QUESTION 750
A disaster relief company is designing a new solution to analyze real-time csv data. The data is collected by a network of thousands of research stations met are distributed across the world. The data volume is consistent and constant, and the size of each data We is 512 KB. The company needs to stream the data and analyze the data in real time.
Which combination of actions should a solutions architect take to meet these requirements? (Select TWO.)

A. Provision an appropriately sized Amazon Simple Queue Service (Amazon SOS) queue.
    Use the AWS SDK at the research stations to write the data into the SOS queue
B. Provision an appropriately sized Amazon Kinesis Data Firehose delivery stream.
    Use the AWS SDK at the research stations to write the data into the delivery stream and then into an Amazon S3 bucket.
C. Provision an appropriately sized Amazon Kinesis Data Analytics application.
    Use the AWS CLI to configure Kinesis Data Analytics with SOL queries
D. Provision an AWS Lambda function to process the data.
    Set up the BatchSize property on the Lambda event source.
E. Provision an AWS Lambda function to process the data.
    Set up an Amazon EventBridge (Amazon CloudWatch Events) cron expression rule to invoke the Lambda function

Answer: AD

QUESTION 751
A company hosts more than 300 global websites and applications.
The company requires a platform to analyze more than 30 TB of clickstream data each day.
What should a solutions architect do to transmit and process the clickstream data?

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

Answer: D

QUESTION 752
A company hosts a popular website in the AWS Cloud.

A solutions architect needs to provide reports about user click behaviour in near-real time as users navigate the website.
Which solution will meet this requirement?

A. Store the clickstream data in Amazon DynamoDB.
Deploy an application that runs on AWS Elastic Beanstalk to process and analyze the data.
B. Push the clickstream data from each session to an Amazon Kinesis data stream
Analyze the dab by using Amazon Kinesis Data Analytics.
C. Store the clickstream data in an Amazon S3 bucket.
Order the data by timestamp
Process the data with an AWS Lambda function that is subscribed to object creation events on the S3 bucket.
D. Forward the clickstream data to Amazon Simple Queue Service (Amazon SOS)
Store the data In an Amazon ROS for MySQL DB instance.
Deploy Amazon FC2 Instances to process and analyze the data

Answer: B

QUESTION 753
A company plant to host a survey website on AWS. The company anticipates an unpredictable amount of traffic. This traffic results m asynchronous updates to the database. The company wants to ensure mat writes to the database hosted on AWS do not gel dropped.
How should the company write its application to hand to these database requests?

A. Configure the application to publish to an Amazon Simple Notification Service (Amazon SNS) topic
Subscribe the database to the SNS topic.
B. Configure the application to subscribe to an Amazon Simple Notification Service (Amazon SNS) topic.
Publish the database updates to the SNS topic
C. Use Amazon Simple Queue Service (Amazon SOS) FIFO queues to queue the database connection until the database has resources to wrist the data.
D. Use Amazon Simple Queue Service (Amazon SOS) FIFO queues tor capturing the writes and draining the queue as each write is made to the database.

Answer: D

QUESTION 754
A company stores can wordings on a monthly basis. Users access lie recorded files randomly within 1 year of recording, but users rarely access the files after 1year. The company wants to optimize its solution by allowing only files that ant newer than 1 year old to be queried and retrieved as quickly as possible. A delay in retrieving older fees is acceptable. Which solution meets these requirements MOST cost-effectively?

A. Store individual files in Amazon S3 Glacier
Store search metadata in object tags that are created in S3 Glacier
Query the S3 Glacier tags to retrieve the files from S3 Glacier.
B. Store individual files in Amazon S3.
Use S3 Lifecycle polices to move the ties to S3 Glacier after 1 year.
Query and retrieve the files that are in Amazon S3 by using Amazon Athena.
Query and retrieve the files that are in S3 Glacier by using S3 Glacier Select.
C. Store Individual files In Amazon S3
Store search metadata for each archive In Amazon S3
Use S3 Lifecycle policies to move the ties to S3 Glacier after 1 year
Query and retrieve tie flies by searching for metadata from Amazon S3.
D. Store individual files in Amazon S3 Use S3
Lifecycle policies to move the files to S3 Glacier after 1 year.

