target value. For example, you can use target tracking scaling to maintain a target CPU utilization or request count per task for your service.
Set the minimum capacity to 3. This ensures that your service always has at least three tasks running across three Availability Zones, providing high availability and fault tolerance for your application.
Set the task placement strategy type to spread with an Availability Zone attribute. This ensures that your tasks are evenly distributed across the Availability Zones in your cluster, maximizing the availability of your service.
This solution will provide high availability across Availability Zones, require minimal changes to the application, and reduce the operational overhead of managing your own cluster infrastructure.
Amazon Elastic Container Service
Amazon ECS Service Auto Scaling
Target Tracking Scaling Policies for Amazon ECS Services
Amazon ECS Task Placement Strategies

QUESTION 159
A company uses high concurrency AWS Lambda functions to process a constantly increasing number of messages in a message queue during marketing events. The Lambda functions use CPU intensive code to process the messages. The company wants to reduce the compute costs and to maintain service latency for its customers.
Which solution will meet these requirements?

A. Configure reserved concurrency for the Lambda functions. Decrease the memory allocated to the Lambda functions.
B. Configure reserved concurrency for the Lambda functions. Increase the memory according to AWS Compute Optimizer recommendations.
C. Configure provisioned concurrency for the Lambda functions. Decrease the memory allocated to the Lambda functions.
D. Configure provisioned concurrency for the Lambda functions. Increase the memory according to AWS Compute Optimizer recommendations.

Correct Answer: D
Section:
Explanation:

The company wants to reduce the compute costs and maintain service latency for its Lambda functions that process a constantly increasing number of messages in a message queue. The Lambda functions use CPU intensive code to process the messages. To meet these requirements, a solutions architect should recommend the following solution:
Configure provisioned concurrency for the Lambda functions. Provisioned concurrency is the number of pre-initialized execution environments that are allocated to the Lambda functions. These execution environments are prepared to respond immediately to incoming function requests, reducing the cold start latency. Configuring provisioned concurrency also helps to avoid throttling errors due to reaching the concurrency limit of the Lambda service.
Increase the memory according to AWS Compute Optimizer recommendations. AWS Compute Optimizer is a service that provides recommendations for optimal AWS resource configurations based on your utilization data. By increasing the memory allocated to the Lambda functions, you can also increase the CPU power and improve the performance of your CPU intensive code. AWS Compute Optimizer can help you find the optimal memory size for your Lambda functions based on your workload characteristics and performance goals.
This solution will reduce the compute costs by avoiding unnecessary over-provisioning of memory and CPU resources, and maintain service latency by using provisioned concurrency and optimal memory size for the Lambda functions.
Provisioned Concurrency
AWS Compute Optimizer

QUESTION 160
A company's ecommerce website has unpredictable traffic and uses AWS Lambda functions to directly access a private Amazon RDS for PostgreSQL DB instance. The company wants to maintain predictable database performance and ensure that the Lambda invocations do not overload the database with too many connections.
What should a solutions architect do to meet these requirements?

A. Point the client driver at an RDS custom endpoint. Deploy the Lambda functions inside a VPC.
B. Point the client driver at an RDS proxy endpoint. Deploy the Lambda functions inside a VPC.
C. Point the client driver at an RDS custom endpoint. Deploy the Lambda functions outside a VPC.
D. Point the client driver at an RDS proxy endpoint. Deploy the Lambda functions outside a VPC.

Correct Answer: B
Section:
Explanation:
To maintain predictable database performance and ensure that the Lambda invocations do not overload the database with too many connections, a solutions architect should point the client driver at an RDS proxy endpoint and deploy the Lambda functions inside a VPC. An RDS proxy is a fully managed database proxy that allows applications to share connections to a database, improving database availability and scalability. By using an RDS proxy, the Lambda functions can reuse existing connections, rather than creating new ones for every invocation, reducing the connection overhead and latency. Deploying the Lambda functions inside a VPC allows them to access the private RDS DB instance securely and efficiently, without exposing it to the public internet.Reference:
Using Amazon RDS Proxy with AWS Lambda
Configuring a Lambda function to access resources in a VPC

QUESTION 161
A development team is collaborating with another company to create an integrated product. The other company needs to access an Amazon Simple Queue Service (Amazon SQS) queue that is contained in the development team's account. The other company wants to poll the queue without giving up its own account permissions to do so.
How should a solutions architect provide access to the SQS queue?

A. Create an instance profile that provides the other company access to the SQS queue.
B. Create an IAM policy that provides the other company access to the SQS queue.
C. Create an SQS access policy that provides the other company access to the SQS queue.
D. Create an Amazon Simple Notification Service (Amazon SNS) access policy that provides the other company access to the SQS queue.

Correct Answer: C
Section:
Explanation:
To provide access to the SQS queue to the other company without giving up its own account permissions, a solutions architect should create an SQS access policy that provides the other company access to the SQS queue. An SQS access policy is a resource-based policy that defines who can access the queue and what actions they can perform. The policy can specify the AWS account ID of the other company as a principal, and grant permissions for actions such assqs:ReceiveMessage,sqs:DeleteMessage, andsqs:GetQueueAttributes. This way, the other company can poll the queue using its own credentials, without needing to assume a role or use cross-account

access keys.Reference:
Using identity-based policies (IAM policies) for Amazon SQS
Using custom policies with the Amazon SQS access policy language

QUESTION 162
A city has deployed a web application running on Amazon EC2 instances behind an Application Load Balancer (ALB). The application's users have reported sporadic performance, which appears to be related to DDoS attacks originating from random IP addresses. The city needs a solution that requires minimal configuration changes and provides an audit trail for the DDoS sources.
Which solution meets these requirements?

A. Enable an AWS WAF web ACL on the ALB, and configure rules to block traffic from unknown sources.
B. Subscribe to Amazon Inspector. Engage the AWS DDoS Response Team (DRT) to integrate mitigating controls into the service.
C. Subscribe to AWS Shield Advanced. Engage the AWS DDoS Response Team (DRT) to integrate mitigating controls into the service.
D. Create an Amazon CloudFront distribution for the application, and set the ALB as the origin. Enable an AWS WAF web ACL on the distribution, and configure rules to block traffic from unknown sources.

Correct Answer: C
Section:
Explanation:
To protect the web application from DDoS attacks originating from random IP addresses, a solutions architect should subscribe to AWS Shield Advanced and engage the AWS DDoS Response Team (DRT) to integrate mitigating controls into the service. AWS Shield Advanced is a managed service that provides protection against large and sophisticated DDoS attacks, with access to 24/7 support and response from the DRT. The DRT can help the city configure proactive and reactive safeguards, such as AWS WAF rules, rate-based rules, and network ACLs, to block malicious traffic and improve the application's resilience. The service also provides an audit trail for the DDoS sources through detailed attack reports and Amazon CloudWatch metrics.

QUESTION 163
A company that uses AWS needs a solution to predict the resources needed for manufacturing processes each month. The solution must use historical values that are currently stored in an Amazon S3 bucket The company has no machine learning (ML) experience and wants to use a managed service for the training and predictions.
Which combination of steps will meet these requirements? (Select TWO.)

A. Deploy an Amazon SageMaker model. Create a SageMaker endpoint for inference.
B. Use Amazon SageMaker to train a model by using the historical data in the S3 bucket.
C. Configure an AWS Lambda function with a function URL that uses Amazon SageMaker endpoints to create predictions based on the inputs.
D. Configure an AWS Lambda function with a function URL that uses an Amazon Forecast predictor to create a prediction based on the inputs.
E. Train an Amazon Forecast predictor by using the historical data in the S3 bucket.

Correct Answer: B, E
Section:
Explanation:
To predict the resources needed for manufacturing processes each month using historical values that are currently stored in an Amazon S3 bucket, a solutions architect should use Amazon SageMaker to train a model by using the historical data in the S3 bucket, and deploy an Amazon SageMaker model and create a SageMaker endpoint for inference. Amazon SageMaker is a fully managed service that provides an easy way to build, train, and deploy machine learning (ML) models. The solutions architect can use the built-in algorithms or frameworks provided by SageMaker, or bring their own custom code, to train a model using the historical data in the S3 bucket as input. The trained model can then be deployed to a SageMaker endpoint, which is a scalable and secure web service that can handle requests for predictions from the application. The solutions architect does not need to have any ML experience or manage any infrastructure to use SageMaker.

QUESTION 164
A company's developers want a secure way to gain SSH access on the company's Amazon EC2 instances that run the latest version of Amazon Linux. The developers work remotely and in the corporate office.
The company wants to use AWS services as a part of the solution. The EC2 instances are hosted in a VPC private subnet and access the internet through a NAT gateway that is deployed in a public subnet.
What should a solutions architect do to meet these requirements MOST cost-effectively?

A. Create a bastion host in the same subnet as the EC2 instances. Grant the ec2: CreateVpnConnection 1AM permission to the developers. Install EC2 Instance Connect so that the developers can connect to the EC2

instances.

B. Create an AWS Site-to-Site VPN connection between the corporate network and the VPC. Instruct the developers to use the Site-to-Site VPN connection to access the EC2 instances when the developers are on the corporate network. Instruct the developers to set up another VPN connection for access when they work remotely.

C. Create a bastion host in the public subnet of the VPC. Configure the security groups and SSH keys of the bastion host to only allow connections and SSH authentication from the developers' corporate and remote networks. Instruct the developers to connect through the bastion host by using SSH to reach the EC2 instances.

D. Attach the AmazonSSMManagedlnstanceCore 1AM policy to an 1AM role that is associated with the EC2 instances. Instruct the developers to use AWS Systems Manager Session Manager to access the EC2 instances.

Correct Answer: D
Section:
Explanation:
AWS Systems Manager Session Manager is a service that enables you to securely connect to your EC2 instances without using SSH keys or bastion hosts. You can use Session Manager to access your instances through the AWS Management Console, the AWS CLI, or the AWS SDKs. Session Manager uses IAM policies and roles to control who can access which instances. By attaching the AmazonSSMManagedlnstanceCore IAM policy to an IAM role that is associated with the EC2 instances, you grant the Session Manager service the necessary permissions to perform actions on your instances. You also need to attach another IAM policy to the developers' IAM users or roles that allows them to start sessions to the instances. Session Manager uses the AWS Systems Manager Agent (SSM Agent) that is installed by default on Amazon Linux 2 and other supported Linux distributions. Session Manager also encrypts all session data between your client and your instances, and streams session logs to Amazon S3, Amazon CloudWatch Logs, or both for auditing purposes. This solution is the most cost-effective, as it does not require any additional resources or services, such as bastion hosts, VPN connections, or NAT gateways. It also simplifies the security and management of SSH access, as it eliminates the need for SSH keys, port opening, or firewall rules.Reference:
What is AWS Systems Manager?
Setting up Session Manager
Getting started with Session Manager
Controlling access to Session Manager
Logging Session Manager activity

QUESTION 165
A company runs a highly available web application on Amazon EC2 instances behind an Application Load Balancer The company uses Amazon CloudWatch metrics
As the traffic to the web application Increases, some EC2 instances become overloaded with many outstanding requests The CloudWatch metrics show that the number of requests processed and the time to receive the responses from some EC2 instances are both higher compared to other EC2 instances The company does not want new requests to be forwarded to the EC2 instances that are already overloaded.
Which solution will meet these requirements?

A. Use the round robin routing algorithm based on the RequestCountPerTarget and Active Connection Count CloudWatch metrics.

B. Use the least outstanding requests algorithm based on the RequestCountPerTarget and ActiveConnectionCount CloudWatch metrics.

C. Use the round robin routing algorithm based on the RequestCount and TargetResponseTime CloudWatch metrics.

D. Use the least outstanding requests algorithm based on the RequestCount and TargetResponseTime CloudWatch metrics.

Correct Answer: D
Section:
Explanation:
The least outstanding requests (LOR) algorithm is a load balancing algorithm that distributes incoming requests to the target with the fewest outstanding requests. This helps to avoid overloading any single target and improves the overall performance and availability of the web application. The LOR algorithm can use the RequestCount and TargetResponseTime CloudWatch metrics to determine the number of outstanding requests and the response time of each target. These metrics measure the number of requests processed by each target and the time elapsed after the request leaves the load balancer until a response from the target is received by the load balancer, respectively. By using these metrics, the LOR algorithm can route new requests to the targets that are less busy and more responsive, and avoid sending requests to the targets that are already overloaded or slow. This solution meets the requirements of the company.
Application Load Balancer now supports Least Outstanding Requests algorithm for load balancing requests
Target groups for your Application Load Balancers
Elastic Load Balancing - Application Load Balancers

QUESTION 166
A pharmaceutical company is developing a new drug. The volume of data that the company generates has grown exponentially over the past few months. The company's researchers regularly require a subset of the entire dataset to be immediately available with minimal lag. However the entire dataset does not need to be accessed on a daily basis. All the data currently resides in on-premises storage arrays, and the company wants to reduce

ongoing capital expenses.
Which storage solution should a solutions architect recommend to meet these requirements?

A. Run AWS DataSync as a scheduled cron job to migrate the data to an Amazon S3 bucket on an ongoing basis.
B. Deploy an AWS Storage Gateway file gateway with an Amazon S3 bucket as the target storage Migrate the data to the Storage Gateway appliance.
C. Deploy an AWS Storage Gateway volume gateway with cached volumes with an Amazon S3 bucket as the target storage. Migrate the data to the Storage Gateway appliance.
D. Configure an AWS Site-to-Site VPN connection from the on-premises environment to AWS. Migrate data to an Amazon Elastic File System (Amazon EFS) file system.

