Number: SAA-C03
Passing Score: 800
Time Limit: 120
File Version: 51.0

Exam Code: SAA-C03

Exam Name: AWS Certified Solutions Architect – Associate

Exam A

QUESTION 1
A company has hired an external vendor to perform work in the company's AWS account. The vendor uses an automated tool that is hosted in an AWS account that the vendor owns. The vendor does not have IAM access to the company's AWS account.
How should a solutions architect grant this access to the vendor?

A. Create an IAM role in the company's account to delegate access to the vendor's IAM role. Attach the appropriate IAM policies to the role for the permissions that the vendor requires.
B. Create an IAM user in the company's account with a password that meets the password complexity requirements. Attach the appropriate IAM policies to the user for the permissions that the vendor requires.
C. Create an IAM group in the company's account. Add the tool's IAM user from the vendor account to the group. Attach the appropriate IAM policies to the group for the permissions that the vendor requires.
D. Create a new identity provider by choosing "AWS account" as the provider type in the IAM console. Supply the vendor's AWS account ID and user name. Attach the appropriate IAM policies to the new provider for the permissions that the vendor requires.

Correct Answer: A
Section:
Explanation:
https://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_common-scenarios_third-party.html

QUESTION 2
A company has deployed a Java Spring Boot application as a pod that runs on Amazon Elastic Kubernetes Service (Amazon EKS) in private subnets. The application needs to write data to an Amazon DynamoDB table. A solutions architect must ensure that the application can interact with the DynamoDB table without exposing traffic to the internet. Which combination of steps should the solutions architect take to accomplish this goal? (Choose two.)

A. Attach an IAM role that has sufficient privileges to the EKS pod.
B. Attach an IAM user that has sufficient privileges to the EKS pod.
C. Allow outbound connectivity to the DynamoDB table through the private subnets' network ACLs.
D. Create a VPC endpoint for DynamoDB.
E. Embed the access keys in the Java Spring Boot code.

Correct Answer: A, D
Section:
Explanation:
https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/vpc-endpointsdynamodb.html
https://aws.amazon.com/about-aws/whats-new/2019/09/amazon-eks-adds-support-to-assign-iampermissions-to-kubernetes-service-accounts/

QUESTION 3
A company recently migrated its web application to AWS by rehosting the application on Amazon EC2 instances in a single AWS Region. The company wants to redesign its application architecture to be highly available and fault tolerant. Traffic must reach all running EC2 instances randomly.
Which combination of steps should the company take to meet these requirements? (Choose two.)

A. Create an Amazon Route 53 failover routing policy.
B. Create an Amazon Route 53 weighted routing policy.
C. Create an Amazon Route 53 multivalue answer routing policy.
D. Launch three EC2 instances: two instances in one Availability Zone and one instance in another Availability Zone.
E. Launch four EC2 instances: two instances in one Availability Zone and two instances in another Availability Zone.

Correct Answer: C, E
Section:
Explanation:
https://aws.amazon.com/premiumsupport/knowledge-center/multivalue-versus-simple-policies/

QUESTION 4
A company is building a three-tier application on AWS. The presentation tier will serve a static website. The logic tier is a containerized application. This application will store data in a relational database. The company wants to simplify deployment and to reduce operational costs.
Which solution will meet these requirements?

A. Use Amazon S3 to host static content. Use Amazon Elastic Container Service (Amazon ECS) with AWS Fargate for compute power. Use a managed Amazon RDS cluster for the database.
B. Use Amazon CloudFront to host static content. Use Amazon Elastic Container Service (Amazon ECS) with Amazon EC2 for compute power. Use a managed Amazon RDS cluster for the database.
C. Use Amazon S3 to host static content. Use Amazon Elastic Kubernetes Service (Amazon EKS) with AWS Fargate for compute power. Use a managed Amazon RDS cluster for the database.
D. Use Amazon EC2 Reserved Instances to host static content. Use Amazon Elastic Kubernetes Service (Amazon EKS) with Amazon EC2 for compute power. Use a managed Amazon RDS cluster for the database.

Correct Answer: A
Section:
Explanation:
Amazon S3 is an object storage service that offers industry-leading scalability, data availability, security, and performance. You can use Amazon S3 to host static content for your website, such as HTML files, images, videos, etc. Amazon Elastic Container Service (Amazon ECS) is a fully managed container orchestration service that allows you to run and scale containerized applications on AWS. AWS Fargate is a serverless compute engine for containers that works with both Amazon ECS and Amazon EKS. Fargate makes it easy for you to focus on building your applications by removing the need to provision and manage servers. You can use Amazon ECS with AWS Fargate for compute power for your containerized application logic tier. Amazon RDS is a managed relational database service that makes it easy to set up, operate, and scale a relational database in the cloud. You can use a managed Amazon RDS cluster for the database tier of your application. This solution will simplify deployment and reduce operational costs for your three-tier application.
Reference: https://docs.aws.amazon.com/AmazonS3/latest/userguide/WebsiteHosting.html https://docs.aws.amazon.com/AmazonECS/latest/developerguide/Welcome.html
https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Welcome.html

QUESTION 5
A company has a workload in an AWS Region. Customers connect to and access the workload by using an Amazon API Gateway REST API. The company uses Amazon Route 53 as its DNS provider. The company wants to provide individual and secure URLs for all customers.
Which combination of steps will meet these requirements with the MOST operational efficiency? (Select THREE.)

A. Register the required domain in a registrar. Create a wildcard custom domain name in a Route 53 hosted zone and record in the zone that points to the API Gateway endpoint.
B. Request a wildcard certificate that matches the domains in AWS Certificate Manager (ACM) in a different Region.
C. Create hosted zones for each customer as required in Route 53. Create zone records that point to the API Gateway endpoint.
D. Request a wildcard certificate that matches the custom domain name in AWS Certificate Manager (ACM) in the same Region.
E. Create multiple API endpoints for each customer in API Gateway.
F. Create a custom domain name in API Gateway for the REST API. Import the certificate from AWS Certificate Manager (ACM).

Correct Answer: A, D, F
Section:
Explanation:
To provide individual and secure URLs for all customers using an API Gateway REST API, you need to do the following steps:
a) Register the required domain in a registrar. Create a wildcard custom domain name in a Route 53 hosted zone and record in the zone that points to the API Gateway endpoint. This step will allow you to use a custom domain name for your API instead of the default one generated by API Gateway. A wildcard custom domain name means that you can use any subdomain under your domain name (such as customer1.example.com or customer2.example.com) to access your API. You need to register your domain name with a registrar (such as Route 53 or a third-party registrar) and create a hosted zone in Route 53 for your domain name. You also need to create a record in the hosted zone that points to the API Gateway endpoint using an alias record.
d) Request a wildcard certificate that matches the custom domain name in AWS Certificate Manager (ACM) in the same Region. This step will allow you to secure your API with HTTPS using a certificate issued by ACM. A wildcard certificate means that it can match any subdomain under your domain name (such as *.example.com). You need to request or import a certificate in ACM that matches your custom domain name and verify that you

own the domain name. You also need to request the certificate in the same Region as your API.
f) Create a custom domain name in API Gateway for the REST API. Import the certificate from AWS Certificate Manager (ACM). This step will allow you to associate your custom domain name with your API and use the certificate from ACM to enable HTTPS. You need to create a custom domain name in API Gateway for the REST API and specify the certificate ARN from ACM. You also need to create a base path mapping that maps a path from your custom domain name to your API stage.

QUESTION 6
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 7
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 8
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 9
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 10
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 11
A media company collects and analyzes user activity data on premises. The company wants to migrate this capability to AWS. The user activity data store will continue to grow and will be petabytes in size. The company needs to build a highly available data ingestion solution that facilitates on-demand analytics of existing data and new data with SQL. Which solution will meet these requirements with the LEAST operational overhead?

A. Send activity data to an Amazon Kinesis data stream. Configure the stream to deliver the data to an Amazon S3 bucket.
B. Send activity data to an Amazon Kinesis Data Firehose delivery stream. Configure the stream to deliver the data to an Amazon Redshift cluster.
C. Place activity data in an Amazon S3 bucket. Configure Amazon S3 to run an AWS Lambda function on the data as the data arrives in the S3 bucket.
D. Create an ingestion service on Amazon EC2 instances that are spread across multiple Availability Zones. Configure the service to forward data to an Amazon RDS Multi-AZ database.

Correct Answer: B
Section:
Explanation:
Amazon Redshift is a fully managed, petabyte-scale data warehouse service in the cloud. You can start with just a few hundred gigabytes of data and scale to a petabyte or more. This allows you to use your data to gain new insights for your business and customers. The first step to create a data warehouse is to launch a set of nodes, called an Amazon Redshift cluster. After you provision your cluster, you can upload your data set and then perform data analysis queries. Regardless of the size of the data set, Amazon Redshift offers fast query performance using the same SQL-based tools and business intelligence applications that you use today.

QUESTION 12
A company has a serverless website with millions of objects in an Amazon S3 bucket. The company uses the S3 bucket as the origin for an Amazon CloudFront distribution. The company did not set encryption on the S3 bucket before the objects were loaded. A solutions architect needs to enable encryption for all existing objects and for all objects that are added to the S3 bucket in the future. Which solution will meet these requirements with the LEAST amount of effort?

A. Create a new S3 bucket. Turn on the default encryption settings for the new S3 bucket. Download all existing objects to temporary local storage. Upload the objects to the new S3 bucket.
B. Turn on the default encryption settings for the S3 bucket. Use the S3 Inventory feature to create a .csv file that lists the unencrypted objects. Run an S3 Batch Operations job that uses the copy command to encrypt those objects.
C. Create a new encryption key by using AWS Key Management Service (AWS KMS). Change the settings on the S3 bucket to use server-side encryption with AWS KMS managed encryption keys (SSE-KMS). Turn on versioning for the S3 bucket.
D. Navigate to Amazon S3 in the AWS Management Console. Browse the S3 bucket's objects. Sort by the encryption field. Select each unencrypted object. Use the Modify button to apply default encryption settings to every unencrypted object in the S3 bucket.

Correct Answer: B
Section:
Explanation:
https://spin.atomicobject.com/2020/09/15/aws-s3-encrypt-existing-objects/

QUESTION 13
A solutions architect is designing a new API using Amazon API Gateway that will receive requests from users. The volume of requests is highly variable; several hours can pass without receiving a single request. The data processing will take place asynchronously, but should be completed within a few seconds after a request is made.
Which compute service should the solutions architect have the API invoke to deliver the requirements at the lowest cost?

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

D. A containerized service hosted in Amazon ECS with Amazon EC2

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

QUESTION 14
A company hosts multiple production applications. One of the applications consists of resources from Amazon EC2, AWS Lambda, Amazon RDS, Amazon Simple Notification Service (Amazon SNS), and Amazon Simple Queue Service (Amazon SQS) across multiple AWS Regions. All company resources are tagged with a tag name of "application" and a value that corresponds to each application. A solutions architect must provide the quickest solution for identifying all of the tagged components.
Which solution meets these requirements?

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

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

QUESTION 15
A company has a data ingestion workflow that consists the following:
An Amazon Simple Notification Service (Amazon SNS) topic for notifications about new data deliveries An AWS Lambda function to process the data and record metadata The company observes that the ingestion workflow fails occasionally because of network connectivity issues. When such a failure occurs, the Lambda function does not ingest the corresponding data unless the company manually reruns the job. Which combination of actions should a solutions architect take to ensure that the Lambda function ingests all data in the future? (Select TWO.)

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

Correct Answer: B, E
Section:
Explanation:
To ensure that the Lambda function ingests all data in the future despite occasional network connectivity issues, the following actions should be taken: Create an Amazon Simple Queue Service (SQS) queue and subscribe it to the SNS topic. This allows

QUESTION 16
A company has an application that provides marketing services to stores. The services are based on previous purchases by store customers. The stores upload transaction data to the company through SFTP, and the data is processed and analyzed to generate new marketing offers. Some of the files can exceed 200 GB in size.
Recently, the company discovered that some of the stores have uploaded files that contain personally identifiable information (PII) that should not have been included. The company wants administrators to be alerted if PII is shared again. The company also wants to automate remediation.
What should a solutions architect do to meet these requirements with the LEAST development effort?

A. Use an Amazon S3 bucket as a secure transfer point. Use Amazon Inspector to scan me objects in the bucket. If objects contain Pll. trigger an S3 Lifecycle policy to remove the objects that contain Pll.
B. Use an Amazon S3 bucket as a secure transfer point. Use Amazon Macie to scan the objects in the bucket. If objects contain Pll. Use Amazon Simple Notification Service (Amazon SNS) to trigger a notification to the administrators to remove the objects mat contain Pll.
C. Implement custom scanning algorithms in an AWS Lambda function. Trigger the function when objects are loaded into the bucket. It objects contain Rll. use Amazon Simple Notification Service (Amazon SNS) to trigger a notification to the administrators to remove the objects that contain Pll.
D. Implement custom scanning algorithms in an AWS Lambda function. Trigger the function when objects are loaded into the bucket. If objects contain Pll. use Amazon Simple Email Service (Amazon STS) to trigger a notification to the administrators and trigger on S3 Lifecycle policy to remove the objects mot contain PII.

Correct Answer: A
Section:
Explanation:

QUESTION 17
A company needs guaranteed Amazon EC2 capacity in three specific Availability Zones in a specific AWS Region for an upcoming event that will last 1 week. What should the company do to guarantee the EC2 capacity?

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

Correct Answer: D
Section:
Explanation:
https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-capacity-reservations.html
Reserve instances: You will have to pay for the whole term (1 year or 3years) which is not cost effective

QUESTION 18
A company's website uses an Amazon EC2 instance store for its catalog of items. The company wants to make sure that the catalog is highly available and that the catalog is stored in a durable location. What should a solutions architect do to meet these requirements?

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

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

QUESTION 19
A company stores call transcript files on a monthly basis. Users access the files randomly within 1 year of the call, but users access the files infrequently after 1 year. The company wants to optimize its solution by giving users the ability to query and retrieve files that are less than 1-year-old as quickly as possible. A delay in retrieving older files is acceptable. Which solution will meet these requirements MOST cost-effectively?

A. Store individual files with tags in Amazon S3 Glacier Instant Retrieval. Query the tags to retrieve the files from S3 Glacier Instant Retrieval.

B. Store individual files in Amazon S3 Intelligent-Tiering. Use S3 Lifecycle policies to move the files to S3 Glacier Flexible Retrieval after 1 year. Query and retrieve the files that are in Amazon S3 by using Amazon Athena. Query and retrieve the files that are in S3 Glacier by using S3 Glacier Select.
C. Store individual files with tags in Amazon S3 Standard storage. Store search metadata for each archive in Amazon S3 Standard storage. Use S3 Lifecycle policies to move the files to S3 Glacier Instant Retrieval after 1 year. Query and retrieve the files by searching for metadata from Amazon S3.
D. Store individual files in Amazon S3 Standard storage. Use S3 Lifecycle policies to move the files to S3 Glacier Deep Archive after 1 year. Store search metadata in Amazon RDS. Query the files from Amazon RDS. Retrieve the files from S3 Glacier Deep Archive.

Correct Answer: B
Section:
Explanation:
"For archive data that needs immediate access, such as medical images, news media assets, or genomics data, choose the S3 Glacier Instant Retrieval storage class, an archive storage class that delivers the lowest cost storage with milliseconds retrieval. For archive data that does not require immediate access but needs the flexibility to retrieve large sets of data at no cost, such as backup or disaster recovery use cases, choose S3 Glacier Flexible Retrieval (formerly S3 Glacier), with retrieval in minutes or free bulk retrievals in 5-12 hours." https://aws.amazon.com/about-aws/whats-new/2021/11/amazon-s3-glacier-instant-retrieval-storage-class/

QUESTION 20
A company has a production workload that runs on 1,000 Amazon EC2 Linux instances. The workload is powered by third-party software. The company needs to patch the third-party software on all EC2 instances as quickly as possible to remediate a critical security vulnerability.
What should a solutions architect do to meet these requirements?

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

Correct Answer: D
Section:
Explanation:

QUESTION 21
A company is developing an application that provides order shipping statistics for retrieval by a REST API. The company wants to extract the shipping statistics, organize the data into an easy-to-read HTML format, and send the report to several email addresses at the same time every morning.
Which combination of steps should a solutions architect take to meet these requirements? (Choose two.)

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

Correct Answer: B, D
Section:
Explanation:
https://docs.aws.amazon.com/ses/latest/dg/send-email-formatted.html D. Create an Amazon EventBridge (Amazon CloudWatch Events) scheduled event that invokes an AWS Lambda function to query the application's API for the data. This step can be done using AWS Lambda to extract the shipping statistics and organize the data into an HTML format.B. Use Amazon Simple Email Service (Amazon SES) to format the data and send the report by email.This step can be done by using Amazon SES to send the report to multiple email addresses at the same time every morning.
Therefore, options D and B are the correct choices for this question. Option A is incorrect because Kinesis Data Firehose is not necessary for this use case. Option C is incorrect because AWS Glue is not required to query the application's API. Option E is incorrect because S3 event notifications cannot be used to send the report by email.

QUESTION 22
A company wants to migrate its on-premises application to AWS. The application produces output files that vary in size from tens of gigabytes to hundreds of terabytes The application data must be stored in a standard file system structure The company wants a solution that scales automatically, is highly available, and requires minimum operational overhead. Which solution will meet these requirements?

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

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

QUESTION 23
A company needs to store its accounting records in Amazon S3. The records must be immediately accessible for 1 year and then must be archived for an additional 9 years. No one at the company, including administrative users and root users, can be able to delete the records during the entire 10- year period. The records must be stored with maximum resiliency. Which solution will meet these requirements?

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

Correct Answer: C
Section:
Explanation:
To meet the requirements of immediately accessible records for 1 year and then archived for an additional 9 years with maximum resiliency, we can use S3 Lifecycle policy to transition records from S3 Standard to S3 Glacier Deep Archive after 1 year. And to ensure that the records cannot be deleted by anyone, including administrative and root users, we can use S3 Object Lock in compliance mode for a period of 10 years. Therefore, the correct answer is option C.Reference: https://docs.aws.amazon.com/AmazonS3/latest/userguide/object-lock.htmls

QUESTION 24
A company runs multiple Windows workloads on AWS. The company's employees use Windows file shares that are hosted on two Amazon EC2 instances. The file shares synchronize data between themselves and maintain duplicate copies. The company wants a highly available and durable storage solution that preserves how users currently access the files. What should a solutions architect do to meet these requirements?

A. Migrate all the data to Amazon S3 Set up IAM authentication for users to access files
B. Set up an Amazon S3 File Gateway. Mount the S3 File Gateway on the existing EC2 Instances.
C. Extend the file share environment to Amazon FSx for Windows File Server with a Multi-AZ configuration. Migrate all the data to FSx for Windows File Server.
D. Extend the file share environment to Amazon Elastic File System (Amazon EFS) with a Multi-AZ configuration. Migrate all the data to Amazon EFS.

Correct Answer: C
Section:
Explanation:
https://docs.aws.amazon.com/AWSEC2/latest/WindowsGuide/AmazonEFS.html Amazon FSx for Windows File Server provides fully managed Microsoft Windows file servers, backed by a fully nativeWindows file system.
https://docs.aws.amazon.com/fsx/latest/WindowsGuide/what-is.html

QUESTION 25
A solutions architect is developing a multiple-subnet VPC architecture. The solution will consist of six subnets in two Availability Zones. The subnets are defined as public, private and dedicated for databases. Only the Amazon EC2 instances running in the private subnets should be able to access a database.

Which solution meets these requirements?

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

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

QUESTION 26
A company has registered its domain name with Amazon Route 53. The company uses Amazon API Gateway in the ca-central-1 Region as a public interface for its backend microservice APIs. Thirdparty services consume the APIs securely. The company wants to design its API Gateway URL with the company's domain name and corresponding certificate so that the third-party services can use HTTPS. Which solution will meet these requirements?

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

Correct Answer: C
Section:
Explanation:
To design the API Gateway URL with the company's domain name and corresponding certificate, the company needs to do the following: 1. Create a Regional API Gateway endpoint: This will allow the company to create an endpoint that is specific to a region. 2. Associate the API Gateway endpoint with the company's domain name: This will allow the company to use its own domain name for the API Gateway URL. 3. Import the public certificate associated with the company's domain name into AWS Certificate Manager (ACM) in the same Region: This will allow the company to use HTTPS for secure communication with its APIs. 4. Attach the certificate to the API Gateway endpoint: This will allow the company to use the certificate for securing the API Gateway URL. 5. Configure Route 53 to route traffic to the API Gateway endpoint: This will allow the company to use Route 53 to route traffic to the API Gateway URL using the company's domain name.

QUESTION 27
A company is running a popular social media website. The website gives users the ability to upload images to share with other users. The company wants to make sure that the images do not contain inappropriate content. The company needs a solution that minimizes development effort.
What should a solutions architect do to meet these requirements?

A. Use Amazon Comprehend to detect inappropriate content. Use human review for low-confidence predictions.
B. Use Amazon Rekognition to detect inappropriate content. Use human review for low-confidence predictions.
C. Use Amazon SageMaker to detect inappropriate content. Use ground truth to label low-confidence predictions.

D. Use AWS Fargate to deploy a custom machine learning model to detect inappropriate content.
Use ground truth to label low-confidence predictions.

Correct Answer: B
Section:
Explanation:
https://docs.aws.amazon.com/rekognition/latest/dg/moderation.html?pg=ln&sec=fthttps://docs.aws.amazon.com/rekognition/latest/dg/a2i-rekognition.html

QUESTION 28
A company wants to run its critical applications in containers to meet requirements tor scalability and availability The company prefers to focus on maintenance of the critical applications The company does not want to be responsible for provisioning and managing the underlying infrastructure that runs the containerized workload What should a solutions architect do to meet those requirements?

A. Use Amazon EC2 Instances, and Install Docker on the Instances
B. Use Amazon Elastic Container Service (Amazon ECS) on Amazon EC2 worker nodes
C. Use Amazon Elastic Container Service (Amazon ECS) on AWS Fargate
D. Use Amazon EC2 instances from an Amazon Elastic Container Service (Amazon ECS)-op6mized Amazon Machine Image (AMI).

Correct Answer: C
Section:
Explanation:
Explanation: using AWS ECS on AWS Fargate since they requirements are for scalability and availability without having to provision and manage the underlying infrastructure to run the containerized workload.
https://docs.aws.amazon.com/AmazonECS/latest/userguide/what-is-fargate.html

QUESTION 29
A company hosts more than 300 global websites and applications. The company requires a platform to analyze more than 30 TB of clickstream data each day. What should a solutions architect do to transmit and process the clickstream data?

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

Correct Answer: D
Section:
Explanation:
https://aws.amazon.com/es/blogs/big-data/real-time-analytics-with-amazon-redshift-streaming- ingestion/

QUESTION 30
A company has a website hosted on AWS The website is behind an Application Load Balancer (ALB) that is configured to handle HTTP and HTTPS separately. The company wants to forward all requests to the website so that the requests will use HTTPS.
What should a solutions architect do to meet this requirement?

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

Correct Answer: C

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

QUESTION 31
A company is developing a two-tier web application on AWS. The company's developers have deployed the application on an Amazon EC2 instance that connects directly to a backend Amazon RDS database. The company must not hardcode database credentials in the application. The company must also implement a solution to automatically rotate the database credentials on a regular basis. Which solution will meet these requirements with the LEAST operational overhead?

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

Correct Answer: C
Section:
Explanation:
https://docs.aws.amazon.com/secretsmanager/latest/userguide/create_database_secret.html

QUESTION 32
A company is deploying a new public web application to AWS. The application will run behind an Application Load Balancer (ALB). The application needs to be encrypted at the edge with an SSL/TLS certificate that is issued by an external certificate authority (CA). The certificate must be rotated each year before the certificate expires. What should a solutions architect do to meet these requirements?

A. Use AWS Certificate Manager (ACM) to issue an SSL/TLS certificate. Apply the certificate to the ALB. Use the managed renewal feature to automatically rotate the certificate.
B. Use AWS Certificate Manager (ACM) to issue an SSL/TLS certificate. Import the key material from the certificate. Apply the certificate to the ALB. Use the managed renewal feature to automatically rotate the certificate.
C. Use AWS Certificate Manager (ACM) Private Certificate Authority to issue an SSL/TLS certificate from the root CA. Apply the certificate to the ALB. Use the managed renewal feature to automatically rotate the certificate.
D. Use AWS Certificate Manager (ACM) to import an SSL/TLS certificate. Apply the certificate to the ALB. Use Amazon EventBridge (Amazon CloudWatch Events) to send a notification when the certificate is nearing expiration. Rotate the certificate manually.