Store search metadata in Amazon RDS
Query the Sea from Amazon RDS Retrieve the files from Amazon S3 or S3 Glacier

Answer: D

QUESTION 755
A company uses a combination of Amazon EC2 instances and AWS Fargate tasks to process daily transactions. The company faces unpredictable and sudden increases in transaction volume. The company needs a solution that will process the transactions immediately.
Which solution meets these requirement MOST cost-effectively?

A. Purchase a Compute Savings Plan
B. Purchase an EC2 Instance Savings Plan.
C. Purchase Reserved Instances tor existing EC2 workloads.
D. Use Spot Instances for existing EC2 workloads.
E. Use Far gale Spot capacity for the tasks.

Answer: B

QUESTION 756
A meteorological startup company has a custom web application to sell weather data to its users online. The company uses Amazon DynamoDB to store is data and wants to build a new service that sends an alert to the managers of four Internal teams every time a new weather event is recorded. The company does not want true new service to affect the performance of the current application.
What should a solutions architect do to meet these requirement with the LEAST amount of operational overhead?

A. Use DynamoDB transactions to write new event data to the table
Configure the transactions to notify internal teams.
B. Have the current application publish a message to four Amazon Simple Notification Service (Amazon SNS) topics.
Have each team subscribe to one topic.
C. Enable Amazon DynamoDB Streams on the table.
Use triggers to write to a mingle Amazon Simple Notification Service (Amazon SNS) topic to which the teams can subscribe.
D. Add a custom attribute to each record to flag new items.
Write a cron job that scans the table every minute for items that are new and notifies an Amazon Simple Queue Service (Amazon SOS) queue to which the teams can subscribe.

Answer: C

QUESTION 757
A company has a popular gaming platform running on AWS. The application is sensitive to latency because latency can impact the user experience and Introduce unfair advantages to some prayers. The application la deployed In a very AWS Region. It runs on Amazon FC2 Instances Vial are part of Auto Scaling groups configured behind Application Load Balancers (ALBs).
A solutions architect needs to implement a mechanism to monitor the hearth of the application and redirect traffic to healthy endpoints.
Which solution meets these requirements?

A. Configure an accelerator In AWS Global Accelerator
Add a listens for the port that the application listens on and attach it to a Regional endpoint in each Region
Add the ALB as the endpoint
B. Create an Amazon CloudFron4t distribution and specify the ALB as the origin server
Configure the cache behaviour to use origin cache headers
Use AWS Lambda functions to optimize the traffic

C. Create an Amazon CloudFront distribution and specify Amazon S3 as the origin server.
Configure tie cache behaviour to use origin cache headers
Use AWS Lambda functions to optimize the traffic
D. Configure an Amazon DynamoDB database to serve as the data store tor the application
Create a DynamoDB Accelerator (DAX) cluster to act as the in-memory cache for DynamoDB
hosting the application data

Answer: A

QUESTION 758
A company has NFS servers in an on-premises data center that need to periodically back up small amounts of data to Amazon S3.
Which solution marts these requirement and is MOST cost-effective?

A. Set up AWS Glue lo copy the data from the on-premises servers to Amazon S3.
B. Set up an AWS DataSync agent on Vie on-premises servers, and sync the data lo Amazon S3
C. Set up an SFTP sync using AWS Transfer for SFTP lo sync data from on premises lo Amazon S3
D. Set up an AWS Direct Connect connection between the on-premises data center and a VPC, and copy the data to Amazon S3

Answer: B

QUESTION 759
A company is hosting a web application from an Amazon S3 bucket.
The application uses Amazon Cognito as an identity provider lo authenticate users and return a JSON Web Token (JWT) that provides access to protected resources that am restored in another S3 bucket.
Upon deployment of the application, users report errors and are unable to access the protected content.
A solutions architect must resolve this issue by providing proper permissions so that users can access the protected content.
Which solution meets these requirements?