Correct Answer: C
Section:
Explanation:
AWS Storage Gateway is a hybrid cloud storage service that allows you to seamlessly integrate your on-premises applications with AWS cloud storage. Volume Gateway is a type of Storage Gateway that presents cloud-backed iSCSI block storage volumes to your on-premises applications. Volume Gateway operates in either cache mode or stored mode. In cache mode, your primary data is stored in Amazon S3, while retaining your frequently accessed data locally in the cache for low latency access. In stored mode, your primary data is stored locally and your entire dataset is available for low latency access on premises while also asynchronously getting backed up to Amazon S3.
For the pharmaceutical company's use case, cache mode is the most suitable option, as it meets the following requirements:
It reduces the need to scale the on-premises storage infrastructure, as most of the data is stored in Amazon S3, which is scalable, durable, and cost-effective.
It provides low latency access to the subset of the data that the researchers regularly require, as it is cached locally in the Storage Gateway appliance.
It does not require the entire dataset to be accessed on a daily basis, as it is stored in Amazon S3 and can be retrieved on demand.
It offers flexible data protection and recovery options, as it allows taking point-in-time copies of the volumes using AWS Backup, which are stored in AWS as Amazon EBS snapshots.
Therefore, the solutions architect should recommend deploying an AWS Storage Gateway volume gateway with cached volumes with an Amazon S3 bucket as the target storage and migrating the data to the Storage Gateway appliance.
Volume Gateway | Amazon Web Services
How Volume Gateway works (architecture) - AWS Storage Gateway
AWS Storage Volume Gateway - Cached volumes - Stack Overflow

QUESTION 167
A company hosts an application used to upload files to an Amazon S3 bucket Once uploaded, the files are processed to extract metadata which takes less than 5 seconds The volume and frequency of the uploads varies from a few files each hour to hundreds of concurrent uploads The company has asked a solutions architect to design a cost-effective architecture that will meet these requirements.
What should the solutions architect recommend?

A. Configure AWS CloudTrail trails to tog S3 API calls Use AWS AppSync to process the files.
B. Configure an object-created event notification within the S3 bucket to invoke an AWS Lambda function to process the files.
C. Configure Amazon Kinesis Data Streams to process and send data to Amazon S3. Invoke an AWS Lambda function to process the files.
D. Configure an Amazon Simple Notification Service (Amazon SNS) topic to process the files uploaded to Amazon S3 Invoke an AWS Lambda function to process the files.

Correct Answer: B
Section:
Explanation:
This option is the most cost-effective and scalable way to process the files uploaded to S3. AWS CloudTrail is used to log API calls, not to trigger actions based on them. AWS AppSync is a service for building GraphQL APIs, not for processing files. Amazon Kinesis Data Streams is used to ingest and process streaming data, not to send data to S3. Amazon SNS is a pub/sub service that can be used to notify subscribers of events, not to process files.Reference:
Using AWS Lambda with Amazon S3
AWS CloudTrail FAQs
What Is AWS AppSync?
[What Is Amazon Kinesis Data Streams?]
[What Is Amazon Simple Notification Service?]

QUESTION 168

A company uses Amazon S3 to store high-resolution pictures in an S3 bucket. To minimize application changes, the company stores the pictures as the latest version of an S3 object
The company needs to retain only the two most recent versions ot the pictures.
The company wants to reduce costs. The company has identified the S3 bucket as a large expense.
Which solution will reduce the S3 costs with the LEAST operational overhead?

A. Use S3 Lifecycle to delete expired object versions and retain the two most recent versions.
B. Use an AWS Lambda function to check for older versions and delete all but the two most recent versions
C. Use S3 Batch Operations to delete noncurrent object versions and retain only the two most recent versions
D. Deactivate versioning on the S3 bucket and retain the two most recent versions.

Correct Answer: A
Section:
Explanation:
S3 Lifecycle is a feature that allows you to automate the management of your S3 objects based on predefined rules. You can use S3 Lifecycle to delete expired object versions and retain the two most recent versions by creating a lifecycle configuration rule that applies to all objects in the bucket and specifies the expiration action for noncurrent versions. This way, you can reduce the storage costs of your S3 bucket without requiring any application changes or manual intervention. S3 Lifecycle runs once a day and marks the eligible object versions for deletion. You are no longer charged for the objects that are marked for deletion. S3 Lifecycle is the most cost-effective and simple solution among the options.
B) Use an AWS Lambda function to check for older versions and delete all but the two most recent versions. This option is not optimal because it requires you to write, test, and maintain a custom Lambda function that scans the S3 bucket for older versions and deletes them. This can incur additional costs for Lambda invocations and increase the operational complexity and overhead. Moreover, you need to ensure that your Lambda function has the appropriate permissions and error handling mechanisms to perform the deletion operation.
C) Use S3 Batch Operations to delete noncurrent object versions and retain only the two most recent versions. This option is not ideal because S3 Batch Operations is designed for performing large-scale operations on S3 objects, such as copying, tagging, restoring, or invoking a Lambda function. To use S3 Batch Operations to delete noncurrent object versions, you need to provide a manifest file that lists the object versions that you want to delete. This can be challenging and time-consuming to generate and update. Moreover, S3 Batch Operations charges you for each operation that you perform, which can increase your costs.
D) Deactivate versioning on the S3 bucket and retain the two most recent versions. This option is not feasible because deactivating versioning on an S3 bucket does not delete the existing object versions. Instead, it prevents new versions from being created. Therefore, you still need to delete the older versions manually or use another method to do so. Additionally, deactivating versioning can compromise the data protection and recovery capabilities of your S3 bucket.
1Considering four different replication options for data in Amazon S3 | AWS Storage Blog
2Using AWS Lambda with Amazon S3 batch operations - AWS Lambda
3Empty an Amazon S3 bucket with a lifecycle configuration rule
4Amazon S3 - Lifecycle Management - GeeksforGeeks

QUESTION 169
A solutions architect needs to design the architecture for an application that a vendor provides as a Docker container image The container needs 50 GB of storage available for temporary files The infrastructure must be serverless.
Which solution meets these requirements with the LEAST operational overhead?

A. Create an AWS Lambda function that uses the Docker container image with an Amazon S3 mounted volume that has more than 50 GB of space
B. Create an AWS Lambda function that uses the Docker container image with an Amazon Elastic Block Store (Amazon EBS) volume that has more than 50 GB of space
C. Create an Amazon Elastic Container Service (Amazon ECS) cluster that uses the AWS Fargate launch type Create a task definition for the container image with an Amazon Elastic File System (Amazon EFS) volume. Create a service with that task definition.
D. Create an Amazon Elastic Container Service (Amazon ECS) cluster that uses the Amazon EC2 launch type with an Amazon Elastic Block Store (Amazon EBS) volume that has more than 50 GB of space Create a task definition for the container image. Create a service with that task definition.

Correct Answer: C
Section:
Explanation:
The AWS Fargate launch type is a serverless way to run containers on Amazon ECS, without having to manage any underlying infrastructure. You only pay for the resources required to run your containers, and AWS handles the provisioning, scaling, and security of the cluster. Amazon EFS is a fully managed, elastic, and scalable file system that can be mounted to multiple containers, and provides high availability and durability. By using AWS Fargate and Amazon EFS, you can run your Docker container image with 50 GB of storage available for temporary files, with the least operational overhead. This solution meets the requirements of the question.

AWS Fargate
Amazon Elastic File System
Using Amazon EFS file systems with Amazon ECS

QUESTION 170
A company needs to extract the names of ingredients from recipe records that are stored as text files in an Amazon S3 bucket A web application will use the ingredient names to query an Amazon DynamoDB table and determine a nutrition score.
The application can handle non-food records and errors The company does not have any employees who have machine learning knowledge to develop this solution
Which solution will meet these requirements MOST cost-effectively?

A. Use S3 Event Notifications to invoke an AWS Lambda function when PutObject requests occur Program the Lambda function to analyze the object and extract the ingredient names by using Amazon Comprehend Store the Amazon Comprehend output in the DynamoDB table.
B. Use an Amazon EventBridge rule to invoke an AWS Lambda function when PutObject requests occur. Program the Lambda function to analyze the object by using Amazon Forecast to extract the ingredient names Store the Forecast output in the DynamoDB table.
C. Use S3 Event Notifications to invoke an AWS Lambda function when PutObject requests occur Use Amazon Polly to create audio recordings of the recipe records. Save the audio files in the S3 bucket Use Amazon Simple Notification Service (Amazon SNS) to send a URL as a message to employees Instruct the employees to listen to the audio files and calculate the nutrition score Store the ingredient names in the DynamoDB table.
D. Use an Amazon EventBridge rule to invoke an AWS Lambda function when a PutObject request occurs Program the Lambda function to analyze the object and extract the ingredient names by using Amazon SageMaker Store the inference output from the SageMaker endpoint in the DynamoDB table.

Correct Answer: A
Section:
Explanation:
This solution meets the following requirements:
It is cost-effective, as it only uses serverless components that are charged based on usage and do not require any upfront provisioning or maintenance.
It is scalable, as it can handle any number of recipe records that are uploaded to the S3 bucket without any performance degradation or manual intervention.
It is easy to implement, as it does not require any machine learning knowledge or complex data processing logic. Amazon Comprehend is a natural language processing service that can automatically extract entities such as ingredients from text files. The Lambda function can simply invoke the Comprehend API and store the results in the DynamoDB table.
It is reliable, as it can handle non-food records and errors gracefully. Amazon Comprehend can detect the language and domain of the text files and return an appropriate response. The Lambda function can also implement error handling and logging mechanisms to ensure the data quality and integrity.
Using AWS Lambda with Amazon S3 - AWS Lambda
What Is Amazon Comprehend? - Amazon Comprehend
Working with Tables - Amazon DynamoDB

QUESTION 171
A company needs to create an AWS Lambda function that will run in a VPC in the company's primary AWS account. The Lambda function needs to access files that the company stores in an Amazon Elastic File System (Amazon EFS) file system. The EFS file system is located in a secondary AWS account. As the company adds files to the file system the solution must scale to meet the demand.
Which solution will meet these requirements MOST cost-effectively?

A. Create a new EPS file system in the primary account Use AWS DataSync to copy the contents of the original EPS file system to the new EPS file system
B. Create a VPC peering connection between the VPCs that are in the primary account and the secondary account
C. Create a second Lambda function In the secondary account that has a mount that is configured for the file system. Use the primary account's Lambda function to invoke the secondary account's Lambda function
D. Move the contents of the file system to a Lambda Layer's Configure the Lambda layer's permissions to allow the company's secondary account to use the Lambda layer.

Correct Answer: B
Section:
Explanation:
This option is the most cost-effective and scalable way to allow the Lambda function in the primary account to access the EFS file system in the secondary account. VPC peering enables private connectivity between two VPCs without requiring gateways, VPN connections, or dedicated network connections. The Lambda function can use the VPC peering connection to mount the EFS file system as a local file system and access the files as needed. The solution does not incur additional data transfer or storage costs, and it leverages the existing EFS file system without duplicating or moving the data.

Option A is not cost-effective because it requires creating a new EFS file system and using AWS DataSync to copy the data from the original EFS file system. This would incur additional storage and data transfer costs, and it would not provide real-time access to the files.
Option C is not scalable because it requires creating a second Lambda function in the secondary account and configuring cross-account permissions to invoke it from the primary account. This would add complexity and latency to the solution, and it would increase the Lambda invocation costs.
Option D is not feasible because Lambda layers are not designed to store large amounts of data or provide file system access. Lambda layers are used to share common code or libraries across multiple Lambda functions. Moving the contents of the EFS file system to a Lambda layer would exceed the size limit of 250 MB for a layer, and it would not allow the Lambda function to read or write files to the layer.Reference:
What Is VPC Peering?
Using Amazon EFS file systems with AWS Lambda
What Are Lambda Layers?

QUESTION 172
A company has migrated a fleet of hundreds of on-premises virtual machines (VMs) to Amazon EC2 instances. The instances run a diverse fleet of Windows Server versions along with several Linux distributions. The company wants a solution that will automate inventory and updates of the operating systems. The company also needs a summary of common vulnerabilities of each instance for regular monthly reviews.
What should a solutions architect recommend to meet these requirements?

A. Set up AWS Systems Manager Patch Manager to manage all the EC2 instances. Configure AWS Security Hub to produce monthly reports.
B. Set up AWS Systems Manager Patch Manager to manage all the EC2 instances Deploy Amazon Inspector, and configure monthly reports
C. Set up AWS Shield Advanced, and configure monthly reports Deploy AWS Config to automate patch installations on the EC2 instances
D. Set up Amazon GuardDuty in the account to monitor all EC2 instances Deploy AWS Config to automate patch installations on the EC2 instances.

