Amazon.SAA-C03.vJan-2024.by.Racrene.225q

Number: SAA-C03
Passing Score: 800
Time Limit: 120
File Version: 23.0

Exam Code: SAA-C03

Exam Name: AWS Certified Solutions Architect – Associate

Exam A

QUESTION 1

A company has an automobile sales website that stores its listings in a database on Amazon RDS When an automobile is sold the listing needs to be removed from the website and the data must be sent to multiple target systems. Which design should a solutions architect recommend?

A. Create an AWS Lambda function triggered when the database on Amazon RDS is updated to send the information to an Amazon Simple Queue Service (Amazon SQS> queue for the targets to consume

B. Create an AWS Lambda function triggered when the database on Amazon RDS is updated to send the information to an Amazon Simple Queue Service (Amazon SQS) FIFO queue for the targets to consume

C. Subscribe to an RDS event notification and send an Amazon Simple Queue Service (Amazon SQS) queue fanned out to multiple Amazon Simple Notification Service (Amazon SNS) topics Use AWS Lambda functions to update the targets

D. Subscribe to an RDS event notification and send an Amazon Simple Notification Service (Amazon SNS) topic fanned out to multiple Amazon Simple Queue Service (Amazon SQS) queues Use AWS Lambda functions to update the targets

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

QUESTION 2

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

A. Create an S3 Glacier vault Apply a write-once, read-many (WORM) vault lock policy to the objects

B. Create an S3 bucket with S3 Object Lock enabled Enable versioning Set a retention period of 100 years Use governance mode as the S3 bucket's default retention mode for new objects

C. Create an S3 bucket Use AWS CloudTrail to (rack any S3 API events that modify the objects Upon notification, restore the modified objects from any backup versions that the company has

D. Create an S3 bucket with S3 Object Lock enabled Enable versioning Add a legal hold to the objects Add the s3 PutObjectLegalHold permission to the IAM policies of users who need to delete the objects

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

QUESTION 3

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

A. Configure the application to upload images to S3 Glacier.

B. Configure the web server to upload the original images to Amazon S3.

C. Configure the application to upload images directly from each user's browser to Amazon S3 through the use of a presigned URL.

D. Configure S3 Event Notifications to invoke an AWS Lambda function when an image is uploaded.
Use the function to resize the image

E. Create an Amazon EventBridge (Amazon CloudWatch Events) rule that invokes an AWS Lambda function on a schedule to resize uploaded images.

Correct Answer: B, D
Section:
Explanation:

QUESTION 4

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

A. Add a second ActiveMQ server to another Availably Zone Add an additional consumer EC2 instance in another Availability Zone. Replicate the MySQL database to another Availability Zone.

B. Use Amazon MO with active/standby brokers configured across two Availability Zones Add an additional consumer EC2 instance in another Availability Zone. Replicate the MySQL database to another Availability Zone.

C. Use Amazon MO with active/standby blotters configured across two Availability Zones. Add an additional consumer EC2 instance in another Availability Zone. Use Amazon ROS tor MySQL with Multi-AZ enabled.

D. Use Amazon MQ with active/standby brokers configured across two Availability Zones Add an Auto Scaling group for the consumer EC2 instances across two Availability Zones. Use Amazon RDS for MySQL with Multi-AZ enabled.

Correct Answer: D
Section:
Explanation:
Amazon S3 is a highly scalable and durable object storage service that can store and retrieve anyamount of data from anywhere on the web1. Users can configure the application to upload images directly from each user's browser to Amazon S3 through the use of a presigned URL. A presigned URLis a URL that gives access to an object in an S3 bucket for a limited time and with a specific action,such as uploading an object2. Users can generate a presigned URL programmatically using the AWS SDKs or AWS CLI. By using a presigned URL, users can reduce coupling within the application and improve website performance, as they do not need to send the images to the web server first.AWS Lambda is a serverless compute service that runs code in response to events and automaticallymanages the underlying compute resources3. Users can configure S3 Event Notifications to invoke an AWS Lambda function when an image is uploaded. S3 Event Notifications is a feature that allows users to receive notifications when certain events happen in an S3 bucket, such as object creation or deletion. Users can configure S3 Event Notifications to invoke a Lambda function that resizes the image and stores it back in the same or a different S3 bucket. This way, users can offload the image resizing task from the web server to Lambda.

QUESTION 5

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

A. Use AWS Fargate on Amazon Elastic Container Service (Amazon ECS) to run the containerized web application with Service Auto Scaling. Use an Application Load Balancer to distribute the incoming requests.

B. Use two Amazon EC2 instances to host the containerized web application. Use an Application Load Balancer to distribute the incoming requests

C. Use AWS Lambda with a new code that uses one of the supported languages. Create multiple Lambda functions to support the load. Use Amazon API Gateway as an entry point to the Lambda functions.

D. Use a high performance computing (HPC) solution such as AWS ParallelClusterto establish an HPC cluster that can process the incoming requests at the appropriate scale.

Correct Answer: A
Section:
Explanation:
AWS Fargate is a serverless compute engine that lets users run containers without having to manage servers or clusters of Amazon EC2 instances1. Users can use AWS Fargate on Amazon Elastic Container Service (Amazon ECS) to run the containerized web application with Service Auto
Scaling. Amazon ECS is a fully managed container orchestration service that supports both Docker and Kubernetes2. Service Auto Scaling is a feature that allows users to adjust the desired number of tasks in an ECS service based on CloudWatch metrics, such as CPU utilization or request count3.Users can use AWS Fargate on Amazon ECS to migrate the application to AWS with minimum code changes and minimum development effort, as they only need to package their application in containers and specify the CPU and memory requirements.Users can also use an Application Load Balancer to distribute the incoming requests. An Application Load Balancer is a load balancer that operates at the application layer and routes traffic to targets based on the content of the request. Users can register their ECS tasks as targets for an Application Load Balancer and configure listener rules to route requests to different target groups based on path or host headers. Users can use an Application Load Balancer to improve the availability and performance of their web application.

QUESTION 6
A company is implementing a shared storage solution for a media application that is hosted m the AWS Cloud The company needs the ability to use SMB clients to access data The solution must he fully managed. Which AWS solution meets these requirements?

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

Correct Answer: D
Section:
Explanation:
Amazon FSx has native support for Windows file system features and for the industry-standard Server Message Block (SMB) protocol to access file storage over a network.https://docs.aws.amazon.com/fsx/latest/WindowsGuide/what- is.html

QUESTION 7
A company's containerized application runs on an Amazon EC2 instance. The application needs to download security certificates before it can communicate with other business applications. The company wants a highly secure solution to encrypt and decrypt the certificates in near real time. The solution also needs to store data in highly available storage after the data is encrypted. Which solution will meet these requirements with the LEAST operational overhead?

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

Correct Answer: C
Section:

QUESTION 8
A solutions architect is designing a VPC with public and private subnets. The VPC and subnets use IPv4 CIDR blocks. There is one public subnet and one private subnet in each of three Availability Zones (AZs) for high availability. An internet gateway is used to provide internet access for the public subnets. The private subnets require access to the internet to allow Amazon EC2 instances to download software updates. What should the solutions architect do to enable Internet access for the private subnets?

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

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

QUESTION 9
A company wants to migrate an on-premises data center to AWS. The data canter hosts an SFTP server that stores its data on an NFS-based file system. The server holds 200 GB of data that needs to be transferred. The server must be hosted on an Amazon EC2 instance that uses an Amazon Elastic File System (Amazon EFS) file system When combination of steps should a solutions architect take to automate this task? (Select TWO)

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

Correct Answer: B, E
Section:
Explanation:
AWS DataSync is an online data movement and discovery service that simplifies data migration and helps users quickly, easily, and securely move their file or object data to, from, and between AWS

QUESTION 10
A company has an AWS Glue extract. transform, and load (ETL) job that runs every day at the same time. The job processes XML data that is in an Amazon S3 bucket. New data is added to the S3 bucket every day. A solutions architect notices that AWS Glue is processing all the data during each run. What should the solutions architect do to prevent AWS Glue from reprocessing old data?

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

Correct Answer: A
Section:
Explanation:

QUESTION 11
A solutions architect must design a highly available infrastructure for a website. The website is powered by Windows web servers that run on Amazon EC2 instances. The solutions architect must implement a solution that can mitigate a large-scale DDoS attack that originates from thousands of IP addresses. Downtime is not acceptable for the website. Which actions should the solutions architect take to protect the website from such an attack? (Select TWO.)

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

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

QUESTION 12
A company is preparing to deploy a new serverless workload. A solutions architect must use the principle of least privilege to configure permissions that will be used to run an AWS Lambda function. An Amazon EventBridge (Amazon CloudWatch Events) rule will invoke the function.
Which solution meets these requirements?

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

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

QUESTION 13
A company is preparing to store confidential data in Amazon S3 For compliance reasons the data must be encrypted at rest Encryption key usage must be logged tor auditing purposes. Keys must be rotated every year. Which solution meets these requirements and «the MOST operationally efferent?

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

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

QUESTION 14
A bicycle sharing company is developing a multi-tier architecture to track the location of its bicycles during peak operating hours The company wants to use these data points in its existing analytics platform A solutions architect must determine the most viable multi-tier option to support this architecture The data points must be accessible from the REST API. Which action meets these requirements for storing and retrieving location data?

A. Use Amazon Athena with Amazon S3
B. Use Amazon API Gateway with AWS Lambda
C. Use Amazon QuickSight with Amazon Redshift.
D. Use Amazon API Gateway with Amazon Kinesis Data Analytics

Correct Answer: B
Section:
Explanation:

QUESTION 15
A company needs to retain application logs files for a critical application for 10 years. The application team regularly accesses logs from the past month for troubleshooting, but logs older than 1 month are rarely accessed. The application generates more than 10 TB of logs per month.
Which storage option meets these requirements MOST cost-effectively?

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

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

C. Store the logs in Amazon CloudWatch Logs Use AWS Backup to move logs more then 1 month old to S3 Glacier Deep Archive

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

Correct Answer: B
Section:
Explanation:
You need S3 to be able to archive the logs after one month. Cannot do that with CloudWatch Logs.

QUESTION 16

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

A. Configure the Lambda function (or deployment across multiple Availability Zones

B. Modify me Lambda functions configuration to increase the CPU and memory allocations tor the (unction

C. Configure the SNS topic's retry strategy to increase both the number of retries and the wait time between retries

D. Configure an Amazon Simple Queue Service (Amazon SQS) queue as the on failure destination Modify the Lambda function to process messages in the queue

Correct Answer: D
Section:
Explanation:
https://docs.aws.amazon.com/sns/latest/dg/sns-dead-letter-queues.html

QUESTION 17

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

A. Create an Amazon Simple Queue Service (Amazon SQS) FIFO queue to hold messages Set up an AWS Lambda function to process messages from the queue

B. Create an Amazon Simple Notification Service (Amazon SNS) topic to deliver notifications containing payloads to process Configure an AWS Lambda function as a subscriber.

C. Create an Amazon Simple Queue Service (Amazon SQS) standard queue to hold messages. Set up an AWS Lambda function to process messages from the queue independently

D. Create an Amazon Simple Notification Service (Amazon SNS) topic to deliver notifications containing payloads to process. Configure an Amazon Simple Queue Service (Amazon SQS) queue as a subscriber.

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

QUESTION 18

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

What should the solutions architect do to meet these requirements?

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

Correct Answer: A
Section:

QUESTION 19
A company wants to migrate its on-premises data center to AWS. According to the company's compliance requirements, the company can use only the ap-northeast-3 Region. Company administrators are not permitted to connect VPCs to the internet.
Which solutions will meet these requirements? (Choose two.)

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

Correct Answer: A, C
Section:

QUESTION 20
A company uses a three-tier web application to provide training to new employees. The application is accessed for only 12 hours every day. The company is using an Amazon RDS for MySQL DB instance to store information and wants to minimize costs.
What should a solutions architect do to meet these requirements?

A. Configure an IAM policy for AWS Systems Manager Session Manager. Create an IAM role for the policy. Update the trust relationship of the role. Set up automatic start and stop for the DB instance.
B. Create an Amazon ElastiCache for Redis cache cluster that gives users the ability to access the data from the cache when the DB instance is stopped. Invalidate the cache after the DB instance is started.
C. Launch an Amazon EC2 instance. Create an IAM role that grants access to Amazon RDS. Attach the role to the EC2 instance. Configure a cron job to start and stop the EC2 instance on the desired schedule.
D. Create AWS Lambda functions to start and stop the DB instance. Create Amazon EventBridge (Amazon CloudWatch Events) scheduled rules to invoke the Lambda functions. Configure the Lambda functions as event targets for the rules

Correct Answer: D
Section:
Explanation:

QUESTION 21
A company sells ringtones created from clips of popular songs. The files containing the ringtones are stored in Amazon S3 Standard and are at least 128 KB in size. The company has millions of files, but downloads are infrequent for ringtones older than 90 days. The company needs to save money on storage while keeping the most accessed files readily available for its users. Which action should the company take to meet these requirements MOST cost-effectively?

A. Configure S3 Standard-Infrequent Access (S3 Standard-IA) storage for the initial storage tier of the objects.
B. Move the files to S3 Intelligent-Tiering and configure it to move objects to a less expensive storage tier after 90 days.
C. Configure S3 inventory to manage objects and move them to S3 Standard-Infrequent Access (S3 Standard-1A) after 90 days.

D. Implement an S3 Lifecycle policy that moves the objects from S3 Standard to S3 Standard- Infrequent Access (S3 Standard-1A) after 90 days.

Correct Answer: D
Section:
Explanation:
This solution meets the requirements of saving money on storage while keeping the most accessed files readily available for the users. S3 Lifecycle policy can automatically move objects from one storage class to another based on predefined rules. S3 Standard-IA is a lower-cost storage class for data that is accessed less frequently, but requires rapid access when needed. It is suitable for ringtones older than 90 days that are downloaded infrequently.

QUESTION 22
A company needs to save the results from a medical trial to an Amazon S3 repository. The repository must allow a few scientists to add new files and must restrict all other users to read-only access. No users can have the ability to modify or delete any files in the repository. The company must keep every file in the repository for a minimum of 1 year after its creation date. Which solution will meet these requirements?

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

Correct Answer: B
Section:
Explanation:
n compliance mode, a protected object version can't be overwritten or deleted by any user, including the root user in your AWS account. When an object is locked in compliance mode, its retention mode can't be changed, and its retention period can't be shortened. Compliance mode helps ensure that an object version can't be overwritten or deleted for the duration of the retention period. In governance mode, users can't overwrite or delete an object version or alter its lock settings unless they have special permissions. With governance mode, you protect objects against being deleted by most users, but you can still grant some users permission to alter the retention settings or delete the object if necessary. In Governance mode, Objects can be deleted by some users with special permissions, this is against the requirement.Compliance:- Object versions can't be overwritten or deleted by any user, including the root user- Objects retention modes can't be changed, and retention periods can't be shortenedGovernance:

QUESTION 23
A large media company hosts a web application on AWS. The company wants to start caching confidential media files so that users around the world will have reliable access to the files. The content is stored in Amazon S3 buckets. The company must deliver the content quickly, regardless of where the requests originate geographically.
Which solution will meet these requirements?

A. Use AWS DataSync to connect the S3 buckets to the web application.
B. Deploy AWS Global Accelerator to connect the S3 buckets to the web application.
C. Deploy Amazon CloudFront to connect the S3 buckets to CloudFront edge servers.
D. Use Amazon Simple Queue Service (Amazon SQS) to connect the S3 buckets to the web application.

Correct Answer: C
Section:
Explanation:
CloudFront uses a local cache to provide the response, AWS Global accelerator proxies requests and connects to the application all the time for the response.
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/private-contentrestricting- access-to-s3.html#private-content-granting-permissions-to-oai

QUESTION 24
A company produces batch data that comes from different databases. The company also produces live stream data from network sensors and application APIs. The company needs to consolidate all the data into one place for business analytics. The company needs to process the incoming data and then stage the data in different Amazon S3 buckets. Teams will later run one-time queries and import the data into a business intelligence tool to show key performance indicators (KPIs).
Which combination of steps will meet these requirements with the LEAST operational overhead?
(Choose two.)

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

Correct Answer: A, E
Section:
Explanation:

QUESTION 25
A gaming company has a web application that displays scores. The application runs on Amazon EC2 instances behind an Application Load Balancer. The application stores data in an Amazon RDS for MySQL database. Users are starting to experience long delays and interruptions that are caused by database read performance. The company wants to improve the user experience while minimizing changes to the application's architecture. What should a solutions architect do to meet these requirements?

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

Correct Answer: B
Section:
Explanation:

QUESTION 26
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 D. 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:

QUESTION 27
A company has created an image analysis application in which users can upload photos and add photo frames to their images. The users upload images and metadata to indicate which photo frames they want to add to their images. The application uses a single Amazon EC2 instance and Amazon DynamoDB to store the metadata.
The application is becoming more popular, and the number of users is increasing. The company expects the number of concurrent users to vary significantly depending on the time of day and day of week. The company must ensure that the application can scale to meet the needs of the growing user base.
Which solution meats these requirements?

A. Use AWS Lambda to process the photos. Store the photos and metadata in DynamoDB.

B. Use Amazon Kinesis Data Firehose to process the photos and to store the photos and metadata.

C. Use AWS Lambda to process the photos. Store the photos in Amazon S3. Retain DynamoDB to store the metadata.

D. Increase the number of EC2 instances to three. Use Provisioned IOPS SSD (io2) Amazon Elastic Block Store (Amazon EBS) volumes to store the photos and metadata.

Correct Answer: C
Section:
Explanation:
https://www.quora.com/How-can-I-use-DynamoDB-for-storing-metadata-for-Amazon-S3-objectsThis solution meets the requirements of scalability, performance, and availability. AWS Lambda can process the photos in parallel and scale up or down automatically depending on the demand.Amazon S3 can store the photos and metadata reliably and durably, and provide high availability and low latency. DynamoDB can store the metadata efficiently and provide consistent performance. This solution also reduces the cost and complexity of managing EC2 instances and EBS volumes.Option A is incorrect because storing the photos in DynamoDB is not a good practice, as it can increase the storage cost and limit the throughput. Option B is incorrect because Kinesis Data Firehose is not designed for processing photos, but for streaming data to destinations such as S3 or Redshift. Option D is incorrect because increasing the number of EC2 instances and using Provisioned IOPS SSD volumes does not guarantee scalability, as it depends on the load balancer and the application code. It also increases the cost and complexity of managing the infrastructure.Reference:https://aws.amazon.com/certification/certified-solutions-architect-professional/ https://www.examtopics.com/discussions/amazon/view/7193-exam-aws-certified-solutions- architect-professional-topic-1/ https://aws.amazon.com/architecture/

QUESTION 28
A medical records company is hosting an application on Amazon EC2 instances. The application processes customer data files that are stored on Amazon S3. The EC2 instances are hosted in public subnets. The EC2 instances access Amazon S3 over the internet, but they do not require any other network access.
A new requirement mandates that the network traffic for file transfers take a private route and not be sent over the internet. Which change to the network architecture should a solutions architect recommend to meet this requirement?