Correct Answer: D
Section:
Explanation:

QUESTION 33
A company runs its Infrastructure on AWS and has a registered base of 700.000 users for res document management application The company intends to create a product that converts large pdf files to jpg Imago files. The .pdf files average 5 MB in size. The company needs to store the original files and the converted files. A solutions architect must design a scalable solution to accommodate demand that will grow rapidly over lime. Which solution meets these requirements MOST cost-effectively?

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

D. Upload the .pdf files to an AWS Elastic Beanstalk application that includes Amazon EC2 instances, Amazon Elastic File System (Amazon EPS) storage, and an Auto Scaling group. Use a program in the EC2 instances to convert the file to jpg format Save the pdf files and the jpg files in the EBS store.

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

QUESTION 34
A company has more than 5 TB of file data on Windows file servers that run on premises Users and applications interact with the data each day The company is moving its Windows workloads to AWS. As the company continues this process, the company requires access to AWS and on-premises file storage with minimum latency The company needs a solution that minimizes operational overhead and requires no significant changes to the existing file access patterns. The company uses an AWS Site-to-Site VPN connection for connectivity to AWS What should a solutions architect do to meet these requirements?

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

Correct Answer: D
Section:
Explanation:

QUESTION 35
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 36
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 37
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 38
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 39
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 40
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 41
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 42
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 43
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 44
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 45
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 46
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 47
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 48
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 49
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 50
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 51
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 52
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 53
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 54
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 55
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 56
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 57
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 58
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 59
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 60
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 61
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 62
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 63
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 64
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 65
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 66
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 67
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 68
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 69
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 70
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 71
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 72
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 73
A medical research lab produces data that is related to a new study. The lab wants to make the data available with minimum latency to clinics across the country for their on-premises, file-based applications. The data files are stored in an Amazon S3 bucket that has read-only permissions for each clinic.
What should a solutions architect recommend to meet these requirements?

A. Deploy an AWS Storage Gateway file gateway as a virtual machine (VM) on premises at each clinic
B. Migrate the files to each clinic's on-premises applications by using AWS DataSync for processing.
C. Deploy an AWS Storage Gateway volume gateway as a virtual machine (VM) on premises at each clinic.
D. Attach an Amazon Elastic File System (Amazon EFS) file system to each clinic's on-premises servers.

Correct Answer: A
Section:
Explanation:
AWS Storage Gateway is a service that connects an on-premises software appliance with cloud-based storage to provide seamless and secure integration between an organization's on-premises IT environment and AWS's storage infrastructure. By deploying a file gateway as a virtual machine on each clinic's premises, the medical research lab can provide low-latency access to the data stored in the S3 bucket while maintaining read-only permissions for each clinic. This solution allows the clinics to access the data files directly from their on-premises file-based applications without the need for data transfer or migration.

QUESTION 74
A company is using a content management system that runs on a single Amazon EC2 instance. The EC2 instance contains both the web server and the database software. The company must make its website platform highly available and must enable the website to scale to meet user demand.
What should a solutions architect recommend to meet these requirements?

A. Move the database to Amazon RDS, and enable automatic backups. Manually launch another EC2 instance in the same Availability Zone. Configure an Application Load Balancer in the Availability Zone, and set the two instances as targets.
B. Migrate the database to an Amazon Aurora instance with a read replica in the same Availability Zone as the existing EC2 instance. Manually launch another EC2 instance in the same Availability Zone. Configure an Application Load Balancer, and set the two EC2 instances as targets.
C. Move the database to Amazon Aurora with a read replica in another Availability Zone. Create an Amazon Machine Image (AMI) from the EC2 instance. Configure an Application Load Balancer in two Availability Zones. Attach an Auto Scaling group that uses the AMI across two Availability Zones.
D. Move the database to a separate EC2 instance, and schedule backups to Amazon S3. Create an Amazon Machine Image (AMI) from the original EC2 instance. Configure an Application Load Balancer in two Availability Zones. Attach an Auto Scaling group that uses the AMI across two Availability Zones.

Correct Answer: C
Section:
Explanation:
This approach will provide both high availability and scalability for the website platform. By moving the database to Amazon Aurora with a read replica in another availability zone, it will provide a failover option for the database. The use of an Application Load Balancer and an Auto Scaling group across two availability zones allows for automatic scaling of the website to meet increased user demand. Additionally, creating an AMI from the original EC2 instance allows for easy replication of the instance in case of failure.

QUESTION 75
A company has a three-tier application for image sharing. The application uses an Amazon EC2 instance for the front-end layer, another EC2 instance for the application layer, and a third EC2 instance for a MySQL database. A solutions architect must design a scalable and highly available solution that requires the least amount of change to the application. Which solution meets these requirements?

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

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

QUESTION 76
An application running on an Amazon EC2 instance in VPC-A needs to access files in another EC2 instance in VPC-B. Both VPCs are in separate AWS accounts. The network administrator needs to design a solution to configure secure access to EC2 instance in VPC-B from VPCA.
The connectivity should not have a single point of failure or bandwidth concerns.
Which solution will meet these requirements?

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

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

QUESTION 77
A company wants to experiment with individual AWS accounts for its engineer team. The company wants to be notified as soon as the Amazon EC2 instance usage for a given month exceeds a specific threshold for each account. What should a solutions architect do to meet this requirement MOST cost-effectively?

A. Use Cost Explorer to create a daily report of costs by service. Filter the report by EC2 instances. Configure Cost Explorer to send an Amazon Simple Email Service (Amazon SES) notification when a threshold is exceeded.
B. Use Cost Explorer to create a monthly report of costs by service. Filter the report by EC2 instances. Configure Cost Explorer to send an Amazon Simple Email Service (Amazon SES) notification when a threshold is exceeded.
C. Use AWS Budgets to create a cost budget for each account. Set the period to monthly. Set the scope to EC2 instances. Set an alert threshold for the budget. Configure an Amazon Simple Notification Service (Amazon SNS) topic to receive a notification when a threshold is exceeded.
D. Use AWS Cost and Usage Reports to create a report with hourly granularity. Integrate the report data with Amazon Athena. Use Amazon EventBridge to schedule an Athena query. Configure an Amazon Simple Notification Service (Amazon SNS) topic to receive a notification when a threshold is exceeded.

Correct Answer: C
Section:

Explanation:
AWS Budgets allows you to create budgets for your AWS accounts and set alerts when usage exceeds a certain threshold. By creating a budget for each account, specifying the period as monthly and the scope as EC2 instances, you can effectively track the EC2 usage for each account and be notified when a threshold is exceeded. This solution is the most cost-effective option as it does not require additional resources such as Amazon Athena or Amazon EventBridge.

QUESTION 78
A company previously migrated its data warehouse solution to AWS. The company also has an AWS Direct Connect connection. Corporate office users query the data warehouse using a visualization tool. The average size of a query returned by the data warehouse is 50 MB and each webpage sent by the visualization tool is approximately 500 KB. Result sets returned by the data warehouse are not cached. Which solution provides the LOWEST data transfer egress cost for the company?

A. Host the visualization tool on premises and query the data warehouse directly over the internet.
B. Host the visualization tool in the same AWS Region as the data warehouse. Access it over the internet.
C. Host the visualization tool on premises and query the data warehouse directly over a Direct Connect connection at a location in the same AWS Region.
D. Host the visualization tool in the same AWS Region as the data warehouse and access it over a Direct Connect connection at a location in the same Region.

Correct Answer: D
Section:
Explanation:
https://aws.amazon.com/directconnect/pricing/
https://aws.amazon.com/blogs/aws/aws-data-transfer-prices-reduced/

QUESTION 79
An application runs on an Amazon EC2 instance that has an Elastic IP address in VPC

A. The application requires access to a database in VPC B. Both VPCs are in the same AWS account.
Which solution will provide the required access MOST securely?
B. Create a DB instance security group that allows all traffic from the public IP address of the application server in VPC A.
C. Configure a VPC peering connection between VPC A and VPC B.
D. Make the DB instance publicly accessible. Assign a public IP address to the DB instance.
E. Launch an EC2 instance with an Elastic IP address into VPC B. Proxy all requests through the new EC2 instance.

Correct Answer: B
Section:

QUESTION 80
A company runs demonstration environments for its customers on Amazon EC2 instances. Each environment is isolated in its own VPC. The company's operations team needs to be notified when RDP or SSH access to an environment has been established.

A. Configure Amazon CloudWatch Application Insights to create AWS Systems Manager OpsItems when RDP or SSH access is detected.
B. Configure the EC2 instances with an IAM instance profile that has an IAM role with the AmazonSSMManagedInstanceCore policy attached.
C. Publish VPC flow logs to Amazon CloudWatch Logs. Create required metric filters. Create an Amazon CloudWatch metric alarm with a notification action for when the alarm is in the ALARM state.
D. Configure an Amazon EventBridge rule to listen for events of type EC2 Instance State-change Notification. Configure an Amazon Simple Notification Service (Amazon SNS) topic as a target. Subscribe the operations team to the topic.

Correct Answer: C
Section:

QUESTION 81

A solutions architect has created a new AWS account and must secure AWS account root user access.
Which combination of actions will accomplish this? (Choose two.)

A. Ensure the root user uses a strong password.
B. Enable multi-factor authentication to the root user.
C. Store root user access keys in an encrypted Amazon S3 bucket.
D. Add the root user to a group containing administrative permissions.
E. Apply the required permissions to the root user with an inline policy document.

Correct Answer: A, B
Section:

QUESTION 82
A company is moving its on-premises Oracle database to Amazon Aurora PostgreSQL. The database has several applications that write to the same tables. The applications need to be migrated one by one with a month in between each migration. Management has expressed concerns that the database has a high number of reads and writes. The data must be kept in sync across both databases throughout the migration. What should a solutions architect recommend?

A. Use AWS DataSync for the initial migration. Use AWS Database Migration Service (AWS DMS) to create a change data capture (CDC) replication task and a table mapping to select all tables.
B. Use AWS DataSync for the initial migration. Use AWS Database Migration Service (AWS DMS) to create a full load plus change data capture (CDC) replication task and a table mapping to select all tables.
C. Use the AWS Schema Conversion Tool with AWS Database Migration Service (AWS DMS) using a memory optimized replication instance. Create a full load plus change data capture (CDC) replication task and a table mapping to select all tables.
D. Use the AWS Schema Conversion Tool with AWS Database Migration Service (AWS DMS) using a compute optimized replication instance. Create a full load plus change data capture (CDC) replication task and a table mapping to select the largest tables.

Correct Answer: C
Section:

QUESTION 83
A company needs to retain its AWS CloudTrail logs for 3 years. The company is enforcing CloudTrail across a set of AWS accounts by using AWS Organizations from the parent account. The CloudTrail target S3 bucket is configured with S3 Versioning enabled. An S3 Lifecycle policy is in place to delete current objects after 3 years.
After the fourth year of use of the S3 bucket, the S3 bucket metrics show that the number of objects has continued to rise. However, the number of new CloudTrail logs that are delivered to the S3 bucket has remained consistent. Which solution will delete objects that are older than 3 years in the MOST cost-effective manner?

A. Configure the organization's centralized CloudTrail trail to expire objects after 3 years.
B. Configure the S3 Lifecycle policy to delete previous versions as well as current versions.
C. Create an AWS Lambda function to enumerate and delete objects from Amazon S3 that are older than 3 years.
D. Configure the parent account as the owner of all objects that are delivered to the S3 bucket.

Correct Answer: B
Section:
Explanation:
https://docs.aws.amazon.com/awscloudtrail/latest/userguide/best-practicessecurity.html#:~:text=The%20CloudTrail%20trail,time%20has%20passed.

QUESTION 84
A company manages its own Amazon EC2 instances that run MySQL databases. The company is manually managing replication and scaling as demand increases or decreases. The company needs a new solution that simplifies the process of adding or removing compute capacity to or from its database tier as needed. The solution also must offer improved performance, scaling, and durability with minimal effort from operations. Which solution meets these requirements?

A. Migrate the databases to Amazon Aurora Serverless for Aurora MySQL.
B. Migrate the databases to Amazon Aurora Serverless for Aurora PostgreSQL.
C. Combine the databases into one larger MySQL database. Run the larger database on larger EC2 instances.
D. Create an EC2 Auto Scaling group for the database tier. Migrate the existing databases to the new environment.

Correct Answer: A
Section:
Explanation:
https://aws.amazon.com/rds/aurora/serverless/

QUESTION 85
A company is concerned that two NAT instances in use will no longer be able to support the traffic needed for the company's application. A solutions architect wants to implement a solution that is highly available, fault tolerant, and automatically scalable.
What should the solutions architect recommend?

A. Remove the two NAT instances and replace them with two NAT gateways in the same Availability Zone.
B. Use Auto Scaling groups with Network Load Balancers for the NAT instances in different Availability Zones.
C. Remove the two NAT instances and replace them with two NAT gateways in different Availability Zones.
D. Replace the two NAT instances with Spot Instances in different Availability Zones and deploy a Network Load Balancer.

Correct Answer: C
Section:
Explanation:
If you have resources in multiple Availability Zones and they share one NAT gateway, and if the NAT gateway's Availability Zone is down, resources in the other Availability Zones lose internet access. To create an Availability Zone- independent architecture, create a NAT gateway in each Availability Zone and configure your routing to ensure that resources use the NAT gateway in the same Availability Zone.
https://docs.aws.amazon.com/vpc/latest/userguide/vpc-nat- gateway.html#nat-gatewaybasics

QUESTION 86
A company runs an application on a group of Amazon Linux EC2 instances. For compliance reasons, the company must retain all application log files for 7 years. The log files will be analyzed by a reporting tool that must be able to access all the files concurrently.
Which storage solution meets these requirements MOST cost-effectively?

A. Amazon Elastic Block Store (Amazon EBS)
B. Amazon Elastic File System (Amazon EFS)
C. Amazon EC2 instance store
D. Amazon S3

Correct Answer: D
Section:

QUESTION 87
A company's reporting system delivers hundreds of .csv files to an Amazon S3 bucket each day. The company must convert these files to Apache Parquet format and must store the files in a transformed data bucket. Which solution will meet these requirements with the LEAST development effort?

A. Create an Amazon EMR cluster with Apache Spark installed. Write a Spark application to transform the data. Use EMR File System (EMRFS) to write files to the transformed data bucket.
B. Create an AWS Glue crawler to discover the data. Create an AWS Glue extract, transform, and load (ETL) job to transform the data. Specify the transformed data bucket in the output step.
C. Use AWS Batch to create a job definition with Bash syntax to transform the data and output the data to the transformed data bucket. Use the job definition to submit a job. Specify an array job as the job type.
D. Create an AWS Lambda function to transform the data and output the data to the transformed data bucket. Configure an event notification for the S3 bucket. Specify the Lambda function as the destination for the event

notification.

Correct Answer: B
Section:
Explanation:
https://docs.aws.amazon.com/prescriptive-guidance/latest/patterns/three-aws-glue-etl-job-typesfor-converting-data-to-apache-parquet.html

QUESTION 88
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 CreateImage 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 CreateImage 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 CreateImage when an API call is detected.
C. Create an Amazon EventBridge (Amazon CloudWatch Events) rule for the CreateImage API call.
Configure the target as an Amazon Simple Notification Service (Amazon SNS) topic to send an alert when a CreateImage 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 CreateImage API call is detected.

Correct Answer: C
Section:
Explanation:
https://docs.aws.amazon.com/AWSEC2/latest/WindowsGuide/monitor-amievents.html#:~:text=For%20example%2C%20you%20can%20create%20an%20EventBridge%20rule%20that%20detects%20when%20the%20AMI%20creation%20process%20has%20completed%20and%20then%20invokes%20an%20Amazon%20SNS%20topic%20to%20send%20an%20email%20notification%20to%20you.

QUESTION 89
A company offers a food delivery service that is growing rapidly. Because of the growth, the company's order processing system is experiencing scaling problems during peak traffic hours. The current architecture includes the following:
• A group of Amazon EC2 instances that run in an Amazon EC2 Auto Scaling group to collect orders from the application • Another group of EC2 instances that run in an Amazon EC2 Auto Scaling group to fulfill orders The order collection process occurs quickly, but the order fulfillment process can take longer. Data must not be lost because of a scaling event. A solutions architect must ensure that the order collection process and the order fulfillment process can both scale properly during peak traffic hours. The solution must optimize utilization of the company's AWS resources. Which solution meets these requirements?

A. Use Amazon CloudWatch metrics to monitor the CPU of each instance in the Auto Scaling groups.
Configure each Auto Scaling group's minimum capacity according to peak workload values.
B. Use Amazon CloudWatch metrics to monitor the CPU of each instance in the Auto Scaling groups.
Configure a CloudWatch alarm to invoke an Amazon Simple Notification Service (Amazon SNS) topic that creates additional Auto Scaling groups on demand.
C. Provision two Amazon Simple Queue Service (Amazon SQS) queues: one for order collection and another for order fulfillment. Configure the EC2 instances to poll their respective queue. Scale the Auto Scaling groups based on notifications that the queues send.
D. Provision two Amazon Simple Queue Service (Amazon SQS) queues: one for order collection and another for order fulfillment. Configure the EC2 instances to poll their respective queue. Create a metric based on a backlog per instance calculation. Scale the Auto Scaling groups based on this metric.

Correct Answer: D
Section:
Explanation:
The number of instances in your Auto Scaling group can be driven by how long it takes to process a message and the acceptable amount of latency (queue delay). The solution is to use a backlog per instance metric with the target value being the acceptable backlog per instance to maintain.

QUESTION 90
A company is developing a marketing communications service that targets mobile app users. The company needs to send confirmation messages with Short Message Service (SMS) to its users. The users must be able to reply

to the SMS messages. The company must store the responses for a year for analysis.
What should a solutions architect do to meet these requirements?

A. Create an Amazon Connect contact flow to send the SMS messages. Use AWS Lambda to process the responses.
B. Build an Amazon Pinpoint journey. Configure Amazon Pinpoint to send events to an Amazon Kinesis data stream for analysis and archiving.
C. Use Amazon Simple Queue Service (Amazon SQS) to distribute the SMS messages. Use AWS Lambda to process the responses.
D. Create an Amazon Simple Notification Service (Amazon SNS) FIFO topic. Subscribe an Amazon Kinesis data stream to the SNS topic for analysis and archiving.

Correct Answer: B
Section:
Explanation:
https://aws.amazon.com/pinpoint/product-details/sms/ Two-Way Messaging: Receive SMSmessages from your customers and reply back to them in a chat-like interactive experience. With Amazon Pinpoint, you can create automatic responses when customers send you messages that contain certain keywords. You can even use Amazon Lex to create conversational bots. A majority of mobile phone users read incoming SMS messages almost immediately after receiving them. If you need to be able to provide your customers with urgent or important information, SMS messaging may be the right solution for you. You can use Amazon Pinpoint to create targeted groups of customers, and then send them campaign-based messages. You can also use Amazon Pinpoint to send direct messages, such as appointment confirmations, order updates, and one-time passwords.

QUESTION 91
The customers of a finance company request appointments with financial advisors by sending text messages. A web application that runs on Amazon EC2 instances accepts the appointment requests. The text messages are published to an Amazon Simple Queue Service (Amazon SQS) queue through the web application. Another application that runs on EC2 instances then sends meeting invitations and meeting confirmation email messages to the customers. After successful scheduling, this application stores the meeting information in an Amazon DynamoDB database. As the company expands, customers report that their meeting invitations are taking longer to arrive.
What should a solutions architect recommend to resolve this issue?

A. Add a DynamoDB Accelerator (DAX) cluster in front of the DynamoDB database.
B. Add an Amazon API Gateway API in front of the web application that accepts the appointment requests.
C. Add an Amazon CloudFront distribution. Set the origin as the web application that accepts the appointment requests.
D. Add an Auto Scaling group for the application that sends meeting invitations. Configure the Auto Scaling group to scale based on the depth of the SQS queue.

Correct Answer: D
Section:
Explanation:
To resolve the issue of longer delivery times for meeting invitations, the solutions architect can recommend adding an Auto Scaling group for the application that sends meeting invitations and configuring the Auto Scaling group to scale based on the depth of the SQS queue. This will allow the application to scale up as the number of appointment requests increases, improving the performance and delivery times of the meeting invitations.

QUESTION 92

An online retail company has more than 50 million active customers and receives more than 25,000 orders each day. The company collects purchase data for customers and stores this data in Amazon S3. Additional customer data is stored in Amazon RDS.

The company wants to make all the data available to various teams so that the teams can perform analytics. The solution must provide the ability to manage fine-grained permissions for the data and must minimize operational overhead. Which solution will meet these requirements?

A. Migrate the purchase data to write directly to Amazon RDS. Use RDS access controls to limit access.
B. Schedule an AWS Lambda function to periodically copy data from Amazon RDS to Amazon S3.
Create an AWS Glue crawler. Use Amazon Athena to query the data. Use S3 policies to limit access.
C. Create a data lake by using AWS Lake Formation. Create an AWS Glue JDBC connection to Amazon RDS. Register the S3 bucket in Lake Formation. Use Lake Formation access controls to limit access.
D. Create an Amazon Redshift cluster. Schedule an AWS Lambda function to periodically copy data from Amazon S3 and Amazon RDS to Amazon Redshift. Use Amazon Redshift access controls to limit access.

Correct Answer: C
Section:
Explanation:
https://aws.amazon.com/blogs/big-data/manage-fine-grained-access-control-using-aws-lakeformation/

QUESTION 93
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:

QUESTION 94
An ecommerce company stores terabytes of customer data in the AWS Cloud. The data contains personally identifiable information (Pll). The company wants to use the data in three applications. Only one of the applications needs to process the Pll. The Pll must be removed before the other two applications process the data. Which solution will meet these requirements with the LEAST operational overhead?

A. Store the data in an Amazon DynamoDB table. Create a proxy application layer to intercept and process the data that each application requests.
B. Store the data in an Amazon S3 bucket. Process and transform the data by using S3 Object Lambda before returning the data to the requesting application.
C. Process the data and store the transformed data in three separate Amazon S3 buckets so that each application has its own custom dataset. Point each application to its respective S3 bucket.
D. Process the data and store the transformed data in three separate Amazon DynamoDB tables so that each application has its own custom dataset. Point each application to its respective DynamoDB table.

Correct Answer: B
Section:
Explanation:
https://aws.amazon.com/blogs/aws/introducing-amazon-s3-object-lambda-use-your-code-toprocess-data-as-it-is-being-retrieved-from-s3/ S3 Object Lambda is a new feature of Amazon S3 that enables customers to add their own code to process data retrieved from S3 before returning it to the application. By using S3 Object Lambda, the data can be processed and transformed in real-time, without the need to store multiple copies of the data in separate S3 buckets or DynamoDB tables. In this case, the Pll can be removed from the data by the code added to S3 Object Lambda before returning the data to the two applications that do not need to process Pll. The one application that requires Pll can be pointed to the original S3 bucket where the Pll is still stored. Using S3 Object Lambda is the simplest and most cost-effective solution, as it eliminates the need to maintain multiple copies of the same data in different buckets or tables, which can result in additional storage costs and operational overhead.

QUESTION 95
A company uses Amazon API Gateway to run a private gateway with two REST APIs in the same VPC.
The BuyStock RESTful web service calls the CheckFunds RESTful web service to ensure that enough funds are available before a stock can be purchased. The company has noticed in the VPC flow logs that the BuyStock RESTful web service calls the CheckFunds RESTful web service over the internet instead of through the VPC. A solutions architect must implement a solution so that the APIs communicate through the VPC.

Which solution will meet these requirements with the FEWEST changes to the code?
(Select Correct Option/s and give detailed explanation from AWS Certified Solutions Architect Associate (SAA-C03) Study Manual or documents)

A. Add an X-APl-Key header in the HTTP header for authorization.
B. Use an interface endpoint.
C. Use a gateway endpoint.
D.