Correct Answer: B
Section:
Explanation:
Understanding the Requirement: The company needs to automate inventory and updates of diverse OS versions on EC2 instances and summarize common vulnerabilities for monthly reviews.
Analysis of Options:
Systems Manager Patch Manager and Security Hub: Patch Manager automates patching, but Security Hub is more focused on compliance and security posture rather than inventory and vulnerability management.
Systems Manager Patch Manager and Amazon Inspector: Patch Manager automates OS updates, and Amazon Inspector provides vulnerability assessments, making this a comprehensive solution for the requirements.
AWS Shield Advanced and AWS Config: Shield Advanced is for DDoS protection, not suitable for OS patch management and vulnerability reporting.
Amazon GuardDuty and AWS Config: GuardDuty is for threat detection and monitoring, not specifically for patch management and vulnerability assessments.
Best Solution:
Systems Manager Patch Manager and Amazon Inspector: This combination automates OS updates and provides detailed vulnerability assessments, meeting both the inventory and security reporting needs effectively.
AWS Systems Manager Patch Manager
Amazon Inspector

QUESTION 173
A company wants to use Amazon Elastic Container Service (Amazon ECS) to run its on-premises application in a hybrid environment The application currently runs on containers on premises.
The company needs a single container solution that can scale in an on-premises, hybrid, or cloud environment The company must run new application containers in the AWS Cloud and must use a load balancer for HTTP traffic.
Which combination of actions will meet these requirements? (Select TWO.)

A. Set up an ECS cluster that uses the AWS Fargate launch type for the cloud application containers Use an Amazon ECS Anywhere external launch type for the on-premises application containers.
B. Set up an Application Load Balancer for cloud ECS services
C. Set up a Network Load Balancer for cloud ECS services.
D. Set up an ECS cluster that uses the AWS Fargate launch type Use Fargate for the cloud application containers and the on-premises application containers.
E. Set up an ECS cluster that uses the Amazon EC2 launch type for the cloud application containers. Use Amazon ECS Anywhere with an AWS Fargate launch type for the on-premises application containers.

Correct Answer: A, B
Section:

Explanation:
Understanding the Requirement: The company needs a container solution that can scale across on-premises, hybrid, and cloud environments, with a load balancer for HTTP traffic.
Analysis of Options:
Fargate Launch Type and ECS Anywhere: Using Fargate for cloud-based containers and ECS Anywhere for on-premises containers provides a unified management experience across environments without needing to manage infrastructure.
Application Load Balancer: Suitable for HTTP traffic and can distribute requests to the ECS services, ensuring scalability and performance.
Network Load Balancer: Typically used for TCP/UDP traffic, not specifically optimized for HTTP traffic.
EC2 Launch Type for ECS and ECS Anywhere with Fargate: Involves managing infrastructure for EC2 instances, increasing operational overhead.
Best Combination of Solutions:
ECS with Fargate Launch Type and ECS Anywhere: This provides flexibility and scalability across hybrid environments with minimal operational overhead.
Application Load Balancer: Optimized for HTTP traffic, ensuring efficient load distribution and scaling for the ECS services.
Amazon ECS on AWS Fargate
Amazon ECS Anywhere
Application Load Balancer

QUESTION 174
A company is deploying a critical application by using Amazon RDS for MySQL. The application must be highly available and must recover automatically. The company needs to support interactive users (transactional queries) and batch reporting (analytical queries) with no more than a 4-hour lag. The analytical queries must not affect the performance of the transactional queries.

A. Configure Amazon RDS for MySQL in a Multi-AZ DB instance deployment with one standby instance. Point the transactional queries to the primary DB instance. Point the analytical queries to a secondary DB instance that runs in a different Availability Zone.
B. Configure Amazon RDS for MySQL in a Multi-AZ DB cluster deployment with two standby instances. Point the transactional queries to the primary DB instance. Point the analytical queries to the reader endpoint.
C. Configure Amazon RDS for MySQL to use multiple read replicas across multiple Availability Zones. Point the transactional queries to the primary DB instance. Point the analytical queries to one of the replicas in a different Availability Zone.
D. Configure Amazon RDS for MySQL as the primary database for the transactional queries with automated backups enabled. Configure automated backups. Each night, create a read-only database from the most recent snapshot to support the analytical queries. Terminate the previously created database.

Correct Answer: C
Section:
Explanation:
Key Requirements:
High availability and automatic recovery.
Separate transactional and analytical queries with minimal performance impact.
Allow up to a 4-hour lag for analytical queries.
Analysis of Options:
Option A:
Multi-AZ deployments provide high availability but do not include read replicas for separating transactional and analytical queries.
Analytical queries on the secondary DB instance would impact the transactional workload.
Incorrect Approach: Does not meet the requirement of query separation.
Option B:
Multi-AZ DB clusters provide high availability and include a reader endpoint. However, these are better suited for Aurora and not RDS for MySQL.
Incorrect Approach: Not applicable to standard RDS for MySQL.
Option C:
Multiple read replicas allow separation of transactional and analytical workloads.
Queries can be pointed to a replica in a different AZ, ensuring no impact on transactional queries.
Correct Approach: Meets all requirements with high availability and query separation.
Option D:
Creating nightly snapshots and read-only databases adds significant operational overhead and does not support the 4-hour lag requirement.
Incorrect Approach: Not practical for dynamic query separation.

AWS Solution Architect
Reference:
Amazon RDS Read Replicas
Multi-AZ Deployments

QUESTION 175
A financial service company has a two-tier consumer banking application. The frontend serves static web content. The backend consists of APIs. The company needs to migrate the frontend component to AWS. The backend of the application will remain on premises. The company must protect the application from common web vulnerabilities and attacks.
Which solution will meet these requirements with the LEAST operational overhead?

A. Migrate the frontend to Amazon EC2 instances. Deploy an Application Load Balancer (ALB) in front of the instances. Use the instances to invoke the on-premises APIs. Associate AWS WAF rules with the instances.
B. Deploy the frontend as an Amazon CloudFront distribution that has multiple origins. Configure one origin to be an Amazon S3 bucket that serves the static web content. Configure a second origin to route traffic to the on-premises APIs based on the URL pattern. Associate AWS WAF rules with the distribution.
C. Migrate the frontend to Amazon EC2 instances. Deploy a Network Load Balancer (NLB) in front of the instances. Use the instances to invoke the on-premises APIs. Create an AWS Network Firewall instance. Route all traffic through the Network Firewall instance.
D. Deploy the frontend as a static website based on an Amazon S3 bucket. Use an Amazon API Gateway REST API and a set of Amazon EC2 instances to invoke the on-premises APIs. Associate AWS WAF rules with the REST API and the S3 bucket.

Correct Answer: B
Section:
Explanation:
Comprehensive
Deploying the frontend as a CloudFront distribution with multiple origins provides an efficient and scalable solution. Using WAF rules with CloudFront protects against web vulnerabilities, while the multi-origin configuration allows traffic routing to the on-premises backend APIs. This approach minimizes operational overhead compared to managing EC2 instances.
Amazon CloudFront Features
AWS WAF Integration with CloudFront

QUESTION 176
A company is deploying a new application to a VPC on existing Amazon EC2 instances. The application has a presentation tier that uses an Auto Scaling group of EC2 instances. The application also has a database tier that uses an Amazon RDS Multi-AZ database.
The VPC has two public subnets that are split between two Availability Zones. A solutions architect adds one private subnet to each Availability Zone for the RDS database. The solutions architect wants to restrict network access to the RDS database to block access from EC2 instances that do not host the new application.
Which solution will meet this requirement?

A. Modify the RDS database security group to allow traffic from a CIDR range that includes IP addresses of the EC2 instances that host the new application.
B. Associate a new ACL with the private subnets. Deny all incoming traffic from IP addresses that belong to any EC2 instance that does not host the new application.
C. Modify the RDS database security group to allow traffic from the security group that is associated with the EC2 instances that host the new application.
D. Associate a new ACL with the private subnets. Deny all incoming traffic except for traffic from a CIDR range that includes IP addresses of the EC2 instances that host the new application.

Correct Answer: C
Section:
Explanation:
Correct Approach:
AWS Security Groups:
Security groups operate at the instance level, making them the ideal tool for controlling access to specific resources such as an Amazon RDS database.
By default, security groups deny all incoming traffic. You can allow access by explicitly specifying another security group.
Associating an RDS database security group with the EC2 instances' security group ensures only the specified EC2 instances can access the RDS database.

Incorrect Options Analysis:
Option A: Using CIDR blocks for IP-based access is less secure and more difficult to manage. Additionally, Auto Scaling groups dynamically allocate IP addresses, making this approach impractical.
Option B: Network ACLs (NACLs) operate at the subnet level and are stateless. While NACLs can deny or allow traffic, they are not suited to application-specific access control.
Option D: Similar to Option B, using a NACL with CIDR ranges for EC2 IPs is difficult to manage and not application-specific.
Amazon RDS Security Groups
Security Group Best Practices
Differences Between Security Groups and NACLs

QUESTION 177
A company runs analytics software on Amazon EC2 instances The software accepts job requests from users to process data that has been uploaded to Amazon S3 Users report that some submitted data is not being processed Amazon CloudWatch reveals that the EC2 instances have a consistent CPU utilization at or near 100% The company wants to improve system performance and scale the system based on user load.
What should a solutions architect do to meet these requirements?

A. Create a copy of the instance Place all instances behind an Application Load Balancer
B. Create an S3 VPC endpoint for Amazon S3 Update the software to reference the endpoint
C. Stop the EC2 instances. Modify the instance type to one with a more powerful CPU and more memory. Restart the instances.
D. Route incoming requests to Amazon Simple Queue Service (Amazon SQS) Configure an EC2 Auto Scaling group based on queue size Update the software to read from the queue.

Correct Answer: D
Section:
Explanation:
This option is the best solution because it allows the company to decouple the analytics software from the user requests and scale the EC2 instances dynamically based on the demand. By using Amazon SQS, the company can create a queue that stores the user requests and acts as a buffer between the users and the analytics software. This way, the software can process the requests at its own pace without losing any data or overloading the EC2 instances. By using EC2 Auto Scaling, the company can create an Auto Scaling group that launches or terminates EC2 instances automatically based on the size of the queue. This way, the company can ensure that there are enough instances to handle the load and optimize the cost and performance of the system. By updating the software to read from the queue, the company can enable the analytics software to consume the requests from the queue and process the data from Amazon S3.
A) Create a copy of the instance Place all instances behind an Application Load Balancer. This option is not optimal because it does not address the root cause of the problem, which is the high CPU utilization of the EC2 instances. An Application Load Balancer can distribute the incoming traffic across multiple instances, but it cannot scale the instances based on the load or reduce the processing time of the analytics software. Moreover, this option can incur additional costs for the load balancer and the extra instances.
B) Create an S3 VPC endpoint for Amazon S3 Update the software to reference the endpoint. This option is not effective because it does not solve the issue of the high CPU utilization of the EC2 instances. An S3 VPC endpoint can enable the EC2 instances to access Amazon S3 without going through the internet, which can improve the network performance and security. However, it cannot reduce the processing time of the analytics software or scale the instances based on the load.
C) Stop the EC2 instances. Modify the instance type to one with a more powerful CPU and more memory. Restart the instances. This option is not scalable because it does not account for the variability of the user load. Changing the instance type to a more powerful one can improve the performance of the analytics software, but it cannot adjust the number of instances based on the demand. Moreover, this option can increase the cost of the system and cause downtime during the instance modification.
1Using Amazon SQS queues with Amazon EC2 Auto Scaling - Amazon EC2 Auto Scaling
2Tutorial: Set up a scaled and load-balanced application - Amazon EC2 Auto Scaling
3Amazon EC2 Auto Scaling FAQs

QUESTION 178
A company is developing a mobile game that streams score updates to a backend processor and then posts results on a leaderboard A solutions architect needs to design a solution that can handle large traffic spikes process the mobile game updates in order of receipt, and store the processed updates in a highly available database The company also wants to minimize the management overhead required to maintain the solution
What should the solutions architect do to meet these requirements?

A. Push score updates to Amazon Kinesis Data Streams Process the updates in Kinesis Data Streams with AWS Lambda Store the processed updates in Amazon DynamoDB.
B. Push score updates to Amazon Kinesis Data Streams. Process the updates with a fleet of Amazon EC2 instances set up for Auto Scaling Store the processed updates in Amazon Redshift.
C. Push score updates to an Amazon Simple Notification Service (Amazon SNS) topic Subscribe an AWS Lambda function to the SNS topic to process the updates. Store the processed updates in a SQL database running on Amazon EC2.
D. Push score updates to an Amazon Simple Queue Service (Amazon SQS) queue. Use a fleet of Amazon EC2 instances with Auto Scaling to process the updates in the SQS queue. Store the processed updates in an Amazon

RDS Multi-AZ DB instance.

Correct Answer: A
Section:
Explanation:
Amazon Kinesis Data Streams is a scalable and reliable service that can ingest, buffer, and process streaming data in real-time. It can handle large traffic spikes and preserve the order of the incoming data records. AWS Lambda is a serverless compute service that can process the data streams from Kinesis Data Streams without requiring any infrastructure management. It can also scale automatically to match the throughput of the data stream. Amazon DynamoDB is a fully managed, highly available, and fast NoSQL database that can store the processed updates from Lambda. It can also handle high write throughput and provide consistent performance. By using these services, the solutions architect can design a solution that meets the requirements of the company with the least operational overhead.

QUESTION 179
A company has a three-tier environment on AWS that ingests sensor data from its users' devices The traffic flows through a Network Load Balancer (NIB) then to Amazon EC2 instances for the web tier and finally to EC2 instances for the application tier that makes database calls
What should a solutions architect do to improve the security of data in transit to the web tier?