A. Update the Amazon Cognito identity pool to assume the proper IAM role for access to the protected consent.
B. Update the S3 ACL to allow the application to access the protected content
C. Redeploy the application to Amazon S3 to prevent eventually consistent reads m the S3 bucket from affecting the ability of users to access the protected content.
D. Update the Amazon Cognito pool to use custom attribute mappings within tie Identity pool and grant users the proper permissions to access the protected content

Answer: B

QUESTION 760
A company has several Amazon EC2 instances set up m a private subnet for security reasons. These instances host applications that read and write large amounts of data to end from Amazon S3 regularly. Currently subnet routing directs all the traffic destined for the internet through a NAT gateway. The company wants to optimize the overall coat without impacting the ability of the application to communication Amazon S3 or the outside internet. What should a solutions architect do to optimize costs?

A. Create an additional NAT gateway.
Update the route table to route to the NAT gateway
Update the network ACL lo allow S3 traffic
B. Create an internet gateway
Update the route table to route traffic to the internet gateway
Update the network ACL to allow S3 traffic
C. Create a VPC endpoint for Amazon S3 Attach an endpoint policy to the endpoint
Update the route table lo direct traffic to the VPC endpoint.

D. Create an AWS Lambda function outside of the VPC to handle S3 requests
Attach an IAM policy to the EC2 instances, allowing them to invoke the Lambda function.

Answer: C

QUESTION 761
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 dies 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 heap every lie in the repository for a minimum of 1 year after its creation date.
Which solution will meet these requirements?

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

Answer: B

QUESTION 762
A company is rebelling its data canter and wants to securely transfer 50 TB of data to AWS within 2 weeks.
The existing data center has a Site-to-Site VPN connection to AWS that is 90 % utilized.
Which AWS service should a solutions architect use to meet these requirements?

A. AWS DataSync with a VPC endpoint
B. AWS Direct Conned
C. AWS Snowball Edge Storage Optimized
D. AWS Storage Gateway

Answer: C

QUESTION 763
A company runs an internet-facing web application on AWS. The company uses Amazon Route 53 for DNS management and has a public hosted zone lo route traffic from the internet to the application. The company wants to tog DNS response codes to help system administrators perform any root cause analysis in the future.
Which solution will meet these requirements?

A. Use Route 53 to configure query togging
B. Use AWS CloudTrail lo record ail Route 53 queries
C. Use Amazon CloudWatch to record and process Route 53 metrics
D. Use AWS Trusted Advisor to perform on-demand root cause analysis

Answer: A

QUESTION 764
A company used an AWS Direct Connect connection to copy 1 TB of data from a colocation facility to an Amazon S3 bucket in the us-east-1 Region. The company now wants to copy the data to another S3 bucket in the us-weet-2 Region.
Which solution will meet this requirement?

A. Use an AWS Snowball Edge Storage Optimized device to copy the data from the colocation facility to ua-weet-2
B. Use the S3 console to copy the data horn the source S3 bucket to the target S3 bucket.

C. Use S3 Transfer Acceleration and the S3 copy-object command to copy the data from the source S3 bucket to the target S3 bucket
D. Add an S3 Cross-Region Replication configuration to copy the data from the source S3 bucket to the target S3 bucket.

Answer: D

QUESTION 765
A solution architect is designing he architect of a new application being deployed to the AWS Cloud. The application will run on Amazon EC2 On-Demand instances and will automatically scale across multiple Availability Zones. The EC2 instances will scale up and down frequently the day. An Application load balancer (ALB) will handle the load distribution. The architecture needs to support distributed session data management. The company is willing to make charges to code if needed.
What should the solutions architect do to ensure that the architecture supports distributed session data management?