Correct Answer: B
Section:
Explanation:
A. Add an X-APl-Key header in the HTTP header for authorization.
B. Use an interface endpoint.
C. Use a gateway endpoint.
D. Add an Amazon Simple Queue Service (Amazon SQS) queue between the two REST APIs.
Answer: B
Explanation:
Using an interface endpoint will allow the BuyStock RESTful web service and the CheckFunds RESTful web service to communicate through the VPC without any changes to the code. An interface endpoint creates an elastic network interface (ENI) in the customer's VPC, and then configures the route tables to route traffic from the APIs to the ENI. This will ensure that the two APIs will communicate through the VPC without any changes to the code.

QUESTION 96
A solutions architect needs to optimize storage costs. The solutions architect must identify any Amazon S3 buckets that are no longer being accessed or are rarely accessed. Which solution will accomplish this goal with the LEAST operational overhead?

A. Analyze bucket access patterns by using the S3 Storage Lens dashboard for advanced activity metrics.
B. Analyze bucket access patterns by using the S3 dashboard in the AWS Management Console.
C. Turn on the Amazon CloudWatch BucketSizeBytes metric for buckets. Analyze bucket access patterns by using the metrics data with Amazon Athena.
D. Turn on AWS CloudTrail for S3 object monitoring. Analyze bucket access patterns by using CloudTrail logs that are integrated with Amazon CloudWatch Logs.

Correct Answer: A
Section:
Explanation:
S3 Storage Lens is a fully managed S3 storage analytics solution that provides a comprehensive view of object storage usage, activity trends, and recommendations to optimize costs. Storage Lens allows you to analyze object access patterns across all of your S3 buckets and generate detailed metrics and reports.

QUESTION 97
A company has multiple AWS accounts that use consolidated billing. The company runs several active high performance Amazon RDS for Oracle On-Demand DB instances for 90 days. The company's finance team has access to AWS Trusted Advisor in the consolidated billing account and all other AWS accounts.
The finance team needs to use the appropriate AWS account to access the Trusted Advisor check recommendations for RDS. The finance team must review the appropriate Trusted Advisor check to reduce RDS costs. Which combination of steps should the finance team take to meet these requirements? (Select TWO.)

A. Use the Trusted Advisor recommendations from the account where the RDS instances are running.
B. Use the Trusted Advisor recommendations from the consolidated billing account to see all RDS instance checks at the same time.
C. Review the Trusted Advisor check for Amazon RDS Reserved Instance Optimization.
D. Review the Trusted Advisor check for Amazon RDS Idle DB Instances.
E. Review the Trusted Advisor check for Amazon Redshift Reserved Node Optimization.

Correct Answer: B, C

Section:
Explanation:
1. Use the Trusted Advisor recommendations from the consolidated billing account to see all RDS instance checks at the same time. The consolidated billing account has access to all the other AWS accounts that use consolidated billing. Using the Trusted Advisor recommendations from the consolidated billing account will allow the finance team to see all RDS instance checks for all accounts at the same time.
2. Review the Trusted Advisor check for Amazon RDS Reserved Instance Optimization.
The Trusted Advisor check for Amazon RDS Reserved Instance Optimization provides recommendations for purchasing reserved instances to reduce RDS costs. By reviewing this check, the finance team can identify which RDS instances can be converted to reserved instances to save costs.

QUESTION 98
A company is designing the network for an online multi-player game. The game uses the UDP networking protocol and will be deployed in eight AWS Regions. The network architecture needs to minimize latency and packet loss to give end users a high-quality gaming experience.
Which solution will meet these requirements?

A. Set up a transit gateway in each Region. Create inter-Region peering attachments between each transit gateway.
B. Set up AWS Global Accelerator with UDP listeners and endpoint groups in each Region.
C. Set up Amazon CloudFront with UDP turned on. Configure an origin in each Region.
D. Set up a VPC peering mesh between each Region. Turn on UDP for each VPC.

Correct Answer: B
Section:
Explanation:
The best solution for this situation is option B, setting up AWS Global Accelerator with UDP listeners and endpoint groups in each Region. AWS Global Accelerator is a networking service that improves the availability and performance of internet applications by routing user requests to the nearest AWS
RegionÂ [1]. It also improves the performance of UDP applications by providing faster, more reliable data transfers with lower latency and fewer packet losses. By setting up UDP listeners and endpoint groups in each Region, Global Accelerator will route traffic to the nearest Region for faster response times and a better user experience.

QUESTION 99
A company hosts a serverless application on AWS. The application uses Amazon API Gateway, AWS Lambda, and an Amazon RDS for PostgreSQL database. The company notices an increase in application errors that result from database connection timeouts during times Of peak traffic or unpredictable traffic. The company needs a solution that reduces the application failures with the least amount of change to the code. What should a solutions architect do to meet these requirements?

A. Reduce the Lambda concurrency rate.
B. Enable RDS Proxy on the RDS DB instance.
C. Resize the RDS DB instance class to accept more connections.
D. Migrate the database to Amazon DynamoDB with on-demand scaling.

Correct Answer: B
Section:
Explanation:
Using RDS Proxy, you can handle unpredictable surges in database traffic. Otherwise, these surges might cause issues due to oversubscribing connections or creating new connections at a fast rate. RDS Proxy establishes a database connection pool and reuses connections in this pool. This approach avoids the memory and CPU overhead of opening a new database connection each time. To protect the database against oversubscription, you can control the number of database connections that are created. https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/rds-proxy.html

QUESTION 100
A hospital needs to store patient records in an Amazon S3 bucket. The hospital's compliance team must ensure that all protected health information (PHI) is encrypted in transit and at rest. The compliance team must administer the encryption key for data at rest.
Which solution will meet these requirements?

A. Create a public SSL/TLS certificate in AWS Certificate Manager (ACM). Associate the certificate with Amazon S3. Configure default encryption for each S3 bucket to use server-side encryption with AWS KMS keys (SSE-KMS). Assign the compliance team to manage the KMS keys.

B. Use the aws:SecureTransport condition on S3 bucket policies to allow only encrypted connections over HTTPS (TLS). Configure default encryption for each S3 bucket to use server-side encryption with S3 managed encryption keys (SSE-S3). Assign the compliance team to manage the SSE-S3 keys.
C. Use the aws:SecureTransport condition on S3 bucket policies to allow only encrypted connections over HTTPS (TLS). Configure default encryption for each S3 bucket to use server-side encryption with AWS KMS keys (SSE-KMS). Assign the compliance team to manage the KMS keys.
D. Use the aws:SecureTransport condition on S3 bucket policies to allow only encrypted connections over HTTPS (TLS). Use Amazon Macie to protect the sensitive data that is stored in Amazon S3. Assign the compliance team to manage Macie.

Correct Answer: C
Section:
Explanation:
it allows the compliance team to manage the KMS keys used for server-side encryption, thereby providing the necessary control over the encryption keys. Additionally, the use of the "aws:SecureTransport" condition on the bucket policy ensures that all connections to the S3 bucket are encrypted in transit.

QUESTION 101
A social media company runs its application on Amazon EC2 instances behind an Application Load Balancer (ALB). The ALB is the origin for an Amazon CloudFront distribution. The application has more than a billion images stored in an Amazon S3 bucket and processes thousands of images each second. The company wants to resize the images dynamically and serve appropriate formats to clients.Which solution will meet these requirements with the LEAST operational overhead?

A. Install an external image management library on an EC2 instance. Use the image management library to process the images.
B. Create a CloudFront origin request policy. Use the policy to automatically resize images and to serve the appropriate format based on the User-Agent HTTP header in the request.
C. Use a Lambda@Edge function with an external image management library. Associate the Lambda@Edge function with the CloudFront behaviors that serve the images.
D. Create a CloudFront response headers policy. Use the policy to automatically resize images and to serve the appropriate format based on the User-Agent HTTP header in the request.

Correct Answer: C
Section:
Explanation:
Lambda@Edge is a service that allows you to run Lambda functions at CloudFront edge locations. It can be used to modify requests and responses that flow through CloudFront. CloudFront origin request policy is a policy that controls the values (URL query strings, HTTP headers, and cookies) that are included in requests that CloudFront sends to the origin. It can be used to collect additional information at the origin or to customize the origin response. CloudFront response headers policy is a policy that specifies the HTTP headers that CloudFront removes or adds in responses that it sends to viewers. It can be used to add security or custom headers to responses.Based on these definitions, the solution that will meet the requirements with the least operational overhead is:1. Use a Lambda@Edge function with an external image management library. Associate the Lambda@Edge function with the CloudFront behaviors that serve the images.This solution would allow the application to use a Lambda@Edge function to resize the images dynamically and serve appropriate formats to clients based on the User-Agent HTTP header in the request. The Lambda@Edge function would run at the edge locations, reducing latency and load onthe origin. The application code would only need to include an external image management library that can perform image manipulation tasks1.

QUESTION 102
A company has a production workload that is spread across different AWS accounts in various AWS Regions. The company uses AWS Cost Explorer to continuously monitor costs and usage. The company wants to receive notifications when the cost and usage spending of the workload is unusual.
Which combination of steps will meet these requirements? (Select TWO.)

A. In the AWS accounts where the production workload is running, create a linked account budget by using Cost Explorer in the AWS Cost Management console
B. In ys AWS accounts where the production workload is running, create a linked account monitor by using AWS Cost Anomaly Detection in the AWS Cost Management console
C. In the AWS accounts where the production workload is running, create a Cost and Usage Report by using Cost Anomaly Detection in the AWS Cost Management console.
D. Create a report and send email messages to notify the company on a weekly basis.
E. Create a subscription with the required threshold and notify the company by using weekly summaries.

Correct Answer: B, E

Section:
Explanation:
AWS Cost Anomaly Detection allows you to create monitors that track the cost and usage of your AWS resources and alert you when there is an unusual spending pattern. You can create monitors based on different dimensions, such as AWS services, accounts, tags, or cost categories. You can also create alert subscriptions that notify you by email or Amazon SNS when an anomaly is detected. You can specify the threshold and frequency of the alerts, and choose to receive weekly summaries of your anomalies.
Reference URLs:
1https://aws.amazon.com/aws-cost-management/aws-cost-anomaly-detection/
2https://docs.aws.amazon.com/cost-management/latest/userguide/getting-started-ad.html
3https://docs.aws.amazon.com/cost-management/latest/userguide/manage-ad.html

QUESTION 103
A company stores raw collected data in an Amazon S3 bucket. The data is used for several types of analytics on behalf of the company's customers. The type of analytics requested to determines the access pattern on the S3 objects.
The company cannot predict or control the access pattern. The company wants to reduce its S3 costs.
which solution will meet these requirements?

A. Use S3 replication to transition infrequently accessed objects to S3 Standard-Infrequent Access (S3 Standard-1A)
B. Use S3 Lifecycle rules to transition objects from S3 Standard to Standard-Infrequent Access (S3 Standard-1A).
C. Use S3 Lifecycle rules for transition objects from S3 Standard to S3 Intelligent-Tiering.
D. Use S3 Inventory to identify and transition objects that have not been accessed from S3 Standard to S3 Intelligent-Tiering.

Correct Answer: C
Section:
Explanation:
S3 Intelligent-Tiering is a storage class that automatically reduces storage costs by moving data to the most cost-effective access tier based on access frequency. It has two access tiers: frequent access and infrequent access. Data is stored in the frequent access tier by default, and moved to the infrequent access tier after 30 consecutive days of no access. If the data is accessed again, it is moved back to the frequent access tier1. By using S3 Lifecycle rules to transition objects from S3 Standard to S3 Intelligent-Tiering, the solution can reduce S3 costs for data with unknown or changing access patterns.
a) Use S3 replication to transition infrequently accessed objects to S3 Standard-Infrequent Access (S3 Standard-IA). This solution will not meet the requirement of reducing S3 costs for data with unknown or changing access patterns, as S3 replication is a feature that copies objects across buckets or Regions for redundancy or compliance purposes. It does not automatically move objects to a different storage class based on access frequency2.
b) Use S3 Lifecycle rules to transition objects from S3 Standard to Standard-Infrequent Access (S3 Standard-IA). This solution will not meet the requirement of reducing S3 costs for data with unknown or changing access patterns, as S3 Standard-IA is a storage class that offers lower storage costs than S3 Standard, but charges a retrieval fee for accessing the data. It is suitable for long-lived and infrequently accessed data, not for data with changing access patterns1.
d) Use S3 Inventory to identify and transition objects that have not been accessed from S3 Stand-ard to S3 Intelligent-Tiering. This solution will not meet the requirement of reducing S3 costs for data with unknown or changing access patterns, as S3 Inventory is a feature that provides a report of the objects in a bucket and their metadata on a daily or weekly basis. It does not automatically move objects to a different storage class based on access frequency3.
Reference URL: https://aws.amazon.com/s3/storage-classes/intelligent-tiering/
S3 Intelligent-Tiering is the best solution for reducing S3 costs when the access pattern is unpredictable or changing. S3 Intelligent-Tiering automatically moves objects between two access tiers (frequent and infrequent) based on the access frequency, without any performance impact or retrieval fees. S3 Intelligent-Tiering also has an optional archive tier for objects that are rarely accessed. S3 Lifecycle rules can be used to transition objects from S3 Standard to S3 Intelligent-Tiering.
Reference URLs:
1https://aws.amazon.com/s3/storage-classes/intelligent-tiering/
2https://docs.aws.amazon.com/AmazonS3/latest/userguide/using-intelligent-tiering.html
3https://docs.aws.amazon.com/AmazonS3/latest/userguide/intelligent-tiering-overview.html

QUESTION 104
A company has applications hosted on Amazon EC2 instances with IPv6 addresses. The applications must initiate communications with other external applications using the internet.
However, the company's security policy states that any external service cannot initiate a connection to the EC2 instances.
What should a solutions architect recommend to resolve this issue?

A. Create a NAT gateway and make it the destination of the subnet's route table.

B. Create an internet gateway and make it the destination of the subnet's route table
C. Create a virtual private gateway and make it the destination of the subnet's route table.
D. Create an egress-only internet gateway and make it the destination of the subnet's route table.

Correct Answer: D
Section:
Explanation:
An egress-only internet gateway is a VPC component that allows outbound communication over IPv6 from instances in your VPC to the internet, and prevents the internet from initiating an IPv6 connection with your instances. This meets the company's security policy and requirements. To use an egress-only internet gateway, you need to add a route in the subnet's route table that routes IPv6 internet traffic (::/0) to the egress-only internet gateway.
Reference URLs:
1https://docs.aws.amazon.com/vpc/latest/userguide/egress-only-internet-gateway.html
2https://dev.to/aws-builders/what-is-an-egress-only-internet-gateways-in-aws-7gp
3https://docs.aws.amazon.com/vpc/latest/userguide/route-table-options.html

QUESTION 105
A company is making a prototype of the infrastructure for its new website by manually provisioning the necessary infrastructure. This infrastructure includes an Auto Scaling group, an Application Load Balancer, and an Amazon RDS database. After the configuration has been thoroughly validated, the company wants the capability to immediately deploy the infrastructure for development and production use in two Availability Zones in an automated fashion.
What should a solutions architect recommend to meet these requirements?

A. Use AWS Systems Manager to replicate and provision the prototype infrastructure in two Availability Zones.
B. Define the infrastructure as a template by using the prototype infrastructure as a guide. Deploy the infrastructure with AWS CloudFormation
C. Use AWS Config to record the inventory of resources that are used in the prototype infrastructure. Use AWS Config to deploy the prototype infrastructure into two Availability Zones.
D. Use AWS Elastic Beanstalk and configure it to use an automated reference to the prototype infrastructure to automatically deploy new environments in two Availability Zones

Correct Answer: B
Section:
Explanation:
AWS CloudFormation is a service that helps you model and set up your AWS resources by using templates that describe all the resources that you want, such as Auto Scaling groups, load balancers, and databases. You can use AWS CloudFormation to deploy your infrastructure in an automated and consistent way across multiple environments and regions. You can also use AWS CloudFormation to update or delete your infrastructure as a single unit.
Reference URLs:
1 https://aws.amazon.com/cloudformation/
2 https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/Welcome.html
3 https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/cfn-whatis-concepts.html

QUESTION 106
A company is developing a mobile gaming app in a single AWS Region. The app runs on multiple Amazon EC2 instances in an Auto Scaling group. The company stores the app data in Amazon DynamoDB. The app communicates by using TCP traffic and UDP traffic between the users and the servers. The application will be used globally. The company wants to ensure the lowest possible latency for all users.
Which solution will meet these requirements?

A. Use AWS Global Accelerator to create an accelerator. Create an Application Load Balancer (ALB) behind an accelerator endpoint that uses Global Accelerator integration and listening on the TCP and UDP ports. Update the Auto Scaling group to register instances on the ALB.
B. Use AWS Global Accelerator to create an accelerator. Create a Network Load Balancer (NLB) behind an accelerator endpoint that uses Global Accelerator integration and listening on the TCP and UDP ports. Update the Auto Scaling group to register instances on the NLB
C. Create an Amazon CloudFront content delivery network (CDN) endpoint. Create a Network Load Balancer (NLB) behind the endpoint and listening on the TCP and UDP ports. Update the Auto Scaling group to register instances on the NLB. Update CloudFront to use the NLB as the origin.

D. Create an Amazon Cloudfront content delivery network (CDN) endpoint. Create an Application Load Balancer (ALB) behind the endpoint and listening on the TCP and UDP ports. Update the Auto Scaling group to register instances on the ALB. Update CloudFront to use the ALB as the origin

Correct Answer: B
Section:
Explanation:
AWS Global Accelerator is a networking service that improves the performance and availability of applications for global users. It uses the AWS global network to route user traffic to the optimal endpoint based on performance and health. It also provides static IP addresses that act as a fixed entry point to the applications and support both TCP and UDP protocols1. By using AWS Global Accelerator, the solution can ensure the lowest possible latency for all users.
a) Use AWS Global Accelerator to create an accelerator. Create an Application Load Balancer (ALB) behind an accelerator endpoint that uses Global Accelerator integration and listening on the TCP and UDP ports. Update the Auto Scaling group to register instances on the ALB. This solution will not work, as ALB does not support UDP protocol2.
c) Create an Amazon CloudFront content delivery network (CDN) endpoint. Create a Network Load Balancer (NLB) behind the endpoint and listening on the TCP and UDP ports. Update the Auto Scaling group to register instances on the NLB. Update CloudFront to use the NLB as the origin. This solution will not work, as CloudFront does not support UDP protocol3.
d) Create an Amazon Cloudfront content delivery network (CDN) endpoint. Create an Application Load Balancer (ALB) behind the endpoint and listening on the TCP and UDP ports. Update the Auto Scaling group to register instances on the ALB. Update CloudFront to use the ALB as the origin. This solution will not work, as CloudFront and ALB do not support UDP protocol23.
Reference URL: https://aws.amazon.com/global-accelerator/

QUESTION 107
A business application is hosted on Amazon EC2 and uses Amazon S3 for encrypted object storage. The chief information security officer has directed that no application traffic between the two services should traverse the public internet.
Which capability should the solutions architect use to meet the compliance requirements?

A. AWS Key Management Service (AWS KMS)
B. VPC endpoint
C. Private subnet
D. Virtual private gateway

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

QUESTION 108
A company wants to securely exchange data between its software as a service (SaaS) application Salesforce account and Amazon S3. The company must encrypt the data at rest by using AWS Key Management Service (AWS KMS) customer managed keys (CMKs). The company must also encrypt the data in transit. The company has enabled API access for the Salesforce account.
Which solution will meet these requirements with the LEAST development effort?

A. Create AWS Lambda functions to transfer the data securely from Salesforce to Amazon S3.
B. Create an AWS Step Functions workflow Define the task to transfer the data securely from Salesforce to Amazon S3.
C. Create Amazon AppFlow flows to transfer the data securely from Salesforce to Amazon S3.
D. Create a custom connector for Salesforce to transfer the data securely from Salesforce to Amazon S3.

Correct Answer: C
Section:
Explanation:
Amazon AppFlow is a fully managed integration service that enables users to transfer data securely between SaaS applications and AWS services. It supports Salesforce as a source and Amazon S3 as a destination. It also supports encryption of data at rest using AWS KMS CMKs and encryption of data in transit using SSL/TLS1. By using Amazon AppFlow, the solution can meet the requirements with the least development effort.
a) Create AWS Lambda functions to transfer the data securely from Salesforce to Amazon S3. This solution will not meet the requirement of the least development effort, as it involves writing custom code to interact with Salesforce and Amazon S3 APIs, handle authentication, encryption, error handling, and monitoring2.

b) Create an AWS Step Functions workflow Define the task to transfer the data securely from Salesforce to Amazon S3. This solution will not meet the requirement of the least development effort, as it involves creating a state machine definition to orchestrate the data transfer task, and invoking Lambda functions or other services to perform the actual data transfer3.
d) Create a custom connector for Salesforce to transfer the data securely from Salesforce to Ama-zon S3. This solution will not meet the requirement of the least development effort, as it involves using the Amazon AppFlow Custom Connector SDK to build and deploy a custom connector for Salesforce, which requires additional configuration and management.
Reference URL: https://aws.amazon.com/appflow/

QUESTION 109
A company has multiple AWS accounts for development work. Some staff consistently use oversized Amazon EC2 instances, which causes the company to exceed the yearly budget for the development accounts The company wants to centrally restrict the creation of AWS resources in these accounts
Which solution will meet these requirements with the LEAST development effort?

A. Develop AWS Systems Manager templates that use an approved EC2 creation process. Use the approved Systems Manager templates to provision EC2 instances.
B. Use AWS Organizations to organize the accounts into organizational units (OUs). Define and attach a service control policy (SCP) to control the usage of EC2 instance types.
C. Configure an Amazon EventBridge rule that invokes an AWS Lambda function when an EC2 instance is created. Stop disallowed EC2 instance types.
D. Set up AWS Service Catalog products for the staff to create the allowed EC2 instance types Ensure that staff can deploy EC2 instances only by using the Service Catalog products.

Correct Answer: B
Section:
Explanation:
AWS Organizations is a service that helps users centrally manage and govern multiple AWS accounts. It allows users to create organizational units (OUs) to group accounts based on business needs or other criteria. It also allows users to define and attach service control policies (SCPs) to OUs or accounts to restrict the actions that can be performed by the accounts1. By using AWS Organizations, the solution can centrally restrict the creation of AWS resources in the development accounts.
a) Develop AWS Systems Manager templates that use an approved EC2 creation process. Use the approved Systems Manager templates to provision EC2 instances. This solution will not meet the requirement of the least development effort, as it involves developing and maintaining custom templates for EC2 creation, and relying on the staff to use the approved templates instead of enforcing a restriction2.
c) Configure an Amazon EventBridge rule that invokes an AWS Lambda function when an EC2 instance is created. Stop disallowed EC2 instance types. This solution will not meet the requirement of the least development effort, as it involves writing custom code for Lambda functions, and handling events and errors for EC2 creation3.
d) Set up AWS Service Catalog products for the staff to create the allowed EC2 instance types En-sure that staff can deploy EC2 instances only by using the Service Catalog products. This solution will not meet the requirement of the least development effort, as it involves setting up and managing Service Catalog products for EC2 creation, and ensuring that staff can only use Service Catalog products instead of enforcing a restriction.
Reference URL: https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps.html

QUESTION 110
A company runs container applications by using Amazon Elastic Kubernetes Service (Amazon EKS). The company's workload is not consistent throughout the day The company wants Amazon EKS to scale in and out according to the workload.
Which combination of steps will meet these requirements with the LEAST operational overhead? {Select TWO.)

A. Use an AWS Lambda function to resize the EKS cluster
B. Use the Kubernetes Metrics Server to activate horizontal pod autoscaling.
C. Use the Kubernetes Cluster Autoscaler to manage the number of nodes in the cluster.
D. Use Amazon API Gateway and connect it to Amazon EKS
E. Use AWS App Mesh to observe network activity.

Correct Answer: B, C
Section:
Explanation:
https://docs.aws.amazon.com/eks/latest/userguide/horizontal-pod-autoscaler.html
https://docs.aws.amazon.com/eks/latest/userguide/autoscaling.html
Horizontal pod autoscaling is a feature of Kubernetes that automatically scales the number of pods in a deployment, replication controller, or replica set based on that resource's CPU utilization. It requires a metrics source such as the Kubernetes Metrics Server to provide CPU usage data1. Cluster autoscaling is a feature of Kubernetes that automatically adjusts the number of nodes in a cluster when pods fail or are rescheduled onto other nodes. It requires an integration with AWS Auto Scaling groups to manage the EC2 instances that join the cluster2. By using both horizontal pod autoscaling and cluster autoscaling, the solution can ensure that Amazon EKS scales in and out according to the workload.