A. Configure a TLS listener and add the server certificate on the NLB
B. Configure AWS Shield Advanced and enable AWS WAF on the NLB
C. Change the load balancer to an Application Load Balancer and attach AWS WAF to it
D. Encrypt the Amazon Elastic Block Store (Amazon EBS) volume on the EC2 instances using AWS Key Management Service (AWS KMS)

Correct Answer: A
Section:
Explanation:
A: How do you protect your data in transit?
Best Practices:
Implement secure key and certificate management: Store encryption keys and certificates securely and rotate them at appropriate time intervals while applying strict access control; for example, by using a certificate management service, such as AWS Certificate Manager (ACM).
Enforce encryption in transit: Enforce your defined encryption requirements based on appropriate standards and recommendations to help you meet your organizational, legal, and compliance requirements.
Automate detection of unintended data access: Use tools such as GuardDuty to automatically detect attempts to move data outside of defined boundaries based on data classification level, for example, to detect a trojan that is copying data to an unknown or untrusted network using the DNS protocol.
Authenticate network communications: Verify the identity of communications by using protocols that support authentication, such as Transport Layer Security (TLS) or IPsec.
https://wa.aws.amazon.com/wat.question.SEC_9.en.html

QUESTION 180
A company maintains about 300 TB in Amazon S3 Standard storage month after month The S3 objects are each typically around 50 GB in size and are frequently replaced with multipart uploads by their global application The number and size of S3 objects remain constant but the company's S3 storage costs are increasing each month.
How should a solutions architect reduce costs in this situation?

A. Switch from multipart uploads to Amazon S3 Transfer Acceleration.
B. Enable an S3 Lifecycle policy that deletes incomplete multipart uploads.
C. Configure S3 inventory to prevent objects from being archived too quickly.
D. Configure Amazon CloudFront to reduce the number of objects stored in Amazon S3.

Correct Answer: B
Section:

Explanation:
This option is the most cost-effective way to reduce the S3 storage costs in this situation. Incomplete multipart uploads are parts of objects that are not completed or aborted by the application. They consume storage space and incur charges until they are deleted. By enabling an S3 Lifecycle policy that deletes incomplete multipart uploads, you can automatically remove them after a specified period of time (such as one day) and free up the storage space. This will reduce the S3 storage costs and also improve the performance of the application by avoiding unnecessary retries or errors.
Option A is not correct because switching from multipart uploads to Amazon S3 Transfer Acceleration will not reduce the S3 storage costs. Amazon S3 Transfer Acceleration is a feature that enables faster data transfers to and from S3 by using the AWS edge network. It is useful for improving the upload speed of large objects over long distances, but it does not affect the storage space or charges. In fact, it may increase the costs by adding a data transfer fee for using the feature.
Option C is not correct because configuring S3 inventory to prevent objects from being archived too quickly will not reduce the S3 storage costs. Amazon S3 Inventory is a feature that provides a report of the objects and their metadata in an S3 bucket. It is useful for managing and auditing the S3 objects, but it does not affect the storage space or charges. In fact, it may increase the costs by generating additional S3 objects for the inventory reports.
Option D is not correct because configuring Amazon CloudFront to reduce the number of objects stored in Amazon S3 will not reduce the S3 storage costs. Amazon CloudFront is a content delivery network (CDN) that distributes the S3 objects to edge locations for faster and lower latency access. It is useful for improving the download speed and availability of the S3 objects, but it does not affect the storage space or charges. In fact, it may increase the costs by adding a data transfer fee for using the service.Reference:
Managing your storage lifecycle
Using multipart upload
Amazon S3 Transfer Acceleration
Amazon S3 Inventory
What Is Amazon CloudFront?

QUESTION 181
A financial company needs to handle highly sensitive data The company will store the data in an Amazon S3 bucket The company needs to ensure that the data is encrypted in transit and at rest The company must manage the encryption keys outside the AWS Cloud
Which solution will meet these requirements?

A. Encrypt the data in the S3 bucket with server-side encryption (SSE) that uses an AWS Key Management Service (AWS KMS) customer managed key
B. Encrypt the data in the S3 bucket with server-side encryption (SSE) that uses an AWS Key Management Service (AWS KMS) AWS managed key
C. Encrypt the data in the S3 bucket with the default server-side encryption (SSE)
D. Encrypt the data at the company's data center before storing the data in the S3 bucket

Correct Answer: D
Section:
Explanation:
This option is the only solution that meets the requirements because it allows the company to encrypt the data with its own encryption keys and tools outside the AWS Cloud. By encrypting the data at the company's data center before storing the data in the S3 bucket, the company can ensure that the data is encrypted in transit and at rest, and that the company has full control over the encryption keys and processes. This option also avoids the need to use any AWS encryption services or features, which may not be compatible with the company's security policies or compliance standards.
A) Encrypt the data in the S3 bucket with server-side encryption (SSE) that uses an AWS Key Management Service (AWS KMS) customer managed key. This option does not meet the requirements because it does not allow the company to manage the encryption keys outside the AWS Cloud. Although the company can create and use its own customer managed key in AWS KMS, the key is still stored and managed by AWS KMS, which is a service within the AWS Cloud. Moreover, the company still needs to use the AWS encryption features and APIs to encrypt and decrypt the data in the S3 bucket, which may not be compatible with the company's security policies or compliance standards.
B) Encrypt the data in the S3 bucket with server-side encryption (SSE) that uses an AWS Key Management Service (AWS KMS) AWS managed key. This option does not meet the requirements because it does not allow the company to manage the encryption keys outside the AWS Cloud. In this option, the company uses the default AWS managed key in AWS KMS, which is created and managed by AWS on behalf of the company. The company has no control over the key rotation, deletion, or recovery policies. Moreover, the company still needs to use the AWS encryption features and APIs to encrypt and decrypt the data in the S3 bucket, which may not be compatible with the company's security policies or compliance standards.
C) Encrypt the data in the S3 bucket with the default server-side encryption (SSE). This option does not meet the requirements because it does not allow the company to manage the encryption keys outside the AWS Cloud. In this option, the company uses the default server-side encryption with Amazon S3 managed keys (SSE-S3), which is applied to every bucket in Amazon S3. The company has no visibility or control over the encryption keys, which are managed by Amazon S3. Moreover, the company still needs to use the AWS encryption features and APIs to encrypt and decrypt the data in the S3 bucket, which may not be compatible with the company's security policies or compliance standards.
1Protecting data with encryption - Amazon Simple Storage Service
2Protecting data with server-side encryption - Amazon Simple Storage Service
3Protecting data by using client-side encryption - Amazon Simple Storage Service

QUESTION 182
A company website hosted on Amazon EC2 instances processes classified data stored in The application writes data to Amazon Elastic Block Store (Amazon EBS) volumes The company needs to ensure that all data that is written to the EBS volumes is encrypted at rest.
Which solution will meet this requirement?

A. Create an 1AM role that specifies EBS encryption Attach the role to the EC2 instances
B. Create the EBS volumes as encrypted volumes Attach the EBS volumes to the EC2 instances
C. Create an EC2 instance tag that has a key of Encrypt and a value of True Tag all instances that require encryption at the EBS level
D. Create an AWS Key Management Service (AWS KMS) key policy that enforces EBS encryption in the account Ensure that the key policy is active

Correct Answer: B
Section:
Explanation:
The simplest and most effective way to ensure that all data that is written to the EBS volumes is encrypted at rest is to create the EBS volumes as encrypted volumes. You can do this by selecting the encryption option when you create a new EBS volume, or by copying an existing unencrypted volume to a new encrypted volume. You can also specify the AWS KMS key that you want to use for encryption, or use the default AWS-managed key. When you attach the encrypted EBS volumes to the EC2 instances, the data will be automatically encrypted and decrypted by the EC2 host. This solution does not require any additional IAM roles, tags, or policies.
Amazon EBS encryption
Creating an encrypted EBS volume
Encrypting an unencrypted EBS volume

QUESTION 183
A company's website hosted on Amazon EC2 instances processes classified data stored in Amazon S3 Due to security concerns, the company requires a pnvate and secure connection between its EC2 resources and Amazon S3.
Which solution meets these requirements?

A. Set up S3 bucket policies to allow access from a VPC endpomt.
B. Set up an 1AM policy to grant read-write access to the S3 bucket.
C. Set up a NAT gateway to access resources outside the private subnet.
D. Set up an access key ID and a secret access key to access the S3 bucket.

Correct Answer: A
Section:
Explanation:
This solution meets the following requirements:
It is private and secure, as it allows the EC2 instances to access the S3 bucket without using the public internet. A VPC endpoint is a gateway that enables you to create a private connection between your VPC and another AWS service, such as S3, within the same Region. A VPC endpoint for S3 provides secure and direct access to S3 buckets and objects using private IP addresses from your VPC. You can also use VPC endpoint policies and S3 bucket policies to control the access to the S3 resources based on the endpoint, the IAM user, the IAM role, or the source IP address.
It is simple and scalable, as it does not require any additional AWS services, gateways, or NAT devices. A VPC endpoint for S3 is a fully managed service that scales automatically with the network traffic. You can create a VPC endpoint for S3 with a few clicks in the VPC console or with a simple API call. You can also use the same VPC endpoint to access multiple S3 buckets in the same Region.
VPC Endpoints - Amazon Virtual Private Cloud
Gateway VPC endpoints - Amazon Virtual Private Cloud
Using Amazon S3 with interface VPC endpoints - Amazon Simple Storage Service
Using Amazon S3 with gateway VPC endpoints - Amazon Simple Storage Service

QUESTION 184
A company has a three-tier environment on AWS that ingests sensor data from its users' devices The traffic flows through a Network Load Balancer (NIB) then to Amazon EC2 instances for the web tier and finally to EC2 instances for the application tier that makes database calls

What should a solutions architect do to improve the security of data in transit to the web tier?

A. Configure a TLS listener and add the server certificate on the NLB
B. Configure AWS Shield Advanced and enable AWS WAF on the NLB
C. Change the load balancer to an Application Load Balancer and attach AWS WAF to it
D. Encrypt the Amazon Elastic Block Store (Amazon EBS) volume on the EC2 instances using AWS Key Management Service (AWS KMS)

Correct Answer: A
Section:
Explanation:
A: How do you protect your data in transit?
Best Practices:
Implement secure key and certificate management: Store encryption keys and certificates securely and rotate them at appropriate time intervals while applying strict access control; for example, by using a certificate management service, such as AWS Certificate Manager (ACM).
Enforce encryption in transit: Enforce your defined encryption requirements based on appropriate standards and recommendations to help you meet your organizational, legal, and compliance requirements.
Automate detection of unintended data access: Use tools such as GuardDuty to automatically detect attempts to move data outside of defined boundaries based on data classification level, for example, to detect a trojan that is copying data to an unknown or untrusted network using the DNS protocol.
Authenticate network communications: Verify the identity of communications by using protocols that support authentication, such as Transport Layer Security (TLS) or IPsec.
https://wa.aws.amazon.com/wat.question.SEC_9.en.html

QUESTION 185
A company wants to use NAT gateways in its AWS environment. The company's Amazon EC2 instances in private subnets must be able to connect to the public internet through the NAT gateways.
Which solution will meet these requirements'?

A. Create public NAT gateways in the same private subnets as the EC2 instances
B. Create private NAT gateways in the same private subnets as the EC2 instances
C. Create public NAT gateways in public subnets in the same VPCs as the EC2 instances
D. Create private NAT gateways in public subnets in the same VPCs as the EC2 instances

Correct Answer: C
Section:
Explanation:
A public NAT gateway enables instances in a private subnet to send outbound traffic to the internet, while preventing the internet from initiating connections with the instances. A public NAT gateway requires an elastic IP address and a route to the internet gateway for the VPC. A private NAT gateway enables instances in a private subnet to connect to other VPCs or on-premises networks through a transit gateway or a virtual private gateway. A private NAT gateway does not require an elastic IP address or an internet gateway. Both private and public NAT gateways map the source private IPv4 address of the instances to the private IPv4 address of the NAT gateway, but in the case of a public NAT gateway, the internet gateway then maps the private IPv4 address of the public NAT gateway to the elastic IP address associated with the NAT gateway. When sending response traffic to the instances, whether it's a public or private NAT gateway, the NAT gateway translates the address back to the original source IP address.
Creating public NAT gateways in the same private subnets as the EC2 instances (option A) is not a valid solution, as the NAT gateways would not have a route to the internet gateway. Creating private NAT gateways in the same private subnets as the EC2 instances (option B) is also not a valid solution, as the instances would not be able to access the internet through the private NAT gateways. Creating private NAT gateways in public subnets in the same VPCs as the EC2 instances (option D) is not a valid solution either, as the internet gateway would drop the traffic from the private NAT gateways.
Therefore, the only valid solution is to create public NAT gateways in public subnets in the same VPCs as the EC2 instances (option C), as this would allow the instances to access the internet through the public NAT gateways and the internet gateway.Reference:
NAT gateways - Amazon Virtual Private Cloud
NAT gateway use cases - Amazon Virtual Private Cloud
Amazon Web Services -- Introduction to NAT Gateways

What is AWS NAT Gateway? - KnowledgeHut

QUESTION 186
Asocial media company has workloads that collect and process data The workloads store the data in on-premises NFS storage The data store cannot scale fast enough to meet the company's expanding business needs The company wants to migrate the current data store to AWS
Which solution will meet these requirements MOST cost-effectively?