A. Create a NAT gateway. Configure the route table for the public subnets to send traffic to Amazon S3 through the NAT gateway.

B. Configure the security group for the EC2 instances to restrict outbound traffic so that only traffic to the S3 prefix list is permitted.

C. Move the EC2 instances to private subnets. Create a VPC endpoint for Amazon S3, and link the endpoint to the route table for the private subnets

D. Remove the internet gateway from the VPC. Set up an AWS Direct Connect connection, and route traffic to Amazon S3 over the Direct Connect connection.

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

QUESTION 29
A company uses a popular content management system (CMS) for its corporate website. However, the required patching and maintenance are burdensome. The company is redesigning its website and wants anew solution. The website will be updated four times a year and does not need to have any dynamic content available. The solution must provide high scalability and enhanced security. Which combination of changes will meet these requirements with the LEAST operational overhead?
(Choose two.)

A. Deploy an AWS WAF web ACL in front of the website to provide HTTPS functionality

B. Create and deploy an AWS Lambda function to manage and serve the website content

C. Create the new website and an Amazon S3 bucket Deploy the website on the S3 bucket with static website hosting enabled

D. Create the new website. Deploy the website by using an Auto Scaling group of Amazon EC2 instances behind an Application Load Balancer.

Correct Answer: A, D
Section:

Explanation:

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

QUESTION 30

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

A. Configure a CloudWatch Logs subscription to stream the logs to Amazon OpenSearch Service (Amazon Elasticsearch Service).

B. Create an AWS Lambda function. Use the log group to invoke the function to write the logs to Amazon OpenSearch Service (Amazon Elasticsearch Service).

C. Create an Amazon Kinesis Data Firehose delivery stream Configure the log group as the delivery stream's source. Configure Amazon OpenSearch Service (Amazon Elasticsearch Service) as the delivery stream's destination.

D. Install and configure Amazon Kinesis Agent on each application server to deliver the logs to Amazon Kinesis Data Streams. Configure Kinesis Data Streams to deliver the logs to Amazon OpenSearch Service (Amazon Elasticsearch Service)

Correct Answer: A
Section:
Explanation:

QUESTION 31

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

A. Migrate the database to an Amazon RDS for MySQL DB instance. Create an AMI of the web application. Use the AMI to launch a second EC2 On-Demand Instance. Use an Application Load Balancer to distribute the load to each EC2 instance.

B. Migrate the database to an Amazon RDS for MySQL DB instance. Create an AMI of the web application. Use the AMI to launch a second EC2 On-Demand Instance. Use Amazon Route 53 weighted routing to distribute the load across the two EC2 instances.

C. Migrate the database to an Amazon Aurora MySQL DB instance. Create an AWS Lambda function to stop the EC2 instance and change the instance type. Create an Amazon CloudWatch alarm to invoke the Lambda function when CPU utilization surpasses 75%.

D. Migrate the database to an Amazon Aurora MySQL DB instance. Create an AMI of the web application. Apply the AMI to a launch template. Create an Auto Scaling group with the launch template Configure the launch template to use a Spot Fleet. Attach an Application Load Balancer to the Auto Scaling group.

Correct Answer: D
Section:

QUESTION 32

A company runs a stateless web application in production on a group of Amazon EC2 On-Demand Instances behind an Application Load Balancer. The application experiences heavy usage during an 8- hour period each business day. Application usage is moderate and steady overnight Application usage is low during weekends.

The company wants to minimize its EC2 costs without affecting the availability of the application.

Which solution will meet these requirements?

A. Use Spot Instances for the entire workload.

B. Use Reserved instances for the baseline level of usage Use Spot Instances for any additional capacity that the application needs.

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

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

Correct Answer: B
Section:

QUESTION 33

A company collects temperature, humidity, and atmospheric pressure data in cities across multiple continents. The average volume of data collected per site each day is 500 GB. Each site has a highspeed internet connection.
The company's weather forecasting applications are based in a single Region and analyze the data daily.
What is the FASTEST way to aggregate data from all of these global sites?

A. Enable Amazon S3 Transfer Acceleration on the destination bucket. Use multipart uploads to directly upload site data to the destination bucket.

B. Upload site data to an Amazon S3 bucket in the closest AWS Region. Use S3 cross-Region replication to copy objects to the destination bucket.

C. Schedule AWS Snowball jobs daily to transfer data to the closest AWS Region. Use S3 cross-Region replication to copy objects to the destination bucket.

D. Upload the data to an Amazon EC2 instance in the closest Region. Store the data in an Amazon Elastic Block Store (Amazon EBS) volume. Once a day take an EBS snapshot and copy it to the centralized Region. Restore the EBS volume in the centralized Region and run an analysis on the data daily.

Correct Answer: A
Section:
Explanation:

You might want to use Transfer Acceleration on a bucket for various reasons, including the following:
You have customers that upload to a centralized bucket from all over the world.
You transfer gigabytes to terabytes of data on a regular basis across continents.
You are unable to utilize all of your available bandwidth over the Internet when uploading to Amazon S3. https://docs.aws.amazon.com/AmazonS3/latest/dev/transfer-acceleration.html
https://aws.amazon.com/s3/transferacceleration/#:~:text=S3%20Transfer%20Acceleration%20(S3TA)%20reduces,to%20S3%20for%20remote%20applications:
"Amazon S3 Transfer Acceleration can speed up content transfers to and from Amazon S3 by as much as 50-500% for long-distance transfer of larger objects. Customers who have either web or mobile applications with widespread users or applications hosted far away from their S3 bucket can experience long and variable upload and download speeds over the Internet"
https://docs.aws.amazon.com/AmazonS3/latest/userguide/mpuoverview.html "Improved throughput - You can upload parts in parallel to improve throughput."

QUESTION 34

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

A. Use Amazon Redshift to load all the content into one place and run the SQL queries as needed

B. Use Amazon CloudWatch Logs to store the logs Run SQL queries as needed from the Amazon CloudWatch console

C. Use Amazon Athena directly with Amazon S3 to run the queries as needed

D. Use AWS Glue to catalog the logs Use a transient Apache Spark cluster on Amazon EMR to run the SQL queries as needed

Correct Answer: C
Section:
Explanation:

Amazon Athena can be used to query JSON in S3

QUESTION 35

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

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

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

QUESTION 36
An application runs on an Amazon EC2 instance in a VPC. The application processes logs that are stored in an Amazon S3 bucket. The EC2 instance needs to access the S3 bucket without connectivity to the internet. Which solution will provide private network connectivity to Amazon S3?

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

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

QUESTION 37
A company is hosting a web application on AWS using a single Amazon EC2 instance that stores useruploaded documents in an Amazon EBS volume. For better scalability and availability, the company duplicated the architecture and created a second EC2 instance and EBS volume in another Availability Zone placing both behind an Application Load Balancer After completing this change, users reported that, each time they refreshed the website, they could see one subset of their documents or the other, but never all of the documents at the same time.
What should a solutions architect propose to ensure users see all of their documents at once?

A. Copy the data so both EBS volumes contain all the documents.
B. Configure the Application Load Balancer to direct a user to the server with the documents
C. Copy the data from both EBS volumes to Amazon EFS Modify the application to save new documents to Amazon EFS

D. Configure the Application Load Balancer to send the request to both servers Return each document from the correct server.

Correct Answer: C
Section:
Explanation:

QUESTION 38
A company uses NFS to store large video files in on-premises network attached storage. Each video file ranges in size from 1MB to 500 GB. The total storage is 70 TB and is no longer growing. The company decides to migrate the video files to Amazon S3. The company must migrate the video files as soon as possible while using the least possible network bandwidth. Which solution will meet these requirements?

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

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

QUESTION 39
A company has an application that ingests incoming messages. These messages are then quickly consumed by dozens of other applications and microservices. The number of messages varies drastically and sometimes spikes as high as 100,000 each second.
The company wants to decouple the solution and increase scalability.
Which solution meets these requirements?

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

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

QUESTION 40
A company is migrating a distributed application to AWS The application serves variable workloads The legacy platform consists of a primary server trial coordinates jobs across multiple compute nodes The company wants to modernize the application with a solution that maximizes resiliency and scalability How should a solutions architect design the architecture to meet these requirements?

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

Correct Answer: B
Section:

QUESTION 41
A company is running an SMB file server in its data center. The file server stores large files that are accessed frequently for the first few days after the files are created. After 7 days the files are rarely accessed. The total data size is increasing and is close to the company's total storage capacity. A solutions architect must increase the company's available storage space without losing low-latency access to the most recently accessed files. The solutions architect must also provide file lifecycle management to avoid future storage issues.
Which solution will meet these requirements?

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

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

QUESTION 42
A company is building an ecommerce web application on AWS. The application sends information about new orders to an Amazon API Gateway REST API to process. The company wants to ensure that orders are processed in the order that they are received.
Which solution will meet these requirements?

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

Correct Answer: B
Section:

QUESTION 43

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

A. Use AWS Secrets Manager. Turn on automatic rotation.

B. Use AWS Systems Manager Parameter Store. Turn on automatic rotation.

C. Create an Amazon S3 bucket lo store objects that are encrypted with an AWS Key C. Management Service (AWS KMS) encryption key. Migrate the credential file to the S3 bucket. Point the application to the S3 bucket.

D. Create an encrypted Amazon Elastic Block Store (Amazon EBS) volume (or each EC2 instance.
Attach the new EBS volume to each EC2 instance. Migrate the credential file to the new EBS volume.
Point the application to the new EBS volume.

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

QUESTION 44
A global company hosts its web application on Amazon EC2 instances behind an Application Load Balancer (ALB). The web application has static data and dynamic dat a. The company stores its static data in an Amazon S3 bucket. The company wants to improve performance and reduce latency for the static data and dynamic data. The company is using its own domain name registered with Amazon Route 53. What should a solutions architect do to meet these requirements?

A. Create an Amazon CloudFront distribution that has the S3 bucket and the ALB as origins Configure Route 53 to route traffic to the CloudFront distribution.

B. Create an Amazon CloudFront distribution that has the ALB as an origin Create an AWS Global Accelerator standard accelerator that has the S3 bucket as an endpoint. Configure Route 53 to route traffic to the CloudFront distribution.

C. Create an Amazon CloudFront distribution that has the S3 bucket as an origin Create an AWS Global Accelerator standard accelerator that has the ALB and the CloudFront distribution as endpoints Create a custom domain name that points to the accelerator DNS name Use the custom domain name as an endpoint for the web application.

D. Create an Amazon CloudFront distribution that has the ALB as an origin C. Create an AWS Global Accelerator standard accelerator that has the S3 bucket as an endpoint Create two domain names. Point one domain name to the CloudFront DNS name for dynamic content, Point the other domain name to the accelerator DNS name for static content Use the domain names as endpoints for the web application.

Correct Answer: A
Section:
Explanation:

QUESTION 45
A company performs monthly maintenance on its AWS infrastructure. During these maintenance activities, the company needs to rotate the credentials tor its Amazon ROS tor MySQL databases across multiple AWS Regions Which solution will meet these requirements with the LEAST operational overhead?

A. Store the credentials as secrets in AWS Secrets Manager Use multi-Region secret replication for the required Regions Configure Secrets Manager to rotate the secrets on a schedule

B. Store the credentials as secrets in AWS Systems Manager by creating a secure string parameter Use multi-Region secret replication for the required Regions Configure Systems Manager to rotate the secrets on a schedule

C. Store the credentials in an Amazon S3 bucket that has server-side encryption (SSE) enabled Use Amazon EventBridge (Amazon CloudWatch Events) to invoke an AWS Lambda function to rotate the credentials

D. Encrypt the credentials as secrets by using AWS Key Management Service (AWS KMS) multi- Region customer managed keys Store the secrets in an Amazon DynamoDB global table Use an AWS Lambda function to retrieve the secrets from DynamoDB Use the RDS API to rotate the secrets.

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

QUESTION 46
A company runs an ecommerce application on Amazon EC2 instances behind an Application Load Balancer. The instances run in an Amazon EC2 Auto Scaling group across multiple Availability Zones. The Auto Scaling group scales based on CPU utilization metrics. The ecommerce application stores the transaction data in a MySQL 8.0 database that is hosted on a large EC2 instance. The database's performance degrades quickly as application load increases. The application handles more read requests than write transactions. The company wants a solution that will automatically scale the database to meet the demand of unpredictable read workloads while maintaining high availability.
Which solution will meet these requirements?

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

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

QUESTION 47
A company wants to improve its ability to clone large amounts of production data into a test environment in the same AWS Region. The data is stored in Amazon EC2 instances on Amazon Elastic Block Store (Amazon EBS) volumes. Modifications to the cloned data must not affect the production environment. The software that accesses this data requires consistently high I/O performance. A solutions architect needs to minimize the time that is required to clone the production data into the test environment. Which solution will meet these requirements?

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

Correct Answer: D
Section:
Explanation:

QUESTION 48
An ecommerce company wants to launch a one-deal-a-day website on AWS. Each day will feature exactly one product on sale for a period of 24 hours. The company wants to be able to handle millions of requests each hour with millisecond latency during peak hours.
Which solution will meet these requirements with the LEAST operational overhead?

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

Correct Answer: D
Section:

Explanation:

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

QUESTION 49

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

A. S3 Standard

B. S3 Intelligent-Tiering

C. S3 Standard-Infrequent Access {S3 Standard-IA)

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

Correct Answer: B

Section:

Explanation:

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

QUESTION 50

A company is storing backup files by using Amazon S3 Standard storage. The files are accessed frequently for 1 month. However, the files are not accessed after 1 month. The company must keep the files indefinitely. Which storage solution will meet these requirements MOST cost-effectively?

A. Configure S3 Intelligent-Tiering to automatically migrate objects.

B. Create an S3 Lifecycle configuration to transition objects from S3 Standard to S3 Glacier Deep Archive after 1 month.

C. Create an S3 Lifecycle configuration to transition objects from S3 Standard to S3 Standard- Infrequent Access (S3 Standard-IA) after 1 month.

D. Create an S3 Lifecycle configuration to transition objects from S3 Standard to S3 One Zone- Infrequent Access (S3 One Zone-IA) after 1 month.

Correct Answer: B

Section:

Explanation:

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

QUESTION 51

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

A. Use AWS Budgets to create a budget report and compare EC2 costs based on instance types

B. Use Cost Explorer's granular filtering feature to perform an in-depth analysis of EC2 costs based on instance types

C. Use graphs from the AWS Billing and Cost Management dashboard to compare EC2 costs based on instance types for the last 2 months

D. Use AWS Cost and Usage Reports to create a report and send it to an Amazon S3 bucket Use Amazon QuickSight with Amazon S3 as a source to generate an interactive graph based on instance types.

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

QUESTION 52
A company is designing an application. The application uses an AWS Lambda function to receive information through Amazon API Gateway and to store the information in an Amazon Aurora PostgreSQL database. During the proof-of-concept stage, the company has to increase the Lambda quotas significantly to handle the high volumes of data that the company needs to load into the database. A solutions architect must recommend a new design to improve scalability and minimize the configuration effort.
Which solution will meet these requirements?

A. Refactor the Lambda function code to Apache Tomcat code that runs on Amazon EC2 instances.
Connect the database by using native Java Database Connectivity (JDBC) drivers.

B. Change the platform from Aurora to Amazon DynamoDB. Provision a DynamoDB Accelerator (DAX) cluster. Use the DAX client SDK to point the existing DynamoDB API calls at the DAX cluster.

C. Set up two Lambda functions. Configure one function to receive the information. Configure the other function to load the information into the database. Integrate the Lambda functions by using Amazon Simple Notification Service (Amazon SNS).

D. Set up two Lambda functions. Configure one function to receive the information. Configure the other function to load the information into the database. Integrate the Lambda functions by using an Amazon Simple Queue Service (Amazon SQS) queue.

Correct Answer: D
Section:
Explanation:

QUESTION 53
A company needs to review its AWS Cloud deployment to ensure that its Amazon S3 buckets do not have unauthorized configuration changes. What should a solutions architect do to accomplish this goal?

A. Turn on AWS Config with the appropriate rules.

B. Turn on AWS Trusted Advisor with the appropriate checks.

C. Turn on Amazon Inspector with the appropriate assessment template.

D. Turn on Amazon S3 server access logging. Configure Amazon EventBridge (Amazon Cloud Watch Events).

Correct Answer: A
Section:
Explanation:

QUESTION 54
A company is launching a new application and will display application metrics on an Amazon CloudWatch dashboard. The company's product manager needs to access this dashboard periodically. The product manager does not have an AWS account. A solution architect must provide access to the product manager by following the principle of least privilege. Which solution will meet these requirements?

A. Share the dashboard from the CloudWatch console. Enter the product manager's email address, and complete the sharing steps. Provide a shareable link for the dashboard to the product manager.

B. Create an IAM user specifically for the product manager. Attach the CloudWatch Read Only Access managed policy to the user. Share the new login credential with the product manager. Share the browser URL of the correct dashboard with the product manager.

C. Create an IAM user for the company’s employees, Attach the View Only Access AWS managed policy to the IAM user. Share the new login credentials with the product manager. Ask the product manager to navigate to the CloudWatch console and locate the dashboard by name in the Dashboards section.

D. Deploy a bastion server in a public subnet. When the product manager requires access to the dashboard, start the server and share the RDP credentials. On the bastion server, ensure that the browser is configured to open the dashboard URL with cached AWS credentials that have appropriate permissions to view the dashboard.

Correct Answer: A
Section:
Explanation:

QUESTION 55

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

A. Enable AWS Single Sign-On (AWS SSO) from the AWS SSO console. Create a one-way forest trust or a one-way domain trust to connect the company's self-managed Microsoft Active Directory with AWS SSO by using AWS Directory Service for Microsoft Active Directory.

B. Enable AWS Single Sign-On (AWS SSO) from the AWS SSO console. Create a two-way forest trust to connect the company's self-managed Microsoft Active Directory with AWS SSO by using AWS Directory Service for Microsoft Active Directory.

C. Use AWS Directory Service. Create a two-way trust relationship with the company's self-managed Microsoft Active Directory.

D. Deploy an identity provider (IdP) on premises. Enable AWS Single Sign-On (AWS SSO) from the AWS SSO console.

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

QUESTION 56

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

A. Deploy a Network Load Balancer (NLB) and an associated target group. Associate the target group with the Auto Scaling group. Use the NLB as an AWS Global Accelerator endpoint in each Region.

B. Deploy an Application Load Balancer (ALB) and an associated target group. Associate the target group with the Auto Scaling group. Use the ALB as an AWS Global Accelerator endpoint in each Region.

C. Deploy a Network Load Balancer (NLB) and an associated target group. Associate the target group with the Auto Scaling group. Create an Amazon Route 53 latency record that points to aliases for each NLB. Create an Amazon CloudFront distribution that uses the latency record as an origin.

D. Deploy an Application Load Balancer (ALB) and an associated target group. Associate the target group with the Auto Scaling group. Create an Amazon Route 53 weighted record that points to aliases for each ALB. Deploy an Amazon CloudFront distribution that uses the weighted record as an origin.