QUESTION 111
A company has a mobile chat application with a data store based in Amazon uynamoUb. users would like new messages to be read with as little latency as possible A solutions architect needs to design an optimal solution that requires minimal application changes.
Which method should the solutions architect select?

A. Configure Amazon DynamoDB Accelerator (DAX) for the new messages table. Update the code to use the DAXendpoint.
B. Add DynamoDB read repticas to handle the increased read load. Update the application to point to the read endpoint for the read replicas.
C. Double the number of read capacity units for the new messages table in DynamoDB. Continue to use the existing DynamoDB endpoint.
D. Add an Amazon ElastiCache for Redis cache to the application stack. Update the application to point to the Redis cache endpoint instead of DynamoDB.

Correct Answer: A
Section:
Explanation:
https://aws.amazon.com/premiumsupport/knowledge-center/dynamodb-high-latency/
Amazon DynamoDB Accelerator (DAX) is a fully managed in-memory cache for DynamoDB that improves the performance of DynamoDB tables by up to 10 times and provides microsecond level of response time at any scale. It is compatible with DynamoDB API operations and requires minimal code changes to use1. By configuring DAX for the new messages table, the solution can reduce the latency for reading new messages with minimal application changes.
b) Add DynamoDB read repticas to handle the increased read load. Update the application to point to the read endpoint for the read replicas. This solution will not work, as DynamoDB does not support read replicas as a feature. Read replicas are available for Amazon RDS, not for DynamoDB2.
c) Double the number of read capacity units for the new messages table in DynamoDB. Continue to use the existing DynamoDB endpoint. This solution will not meet the requirement of reading new messages with as little latency as possible, as increasing the read capacity units will only increase the throughput of DynamoDB, not the performance or latency3.
d) Add an Amazon ElastiCache for Redis cache to the application stack. Update the application to point to the Redis cache endpoint instead of DynamoDB. This solution will not meet the requirement of minimal application changes, as adding ElastiCache for Redis will require significant code changes to implement caching logic, such as querying cache first, updating cache after writing to DynamoDB, and invalidating cache when needed.
Reference URL: https://aws.amazon.com/dynamodb/dax/

QUESTION 112
A company needs to integrate with a third-party data feed. The data feed sends a webhook to notifyan external service when new data is ready for consumption A developer wrote an AWS Lambfefunction to retrieve data when the company receives a webhook callback The developer must makethe Lambda function available for the third party to call.Which solution will meet these requirements with the MOST operational efficiency?

A. Create a function URL for the Lambda function. Provide the Lambda function URL to the third party for the webhook.
B. Deploy an Application Load Balancer (ALB) in front of the Lambda function. Provide the ALB URL to the third party for the webhook
C. Create an Amazon Simple Notification Service (Amazon SNS) topic. Attach the topic to the Lambda function. Provide the public hostname of the SNS topic to the third party for the webhook.
D. Create an Amazon Simple Queue Service (Amazon SQS) queue. Attach the queue to the Lambda function. Provide the public hostname of the SQS queue to the third party for the webhook.

Correct Answer: A
Section:
Explanation:
A function URL is a unique identifier for a Lambda function that can be used to invoke the function over HTTPS. It is composed of the API endpoint of the AWS Region where the function is deployed, and the name or ARN of the function1. By creating a function URL for the Lambda function, the solution can make the Lambda function available for the third party to call with the most operational efficiency.
b) Deploy an Application Load Balancer (ALB) in front of the Lambda function. Provide the ALB URL to the third party for the webhook. This solution will not meet the requirement of the most operational efficiency, as it involves creating and managing an additional resource (ALB) that is not necessary for invoking a Lambda function over HTTPS2.
c) Create an Amazon Simple Notification Service (Amazon SNS) topic. Attach the topic to the Lambda function. Provide the public hostname of the SNS topic to the third party for the webhook. This solution will not work, as Amazon SNS topics do not have public hostnames that can be used as webhooks. SNS topics are used to publish messages to subscribers, not to receive messages from external sources3.
d) Create an Amazon Simple Queue Service (Amazon SQS) queue. Attach the queue to the Lamb-da function. Provide the public hostname of the SQS queue to the third party for the webhook. This solution will not work, as Amazon SQS queues do not have public hostnames that can be used as webhooks. SQS queues are used to send, store, and receive messages between AWS services, not to receive messages from external sources.
Reference URL: https://docs.aws.amazon.com/lambda/latest/dg/lambda-api-permissions-ref.html

QUESTION 113
A company wants to move from many standalone AWS accounts to a consolidated, multi-account architecture The company plans to create many new AWS accounts for different business units. The company needs to

authenticate access to these AWS accounts by using a centralized corporate directory service.
Which combination of actions should a solutions architect recommend to meet these requirements? (Select TWO.)

A. Create a new organization in AWS Organizations with all features turned on. Create the new AWS accounts in the organization.
B. Set up an Amazon Cognito identity pool. Configure AWS IAM Identity Center (AWS Single Sign-On) to accept Amazon Cognito authentication.
C. Configure a service control policy (SCP) to manage the AWS accounts. Add AWS IAM Identity Center (AWS Single Sign-On) to AWS Directory Service.
D. Create a new organization in AWS Organizations. Configure the organization's authentication mechanism to use AWS Directory Service directly.
E. Set up AWS IAM Identity Center (AWS Single Sign-On) in the organization. Configure IAM Identity Center, and integrate it with the company's corporate directory service.

Correct Answer: A, E
Section:
Explanation:
AWS Organizations is a service that helps users centrally manage and govern multiple AWS accounts. It allows users to create organizational units (OUs) to group accounts based on business needs or other criteria. It also allows users to define and attach service control policies (SCPs) to OUs or accounts to restrict the actions that can be performed by the accounts1. By creating a new organization in AWS Organizations with all features turned on, the solution can consolidate and manage the new AWS accounts for different business units.
AWS IAM Identity Center (formerly known as AWS Single Sign-On) is a service that provides single sign-on access for all of your AWS accounts and cloud applications. It connects with Microsoft Active Directory through AWS Directory Service to allow users in that directory to sign in to a personalized AWS access portal using their existing Active Directory user names and passwords. From the AWS access portal, users have access to all the AWS accounts and cloud applications that they have permissions for2. By setting up IAM Identity Center in the organization and integrating it with the company's corporate directory service, the solution can authenticate access to these AWS accounts using a centralized corporate directory service.
b) Set up an Amazon Cognito identity pool. Configure AWS IAM Identity Center (AWS Single Sign-On) to accept Amazon Cognito authentication. This solution will not meet the requirement of authenticating access to these AWS accounts by using a centralized corporate directory service, as Amazon Cognito is a service that provides user sign-up, sign-in, and access control for web and mobile applications, not for corporate directory services3.
c) Configure a service control policy (SCP) to manage the AWS accounts. Add AWS IAM Identi-ty Center (AWS Single Sign-On) to AWS Directory Service. This solution will not work, as SCPs are used to restrict the actions that can be performed by the accounts in an organization, not to manage the accounts themselves1. Also, IAM Identity Center cannot be added to AWS Directory Service, as it is a separate service that connects with Microsoft Active Directory through AWS Directory Service2.
d) Create a new organization in AWS Organizations. Configure the organization's authentication mechanism to use AWS Directory Service directly. This solution will not work, as AWS Organizations does not have an authentication mechanism that can use AWS Directory Service directly. AWS Organizations relies on IAM Identity Center to provide single sign-on access for the accounts in an organization.
Reference URL: https://docs.aws.amazon.com/organizations/latest/userguide/orgs_integrate_services.html

QUESTION 114
A group requires permissions to list an Amazon S3 bucket and delete objects from that bucket An administrator has created the following IAM policy to provide access to the bucket and applied that policy to the group. The group is not able to delete objects in the bucket. The company follows least-privilege access rules.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Action": [
                "s3:ListBucket",
                "s3:DeleteObject"
            ],
            "Resource": [
                "arn:aws:s3:::bucket-name"
            ],
            "Effect": "Allow"
        }
    ]
}
```

A)

```
"Action": [
    "s3:*Object"
],
"Resource": [
    "arn:aws:s3:::bucket-name/*"
],
"Effect": "Allow"
```

B)

```
"Action": [
    "s3:DeleteObject"
],
"Resource": [
    "arn:aws:s3:::bucket-name*"
],
"Effect": "Allow"
```

C)