A. Set up an AWS Storage Gateway Volume Gateway Use an Amazon S3 Lifecycle policy to transition the data to the appropnate storage class
B. Set up an AWS Storage Gateway Amazon S3 File Gateway Use an Amazon S3 Lifecycle policy to transition the data to the appropriate storage class
C. Use the Amazon Elastic File System (Amazon EFS) Standard-Infrequent Access (Standard-IA) storage class Activate the infrequent access lifecycle policy
D. Use the Amazon Elastic File System (Amazon EFS) One Zone-Infrequent Access (One Zone-IA) storage class Activate the infrequent access lifecycle policy

Correct Answer: B
Section:
Explanation:
This solution meets the requirements most cost-effectively because it enables the company to migrate its on-premises NFS data store to AWS without changing the existing applications or workflows. AWS Storage Gateway is a hybrid cloud storage service that provides seamless and secure integration between on-premises and AWS storage. Amazon S3 File Gateway is a type of AWS Storage Gateway that provides a file interface to Amazon S3, with local caching for low-latency access. By setting up an Amazon S3 File Gateway, the company can store and retrieve files as objects in Amazon S3 using standard file protocols such as NFS. The company can also use an Amazon S3 Lifecycle policy to automatically transition the data to the appropriate storage class based on the frequency of access and the cost of storage. For example, the company can use S3 Standard for frequently accessed data, S3 Standard-Infrequent Access (S3 Standard-IA) or S3 One Zone-Infrequent Access (S3 One Zone-IA) for less frequently accessed data, and S3 Glacier or S3 Glacier Deep Archive for long-term archival data.
Option A is not a valid solution because AWS Storage Gateway Volume Gateway is a type of AWS Storage Gateway that provides a block interface to Amazon S3, with local caching for low-latency access. Volume Gateway is not suitable for migrating an NFS data store, as it requires attaching the volumes to EC2 instances or on-premises servers using the iSCSI protocol. Option C is not a valid solution because Amazon Elastic File System (Amazon EFS) is a fully managed elastic NFS file system that is designed for workloads that require high availability, scalability, and performance. Amazon EFS Standard-Infrequent Access (Standard-IA) is a storage class within Amazon EFS that is optimized for infrequently accessed files, with a lower price per GB and a higher price per access. Using Amazon EFS Standard-IA for migrating an NFS data store would not be cost-effective, as it would incur higher access charges and require additional configuration to enable lifecycle management. Option D is not a valid solution because Amazon EFS One Zone-Infrequent Access (One Zone-IA) is a storage class within Amazon EFS that is optimized for infrequently accessed files that do not require the availability and durability of Amazon EFS Standard or Standard-IA. Amazon EFS One Zone-IA stores data in a single Availability Zone, which reduces the cost by 47% compared to Amazon EFS Standard-IA, but also increases the risk of data loss in the event of an Availability Zone failure. Using Amazon EFS One Zone-IA for migrating an NFS data store would not be cost-effective, as it would incur higher access charges and require additional configuration to enable lifecycle management. It would also compromise the availability and durability of the data.
AWS Storage Gateway - Amazon Web Services
Amazon S3 File Gateway - AWS Storage Gateway
Object Lifecycle Management - Amazon Simple Storage Service
[AWS Storage Gateway Volume Gateway - AWS Storage Gateway]
[Amazon Elastic File System - Amazon Web Services]
[Using EFS storage classes - Amazon Elastic File System]

QUESTION 187
A company has an internal application that runs on Amazon EC2 instances in an Auto Scaling group. The EC2 instances are compute optimized and use Amazon Elastic Block Store (Amazon EBS) volumes.
The company wants to identify cost optimizations across the EC2 instances, the Auto Scaling group, and the EBS volumes.
Which solution will meet these requirements with the MOST operational efficiency?

A. Create a new AWS Cost and Usage Report. Search the report for cost recommendations for the EC2 instances, the Auto Scaling group, and the EBS volumes.
B. Create new Amazon CloudWatch billing alerts. Check the alert statuses for cost recommendations for the EC2 instances, the Auto Scaling group, and the EBS volumes.
C. Configure AWS Compute Optimizer for cost recommendations for the EC2 instances, the Auto Scaling group, and the EBS volumes.
D. Configure AWS Compute Optimizer for cost recommendations for the EC2 instances. Create a new AWS Cost and Usage Report. Search the report for cost recommendations for the Auto Scaling group and the EBS volumes.

Correct Answer: C
Section:
Explanation:

Requirement Analysis: The company wants to identify cost optimizations for EC2 instances, the Auto Scaling group, and EBS volumes with high operational efficiency.
AWS Compute Optimizer: This service provides actionable recommendations to help optimize your AWS resources, including EC2 instances, Auto Scaling groups, and EBS volumes.
Cost Recommendations: Compute Optimizer analyzes the utilization of resources and provides specific recommendations for rightsizing or optimizing the configurations.
Operational Efficiency: Using Compute Optimizer automates the process of identifying cost-saving opportunities, reducing the need for manual analysis.
Implementation:
Enable AWS Compute Optimizer for your AWS account.
Review the recommendations provided for EC2 instances, Auto Scaling groups, and EBS volumes.
Conclusion: This solution provides a comprehensive, automated approach to identifying cost optimizations with minimal operational effort.
Reference
AWS Compute Optimizer: AWS Compute Optimizer Documentation

QUESTION 188
A company uses GPS trackers to document the migration patterns of thousands of sea turtles. The trackers check every 5 minutes to see if a turtle has moved more than 100 yards (91.4 meters). If a turtle has moved, its tracker sends the new coordinates to a web application running on three Amazon EC2 instances that are in multiple Availability Zones in one AWS Region.
Jgpently. the web application was overwhelmed while processing an unexpected volume of tracker data. Data was lost with no way to replay the events. A solutions ftitect must prevent this problem from happening again and needs a solution with the least operational overhead.
at should the solutions architect do to meet these requirements?

A. Create an Amazon S3 bucket to store the data. Configure the application to scan for new data in the bucket for processing.
B. Create an Amazon API Gateway endpoint to handle transmitted location coordinates. Use an AWS Lambda function to process each item concurrently.
C. Create an Amazon Simple Queue Service (Amazon SOS) queue to store the incoming data. Configure the application to poll for new messages for processing.
D. Create an Amazon DynamoDB table to store transmitted location coordinates. Configure the application to query the table for new data for processing. Use TTL to remove data that has been processed.

Correct Answer: C
Section:
Explanation:
Requirement Analysis: The application was overwhelmed with unexpected data volume, leading to data loss and the need for a replay mechanism.
Amazon SQS Overview: SQS is a fully managed message queuing service that decouples and scales microservices, distributed systems, and serverless applications.
Data Decoupling: By using an SQS queue, the application can store incoming tracker data reliably and process it asynchronously, preventing data loss.
Implementation:
Create an SQS queue.
Modify the web application to send incoming data to the SQS queue.
Configure the application instances to poll the SQS queue and process the messages.
Conclusion: This solution meets the requirements with minimal operational overhead, ensuring data is not lost and can be processed at the application's own pace.
Reference
Amazon SQS: Amazon SQS Documentation

QUESTION 189
A solutions architect is creating an application that will handle batch processing of large amounts of data. The input data will be held in Amazon S3 and the ou data will be stored in a different S3 bucket. For processing, the application will transfer the data over the network between multiple Amazon EC2 instances.
What should the solutions architect do to reduce the overall data transfer costs?

A. Place all the EC2 instances in an Auto Scaling group.
B. Place all the EC2 instances in the same AWS Region.
C. Place all the EC2 instances in the same Availability Zone.
D. Place all the EC2 instances in private subnets in multiple Availability Zones.

Correct Answer: C
Section:

Explanation:
Requirement Analysis: The application involves batch processing of large data transfers between EC2 instances.
Data Transfer Costs: Data transfer within the same Availability Zone (AZ) is typically free, while cross-AZ transfers incur additional costs.
Implementation:
Launch all EC2 instances within the same Availability Zone.
Ensure the instances are part of the same subnet to facilitate seamless data transfer.
Conclusion: Placing all EC2 instances in the same AZ reduces data transfer costs significantly without affecting the application's functionality.
Reference
AWS Pricing: AWS Data Transfer Pricing

QUESTION 190
A company hosts an application in a private subnet. The company has already integrated the application with Amazon Cognito. The company uses an Amazon Cognito user pool to authenticate users.
The company needs to modify the application so the application can securely store user documents in an Amazon S3 bucket.
Which combination of steps will securely integrate Amazon S3 with the application? (Select TWO.)

A. Create an Ama2on Cognito identity pool to generate secure Amazon S3 access tokens for users when they successfully log in.
B. Use the existing Amazon Cognito user pool to generate Amazon S3 access tokens for users when they successfully log in.
C. Create an Amazon S3 VPC endpoint in the same VPC where the company hosts the application.
D. Create a NAT gateway in the VPC where the company hosts the application. Assign a policy to the S3 bucket to deny any request that is not initiated from Amazon Cognito.
E. Attach a policy to the S3 bucket that allows access only from the users' IP addresses.

Correct Answer: A, C
Section:
Explanation:
To securely integrate Amazon S3 with an application that uses Amazon Cognito for user authentication, the following two steps are essential:
Detailed Explanation:
Step 1: Create an Amazon Cognito Identity Pool (Option A)
Amazon Cognito Identity Pools allow users to obtain temporary AWS credentials to access AWS resources, such as Amazon S3, after successfully authenticating with the Cognito user pool. The identity pool bridges the gap between user authentication and AWS service access by generating temporary credentials using AWS Identity and Access Management (IAM).
Once a user logs in using the Cognito User Pool, the identity pool provides IAM roles with specific permissions that the application can use to access S3 securely. This ensures that each user has appropriate access controls while accessing the S3 bucket.
This is a secure way to ensure that users only have temporary and least-privilege access to the S3 bucket for their documents.
Step 2: Create an Amazon S3 VPC Endpoint (Option C)
By creating an Amazon S3 VPC endpoint, the company ensures that communication between the application (which is hosted in a private subnet) and the S3 bucket occurs over the AWS private network, without the need to traverse the internet. This enhances security and prevents exposure of data to public networks.
The VPC endpoint allows the application to access the S3 bucket privately and securely within the VPC. It also ensures that traffic stays within the AWS network, reducing attack surface and improving overall security.
Why the Other Options Are Incorrect:
Option B: This is incorrect because Amazon Cognito User Pools are used for user authentication, not for generating S3 access tokens. To provide S3 access, you need to use Amazon Cognito Identity Pools, which offer AWS credentials.
Option D: A NAT gateway is unnecessary in this scenario. Using a VPC endpoint for S3 access provides a more secure and cost-effective solution by keeping traffic within AWS.
Option E: Attaching a policy to restrict access based on IP addresses is not scalable or efficient. It would require managing users' dynamic IP addresses, which is not an effective security measure for this use case.
AWS
Reference:
Amazon Cognito Identity Pools
Amazon VPC Endpoints for S3

QUESTION 191
A company is using AWS DataSync to migrate millions of files from an on-premises system to AWS. The files are 10 KB in size on average.
The company wants to use Amazon S3 for file storage. For the first year after the migration the files will be accessed once or twice and must be immediately available. After 1 year the files must be archived for at least 7 years.

Which solution will meet these requirements MOST cost-effectively?

A. Use an archive tool lo group the files into large objects. Use DataSync to migrate the objects. Store the objects in S3 Glacier Instant Retrieval for the first year. Use a lifecycle configuration to transition the files to S3 Glacier Deep Archive after 1 year with a retention period of 7 years.
B. Use an archive tool to group the files into large objects. Use DataSync to copy the objects to S3 Standard-Infrequent Access (S3 Standard-IA). Use a lifecycle configuration to transition the files to S3 Glacier Instant Retrieval after 1 year with a retention period of 7 years.
C. Configure the destination storage class for the files as S3 Glacier Instant. Retrieval Use a lifecycle policy to transition the files to S3 Glacier Flexible Retrieval after 1 year with a retention period of 7 years.
D. Configure a DataSync task to transfer the files to S3 Standard-Infrequent Access (S3 Standard-IA) Use a lifecycle configuration to transition the files to S3. Deep Archive after 1 year with a retention period of 7 years.

Correct Answer: A
Section:

QUESTION 192
A company sets up an organization in AWS Organizations that contains 10AWS accounts. A solutions architect must design a solution to provide access to the accounts for several thousand employees. The company has an existing identity provider (IdP). The company wants to use the existing IdP for authentication to AWS.
Which solution will meet these requirements?

A. Create 1AM users for the employees in the required AWS accounts. Connect 1AM users to the existing IdP. Configure federated authentication for the 1AM users.
B. Set up AWS account root users with user email addresses and passwords that are synchronized from the existing IdP.
C. Configure AWS 1AM Identity Center Connect 1AM Identity Center to the existing IdP Provision users and groups from the existing IdP
D. Use AWS Resource Access Manager (AWS RAM) to share access to the AWS accounts with the users in the existing IdP.

