QUESTION 241
A company recently implemented hybrid cloud connectivity using AWS Direct Connect and is migrating data to Amazon S3.
The company is looking for a fully managed solution that will automate and accelerate the replication of data between the on-premises storage systems and AWS storage services.
Which solution should a solutions architect recommend to keep the data private?

A. Deploy an AWS DataSync agent tor the on-premises environment.
Configure a sync job to replicate the data and connect it with an AWS service endpoint.
B. Deploy an AWS DataSync agent for the on-premises environment.
Schedule a batch job to replicate point-ln-time snapshots to AWS.
C. Deploy an AWS Storage Gateway volume gateway for the on-premises environment.
Configure it to store data locally, and asynchronously back up point-in-time snapshots to AWS.
D. Deploy an AWS Storage Gateway file gateway for the on-premises environment.
Configure it to store data locally, and asynchronously back up point-in-lime snapshots to AWS.

Answer: C

QUESTION 242
A company has an on-premises data center that is running out of storage capacity.
The company wants to migrate its storage infrastructure to AWS while minimizing bandwidth costs.
The solution must allow for immediate retrieval of data at no additional cost.
How can these requirements be met?

A. Deploy Amazon S3 Glacier Vault and enable expedited retrieval.
Enable provisioned retrieval capacity for the workload
B. Deploy AWS Storage Gateway using cached volumes.
Use Storage Gateway to store data in Amazon S3 while retaining copies of frequently accessed data subsets locally.
C. Deploy AWS Storage Gateway using stored volumes to store data locally.
Use Storage Gateway to asynchronously back up point-in-time snapshots of the data to Amazon S3
D. Deploy AWS Direct Connect to connect with the on-premises data center.
Configure AWS Storage Gateway to store data locally.
Use Storage Gateway to asynchronously backK up potnt-tn-time snapshots of the data to Amazon S3.

Answer: B

QUESTION 243

A company Is reviewing Its AWS Cloud deployment to ensure its data is not accessed by anyone without appropriate authorization.
A solutions architect is tasked with identifying all open Amazon S3 buckets and recording any S3 bucket configuration changes.
What should the solutions architect do to accomplish this?

A. Enable AWS Config service with the appropriate rules
B. Enable AWS Trusted Advisor with the appropriate checks.
C. Write a script using an AWS SDK to generate a bucket report
D. Enable Amazon S3 server access logging and configure Amazon CloudWatch Events.

Answer: A

QUESTION 244
A company built an application that lets users check in to places they visit, rank the places, and add reviews about their experiences.
The application is successful with a rapid increase in the number of users every month.
The chief technology officer fears the database supporting the current Infrastructure may not handle the new load the following month because the single Amazon RDS for MySQL instance has triggered alarms related lo resource exhaustion due to read requests.
What can a solutions architect recommend to prevent service Interruptions at the database layer with minimal changes to code?

A. Create RDS read replicas and redirect read-only traffic to the read replica endpoints.
Enable a Multi-AZ deployment.
B. Create an Amazon EMR cluster and migrate the data to a Hadoop Distributed File System (HDFS) with a replication factor of 3.
C. Create an Amazon ElastiCache cluster and redirect all read-only traffic to the cluster.
Set up the cluster to be deployed m three Availability Zones.
D. Create an Amazon DynamoDB table to replace the RDS instance and redirect all read-only traffic to the DynamoDB table.
Enable DynamoDB Accelerator to offload traffic from the main table.

Answer: A

QUESTION 245
A company runs an application on Amazon EC2 Instances.
The application is deployed in private subnets in three Availability Zones of the us-east-1 Region.
The instances must be able to connect to the internet to download files.
The company wants a design that Is highly available across the Region.
Which solution should be implemented to ensure that there are no disruptions to Internet connectivity?

A. Deploy a NAT Instance In a private subnet of each Availability Zone.
B. Deploy a NAT gateway in a public subnet of each Availability Zone.
C. Deploy a transit gateway in a private subnet of each Availability Zone.
D. Deploy an internet gateway in a public subnet of each Availability Zone.

Answer: B

QUESTION 246
A company has migrated an on-premises Oracle database to an Amazon RDS (or Oracle Multi-AZ DB instance In the us-east-l Region.
A solutions architect is designing a disaster recovery strategy to have the database provisioned In the us-west-2 Region In case the database becomes unavailable in the us-east-1 Region.
The design must ensure the database is provisioned in the us-west-2 Region in a maximum of 2 hours, with a data loss window of no more than 3 hours.