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

QUESTION 57

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

A. Stop the DB instance when tests are completed. Restart the DB instance when required.

B. Use an Auto Scaling policy with the DB instance to automatically scale when tests are completed.

C. Create a snapshot when tests are completed. Terminate the DB instance and restore the snapshot when required.

D. Modify the DB instance to a low-capacity instance when tests are completed. Modify the DB instance again when required.

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

QUESTION 58

A company that hosts its web application on AWS wants to ensure all Amazon EC2 instances. Amazon RDS DB instances. and Amazon Redshift clusters are configured with tags. The company wants to minimize the effort of configuring and operating this check.

What should a solutions architect do to accomplish this?

A. Use AWS Config rules to define and detect resources that are not properly tagged.
B. Use Cost Explorer to display resources that are not properly tagged. Tag those resources manually.
C. Write API calls to check all resources for proper tag allocation. Periodically run the code on an EC2 instance.
D. Write API calls to check all resources for proper tag allocation. Schedule an AWS Lambda function through Amazon CloudWatch to periodically run the code.

Correct Answer: A
Section:
Explanation:
To ensure all Amazon EC2 instances, Amazon RDS DB instances, and Amazon Redshift clusters are configured with tags, a solutions architect should use AWS Config rules to define and detect resources that are not properly tagged. AWS Config rules are a set of customizable rules that AWS Config uses to evaluate AWS resource configurations for compliance with best practices and company policies. Using AWS Config rules can minimize the effort of configuring and operating this check because it automates the process of identifying non-compliant resources and notifying the responsible teams.Reference: AWS Config Developer Guide: AWS Config Rules(https://docs.aws.amazon.com/config/latest/developerguide/evaluate-config_use-managed- rules.html)

QUESTION 59
A development team needs to host a website that will be accessed by other teams. The website contents consist of HTML, CSS, client-side JavaScript, and images Which method is the MOST costeffective for hosting the website?

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

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

QUESTION 60
A company runs an online marketplace web application on AWS. The application serves hundreds of thousands of users during peak hours. The company needs a scalable, near-real-time solution to share the details of millions of financial transactions with several other internal applications Transactions also need to be processed to remove sensitive data before being stored in a document database for low-latency retrieval. What should a solutions architect recommend to meet these requirements?

A. Store the transactions data into Amazon DynamoDB Set up a rule in DynamoDB to remove sensitive data from every transaction upon write Use DynamoDB Streams to share the transactions data with other applications
B. Stream the transactions data into Amazon Kinesis Data Firehose to store data in Amazon DynamoDB and Amazon S3 Use AWS Lambda integration with Kinesis Data Firehose to remove sensitive data. Other applications can consume the data stored in Amazon S3
C. Stream the transactions data into Amazon Kinesis Data Streams Use AWS Lambda integration to remove sensitive data from every transaction and then store the transactions data in Amazon DynamoDB Other applications can consume the transactions data off the Kinesis data stream.
D. Store the batched transactions data in Amazon S3 as files. Use AWS Lambda to process every file and remove sensitive data before updating the files in Amazon S3 The Lambda function then stores the data in Amazon DynamoDB Other applications can consume transaction files stored in Amazon S3.

Correct Answer: C

Section:
Explanation:
The destination of your Kinesis Data Firehose delivery stream. Kinesis Data Firehose can send data records to various destinations, including Amazon Simple Storage Service (Amazon S3), Amazon Redshift, Amazon OpenSearch Service, and any HTTP endpoint that is owned by you or any of your third-party service providers. The following are the supported destinations:
* Amazon OpenSearch Service
* Amazon S3
* Datadog
* Dynatrace
* Honeycomb
* HTTP Endpoint
* Logic Monitor
* MongoDB Cloud
* New Relic
* Splunk
* Sumo Logic
https://docs.aws.amazon.com/firehose/latest/dev/create-name.html
https://aws.amazon.com/kinesis/data-streams/
Amazon Kinesis Data Streams (KDS) is a massively scalable and durable real-time data streaming service. KDS can continuously capture gigabytes of data per second from hundreds of thousands of sources such as website clickstreams, database event streams, financial transactions, social media feeds, IT logs, and location-tracking events.

QUESTION 61
A company hosts its multi-tier applications on AWS. For compliance, governance, auditing, and security, the company must track configuration changes on its AWS resources and record a history of API calls made to these resources. What should a solutions architect do to meet these requirements?

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

Correct Answer: B
Section:
Explanation:
AWS Config is a fully managed service that allows the company to assess, audit, and evaluate the configurations of its AWS resources. It provides a detailed inventory of the resources in use and tracks changes to resource configurations. AWS Config can detect configuration changes and alert the company when changes occur. It also provides a historical view of changes, which is essential for compliance and governance purposes. AWS CloudTrail is a fully managed service that provides a detailed history of API calls made to the company's AWS resources. It records all API activity in the AWS account, including who made the API call, when the call was made, and what resources were affected by the call. This information is critical for security and auditing purposes, as it allows the company to investigate any suspicious activity that might occur on its AWS resources.

QUESTION 62
A company is preparing to launch a public-facing web application in the AWS Cloud. The architecture consists of Amazon EC2 instances within a VPC behind an Elastic Load Balancer (ELB). A third-party service is used for the DNS. The company's solutions architect must recommend a solution to detect and protect against large-scale DDoS attacks. Which solution meets these requirements?

A. Enable Amazon GuardDuty on the account.
B. Enable Amazon Inspector on the EC2 instances.
C. Enable AWS Shield and assign Amazon Route 53 to it.
D. Enable AWS Shield Advanced and assign the ELB to it.

Correct Answer: D
Section:
Explanation:

https://aws.amazon.com/shield/faqs/

QUESTION 63

A company is building an application in the AWS Cloud. The application will store data in Amazon S3 buckets in two AWS Regions. The company must use an AWS Key Management Service (AWS KMS) customer managed key to encrypt all data that is stored in the S3 buckets. The data in both S3 buckets must be encrypted and decrypted with the same KMS key. The data and the key must be stored in each of the two Regions. Which solution will meet these requirements with the LEAST operational overhead?

A. Create an S3 bucket in each Region Configure the S3 buckets to use server-side encryption with Amazon S3 managed encryption keys (SSE-S3) Configure replication between the S3 buckets.

B. Create a customer managed multi-Region KMS key. Create an S3 bucket in each Region. Configure replication between the S3 buckets. Configure the application to use the KMS key with client-side encryption.

C. Create a customer managed KMS key and an S3 bucket in each Region Configure the S3 buckets to use server-side encryption with Amazon S3 managed encryption keys (SSE-S3) Configure replication between the S3 buckets.

D. Create a customer managed KMS key and an S3 bucket m each Region Configure the S3 buckets to use server-side encryption with AWS KMS keys (SSE-KMS) Configure replication between the S3 buckets.

Correct Answer: B
Section:
Explanation:

From https://docs.aws.amazon.com/kms/latest/developerguide/custom-key-store-overview.htmlFor most users, the default AWS KMS key store, which is protected by FIPS 140-2 validatedcryptographic modules, fulfills their security requirements. There is no need to add an extra layer ofmaintenance responsibility or a dependency on an additional service. However, you might considercreating a custom key store if your organization has any of the following requirements: Key materialcannot be stored in a shared environment. Key material must be subject to a secondary, independentaudit path. The HSMs that generate and store key material must be certified at FIPS 140-2 Level 3. https://docs.aws.amazon.com/kms/latest/developerguide/custom-key-store-overview.html

QUESTION 64

A company recently launched a variety of new workloads on Amazon EC2 instances in its AWS account. The company needs to create a strategy to access and administer the instances remotely and securely. The company needs to implement a repeatable process that works with native AWS services and follows the AWS Well-Architected Framework. Which solution will meet these requirements with the LEAST operational overhead?

A. Use the EC2 serial console to directly access the terminal interface of each instance for administration.

B. Attach the appropriate IAM role to each existing instance and new instance. Use AWS Systems Manager Session Manager to establish a remote SSH session.

C. Create an administrative SSH key pair. Load the public key into each EC2 instance. Deploy a bastion host in a public subnet to provide a tunnel for administration of each instance.

D. Establish an AWS Site-to-Site VPN connection. Instruct administrators to use their local onpremises machines to connect directly to the instances by using SSH keys across the VPN tunnel.

Correct Answer: B
Section:
Explanation:

https://docs.aws.amazon.com/systems-manager/latest/userguide/setup-launch-managedinstance. html

QUESTION 65

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

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

B. Provision accelerators in AWS Global Accelerator. Associate the supplied IP addresses with the S3 bucket. Edit the Route 53 entries to point to the IP addresses of the accelerators.

C. Add an Amazon CloudFront distribution in front of the S3 bucket. Edit the Route 53 entries to point to the CloudFront distribution.

D. Enable S3 Transfer Acceleration on the bucket. Edit the Route 53 entries to point to the new endpoint.

Correct Answer: C
Section:
Explanation:

Amazon CloudFront is a content delivery network (CDN) that caches content at edge locations around the world, providing low latency and high transfer speeds to users accessing the content.Adding a CloudFront distribution in front of the S3 bucket will cache the static website's content at edge locations around the world, decreasing latency for users accessing the website. This solution is

QUESTION 66

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

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

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

QUESTION 67

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

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

Correct Answer: A
Section:
Explanation:
https://aws.amazon.com/kinesis/datafirehose/ features/?nc=sn&loc=2#:~:text=into%20Amazon%20S3%2C%20Amazon%20Redshift%2C%20Amazon%20OpenSearch%20Service%2C%20Kinesis,Delivery%20streams

QUESTION 68

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 69
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 70
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 71
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 72
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 73
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 74
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 75

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 76

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 77

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 78
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 79
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 80
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 81
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 82
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 83
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 84
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:

QUESTION 85
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 86
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 87

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 88

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 89

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

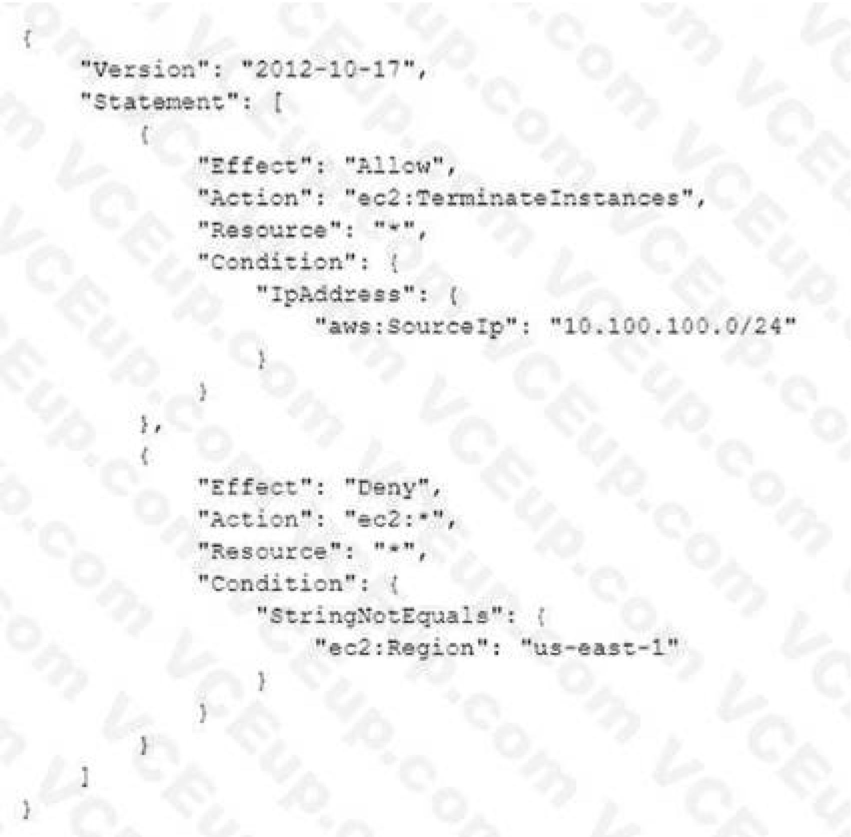

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": "ec2:TerminateInstances",
            "Resource": "*",
            "Condition": {
                "IpAddress": {
                    "aws:SourceIp": "10.100.100.0/24"
                }
            }
        },
        {
            "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 90
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 91
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 92
A business's backup data totals 700 terabytes (TB) and is kept in network attached storage (NAS) at its data center. This backup data must be available in the event of occasional regulatory inquiries and preserved for a period of seven years. The organization has chosen to relocate its backup data from its on-premises data center to Amazon Web Services (AWS). Within one month, the migration must be completed. The company's public internet connection provides 500 Mbps of dedicated capacity for data transport.
What should a solutions architect do to ensure that data is migrated and stored at the LOWEST possible cost?

A. Order AWS Snowball devices to transfer the data. Use a lifecycle policy to transition the files to Amazon S3 Glacier Deep Archive.
B. Deploy a VPN connection between the data center and Amazon VPC. Use the AWS CLI to copy the data from on premises to Amazon S3 Glacier.
C. Provision a 500 Mbps AWS Direct Connect connection and transfer the data to Amazon S3. Use a lifecycle policy to transition the files to Amazon S3 Glacier Deep Archive.
D. Use AWS DataSync to transfer the data and deploy a DataSync agent on premises. Use the DataSync task to copy files from the on-premises NAS storage to Amazon S3 Glacier.

Correct Answer: A
Section:
Explanation:
https://www.omnicalculator.com/other/data-transfer

QUESTION 93

A company wants to direct its users to a backup static error page if the company's primary website is unavailable. The primary website's DNS records are hosted in Amazon Route 53. The domain is pointing to an Application Load Balancer (ALB). The company needs a solution that minimizes changes and infrastructure overhead.

Which solution will meet these requirements?

A. Update the Route 53 records to use a latency routing policy. Add a static error page that is hosted in an Amazon S3 bucket to the records so that the traffic is sent to the most responsive endpoints.

B. Set up a Route 53 active-passive failover configuration. Direct traffic to a static error page that is hosted in an Amazon S3 bucket when Route 53 health checks determine that the ALB endpoint is unhealthy.

C. Set up a Route 53 active-active configuration with the ALB and an Amazon EC2 instance that hosts a static error page as endpoints. Configure Route 53 to send requests to the instance only if the health checks fail for the ALB.

D. Update the Route 53 records to use a multivalue answer routing policy. Create a health check.
Direct traffic to the website if the health check passes. Direct traffic to a static error page that is hosted in Amazon S3 if the health check does not pass.

Correct Answer: B
Section:
Explanation:
This solution meets the requirements of directing users to a backup static error page if the primary website is unavailable, minimizing changes and infrastructure overhead. Route 53 active-passive failover configuration can route traffic to a primary resource when it is healthy or to a secondary resource when the primary resource is unhealthy. Route 53 health checks can monitor the health of the ALB endpoint and trigger the failover when needed. The static error page can be hosted in an S3 bucket that is configured as a website, which is a simple and cost-effective way to serve static content.Option A is incorrect because using a latency routing policy can route traffic based on the lowest network latency for users, but it does not provide failover functionality. Option C is incorrect because using an active-active configuration with the ALB and an EC2 instance can increase the infrastructure overhead and complexity, and it does not guarantee that the EC2 instance will always be healthy.Option D is incorrect because using a multivalue answer routing policy can return multiple values for a query, but it does not provide failover functionality.Reference:https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-policy-failover.html https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/dns-failover.html

QUESTION 94
A corporation has recruited a new cloud engineer who should not have access to the CompanyConfidential Amazon S3 bucket. The cloud engineer must have read and write permissions on an S3 bucket named AdminTools.
Which IAM policy will satisfy these criteria?

A.

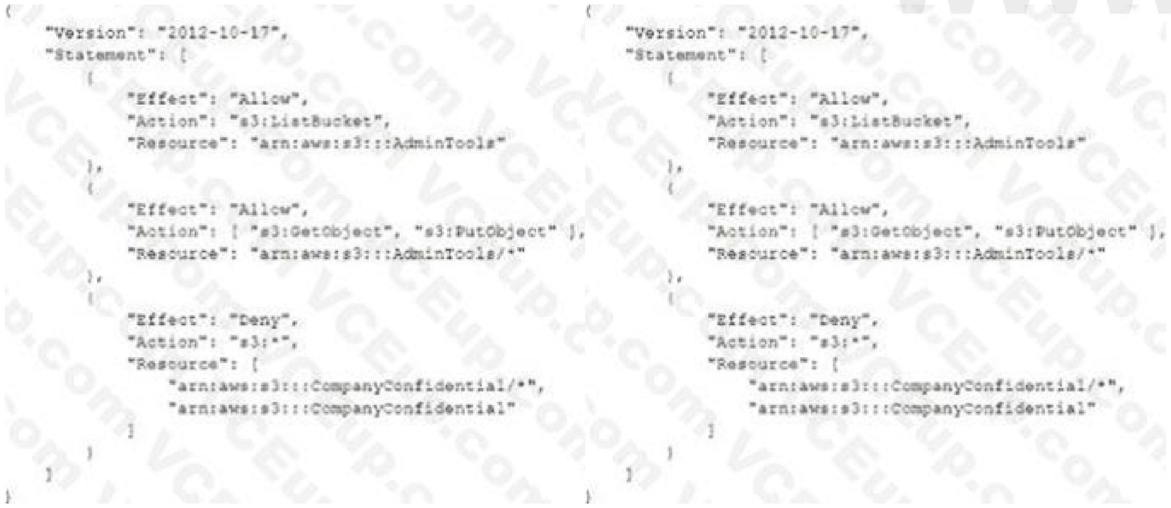

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": "s3:ListBucket",
            "Resource": "arn:aws:s3:::AdminTools"
        },
        {
            "Effect": "Allow",
            "Action": [ "s3:GetObject", "s3:PutObject" ],
            "Resource": "arn:aws:s3:::AdminTools/*"
        },
        {
            "Effect": "Deny",
            "Action": "s3:*",
            "Resource": [
                "arn:aws:s3:::CompanyConfidential/*",
                "arn:aws:s3:::CompanyConfidential"
            ]
        }
    ]
}
```

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": "s3:ListBucket",
            "Resource": "arn:aws:s3:::AdminTools"
        },
        {
            "Effect": "Allow",
            "Action": [ "s3:GetObject", "s3:PutObject" ],
            "Resource": "arn:aws:s3:::AdminTools/*"
        },
        {
            "Effect": "Deny",
            "Action": "s3:*",
            "Resource": [
                "arn:aws:s3:::CompanyConfidential/*",
                "arn:aws:s3:::CompanyConfidential"
            ]
        }
    ]
}
```

B.