Correct Answer: C
Section:
Explanation:
AWS IAM Identity Center:
IAM Identity Center provides centralized access management for multiple AWS accounts within an organization and integrates seamlessly with existing identity providers (IdPs) through SAML 2.0 federation.
It allows users to authenticate using their existing IdP credentials and gain access to AWS resources without the need to create and manage separate IAM users in each account.
IAM Identity Center also simplifies provisioning and de-provisioning users, as it can automatically synchronize users and groups from the external IdP to AWS, ensuring secure and managed access.
Integration with Existing IdP:
The solution involves configuring IAM Identity Center to connect to the company's IdP using SAML. This setup allows employees to log in with their existing credentials, reducing the complexity of managing separate AWS credentials.
Once connected, IAM Identity Center handles authentication and authorization, granting users access to the AWS accounts based on their assigned roles and permissions.
Why the Other Options Are Incorrect:
Option A: Creating separate IAM users for each employee is not scalable or efficient. Managing thousands of IAM users across multiple AWS accounts introduces unnecessary complexity and operational overhead.
Option B: Using AWS root users with synchronized passwords is a security risk and goes against AWS best practices. Root accounts should never be used for day-to-day operations.
Option D: AWS Resource Access Manager (RAM) is used for sharing AWS resources between accounts, not for federating access for users across accounts. It doesn't provide a solution for authentication via an external IdP.
AWS
Reference:
AWS IAM Identity Center
SAML 2.0 Integration with AWS IAM Identity Center
By setting up IAM Identity Center and connecting it to the existing IdP, the company can efficiently manage access for thousands of employees across multiple AWS accounts with a high degree of operational efficiency and security. Therefore, Option C is the best solution.

QUESTION 193
A company regularly uploads GB-sized files to Amazon S3. After Ihe company uploads the files, the company uses a fleet of Amazon EC2 Spot Instances to transcode the file format. The company needs to scale throughput when the company uploads data from the on-premises data center to Amazon S3 and when Ihe company downloads data from Amazon S3 to the EC2 instances.
gUkicn solutions will meet these requirements? (Select TWO.)

A. Use the S3 bucket access point instead of accessing the S3 bucket directly.
B. Upload the files into multiple S3 buckets.
C. Use S3 multipart uploads.
D. Fetch multiple byte-ranges of an object in parallel. fe
E. Add a random prefix to each object when uploading the files.

Correct Answer: C, D
Section:
Explanation:
Requirement Analysis: The company needs to scale throughput for uploading large files to S3 and downloading them to EC2 instances.
S3 Multipart Uploads: This method allows for the parallel upload of parts of a file, improving upload efficiency and reliability.
Parallel Fetching: Fetching multiple byte-ranges in parallel from S3 improves download performance.
Implementation:
For uploads, use the S3 multipart upload API to upload files in parallel.
For downloads, use the S3 API to request multiple byte-ranges concurrently.
Conclusion: These solutions effectively scale throughput and improve the performance of both uploads and downloads.
Reference
S3 Multipart Upload: Amazon S3 Multipart Upload
Parallel Fetching: S3 Byte-Range Fetches

QUESTION 194
A company is creating a prototype of an ecommerce website on AWS. The website consists of an Application Load Balancer, an Auto Scaling group of Amazon EC2 instances for web servers, and an Amazon RDS for MySQL DB instance that runs with the Single-AZ configuration.
The website is slow to respond during searches of the product catalog. The product catalog is a group of tables in the MySQL database that the company does not ate frequently. A solutions architect has determined that the CPU utilization on the DB instance is high when product catalog searches occur.
What should the solutions architect recommend to improve the performance of the website during searches of the product catalog?

A. Migrate the product catalog to an Amazon Redshift database. Use the COPY command to load the product catalog tables.
B. Implement an Amazon ElastiCache for Redis cluster to cache the product catalog. Use lazy loading to populate the cache.
C. Add an additional scaling policy to the Auto Scaling group to launch additional EC2 instances when database response is slow.
D. Turn on the Multi-AZ configuration for the DB instance. Configure the EC2 instances to throttle the product catalog queries that are sent to the database.

Correct Answer: B
Section:
Explanation:
Requirement Analysis: The product catalog search is causing high CPU utilization on the MySQL DB instance, slowing down the website.
ElastiCache Overview: Amazon ElastiCache for Redis can be used to cache frequently accessed data, reducing load on the database.
Lazy Loading: This caching strategy loads data into the cache only when it is requested, improving response times for repeated queries.
Implementation:
Set up an ElastiCache for Redis cluster.
Modify the application to check the cache before querying the database.
Use lazy loading to populate the cache on cache misses.
Conclusion: This approach reduces database load and improves website performance during product catalog searches.
Reference
Amazon ElastiCache: ElastiCache Documentation
Caching Strategies: ElastiCache Caching Strategies

QUESTION 195

A media company has a multi-account AWS environment in the us-east-1 Region. The company has an Amazon Simple Notification Service {Amazon SNS) topic in a production account that publishes performance metrics. The company has an AWS Lambda function in an administrator account to process and analyze log data.
The Lambda function that is in the administrator account must be invoked by messages from the SNS topic that is in the production account when significant metrics tM* reported.
Which combination of steps will meet these requirements? (Select TWO.)

A. Create an IAM resource policy for the Lambda function that allows Amazon SNS to invoke the function. Implement an Amazon Simple Queue Service (Amazon SQS) queue in the administrator account to buffer messages from the SNS topic that is in the production account. Configure the SOS queue to invoke the Lambda function.
B. Create an IAM policy for the SNS topic that allows the Lambda function to subscribe to the topic.
C. Use an Amazon EventBridge rule in the production account to capture the SNS topic notifications. Configure the EventBridge rule to forward notifications to the Lambda function that is in the administrator account.
D. Store performance metrics in an Amazon S3 bucket in the production account. Use Amazon Athena to analyze the metrics from the administrator account.

Correct Answer: A, B
Section:
Explanation:
Requirement Analysis: The Lambda function in the administrator account needs to process messages from an SNS topic in the production account.
IAM Policy for SNS Topic: Allows the Lambda function to subscribe and be invoked by the SNS topic.
SQS Queue for Buffering: Using an SQS queue provides reliable message delivery and buffering between SNS and Lambda, ensuring all messages are processed.
Implementation:
Create an SQS queue in the administrator account.
Set an IAM policy to allow the Lambda function to subscribe to and be invoked by the SNS topic.
Configure the SNS topic to send messages to the SQS queue.
Set up the SQS queue to trigger the Lambda function.
Conclusion: This solution ensures reliable message delivery and processing with appropriate permissions.
Reference
Amazon SNS: Amazon SNS Documentation
Amazon SQS: Amazon SQS Documentation
AWS Lambda: AWS Lambda Documentation

QUESTION 196
A company has an application that is running on Amazon EC2 instances A solutions architect has standardized the company on a particular instance family and various instance sizes based on the current needs of the company.
The company wants to maximize cost savings for the application over the next 3 years. The company needs to be able to change the instance family and sizes in the next 6 months based on application popularity and usage
Which solution will meet these requirements MOST cost-effectively?

A. Compute Savings Plan
B. EC2 Instance Savings Plan
C. Zonal Reserved Instances
D. Standard Reserved Instances

Correct Answer: A
Section:
Explanation:
Understanding the Requirement: The company wants to maximize cost savings for their application over the next three years, with the flexibility to change the instance family and sizes within the next six months based on application popularity and usage.
Analysis of Options:
Compute Savings Plan: This plan offers the most flexibility, allowing the company to change instance families, sizes, and regions. It applies to EC2, AWS Fargate, and AWS Lambda, offering significant cost savings with this flexibility.
EC2 Instance Savings Plan: This plan is less flexible than the Compute Savings Plan, as it only applies to EC2 instances and allows changes within a specific instance family.
Zonal Reserved Instances: These provide a discount on EC2 instances but are tied to a specific availability zone and instance type, offering the least flexibility.

Standard Reserved Instances: These offer discounts on EC2 instances but with more restrictions compared to Savings Plans, particularly when changing instance types and families.
Best Option for Flexibility and Savings:
The Compute Savings Plan is the most cost-effective solution because it allows the company to maintain flexibility while still achieving significant cost savings. This is critical for adapting to changing application demands without being locked into specific instance types or families.
AWS Savings Plans
EC2 Instance Types

QUESTION 197
A company runs a critical data analysis job each week before the first day of the work week The job requires at least 1 hour to complete the analysis The job is stateful and cannot tolerate interruptions. The company needs a solution to run the job on AWS.
Which solution will meet these requirements?

A. Create a container for the job. Schedule the job to run as an AWS Fargate task on an Amazon Elastic Container Service (Amazon ECS) cluster by using Amazon EventBridge Scheduler.
B. Configure the job to run in an AWS Lambda function. Create a scheduled rule in Amazon EventBridge to invoke the Lambda function.
C. Configure an Auto Scaling group of Amazon EC2 Spot Instances that run Amazon Linux Configure a crontab entry on the instances to run the analysis.
D. Configure an AWS DataSync task to run the job Configure a cron expression to run the task on a schedule.

Correct Answer: A
Section:
Explanation:
Understanding the Requirement: The job is stateful, cannot tolerate interruptions, and needs to run reliably for at least one hour each week.
Analysis of Options:
AWS Fargate with Amazon ECS and EventBridge: This option provides a serverless compute engine for containers that can run stateful tasks reliably. Using EventBridge Scheduler, the job can be triggered automatically at the specified time without manual intervention.
AWS Lambda with EventBridge: Lambda functions are not suitable for long-running stateful jobs since they have a maximum execution time of 15 minutes.
EC2 Spot Instances: Spot Instances can be interrupted, making them unsuitable for a stateful job that cannot tolerate interruptions.
AWS DataSync: This service is primarily for moving large amounts of data and is not designed to run stateful analysis jobs.
Best Option for Reliable, Scheduled Execution:
The Fargate task on ECS with EventBridge Scheduler meets all requirements, providing the necessary reliability and scheduling capabilities without interruption risks.
Amazon ECS
AWS Fargate
Amazon EventBridge

QUESTION 198
A company runs workloads in the AWS Cloud The company wants to centrally collect security data to assess security across the entire company and to improve workload protection.
Which solution will meet these requirements with the LEAST development effort?

A. Configure a data lake in AWS Lake Formation Use AWS Glue crawlers to ingest the security data into the data lake.
B. Configure an AWS Lambda function to collect the security data in csv format. Upload the data to an Amazon S3 bucket
C. Configure a data lake in Amazon Security Lake to collect the security data Upload the data to an Amazon S3 bucket.
D. Configure an AWS Database Migration Service (AWS DMS) replication instance to load the security data into an Amazon RDS cluster

Correct Answer: C
Section:
Explanation:
Understanding the Requirement: The company wants to centrally collect security data with minimal development effort to assess and improve security across all workloads.
Analysis of Options:
Amazon Security Lake: This is a purpose-built service for centralizing security data from across AWS services and third-party sources into a data lake. It provides native integration and requires minimal development effort to set up.

AWS Lake Formation with AWS Glue: While this can be used to create a data lake, it requires more development effort to set up and configure Glue crawlers for ingestion.
AWS Lambda with S3: This approach involves custom development to collect and process security data before storing it in S3, which requires more effort.
AWS DMS to RDS: AWS Database Migration Service is typically used for database migrations and is not suited for collecting and analyzing security data.
Best Option for Minimal Development Effort:
Amazon Security Lake provides the least development effort for setting up a centralized repository for security data. It simplifies data ingestion and management, making it the most efficient solution for this use case.
Amazon Security Lake
AWS Lake Formation
AWS Glue

QUESTION 199
A company is storing petabytes of data in Amazon S3 Standard The data is stored in multiple S3 buckets and is accessed with varying frequency The company does not know access patterns for all the data. The company needs to implement a solution for each S3 bucket to optimize the cost of S3 usage.
Which solution will meet these requirements with the MOST operational efficiency?

A. Create an S3 Lifecycle configuration with a rule to transition the objects in the S3 bucket to S3 Intelligent-Tiering.
B. Use the S3 storage class analysis tool to determine the correct tier for each object in the S3 bucket. Move each object to the identified storage tier.
C. Create an S3 Lifecycle configuration with a rule to transition the objects in the S3 bucket to S3 Glacier Instant Retrieval.
D. Create an S3 Lifecycle configuration with a rule to transition the objects in the S3 bucket to S3 One Zone-Infrequent Access (S3 One Zone-IA).

Correct Answer: A
Section:
Explanation:
Understanding the Requirement: The company has petabytes of data in S3 Standard across multiple buckets with varying access frequencies. They do not know the access patterns and need a cost-optimized storage solution with minimal operational effort.
Analysis of Options:
S3 Intelligent-Tiering: This storage class automatically moves data between two access tiers (frequent and infrequent) based on changing access patterns. It incurs a small monitoring and automation charge but eliminates the need to manually move data between storage classes.
S3 Storage Class Analysis Tool: While useful for determining access patterns, this tool requires manual intervention to move objects to the appropriate storage class, which increases operational overhead.
S3 Glacier Instant Retrieval: This storage class is designed for data that is rarely accessed but requires instant retrieval when needed. It may not be suitable for data with unknown and varying access patterns.
S3 One Zone-IA: This is a lower-cost option for infrequently accessed data stored in a single availability zone. It does not provide the same level of durability and availability as other options and requires knowledge of access patterns.
Best Option for Operational Efficiency:
S3 Intelligent-Tiering provides the best balance of cost savings and operational efficiency. It dynamically adjusts to access patterns without manual intervention, ensuring the company is only paying for what they need without the risk of incurring high costs for infrequent access data.
Amazon S3 Intelligent-Tiering
Managing your storage lifecycle

QUESTION 200
A company is planning to migrate data to an Amazon S3 bucket The data must be encrypted at rest within the S3 bucket The encryption key must be rotated automatically every year.
Which solution will meet these requirements with the LEAST operational overhead?