How can these requirements be met?

A. Edit the DB instance and create a read replica in us-west-2.
Promote the read replica to master In us- west-2 in case the disaster recovery environment needs to be activated.
B. Select the multi-Region option to provision a standby instance in us-west-2.
The standby Instance will be automatically promoted to master In us-west-2 in case the disaster recovery environment needs to be created.
C. Take automated snapshots of the database instance and copy them to us-west-2 every 3 hours.
Restore the latest snapshot to provision another database instance in us-west-2 in case the disaster recovery environment needs to be activated.
D. Create a multimaster read/write instances across multiple AWS Regions Select VPCs in us-east-1 and us-west-2 lo make that deployment.
Keep the master read/write instance in us-west-2 available to avoid having to activate a disaster recovery environment.

Answer: A

QUESTION 247
A company has an application with a REST-based Interface that allows data to be received in near-real time from a third-party vendor.
Once received, the application processes and stores the data for further analysis.
The application Is running on Amazon EC2 instances.
The third-party vendor has received many 503 Service Unavailable Errors when sending data to the application.
When the data volume spikes, the compute capacity reaches its maximum limit and the application is unable to process all requests.
Which design should a solutions architect recommend to provide a more scalable solution?

A. Use Amazon Kinesis Data Streams to ingest the data.
Process the data using AWS Lambda functions.
B. Use Amazon API Gateway on top of the existing application.
Create a usage plan with a quota limit for the third-party vendor.
C. Use Amazon Simple Notification Service (Amazon SNS) to ingest the data.
Put the EC2 instances in an Auto Scaling group behind an Application Load Balancer.
D. Repackage the application as a container.
Deploy the application using Amazon Elastic Container Service (Amazon ECS) using the EC2 launch type with an Auto Scaling group.

Answer: A

QUESTION 248
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.

Answer: A

QUESTION 249
A company recently deployed a two-tier application in two Availability Zones in the us-east-1 Region.
The databases are deployed in a private subnet while the web servers are deployed in a public subnet.
An internet gateway is attached to the VPC. The application and database run on Amazon EC2 instances.

The database servers are unable to access patches on the internet.
A solutions architect needs to design a solution that maintains database security with the least operational overhead.
Which solution meets these requirements?

A. Deploy a NAT gateway inside the public subnet for each Availability Zone and associate it with an Elastic IP address.
Update the routing table of the private subnet to use it as the default route.
B. Deploy a NAT gateway inside the private subnet for each Availability Zone and associate it with an Elastic IP address.
Update the routing table of the private subnet to use it as the default route.
C. Deploy two NAT instances inside the public subnet for each Availability Zone and associate them with Elastic IP addresses.
Update the routing table of the private subnet to use it as the default route.
D. Deploy two NAT instances inside the private subnet for each Availability Zone and associate them with Elastic IP addresses.
Update the routing table of the private subnet to use it as the default route.

Answer: B

QUESTION 250
A solutions architect must design a solution for a persistent database that is being migrated from on- premises to AWS.
The database requires 64,000 IOPS according to the database administrator.
If possible, the database administrator wants to use a single Amazon Elastic Block Store (Amazon EBS) volume to host the database instance.
Which solution effectively meets the database administrator's criteria?

A. Use an instance from the 13 I/O optimized family and leverage local ephemeral storage to achieve the IOPS requirement.
B. Create an Nitro-based Amazon EC2 instance with an Amazon EBS Provisioned IOPS SSD (io1) volume attached. Configure the volume to have 64,000 IOPS.
C. Create and map an Amazon Elastic File System (Amazon EFS) volume to the database instance and use the volume to achieve the required IOPS for the database.
D. Provision two volumes and assign 32,000 IOPS to each. Create a logical volume at the operating system level that aggregates both volumes to achieve the IOPS requirements.

Answer: B

QUESTION 251
A company recently launched its website to serve content to its global user base.
The company wants to store and accelerate the delivery of static content to its users by leveraging Amazon CloudFront with an Amazon EC2 instance attached as its origin.
How should a solutions architect optimize high availability for the application?

A. Use Lambda@Edge for CloudFront.
B. Use Amazon S3 Transfer Acceleration for CloudFront.
C. Configure another EC2 instance in a different Availability Zone as part of the origin group.
D. Configure another EC2 instance as part of the origin server cluster in the same Availability Zone.

Answer: A