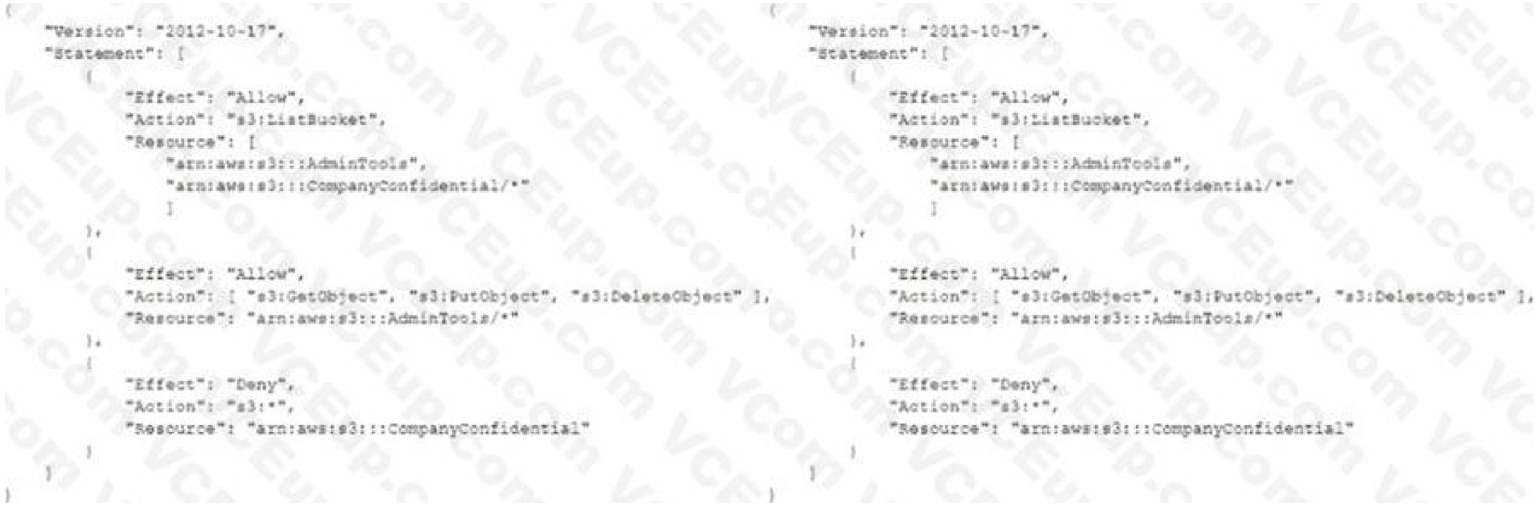

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": "s3:ListBucket",
            "Resource": [
                "arn:aws:s3:::AdminTools",
                "arn:aws:s3:::CompanyConfidential/*"
                ]
        },
        {
            "Effect": "Allow",
            "Action": [ "s3:GetObject", "s3:PutObject", "s3:DeleteObject" ],
            "Resource": "arn:aws:s3:::AdminTools/*"
        },
        {
            "Effect": "Deny",
            "Action": "s3:*",
            "Resource": "arn:aws:s3:::CompanyConfidential"
        }
    ]
}
```

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": "s3:ListBucket",
            "Resource": [
                "arn:aws:s3:::AdminTools",
                "arn:aws:s3:::CompanyConfidential/*"
                ]
        },
        {
            "Effect": "Allow",
            "Action": [ "s3:GetObject", "s3:PutObject", "s3:DeleteObject" ],
            "Resource": "arn:aws:s3:::AdminTools/*"
        },
        {
            "Effect": "Deny",
            "Action": "s3:*",
            "Resource": "arn:aws:s3:::CompanyConfidential"
        }
    ]
}
```

C.

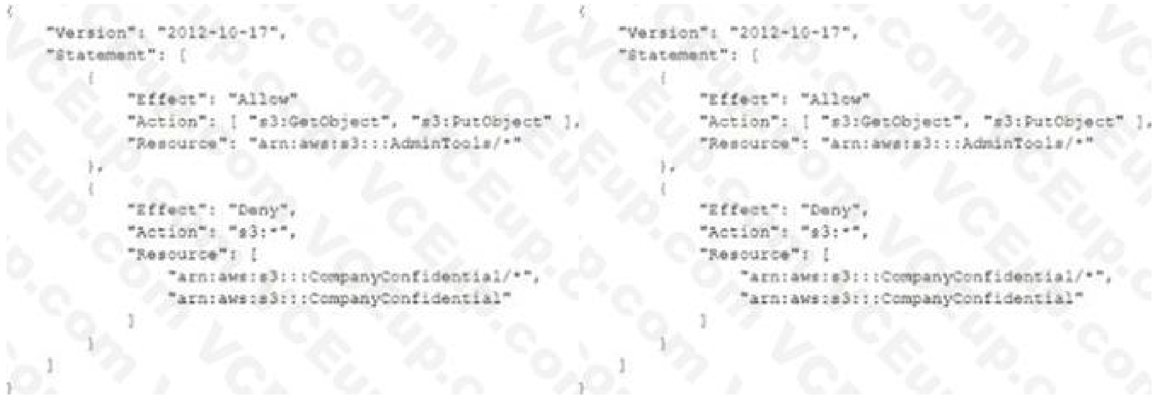

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow"
            "Action": [ "s3:GetObject", "s3:PutObject" ],
            "Resource": "arn:aws:s3:::AdminTools/*"
        },
        {
            "Effect": "Deny",
            "Action": "s3:*",
            "Resource": [
                "arn:aws:s3:::CompanyConfidential/*",
                "arn:aws:s3:::CompanyConfidential"
            ]
        }
    ]
}
```

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow"
            "Action": [ "s3:GetObject", "s3:PutObject" ],
            "Resource": "arn:aws:s3:::AdminTools/*"
        },
        {
            "Effect": "Deny",
            "Action": "s3:*",
            "Resource": [
                "arn:aws:s3:::CompanyConfidential/*",
                "arn:aws:s3:::CompanyConfidential"
            ]
        }
    ]
}
```

D.

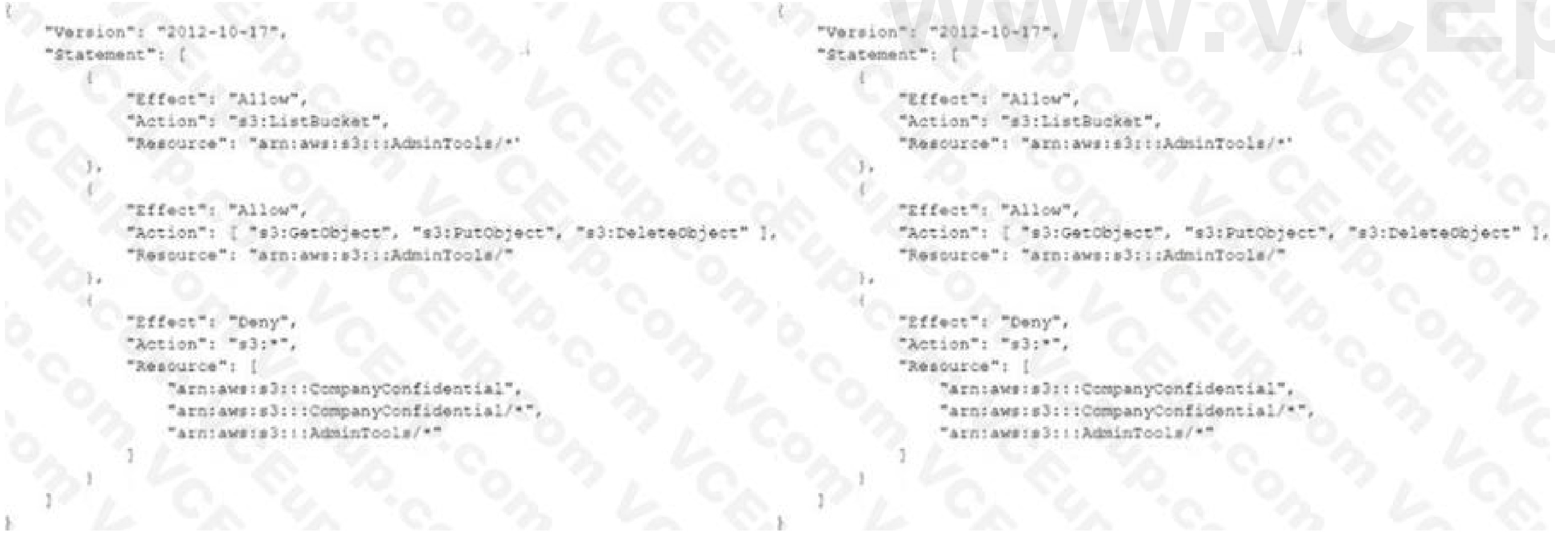

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": "s3:ListBucket",
            "Resource": "arn:aws:s3:::AdminTools/*'
        },
        {
            "Effect": "Allow",
            "Action": [ "s3:GetObject", "s3:PutObject", "s3:DeleteObject" ],
            "Resource": "arn:aws:s3:::AdminTools/"
        },
        {
            "Effect": "Deny",
            "Action": "s3:*",
            "Resource": [
                "arn:aws:s3:::CompanyConfidential",
                "arn:aws:s3:::CompanyConfidential/*",
                "arn:aws:s3:::AdminTools/*"
            ]
        }
    ]
}
```

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": "s3:ListBucket",
            "Resource": "arn:aws:s3:::AdminTools/*'
        },
        {
            "Effect": "Allow",
            "Action": [ "s3:GetObject", "s3:PutObject", "s3:DeleteObject" ],
            "Resource": "arn:aws:s3:::AdminTools/"
        },
        {
            "Effect": "Deny",
            "Action": "s3:*",
            "Resource": [
                "arn:aws:s3:::CompanyConfidential",
                "arn:aws:s3:::CompanyConfidential/*",
                "arn:aws:s3:::AdminTools/*"
            ]
        }
    ]
}
```

Correct Answer: C
Section:
Explanation:

QUESTION 95

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

A. Have the deployment engineer use AWS account roof user credentials for performing AWS CloudFormation stack operations.

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

Correct Answer: D, E
Section:
Explanation:
https://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles.html
https://docs.aws.amazon.com/IAM/latest/UserGuide/id_users.html

QUESTION 96
A company runs a high performance computing (HPC) workload on AWS. The workload required lowlatency network performance and high network throughput with tightly coupled node-to-node communication. The Amazon EC2 instances are properly sized for compute and storage capacity, and are launched using default options.
What should a solutions architect propose to improve the performance of the workload?

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

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

QUESTION 97
A company wants to use the AWS Cloud to make an existing application highly available and resilient. The current version of the application resides in the company's data center. The application recently experienced data loss after a database server crashed because of an unexpected power outage. The company needs a solution that avoids any single points of failure. The solution must give the application the ability to scale to meet user demand.
Which solution will meet these requirements?

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

Correct Answer: A
Section:

QUESTION 98
A company wants to run a gaming application on Amazon EC2 instances that are part of an Auto Scaling group in the AWS Cloud. The application will transmit data by using UDP packets. The company wants to ensure that the application can scale out and in as traffic increases and decreases.
What should a solutions architect do to meet these requirements?

A. Attach a Network Load Balancer to the Auto Scaling group

B. Attach an Application Load Balancer to the Auto Scaling group.
C. Deploy an Amazon Route 53 record set with a weighted policy to route traffic appropriately
D. Deploy a NAT instance that is configured with port forwarding to the EC2 instances in the Auto Scaling group.

Correct Answer: A
Section:
Explanation:
This solution meets the requirements of running a gaming application that transmits data by using UDP packets and scaling out and in as traffic increases and decreases. A Network Load Balancer can handle millions of requests per second while maintaining high throughput at ultra low latency, and it supports both TCP and UDP protocols. An Auto Scaling group can automatically adjust the number of EC2 instances based on the demand and the scaling policies. Option B is incorrect because an Application Load Balancer does not support UDP protocol, only HTTP and HTTPS. Option C is incorrect because Amazon Route 53 is a DNS service that can route traffic based on different policies, but it does not provide load balancing or scaling capabilities. Option D is incorrect because a NAT instance is used to enable instances in a private subnet to connect to the internet or other AWS services, but it does not provide load balancing or scaling capabilities.
https://aws.amazon.com/blogs/aws/new-udp-load-balancing-for-network-load-balancer/ https://docs.aws.amazon.com/autoscaling/ec2/userguide/AutoScalingGroup.html

QUESTION 99
A solutions architect is designing a customer-facing application for a company. The application's database will have a clearly defined access pattern throughout the year and will have a variable number of reads and writes that depend on the time of year. The company must retain audit records for the database for 7 days. The recovery point objective (RPO) must be less than 5 hours. Which solution meets these requirements?

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

Correct Answer: A
Section:
Explanation:

QUESTION 100
A company hosts a two-tier application on Amazon EC2 instances and Amazon RDS. The application's demand varies based on the time of day. The load is minimal after work hours and on weekends. The EC2 instances run in an EC2 Auto Scaling group that is configured with a minimum of two instances and a maximum of five instances. The application must be available at all times, but the company is concerned about overall cost. Which solution meets the availability requirement MOST cost-effectively?

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

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

in a fixed amount of compute usage per hour, which may not match the actual usage pattern of the application. Using up to three additional EC2 On-Demand Instances as needed can incur higher costs than using Spot Instances.

QUESTION 101

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

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

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

QUESTION 102

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

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

Correct Answer: A
Section:

QUESTION 103

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

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

Correct Answer: D
Section:

QUESTION 104

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

A. Store container images In an Amazon Elastic Container Registry (Amazon ECR) repository. Use an Amazon Elastic Container Service (Amazon ECS) cluster with the AWS Fargate launch type to run the containers. Use target tracking to scale automatically based on demand.

B. Store container images in an Amazon Elastic Container Registry (Amazon ECR) repository. Use an Amazon Elastic Container Service (Amazon ECS) cluster with the Amazon EC2 launch type to run the containers. Use target tracking to scale automatically based on demand.

C. Store container images in a repository that runs on an Amazon EC2 instance. Run the containers on EC2 instances that are spread across multiple Availability Zones. Monitor the average CPU utilization in Amazon CloudWatch. Launch new EC2 instances as needed

D. Create an Amazon EC2 Amazon Machine Image (AMI) that contains the container image Launch EC2 Instances in an Auto Scaling group across multiple Availability Zones. Use an Amazon CloudWatch alarm to scale out EC2 instances when the average CPU utilization threshold is breached.

Correct Answer: A
Section:

QUESTION 105

A company's application Is having performance issues The application staleful and needs to complete m-memory tasks on Amazon EC2 instances. The company used AWS CloudFormation to deploy infrastructure and used the M5 EC2 Instance family As traffic increased, the application performance degraded Users are reporting delays when the users attempt to access the application. Which solution will resolve these issues in the MOST operationally efficient way?

A. Replace the EC2 Instances with T3 EC2 instances that run in an Auto Scaling group. Made the changes by using the AWS Management Console.

B. Modify the CloudFormation templates to run the EC2 instances in an Auto Scaling group. Increase the desired capacity and the maximum capacity of the Auto Scaling group manually when an increase is necessary

C. Modify the CloudFormation templates. Replace the EC2 instances with R5 EC2 instances. Use Amazon CloudWatch built-in EC2 memory metrics to track the application performance for future capacity planning.

D. Modify the CloudFormation templates. Replace the EC2 instances with R5 EC2 instances. Deploy the Amazon CloudWatch agent on the EC2 instances to generate custom application latency metrics for future capacity planning.

Correct Answer: C
Section:
Explanation:

QUESTION 106

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

A. Configure provisioned concurrency for the Lambda function Modify the database to be a global database in multiple AWS Regions

B. Use Amazon RDS Proxy to create a proxy for the database Modify the Lambda function to use the RDS Proxy endpoint instead of the database endpoint

C. Create a read replica for the database in a different AWS Region Use query string parameters in API Gateway to route traffic to the read replica

D. Migrate the data from Aurora PostgreSQL to Amazon DynamoDB by using AWS Database Migration Service (AWS DMS| Modify the Lambda function to use the OynamoDB table

Correct Answer: B
Section:
Explanation:

QUESTION 107

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

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

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

QUESTION 108
A company wants to measure the effectiveness of its recent marketing campaigns. The company performs batch processing on csv files of sales data and stores the results «i an Amazon S3 bucket once every hour. The S3 bi petabytes of objects. The company runs one-time queries in Amazon Athena to determine which products are most popular on a particular date for a particular region Queries sometimes fail or take longer than expected to finish. Which actions should a solutions architect take to improve the query performance and reliability?
(Select TWO.)

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

Correct Answer: B, E
Section:
Explanation:

QUESTION 109
A company is running several business applications in three separate VPCs within me us-east-1 Region. The applications must be able to communicate between VPCs. The applications also must be able to consistently send hundreds to gigabytes of data each day to a latency-sensitive application that runs in a single on-premises data center. A solutions architect needs to design a network connectivity solution that maximizes costeffectiveness Which solution moots those requirements?

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

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

QUESTION 110
A company's web application is running on Amazon EC2 instances behind an Application Load Balancer. The company recently changed its policy, which now requires the application to be accessed from one specific country

only. Which configuration will meet this requirement?

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

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

QUESTION 111
A company has a mulli-tier application that runs six front-end web servers in an Amazon EC2 Auto Scaling group in a single Availability Zone behind an Application Load Balancer (ALB). A solutions architect needs lo modify the infrastructure to be highly available without modifying the application.
Which architecture should the solutions architect choose that provides high availability?

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

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

QUESTION 112
Organizers for a global event want to put daily reports online as static HTML pages. The pages are expected to generate millions of views from users around the world. The files are stored In an Amazon S3 bucket. A solutions architect has been asked to design an efficient and effective solution.
Which action should the solutions architect take to accomplish this?

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

Correct Answer: D
Section:

QUESTION 113
A company runs an application using Amazon ECS. The application creates esi/ed versions of an original image and then makes Amazon S3 API calls to store the resized images in Amazon S3. How can a solutions architect ensure that the application has permission to access Amazon S3?

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

D. Create an IAM user with S3 permissions, and then relaunch the Amazon EC2 instances for the ECS cluster while logged in as this account.

Correct Answer: B
Section:

QUESTION 114
A solutions architect needs to securely store a database user name and password that an application uses to access an Amazon RDS DB instance. The application that accesses the database runs on an Amazon EC2 instance. The solutions architect wants to create a secure parameter in AWS Systems Manager Parameter Store.
What should the solutions architect do to meet this requirement?

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

Correct Answer: A
Section:
Explanation:

QUESTION 115
An entertainment company is using Amazon DynamoDB to store media metadat a. The application is read intensive and experiencing delays. The company does not have staff to handle additional operational overhead and needs to improve the performance efficiency of DynamoDB without reconfiguring the application.
What should a solutions architect recommend to meet this requirement?

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

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

QUESTION 116
A security team wants to limit access to specific services or actions in all of the team's AWS accounts.
All accounts belong to a large organization in AWS Organizations. The solution must be scalable and there must be a single point where permissions can be maintained.
What should a solutions architect do to accomplish this?

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

Correct Answer: D

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

QUESTION 117
A company is concerned about the security of its public web application due to recent web attacks.
The application uses an Application Load Balancer (ALB). A solutions architect must reduce the risk of DDoS attacks against the application. What should the solutions architect do to meet this requirement?

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

Correct Answer: C
Section:

QUESTION 118
A company runs a production application on a fleet of Amazon EC2 instances. The application reads the data from an Amazon SQS queue and processes the messages in parallel. The message volume is unpredictable and often has intermittent traffic. This application should continually process messages without any downtime.
Which solution meets these requirements MOST cost-effectively?

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

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

QUESTION 119
A solutions architect must design a solution that uses Amazon CloudFront with an Amazon S3 origin to store a static website. The company's security policy requires that all website traffic be inspected by AWS WAR How should the solutions architect comply with these requirements?