A. Migrate the data to the S3 bucket. Use server-side encryption with Amazon S3 managed keys (SSE-S3). Use the built-in key rotation behavior of SSE-S3 encryption keys.
B. Create an AWS Key Management Service (AWS KMS) customer managed key Enable automatic key rotation Set the S3 bucket's default encryption behavior to use the customer managed KMS key. Migrate the data to the S3 bucket.
C. Create an AWS Key Management Service (AWS KMS) customer managed key Set the S3 bucket's default encryption behavior to use the customer managed KMS key. Migrate the data to the S3 bucket. Manually rotate the KMS key every year.
D. Use customer key material to encrypt the data Migrate the data to the S3 bucket. Create an AWS Key Management Service (AWS KMS) key without key material Import the customer key material into the KMS key. Enable automatic key rotation.

Correct Answer: B
Section:
Explanation:
Understanding the Requirement: The data must be encrypted at rest with automatic key rotation every year, with minimal operational overhead.
Analysis of Options:
SSE-S3: This option provides encryption with S3 managed keys and automatic key rotation but offers less control and flexibility compared to KMS keys.
AWS KMS with Customer Managed Key (automatic rotation): This option offers full control over encryption keys, with AWS KMS handling automatic key rotation, minimizing operational overhead.
AWS KMS with Customer Managed Key (manual rotation): This requires manual intervention for key rotation, increasing operational overhead.
Customer Key Material: This involves more complex management, including importing key material and setting up automatic rotation, which increases operational overhead.
Best Option for Minimal Operational Overhead:
AWS KMS with a customer managed key and automatic rotation provides the needed security and key rotation with minimal operational effort. Setting the S3 bucket's default encryption to use this key ensures all data is encrypted as required.
AWS Key Management Service (KMS)
Amazon S3 default encryption

QUESTION 201
A company wants to build a map of its IT infrastructure to identify and enforce policies on resources that pose security risks. The company's security team must be able to query data in the IT infrastructure map and quickly identify security risks.
Which solution will meet these requirements with the LEAST operational overhead?

A. Use Amazon RDS to store the data. Use SQL to query the data to identify security risks.
B. Use Amazon Neptune to store the data. Use SPARQL to query the data to identify security risks.
C. Use Amazon Redshift to store the data. Use SQL to query the data to identify security risks.
D. Use Amazon DynamoDB to store the data. Use PartiQL to query the data to identify security risks.

Correct Answer: B
Section:
Explanation:
Understanding the Requirement: The company needs to map its IT infrastructure to identify and enforce security policies, with the ability to quickly query and identify security risks.
Analysis of Options:
Amazon RDS: While suitable for relational data, it is not optimized for handling complex relationships and querying those relationships, which is essential for an IT infrastructure map.
Amazon Neptune: A graph database service designed for handling highly connected data. It uses SPARQL to query graph data efficiently, making it ideal for mapping IT infrastructure and identifying relationships that pose security risks.
Amazon Redshift: A data warehouse solution optimized for complex queries on large datasets but not specifically for graph data.
Amazon DynamoDB: A NoSQL database that uses PartiQL for querying, but it is not optimized for complex relationships in graph data.
Best Option for Mapping and Querying IT Infrastructure:
Amazon Neptune provides the most suitable solution with the least operational overhead. It is purpose-built for graph data and enables efficient querying of complex relationships to identify security risks.
Amazon Neptune
Querying with SPARQL

QUESTION 202
A company wants to add its existing AWS usage cost to its operation cost dashboard A solutions architect needs to recommend a solution that will give the company access to its usage cost programmatically. The company must be able to access cost data for the current year and forecast costs for the next 12 months.
Which solution will meet these requirements with the LEAST operational overhead?

A. Access usage cost-related data by using the AWS Cost Explorer API with pagination.
B. Access usage cost-related data by using downloadable AWS Cost Explorer report csv files.
C. Configure AWS Budgets actions to send usage cost data to the company through FTP.
D. Create AWS Budgets reports for usage cost data Send the data to the company through SMTP.

Correct Answer: A
Section:
Explanation:
Understanding the Requirement: The company needs programmatic access to its AWS usage costs for the current year and cost forecasts for the next 12 months, with minimal operational overhead.
Analysis of Options:
AWS Cost Explorer API: Provides programmatic access to detailed usage and cost data, including forecast costs. It supports pagination for handling large datasets, making it an efficient solution.
Downloadable AWS Cost Explorer report csv files: While useful, this method requires manual handling of files and does not provide real-time access.
AWS Budgets actions via FTP: This is less suitable as it involves setting up FTP transfers and does not provide the same level of detail and real-time access as the API.
AWS Budgets reports via SMTP: Similar to FTP, this method involves additional setup and lacks the real-time access and detail provided by the API.
Best Option for Minimal Operational Overhead:
AWS Cost Explorer API provides direct, programmatic access to cost data, including detailed usage and forecasting, with minimal setup and operational effort. It is the most efficient solution for integrating cost data into an operational cost dashboard.
AWS Cost Explorer API
AWS Cost and Usage Reports

QUESTION 203
A company wants to create a mobile app that allows users to stream slow-motion video clips on their mobile devices. Currently, the app captures video clips and uploads the video clips in raw format into an Amazon S3 bucket. The app retrieves these video clips directly from the S3 bucket. However, the videos are large in their raw format.
Users are experiencing issues with buffering and playback on mobile devices. The company wants to implement solutions to maximize the performance and scalability of the app while minimizing operational overhead.
Which combination of solutions will meet these requirements? (Select TWO.)

A. Deploy Amazon CloudFront for content delivery and caching
B. Use AWS DataSync to replicate the video files across AWS Regions in other S3 buckets
C. Use Amazon Elastic Transcoder to convert the video files to more appropriate formats.
D. Deploy an Auto Scaling group of Amazon EC2 instances in Local Zones for content delivery and caching
E. Deploy an Auto Scaling group of Amazon EC2 Instances to convert the video files to more appropriate formats.

Correct Answer: A, C
Section:
Explanation:
Understanding the Requirement: The mobile app captures and uploads raw video clips to S3, but users experience buffering and playback issues due to the large size of these videos.
Analysis of Options:
Amazon CloudFront: A content delivery network (CDN) that can cache and deliver content globally with low latency. It helps reduce buffering by delivering content from edge locations closer to the users.
AWS DataSync: Primarily used for data transfer and replication across AWS Regions, which does not directly address the video size and buffering issue.
Amazon Elastic Transcoder: A media transcoding service that can convert raw video files into formats and resolutions more suitable for streaming, reducing the size and improving playback performance.
EC2 Instances in Local Zones: While this could provide content delivery and caching, it involves more operational overhead compared to using CloudFront.
EC2 Instances for Transcoding: Involves setting up and maintaining infrastructure, leading to higher operational overhead compared to using Elastic Transcoder.
Best Combination of Solutions:
Deploy Amazon CloudFront: This optimizes the performance by caching content at edge locations, reducing latency and buffering for users.
Use Amazon Elastic Transcoder: This reduces the file size and converts videos into formats better suited for streaming on mobile devices.
Amazon CloudFront
Amazon Elastic Transcoder

QUESTION 204
A company's application is running on Amazon EC2 instances within an Auto Scaling group behind an Elastic Load Balancing (ELB) load balancer Based on the application's history, the company anticipates a spike in traffic during a holiday each year. A solutions architect must design a strategy to ensure that the Auto Scaling group proactively increases capacity to minimize any performance impact on application users.
Which solution will meet these requirements?

A. Create an Amazon CloudWatch alarm to scale up the EC2 instances when CPU utilization exceeds 90%.

B. Create a recurring scheduled action to scale up the Auto Scaling group before the expected period of peak demand
C. Increase the minimum and maximum number of EC2 instances in the Auto Scaling group during the peak demand period
D. Configure an Amazon Simple Notification Service (Amazon SNS) notification to send alerts when there are autoscaling:EC2_INSTANCE_LAUNCH events.

Correct Answer: B
Section:
Explanation:
Understanding the Requirement: The company anticipates a spike in traffic during a holiday and wants to ensure the Auto Scaling group can handle the increased load without impacting performance.
Analysis of Options:
CloudWatch Alarm: This reacts to spikes based on metrics like CPU utilization but does not proactively scale before the anticipated demand.
Recurring Scheduled Action: This allows the Auto Scaling group to scale up based on a known schedule, ensuring additional capacity is available before the expected spike.
Increase Min/Max Instances: This could result in unnecessary costs by maintaining higher capacity even when not needed.
SNS Notification: Alerts on scaling events but does not proactively manage scaling to prevent performance issues.
Best Solution for Proactive Scaling:
Create a recurring scheduled action: This approach ensures that the Auto Scaling group scales up before the peak demand, providing the necessary capacity proactively without manual intervention.
Scheduled Scaling for Auto Scaling

QUESTION 205
A company is hosting a high-traffic static website on Amazon S3 with an Amazon CloudFront distribution that has a default TTL of 0 seconds The company wants to implement caching to improve performance for the website However, the company also wants to ensure that stale content Is not served for more than a few minutes after a deployment
Which combination of caching methods should a solutions architect implement to meet these requirements? (Select TWO.)

A. Set the CloudFront default TTL to 2 minutes.
B. Set a default TTL of 2 minutes on the S3 bucket
C. Add a Cache-Control private directive to the objects in Amazon S3.
D. Create an AWS Lambda@Edge function to add an Expires header to HTTP responses Configure the function to run on viewer response.
E. Add a Cache-Control max-age directive of 24 hours to the objects in Amazon S3. On deployment, create a CloudFront invalidation to clear any changed files from edge caches

Correct Answer: A, E
Section:
Explanation:
Understanding the Requirement: The company wants to improve caching to enhance website performance while ensuring that stale content is not served for more than a few minutes after a deployment.
Analysis of Options:
Set CloudFront TTL: Setting a short TTL (e.g., 2 minutes) ensures that cached content is refreshed frequently, reducing the risk of serving stale content.
S3 Bucket TTL: This would not control the cache duration for the CloudFront distribution.
Cache-Control Private: This directive is for controlling caching by private caches (e.g., browsers) and is not applicable for CloudFront.
Lambda@Edge: While this can add headers dynamically, it adds complexity and operational overhead.
Cache-Control max-age and CloudFront Invalidation: Setting a longer max-age for objects ensures they are cached longer, reducing load on the origin. Invalidation ensures that updated content is refreshed immediately after deployment.
Best Combination of Caching Methods:
Set the CloudFront default TTL to 2 minutes: This balances caching and freshness of content.
Add a Cache-Control max-age directive of 24 hours and use CloudFront invalidation: This ensures efficient caching while providing a mechanism to clear outdated content immediately after a deployment.
Amazon CloudFront Caching
Invalidating Files in CloudFront

QUESTION 206
A company that uses AWS Organizations runs 150 applications across 30 different AWS accounts The company used AWS Cost and Usage Report to create a new report in the management account The report is delivered to an Amazon S3 bucket that is replicated to a bucket in the data collection account.
The company's senior leadership wants to view a custom dashboard that provides NAT gateway costs each day starting at the beginning of the current month.

Which solution will meet these requirements?

A. Share an Amazon QuickSight dashboard that includes the requested table visual. Configure QuickSight to use AWS DataSync to query the new report
B. Share an Amazon QuickSight dashboard that includes the requested table visual. Configure QuickSight to use Amazon Athena to query the new report.
C. Share an Amazon CloudWatch dashboard that includes the requested table visual Configure CloudWatch to use AWS DataSync to query the new report
D. Share an Amazon CloudWatch dashboard that includes the requested table visual. Configure CloudWatch to use Amazon Athena to query the new report

Correct Answer: B
Section:
Explanation:
Understanding the Requirement: Senior leadership wants a custom dashboard displaying NAT gateway costs daily, starting from the beginning of the current month.
Analysis of Options:
QuickSight with DataSync: While QuickSight is suitable for dashboards, DataSync is not designed for querying and analyzing data reports.
QuickSight with Athena: QuickSight can visualize data queried by Athena, which is designed to analyze data directly from S3.
CloudWatch with DataSync: CloudWatch is primarily for monitoring metrics, not for creating detailed cost analysis dashboards.
CloudWatch with Athena: Similarly, using CloudWatch with Athena does not align well with the requirement for a visual dashboard.
Best Solution for Visualization and Querying:
Amazon QuickSight with Athena: This combination allows for powerful data visualization and querying capabilities. QuickSight can create dynamic dashboards, while Athena efficiently queries the cost and usage report data stored in S3.
Amazon QuickSight
Amazon Athena

QUESTION 207
A company has an application that runs on Amazon EC2 instances in a private subnet The application needs to process sensitive information from an Amazon S3 bucket The application must not use the internet to connect to the S3 bucket.
Which solution will meet these requirements?