```
"Action": [
    "s3:DeleteObject"
],
"Resource": [
    "arn:aws:s3:::bucket-name/*"
],
"Effect": "Allow"
```

A. Option A
B. Option B
C. Option C
D. Option D

Correct Answer: D
Section:
Explanation:
{
'Version': '2012-10-17',
'Statement': [
{
'Action': [
's3:ListBucket',
's3:DeleteObject'
],
'Resource': [
'arn:aws:s3:::<bucket-name>'
],
'Effect': 'Allow',
},
{

'Action': 's3:*DeleteObject',
'Resource': [
'arn:aws:s3:::<bucket-name>/*' # <- The policy clause kludge 'added' to match the solution (Q248.1) example
],
'Effect': 'Allow'
}
]
}

QUESTION 115
A solutions architect is designing a REST API in Amazon API Gateway for a cash payback service The application requires 1 GB of memory and 2 GB of storage for its computation resources. The application will require that the data is in a relational format.
Which additional combination of AWS services will meet these requirements with the LEAST administrative effort? {Select TWO.)

A. Amazon EC2
B. AWS Lambda
C. Amazon RDS
D. Amazon DynamoDB
E. Amazon Elastic Kubernetes Services (Amazon EKS)

Correct Answer: B, C
Section:
Explanation:
AWS Lambda is a service that lets users run code without provisioning or managing servers. It automatically scales and manages the underlying compute resources for the code. It supports multiple languages, such as Java, Python, Node.js, and Go1. By using AWS Lambda for the REST API, the solution can meet the requirements of 1 GB of memory and minimal administrative effort.
Amazon RDS is a service that makes it easy to set up, operate, and scale a relational database in the cloud. It provides cost-efficient and resizable capacity while automating time-consuming administration tasks such as hardware provisioning, database setup, patching and backups. It supports multiple database engines, such as MySQL, PostgreSQL, Oracle, and SQL Server2. By using Amazon RDS for the data store, the solution can meet the requirements of 2 GB of storage and a relational format.
a) Amazon EC2. This solution will not meet the requirement of minimal administrative effort, as Amazon EC2 is a service that provides virtual servers in the cloud that users have to configure and manage themselves. It requires users to choose an instance type, an operating system, a security group, and other options3.
d) Amazon DynamoDB. This solution will not meet the requirement of a relational format, as Amazon DynamoDB is a service that provides a key-value and document database that delivers single-digit millisecond performance at any scale. It is a non-relational or NoSQL database that does not support joins or transactions.
e) Amazon Elastic Kubernetes Services (Amazon EKS). This solution will not meet the requirement of minimal administrative effort, as Amazon EKS is a service that provides a fully managed Kubernetes service that users have to configure and manage themselves. It requires users to create clusters, nodes groups, pods, services, and other Kubernetes resources.
Reference URL: https://aws.amazon.com/lambda/

QUESTION 116
A company has resources across multiple AWS Regions and accounts. A newly hired solutions architect discovers a previous employee did not provide details about the resources invent^. The solutions architect needs to build and map the relationship details of the various workloads across all accounts.
Which solution will meet these requirements in the MOST operationally efficient way?

A. Use AWS Systems Manager Inventory to generate a map view from the detailed view report.
B. Use AWS Step Functions to collect workload details Build architecture diagrams of the workloads manually.
C. Use Workload Discovery on AWS to generate architecture diagrams of the workloads.
D. Use AWS X-Ray to view the workload details Build architecture diagrams with relationships

Correct Answer: C
Section:

Explanation:
Workload Discovery on AWS (formerly called AWS Perspective) is a tool that visualizes AWS Cloud workloads. It maintains an inventory of the AWS resources across your accounts and Regions, mapping relationships between them, and displaying them in a web UI. It also allows you to query AWS Cost and Usage Reports, search for resources, save and export architecture diagrams, and more1. By using Workload Discovery on AWS, the solution can build and map the relationship details of the various workloads across all accounts with the least operational effort.
a) Use AWS Systems Manager Inventory to generate a map view from the detailed view report. This solution will not meet the requirement of building and mapping the relationship details of the various workloads across all accounts, as AWS Systems Manager Inventory is a feature that collects metadata from your managed instances and stores it in a central Amazon S3 bucket. It does not provide a map view or architecture diagrams of the workloads2.
b) Use AWS Step Functions to collect workload details Build architecture diagrams of the work-loads manually. This solution will not meet the requirement of the least operational effort, as it involves creating and managing state machines to orchestrate the workload details collection, and building architecture diagrams manually.
d) Use AWS X-Ray to view the workload details Build architecture diagrams with relationships. This solution will not meet the requirement of the least operational effort, as it involves instrumenting your applications with X-Ray SDKs to collect workload details, and building architecture diagrams manually.
Reference URL: https://aws.amazon.com/solutions/implementations/workload-discovery-on-aws/

QUESTION 117
A company's applications run on Amazon EC2 instances in Auto Scaling groups. The company notices that its applications experience sudden traffic increases on random days of the week The company wants to maintain application performance during sudden traffic increases.
Which solution will meet these requirements MOST cost-effectively?

A. Use manual scaling to change the size of the Auto Scaling group.
B. Use predictive scaling to change the size of the Auto Scaling group.
C. Use dynamic scaling to change the size of the Auto Scaling group.
D. Use schedule scaling to change the size of the Auto Scaling group

Correct Answer: C
Section:
Explanation:
Dynamic scaling is a type of autoscaling that automatically adjusts the number of EC2 instances in an Auto Scaling group based on demand or load. It uses CloudWatch alarms to trigger scaling actions when a specified metric crosses a threshold. It can scale out (add instances) or scale in (remove instances) as needed1. By using dynamic scaling, the solution can maintain application performance during sudden traffic increases most cost-effectively.
a) Use manual scaling to change the size of the Auto Scaling group. This solution will not meet the requirement of maintaining application performance during sudden traffic increases, as manual scaling requires users to manually increase or decrease the number of instances through a CLI or console. It does not respond automatically to changes in demand or load2.
b) Use predictive scaling to change the size of the Auto Scaling group. This solution will not meet the requirement of most cost-effectiveness, as predictive scaling uses machine learning and artificial intelligence tools to evaluate traffic loads and anticipate when more or fewer resources are needed. It performs scheduled scaling actions based on the prediction, which may not match the actual demand or load at any given time. Predictive scaling is more suitable for scenarios where there are predictable traffic patterns or known changes in traffic loads3.
d) Use schedule scaling to change the size of the Auto Scaling group. This solution will not meet the requirement of maintaining application performance during sudden traffic increases, as schedule scaling performs scaling actions at specific times that users schedule. It does not respond automatically to changes in demand or load. Schedule scaling is more suitable for scenarios where there are predictable traffic drops or spikes at specific times of the day.
Reference URL: https://docs.aws.amazon.com/autoscaling/ec2/userguide/as-scale-based-on-demand.html

QUESTION 118
A company has an on-premises server that uses an Oracle database to process and store customer information The company wants to use an AWS database service to achieve higher availability and to improve application performance. The company also wants to offload reporting from its primary database system.
Which solution will meet these requirements in the MOST operationally efficient way?

A. Use AWS Database Migration Service (AWS DMS) to create an Amazon RDS DB instance in multiple AWS Regions Point the reporting functions toward a separate DB instance from the primary DB instance.
B. Use Amazon RDS in a Single-AZ deployment to create an Oracle database Create a read replica in the same zone as the primary DB instance. Direct the reporting functions to the read replica.
C. Use Amazon RDS deployed in a Multi-AZ cluster deployment to create an Oracle database Direct the reporting functions to use the reader instance in the cluster deployment
D. Use Amazon RDS deployed in a Multi-AZ instance deployment to create an Amazon Aurora database. Direct the reporting functions to the reader instances.

Correct Answer: D

Section:
Explanation:
Amazon Aurora is a fully managed relational database that is compatible with MySQL and PostgreSQL. It provides up to five times better performance than MySQL and up to three times better performance than PostgreSQL. It also provides high availability and durability by replicating data across multiple Availability Zones and continuously backing up data to Amazon S31. By using Amazon RDS deployed in a Multi-AZ instance deployment to create an Amazon Aurora database, the solution can achieve higher availability and improve application performance.
Amazon Aurora supports read replicas, which are separate instances that share the same underlying storage as the primary instance. Read replicas can be used to offload read-only queries from the primary instance and improve performance. Read replicas can also be used for reporting functions2. By directing the reporting functions to the reader instances, the solution can offload reporting from its primary database system.
a) Use AWS Database Migration Service (AWS DMS) to create an Amazon RDS DB instance in multiple AWS Regions Point the reporting functions toward a separate DB instance from the pri-mary DB instance. This solution will not meet the requirement of using an AWS database service, as AWS DMS is a service that helps users migrate databases to AWS, not a database service itself. It also involves creating multiple DB instances in different Regions, which may increase complexity and cost.
b) Use Amazon RDS in a Single-AZ deployment to create an Oracle database Create a read replica in the same zone as the primary DB instance. Direct the reporting functions to the read replica. This solution will not meet the requirement of achieving higher availability, as a Single-AZ deployment does not provide failover protection in case of an Availability Zone outage. It also involves using Oracle as the database engine, which may not provide better performance than Aurora.
c) Use Amazon RDS deployed in a Multi-AZ cluster deployment to create an Oracle database Di-rect the reporting functions to use the reader instance in the cluster deployment. This solution will not meet the requirement of improving application performance, as Oracle may not provide better performance than Aurora. It also involves using a cluster deployment, which is only supported for Aurora, not for Oracle.
Reference URL: https://aws.amazon.com/rds/aurora/

QUESTION 119
A law firm needs to share information with the public The information includes hundreds of files that must be publicly readable Modifications or deletions of the files by anyone before a designated future date are prohibited. Which solution will meet these requirements in the MOST secure way?

A. Upload all files to an Amazon S3 bucket that is configured for static website hosting. Grant read-only IAM permissions to any AWS principals that access the S3 bucket until the designated date.
B. Create a new Amazon S3 bucket with S3 Versioning enabled Use S3 Object Lock with a retention period in accordance with the designated date Configure the S3 bucket for static website hosting. Set an S3 bucket policy to allow read-only access to the objrcts.
C. Create a new Amazon S3 bucket with S3 Versioning enabled Configure an event trigger to run an AWS Lambda function in case of object modification or deletion. Configure the Lambda function to replace the objects with the original versions from a private S3 bucket.
D. Upload all files to an Amazon S3 bucket that is configured for static website hosting. Select the folder that contains the files. Use S3 Object Lock with a retention period in accordance with the designated date. Grant read-only IAM permissions to any AWS principals that access the S3 bucket.

Correct Answer: B
Section:
Explanation:
Amazon S3 is a service that provides object storage in the cloud. It can be used to store and serve static web content, such as HTML, CSS, JavaScript, images, and videos1. By creating a new Amazon S3 bucket and configuring it for static website hosting, the solution can share information with the public.
Amazon S3 Versioning is a feature that keeps multiple versions of an object in the same bucket. It helps protect objects from accidental deletion or overwriting by preserving, retrieving, and restoring every version of every object stored in an S3 bucket2. By enabling S3 Versioning on the new bucket, the solution can prevent modifications or deletions of the files by anyone.
Amazon S3 Object Lock is a feature that allows users to store objects using a write-once-read-many (WORM) model. It can help prevent objects from being deleted or overwritten for a fixed amount of time or indefinitely. It requires S3 Versioning to be enabled on the bucket3. By using S3 Object Lock with a retention period in accordance with the designated date, the solution can prohibit modifications or deletions of the files by anyone before that date.
Amazon S3 bucket policies are JSON documents that define access permissions for a bucket and its objects. They can be used to grant or deny access to specific users or groups based on conditions such as IP address, time of day, or source bucket. By setting an S3 bucket policy to allow read-only access to the objects, the solution can ensure that the files are publicly readable.
a) Upload all files to an Amazon S3 bucket that is configured for static website hosting. Grant read-only IAM permissions to any AWS principals that access the S3 bucket until the designated date. This solution will not meet the requirement of prohibiting modifications or deletions of the files by anyone before a designated future date, as IAM permissions only apply to AWS principals, not to public users. It also does not use any feature to prevent accidental or intentional deletion or overwriting of the files.
c) Create a new Amazon S3 bucket with S3 Versioning enabled Configure an event trigger to run an AWS Lambda function in case of object modification or deletion. Configure the Lambda func-tion to replace the objects with the original versions from a private S3 bucket. This solution will not meet the requirement of prohibiting modifications or deletions of the files by anyone before a designated future date, as it only reacts to object modification or deletion events after they occur. It also involves creating and managing an additional resource (Lambda function) and a private S3 bucket.
d) Upload all files to an Amazon S3 bucket that is configured for static website hosting. Select the folder that contains the files. Use S3 Object Lock with a retention period in accordance with the designated date. Grant read-only IAM permissions to any AWS principals that access the S3 bucket. This solution will not meet the requirement of prohibiting modifications or deletions of the files by anyone before a designated future date, as it does not

enable S3 Versioning on the bucket, which is required for using S3 Object Lock. It also does not allow read-only access to public users.
Reference URL: https://docs.aws.amazon.com/AmazonS3/latest/userguide/WebsiteHosting.html

QUESTION 120
A company is looking for a solution that can store video archives in AWS from old news footage. The company needs to minimize costs and will rarely need to restore these files. When the h|es are needed, they must be available in a maximum of five minutes.
What is the MOST cost-effective solution?

A. Store the video archives in Amazon S3 Glacier and use Expedited retrievals.
B. Store the video archives in Amazon S3 Glacier and use Standard retrievals.
C. Store the video archives in Amazon S3 Standard-Infrequent Access (S3 Standard-IA).
D. Store the video archives in Amazon S3 One Zone-Infrequent Access (S3 One Zone-IA)

Correct Answer: A
Section:
Explanation:
Amazon S3 Glacier is a storage class that provides secure, durable, and extremely low-cost storage for data archiving and long-term backup. It is designed for data that is rarely accessed and for which retrieval times of several hours are suitable1. By storing the video archives in Amazon S3 Glacier, the solution can minimize costs.
Amazon S3 Glacier offers three options for data retrieval: Expedited, Standard, and Bulk. Expedited retrievals typically return data in 1--5 minutes and are suitable for Active Archive use cases. Standard retrievals typically complete within 3--5 hours and are suitable for less urgent needs. Bulk retrievals typically complete within 5--12 hours and are the lowest-cost retrieval option2. By using Expedited retrievals, the solution can meet the requirement of restoring the files in a maximum of five minutes.
b) Store the video archives in Amazon S3 Glacier and use Standard retrievals. This solution will not meet the requirement of restoring the files in a maximum of five minutes, as Standard retrievals typically complete within 3--5 hours.
c) Store the video archives in Amazon S3 Standard-Infrequent Access (S3 Standard-IA). This solution will not meet the requirement of minimizing costs, as S3 Standard-IA is a storage class that provides low-cost storage for data that is accessed less frequently but requires rapid access when needed. It has a higher storage cost than S3 Glacier.
d) Store the video archives in Amazon S3 One Zone-Infrequent Access (S3 One Zone-IA). This solution will not meet the requirement of minimizing costs, as S3 One Zone-IA is a storage class that provides low-cost storage for data that is accessed less frequently but requires rapid access when needed. It has a higher storage cost than S3 Glacier.
Reference URL: https://aws.amazon.com/s3/glacier/

QUESTION 121
A company is creating an application that runs on containers in a VPC. The application stores and accesses data in an Amazon S3 bucket During the development phase, the application will store and access 1 TB of data in Amazon S3 each day. The company wants to minimize costs and wants to prevent traffic from traversing the internet whenever possible.
Which solution will meet these requirements?

A. Enable S3 Intelligent-Tiering for the S3 bucket.
B. Enable S3 Transfer Acceleration for the S3 bucket.
C. Create a gateway VPC endpoint for Amazon S3. Associate this endpoint with all route tables in the VPC.
D. Create an interface endpoint for Amazon S3 in the VPC. Associate this endpoint with all route tables in the VPC.

Correct Answer: C
Section:
Explanation:
A gateway VPC endpoint for Amazon S3 enables private connections between the VPC and Amazon S3 that do not require an internet gateway or NAT device. This minimizes costs and prevents traffic from traversing the internet. A gateway VPC endpoint uses a prefix list as the route target in a VPC route table to route traffic privately to Amazon S31. Associating the endpoint with all route tables in the VPC ensures that all subnets can access Amazon S3 through the endpoint.
Option A is incorrect because S3 Intelligent-Tiering is a storage class that optimizes storage costs by automatically moving objects between two access tiers based on changing access patterns. It does not affect the network traffic between the VPC and Amazon S32.
Option B is incorrect because S3 Transfer Acceleration is a feature that enables fast, easy, and secure transfers of files over long distances between clients and an S3 bucket. It does not prevent traffic from traversing the internet3.

Option D is incorrect because an interface VPC endpoint for Amazon S3 is powered by AWS PrivateLink, which requires an elastic network interface (ENI) with a private IP address in each subnet. This adds complexity and cost to the solution. Moreover, an interface VPC endpoint does not support cross-Region access to Amazon S3. Reference URL:1: https://docs.aws.amazon.com/vpc/latest/privatelink/vpc-endpoints-s3.html2: https://docs.aws.amazon.com/AmazonS3/latest/userguide/storage-class-intro.html#sc-dynamic-data-access3: https://docs.aws.amazon.com/AmazonS3/latest/userguide/transfer-acceleration.html : https://aws.amazon.com/blogs/architecture/choosing-your-vpc-endpoint-strategy-for-amazon-s3/

QUESTION 122
A company stores data in PDF format in an Amazon S3 bucket The company must follow a legal requirement to retain all new and existing data in Amazon S3 for 7 years.
Which solution will meet these requirements with the LEAST operational overhead?

A. Turn on the S3 Versionmg feature for the S3 bucket Configure S3 Lifecycle to delete the data after 7 years. Configure multi-factor authentication (MFA) delete for all S3 objects.
B. Turn on S3 Object Lock with governance retention mode for the S3 bucket Set the retention period to expire after 7 years. Recopy all existing objects to bring the existing data into compliance
C. Turn on S3 Object Lock with compliance retention mode for the S3 bucket. Set the retention period to expire after 7 years. Recopy all existing objects to bring the existing data into compliance
D. Turn on S3 Object Lock with compliance retention mode for the S3 bucket. Set the retention period to expire after 7 years. Use S3 Batch Operations to bring the existing data into compliance

Correct Answer: C
Section:
Explanation:
S3 Object Lock enables a write-once-read-many (WORM) model for objects stored in Amazon S3. It can help prevent objects from being deleted or overwritten for a fixed amount of time or indefinitely1. S3 Object Lock has two retention modes: governance mode and compliance mode. Compliance mode provides the highest level of protection and prevents any user, including the root user, from deleting or modifying an object version until the retention period expires. To use S3 Object Lock, a new bucket with Object Lock enabled must be created, and a default retention period can be optionally configured for objects placed in the bucket2. To bring existing objects into compliance, they must be recopied into the bucket with a retention period specified.
Option A is incorrect because S3 Versioning and S3 Lifecycle do not provide WORM protection for objects. Moreover, MFA delete only applies to deleting object versions, not modifying them.
Option B is incorrect because governance mode allows users with special permissions to override or remove the retention settings or delete the object if necessary. This does not meet the legal requirement of retaining all data for 7 years.
Option D is incorrect because S3 Batch Operations cannot be used to apply compliance mode retention periods to existing objects. S3 Batch Operations can only apply governance mode retention periods or legal holds.
Reference URL:2: https://docs.aws.amazon.com/AmazonS3/latest/userguide/object-lock-console.html3: https://docs.aws.amazon.com/AmazonS3/latest/userguide/storage-class-intro.html#sc-dynamic-data-access4: https://docs.aws.amazon.com/AmazonS3/latest/userguide/transfer-acceleration.html1: https://docs.aws.amazon.com/AmazonS3/latest/userguide/object-lock.html : https://docs.aws.amazon.com/AmazonS3/latest/userguide/object-lock-overview.html : https://docs.aws.amazon.com/AmazonS3/latest/userguide/object-lock-managing.html : https://aws.amazon.com/blogs/storage/managing-amazon-s3-access-with-vpc-endpoints-and-s3-access-points/

QUESTION 123
An image hosting company uploads its large assets to Amazon S3 Standard buckets The company uses multipart upload in parallel by using S3 APIs and overwrites if the same object is uploaded again. For the first 30 days after upload, the objects will be accessed frequently. The objects will be used less frequently after 30 days, but the access patterns for each object will be inconsistent The company must optimize its S3 storage costs while maintaining high availability and resiliency of stored assets.
Which combination of actions should a solutions architect recommend to meet these requirements? (Select TWO.)

A. Move assets to S3 Intelligent-Tiering after 30 days.
B. Configure an S3 Lifecycle policy to clean up incomplete multipart uploads.
C. Configure an S3 Lifecycle policy to clean up expired object delete markers.
D. Move assets to S3 Standard-Infrequent Access (S3 Standard-IA) after 30 days
E. Move assets to S3 One Zone-Infrequent Access (S3 One Zone-IA) after 30 days.

Correct Answer: A, B
Section:
Explanation:
S3 Intelligent-Tiering is a storage class that automatically moves data to the most cost-effective access tier based on access frequency, without performance impact, retrieval fees, or operational overhead1. It is ideal for data with unknown or changing access patterns, such as the company's assets. By moving assets to S3 Intelligent-Tiering after 30 days, the company can optimize its storage costs while maintaining high availability and resilience of stored assets.
S3 Lifecycle is a feature that enables you to manage your objects so that they are stored cost effectively throughout their lifecycle2. You can create lifecycle rules to define actions that Amazon S3 applies to a group of objects. One of the actions is to abort incomplete multipart uploads that can occur when an upload is interrupted. By configuring an S3 Lifecycle policy to clean up incomplete multipart uploads, the company can reduce its storage

costs and avoid paying for parts that are not used.
Option C is incorrect because expired object delete markers are automatically deleted by Amazon S3 and do not incur any storage costs3. Therefore, configuring an S3 Lifecycle policy to clean up expired object delete markers will not have any effect on the company's storage costs.
Option D is incorrect because S3 Standard-IA is a storage class for data that is accessed less frequently, but requires rapid access when needed1. It has a lower storage cost than S3 Standard, but it has a higher retrieval cost and a minimum storage duration charge of 30 days. Therefore, moving assets to S3 Standard-IA after 30 days may not optimize the company's storage costs if the assets are still accessed occasionally.
Option E is incorrect because S3 One Zone-IA is a storage class for data that is accessed less frequently, but requires rapid access when needed1. It has a lower storage cost than S3 Standard-IA, but it stores data in only one Availability Zone and has less resilience than other storage classes. It also has a higher retrieval cost and a minimum storage duration charge of 30 days. Therefore, moving assets to S3 One Zone-IA after 30 days may not optimize the company's storage costs if the assets are still accessed occasionally or require high availability. Reference URL:1: https://docs.aws.amazon.com/AmazonS3/latest/userguide/storage-class-intro.html2: https://docs.aws.amazon.com/AmazonS3/latest/userguide/object-lifecycle-mgmt.html3: https://docs.aws.amazon.com/AmazonS3/latest/userguide/delete-or-empty-bucket.html#delete-bucket-considerations : https://docs.aws.amazon.com/AmazonS3/latest/userguide/mpuoverview.html : https://aws.amazon.com/certification/certified-solutions-architect-associate/

QUESTION 124
A company designed a stateless two-tier application that uses Amazon EC2 in a single Availability Zone and an Amazon RDS Multi-AZ DB instance New company management wants to ensure the application is highly available. What should a solutions architect do to meet this requirement?

A. Configure the application to use Multi-AZ EC2 Auto Scaling and create an Application Load Balancer
B. Configure the application to take snapshots of the EC2 instances and send them to a different AWS Region.
C. Configure the application to use Amazon Route 53 latency-based routing to feed requests to the application.
D. Configure Amazon Route 53 rules to handle incoming requests and create a Multi-AZ Application Load Balancer

Correct Answer: A
Section:
Explanation:
https://docs.aws.amazon.com/autoscaling/ec2/userguide/as-add-availability-zone.html

QUESTION 125
A company hosts an online shopping application that stores all orders in an Amazon RDS for PostgreSQL Singfe-AZ DB instance. Management wants to eliminate single points of C^ilure and has asked a solutions architect to recommend an approach to minimize database downtime without requiring any changes to the application code.
Which solution meets these requirements?

A. Convert the existing database instance to a Multi-AZ deployment by modifying the database instance and specifying the Multi-AZ option.
B. Create a new RDS Multi-AZ deployment. Take a snapshot of the current RDS instance and restore the new Multi-AZ deployment with the snapshot.
C. Create a read-only replica of the PostgreSQL database in another Availability Zone. Use Amazon Route 53 weighted record sets to distribute requests across the databases.
D. Place the RDS for PostgreSQL database in an Amazon EC2 Auto Scaling group with a minimum group size of two. Use Amazon Route 53 weighted record sets to distribute requests across instances.

Correct Answer: A
Section:
Explanation:
https://aws.amazon.com/rds/features/multi-az/ To convert an existing Single-AZ DB Instance to a Multi-AZ deployment, use the 'Modify' option corresponding to your DB Instance in the AWS Management Console.

QUESTION 126
A company stores data in Amazon S3. According to regulations, the data must not contain personally identifiable information (Pll). The company recently discovered that S3 buckets have some objects that contain Pll. The company needs to automatically detect Pll in S3 buckets and to notify the company's security team.
Which solution will meet these requirements?

A. Use Amazon Macie. Create an Amazon EventBridge rule to filter the SensitiveData event type from Macie findings and to send an Amazon Simple Notification Service (Amazon SNS) notification to the security team.
B. Use Amazon GuardDuty. Create an Amazon EventBridge rule to filter the CRITICAL event type from GuardDuty findings and to send an Amazon Simple Notification Service (Amazon SNS) notification to the security team.
C. Use Amazon Macie. Create an Amazon EventBridge rule to filter the SensitiveData:S30bject/Personal event type from Macie findings and to send an Amazon Simple Queue Service (Amazon SQS) notification to the

security team.

D. Use Amazon GuardDuty. Create an Amazon EventBridge rule to filter the CRITICAL event type from GuardDuty findings and to send an Amazon Simple Queue Service (Amazon SQS) notification to the security team.

Correct Answer: A
Section:
Explanation:
Amazon Macie can also send its findings to Amazon EventBridge, which is a serverless event bus that makes it easy to connect applications using data from a variety of sources. You can create an EventBridge rule that filters the SensitiveData event type from Macie findings and sends an Amazon SNS notification to the security team. Amazon SNS is a fully managed messaging service that enables you to send messages to subscribers or other applications.
Reference: https://docs.aws.amazon.com/macie/latest/userguide/macie-findings.html#macie-findings-eventbridge

QUESTION 127
A company provides an API interface to customers so the customers can retrieve their financial information. The company expects a larger number of requests during peak usage times of the year.
The company requires the API to respond consistently with low latency to ensure customer satisfaction. The company needs to provide a compute host for the API.
Which solution will meet these requirements with the LEAST operational overhead?

A. Use an Application Load Balancer and Amazon Elastic Container Service (Amazon ECS).
B. Use Amazon API Gateway and AWS Lambda functions with provisioned concurrency.
C. Use an Application Load Balancer and an Amazon Elastic Kubernetes Service (Amazon EKS) cluster.
D. Use Amazon API Gateway and AWS Lambda functions with reserved concurrency.

Correct Answer: B
Section:
Explanation:
Amazon API Gateway is a fully managed service that makes it easy for developers to create, publish, maintain, monitor, and secure APIs at any scale. AWS Lambda is a serverless compute service that lets you run code without provisioning or managing servers. Lambda scales automatically based on the incoming requests, but it may take some time to initialize new instances of your function if there is a sudden increase in demand. This may result in high latency or cold starts for your API. To avoid this, you can use provisioned concurrency, which ensures that your function is initialized and ready to respond at any time. Provisioned concurrency also helps you achieve consistent low latency for your API by reducing the impact of scaling on performance.
Reference: https://docs.aws.amazon.com/apigateway/latest/developerguide/http-api-develop-integrations-lambda.html https://docs.aws.amazon.com/lambda/latest/dg/configuration-concurrency.html

QUESTION 128
A company seeks a storage solution for its application The solution must be highly available and scalable. The solution also must function as a file system, be mountable by multiple Linux instances in AWS and on premises through native protocols, and have no minimum size requirements. The company has set up a Site-to-Site VPN for access from its on-premises network to its VPC.
Which storage solution meets these requirements?

A. Amazon FSx Multi-AZ deployments
B. Amazon Elastic Block Store (Amazon EBS) Multi-Attach volumes
C. Amazon Elastic File System (Amazon EFS) with multiple mount targets
D. Amazon Elastic File System (Amazon EFS) with a single mount target and multiple access points

Correct Answer: C
Section:
Explanation:
Amazon EFS is a fully managed file system that can be mounted by multiple Linux instances in AWS and on premises through native protocols such as NFS and SMB. Amazon EFS has no minimum size requirements and can scale up and down automatically as files are added and removed. Amazon EFS also supports high availability and durability by allowing multiple mount targets in different Availability Zones within a region. Amazon EFS meets all the requirements of the question, while the other options do not.
Reference:
https://aws.amazon.com/efs/

https://docs.aws.amazon.com/wellarchitected/latest/performance-efficiency-pillar/storage-architecture-selection.html
https://aws.amazon.com/blogs/storage/from-on-premises-to-aws-hybrid-cloud-architecture-for-network-file-shares/

QUESTION 129
A company hosts a website on Amazon EC2 instances behind an Application Load Balancer (ALB) The website serves static content Website traffic is increasing and the company is concerned about a potential increase in cost.
What should a solutions architect do to reduce the cost of the website?

A. Create an Amazon CloudFront distribution to cache static files at edge locations.
B. Create an Amazon ElastiCache cluster Connect the ALB to the ElastiCache cluster to serve cached files.
C. Create an AWS WAF web ACL and associate it with the ALB. Add a rule to the web ACL to cache static files.
D. Create a second ALB in an alternative AWS Region Route user traffic to the closest Region to minimize data transfer costs

Correct Answer: A
Section:
Explanation:
Amazon CloudFront is a content delivery network (CDN) that can improve the performance and reduce the cost of serving static content from a website. CloudFront can cache static files at edge locations closer to the users, reducing the latency and data transfer costs. CloudFront can also integrate with Amazon S3 as the origin for the static content, eliminating the need for EC2 instances to host the website. CloudFront meets all the requirements of the question, while the other options do not.
Reference:
https://aws.amazon.com/blogs/architecture/architecting-a-low-cost-web-content-publishing-system/
https://nodeployfriday.com/posts/static-website-hosting/
https://aws.amazon.com/cloudfront/

QUESTION 130
A company uses multiple vendors to distribute digital assets that are stored in Amazon S3 buckets The company wants to ensure that its vendor AWS accounts have the minimum access that is needed to download objects in these S3 buckets
Which solution will meet these requirements with the LEAST operational overhead?

A. Design a bucket policy that has anonymous read permissions and permissions to list ail buckets.
B. Design a bucket policy that gives read-only access to users. Specify IAM entities as principals
C. Create a cross-account IAM role that has a read-only access policy specified for the IAM role.
D. Create a user policy and vendor user groups that give read-only access to vendor users

Correct Answer: C
Section:
Explanation:
A cross-account IAM role is a way to grant users from one AWS account access to resources in another AWS account. The cross-account IAM role can have a read-only access policy attached to it, which allows the users to download objects from the S3 buckets without modifying or deleting them. The cross-account IAM role also reduces the operational overhead of managing multiple IAM users and policies in each account. The cross-account IAM role meets all the requirements of the question, while the other options do not.
Reference:
https://docs.aws.amazon.com/AmazonS3/latest/userguide/example-walkthroughs-managing-access-example2.html
https://aws.amazon.com/blogs/storage/setting-up-cross-account-amazon-s3-access-with-s3-access-points/
https://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_create_for-user_externalid.html

QUESTION 131
A company runs a microservice-based serverless web application. The application must be able to retrieve data from multiple Amazon DynamoDB tables. A solutions architect needs to give the application the ability to retrieve the data with no impact on the baseline performance of the application.
Which solution will meet these requirements in the MOST operationally efficient way?

A. AWSAppSync pipeline resolvers
B. Amazon CloudFront with Lambda@Edge functions
C. Edge-optimized Amazon API Gateway with AWS Lambda functions
D. Amazon Athena Federated Query with a DynamoDB connector

Correct Answer: C
Section:
Explanation:
An edge-optimized API Gateway is a way to create RESTful APIs that can access multiple DynamoDB tables through AWS Lambda functions. The edge-optimized API Gateway provides low latency and high performance by caching API responses at CloudFront edge locations. The AWS Lambda functions can use the AWS SDK to query or scan the DynamoDB tables and return the data to the API Gateway. This solution meets all the requirements of the question, while the other options do not.
Reference:
https://aws.amazon.com/blogs/compute/understanding-database-options-for-your-serverless-web-applications/
https://aws.amazon.com/getting-started/hands-on/build-serverless-web-app-lambda-apigateway-s3-dynamodb-cognito/module-3/
https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/best-practices.html

QUESTION 132
A company has an application that processes customer orders. The company hosts the application on an Amazon EC2 instance that saves the orders to an Amazon Aurora database. Occasionally when traffic is high: the workload does not process orders fast enough.
What should a solutions architect do to write the orders reliably to the database as quickly as possible?

A. Increase the instance size of the EC2 instance when traffic is high. Write orders to Amazon Simple Notification Service (Amazon SNS). Subscribe the database endpoint to the SNS topic.
B. Write orders to an Amazon Simple Queue Service (Amazon SQS) queue. Use EC2 instances in an Auto Scaling group behind an Application Load Balancer to read from the SQS queue and process orders into the database.
C. Write orders to Amazon Simple Notification Service (Amazon SNS) Subscribe the database endpoint to the SNS topic Use EC2 instances in an Auto Scaling group behind an Application Load Balancer to read from the SNS topic.
D. Write orders to an Amazon Simple Queue Service (Amazon SQS) queue when the EC2 instance reaches CPU threshold limits. Use scheduled scaling of EC2 instances in an Auto Scaling group behind an Application Load Balancer to read from the SQS queue and process orders into the database

Correct Answer: B
Section:
Explanation:
Amazon SQS is a fully managed message queuing service that can decouple and scale microservices, distributed systems, and serverless applications. By writing orders to an SQS queue, the application can handle spikes in traffic without losing any orders. The EC2 instances in an Auto Scaling group can read from the SQS queue and process orders into the database at a steady pace. The Application Load Balancer can distribute the load across the EC2 instances and provide health checks. This solution meets all the requirements of the question, while the other options do not.
Reference:
https://docs.aws.amazon.com/wellarchitected/latest/reliability-pillar/welcome.html
https://aws.amazon.com/architecture/serverless/
https://aws.amazon.com/sqs/

QUESTION 133
A company is conducting an internal audit. The company wants to ensure that the data in an Amazon S3 bucket that is associated with the company's AWS Lake Formation data lake does not contain sensitive customer or employee data. The company wants to discover personally identifiable information (PII) or financial information, including passport numbers and credit card numbers.
Which solution will meet these requirements?

A. Configure AWS Audit Manager on the account. Select the Payment Card Industry Data Security Standards (PCI DSS) for auditing.
B. Configure Amazon S3 Inventory on the S3 bucket. Configure Amazon Athena to query the inventory.
C. Configure Amazon Macie to run a data discovery job that uses managed identifiers for the required data types.
D. Use Amazon S3 Select to run a report across the S3 bucket.

Correct Answer: C
Section:
Explanation:
Amazon Macie is a fully managed data security and data privacy service that uses machine learning and pattern matching to discover and protect your sensitive data in AWS. Macie can run data discovery jobs that use managed identifiers for various types of PII or financial information, such as passport numbers and credit card numbers. Macie can also generate findings that alert you to potential issues or risks with your data.
Reference: https://docs.aws.amazon.com/macie/latest/userguide/macie-identifiers.html

QUESTION 134
A company containerized a Windows job that runs on .NET 6 Framework under a Windows container. The company wants to run this job in the AWS Cloud. The job runs every 10 minutes. The job's runtime varies between 1 minute and 3 minutes.
Which solution will meet these requirements MOST cost-effectively?

A. Create an AWS Lambda function based on the container image of the job. Configure Amazon EventBridge to invoke the function every 10 minutes.
B. Use AWS Batch to create a job that uses AWS Fargate resources. Configure the job scheduling to run every 10 minutes.
C. Use Amazon Elastic Container Service (Amazon ECS) on AWS Fargate to run the job. Create a scheduled task based on the container image of the job to run every 10 minutes.
D. Use Amazon Elastic Container Service (Amazon ECS) on AWS Fargate to run the job. Create a standalone task based on the container image of the job. Use Windows task scheduler to run the job every 10 minutes.

Correct Answer: A
Section:
Explanation:
AWS Lambda supports container images as a packaging format for functions. You can use existing container development workflows to package and deploy Lambda functions as container images of up to 10 GB in size. You can also use familiar tools such as Docker CLI to build, test, and push your container images to Amazon Elastic Container Registry (Amazon ECR). You can then create an AWS Lambda function based on the container image of your job and configure Amazon EventBridge to invoke the function every 10 minutes using a cron expression. This solution will be cost-effective as you only pay for the compute time you consume when your function runs.
Reference: https://docs.aws.amazon.com/lambda/latest/dg/images-create.html https://docs.aws.amazon.com/eventbridge/latest/userguide/run-lambda-schedule.html

QUESTION 135
A company uses AWS Organizations with all features enabled and runs multiple Amazon EC2 workloads in the ap-southeast-2 Region. The company has a service control policy (SCP) that prevents any resources from being created in any other Region. A security policy requires the company to encrypt all data at rest.
An audit discovers that employees have created Amazon Elastic Block Store (Amazon EBS) volumes for EC2 instances without encrypting the volumes. The company wants any new EC2 instances that any IAM user or root user launches in ap-southeast-2 to use encrypted EBS volumes. The company wants a solution that will have minimal effect on employees who create EBS volumes.
Which combination of steps will meet these requirements? (Select TWO.)

A. In the Amazon EC2 console, select the EBS encryption account attribute and define a default encryption key.
B. Create an IAM permission boundary. Attach the permission boundary to the root organizational unit (OU). Define the boundary to deny the ec2:CreateVolume action when the ec2:Encrypted condition equals false.
C. Create an SCR Attach the SCP to the root organizational unit (OU). Define the SCP to deny the ec2:CreateVolume action when the ec2:Encrypted condition equals false.
D. Update the IAM policies for each account to deny the ec2:CreateVolume action when the ec2:Encrypted condition equals false.
E. In the Organizations management account, specify the Default EBS volume encryption setting.

Correct Answer: C, E
Section:
Explanation:

QUESTION 136
A company runs a highly available SFTP service. The SFTP service uses two Amazon EC2 Linux instances that run with elastic IP addresses to accept traffic from trusted IP sources on the internet. The SFTP service is backed by shared storage that is attached to the instances. User accounts are created and managed as Linux users in the SFTP servers.
The company wants a serverless option that provides high IOPS performance and highly configurable security. The company also wants to maintain control over user permissions.
Which solution will meet these requirements?

A. Create an encrypted Amazon Elastic Block Store (Amazon EBS) volume. Create an AWS Transfer Family SFTP service with a public endpoint that allows only trusted IP addresses. Attach the EBS volume to the SFTP service endpoint. Grant users access to the SFTP service.
B. Create an encrypted Amazon Elastic File System (Amazon EFS) volume. Create an AWS Transfer Family SFTP service with elastic IP addresses and a VPC endpoint that has internet-facing access. Attach a security group to the endpoint that allows only trusted IP addresses. Attach the EFS volume to the SFTP service endpoint. Grant users access to the SFTP service.
C. Create an Amazon S3 bucket with default encryption enabled. Create an AWS Transfer Family SFTP service with a public endpoint that allows only trusted IP addresses. Attach the S3 bucket to the SFTP service endpoint. Grant users access to the SFTP service.
D. Create an Amazon S3 bucket with default encryption enabled. Create an AWS Transfer Family SFTP service with a VPC endpoint that has internal access in a private subnet. Attach a security group that allows only trusted IP addresses. Attach the S3 bucket to the SFTP service endpoint. Grant users access to the SFTP service.

Correct Answer: C
Section:
Explanation:
AWS Transfer Family is a secure transfer service that enables you to transfer files into and out of AWS storage services using SFTP, FTPS, FTP, and AS2 protocols. You can use AWS Transfer Family to create an SFTP-enabled server with a public endpoint that allows only trusted IP addresses. You can also attach an Amazon S3 bucket with default encryption enabled to the SFTP service endpoint, which will provide high IOPS performance and highly configurable security for your data at rest. You can also maintain control over user permissions by granting users access to the SFTP service using IAM roles or service-managed identities.
Reference: https://docs.aws.amazon.com/transfer/latest/userguide/what-is-aws-transfer-family.html https://docs.aws.amazon.com/transfer/latest/userguide/create-server-s3.html

QUESTION 137
A company is expecting rapid growth in the near future. A solutions architect needs to configure existing users and grant permissions to new users on AWS. The solutions architect has decided to create IAM groups. The solutions architect will add the new users to IAM groups based on department.
Which additional action is the MOST secure way to grant permissions to the new users?

A. Apply service control policies (SCPs) to manage access permissions.
B. Create IAM roles that have least privilege permission. Attach the roles to the IAM groups.
C. Create an IAM policy that grants least privilege permission. Attach the policy to the IAM groups.
D. Create IAM roles. Associate the roles with a permissions boundary that defines the maximum permissions.

Correct Answer: C
Section:
Explanation:
An IAM policy is a document that defines the permissions for an IAM identity (such as a user, group, or role). You can use IAM policies to grant permissions to existing users and groups based on department. You can create an IAM policy that grants least privilege permission, which means that you only grant the minimum permissions required for the users to perform their tasks. You can then attach the policy to the IAM groups, which will apply the policy to all the users in those groups. This solution will reduce operational costs and simplify configuration and management of permissions.
Reference: https://docs.aws.amazon.com/IAM/latest/UserGuide/access_policies.html

QUESTION 138
A company has a serverless application on AWS that uses Amazon RDS as a backend database. The application sometimes experiences a sudden unpredictable increase in traffic. During traffic increases, the application frequently opens and closes connections to the database, which causes the application to receive errors from the database or run out of connections. The company needs to ensure that the application is always scalable and highly available.
Which solution will meet these requirements WITHOUT any code changes to the application?

A. Increase the maximum number of connections in the option group of the RDS database of the serverless application.
B. Increase the instance size of the RDS DB instance to meet the peak load traffic.
C. Deploy Amazon RDS Proxy between the serverless application and Amazon RDS.
D. Purchase Reserved Instances for Amazon RDS to ensure that the database is highly available during peak load traffic.

Correct Answer: C
Section:

Explanation:
Amazon RDS Proxy is a fully managed database proxy that makes applications more scalable, more resilient to database failures, and more secure. RDS Proxy sits between your application and your relational database to pool and share established database connections, improving database efficiency and application scalability. RDS Proxy also reduces the load on the database by handling connection management and query retries for transient errors. By deploying RDS Proxy between your serverless application and Amazon RDS, you can avoid opening and closing connections to the database frequently, which can cause errors or run out of connections. This solution will also reduce operational costs and improve availability of your application.
Reference: https://aws.amazon.com/rds/proxy/

QUESTION 139
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 140
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 141
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 142
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 143
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 144
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 145
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 146
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 147
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 148
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 149
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 150
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 151
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 152
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 153
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 154
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 155
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 156
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 157
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 158
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 159
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 160
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 161
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 162
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 163
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 164
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 165
A company's website is used to sell products to the public. The site runs on Amazon EC2 instances in an Auto Scaling group behind an Application Load Balancer (ALB). There is also an Amazon CloudFront distribution, and AWS WAF is being used to protect against SQL injection attacks. The ALB is the origin for the CloudFront distribution. A recent review of security logs revealed an external malicious IP that needs to be blocked from accessing the website.
What should a solutions architect do to protect the application?

A. Modify the network ACL on the CloudFront distribution to add a deny rule for the malicious IP address.
B. Modify the configuration of AWS WAF to add an IP match condition to block the malicious IP address.
C. Modify the network ACL for the EC2 instances in the target groups behind the ALB to deny the malicious IP address.
D. Modify the security groups for the EC2 instances in the target groups behind the ALB to deny the malicious IP address.

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 IP match rule, which allows users to specify a list of IP addresses or IP address ranges that they want to allow or block. By modifying the configuration of AWS WAF to add an IP match condition to block the malicious IP address, the solution architect can prevent the attacker from accessing the website through the CloudFront distribution and the ALB.
The other options are not correct because they do not effectively block the malicious IP address from accessing the website. Modifying the network ACL on the CloudFront distribution or the EC2 instances in the target groups behind the ALB will not work because network ACLs are stateless and do not evaluate traffic at the application layer. Modifying the security groups for the EC2 instances in the target groups behind the ALB will not work because security groups are stateful and only evaluate traffic at the instance level, not at the load balancer level.
AWS WAF
How AWS WAF works
Working with IP match conditions

QUESTION 166
A company has a new mobile app. Anywhere in the world, users can see local news on topics they choose. Users also can post photos and videos from inside the app.
Users access content often in the first minutes after the content is posted. New content quickly replaces older content, and then the older content disappears. The local nature of the news means that users consume 90% of the content within the AWS Region where it is uploaded.
Which solution will optimize the user experience by providing the LOWEST latency for content uploads?

A. Upload and store content in Amazon S3. Use Amazon CloudFront for the uploads.
B. Upload and store content in Amazon S3. Use S3 Transfer Acceleration for the uploads.
C. Upload content to Amazon EC2 instances in the Region that is closest to the user. Copy the data to Amazon S3.

D. Upload and store content in Amazon S3 in the Region that is closest to the user. Use multiple distributions of Amazon CloudFront.

Correct Answer: B
Section:
Explanation:
The most suitable solution for optimizing the user experience by providing the lowest latency for content uploads is to upload and store content in Amazon S3 and use S3 Transfer Acceleration for the uploads. This solution will enable the company to leverage the AWS global network and edge locations to speed up the data transfer between the users and the S3 buckets.
Amazon S3 is a storage service that provides scalable, durable, and highly available object storage for any type of data. Amazon S3 allows users to store and retrieve data from anywhere on the web, and offers various features such as encryption, versioning, lifecycle management, and replication1.
S3 Transfer Acceleration is a feature of Amazon S3 that helps users transfer data to and from S3 buckets more quickly. S3 Transfer Acceleration works by using optimized network paths and Amazon's backbone network to accelerate data transfer speeds. Users can enable S3 Transfer Acceleration for their buckets and use a distinct URL to access them, such as <bucket>.s3-accelerate.amazonaws.com2.
The other options are not correct because they either do not provide the lowest latency or are not suitable for the use case. Uploading and storing content in Amazon S3 and using Amazon CloudFront for the uploads is not correct because this solution is not designed for optimizing uploads, but rather for optimizing downloads. Amazon CloudFront is a content delivery network (CDN) that helps users distribute their content globally with low latency and high transfer speeds. CloudFront works by caching the content at edge locations around the world, so that users can access it quickly and easily from anywhere3. Uploading content to Amazon EC2 instances in the Region that is closest to the user and copying the data to Amazon S3 is not correct because this solution adds unnecessary complexity and cost to the process. Amazon EC2 is a computing service that provides scalable and secure virtual servers in the cloud. Users can launch, stop, or terminate EC2 instances as needed, and choose from various instance types, operating systems, and configurations4. Uploading and storing content in Amazon S3 in the Region that is closest to the user and using multiple distributions of Amazon CloudFront is not correct because this solution is not cost-effective or efficient for the use case. As mentioned above, Amazon CloudFront is a CDN that helps users distribute their content globally with low latency and high transfer speeds. However, creating multiple CloudFront distributions for each Region would incur additional charges and management overhead, and would not be necessary since 90% of the content is consumed within the same Region where it is uploaded3.
What Is Amazon Simple Storage Service? - Amazon Simple Storage Service
Amazon S3 Transfer Acceleration - Amazon Simple Storage Service
What Is Amazon CloudFront? - Amazon CloudFront
What Is Amazon EC2? - Amazon Elastic Compute Cloud

QUESTION 167
A company has applications that run on Amazon EC2 instances. The EC2 instances connect to Amazon RDS databases by using an IAM role that has associated policies. The company wants to use AWS Systems Manager to patch the EC2 instances without disrupting the running applications.
Which solution will meet these requirements?

A. Create a new IAM role. Attach the AmazonSSMManagedInstanceCore policy to the new IAM role. Attach the new IAM role to the EC2 instances and the existing IAM role.
B. Create an IAM user. Attach the AmazonSSMManagedInstanceCore policy to the IAM user. Configure Systems Manager to use the IAM user to manage the EC2 instances.
C. Enable Default Host Configuration Management in Systems Manager to manage the EC2 instances.
D. Remove the existing policies from the existing IAM role. Add the AmazonSSMManagedInstanceCore policy to the existing IAM role.

Correct Answer: C
Section:
Explanation:
The most suitable solution for the company's requirements is to enable Default Host Configuration Management in Systems Manager to manage the EC2 instances. This solution will allow the company to patch the EC2 instances without disrupting the running applications and without manually creating or modifying IAM roles or users.
Default Host Configuration Management is a feature of AWS Systems Manager that enables Systems Manager to manage EC2 instances automatically as managed instances. A managed instance is an EC2 instance that is configured for use with Systems Manager. The benefits of managing instances with Systems Manager include the following:
Connect to EC2 instances securely using Session Manager.
Perform automated patch scans using Patch Manager.
View detailed information about instances using Systems Manager Inventory.
Track and manage instances using Fleet Manager.
Keep SSM Agent up to date automatically.
Default Host Configuration Management makes it possible to manage EC2 instances without having to manually create an IAM instance profile. Instead, Default Host Configuration Management creates and applies a default IAM role to ensure that Systems Manager has permissions to manage all instances in the Region and account where it is activated. If the permissions provided are not sufficient for the use case, the default IAM role can be modified or replaced with a custom role1.

The other options are not correct because they either have more operational overhead or do not meet the requirements. Creating a new IAM role, attaching the AmazonSSMManagedInstanceCore policy to the new IAM role, and attaching the new IAM role and the existing IAM role to the EC2 instances is not correct because this solution requires manual creation and management of IAM roles, which adds complexity and cost to the solution. The AmazonSSMManagedInstanceCore policy is a managed policy that grants permissions for Systems Manager core functionality2. Creating an IAM user, attaching the AmazonSSMManagedInstanceCore policy to the IAM user, and configuring Systems Manager to use the IAM user to manage the EC2 instances is not correct because this solution requires manual creation and management of IAM users, which adds complexity and cost to the solution. An IAM user is an identity within an AWS account that has specific permissions for a single person or application3. Removing the existing policies from the existing IAM role and adding the AmazonSSMManagedInstanceCore policy to the existing IAM role is not correct because this solution may disrupt the running applications that rely on the existing policies for accessing RDS databases. An IAM role is an identity within an AWS account that has specific permissions for a service or entity4.
AWS managed policy: AmazonSSMManagedInstanceCore
IAM users
IAM roles
Default Host Management Configuration - AWS Systems Manager

QUESTION 168
A company hosts a data lake on Amazon S3. The data lake ingests data in Apache Parquet format from various data sources. The company uses multiple transformation steps to prepare the ingested data. The steps include filtering of anomalies, normalizing of data to standard date and time values, and generation of aggregates for analyses.
The company must store the transformed data in S3 buckets that data analysts access. The company needs a prebuilt solution for data transformation that does not require code. The solution must provide data lineage and data profiling. The company needs to share the data transformation steps with employees throughout the company.
Which solution will meet these requirements?

A. Configure an AWS Glue Studio visual canvas to transform the data. Share the transformation steps with employees by using AWS Glue jobs.
B. Configure Amazon EMR Serverless to transform the data. Share the transformation steps with employees by using EMR Serveriess jobs.
C. Configure AWS Glue DataBrew to transform the data. Share the transformation steps with employees by using DataBrew recipes.
D. Create Amazon Athena tables for the data. Write Athena SQL queries to transform the data. Share the Athena SQL queries with employees.

Correct Answer: C
Section:
Explanation:
The most suitable solution for the company's requirements is to configure AWS Glue DataBrew to transform the data and share the transformation steps with employees by using DataBrew recipes. This solution will provide a prebuilt solution for data transformation that does not require code, and will also provide data lineage and data profiling. The company can easily share the data transformation steps with employees throughout the company by using DataBrew recipes.
AWS Glue DataBrew is a visual data preparation tool that makes it easy for data analysts and data scientists to clean and normalize data for analytics or machine learning by up to 80% faster. Users can upload their data from various sources, such as Amazon S3, Amazon RDS, Amazon Redshift, Amazon Aurora, or Glue Data Catalog, and use a point-and-click interface to apply over 250 built-in transformations.Users can also preview the results of each transformation step and see how it affects the quality and distribution of the data1.
A DataBrew recipe is a reusable set of transformation steps that can be applied to one or more datasets. Users can create recipes from scratch or use existing ones from the DataBrew recipe library.Users can also export, import, or share recipes with other users or groups within their AWS account or organization2.
DataBrew also provides data lineage and data profiling features that help users understand and improve their data quality. Data lineage shows the source and destination of each dataset and how it is transformed by each recipe step. Data profiling shows various statistics and metrics about each dataset, such as column

QUESTION 169
A company has an application that uses an Amazon DynamoDB table for storage. A solutions architect discovers that many requests to the table are not returning the latest data. The company's users have not reported any other issues with database performance. Latency is in an acceptable range.
Which design change should the solutions architect recommend?

A. Add read replicas to the table.
B. Use a global secondary index (GSI).
C. Request strongly consistent reads for the table.
D. Request eventually consistent reads for the table.

Correct Answer: C

Section:
Explanation:
The most suitable design change for the company's application is to request strongly consistent reads for the table. This change will ensure that the requests to the table return the latest data, reflecting the updates from all prior write operations.
Amazon DynamoDB is a fully managed NoSQL database service that provides fast and predictable performance with seamless scalability. DynamoDB supports two types of read consistency: eventually consistent reads and strongly consistent reads. By default, DynamoDB uses eventually consistent reads, unless users specify otherwise1.
Eventually consistent reads are reads that may not reflect the results of a recently completed write operation. The response might not include the changes because of the latency of propagating the data to all replicas. If users repeat their read request after a short time, the response should return the updated data. Eventually consistent reads are suitable for applications that do not require up-to-date data or can tolerate eventual consistency1.
Strongly consistent reads are reads that return a result that reflects all writes that received a successful response prior to the read. Users can request a strongly consistent read by setting the ConsistentRead parameter to true in their read operations, such as GetItem, Query, or Scan. Strongly consistent reads are suitable for applications that require up-to-date data or cannot tolerate eventual consistency1.
The other options are not correct because they do not address the issue of read consistency or are not relevant for the use case. Adding read replicas to the table is not correct because this option is not supported by DynamoDB. Read replicas are copies of a primary database instance that can serve read-only traffic and improve availability and performance. Read replicas are available for some relational database services, such as Amazon RDS or Amazon Aurora, but not for DynamoDB2. Using a global secondary index (GSI) is not correct because this option is not related to read consistency. A GSI is an index that has a partition key and an optional sort key that are different from those on the base table. A GSI allows users to query the data in different ways, with eventual consistency3. Requesting eventually consistent reads for the table is not correct because this option is already the default behavior of DynamoDB and does not solve the problem of requests not returning the latest data.