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

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

QUESTION 120

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

A. Create an AWS Lambda function to query AWS CloudTrail logs and to send an alert when a Createlmage API call is detected.

B. Configure AWS CloudTrail with an Amazon Simple Notification Service {Amazon SNS) notification that occurs when updated logs are sent to Amazon S3. Use Amazon Athena to create a new table and to query on Createlmage when an API call is detected.

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

D. Configure an Amazon Simple Queue Service (Amazon SQS) FIFO queue as a target for AWS CloudTrail logs. Create an AWS Lambda function to send an alert to an Amazon Simple Notification Service (Amazon SNS) topic when a Createlmage API call is detected.

Correct Answer: C
Section:
Explanation:

QUESTION 121

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 (he 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: D
Section:
Explanation:

QUESTION 122

A company owns an asynchronous API that is used to ingest user requests and, based on the request type, dispatch requests to the appropriate microservice for processing. The company is using Amazon API Gateway to deploy the API front end, and an AWS Lambda function that invokes Amazon DynamoDB to store user requests before dispatching them to the processing microservices. The company provisioned as much DynamoDB throughput as its budget allows, but the company is still experiencing availability issues and is losing user requests. What should a solutions architect do to address this issue without impacting existing users?

A. Add throttling on the API Gateway with server-side throttling limits.

B. Use DynamoDB Accelerator (DAX) and Lambda to buffer writes to DynamoDB.

C. Create a secondary index in DynamoDB for the table with the user requests.

D. Use the Amazon Simple Queue Service (Amazon SQS) queue and Lambda to buffer writes to DynamoDB.

Correct Answer: D
Section:
Explanation:
By using an SQS queue and Lambda, the solutions architect can decouple the API front end from the processing microservices and improve the overall scalability and availability of the system. The SQS queue acts as a buffer, allowing the API front end to continue accepting user requests even if the processing microservices are experiencing high workloads or are temporarily unavailable. The Lambda function can then retrieve requests from the

SQS queue and write them to DynamoDB, ensuring that all user requests are stored and processed. This approach allows the company to scale the processing microservices independently from the API front end, ensuring that the API remains available to users even during periods of high demand.

QUESTION 123

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

A. Create an interface VPC endpoint for Amazon S3 in the subnet where the EC2 instance is located.
Attach a resource policy to the S3 bucket to only allow the EC2 instance's IAM role for access.

B. Create a gateway VPC endpoint for Amazon S3 in the Availability Zone where the EC2 instance is located. Attach appropriate security groups to the endpoint. Attach a resource policy lo the S3 bucket to only allow the EC2 instance's IAM role for access.

C. Run the nslookup tool from inside the EC2 instance to obtain the private IP address of the S3 bucket's service API endpoint. Create a route in the VPC route table to provide the EC2 instance with access to the S3 bucket. Attach a resource policy to the S3 bucket to only allow the EC2 instance's IAM role for access.

D. Use the AWS provided, publicly available ip-ranges.json tile to obtain the private IP address of the S3 bucket's service API endpoint. Create a route in the VPC route table to provide the EC2 instance with access to the S3 bucket. Attach a resource policy to the S3 bucket to only allow the EC2 instance's IAM role for access.

Correct Answer: B
Section:
Explanation:

QUESTION 124

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

A. Deploy an AWS Global Accelerator accelerator in front of the web servers.

B. Deploy an Amazon CloudFront web distribution in front of the S3 bucket.

C. Deploy an Amazon ElastiCache for Redis instance in front of the web servers.

D. Deploy an Amazon ElastiCache for Memcached instance in front of the web servers.

Correct Answer: B
Section:

QUESTION 125

A company hosts a marketing website in an on-premises data center. The website consists of static documents and runs on a single server. An administrator updates the website content infrequently and uses an SFTP client to upload new documents.
The company decides to host its website on AWS and to use Amazon CloudFront. The company's solutions architect creates a CloudFront distribution. The solutions architect must design the most cost-effective and resilient architecture for website hosting to serve as the CloudFront origin.
Which solution will meet these requirements?

A. Create a virtual server by using Amazon Lightsail. Configure the web server in the Lightsail instance. Upload website content by using an SFTP client.

B. Create an AWS Auto Scaling group for Amazon EC2 instances. Use an Application Load Balancer.
Upload website content by using an SFTP client.

C. Create a private Amazon S3 bucket. Use an S3 bucket policy to allow access from a CloudFront origin access identity (OAI). Upload website content by using theAWSCLI.

D. Create a public Amazon S3 bucket. Configure AWS Transfer for SFTP. Configure the S3 bucket for website hosting. Upload website content by using the SFTP client.

Correct Answer: C

Section:
Explanation:
https://docs.aws.amazon.com/cli/latest/reference/transfer/describe-server.html

QUESTION 126
A company is developing an ecommerce application that will consist of a load-balanced front end, a container-based application, and a relational database. A solutions architect needs to create a highly available solution that operates with as little manual intervention as possible.
Which solutions meet these requirements? (Select TWO.)

A. Create an Amazon RDS DB instance in Multi-AZ mode.
B. Create an Amazon RDS DB instance and one or more replicas in another Availability Zone.
C. Create an Amazon EC2 in stance-based Docker cluster to handle the dynamic application load.
D. Create an Amazon Elastic Container Service (Amazon ECS) cluster with a Fargate launch type to handle the dynamic application load.
E. Create an Amazon Elastic Container Service (Amazon ECS) cluster with an Amazon EC2 launch type to handle the dynamic application load.

Correct Answer: A, D
Section:
Explanation:
https://docs.aws.amazon.com/AmazonECS/latest/developerguide/Welcome.html
1. Relational database: RDS
2. Container-based applications: ECS
"Amazon ECS enables you to launch and stop your container-based applications by using simple API calls. You can also retrieve the state of your cluster from a centralized service and have access to many familiar Amazon EC2 features." 3. Little manual intervention: Fargate
You can run your tasks and services on a serverless infrastructure that is managed by AWS Fargate.
Alternatively, for more control over your infrastructure, you can run your tasks and services on a cluster of Amazon EC2 instances that you manage.

QUESTION 127
A company is designing a cloud communications platform that is driven by APIs. The application is hosted on Amazon EC2 instances behind a Network Load Balancer (NLB). The company uses Amazon API Gateway to provide external users with access to the application through APIs. The company wants to protect the platform against web exploits like SQL injection and also wants to detect and mitigate large, sophisticated DDoS attacks.
Which combination of solutions provides the MOST protection? (Select TWO.)

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

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

QUESTION 128

A company is running a batch application on Amazon EC2 instances. The application consists of a backend with multiple Amazon RDS databases. The application is causing a high number of leads on the databases. A solutions architect must reduce the number of database reads while ensuring high availability.

What should the solutions architect do to meet this requirement?

A. Add Amazon RDS read replicas
B. Use Amazon ElasbCache for Redls
C. Use Amazon Route 53 DNS caching
D. Use Amazon ElastiCache for Memcached

Correct Answer: A
Section:

QUESTION 129
A company has a web application that is based on Java and PHP. The company plans to move the application from on premises to AWS. The company needs the ability to test new site features frequently. The company also needs a highly available and managed solution that requires minimum operational overhead.
Which solution will meet these requirements?

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

Correct Answer: B
Section:

QUESTION 130
A company has a Microsoft NET application that runs on an on-premises Windows Server Trie application stores data by using an Oracle Database Standard Edition server The company is planning a migration to AWS and wants to minimize development changes while moving the application The
AWS application environment should be highly available. Which combination of actions should the company take to meet these requirements? (Select TWO)

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

Correct Answer: B, E
Section:

QUESTION 131
A rapidly growing ecommerce company is running its workloads in a single AWS Region. A solutions architect must create a disaster recovery (DR) strategy that includes a different AWS Region. The company wants its database to be up to date in the DR Region with the least possible latency. The remaining infrastructure in the DR Region needs to run at reduced capacity and must be able to scale up if necessary. Which solution will meet these requirements with the LOWEST recovery time objective (RTO)?

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

Correct Answer: B

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

QUESTION 132
A company's order system sends requests from clients to Amazon EC2 instances. The EC2 instances process the orders and then store the orders in a database on Amazon RDS. Users report that they must reprocess orders when the system fails. The company wants a resilient solution that can process orders automatically if a system outage occurs. What should a solutions architect do to meet these requirements?

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

Correct Answer: C
Section:

QUESTION 133
A company runs an application on a large fleet of Amazon EC2 instances. The application reads and write entries into an Amazon DynamoDB table. The size of the DynamoDB table continuously grows, but the application needs only data from the last 30 days. The company needs a solution that minimizes cost and development effort.
Which solution meets these requirements?

A. Use an AWS CloudFormation template to deploy the complete solution. Redeploy the CloudFormation stack every 30 days, and delete the original stack.
B. Use an EC2 Instance that runs a monitoring application from AWS Marketplace Configure the monitoring application to use Amazon DynamoDB Streams to store the timestamp when a new item is created in the table Use a script that runs on the EC2 instance to delele items that have a timestamp that is older than 30 days
C. Configure Amazon DynamoDB Streams to invoke an AWS Lambda function when a new item is created in the table Configure the Lambda function to delete items in the table that are older than 30 days
D. Extend the application to add an attribute that has a value of the current timestamp plus 30 days to each new item that is created in the (able Configure DynamoDB to use the attribute as (he TTL attribute

Correct Answer: D
Section:
Explanation:
Amazon DynamoDB Time to Live (TTL) allows you to define a per-item timestamp to determine when an item is no longer needed. Shortly after the date and time of the specified timestamp, DynamoDB deletes the item from your table without consuming any write throughput. TTL is provided at no extra cost as a means to reduce stored data volumes by retaining only the items that remain current for your workload's needs. TTL is useful if you store items that lose relevance after a specific time. The following are example TTL use cases:
Remove user or sensor data after one year of inactivity in an application.
Archive expired items to an Amazon S3 data lake via Amazon DynamoDB Streams and AWS Lambda.
Retain sensitive data for a certain amount of time according to contractual or regulatory obligations. https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/TTL.html

QUESTION 134
A company runs a containerized application on a Kubernetes cluster in an on-premises data center. The company is using a MongoDB database for data storage. The company wants to migrate some of these environments to AWS, but no code changes or deployment method changes are possible at this time. The company needs a solution that minimizes operational overhead. Which solution meets these requirements?

A. Use Amazon Elastic Container Service (Amazon ECS) with Amazon EC2 worker nodes for compute and MongoOB on EC2 for data storage
B. Use Amazon Elastic Container Service (Amazon ECS) with AWS Fargate for compute and Amazon DynamoDB tor data storage
C. Use Amazon Elastic Kubernetes Service (Amazon EKS) with Amazon EC2 worker nodes for compute and Amazon DynamoDB for data storage
D. Use Amazon Elastic Kubernetes Service (Amazon EKS) with AWS Fargate for compute and Amazon DocumentDB (with MongoDB compatibility) for data storage.

Correct Answer: D

Section:
Explanation:
Answer: D
Amazon DocumentDB (with MongoDB compatibility) is a fast, reliable, and fully managed database service. Amazon DocumentDB makes it easy to set up, operate, and scale MongoDB-compatible databases in the cloud. With Amazon DocumentDB, you can run the same application code and use the same drivers and tools that you use with MongoDB.
https://docs.aws.amazon.com/documentdb/latest/developerguide/what-is.html

QUESTION 135
A company selves a dynamic website from a flee! of Amazon EC2 instances behind an Application Load Balancer (ALB) The website needs to support multiple languages to serve customers around the world The website's architecture is running in the us-west-1 Region and is exhibiting high request latency tor users that are located in other parts of the world The website needs to serve requests quickly and efficiently regardless of a user's location However the company does not want to recreate the existing architecture across multiple Regions
What should a solutions architect do to meet these requirements?

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

Correct Answer: B
Section:

QUESTION 136
A telemarketing company is designing its customer call center functionality on AWS. The company needs a solution that provides multiple speaker recognition and generates transcript files. The company wants to query the transcript files to analyze the business patterns. The transcript files must be stored for 7 years for auditing purposes.
Which solution will meet these requirements?

A. Use Amazon Rekognition for multiple speaker recognition. Store the transcript files in Amazon S3. Use machine learning models for transcript file analysis.
B. Use Amazon Transcribe for multiple speaker recognition. Use Amazon Athena for transcript file analysis.
C. Use Amazon Translate for multiple speaker recognition. Store the transcript files in Amazon Redshift. Use SQL queries for transcript file analysis.
D. Use Amazon Rekognition for multiple speaker recognition. Store the transcript files in Amazon S3. Use Amazon Textract for transcript file analysis.

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

QUESTION 137
A company is building a new dynamic ordering website. The company wants to minimize server maintenance and patching. The website must be highly available and must scale read and write capacity as quickly as possible to meet changes in user demand.
Which solution will meet these requirements?

A. Host static content in Amazon S3 Host dynamic content by using Amazon API Gateway and AWS Lambda Use Amazon DynamoDB with on-demand capacity for the database Configure Amazon CtoudFront to deliver the website content
B. Host static content in Amazon S3 Host dynamic content by using Amazon API Gateway and AWS Lambda Use Amazon Aurora with Aurora Auto Scaling for the database Configure Amazon CloudFront to deliver the website content

C. Host al the website content on Amazon EC2 instances Create an Auto Scaling group to scale the EC2 Instances Use an Application Load Balancer to distribute traffic Use Amazon DynamoDB with provisioned write capacity for the database

D. Host at the website content on Amazon EC2 instances Create an Auto Scaling group to scale the EC2 instances Use an Application Load Balancer to distribute traffic Use Amazon Aurora with Aurora Auto Scaling for the database

Correct Answer: A
Section:

QUESTION 138
A company hosts its application on AWS The company uses Amazon Cognito to manage users When users log in to the application the application fetches required data from Amazon DynamoOB by using a REST API that is hosted in Amazon API Gateway. The company wants an AWS managed solution that will control access to the REST API to reduce development efforts
Which solution will meet these requirements with the LEAST operational overhead?

A. Configure an AWS Lambda function to be an authorize! in API Gateway to validate which user made the request

B. For each user, create and assign an API key that must be sent with each request Validate the key by using an AWS Lambda function

C. Send the user's email address in the header with every request Invoke an AWS Lambda function to validate that the user with that email address has proper access

D. Configure an Amazon Cognito user pool authorizer in API Gateway to allow Amazon Cognito to validate each request

Correct Answer: D
Section:

QUESTION 139
A company must migrate 20 TB of data from a data center to the AWS Cloud within 30 days. The company's network bandwidth is limited to 15 Mbps and cannot exceed 70% utilization. What should a solutions architect do to meet these requirements?

A. Use AWS Snowball.

B. Use AWS DataSync.

C. Use a secure VPN connection.

D. Use Amazon S3 Transfer Acceleration.

Correct Answer: A
Section:
Explanation:
AWS Snowball is a secure data transport solution that accelerates moving large amounts of data into and out of the AWS cloud. It can move up to 80 TB of data at a time, and provides a network bandwidth of up to 50 Mbps, so it is well- suited for the task. Additionally, it is secure and easy to use, making it the ideal solution for this migration.

QUESTION 140
A company is experiencing sudden increases in demand. The company needs to provision large Amazon EC2 instances from an Amazon Machine image (AMI) The instances will run m an Auto Scaling group. The company needs a solution that provides minimum initialization latency to meet the demand.
Which solution meets these requirements?

A. Use the aws ec2 register-image command to create an AMI from a snapshot Use AWS Step Functions to replace the AMI in the Auto Scaling group

B. Enable Amazon Elastic Block Store (Amazon EBS) fast snapshot restore on a snapshot Provision an AMI by using the snapshot Replace the AMI m the Auto Scaling group with the new AMI

C. Enable AMI creation and define lifecycle rules in Amazon Data Lifecycle Manager (Amazon DLM) Create an AWS Lambda function that modifies the AMI in the Auto Scaling group

D. Use Amazon EventBridge (Amazon CloudWatch Events) to invoke AWS Backup lifecycle policies that provision AMIs Configure Auto Scaling group capacity limits as an event source in EventBridge

Correct Answer: B
Section:

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

QUESTION 141
What should a solutions architect do to ensure that all objects uploaded to an Amazon S3 bucket are encrypted?

A. Update the bucket policy to deny if the PutObject does not have an s3 x-amz-acl header set
B. Update the bucket policy to deny if the PutObject does not have an s3:x-amz-aci header set to private.
C. Update the bucket policy to deny if the PutObject does not have an aws SecureTransport header set to true
D. Update the bucket policy to deny if the PutObject does not have an x-amz-server-sideencryption header set.

Correct Answer: D
Section:
Explanation:
https://aws.amazon.com/blogs/security/how-to-prevent-uploads-of-unencrypted-objects-toamazons3/#:~:text=Solution%20overview,console%2C%20CLI%2C%20or%20SDK.&text=To%20encrypt%20an%20object%20at,S3%2C%20or% 20SSE%2DKMS.

QUESTION 142
A company uses a legacy application to produce data in CSV format The legacy application stores the output data In Amazon S3 The company is deploying a new commercial off-the-shelf (COTS) application that can perform complex SQL queries to analyze data that is stored Amazon Redshift and Amazon S3 only However the COTS application cannot process the csv files that the legacy application produces The company cannot update the legacy application to produce data in another format The company needs to implement a solution so that the COTS application can use the data that the legacy applicator produces. Which solution will meet these requirements with the LEAST operational overhead?

A. Create a AWS Glue extract, transform, and load (ETL) job that runs on a schedule. Configure the ETL job to process the .csv files and store the processed data in Amazon Redshit.
B. Develop a Python script that runs on Amazon EC2 instances to convert the. csv files to sql files invoke the Python script on cron schedule to store the output files in Amazon S3.
C. Create an AWS Lambda function and an Amazon DynamoDB table. Use an S3 event to invoke the Lambda function. Configure the Lambda function to perform an extract transform, and load (ETL) job to process the .csv files and store the processed data in the DynamoDB table.
D. Use Amazon EventBridge (Amazon CloudWatch Events) to launch an Amazon EMR cluster on a weekly schedule. Configure the EMR cluster to perform an extract, tractform, and load (ETL) job to process the .csv files and store the processed data in an Amazon Redshift table.

Correct Answer: A
Section:
Explanation:
This solution meets the requirements of implementing a solution so that the COTS application can use the data that the legacy application produces with the least operational overhead. AWS Glue is a fully managed service that provides a serverless ETL platform to prepare and load data for analytics. AWS Glue can process data in various formats, including .csv files, and store the processed data in Amazon Redshift, which is a fully managed data warehouse service that supports complex SQL queries. AWS Glue can run ETL jobs on a schedule, which can automate the data processing and loading process. Option B is incorrect because developing a Python script that runs on Amazon EC2 instances to convert the .csv files to sql files can increase the operational overhead and complexity, and it may not provide consistent data processing and loading for the COTS application. Option C is incorrect because creating an AWS Lambda function and an Amazon DynamoDB table to process the .csv files and store the processed data in the DynamoDB table does not meet the requirement of using Amazon Redshift as the data source for the COTS application. Option D is incorrect because using Amazon EventBridge (Amazon CloudWatch Events) to launch an Amazon EMR cluster on a weekly schedule to process the .csv files and store the processed data in an Amazon Redshift table can increase the operational overhead and complexity, and it may not provide timely data processing and loading for the COTS application. https://aws.amazon.com/glue/

QUESTION 143
A company has an On-premises volume backup solution that has reached its end of life. The company wants to use AWS as part of a new backup solution and wants to maintain local access to all the data while it is backed up on AWS. The company wants to ensure that the data backed up on AWS is automatically and securely transferred.
Which solution meets these requirements?