A. Configure an internet gateway. Update the S3 bucket policy to allow access from the internet gateway Update the application to use the new internet gateway
B. Configure a VPN connection. Update the S3 bucket policy to allow access from the VPN connection. Update the application to use the new VPN connection.
C. Configure a NAT gateway. Update the S3 bucket policy to allow access from the NAT gateway. Update the application to use the new NAT gateway.
D. Configure a VPC endpoint. Update the S3 bucket policy to allow access from the VPC endpoint. Update the application to use the new VPC endpoint.

Correct Answer: D
Section:
Explanation:
Understanding the Requirement: The application running on EC2 instances in a private subnet needs to process sensitive information from an S3 bucket without using the internet.
Analysis of Options:
Internet Gateway: This would expose the application to the internet, which is not suitable for accessing sensitive information securely.
VPN Connection: VPN is primarily used for secure connections between on-premises networks and AWS VPCs, not for direct S3 access within the same VPC.
NAT Gateway: This allows instances in a private subnet to connect to the internet, but the goal is to avoid internet access.
VPC Endpoint: Provides a private connection between the VPC and S3 without using the internet, ensuring secure access to the S3 bucket.
Best Solution:
VPC Endpoint: Configuring a VPC endpoint allows secure, private communication between the EC2 instances and the S3 bucket without using the internet, ensuring data security and compliance.
Amazon VPC Endpoints
Amazon S3 VPC Endpoint

QUESTION 208
A company uses Amazon EC2 instances and Amazon Elastic Block Store (Amazon EBS) to run its self-managed database The company has 350 TB of data spread across all EBS volumes. The company takes daily EBS snapshots and keeps the snapshots for 1 month. The dally change rate is 5% of the EBS volumes.

Because of new regulations, the company needs to keep the monthly snapshots for 7 years. The company needs to change its backup strategy to comply with the new regulations and to ensure that data is available with minimal administrative effort.
Which solution will meet these requirements MOST cost-effectively?

A. Keep the daily snapshot in the EBS snapshot standard tier for 1 month Copy the monthly snapshot to Amazon S3 Glacier Deep Archive with a 7-year retention period.
B. Continue with the current EBS snapshot policy. Add a new policy to move the monthly snapshot to Amazon EBS Snapshots Archive with a 7-year retention period.
C. Keep the daily snapshot in the EBS snapshot standard tier for 1 month Keep the monthly snapshot in the standard tier for 7 years Use incremental snapshots.
D. Keep the daily snapshot in the EBS snapshot standard tier. Use EBS direct APIs to take snapshots of all the EBS volumes every month. Store the snapshots in an Amazon S3 bucket in the Infrequent Access tier for 7 years.

Correct Answer: B
Section:
Explanation:
Understanding the Requirement: The company needs to keep daily EBS snapshots for 1 month and retain monthly snapshots for 7 years due to new regulations.
Analysis of Options:
S3 Glacier Deep Archive: Moving snapshots to S3 Glacier Deep Archive involves additional complexity and might not be the most straightforward approach for EBS snapshots.
EBS Snapshots Archive: This is a cost-effective solution designed specifically for long-term storage of EBS snapshots.
Standard Tier for 7 Years: Keeping snapshots in the standard tier for 7 years is more expensive and does not optimize costs.
EBS Direct APIs to S3: This involves additional operational overhead and is not the most cost-effective approach compared to using EBS Snapshots Archive.
Best Solution:
EBS Snapshots Archive: Adding a policy to move monthly snapshots to the EBS Snapshots Archive for long-term retention is the most cost-effective and administratively simple solution.
Amazon EBS Snapshots
Amazon EBS Snapshots Archive

QUESTION 209
A company is migrating five on-premises applications to VPCs in the AWS Cloud. Each application is currently deployed in isolated virtual networks on premises and should be deployed similarly in the AWS Cloud. The applications need to reach a shared services VPC. All the applications must be able to communicate with each other.
If the migration is successful, the company will repeat the migration process for more than 100 applications.
Which solution will meet these requirements with the LEAST administrative overhead?

A. Deploy software VPN tunnels between the application VPCs and the shared services VPC. Add routes between the application VPCs in their subnets to the shared services VPC.
B. Deploy VPC peering connections between the application VPCs and the shared services VPC. Add routes between the application VPCs in their subnets to the shared services VPC through the peering connection.
C. Deploy an AWS Direct Connect connection between the application VPCs and the shared services VPC. Add routes from the application VPCs in their subnets to the shared services VPC and the applications VPCs. Add routes from the shared services VPC subnets to the applications VPCs.
D. Deploy a transit gateway with associations between the transit gateway and the application VPCs and the shared services VPC Add routes between the application VPCs in their subnets and the application VPCs to the shared services VPC through the transit gateway.

Correct Answer: D
Section:
Explanation:
Understanding the Requirement: The company needs to migrate applications to AWS, maintaining isolated networks while allowing communication with a shared services VPC and among the applications.
Analysis of Options:
Software VPN Tunnels: This approach involves high administrative overhead and complexity in managing multiple VPN connections.
VPC Peering: While suitable for smaller numbers of VPCs, it becomes complex and hard to manage at scale with over 100 applications.
Direct Connect: Primarily used for high-bandwidth, low-latency connections to on-premises networks, not inter-VPC communication.
Transit Gateway: Simplifies network management by acting as a central hub, allowing easy routing and scalability as more applications are migrated.
Best Solution:
Transit Gateway: This provides a scalable, efficient solution with minimal administrative overhead for managing network connections between multiple VPCs and the shared services VPC.
AWS Transit Gateway
Building a Transit Gateway

QUESTION 210
A company has two AWS accounts: Production and Development. The company needs to push code changes in the Development account to the Production account. In the alpha phase, only two senior developers on the development team need access to the Production account. In the beta phase, more developers will need access to perform testing.
Which solution will meet these requirements?

A. Create two policy documents by using the AWS Management Console in each account. Assign the policy to developers who need access.
B. Create an 1AM role in the Development account Grant the 1AM role access to the Production account. Allow developers to assume the role
C. Create an IAM role in the Production account. Define a trust policy that specifies the Development account Allow developers to assume the role
D. Create an IAM group in the Production account. Add the group as a principal in a trust policy that specifies the Production account. Add developers to the group.

Correct Answer: C
Section:
Explanation:
Understanding the Requirement: Developers in the Development account need to push code changes to the Production account, with phased access control for different stages of the project.
Analysis of Options:
Policy Documents in Each Account: This approach increases complexity and is harder to manage compared to role-based access.
IAM Role in Development Account: Roles in the Development account cannot directly control access to resources in the Production account.
IAM Role in Production Account: Creating a role in the Production account with a trust policy that allows the Development account to assume it provides controlled, secure access.
IAM Group in Production Account: This approach does not provide the required cross-account access control.
Best Solution:
IAM Role in the Production Account: This method allows precise control over who can access the Production account from the Development account, with the ability to manage permissions and access levels effectively.
IAM Roles with Cross-Account Access
Creating a Role for Cross-Account Access

QUESTION 211
A robotics company is designing a solution for medical surgery The robots will use advanced sensors, cameras, and AI algorithms to perceive their environment and to complete surgeries.
The company needs a public load balancer in the AWS Cloud that will ensure seamless communication with backend services. The load balancer must be capable of routing traffic based on the query strings to different target groups. The traffic must also be encrypted
Which solution will meet these requirements?

A. Use a Network Load Balancer with a certificate attached from AWS Certificate Manager (ACM) Use query parameter-based routing
B. Use a Gateway Load Balancer. Import a generated certificate in AWS Identity and Access Management (1AM). Attach the certificate to the load balancer. Use HTTP path-based routing.
C. Use an Application Load Balancer with a certificate attached from AWS Certificate Manager (ACM). Use query parameter-based routing.
D. Use a Network Load Balancer. Import a generated certificate in AWS Identity and Access Management (1AM). Attach the certificate to the load balancer. Use query parameter-based routing.

Correct Answer: C
Section:
Explanation:
Understanding the Requirement: The robotics company needs a public load balancer to ensure seamless communication with backend services, route traffic based on query strings, and encrypt traffic.
Analysis of Options:
Network Load Balancer with ACM Certificate: NLBs primarily operate at the transport layer (Layer 4) and do not natively support query parameter-based routing, which is a Layer 7 feature.
Gateway Load Balancer with IAM Certificate: Gateway Load Balancers are designed for deploying, scaling, and managing third-party virtual appliances and do not support HTTP path-based or query parameter-based routing.
Application Load Balancer with ACM Certificate: ALBs operate at the application layer (Layer 7), supporting features like query parameter-based routing and SSL/TLS termination with ACM certificates.
Network Load Balancer with IAM Certificate: As with the first option, NLBs do not support query parameter-based routing, making it unsuitable for this requirement.
Best Solution:
Application Load Balancer with ACM Certificate: This option provides the necessary Layer 7 routing capabilities and SSL/TLS termination to meet the requirements for query parameter-based routing and encrypted communication.
Application Load Balancer
AWS Certificate Manager

QUESTION 212
A company has multiple VPCs across AWS Regions to support and run workloads that are isolated from workloads in other Regions Because of a recent application launch requirement, the company's VPCs must communicate with all other VPCs across all Regions.
Which solution will meet these requirements with the LEAST amount of administrative effort?

A. Use VPC peering to manage VPC communication in a single Region Use VPC peering across Regions to manage VPC communications.
B. Use AWS Direct Connect gateways across all Regions to connect VPCs across regions and manage VPC communications.
C. Use AWS Transit Gateway to manage VPC communication in a single Region and Transit Gateway peering across Regions to manage VPC communications.
D. Use AWS PrivateLink across all Regions to connect VPCs across Regions and manage VPC communications.

Correct Answer: C
Section:
Explanation:
Understanding the Requirement: The company needs to enable communication between VPCs across multiple AWS Regions with minimal administrative effort.
Analysis of Options:
VPC Peering: Managing multiple VPC peering connections across regions is complex and difficult to scale, leading to significant administrative overhead.
AWS Direct Connect Gateways: Primarily used for creating private connections between AWS and on-premises environments, not for inter-VPC communication across regions.
AWS Transit Gateway: Simplifies VPC interconnections within a region and supports Transit Gateway peering for cross-region connectivity, reducing administrative complexity.
AWS PrivateLink: Used for accessing AWS services and third-party services over a private connection, not for inter-VPC communication.
Best Solution:
AWS Transit Gateway with Transit Gateway Peering: This option provides a scalable and efficient solution for managing VPC communications both within a single region and across multiple regions with minimal administrative overhead.
AWS Transit Gateway
Transit Gateway Peering

QUESTION 213
A developer is creating a serverless application that performs video encoding. The encoding process runs as background jobs and takes several minutes to encode each video. The process must not send an immediate result to users.
The developer is using Amazon API Gateway to manage an API for the application. The developer needs to run test invocations and request validations. The developer must distribute API keys to control access to the API.
Which solution will meet these requirements?

A. Create an HTTP API. Create an AWS Lambda function to handle the encoding jobs. Integrate the function with the HTTP API. Use the Event invocation type to call the Lambda function.
B. Create a REST API with the default endpoint type. Create an AWS Lambda function to handle the encoding jobs. Integrate the function with the REST API. Use the Event invocation type to call the Lambda function.
C. Create an HTTP API. Create an AWS Lambda function to handle the encoding jobs. Integrate the function with the HTTP API. Use the RequestResponse invocation type to call the Lambda function.
D. Create a REST API with the default endpoint type. Create an AWS Lambda function to handle the encoding jobs. Integrate the function with the REST API. Use the RequestResponse invocation type to call the Lambda function.

Correct Answer: B
Section:
Explanation:
Background Jobs with Event Invocation Type:
The Event invocation type is asynchronous, meaning the Lambda function does not send an immediate result to the API Gateway and processes the request in the background. This is ideal for video encoding tasks that take time.
REST API vs. HTTP API:
REST APIs support advanced features like API keys, request validation, and throttling that HTTP APIs do not support fully.
Since the developer needs API keys and request validations, a REST API is the correct choice.
Integration with Lambda:
AWS Lambda integration is seamless with REST APIs, and using the Event invocation ensures asynchronous processing.
Incorrect Options Analysis:

Option A: HTTP API lacks full support for API keys and validation.
Option C and D: RequestResponse invocation type requires immediate responses, unsuitable for background jobs.
AWS Lambda Invocation Types
Amazon API Gateway REST APIs

QUESTION 214
A developer needs to export the contents of several Amazon DynamoDB tables into Amazon S3 buckets to comply with company data regulations. The developer uses the AWS CLI to run commands to export from each table to the proper S3 bucket. The developer sets up AWS credentials correctly and grants resources appropriate permissions. However, the exports of some tables fail.
What should the developer do to resolve this issue?

A. Ensure that point-in-time recovery is enabled on the DynamoDB tables.
B. Ensure that the target S3 bucket is in the same AWS Region as the DynamoDB table.
C. Ensure that DynamoDB streaming is enabled for the tables.
D. Ensure that DynamoDB Accelerator (DAX) is enabled.

Correct Answer: A
Section:
Explanation:
Export Requirements:
To export data from DynamoDB to Amazon S3, point-in-time recovery (PITR) must be enabled for the tables. This feature creates a snapshot of the data, which is essential for exports.
Incorrect Options Analysis:
Option B: S3 buckets and DynamoDB tables do not need to be in the same region for exports.
Option C: DynamoDB streams are unrelated to the export functionality.
Option D: DAX accelerates reads but has no role in exports.
Exporting DynamoDB Data to Amazon S3