Read consistency - Amazon DynamoDB
Working with read replicas - Amazon Relational Database Service
Working with global secondary indexes - Amazon DynamoDB

QUESTION 170
A company plans to migrate toAWS and use Amazon EC2 On-Demand Instances for its application. During the migration testing phase, a technical team observes that the application takes a long time to launch and load memory to become fully productive.
Which solution will reduce the launch time of the application during the next testing phase?

A. Launch two or more EC2 On-Demand Instances. Turn on auto scaling features and make the EC2 On-Demand Instances available during the next testing phase.
B. Launch EC2 Spot Instances to support the application and to scale the application so it is available during the next testing phase.
C. Launch the EC2 On-Demand Instances with hibernation turned on. Configure EC2 Auto Scaling warm pools during the next testing phase.
D. Launch EC2 On-Demand Instances with Capacity Reservations. Start additional EC2 instances during the next testing phase.

Correct Answer: C
Section:
Explanation:
The solution that will reduce the launch time of the application during the next testing phase is to launch the EC2 On-Demand Instances with hibernation turned on and configure EC2 Auto Scaling warm pools. This solution allows the application to resume from a hibernated state instead of starting from scratch, which can save time and resources. Hibernation preserves the memory (RAM) state of the EC2 instances to the root EBS volume and then stops the instances. When the instances are resumed, they restore their memory state from the EBS volume and become productive quickly. EC2 Auto Scaling warm pools can be used to maintain a pool of pre-initialized instances that are ready to scale out when needed. Warm pools can also support hibernated instances, which can further reduce the launch time and cost of scaling out.
The other solutions are not as effective as the first one because they either do not reduce the launch time, do not guarantee availability, or do not use On-Demand Instances as required. Launching two or more EC2 On-Demand Instances with auto scaling features does not reduce the launch time of the application, as each instance still has to go through the initialization process. Launching EC2 Spot Instances does not guarantee availability, as Spot Instances can be interrupted by AWS at any time when there is a higher demand for capacity. Launching EC2 On-Demand Instances with Capacity Reservations does not reduce the launch time of the application, as it only ensures that there is enough capacity available for the instances, but does not pre-initialize them.
Hibernating your instance - Amazon Elastic Compute Cloud
Warm pools for Amazon EC2 Auto Scaling - Amazon EC2 Auto Scaling

QUESTION 171
A company has deployed a multiplayer game for mobile devices. The game requires live location tracking of players based on latitude and longitude. The data store for the game must support rapid updates and retrieval of locations.
The game uses an Amazon RDS for PostgreSQL DB instance with read replicas to store the location data. During peak usage periods, the database is unable to maintain the performance that is needed for reading and writing updates. The game's user base is increasing rapidly.
What should a solutions architect do to improve the performance of the data tier?

A. Take a snapshot of the existing DB instance. Restore the snapshot with Multi-AZ enabled.
B. Migrate from Amazon RDS to Amazon OpenSearch Service with OpenSearch Dashboards.
C. Deploy Amazon DynamoDB Accelerator (DAX) in front of the existing DB instance. Modify the game to use DAX.
D. Deploy an Amazon ElastiCache for Redis cluster in front of the existing DB instance. Modify the game to use Redis.

Correct Answer: D
Section:
Explanation:
The solution that will improve the performance of the data tier is to deploy an Amazon ElastiCache for Redis cluster in front of the existing DB instance and modify the game to use Redis. This solution will enable the game to store and retrieve the location data of the players in a fast and scalable way, as Redis is an in-memory data store that supports geospatial data types and commands. By using ElastiCache for Redis, the game can reduce the load on the RDS for PostgreSQL DB instance, which is not optimized for high-frequency updates and queries of location data. ElastiCache for Redis also supports replication, sharding, and auto scaling to handle the increasing user base of the game.
The other solutions are not as effective as the first one because they either do not improve the performance, do not support geospatial data, or do not leverage caching. Taking a snapshot of the existing DB instance and restoring it with Multi-AZ enabled will not improve the performance of the data tier, as it only provides high availability and durability, but not scalability or low latency. Migrating from Amazon RDS to Amazon OpenSearch Service with OpenSearch Dashboards will not improve the performance of the data tier, as OpenSearch Service is mainly designed for full-text search and analytics, not for real-time location tracking. OpenSearch Service also does not support geospatial data types and commands natively, unlike Redis. Deploying Amazon DynamoDB Accelerator (DAX) in front of the existing DB instance and modifying the game to use DAX will not improve the performance of the data tier, as DAX is only compatible with DynamoDB, not with RDS for PostgreSQL. DAX also does not support geospatial data types and commands.
Amazon ElastiCache for Redis
Geospatial Data Support - Amazon ElastiCache for Redis
Amazon RDS for PostgreSQL
Amazon OpenSearch Service
Amazon DynamoDB Accelerator (DAX)

QUESTION 172
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 173
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 174
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 175
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 176
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 177
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 178
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 179
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 180
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 181
A robotics company is designing a solution for medical surgery The robots will use advanced sensors, cameras, and Al 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 182
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 183
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 184
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 185
A company is migrating its workloads to AWS. The company has sensitive and critical data in on-premises relational databases that run on SQL Server instances. The company wants to use the AWS Cloud to increase security and reduce operational overhead for the databases. Which solution will meet these requirements?

A. Migrate the databases to Amazon EC2 instances. Use an AWS Key Management Service (AWS KMS) AWS managed key for encryption.
B. Migrate the databases to a Multi-AZ Amazon RDS for SQL Server DB instance Use an AWS Key Management Service (AWS KMS) AWS managed key for encryption.
C. Migrate the data to an Amazon S3 bucket Use Amazon Macie to ensure data security
D. Migrate the databases to an Amazon DynamoDB table. Use Amazon CloudWatch Logs to ensure data security

Correct Answer: B
Section:
Explanation:
Understanding the Requirement: The company needs to migrate sensitive and critical data from on-premises SQL Server databases to AWS, aiming to increase security and reduce operational overhead.
Analysis of Options:
EC2 Instances with KMS: Running SQL Server on EC2 provides control but requires significant operational overhead for management, backups, patching, and high availability.
Multi-AZ Amazon RDS for SQL Server with KMS: Amazon RDS for SQL Server offers managed database services, reducing operational overhead. Multi-AZ deployment provides high availability, and KMS encryption ensures data security.
Amazon S3 and Macie: S3 is not a suitable replacement for relational databases, and Macie is used for data security and compliance but not for database operations.
Amazon DynamoDB and CloudWatch Logs: DynamoDB is a NoSQL database and does not support SQL Server workloads directly. CloudWatch Logs are used for monitoring, not for ensuring database security.
Best Solution:
Multi-AZ Amazon RDS for SQL Server with KMS: This solution meets the requirements for security, high availability, and reduced operational overhead by using a managed database service with encryption.
Amazon RDS for SQL Server
AWS Key Management Service (KMS)

QUESTION 186
A company uses 50 TB of data for reporting The company wants to move this data from on premises to AWS A custom application in the company's data center runs a weekly data transformation job The company plans to pause the application until the data transfer is complete and needs to begin the transfer process as soon as possible
The data center does not have any available network bandwidth for additional workloads. A solutions architect must transfer the data and must configure the transformation job to continue to run in the AWS Cloud.
Which solution will meet these requirements with the LEAST operational overhead?

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

Correct Answer: C
Section:
Explanation:
Understanding the Requirement: The company needs to transfer 50 TB of data to AWS with minimal operational overhead and no available network bandwidth for the transfer. The transformation job must continue running in the AWS Cloud.
Analysis of Options:
AWS DataSync and AWS Glue: DataSync is suitable for online data transfer, but there is no available network bandwidth. AWS Glue can be used for data transformation but does not solve the bandwidth issue.
AWS Snowcone: Snowcone is a smaller device suitable for smaller data transfers, and deploying the transformation application on it may not be feasible for 50 TB of data.
AWS Snowball Edge Storage Optimized with Glue: This device is designed for large data transfers. Copying the data to the device is straightforward, and AWS Glue can handle data transformation in the cloud.
AWS Snowball Edge Storage Optimized with EC2: This involves setting up EC2 instances for transformation, adding operational complexity compared to using AWS Glue.
Best Solution:
AWS Snowball Edge Storage Optimized with AWS Glue: This provides the least operational overhead for transferring large amounts of data and setting up the transformation job in the cloud.
AWS Snowball Edge
AWS Glue

QUESTION 187
A company has an on-premises business application that generates hundreds of files each day. These files are stored on an SMB file share and require a low-latency connection to the application servers. A new company policy states all application-generated files must be copied to AWS. There is already a VPN connection to AWS.
The application development team does not have time to make the necessary code modifications to move the application to AWS Which service should a solutions architect recommend to allow the application to copy files to AWS?

A. Amazon Elastic File System (Amazon EFS)
B. Amazon FSx for Windows File Server
C. AWS Snowball
D. AWS Storage Gateway

Correct Answer: D
Section:
Explanation:
Understanding the Requirement: The company needs to copy files generated by an on-premises application to AWS without modifying the application code. The files are stored on an SMB file share and require a low-latency connection to the application servers.
Analysis of Options:
Amazon Elastic File System (EFS): EFS is designed for Linux-based workloads and does not natively support SMB file shares.
Amazon FSx for Windows File Server: FSx supports SMB file shares but would require changes to the application or additional infrastructure to connect on-premises systems.
AWS Snowball: Suitable for large data transfers but not for continuous, low-latency file copying.
AWS Storage Gateway: Provides a hybrid cloud storage solution, supporting SMB file shares and enabling seamless copying of files to AWS without requiring changes to the application.
Best Solution:
AWS Storage Gateway: This service meets the requirement for a low-latency, seamless file transfer solution from on-premises to AWS without modifying the application code.
AWS Storage Gateway
Amazon FSx for Windows File Server

QUESTION 188
A company wants to migrate an application to AWS. The company wants to increase the application's current availability The company wants to use AWS WAF in the application's architecture.
Which solution will meet these requirements?

A. Create an Auto Scaling group that contains multiple Amazon EC2 instances that host the application across two Availability Zones. Configure an Application Load Balancer (ALB) and set the Auto Scaling group as the target. Connect a WAF to the ALB.
B. Create a cluster placement group that contains multiple Amazon EC2 instances that hosts the application Configure an Application Load Balancer and set the EC2 instances as the targets. Connect a WAF to the placement group.
C. Create two Amazon EC2 instances that host the application across two Availability Zones. Configure the EC2 instances as the targets of an Application Load Balancer (ALB). Connect a WAF to the ALB.
D. Create an Auto Scaling group that contains multiple Amazon EC2 instances that host the application across two Availability Zones. Configure an Application Load Balancer (ALB) and set the Auto Scaling group as the target Connect a WAF to the Auto Scaling group.

Correct Answer: A
Section:
Explanation:
Understanding the Requirement: The company wants to migrate an application to AWS, increase its availability, and use AWS WAF in the architecture.
Analysis of Options:
Auto Scaling group with ALB and WAF: This option provides high availability by distributing instances across multiple Availability Zones. The ALB ensures even traffic distribution, and AWS WAF provides security at the application layer.
Cluster placement group with ALB and WAF: Cluster placement groups are for low-latency networking within a single AZ, which does not provide the high availability across AZs.
Two EC2 instances with ALB and WAF: This setup provides some availability but does not scale automatically, missing the benefits of an Auto Scaling group.
Auto Scaling group with WAF directly: AWS WAF cannot be directly connected to an Auto Scaling group; it needs to be attached to an ALB, CloudFront distribution, or API Gateway.
Best Solution:
Auto Scaling group with ALB and WAF: This configuration ensures high availability, scalability, and security, meeting all the requirements effectively.
Amazon EC2 Auto Scaling
Application Load Balancer
AWS WAF

QUESTION 189
A company runs a stateful production application on Amazon EC2 instances The application requires at least two EC2 instances to always be running.
A solutions architect needs to design a highly available and fault-tolerant architecture for the application. The solutions architect creates an Auto Scaling group of EC2 instances.
Which set of additional steps should the solutions architect take to meet these requirements?

A. Set the Auto Scaling group's minimum capacity to two. Deploy one On-Demand Instance in one Availability Zone and one On-Demand Instance in a second Availability Zone.
B. Set the Auto Scaling group's minimum capacity to four Deploy two On-Demand Instances in one Availability Zone and two On-Demand Instances in a second Availability Zone
C. Set the Auto Scaling group's minimum capacity to two. Deploy four Spot Instances in one Availability Zone.
D. Set the Auto Scaling group's minimum capacity to four Deploy two On-Demand Instances in one Availability Zone and two Spot Instances in a second Availability Zone.

Correct Answer: A
Section:
Explanation:
Understanding the Requirement: The application is stateful and requires at least two EC2 instances to be running at all times, with a highly available and fault-tolerant architecture.
Analysis of Options:
Minimum capacity of two with instances in separate AZs: Ensures high availability by distributing instances across multiple AZs, fulfilling the requirement of always having two instances running.
Minimum capacity of four: Provides redundancy but is more than what is required and increases cost without additional benefit.
Spot Instances: Not suitable for a stateful application requiring guaranteed availability, as Spot Instances can be terminated at any time.
Combination of On-Demand and Spot Instances: Mixing instance types might provide cost savings but does not ensure the required availability for a stateful application.
Best Solution:
Minimum capacity of two with instances in separate AZs: This setup ensures high availability and meets the requirement with the least cost and complexity.
Amazon EC2 Auto Scaling
High Availability for Amazon EC2

QUESTION 190
A company manages a data lake in an Amazon S3 bucket that numerous applications access The S3 bucket contains a unique prefix for each application The company wants to restrict each application to its specific prefix and to have granular control of the objects under each prefix.
Which solution will meet these requirements with the LEAST operational overhead?

A. Create dedicated S3 access points and access point policies for each application.
B. Create an S3 Batch Operations job to set the ACL permissions for each object in the S3 bucket
C. Replicate the objects in the S3 bucket to new S3 buckets for each application. Create replication rules by prefix
D. Replicate the objects in the S3 bucket to new S3 buckets for each application Create dedicated S3 access points for each application

Correct Answer: A
Section:
Explanation:
Understanding the Requirement: The company wants to restrict each application to its specific prefix in an S3 bucket and have granular control over the objects under each prefix.
Analysis of Options:
Dedicated S3 Access Points: Provides a scalable and flexible way to manage access to S3 buckets, allowing specific policies to be attached to each access point, thereby controlling access at the prefix level.
S3 Batch Operations: Suitable for large-scale changes but involves more operational overhead and does not dynamically control future access.
Replication to new S3 buckets: Involves unnecessary duplication of data and increased storage costs, and operational overhead for managing multiple buckets.
Combination of replication and access points: Adds unnecessary complexity and overhead compared to using access points directly.
Best Solution:
Dedicated S3 Access Points: This provides the least operational overhead while meeting the requirements for prefix-level access control and granular management.
Amazon S3 Access Points

QUESTION 191
A company has released a new version of its production application The company's workload uses Amazon EC2. AWS Lambda. AWS Fargate. and Amazon SageMaker. The company wants to cost optimize the workload now that usage is at a steady state. The company wants to cover the most services with the fewest savings plans. Which combination of savings plans will meet these requirements? (Select TWO.)

A. Purchase an EC2 Instance Savings Plan for Amazon EC2 and SageMaker.
B. Purchase a Compute Savings Plan for Amazon EC2. Lambda, and SageMaker
C. Purchase a SageMaker Savings Plan
D. Purchase a Compute Savings Plan for Lambda, Fargate, and Amazon EC2
E. Purchase an EC2 Instance Savings Plan for Amazon EC2 and Fargate

Correct Answer: B, D
Section:
Explanation:
Understanding the Requirement: The company wants to cost-optimize their workload that uses EC2, Lambda, Fargate, and SageMaker, covering the most services with the fewest savings plans.
Analysis of Options:
EC2 Instance Savings Plan: Limited to EC2 and SageMaker, missing coverage for Lambda and Fargate.
Compute Savings Plan: Provides the most flexibility, covering a broad range of compute services, including EC2, Lambda, Fargate, and SageMaker.
SageMaker Savings Plan: Specifically for SageMaker, missing coverage for EC2, Lambda, and Fargate.
Combination of plans: The Compute Savings Plan is versatile and can be combined to cover different services efficiently.
Best Solution:
Compute Savings Plan for EC2, Lambda, and SageMaker: Covers the primary compute services efficiently.
Compute Savings Plan for Lambda, Fargate, and EC2: Covers the remaining services, ensuring broad coverage with minimal plans.
AWS Savings Plans
Compute Savings Plans

QUESTION 192
A company is designing an event-driven order processing system Each order requires multiple validation steps after the order is created. An independent AWS Lambda function performs each validation step. Each validation step is independent from the other validation steps Individual validation steps need only a subset of the order event information.
The company wants to ensure that each validation step Lambda function has access to only the information from the order event that the function requires The components of the order processing system should be loosely coupled to accommodate future business changes.
Which solution will meet these requirements?

A. Create an Amazon Simple Queue Service (Amazon SQS> queue for each validation step. Create a new Lambda function to transform the order data to the format that each validation step requires and to publish the messages to the appropriate SQS queues Subscribe each validation step Lambda function to its corresponding SQS queue
B. Create an Amazon Simple Notification Service {Amazon SNS) topic. Subscribe the validation step Lambda functions to the SNS topic. Use message body filtering to send only the required data to each subscribed Lambda function.
C. Create an Amazon EventBridge event bus. Create an event rule for each validation step Configure the input transformer to send only the required data to each target validation step Lambda function.
D. Create an Amazon Simple Queue Service {Amazon SQS) queue Create a new Lambda function to subscribe to the SQS queue and to transform the order data to the format that each validation step requires. Use the new Lambda function to perform synchronous invocations of the validation step Lambda functions in parallel on separate threads.

Correct Answer: C
Section:
Explanation:
Understanding the Requirement: The order processing system requires multiple independent validation steps, each handled by separate Lambda functions, with each function accessing only the subset of order information it needs. The system should be loosely coupled to accommodate future changes.
Analysis of Options:
Amazon SQS with a new Lambda function for transformation: This involves additional complexity in creating and managing multiple SQS queues and an extra Lambda function for data transformation.
Amazon SNS with message filtering: While SNS supports message filtering, it is more suited for pub/sub messaging patterns rather than event-driven processing requiring fine-grained control over the data sent to each function.
Amazon EventBridge with input transformers: EventBridge is designed for event-driven architectures, allowing for fine-grained control with input transformers that can modify and filter the event data sent to each target Lambda function, ensuring each function receives only the necessary information.
SQS with synchronous Lambda invocations: This approach adds unnecessary complexity with synchronous invocations and is not ideal for an event-driven, loosely coupled architecture.
Best Solution:
Amazon EventBridge with input transformers: This option provides the most flexible, scalable, and loosely coupled architecture, enabling each Lambda function to receive only the required subset of data.
Amazon EventBridge
EventBridge Input Transformer

QUESTION 193
A large international university has deployed all of its compute services in the AWS Cloud These services include Amazon EC2. Amazon RDS. and Amazon DynamoDB. The university currently relies on many custom scripts to back up its infrastructure. However, the university wants to centralize management and automate data backups as much as possible by using AWS native options.
Which solution will meet these requirements?

A. Use third-party backup software with an AWS Storage Gateway tape gateway virtual tape library.
B. Use AWS Backup to configure and monitor all backups for the services in use
C. Use AWS Config to set lifecycle management to take snapshots of all data sources on a schedule.
D. Use AWS Systems Manager State Manager to manage the configuration and monitoring of backup tasks.

Correct Answer: B
Section:
Explanation:
Understanding the Requirement: The university wants to centralize management and automate backups for its AWS services (EC2, RDS, and DynamoDB), reducing reliance on custom scripts.
Analysis of Options:
Third-party backup software with AWS Storage Gateway: This solution introduces external dependencies and adds complexity compared to using native AWS services.
AWS Backup: Provides a centralized, fully managed service to automate and manage backups across various AWS services, including EC2, RDS, and DynamoDB.

AWS Config: Primarily used for compliance and configuration monitoring, not for backup management.
AWS Systems Manager State Manager: Useful for configuration management but not specifically designed for managing backups.
Best Solution:
AWS Backup: This service offers the necessary functionality to centralize and automate backups, providing a streamlined and integrated solution with minimal effort.
AWS Backup

QUESTION 194
A company stores several petabytes of data across multiple AWS accounts The company uses AWS Lake Formation to manage its data lake The company's data science team wants to securely share selective data from its accounts with the company's engineering team for analytical purposes.
Which solution will meet these requirements with the LEAST operational overhead?

A. Copy the required data to a common account. Create an 1AM access role in that account Grant access by specifying a permission policy that includes users from the engineering team accounts as trusted entities.
B. Use the Lake Formation permissions Grant command in each account where the data is stored to allow the required engineering team users to access the data.
C. Use AWS Data Exchange to privately publish the required data to the required engineering team accounts
D. Use Lake Formation tag-based access control to authorize and grant cross-account permissions for the required data to the engineering team accounts

Correct Answer: D
Section:
Explanation:
Understanding the Requirement: The data science team needs to securely share selective data with the engineering team across multiple AWS accounts with minimal operational overhead.
Analysis of Options:
Copy data to a common account: Involves data duplication and increased storage costs, and requires managing additional permissions.
Lake Formation permissions Grant command: This method can be effective but may involve significant operational overhead if managing permissions across multiple accounts and datasets manually.
AWS Data Exchange: Designed for sharing data externally or between organizations, which adds unnecessary complexity for internal sharing within the same organization.
Lake Formation tag-based access control: Provides a scalable and efficient way to manage access permissions based on tags, allowing fine-grained control and simplified management across accounts.
Best Solution:
Lake Formation tag-based access control: This solution meets the requirements with the least operational overhead, allowing efficient management of cross-account permissions and secure data sharing.
AWS Lake Formation
Tag-based access control

QUESTION 195
A company stores sensitive data in Amazon S3 A solutions architect needs to create an encryption solution The company needs to fully control the ability of users to create, rotate, and disable encryption keys with minimal effort for any data that must be encrypted.
Which solution will meet these requirements?

A. Use default server-side encryption with Amazon S3 managed encryption keys (SSE-S3) to store the sensitive data
B. Create a customer managed key by using AWS Key Management Service (AWS KMS). Use the new key to encrypt the S3 objects by using server-side encryption with AWS KMS keys (SSE-KMS).
C. Create an AWS managed key by using AWS Key Management Service {AWS KMS) Use the new key to encrypt the S3 objects by using server-side encryption with AWS KMS keys (SSE-KMS).
D. Download S3 objects to an Amazon EC2 instance. Encrypt the objects by using customer managed keys. Upload the encrypted objects back into Amazon S3.

Correct Answer: B
Section:
Explanation:
Understanding the Requirement: The company needs to control the creation, rotation, and disabling of encryption keys for data stored in S3 with minimal effort.
Analysis of Options:
SSE-S3: Provides server-side encryption using S3 managed keys but does not offer full control over key management.
Customer managed key with AWS KMS (SSE-KMS): Allows the company to fully control key creation, rotation, and disabling, providing a high level of security and compliance.
AWS managed key with AWS KMS (SSE-KMS): While it provides some control, it does not offer the same level of granularity as customer-managed keys.
EC2 instance encryption and re-upload: This approach is operationally intensive and does not leverage AWS managed services for efficient key management.

Best Solution:
Customer managed key with AWS KMS (SSE-KMS): This solution meets the requirement for full control over encryption keys with minimal operational overhead, leveraging AWS managed services for secure key management.
AWS Key Management Service (KMS)
Amazon S3 Encryption

QUESTION 196
A company runs an application that uses Amazon RDS for PostgreSQL The application receives traffic only on weekdays during business hours The company wants to optimize costs and reduce operational overhead based on this usage.
Which solution will meet these requirements?

A. Use the Instance Scheduler on AWS to configure start and stop schedules.
B. Turn off automatic backups. Create weekly manual snapshots of the database.
C. Create a custom AWS Lambda function to start and stop the database based on minimum CPU utilization.
D. Purchase All Upfront reserved DB instances

Correct Answer: A
Section:
Explanation:
Understanding the Requirement: The company wants to optimize costs and reduce operational overhead for an RDS for PostgreSQL database that only needs to be active during business hours on weekdays.
Analysis of Options:
Instance Scheduler on AWS: Allows for automated start and stop schedules based on specified times, ideal for resources only needed during certain hours. This directly optimizes costs by running the database only when needed.
Turn off automatic backups and create weekly snapshots: Does not address the requirement of reducing operational overhead and optimizing runtime costs.
Custom Lambda function: This could work but adds unnecessary complexity compared to using the Instance Scheduler.
All Upfront Reserved DB Instances: While this reduces costs, it does not optimize for usage patterns that require the database only during specific hours.
Best Solution:
Instance Scheduler on AWS: This option effectively manages the database runtime based on the specified schedule, reducing costs and operational overhead.
Instance Scheduler on AWS

QUESTION 197
A company recently migrated its web application to the AWS Cloud The company uses an Amazon EC2 instance to run multiple processes to host the application. The processes include an Apache web server that serves static content The Apache web server makes requests to a PHP application that uses a local Redis server for user sessions.
The company wants to redesign the architecture to be highly available and to use AWS managed solutions Which solution will meet these requirements?

A. Use AWS Elastic Beanstalk to host the static content and the PHP application. Configure Elastic Beanstalk to deploy its EC2 instance into a public subnet Assign a public IP address.
B. Use AWS Lambda to host the static content and the PHP application. Use an Amazon API Gateway REST API to proxy requests to the Lambda function. Set the API Gateway CORS configuration to respond to the domain name. Configure Amazon ElastiCache for Redis to handle session information
C. Keep the backend code on the EC2 instance. Create an Amazon ElastiCache for Redis cluster that has Multi-AZ enabled Configure the ElastiCache for Redis cluster in cluster mode Copy the frontend resources to Amazon S3 Configure the backend code to reference the EC2 instance
D. Configure an Amazon CloudFront distribution with an Amazon S3 endpoint to an S3 bucket that is configured to host the static content. Configure an Application Load Balancer that targets an Amazon Elastic Container Service (Amazon ECS) service that runs AWS Fargate tasks for the PHP application. Configure the PHP application to use an Amazon ElastiCache for Redis cluster that runs in multiple Availability Zones

Correct Answer: D
Section:
Explanation:
Understanding the Requirement: The company needs to redesign the architecture to be highly available and use AWS managed solutions for hosting a web application with static content, PHP application, and Redis for user sessions.
Analysis of Options:
AWS Elastic Beanstalk: Suitable for simplifying deployment but may not provide the desired flexibility and control for complex architectures.

AWS Lambda and API Gateway: Not ideal for hosting a stateful PHP application and handling static content. Adding complexity without significant benefit.
EC2 instance with ElastiCache and S3: Provides some high availability but involves managing EC2 instances, which increases operational overhead.
CloudFront with S3, ALB, ECS with Fargate, and ElastiCache: This solution leverages fully managed AWS services for each component, ensuring high availability and scalability.
Best Solution:
CloudFront with S3, ALB, ECS with Fargate, and ElastiCache: This combination of services meets the requirements for a highly available and managed solution, ensuring optimal performance and minimal operational overhead.
Amazon CloudFront
Amazon S3
Amazon ECS with Fargate
Amazon ElastiCache for Redis

QUESTION 198
A company has an application that customers use to upload images to an Amazon S3 bucket Each night, the company launches an Amazon EC2 Spot Fleet that processes all the images that the company received that day. The processing for each image takes 2 minutes and requires 512 MB of memory.