A. Use AWS Snowball to migrate data out of the on-premises solution to Amazon S3. Configure onpremises systems to mount the Snowball S3 endpoint to provide local access to the data.

B. Use AWS Snowball Edge to migrate data out of the on-premises solution to Amazon S3.Use the Snowball Edge file interface to provide on-premises systems with local access to the data.

C. Use AWS Storage Gateway and configure a cached volume gateway. Run the Storage Gateway software application on premises and configure a percentage of data to cache locally. Mount the gateway storage volumes to provide local access to the data.

D. Use AWS Storage Gateway and configure a stored volume gateway. Run the Storage software application on premises and map the gateway storage volumes to on-premises storage. Mount the gateway storage volumes to provide local access to the data.

Correct Answer: D
Section:
Explanation:

QUESTION 144

A transaction processing company has weekly scripted batch jobs that run on Amazon EC2 instances.
The EC2 instances are in an Auto Scaling group. The number of transactions can vary but the beseline CPU utilization that is noted on each run is at least 60%. The company needs to provision the capacity 30 minutes before the jobs run. Currently engineering complete this task by manually modifying the Auto Scaling group parameters.
The company does not have the resources to analyze the required capacity trends for the Auto Scaling group counts. The company needs an automated way to modify the Auto Scaling group's capacity. Which solution will meet these requiements with the LEAST operational overhead?

A. Ceate a dynamic scalling policy for the Auto Scaling group. Configure the policy to scale based on the CPU utilization metric to 60%.

B. Create a scheduled scaling polcy for the Auto Scaling group. Set the appropriate desired capacity, minimum capacity, and maximum capacity. Set the recurrence to weekly. Set the start time to 30 minutes. Before the batch jobs run.

C. Create a predictive scaling policy for the Auto Scaling group. Configure the policy to scale based on forecast. Set the scaling metric to CPU utilization. Set the target value for the metric to 60%. In the Policy, set the instances to pre- launch 30 minutes before the jobs run.

D. Create an Amazon EventBridge event to invoke an AWS Lamda function when the CPU utilization metric value for the Auto Scaling group reaches 60%. Configure the Lambda function to increase the Auto Scaling group's desired capacity and maximum capacity by 20%.

Correct Answer: C
Section:

QUESTION 145

A company experienced a breach that affected several applications in its on-premises data center The attacker took advantage of vulnerabilities in the custom applications that were running on the servers The company is now migrating its applications to run on Amazon EC2 instances The company wants to implement a solution that actively scans for vulnerabilities on the EC2 instances and sends a report that details the findings Which solution will meet these requirements?

A. Deploy AWS Shield to scan the EC2 instances for vulnerabilities Create an AWS Lambda function to log any findings to AWS CloudTrail.

B. Deploy Amazon Macie and AWS Lambda functions to scan the EC2 instances for vulnerabilities Log any findings to AWS CloudTrail

C. Turn on Amazon GuardDuty Deploy the GuardDuty agents to the EC2 instances Configure an AWS Lambda function to automate the generation and distribution of reports that detail the findings

D. Turn on Amazon Inspector Deploy the Amazon Inspector agent to the EC2 instances Configure an AWS Lambda function to automate the generation and distribution of reports that detail the findings

Correct Answer: D
Section:

QUESTION 146

A company stores confidential data in an Amazon Aurora PostgreSQL database in the ap-southeast-3 Region The database is encrypted with an AWS Key Management Service (AWS KMS) customer managed key The company was recently acquired and must securely share a backup of the database with the acquiring company's AWS account in ap-southeast-3. What should a solutions architect do to meet these requirements?

A. Create a database snapshot Copy the snapshot to a new unencrypted snapshot Share the new snapshot with the acquiring company's AWS account

B. Create a database snapshot Add the acquiring company's AWS account to the KMS key policy Share the snapshot with the acquiring company's AWS account

C. Create a database snapshot that uses a different AWS managed KMS key Add the acquiring company's AWS account to the KMS key alias. Share the snapshot with the acquiring company's AWS account.

D. Create a database snapshot Download the database snapshot Upload the database snapshot to an Amazon S3 bucket Update the S3 bucket policy to allow access from the acquiring company's AWS account

Correct Answer: B
Section:
Explanation:
https://docs.aws.amazon.com/kms/latest/developerguide/key-policy-modifying-external-accounts.html There's no need to create another custom AWS KMS key. https://aws.amazon.com/premiumsupport/knowledge-center/aurora-share-encrypted-snapshot/ Give target account access to the custom AWS KMS key within the source account 1. Log in to the source account, and go to the AWS KMS console in the same Region as the DB cluster snapshot. 2. Select Customer-managed keys from the navigation pane. 3. Select your custom AWS KMS key (ALREADY CREATED) 4. From the Other AWS accounts section, select Add another AWS account, and then enter the AWS account number of your target account. Then: Copy and share the DB cluster snapshot

QUESTION 147
A company is launching an application on AWS. The application uses an Application Load (ALB) to direct traffic to at least two Amazon EC2 instances in a single target group. The instances are in an Auto Scaling group for each environment. The company requires a development and a production environment. The production environment will have periods of high traffic. Which solution will configure the development environment MOST cost-effectively?

A. Reconfigure the target group in the development environment to have one EC2 instance as a target.
B. Change the ALB balancing algorithm to least outstanding requests.
C. Reduce the size of the EC2 instances in both environments.
D. Reduce the maximum number of EC2 instances in the development environment’s Auto Scaling group

Correct Answer: D
Section:
Explanation:
This option will configure the development environment in the most cost-effective way as it reduces the number of instances running in the development environment and therefore reduces the cost of running the application. The development environment typically requires less resources than the production environment, and it is unlikely that the development environment will have periods of high traffic that would require a large number of instances. By reducing the maximum number of instances in the development environment's Auto Scaling group, the company can save on costs while still maintaining a functional development environment.

QUESTION 148
A company will deployed a web application on AWS. The company hosts the backend database on Amazon RDS for MySQL with a primary DB instance and five read replicas to support scaling needs. The read replicas must log no more than 1 second bahind the primary DB Instance. The database routinely runs scheduled stored procedures. As traffic on the website increases, the replicas experinces addtional lag during periods of peak lead. A solutions architect must reduce the replication lag as much as possible. The solutions architect must minimize changes to the applicatin code and must minimize ongoing overhead. Which solution will meet these requirements? Migrate the database to Amazon Aurora MySQL. Replace the read replicas with Aurora Replicas, and configure Aurora Auto Scaling. Replace the stored procedures with Aurora MySQL native functions. Deploy an Amazon ElasticCache for Redis cluser in front of the database. Modify the application to check the cache before the application queries the database. Repace the stored procedures with AWS Lambda funcions.

A. Migrate the database to a MYSQL database that runs on Amazn EC2 instances. Choose large, compute optimized for all replica nodes. Maintain the stored procedures on the EC2 instances.
B. Deploy an Amazon ElastiCache for Redis cluster in fornt of the database. Modify the application to check the cache before the application queries the database. Replace the stored procedures with AWS Lambda functions.
C. Migrate the database to a MySQL database that runs on Amazon EC2 instances. Choose large, compute optimized EC2 instances for all replica nodes, Maintain the stored procedures on the EC2 instances.
D. Migrate the database to Amazon DynamoDB, Provision number of read capacity units (RCUs) to support the required throughput, and configure on-demand capacity scaling. Replace the stored procedures with DynamoDB streams.

Correct Answer: A
Section:

QUESTION 149
A company hosts a multiplayer gaming application on AWS. The company wants the application to read data with sub-millisecond latency and run one-time queries on historical data. Which solution will meet these requirements with the LEAST operational overhead?

A. Use Amazon RDS for data that is frequently accessed. Run a periodic custom script to export the data to an Amazon S3 bucket.
B. Store the data directly in an Amazon S3 bucket. Implement an S3 Lifecycle policy to move older data to S3 Glacier Deep Archive for long-term storage. Run one-time queries on the data in Amazon S3 by using Amazon

Athena

C. Use Amazon DynamoDB with DynamoDB Accelerator (DAX) for data that is frequently accessed.
Export the data to an Amazon S3 bucket by using DynamoDB table export. Run one-time queries on the data in Amazon S3 by using Amazon Athena.

D. Use Amazon DynamoDB for data that is frequently accessed Turn on streaming to Amazon Kinesis Data Streams. Use Amazon Kinesis Data Firehose to read the data from Kinesis Data Streams. Store the records in an Amazon S3 bucket.

Correct Answer: C
Section:

QUESTION 150
A company collects data from thousands of remote devices by using a RESTful web services application that runs on an Amazon EC2 instance. The EC2 instance receives the raw data, transforms the raw data, and stores all the data in an Amazon S3 bucket. The number of remote devices will increase into the millions soon. The company needs a highly scalable solution that minimizes operational overhead. Which combination of steps should a solutions architect take to meet these requirements9 (Select TWO.)

A. Use AWS Glue to process the raw data in Amazon S3.

B. Use Amazon Route 53 to route traffic to different EC2 instances.

C. Add more EC2 instances to accommodate the increasing amount of incoming data.

D. Send the raw data to Amazon Simple Queue Service (Amazon SOS). Use EC2 instances to process the data.

E. Use Amazon API Gateway to send the raw data to an Amazon Kinesis data stream. Configure Amazon Kinesis Data Firehose to use the data stream as a source to deliver the data to Amazon S3.

Correct Answer: A, E
Section:
Explanation:
"RESTful web services" => API Gateway. "EC2 instance receives the raw data, transforms the raw data, and stores all the data in an Amazon S3 bucket" => GLUE with (Extract - Transform - Load)

QUESTION 151
A company collects data from a large number of participants who use wearabledevices.The company stores the data in an Amazon DynamoDB table and uses applications to analyze the dat a. The data workload is constant and predictable. The company wants to stay at or below its forecasted budget for DynamoDB.
Whihc solution will meet these requirements MOST cost-effectively?

A. Use provisioned mode and DynamoDB Standard-Infrequent Access (DynamoDB Standard-IA).
Reserve capacity for the forecasted workload.

B. Use provisioned mode Specify the read capacity units (RCUs) and write capacity units (WCUs).

C. Use on-demand mode. Set the read capacity unite (RCUs) and write capacity units (WCUs) high enough to accommodate changes in the workload.

D. Use on-demand mode. Specify the read capacity units (RCUs) and write capacity units (WCUs) with reserved capacity.

Correct Answer: B
Section:
Explanation:
This option is the most efficient because it uses provisioned mode, which is a read/write capacity mode for processing reads and writes on your tables that lets you specify how much read and write throughput you expect your application to perform1. It also specifies the read capacity units (RCUs) and write capacity units (WCUs), which are the amount of data your application needs to read or write per second. It also meets the requirement of staying at or below its forecasted budget for DynamoDB, as provisioned mode has lower costs than on-demand mode for predictable workloads. This solution meets the requirement of collecting data from a large number of participants who use wearable devices with a constant and predictable data workload. Option A is less efficient because it uses provisioned mode and DynamoDB Standard-Infrequent Access (DynamoDB Standard-IA), which is a storage class for infrequently accessed items that require milliseconds latency2. However, this does not meet the requirement of collecting data from a large number of participants who use wearable devices with a constant and predictable data workload, as DynamoDB Standard-IA is more suitable for items that are accessed less frequently than once every 30 days. Option C is less efficient because it uses on-demand mode, which is a read/write capacity mode that lets you pay only for what you use by automatically adjusting your table's capacity in response to changing demand3. However, this does not meet the requirement of staying at or below its forecasted budget for DynamoDB, as on-demand mode has higher costs than provisioned mode for predictable workloads. Option D is less efficient because it uses on- demand mode and specifies the RCUs and WCUs with reserved capacity, which is a way to reserve read and write capacity for your tables in exchange for discounted hourly rates. However, this does not meet the requirement of staying at or below its forecasted budget for DynamoDB, as on-demand mode has higher costs than provisioned mode for predictable workloads. Also, specifying RCUs and WCUs with reserved capacity is not possible with on-demand mode, as it only applies to provisioned mode.

QUESTION 152

A company hostss a three application on Amazon EC2 instances in a single Availability Zone. The web application uses a self-managed MySQL database that is hosted on an EC2 instances to store data in an Amazon Elastic Block Store (Amazon EBS) volumn. The MySQL database currently uses a 1 TB Provisioned IOPS SSD (io2) EBS volume. The company expects traffic of 1,000 IOPS for both reads and writes at peak traffic. The company wants to minimize any distruptions, stabilize perperformace, and reduce costs while retaining the capacity for double the IOPS. The company wants to more the database tier to a fully managed solution that is highly available and fault tolerant.

Which solution will meet these requirements MOST cost-effectively?

A. Use a Multi-AZ deployment of an Amazon RDS for MySQL DB instance with an io2 Block Express EBS volume.

B. Use a Multi-AZ deployment of an Amazon RDS for MySQL DB instance with a General Purpose SSD (gp2) EBS volume.

C. Use Amazon S3 Intelligent-Tiering access tiers.

D. Use two large EC2 instances to host the database in active-passive mode.

Correct Answer: B

Section:

Explanation:

RDS supported Storage > https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/CHAP_Storage.html GP2 max IOPS > https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/general-purpose.html#gp2-performance Amazon RDS provides three storage types: General Purpose SSD (also known as gp2 and gp3), Provisioned IOPS SSD (also known as io1), and magnetic (also known as standard). They differ in performance characteristics and price, which means that you can tailor your storage performance and cost to the needs of your database workload. You can create MySQL, MariaDB, Oracle, and PostgreSQL RDS DB instances with up to 64 tebibytes (TiB) of storage. You can create SQL Server RDS DB instances with up to 16 TiB of storage. For this amount of storage, use the Provisioned IOPS SSD and General Purpose SSD storage types. https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/CHAP_Storage.html

QUESTION 153

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

A. Use AWS Key Management Service (AWS KMS) certificates on the ALB to encrypt data in transit.
Use AWS Certificate Manager (ACM) to encrypt the EBS volumes and Aurora database storage at rest.

B. Use the AWS root account to log in to the AWS Management Console. Upload the company's encryption certificates. While in the root account, select the option to turn on encryption for all data at rest and in transit for the account.

C. Use a AWS Key Management Service (AWS KMS) to encrypt the EBS volumes and Aurora database storage at rest. Attach an AWS Certificate Manager (ACM) certificate to the ALB to encrypt data in transit.

D. Use BitLocker to encrypt all data at rest. Import the company's TLS certificate keys to AWS key Management Service (AWS KMS). Attach the KMS keys to the ALB to encrypt data in transit.

Correct Answer: C

Section:

QUESTION 154

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

A. Extend the Auto Scaling groups for the web tier and the application tier to deploy instances in Availability Zones in a second Region. Use an Aurora global database to deploy the database in the primary Region and the second Region. Use Amazon Route 53 health checks with a failover routing policy to the second Region.

B. Deploy the web tier and the application tier to a second Region. Add an Aurora PostgreSQL cross- Region Aurara Replica in the second Region. Use Amazon Route 53 health checks with a failovers routing policy to the second Region, Promote the secondary to primary as needed.

C. Deploy the web tier and the applicatin tier to a second Region. Create an Aurora PostSQL database in the second Region. Use AWS Database Migration Service (AWS DMS) to replicate the primary database to the second Region. Use Amazon Route 53 health checks with a failover routing policy to the second Region.

D. Deploy the web tier and the application tier to a second Region. Use an Amazon Aurora global database to deploy the database in the primary Region and the second Region. Use Amazon Route 53 health checks with a failover routing policy to the second Region. Promote the secondary to primary as needed.

Correct Answer: D

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

QUESTION 155
A company sells datasets to customers who do research in artificial intelligence and machine learning (Al/ML) The datasets are large, formatted files that are stored in an Amazon S3 bucket in the us-east- 1 Region The company hosts a web application that the customers use to purchase access to a given dataset The web application is deployed on multiple Amazon EC2 instances behind an Application Load Balancer After a purchase is made customers receive an S3 signed URL that allows access to the files.
The customers are distributed across North America and Europe The company wants to reduce the cost that is associated with data transfers and wants to maintain or improve performance. What should a solutions architect do to meet these requirements?

A. Configure S3 Transfer Acceleration on the existing S3 bucket Direct customer requests to the S3 Transfer Acceleration endpoint Continue to use S3 signed URLs for access control
B. Deploy an Amazon CloudFront distribution with the existing S3 bucket as the origin Direct customer requests to the CloudFront URL Switch to CloudFront signed URLs for access control
C. Set up a second S3 bucket in the eu-central-1 Region with S3 Cross-Region Replication between the buckets Direct customer requests to the closest Region Continue to use S3 signed URLs for access control
D. Modify the web application to enable streaming of the datasets to end users. Configure the web application to read the data from the existing S3 bucket Implement access control directly in the application

Correct Answer: B
Section:
Explanation:
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/PrivateContent.html

QUESTION 156
A company wants to configure its Amazon CloudFront distribution to use SSL/TLS certificates. The company does not want to use the default domain name for the distribution. Instead, the company wants to use a different domain name for the distribution.
Which solution will deploy the certificate with icurring any additional costs?

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

Correct Answer: C
Section:
Explanation:

QUESTION 157
A solution architect is designing a company’s disaster recovery (DR) architecture. The company has a MySQL database that runs on an Amazon EC2 instance in a private subnet with scheduled backup. The DR design to

include multiple AWS Regions.
Which solution will meet these requiements with the LEAST operational overhead?

A. Migrate the MySQL database to multiple EC2 instances. Configure a standby EC2 instance in the DR Region Turn on replication.
B. Migrate the MySQL database to Amazon RDS. Use a Multi-AZ deployment. Turn on read replication for the primary DB instance in the different Availability Zones.
C. Migrate the MySQL database to an Amazon Aurora global database. Host the primary DB cluster in the primary Region. Host the secondary DB cluster in the DR Region.
D. Store the schedule backup of the MySQL database in an Amazon S3 bucket that is configured for S3 Cross-Region Replication (CRR). Use the data backup to restore the database in the DR Region.

Correct Answer: C
Section:

QUESTION 158
A rapidly growing global ecommerce company is hosting its web application on AWS. The web application includes static content and dynamic content. The website stores online transaction processing (OLTP) data in an Amazon RDS database. The website's users are experiencing slow page loads.
Which combination of actions should a solutions architect take to resolve this issue? (Select TWO.)

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

Correct Answer: B, D
Section:

QUESTION 159
A company has an application that is backed ny an Amazon DynamoDB table. The company's compliance requirements specify that database backups must be taken every month, must be available for 6 months, and must be retained for 7 years.
Which solution will meet these requirements?

A. Create an AWS Backup plan to back up the DynamoDB table on the first day of each month. Specify a lifecycle policy that transitions the backup to cold storage after 6 months. Set the retention period for each backup to 7 years.
B. Create a DynamoDB on-damand backup of the DynamoDB table on the first day of each month Transition the backup to Amazon S3 Glacier Flexible Retrieval after 6 months. Create an S3 Lifecycle policy to delete backups that are older than 7 years.
C. Use the AWS SDK to develop a script that creates an on-demand backup of the DynamoDB table.
Set up an Amzon EvenlBridge rule that runs the script on the first day of each month. Create a second script that will run on the second day of each month to transition DynamoDB backups that are older than 6 months to cold storage and to delete backups that are older than 7 years.
D. Use the AWS CLI to create an on-demand backup of the DynamoDB table Set up an Amazon EventBridge rule that runs the command on the first day of each month with a cron expression Specify in the command to transition the backups to cold storage after 6 months and to delete the backups after 7 years.

Correct Answer: A
Section:

QUESTION 160
A solutions architect wants all new users to have specific complexity requirements and mandatory rotation periods tor IAM user passwords What should the solutions architect do to accomplish this?

A. Set an overall password policy for the entire AWS account
B. Set a password policy for each IAM user in the AWS account

C. Use third-party vendor software to set password requirements
D. Attach an Amazon CloudWatch rule to the Create_newuser event to set the password with the appropriate requirements

Correct Answer: A
Section:

QUESTION 161
A company wants to deploy a new public web application on AWS The application includes a web server tier that uses Amazon EC2 instances The application also includes a database tier that uses an Amazon RDS for MySQL DB instance The application must be secure and accessible for global customers that have dynamic IP addresses How should a solutions architect configure the security groups to meet these requirements'?

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

Correct Answer: A
Section:

QUESTION 162
A company is planning to migrate a commercial off-the-shelf application from is on-premises data center to AWS. The software has a software licensing model using sockets and cores with predictable capacity and uptime requirements. The company wants to use its existing licenses, which were purchased earlier this year.
Which Amazon EC2 pricing option is the MOST cost-effective?

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

Correct Answer: A
Section:

QUESTION 163
An ecommerce company is experiencing an increase in user traffic. The company's store is deployed on Amazon EC2 instances as a two-tier web application consisting of a web tier and a separate database tier. As traffic increases, the company notices that the architecture is causing significant delays in sending timely marketing and order confirmation email to users. The company wants to reduce the time it spends resolving complex email delivery issues and minimize operational overhead.
What should a solutions architect do to meet these requirements?

A. Create a separata application tier using EC2 instances dedicated to email processing.
B. Configure the web instance to send email through Amazon Simple Email Service (Amazon SES).
C. Configure the web instance to send email through Amazon Simple Notification Service (Amazon SNS)
D. Create a separate application tier using EC2 instances dedicated to email processing. Place the instances in an Auto Scaling group.

Correct Answer: B
Section:

QUESTION 164

A company is deploying a two-tier web application in a VPC. The web tier is using an Amazon EC2 Auto Scaling group with public subnets that span multiple Availability Zones. The database tier consists of an Amazon RDS for MySQL DB instance in separate private subnets. The web tier requires access to the database to retrieve product information. The web application is not working as intended. The web application reports that it cannot connect to the database. The database is confirmed to be up and running. All configurations for the network ACLs. security groups, and route tables are still in their default states.

What should a solutions architect recommend to fix the application?

A. Add an explicit rule to the private subnet's network ACL to allow traffic from the web tier's EC2 instances.

B. Add a route in the VPC route table to allow traffic between the web tier's EC2 instances and Ihe database tier.

C. Deploy the web tier's EC2 instances and the database tier's RDS instance into two separate VPCs. and configure VPC peering.

D. Add an inbound rule to the security group of the database tier's RDS instance to allow traffic from the web tier's security group.

Correct Answer: D

Section:

QUESTION 165

A company is running a multi-tier recommence web application in the AWS Cloud. The application runs on Amazon EC2 instances with an Amazon RDS for MySQL Multi-AZ OB instance. Amazon ROS is configured with the latest generation DB instance with 2.000 GB of storage In a General Purpose SSD (gp3) Amazon Elastic Block Store (Amazon EBSl volume. The database performance affects the application during periods high demand. A database administrator analyzes the logs in Amazon CloudWatch Logs and discovers that the application performance always degrades when the number of read and write IOPS is higher than 20.000. What should a solutions architect do to improve the application performance?

A. Replace the volume with a magnetic volume.

B. Increase the number of IOPS on the gp3 volume.

C. Replace the volume with a Provisioned IOPS SSD (Io2) volume.

D. Replace the 2.000 GB gp3 volume with two 1.000 GB gp3 volumes

Correct Answer: C

Section:

QUESTION 166

A company is deploying a new application on Amazon EC2 instances. The application writes data to Amazon Elastic Block Store (Amazon EBS) volumes. The company needs to ensure that all data that is written to the EBS volumes is encrypted at rest.

Which solution wil meet this requirement?

A. Create an IAM role that specifies EBS encryption. Attach the role to the EC2 instances.

B. Create the EBS volumes as encrypted volumes Attach the EBS volumes to the EC2 instances.

C. Create an EC2 instance tag that has a key of Encrypt and a value of True. Tag all instances that require encryption at the ESS level.

D. Create an AWS Key Management Service (AWS KMS) key policy that enforces EBS encryption in the account Ensure that the key policy is active.

Correct Answer: B

Section:

QUESTION 167

A media company hosts its website on AWS. The website application's architecture includes a fleet of Amazon EC2 instances behind an Application Load Balancer (ALB) and a database that is hosted on Amazon Aurora The company's cyber security teem reports that the application is vulnerable to SOL injection.

How should the company resolve this issue?

A. Use AWS WAF in front of the ALB Associate the appropriate web ACLs with AWS WAF.

B. Create an ALB listener rule to reply to SQL injection with a fixed response

C. Subscribe to AWS Shield Advanced to block all SQL injection attempts automatically.

D. Set up Amazon Inspector to block all SOL injection attempts automatically

Correct Answer: A
Section:

QUESTION 168
A company is moving its data management application to AWS. The company wants to transition to an event-driven architecture. The architecture needs to the more distributed and to use serverless concepts whit performing the different aspects of the workflow. The company also wants to minimize operational overhead.
Which solution will meet these requirements?

A. Build out the workflow in AWS Glue Use AWS Glue to invoke AWS Lambda functions to process the workflow slaps

B. Build out the workflow in AWS Step Functions Deploy the application on Amazon EC2 Instances Use Step Functions to invoke the workflow steps on the EC2 instances

C. Build out the workflow in Amazon EventBridge. Use EventBridge to invoke AWS Lambda functions on a schedule to process the workflow steps.

D. Build out the workflow m AWS Step Functions Use Step Functions to create a stale machine Use the stale machine to invoke AWS Lambda functions to process the workflow steps

Correct Answer: D
Section:
Explanation:
This answer is correct because it meets the requirements of transitioning to an event-driven architecture, using serverless concepts, and minimizing operational overhead. AWS Step Functions is a serverless service that lets you coordinate multiple AWS services into workflows using state machines. State machines are composed of tasks and transitions that define the logic and order of execution of the workflow steps. AWS Lambda is a serverless function-as-a-service platform that lets you run code without provisioning or managing servers. Lambda functions can be invoked by Step Functions as tasks in a state machine, and can perform different aspects of the data management workflow, such as data

QUESTION 169

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

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

Correct Answer: D
Section:

QUESTION 170
A financial company hosts a web application on AWS. The application uses an Amazon API Gateway Regional API endpoint to give users the ability to retrieve current stock prices. The company's security team has noticed an increase in the number of API requests. The security team is concerned that HTTP flood attacks might take the application offline. A solutions architect must design a solution to protect the application from this type of attack. Which solution meats these requirements with the LEAST operational overhead?

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

Correct Answer: B
Section:

QUESTION 171
A payment processing company records all voice communication with its customers and stores the audio files in an Amazon S3 bucket. The company needs to capture the text from the audio files. The company must remove from the text any personally identifiable information (Pll) that belongs to customers.
What should a solutions architect do to meet these requirements?

A. Process the audio files by using Amazon Kinesis Video Streams. Use an AWS Lambda function to scan for known Pll patterns.
B. When an audio file is uploaded to the S3 bucket, invoke an AWS Lambda function to start an Amazon Textract task to analyze the call recordings.
C. Configure an Amazon Transcribe transcription job with Pll redaction turned on. When an audio file is uploaded to the S3 bucket, invoke an AWS Lambda function to start the transcription job. Store the output in a separate S3 bucket.
D. Create an Amazon Connect contact flow that ingests the audio files with transcription turned on. Embed an AWS Lambda function to scan for known Pll patterns. Use Amazon EventBridge (Amazon CloudWatch Events) to start the contact flow when an audio file is uploaded to the S3 bucket.

Correct Answer: C
Section:

QUESTION 172
A company is migrating its on-premises workload to the AWS Cloud. The company already uses several Amazon EC2 instances and Amazon RDS DB instances. The company wants a solution that automatically starts and stops the EC2 instances and D6 instances outside of business hours. The solution must minimize cost and infrastructure maintenance. Which solution will meet these requirement?

A. Scale the EC2 instances by using elastic resize Scale the DB instances to zero outside of business hours
B. Explore AWS Marketplace for partner solutions that will automatically start and stop the EC2 Instances and OB instances on a schedule
C. Launch another EC2 instance. Configure a crontab schedule to run shell scripts that will start and stop the existing EC2 instances and DB instances on a schedule.
D. Create an AWS Lambda function that will start and stop the EC2 instances and DB instances Configure Amazon EventBridge to invoke the Lambda function on a schedule

Correct Answer: D
Section:

QUESTION 173
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 174
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 175
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 176
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 177
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 178
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.

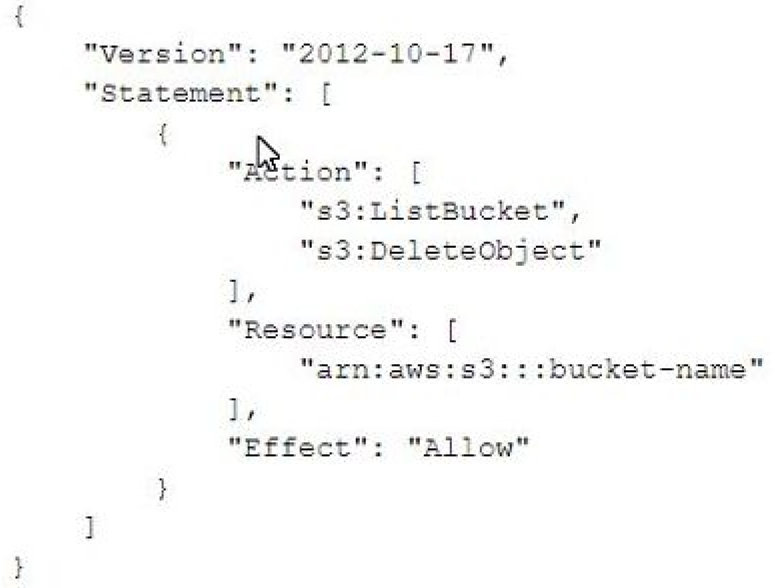

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Action": [
                "s3:ListBucket",
                "s3:DeleteObject"
            ],
            "Resource": [
                "arn:aws:s3:::bucket-name"
            ],
            "Effect": "Allow"
        }
    ]
}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

A)

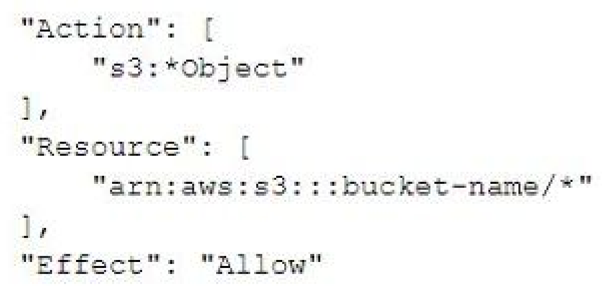

```
"Action": [
    "s3:*Object"
],
"Resource": [
    "arn:aws:s3:::bucket-name/*"
],
"Effect": "Allow"
```

B)

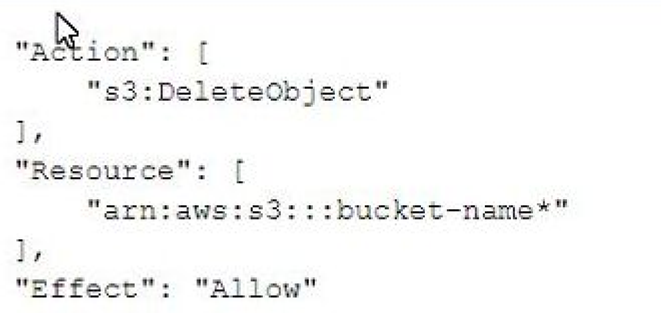

```
"Action": [
    "s3:DeleteObject"
],
"Resource": [
    "arn:aws:s3:::bucket-name*"
],
"Effect": "Allow"
```

C)

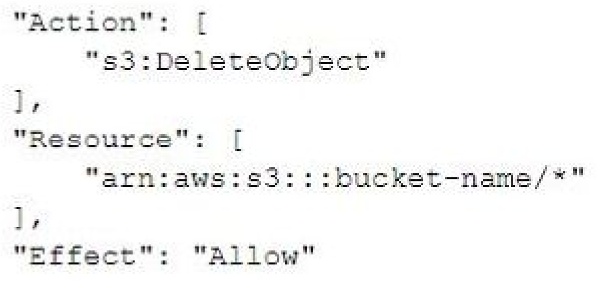

```
"Action": [
    "s3:DeleteObject"
],
"Resource": [
    "arn:aws:s3:::bucket-name/*"
],
"Effect": "Allow"
```

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 179

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 180

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 181

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 182

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 183
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 184

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 185

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 186

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 187

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 188
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 189
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 190
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 191
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 192
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 193

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 194

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 195

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 196
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 197
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 (Pll) 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 198
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 199
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
Section:
Explanation:
A service control policy (SCP) is a type of policy that you can use to manage permissions in your organization. SCPs offer central control over the maximum available permissions for all accounts in your organization, allowing you to ensure your accounts stay within your organization's access control guidelines. You can use an SCP to deny the ec2:CreateVolume action when the ec2:Encrypted condition equals false, which means that any user or role in the accounts under the root OU will not be able to create unencrypted EBS volumes. This solution will have minimal effect on employees who create EBS volumes, as they can still create encrypted volumes as needed.
Reference: https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps.html

QUESTION 200
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 201
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 202
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 203
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 204
A company has developed a new video game as a web application. The application is in a three-tier architecture in a VPC with Amazon RDS for MySQL in the database layer. Several players will compete concurrently online. The game's developers want to display a top-10 scoreboard in near-real time and offer the ability to stop and restore the game while preserving the current scores.
What should a solutions architect do to meet these requirements?

A. Set up an Amazon ElastiCache for Memcached cluster to cache the scores for the web application to display.
B. Set up an Amazon ElastiCache for Redis cluster to compute and cache the scores for the web application to display.
C. Place an Amazon CloudFront distribution in front of the web application to cache the scoreboard in a section of the application.
D. Create a read replica on Amazon RDS for MySQL to run queries to compute the scoreboard and serve the read traffic to the web application.

Correct Answer: B
Section:
Explanation:
This answer is correct because it meets the requirements of displaying a top-10 scoreboard in near-real time and offering the ability to stop and restore the game while preserving the current scores. Amazon ElastiCache for Redis is a blazing fast in-memory data store that provides sub-millisecond latency to power internet-scale real-time applications. You can use Amazon ElastiCache for Redis to set up an ElastiCache for Redis cluster to compute and cache the scores for the web application to display. You can use Redis data structures such as sorted sets and hashes to store and rank the scores of the players, and use Redis commands such as ZRANGE and ZADD to retrieve and update the scores efficiently. You can also use Redis persistence features such as snapshots and append-only files (AOF) to enable point-in-time recovery of your data, which can help you stop and restore the

game while preserving the current scores.
https://docs.aws.amazon.com/AmazonElastiCache/latest/red-ug/WhatIs.html
https://redis.io/topics/data-types
https://redis.io/topics/persistence

QUESTION 205

A company wants to share accounting data with an external auditor. The data is stored in an Amazon RDS DB instance that resides in a private subnet. The auditor has its own AWS account and requires its own copy of the database.
What is the MOST secure way for the company to share the database with the auditor?

A. Create a read replica of the database. Configure IAM standard database authentication to grant the auditor access.
B. Export the database contents to text files. Store the files in an Amazon S3 bucket. Create a new IAM user for the auditor. Grant the user access to the S3 bucket.
C. Copy a snapshot of the database to an Amazon S3 bucket. Create an IAM user. Share the user's keys with the auditor to grant access to the object in the $3 bucket.
D. Create an encrypted snapshot of the database. Share the snapshot with the auditor. Allow access to the AWS Key Management Service (AWS KMS) encryption key.

Correct Answer: D
Section:
Explanation:
This answer is correct because it meets the requirements of sharing the database with the auditor in a secure way. You can create an encrypted snapshot of the database by using AWS Key Management Service (AWS KMS) to encrypt the snapshot with a customer managed key. You can share the snapshot with the auditor by modifying the permissions of the snapshot and specifying the AWS account ID of the auditor. You can also allow access to the AWS KMS encryption key by adding a key policy statement that grants permissions to the auditor's account. This way, you can ensure that only the auditor can access and restore the snapshot in their own AWS account.
https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_ShareSnapshot.html
https://docs.aws.amazon.com/kms/latest/developerguide/key-policies.html#key-policy-default-allow-root-enable-iam

QUESTION 206

A retail company uses a regional Amazon API Gateway API for its public REST APIs. The API Gateway endpoint is a custom domain name that points to an Amazon Route 53 alias record. A solutions architect needs to create a solution that has minimal effects on customers and minimal data loss to release the new version of APIs.
Which solution will meet these requirements?

A. Create a canary release deployment stage for API Gateway. Deploy the latest API version. Point an appropriate percentage of traffic to the canary stage. After API verification, promote the canary stage to the production stage.
B. Create a new API Gateway endpoint with a new version of the API in OpenAPI YAML file format. Use the import-to-update operation in merge mode into the API in API Gateway. Deploy the new version of the API to the production stage.
C. Create a new API Gateway endpoint with a new version of the API in OpenAPI JSON file format. Use the import-to-update operation in overwrite mode into the API in API Gateway. Deploy the new version of the API to the production stage.
D. Create a new API Gateway endpoint with new versions of the API definitions. Create a custom domain name for the new API Gateway API. Point the Route 53 alias record to the new API Gateway API custom domain name.