A solutions architect needs to change the application to process the images when the images are uploaded
Which change will meet these requirements MOST cost-effectively?

A. Use S3 Event Notifications to write a message with image details to an Amazon Simple Queue Service (Amazon SQS) queue. Configure an AWS Lambda function to read the messages from the queue and to process the images
B. Use S3 Event Notifications to write a message with image details to an Amazon Simple Queue Service (Amazon SQS) queue Configure an EC2 Reserved Instance to read the messages from the queue and to process the images.
C. Use S3 Event Notifications to publish a message with image details to an Amazon Simple Notification Service (Amazon SNS) topic. Configure a container instance in Amazon Elastic Container Service (Amazon ECS) to subscribe to the topic and to process the images.
D. Use S3 Event Notifications to publish a message with image details to an Amazon Simple Notification Service (Amazon SNS) topic. to subscribe to the topic and to process the images.

Correct Answer: A
Section:
Explanation:
Understanding the Requirement: The company needs to process images as they are uploaded to S3 in a cost-effective manner, currently using an EC2 Spot Fleet for nightly processing.
Analysis of Options:
S3 Event Notifications to SQS and Lambda: This setup allows for event-driven processing with Lambda, which scales automatically based on the number of messages in the queue. It is cost-effective as Lambda charges are based on the compute time used.
S3 Event Notifications to SQS and EC2 Reserved Instance: Involves managing EC2 instances, which adds operational overhead and is less cost-effective.
S3 Event Notifications to SNS and ECS: More complex and potentially less cost-effective compared to using Lambda for simple processing tasks.
S3 Event Notifications to SNS: Requires additional configuration and management to process messages.
Best Solution:
S3 Event Notifications to SQS and Lambda: This option is the most cost-effective and scalable, leveraging AWS managed services with minimal operational overhead.
Amazon S3 Event Notifications
Amazon SQS
AWS Lambda

QUESTION 199
A company's software development team needs an Amazon RDS Multi-AZ cluster. The RDS cluster will serve as a backend for a desktop client that is deployed on premises. The desktop client requires direct connectivity to the RDS cluster.
The company must give the development team the ability to connect to the cluster by using the client when the team is in the office.
Which solution provides the required connectivity MOST securely?

A. Create a VPC and two public subnets. Create the RDS cluster in the public subnets. Use AWS Site-to-Site VPN with a customer gateway in the company's office.
B. Create a VPC and two private subnets. Create the RDS cluster in the private subnets. Use AWS Site-to-Site VPN with a customer gateway in the company's office.

C. Create a VPC and two private subnets. Create the RDS cluster in the private subnets. Use RDS security groups to allow the company's office IP ranges to access the cluster.
D. Create a VPC and two public subnets. Create the RDS cluster in the public subnets. Create a cluster user for each developer. Use RDS security groups to allow the users to access the cluster.

Correct Answer: B
Section:
Explanation:
Requirement Analysis: Need secure, direct connectivity from an on-premises client to an RDS cluster, accessible only when in the office.
VPC with Private Subnets: Ensures the RDS cluster is not publicly accessible, enhancing security.
Site-to-Site VPN: Provides secure, encrypted connection between on-premises office and AWS VPC.
Implementation:
Create a VPC with two private subnets.
Launch the RDS cluster in the private subnets.
Set up a Site-to-Site VPN connection with a customer gateway in the office.
Conclusion: This setup ensures secure and direct connectivity with minimal exposure, meeting the requirement for secure access from the office.
Reference
AWS Site-to-Site VPN: AWS Site-to-Site VPN Documentation
Amazon RDS: Amazon RDS Documentation

QUESTION 200
A social media company wants to store its database of user profiles, relationships, and interactions in the AWS Cloud. The company needs an application to monitor any changes in the database. The application needs to analyze the relationships between the data entities and to provide recommendations to users.
Which solution will meet these requirements with the LEAST operational overhead?

A. Use Amazon Neptune to store the information. Use Amazon Kinesis Data Streams to process changes in the database.
B. Use Amazon Neptune to store the information. Use Neptune Streams to process changes in the database.
C. Use Amazon Quantum Ledger Database (Amazon QLDB) to store the information. Use Amazon Kinesis Data Streams to process changes in the database.
D. Use Amazon Quantum Ledger Database (Amazon QLDB) to store the information. Use Neptune Streams to process changes in the database.

Correct Answer: B
Section:
Explanation:
Amazon Neptune: Neptune is a fully managed graph database service that is optimized for storing and querying highly connected data. It supports both property graph and RDF graph models, making it suitable for applications that need to analyze relationships between data entities.
Neptune Streams: Neptune Streams captures changes to the graph and streams these changes to other AWS services. This is useful for applications that need to monitor and respond to changes in real-time, such as providing recommendations based on user interactions and relationships.
Least Operational Overhead: Using Neptune Streams directly with Amazon Neptune ensures that the solution is tightly integrated, reducing the need for additional components and minimizing operational overhead. This integration simplifies the architecture by eliminating the need for a separate service like Kinesis for change processing.
Amazon Neptune Documentation
Neptune Streams Documentation

QUESTION 201
A company uses an Amazon S3 bucket as its data lake storage platform The S3 bucket contains a massive amount of data that is accessed randomly by multiple teams and hundreds of applications. The company wants to reduce the S3 storage costs and provide immediate availability for frequently accessed objects
What is the MOST operationally efficient solution that meets these requirements?

A. Create an S3 Lifecycle rule to transition objects to the S3 Intelligent-Tiering storage class
B. Store objects in Amazon S3 Glacier Use S3 Select to provide applications with access to the data.
C. Use data from S3 storage class analysis to create S3 Lifecycle rules to automatically transition objects to the S3 Standard-Infrequent Access (S3 Standard-IA) storage class.

D. Transition objects to the S3 Standard-Infrequent Access (S3 Standard-IA) storage class Create an AWS Lambda function to transition objects to the S3 Standard storage class when they are accessed by an application

Correct Answer: A
Section:
Explanation:
Amazon S3 Intelligent-Tiering: This storage class is designed to optimize costs by automatically moving data between two access tiers (frequent and infrequent) when access patterns change. It provides cost savings without performance impact or operational overhead.
S3 Lifecycle Rules: By creating an S3 Lifecycle rule, the company can automatically transition objects to the Intelligent-Tiering storage class. This eliminates the need for manual intervention and ensures that objects are moved to the most cost-effective storage tier based on their access patterns.
Operational Efficiency: Intelligent-Tiering requires no additional management and delivers immediate availability for frequently accessed objects. This makes it the most operationally efficient solution for the given requirements.
Amazon S3 Intelligent-Tiering
S3 Lifecycle Policies

QUESTION 202
A company needs to optimize its Amazon S3 storage costs for an application that generates many files that cannot be recreated Each file is approximately 5 MB and is stored in Amazon S3 Standard storage.
The company must store the files for 4 years before the files can be deleted The files must be immediately accessible The files are frequently accessed in the first 30 days of object creation, but they are rarely accessed after the first 30 days.
Which solution will meet these requirements MOST cost-effectively?

A. Create an S3 Lifecycle policy to move the files to S3 Glacier Instant Retrieval 30 days after object creation. Delete the files 4 years after object creation.
B. Create an S3 Lifecycle policy to move the files to S3 One Zone-Infrequent Access (S3 One Zone-IA) 30 days after object creation Delete the files 4 years after object creation.
C. Create an S3 Lifecycle policy to move the files to S3 Standard-Infrequent Access (S3 Standard-IA) 30 days after object creation Delete the files 4 years after object creation.
D. Create an S3 Lifecycle policy to move the files to S3 Standard-Infrequent Access (S3 Standard-IA) 30 days after object creation. Move the files to S3 Glacier Flexible Retrieval 4 years after object creation.

Correct Answer: C
Section:
Explanation:
Amazon S3 Standard-IA: This storage class is designed for data that is accessed less frequently but requires rapid access when needed. It offers lower storage costs compared to S3 Standard while still providing high availability and durability.
Access Patterns: Since the files are frequently accessed in the first 30 days and rarely accessed afterward, transitioning them to S3 Standard-IA after 30 days aligns with their access patterns and reduces storage costs significantly.
Lifecycle Policy: Implementing a lifecycle policy to transition the files to S3 Standard-IA ensures automatic management of the data lifecycle, moving files to a lower-cost storage class without manual intervention. Deleting the files after 4 years further optimizes costs by removing data that is no longer needed.
Amazon S3 Storage Classes
S3 Lifecycle Configuration

QUESTION 203
A company runs an AWS Lambda function in private subnets in a VPC. The subnets have a default route to the internet through an Amazon EC2 NAT instance. The Lambda function processes input data and saves its output as an object to Amazon S3.
Intermittently, the Lambda function times out while trying to upload the object because of saturated traffic on the NAT instance's network The company wants to access Amazon S3 without traversing the internet.
Which solution will meet these requirements?

A. Replace the EC2 NAT instance with an AWS managed NAT gateway.
B. Increase the size of the EC2 NAT instance in the VPC to a network optimized instance type
C. Provision a gateway endpoint for Amazon S3 in the VPC. Update the route tables of the subnets accordingly.
D. Provision a transit gateway. Place transit gateway attachments in the private subnets where the Lambda function is running.

Correct Answer: C

Section:
Explanation:
Gateway Endpoint for Amazon S3: A VPC endpoint for Amazon S3 allows you to privately connect your VPC to Amazon S3 without requiring an internet gateway, NAT device, VPN connection, or AWS Direct Connect connection.
Provisioning the Endpoint:
Navigate to the VPC Dashboard.
Select 'Endpoints' and create a new endpoint.
Choose the service name for S3 (com.amazonaws.region.s3).
Select the appropriate VPC and subnets.
Adjust the route tables of the subnets to include the new endpoint.
Update Route Tables: Modify the route tables of the subnets to direct traffic destined for S3 to the newly created endpoint. This ensures that traffic to S3 does not go through the NAT instance, avoiding the saturated network and eliminating timeouts.
Operational Efficiency: This solution minimizes operational overhead by removing dependency on the NAT instance and avoiding internet traffic, leading to more stable and secure S3 interactions.
VPC Endpoints for Amazon S3
Creating a Gateway Endpoint

QUESTION 204
A solutions architect is creating an application. The application will run on Amazon EC2 instances in private subnets across multiple Availability Zones in a VPC. The EC2 instances will frequently access large files that contain confidential information. These files are stored in Amazon S3 buckets for processing. The solutions architect must optimize the network architecture to minimize data transfer costs.
What should the solutions architect do to meet these requirements?

A. Create a gateway endpoint for Amazon S3 in the VPC. In the route tables for the private subnets, add an entry for the gateway endpoint
B. Create a single NAT gateway in a public subnet. In the route tables for the private subnets, add a default route that points to the NAT gateway
C. Create an AWS PrivateLink interface endpoint for Amazon S3 in the VPC. In the route tables for the private subnets, add an entry for the interface endpoint.
D. Create one NAT gateway for each Availability Zone in public subnets. In each of the route labels for the private subnets, add a default route that points lo the NAT gateway in the same Availability Zone

Correct Answer: A
Section:
Explanation:
Understanding the Requirement: The application running on EC2 instances in private subnets needs frequent access to large confidential files stored in S3, minimizing data transfer costs.
Analysis of Options:
Gateway Endpoint for S3: Provides a secure, scalable, and cost-effective way for instances in private subnets to access S3 without using the internet or NAT gateways, thus minimizing data transfer costs.
Single NAT Gateway: Incurs additional costs for data transfer through the NAT gateway, which is not cost-effective.
PrivateLink Interface Endpoint for S3: Primarily used for accessing AWS services over a private connection but is more complex and costly compared to a gateway endpoint for S3.
Multiple NAT Gateways: Increases costs significantly and adds complexity without offering the cost benefits of a gateway endpoint.
Best Solution:
Gateway Endpoint for S3: This solution provides the required access with the least data transfer costs and minimal complexity.
VPC Endpoints for Amazon S3
Gateway Endpoints

QUESTION 205
A company hosts an application on Amazon EC2 On-Demand Instances in an Auto Scaling group. Application peak hours occur at the same time each day. Application users report slow application performance at the start of peak hours. The application performs normally 2-3 hours after peak hours begin. The company wants to ensure that the application works properly at the start o* peak hours.
Which solution will meet these requirements?

A. Configure an Application Load Balancer to distribute traffic properly to the Instances.
B. Configure a dynamic scaling policy for the Auto Scaling group to launch new instances based on memory utilization
C. Configure a dynamic scaling policy for the Auto Scaling group to launch new instances based on CPU utilization.
D. Configure a scheduled scaling policy for the Auto Scaling group to launch new instances before peak hours.

Correct Answer: D
Section:
Explanation:
Understanding the Requirement: The application experiences slow performance at the start of peak hours, but normalizes after a few hours. The goal is to ensure proper performance at the beginning of peak hours.
Analysis of Options:
Application Load Balancer: Ensures proper traffic distribution but does not address the need to have sufficient instances running at the start of peak hours.
Dynamic Scaling Policy Based on Memory or CPU Utilization: While dynamic scaling reacts to usage metrics, it may not preemptively scale in anticipation of peak hours, leading to delays as new instances are launched and become available.
Scheduled Scaling Policy: This allows the Auto Scaling group to launch instances ahead of time, ensuring that enough instances are available and ready to handle the increased load right at the start of peak hours.
Best Solution:
Scheduled Scaling Policy: This approach ensures that new instances are launched and ready before peak hours begin, addressing the slow performance issue at the start of peak periods.
Scheduled Scaling for Amazon EC2 Auto Scaling

QUESTION 206
A company has a web application in the AWS Cloud and wants to collect transaction data in real time. The company wants to prevent data duplication and does not want to manage infrastructure. The company wants to perform additional processing on the data after the data is collected.
Which solution will meet these requirements?

A. Configure an Amazon Simple Queue Service (Amazon SOS) FIFO queue. Configure an AWS Lambda function with an event source mapping for the FIFO queue to process the data.
B. Configure an Amazon Simple Queue Service (Amazon SQS) FIFO queue Use an AWS Batch job to remove duplicate data from the queue Configure an AWS Lambda function to process the data.
C. Use Amazon Kinesis Data Streams to send the Incoming transaction data to an AWS Batch job that removes duplicate data. Launch an Amazon EC2 instance that runs a custom script lo process the data.
D. Set up an AWS Step Functions state machine to send incoming transaction data to an AWS Lambda function to remove duplicate data. Launch an Amazon EC2 instance that runs a custom script to process the data.

Correct Answer: A
Section:
Explanation:
Understanding the Requirement: The company needs to collect transaction data in real time, avoid data duplication, and perform additional processing without managing infrastructure.
Analysis of Options:
SQS FIFO Queue with Lambda: Ensures data is processed in order and prevents duplication. Lambda handles processing without the need to manage servers.
SQS FIFO Queue with AWS Batch: While this ensures no duplicates, it introduces additional complexity and management overhead with AWS Batch.
Kinesis Data Streams with AWS Batch and EC2: Involves more components and infrastructure management, which is against the requirement of not wanting to manage infrastructure.
Step Functions with Lambda and EC2: Involves setting up multiple services and still requires managing EC2 instances, increasing complexity.
Best Solution:
SQS FIFO Queue with Lambda: This combination ensures real-time data processing, prevents duplication, and minimizes infrastructure management, meeting all requirements efficiently.
Amazon SQS FIFO Queues
AWS Lambda and SQS Integration

QUESTION 207
A company uses AWS to host its public ecommerce website. The website uses an AWS Global Accelerator accelerator for traffic from the internet. Tt\e Global Accelerator accelerator forwards the traffic to an Application Load Balancer (ALB) that is the entry point for an Auto Scaling group.
The company recently identified a ODoS attack on the website. The company needs a solution to mitigate future attacks.
Which solution will meet these requirements with the LEAST implementation effort?

A. Configure an AWS WAF web ACL for the Global Accelerator accelerator to block traffic by using rate-based rules.
B. Configure an AWS Lambda function to read the ALB metrics to block attacks by updating a VPC network ACL.
C. Configure an AWS WAF web ACL on the ALB to block traffic by using rate-based rules.
D. Configure an Ama7on CloudFront distribution in front of the Global Accelerator accelerator

Correct Answer: A

Section:
Explanation:
Understanding the Requirement: The company needs to mitigate DDoS attacks on its website, which uses AWS Global Accelerator to route traffic to an Application Load Balancer (ALB).
Analysis of Options:
AWS WAF on Global Accelerator: Allows for centralized protection and can block traffic based on rate-based rules, effectively mitigating DDoS attacks with minimal implementation effort.
Lambda Function and VPC Network ACL: Requires custom implementation and ongoing management, increasing complexity and effort.
AWS WAF on ALB: Provides protection but involves additional configuration and management at the ALB level.
CloudFront Distribution in front of Global Accelerator: Adds unnecessary complexity and changes the current traffic flow setup.
Best Solution:
AWS WAF on Global Accelerator: This provides the required protection with the least implementation effort, ensuring effective DDoS mitigation and maintaining the existing architecture.
AWS WAF
Using AWS WAF with AWS Global Accelerator

QUESTION 208
A company runs an application on Amazon EC2 Instances in a private subnet. The application needs to store and retrieve data in Amazon S3 buckets. According to regulatory requirements, the data must not travel across the public internet.
What should a solutions architect do to meet these requirements MOST cost-effectively?

A. Deploy a NAT gateway to access the S3 buckets.
B. Deploy AWS Storage Gateway to access the S3 buckets.
C. Deploy an S3 interface endpoint to access the S3 buckets.
D. Deploy an S3 gateway endpoint to access the S3 buckets.

Correct Answer: D
Section:
Explanation:
Understanding the Requirement: The application running in a private subnet needs to store and retrieve data from S3 without data traveling over the public internet.
Analysis of Options:
NAT Gateway: Allows private subnets to access the internet but incurs additional costs and still routes traffic through the public internet.
AWS Storage Gateway: Provides hybrid cloud storage solutions but is not the most cost-effective for direct S3 access from within the VPC.
S3 Interface Endpoint: Provides private access to S3 but is generally used for specific use cases where more granular control is required, which might be overkill and more expensive.
S3 Gateway Endpoint: Provides private, cost-effective access to S3 from within the VPC without routing traffic through the public internet.
Best Solution:
S3 Gateway Endpoint: This option meets the requirements for secure, private access to S3 from a private subnet most cost-effectively.
Amazon VPC Endpoints
Gateway Endpoints

QUESTION 209
A development team uses multiple AWS accounts for its development, staging, and production environments. Team members have been launching large Amazon EC2 instances that are underutilized. A solutions architect must prevent large instances from being launched in all accounts.
How can the solutions architect meet this requirement with the LEAST operational overhead?

A. Update the 1AM policies to deny the launch of large EC2 instances. Apply the policies to all users.
B. Define a resource in AWS Resource Access Manager that prevents the launch of large EC2 instances.
C. Create an (AM role in each account that denies the launch of large EC2 instances. Grant the developers 1AM group access to the role.
D. Create an organization in AWS Organizations in the management account with the default policy. Create a service control policy (SCP) that denies the launch of large EC2 Instances, and apply it to the AWS accounts.

Correct Answer: D
Section:

Explanation:
Understanding the Requirement: The development team needs to prevent the launch of large EC2 instances across multiple AWS accounts used for development, staging, and production environments.
Analysis of Options:
IAM Policies: Would need to be applied individually to each user in every account, leading to significant operational overhead.
AWS Resource Access Manager: Used for sharing resources, not for enforcing restrictions on resource creation.
IAM Role in Each Account: Requires creating and managing roles in each account, leading to higher operational overhead compared to using a centralized approach.
Service Control Policy (SCP) with AWS Organizations: Provides a centralized way to enforce policies across multiple AWS accounts, ensuring that large EC2 instances cannot be launched in any account.
Best Solution:
Service Control Policy (SCP) with AWS Organizations: This solution offers the least operational overhead by allowing centralized management and enforcement of policies across all accounts, effectively preventing the launch of large EC2 instances.
AWS Organizations and SCPs

QUESTION 210
A company is developing an application to support customer demands. The company wants to deploy the application on multiple Amazon EC2 Nitro-based instances within the same Availability Zone. The company also wants to give the application the ability to write to multiple block storage volumes in multiple EC2 Nitro-based instances simultaneously to achieve higher application availability.
Which solution will meet these requirements?

A. Use General Purpose SSD (gp3) EBS volumes with Amazon Elastic Block Store (Amazon EBS) Multi-Attach.
B. Use Throughput Optimized HDD (st1) EBS volumes with Amazon Elastic Block Store (Amazon EBS) Multi-Attach
C. Use Provisioned IOPS SSD (io2) EBS volumes with Amazon Elastic Block Store (Amazon EBS) Multi-Attach.
D. Use General Purpose SSD (gp2) EBS volumes with Amazon Elastic Block Store (Amazon E8S) Multi-Attach.

Correct Answer: C
Section:
Explanation:
Understanding the Requirement: The application needs to write to multiple block storage volumes in multiple EC2 Nitro-based instances simultaneously to achieve higher availability.
Analysis of Options:
General Purpose SSD (gp3) with Multi-Attach: Supports Multi-Attach but does not provide the highest performance required for critical applications.
Throughput Optimized HDD (st1) with Multi-Attach: Not suitable for applications requiring high performance and low latency.
Provisioned IOPS SSD (io2) with Multi-Attach: Provides high performance and durability, suitable for applications requiring simultaneous writes and high availability.
General Purpose SSD (gp2) with Multi-Attach: Similar to gp3 but with less flexibility and performance.
Best Solution:
Provisioned IOPS SSD (io2) with Multi-Attach: This solution ensures the highest performance and availability for the application by allowing multiple EC2 instances to attach to and write to the same EBS volume simultaneously.
Amazon EBS Multi-Attach
Provisioned IOPS SSD (io2)

QUESTION 211
An online photo-sharing company stores Hs photos in an Amazon S3 bucket that exists in the us-west-1 Region. The company needs to store a copy of all new photos in the us-east-1 Region.
Which solution will meet this requirement with the LEAST operational effort?

A. Create a second S3 bucket in us-east-1. Use S3 Cross-Region Replication to copy photos from the existing S3 bucket to the second S3 bucket.
B. Create a cross-origin resource sharing (CORS) configuration of the existing S3 bucket. Specify us-east-1 in the CORS rule's AllowedOngm element.
C. Create a second S3 bucket in us-east-1 across multiple Availability Zones. Create an S3 Lifecycle rule to save photos into the second S3 bucket,
D. Create a second S3 bucket In us-east-1. Configure S3 event notifications on object creation and update events to Invoke an AWS Lambda function to copy photos from the existing S3 bucket to the second S3 bucket.

Correct Answer: A
Section:
Explanation:

Understanding the Requirement: The company needs to store a copy of all new photos in the us-east-1 Region from an S3 bucket in the us-west-1 Region.
Analysis of Options:
Cross-Region Replication: Automatically replicates objects across regions with minimal operational effort once configured.
CORS Configuration: Used for allowing resources on a web page to be requested from another domain, not for replication.
S3 Lifecycle Rule: Manages the transition of objects between storage classes within the same bucket, not for cross-region replication.
S3 Event Notifications with Lambda: Requires additional configuration and management compared to Cross-Region Replication.
Best Solution:
S3 Cross-Region Replication: This solution provides an automated and efficient way to replicate objects to another region, meeting the requirement with the least operational effort.
Amazon S3 Cross-Region Replication

QUESTION 212
A company wants to build a logging solution for its multiple AWS accounts. The company currently stores the logs from all accounts in a centralized account. The company has created an Amazon S3 bucket in the centralized account to store the VPC flow logs and AWS CloudTrail logs. All logs must be highly available for 30 days for frequent analysis, retained tor an additional 60 days tor backup purposes, and deleted 90 days after creation.
Which solution will meet these requirements MOST cost-effectively?

A. Transition objects to the S3 Standard storage class 30 days after creation. Write an expiration action that directs Amazon S3 to delete objects after 90 days.
B. Transition objects lo the S3 Standard-Infrequent Access (S3 Standard-IA) storage class 30 days after creation Move all objects to the S3 Glacier Flexible Retrieval storage class after 90 days. Write an expiration action that directs Amazon S3 to delete objects after 90 days.
C. Transition objects to the S3 Glacier Flexible Retrieval storage class 30 days after creation. Write an expiration action that directs Amazon S3 to delete objects alter 90 days.
D. Transition objects to the S3 One Zone-Infrequent Access (S3 One Zone-IA) storage class 30 days after creation. Move all objects to the S3 Glacier Flexible Retrieval storage class after 90 days. Write an expiration action that directs Amazon S3 to delete objects after 90 days.

Correct Answer: D
Section:
Explanation:
Understanding the Requirement: The company needs logs to be highly available for 30 days for frequent analysis, retained for an additional 60 days for backup, and deleted after 90 days.
Analysis of Options:
Transition to S3 Standard after 30 days: Keeps logs in the same high-availability storage, not cost-effective.
Transition to S3 Standard-IA, then Glacier Flexible Retrieval after 90 days: Adds unnecessary cost and complexity since objects need to be accessible for only 30 days and then retained for 60 days.
Transition to Glacier Flexible Retrieval after 30 days: Not suitable for frequent access required in the first 30 days.
Transition to S3 One Zone-IA after 30 days, then Glacier Flexible Retrieval: Provides cost-effective storage for infrequently accessed logs after the initial 30-day period, then moves to the cheapest long-term storage before deletion.
Best Solution:
Transition to S3 One Zone-IA after 30 days, then Glacier Flexible Retrieval: This solution meets the requirements for high availability, cost-effective storage for backup, and scheduled deletion with the least cost.
Amazon S3 Storage Classes
Managing your storage lifecycle

QUESTION 213
A company runs an application in a VPC with public and private subnets. The VPC extends across multiple Availability Zones. The application runs on Amazon EC2 instances in private subnets. The application uses an Amazon Simple Queue Service (Amazon SOS) queue.
A solutions architect needs to design a secure solution to establish a connection between the EC2 instances and the SOS queue
Which solution will meet these requirements?

A. Implement an interface VPC endpoint tor Amazon SOS. Configure the endpoint to use the private subnets. Add to the endpoint a security group that has an inbound access rule that allows traffic from the EC2 instances that are in the private subnets.
B. Implement an interface VPC endpoint tor Amazon SOS. Configure the endpoint to use the public subnets. Attach to the interface endpoint a VPC endpoint policy that allows access from the EC2 Instances that are in the private subnets.
C. Implement an interface VPC endpoint for Ama7on SOS. Configure the endpoint to use the public subnets Attach an Amazon SOS access policy to the interface VPC endpoint that allows requests from only a specified VPC endpoint.

D. Implement a gateway endpoint tor Amazon SOS. Add a NAT gateway to the private subnets. Attach an 1AM role to the EC2 Instances that allows access to the SOS queue.

Correct Answer: A
Section:
Explanation:
Understanding the Requirement: The application running on EC2 instances in private subnets needs to securely connect to an Amazon SQS queue without exposing traffic to the public internet.
Analysis of Options:
Interface VPC Endpoint in Private Subnets: Allows private, secure connectivity to SQS without using the public internet. Configuring security groups ensures controlled access from EC2 instances.
Interface VPC Endpoint in Public Subnets: Not necessary for private EC2 instances and exposes additional security risks.
Gateway Endpoint: Gateway endpoints are not supported for SQS; they are used for services like S3 and DynamoDB.
NAT Gateway with IAM Role: Increases costs and complexity compared to using an interface VPC endpoint directly.
Best Solution:
Interface VPC Endpoint in Private Subnets: This option ensures secure, private connectivity to SQS, meeting the requirement with minimal complexity and optimal security.
VPC Endpoints
Amazon SQS and VPC Endpoints

QUESTION 214
A company deploys Amazon EC2 instances that run in a VPC. The EC2 instances load source data into Amazon S3 buckets so that the data can be processed in the future. According to compliance laws, the data must not be transmitted over the public internet. Servers in the company's on-premises data center will consume the output from an application that runs on the LC2 instances.
Which solution will meet these requirements?

A. Deploy an interface VPC endpoint for Amazon EC2. Create an AWS Site-to-Site VPN connection between the company and the VPC.
B. Deploys gateway VPC endpoint for Amazon S3 Set up an AWS Direct Connect connection between the on-premises network and the VPC.
C. Set up on AWS Transit Gateway connection from the VPC to the S3 buckets. Create an AWS Site-to-Site VPN connection between the company and the VPC.
D. Set up proxy EC2 instances that have routes to NAT gateways. Configure the proxy EC2 instances lo fetch S3 data and feed the application instances.

Correct Answer: B
Section:
Explanation:
Understanding the Requirement: EC2 instances need to upload data to S3 without using the public internet, and on-premises servers need to consume this data.
Analysis of Options:
Interface VPC Endpoint for EC2: Not relevant for accessing S3.
Gateway VPC Endpoint for S3 and Direct Connect: Provides private connectivity from EC2 instances to S3 and from on-premises to AWS, ensuring compliance with the requirement to avoid public internet.
Transit Gateway and Site-to-Site VPN: Adds unnecessary complexity and does not provide the same level of performance as Direct Connect.
Proxy EC2 Instances with NAT Gateways: Increases complexity and costs compared to a direct connection using VPC endpoints and Direct Connect.
Best Solution:
Gateway VPC Endpoint for S3 and Direct Connect: This solution ensures secure, private data transfer both within AWS and between on-premises and AWS, meeting the compliance requirements effectively.
Amazon VPC Endpoints for S3
AWS Direct Connect