Correct Answer: A
Section:
Explanation:
This answer is correct because it meets the requirements of releasing the new version of APIs with minimal effects on customers and minimal data loss. A canary release deployment is a software development strategy in which a new version of an API is deployed for testing purposes, and the base version remains deployed as a production release for normal operations on the same stage. In a canary release deployment, total API traffic is separated at random into a production release and a canary release with a pre-configured ratio. Typically, the canary release receives a small percentage of API traffic and the production release takes up the rest. The updated API features are only visible to API traffic through the canary. You can adjust the canary traffic percentage to optimize test coverage or performance. By keeping canary traffic small and the selection random, most users are not adversely affected at any time by potential bugs in the new version, and no single user is adversely affected all the time. After the test metrics pass your requirements, you can promote the canary release to the production release and disable the canary from the deployment. This makes the new features available in the production stage.
https://docs.aws.amazon.com/apigateway/latest/developerguide/canary-release.html

QUESTION 207

A company has a web application for travel ticketing. The application is based on a database that runs in a single data center in North America. The company wants to expand the application to serve a global user base. The company needs to deploy the application to multiple AWS Regions. Average latency must be less than 1 second on updates to the reservation database.

The company wants to have separate deployments of its web platform across multiple Regions. However the company must maintain a single primary reservation database that is globally consistent.

Which solution should a solutions architect recommend to meet these requirements?

A. Convert the application to use Amazon DynamoDB. Use a global table for the center reservation table. Use the correct Regional endpoint in each Regional deployment.
B. Migrate the database to an Amazon Aurora MySQL database. Deploy Aurora Read Replicas in each Region. Use the correct Regional endpoint in each Regional deployment for access to the database.
C. Migrate the database to an Amazon RDS for MySQL database Deploy MySQL read replicas in each Region. Use the correct Regional endpoint in each Regional deployment for access to the database.
D. Migrate the application to an Amazon Aurora Serverless database. Deploy instances of the database to each Region. Use the correct Regional endpoint in each Regional deployment to access the database. Use AWS Lambda functions to process event streams in each Region to synchronize the databases.

Correct Answer: B
Section:
Explanation:
https://aws.amazon.com/rds/aurora/global-database/
https://aws.amazon.com/blogs/architecture/using-amazon-aurora-global-database-for-low-latency-without-application-changes/

QUESTION 208

An application uses an Amazon RDS MySQL DB instance. The RDS database is becoming low on disk space. A solutions architect wants to increase the disk space without downtime.

Which solution meets these requirements with the LEAST amount of effort?

A. Enable storage autoscaling in RDS.
B. Increase the RDS database instance size.
C. Change the RDS database instance storage type to Provisioned IOPS.
D. Back up the RDS database, increase the storage capacity, restore the database, and stop the previous instance

Correct Answer: A
Section:
Explanation:
https://aws.amazon.com/about-aws/whats-new/2019/06/rds-storage-auto-scaling/

QUESTION 209

A company wants to use high-performance computing and artificial intelligence to improve its fraud prevention and detection technology. The company requires distributed processing to complete a single workload as quickly as possible.

Which solution will meet these requirements?

A. Use Amazon Elastic Kubernetes Service (Amazon EKS) and multiple containers.
B. Use AWS ParallelCluster and the Message Passing Interface (MPI) libraries.
C. Use an Application Load Balancer and Amazon EC2 instances.
D. Use AWS Lambda functions.

Correct Answer: B
Section:
Explanation:

AWS ParallelCluster is a service that allows you to create and manage high-performance computing (HPC) clusters on AWS.It supports multiple schedulers, including AWS Batch, which can run distributed workloads across multiple EC2 instances1.

MPI is a standard for message passing between processes in parallel computing.It provides functions for sending and receiving data, synchronizing processes, and managing communication groups2.

By using AWS ParallelCluster and MPI libraries, you can take advantage of the following benefits:
You can easily create and configure HPC clusters that meet your specific requirements, such as instance type, number of nodes, network configuration, and storage options1.
You can leverage the scalability and elasticity of AWS to run large-scale parallel workloads without worrying about provisioning or managing servers1.
You can use MPI libraries to optimize the performance and efficiency of your parallel applications by enabling inter-process communication and data exchange2.
You can choose from a variety of MPI implementations that are compatible with AWS ParallelCluster, such as Open MPI, Intel MPI, and MPICH3.

QUESTION 210
A company needs to connect several VPCs in the us-east-1 Region that span hundreds of AWS accounts. The company's networking team has its own AWS account to manage the cloud network.
What is the MOST operationally efficient solution to connect the VPCs?

A. Set up VPC peering connections between each VPC. Update each associated subnet's route table.
B. Configure a NAT gateway and an internet gateway in each VPC to connect each VPC through the internet.
C. Create an AWS Transit Gateway in the networking team's AWS account. Configure static routes from each VPC.
D. Deploy VPN gateways in each VPC. Create a transit VPC in the networking team's AWS account to connect to each VPC.

Correct Answer: C
Section:
Explanation:
AWS Transit Gateway is a highly scalable and centralized hub for connecting multiple VPCs, on-premises networks, and remote networks. It simplifies network connectivity by providing a single entry point and reducing the number of connections required. In this scenario, deploying an AWS Transit Gateway in the networking team's AWS account allows for efficient management and control over the network connectivity across multiple VPCs.

QUESTION 211
A company's data platform uses an Amazon Aurora MySQL database. The database has multiple read replicas and multiple DB instances across different Availability Zones. Users have recently reported errors from the database that indicate that there are too many connections. The company wants to reduce the failover time by 20% when a read replica is promoted to primary writer.
Which solution will meet this requirement?

A. Switch from Aurora to Amazon RDS with Multi-AZ cluster deployment.
B. Use Amazon RDS Proxy in front of the Aurora database.
C. Switch to Amazon DynamoDB with DynamoDB Accelerator (DAX) for read connections.
D. Switch to Amazon Redshift with relocation capability.

Correct Answer: B
Section:
Explanation:
Amazon RDS Proxy is a service that provides a fully managed, highly available database proxy for Amazon RDS and Aurora databases. It allows you to pool and share database connections, reduce database load, and improve application scalability and availability.
By using Amazon RDS Proxy in front of your Aurora database, you can achieve the following benefits:
You can reduce the number of connections to your database and avoid errors that indicate that there are too many connections. Amazon RDS Proxy handles the connection management and multiplexing for you, so you can use fewer database connections and resources.
You can reduce the failover time by 20% when a read replica is promoted to primary writer. Amazon RDS Proxy automatically detects failures and routes traffic to the new primary instance without requiring changes to your application code or configuration. According to a benchmark test, using Amazon RDS Proxy reduced the failover time from 66 seconds to 53 seconds, which is a 20% improvement.
You can improve the security and compliance of your database access. Amazon RDS Proxy integrates with AWS Secrets Manager and AWS Identity and Access Management (IAM) to enable secure and granular authentication and authorization for your database connections.

QUESTION 212
A global marketing company has applications that run in the ap-southeast-2 Region and the eu-west-1 Region. Applications that run in a VPC in eu-west-1 need to communicate securely with databases that run in a VPC in ap-southeast-2.
Which network design will meet these requirements?

A. Create a VPC peering connection between the eu-west-1 VPC and the ap-southeast-2 VPC. Create an inbound rule in the eu-west-1 application security group that allows traffic from the database server IP addresses in the ap-southeast-2 security group.
B. Configure a VPC peering connection between the ap-southeast-2 VPC and the eu-west-1 VPC. Update the subnet route tables. Create an inbound rule in the ap-southeast-2 database security group that references the security group ID of the application servers in eu-west-1.
C. Configure a VPC peering connection between the ap-southeast-2 VPC and the eu-west-1 VPC. Update the subnet route tables Create an inbound rule in the ap-southeast-2 database security group that allows traffic from the eu-west-1 application server IP addresses.
D. Create a transit gateway with a peering attachment between the eu-west-1 VPC and the ap-southeast-2 VPC. After the transit gateways are properly peered and routing is configured, create an inbound rule in the database security group that references the security group ID of the application servers in eu-west-1.

Correct Answer: C
Section:
Explanation:
'You cannot reference the security group of a peer VPC that's in a different Region. Instead, use the CIDR block of the peer VPC.' https://docs.aws.amazon.com/vpc/latest/peering/vpc-peering-security-groups.html

QUESTION 213
A solution architect must create a disaster recovery (DR) plan for a high-volume software as a service (SaaS) platform. All data for the platform is stored in an Amazon Aurora MySQL DB cluster.
The DR plan must replcate data to a secondary AWS Region.
Which solution will meet these requirements MOST cost-effectively?
Use MySQL binary log replication to an Aurora cluster

A. Use MySQL binary log replication to an Aurora cluster in the secondary Region Provision one DB instance for the Aurora cluster in the secondary Region.
B. Set up an Aurora global database for the DB cluster. When setup is complete, remove the DB instance from the secondary Region.
C. Use AWS Database Migration Service (AWS QMS) to continuously replicate data to an Aurora cluster in the secondary Region Remove theDB instance from the secondary Region.
D. Set up an Aurora global database for the DB cluster Specify a minimum of one DB instance in the secondary Region

Correct Answer: D
Section:

QUESTION 214
A company has a multi-tier application deployed on several Amazon EC2 instances in an Auto Scaling group. An Amazon RDS for Oracle instance is the application's data layer that uses Oracle-specific PL/SQL functions. Traffic to the application has been steadily increasing. This is causing the EC2 instances to become overloaded and the RDS instance to run out of storage. The Auto Scaling group does not have any scaling metrics and defines the minimum healthy instance count only. The company predicts that traffic will continue to increase at a steady but unpredictable rate before levelling off. What should a solutions architect do to ensure the system can automatically scale for the increased traffic? (Select TWO.)

A. Configure storage Auto Scaling on the RDS for Oracle Instance.
B. Migrate the database to Amazon Aurora to use Auto Scaling storage.
C. Configure an alarm on the RDS for Oracle Instance for low free storage space
D. Configure the Auto Scaling group to use the average CPU as the scaling metric
E. Configure the Auto Scaling group to use the average free memory as the seeing metric

Correct Answer: A, D
Section:
Explanation:
Auto scaling storage RDS will ease storage issues and migrating Oracle Pl/Sql to Aurora is cumbersome. Also Aurora has auto storage scaling by default.
https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_PIOPS.StorageTypes.html#USER_PI OPS.Autoscaling

QUESTION 215
A company has an AWS Lambda function that needs read access to an Amazon S3 bucket that is located in the same AWS account. Which solution will meet these requirement in the MOST secure manner?

A. Apply an S3 bucket pokey that grants road access to the S3 bucket
B. Apply an IAM role to the Lambda function Apply an IAM policy to the role to grant read access to the S3 bucket
C. Embed an access key and a secret key In the Lambda function's coda to grant the required IAM permissions for read access to the S3 bucket
D. Apply an IAM role to the Lambda function. Apply an IAM policy to the role to grant read access to all S3 buckets In the account

Correct Answer: B
Section:

QUESTION 216
A company wants to implement a disaster recovery plan for its primary on-premises file storage volume. The file storage volume is mounted from an Internet Small Computer Systems Interface (iSCSI) device on a local storage server. The file storage volume holds hundreds of terabytes (TB) of data.
The company wants to ensure that end users retain immediate access to all file types from the onpremises systems without experiencing latency. Which solution will meet these requirements with the LEAST amount of change to the company's existing infrastructure?

A. Provision an Amazon S3 File Gateway as a virtual machine (VM) that is hosted on premises. Set the local cache to 10 TB. Modify existing applications to access the files through the NFS protocol. To recover from a disaster, provision an Amazon EC2 instance and mount the S3 bucket that contains the files.
B. Provision an AWS Storage Gateway tape gateway. Use a data backup solution to back up all existing data to a virtual tape library. Configure the data backup solution to run nightly after the initial backup is complete. To recover from a disaster, provision an Amazon EC2 instance and restore the data to an Amazon Elastic Block Store (Amazon EBS) volume from the volumes in the virtual tape library.
C. Provision an AWS Storage Gateway Volume Gateway cached volume. Set the local cache to 10 TB.
Mount the Volume Gateway cached volume to the existing file server by using iSCSI. and copy all files to the storage volume. Configure scheduled snapshots of the storage volume. To recover from a disaster, restore a snapshot to an Amazon Elastic Block Store (Amazon EBS) volume and attach the EBS volume to an Amazon EC2 instance.
D. Provision an AWS Storage Gateway Volume Gateway stored volume with the same amount of disk space as the existing file storage volume. Mount the Volume Gateway stored volume to the existing file server by using iSCSI, and copy all files to the storage volume. Configure scheduled snapshots of the storage volume. To recover from a disaster, restore a snapshot to an Amazon Elastic Block Store (Amazon EBS) volume and attach the EBS volume to an Amazon EC2 instance.

Correct Answer: D
Section:
Explanation:

QUESTION 217
A company is launching a new application deployed on an Amazon Elastic Container Service (Amazon ECS) cluster and is using the Fargate launch type tor ECS tasks The company is monitoring CPU and memory usage because it is expecting high traffic to the application upon its launch However the company wants to reduce costs when utilization decreases What should a solutions architect recommend?

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

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

QUESTION 218
A data analytics company wants to migrate its batch processing system to AWS. The company receives thousands of small data files periodically during the day through FTP. A on-premises batch job processes the data files overnight. However, the batch job takes hours to finish running.

The company wants the AWS solution to process incoming data files are possible with minimal changes to the FTP clients that send the files. The solution must delete the incoming data files the files have been processed successfully. Processing for each file needs to take 3-8 minutes.
Which solution will meet these requirements in the MOST operationally efficient way?

A. Use an Amazon EC2 instance that runs an FTP server to store incoming files as objects in Amazon S3 Glacier Flexible Retrieval. Configure a job queue in AWS Batch. Use Amazon EventBridge rules to invoke the job to process the objects nightly from S3 Glacier Flexible Retrieval. Delete the objects after the job has processed the objects.
B. Use an Amazon EC2 instance that runs an FTP server to store incoming files on an Amazon Elastic Block Store (Amazon EBS) volume. Configure a job queue in AWS Batch. Use Amazon EventBridge rules to invoke the process the files nightly from the EBS volume. Delete the files after the job has processed the files.
C. Use AWS Transfer Family to create an FTP server to store incoming files on an Amazon Elastic Block Store (Amazon EBS) volume. Configure a job queue in AWS Batch. Use an Amazon S3 event notification when each files arrives to invoke the job in AWS Batch. Delete the files after the job has processed the files.
D. Use AWS Transfer Family to create an FTP server to store incoming files in Amazon S3 Standard.
Create an AWS Lambda function to process the files and to delete the files after they are proessed.yse an S3 event notification to invoke the lambda function when the fils arrive

Correct Answer: D
Section:
Explanation:
This option is the most operationally efficient because it uses AWS Transfer Family to create an FTP server that can store incoming files in Amazon S3 Standard12, which is a low-cost and highly available storage service. It also uses AWS Lambda to process the files and delete them after they are processed, which is a serverless and scalable solution that does not require any batch scheduling or infrastructure management. It also uses S3 event notifications to invoke the Lambda function when the files arrive, which enables near real-time processing of the incoming data files3. Option A is less efficient because it uses Amazon S3 Glacier Flexible Retrieval, which is a cold storage class that has higher retrieval costs and longer retrieval times than Amazon S3 Standard. It also uses EventBridge rules to invoke the job nightly, which does not meet the requirement of processing incoming data files as soon as possible. Option B is less efficient because it uses an EBS volume to store incoming files, which is a block storage service that has higher costs and lower durability than Amazon S3. It also uses EventBridge rules to invoke the job nightly, which does not meet the requirement of processing incoming data files as soon as possible. Option C is less efficient because it uses an EBS volume to store incoming files, which is a block storage service that has higher costs and lower durability than Amazon S3. It also uses AWS Batch to process the files, which requires managing compute resources and job queues.

QUESTION 219
A company hosts a three-tier web application that includes a PostgreSQL database The database stores the metadata from documents The company searches the metadata for key terms to retrieve documents that the company reviews in a report each month The documents are stored in Amazon S3 The documents are usually written only once, but they are updated frequency The reporting process takes a few hours with the use of relational queries The reporting process must not affect any document modifications or the addition of new documents.
What are the MOST operationally efficient solutions that meet these requirements? (Select TWO)

A. Set up a new Amazon DocumentDB (with MongoDB compatibility) cluster that includes a read replica Scale the read replica to generate the reports.
B. Set up a new Amazon RDS for PostgreSQL Reserved Instance and an On-Demand read replica Scale the read replica to generate the reports
C. Set up a new Amazon Aurora PostgreSQL DB cluster that includes a Reserved Instance and an Aurora Replica issue queries to the Aurora Replica to generate the reports.
D. Set up a new Amazon RDS for PostgreSQL Multi-AZ Reserved Instance Configure the reporting module to query the secondary RDS node so that the reporting module does not affect the primary node
E. Set up a new Amazon DynamoDB table to store the documents Use a fixed write capacity to support new document entries Automatically scale the read capacity to support the reports

Correct Answer: B, C
Section:

QUESTION 220
A company recently created a disaster recovery site in a Different AWS Region.The company needs to transfer large amounts of data back and forth between NFS file systems in the two Regions on a periods. Which solution will meet these requirements with the LEAST operational overhead?

A. Use AWS DataSync.
B. Use AWS Snowball devices
C. Set up an SFTP server on Amazon EC2
D. Use AWS Database Migration Service (AWS DMS)

Correct Answer: A
Section:

QUESTION 221
A company needs to store contract documents. A contract lasts for 5 years. During the 5-year period, the company must ensure that the documents cannot be overwritten or deleted. The company needs to encrypt the documents at rest and rotate the encryption keys automatically every year.
Which combination of steps should a solutions architect take to meet these requirements with the LEAST operational overhead? (Select TWO.)

A. Store the documents in Amazon S3. Use S3 Object Lock in governance mode.
B. Store the documents in Amazon S3. Use S3 Object Lock in compliance mode.
C. Use server-side encryption with Amazon S3 managed encryption keys (SSE-S3). Configure key rotation.
D. Use server-side encryption with AWS Key Management Service (AWS KMS) customer managed keys. Configure key rotation.
E. Use server-side encryption with AWS Key Management Service (AWS KMS) customer provided (imported) keys. Configure key rotation.

Correct Answer: B, D
Section:
Explanation:
Consider using the default aws/s3 KMS key if: You're uploading or accessing S3 objects using AWS Identity and Access Management (IAM) principals that are in the same AWS account as the AWS KMS key. You don't want to manage policies for the KMS key. Consider using a customer managed key if: You want to create, rotate, disable, or define access controls for the key. You want to grant cross-account access to your S3 objects. You can configure the policy of a customer managed key to allow access from another account. https://repost.aws/knowledge-center/s3-object-encryption-keys

QUESTION 222
A company runs analytics software on Amazon EC2 instances The software accepts job requests from users to process data that has been uploaded to Amazon S3 Users report that some submitted data is not being processed Amazon CloudWatch reveals that the EC2 instances have a consistent CPU utilization at or near 100% The company wants to improve system performance and scale the system based on user load.
What should a solutions architect do to meet these requirements?

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

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

1Using Amazon SQS queues with Amazon EC2 Auto Scaling - Amazon EC2 Auto Scaling
2Tutorial: Set up a scaled and load-balanced application - Amazon EC2 Auto Scaling
3Amazon EC2 Auto Scaling FAQs

QUESTION 223

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

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

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

QUESTION 224

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 225

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

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

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