A. Use Amazon ElastiCache to manage and store session data.
B. Use session affinity (sticky sessions) of the ALB to manage session data.
C. Use Session Manager from AWS Systems Manager to manage the session.
D. Use the GetSessionToken API operation in AWS Security Token Service (AWS STS) to manage the session.

Answer: C

QUESTION 766
A company uses an application to present metrics from sporting events to the public.
The application must scale quickly during live events and must store these metrics for log-term reporting purposes.
The company's architecture includes the following:
- Amazon EC2 instances that run in an Auto Scaling group in private subnets
- A network Load Balancer That runs in public subnets
- A MongoDB database cluster that runs across multiple EC2 instances
A solutions architect mutt implement a solution that minimizes operational overhead.
The solution alto must be able to Kale automatically.
What should the solutions architect set up to meet these requirements?

A. An Amazon DynamoDB database
B. An Amazon ROS for MySQL D6 instance
C. EC2 instances that run MySQL
D. Amazon Redshift

Answer: A

QUESTION 767
A company is running a publicly accessible serverless application that uses Amazon API Gateway and AWS Lambda.
The application's traffic recently spiked due to fraudulent requests from botnets.
Which steps should a solutions architect take to block requests from unauthorized users? (Select TWO.)

A. Create a usage plan with an API key that it shared with genuine users only.
B. Integrate logic within the Lambda function to ignore the requests lion- fraudulent IP addresses
C. Implement an AWS WAF rule to target malicious requests and trigger actions to filler them out
D. Convert the existing public API to a private API Update the DNS records to redirect users to the new API endpoint
E. Create an IAM role tor each user attempting to access the API A user will assume the role when making the API call

Answer: CD

QUESTION 768
An application hosted on AWS is experiencing performance problems, and the application vendor wants to perform an analysis of the log file to troubleshoot further. The log file is stored on Amazon S3 and is 10GB in size.
The application owner will make the log file available to the vendor for a limited time.
Whit it the MOST secure way to do this?

A. Enable public read on the S3 object and provide the link to the vendor
B. Upload the lie to Amazon WorkDocs and share the public link with the vendor.
C. Generate a presigned URL and have the vendor download the log Me before it expiree
D. Create an LAM user for the vendor to provide access to the S3 bucket and tie application.
Enforce multifactor authentication

Answer: C

QUESTION 769
Some of the company's customers are retrieving records frequently, leading to an increase in costs for the company.
The company wants to limit retrieved requests in the future.
The company also wants to ensure that if one customer reaches its retrieval limit other customers will not affected.
Which solution will meet these requirements?

A. Set up server-side throttling limits for API Gateway.
B. Limit DynamoDB read throughput on the table lo an amount that results m the maximum cost that
the company is willing to incur.
C. Set up a usage plan for API Gateway Implement throttling limits tor each customer and distribute
API keys to each customer
D. Set up AWS Budgets.
Monitor the usage of API Gateway and DynamoDB
Configure an alarm to provide an alert when the cost exceeds a certain threshold each month

Answer: D

QUESTION 770
A company collects 10 GB of telemetry data dairy from various machines.
The company stores the data in an Amazon S3 bucket in a source data account.
The company has hired several consuming agencies to use this data for analysis. Each agency needs read access to the data for its analysis. The company must share the data from tie source data account by choosing a solution that maximizes security and operational efficiency.
Which solution will meet these requirements?

A. Configure S3 global tables to replicate data tor each agency
B. Make the S3 bucket public for a limited time Inform only the agencies
C. Configure cross-account access for the S3 bucket to the accounts that the agencies own.
D. Set up an IAM user for each analyst In the source data account Grant each user access to the S3
bucket

Answer: A

QUESTION 771
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

Answer: D

QUESTION 772
A company has a corporate network on premises and has three VPCs in the AWS Cloud.
The company has one VPC each for development, test, and, production.
The company wants its system administrators to security gain remote command-line access from the corporate network to Amazon EC2 instances in the VPCs.
Which solution meets these requirements MOST cost-effectively?

A. Set up a VPN connection between the corporate network and each of the three VPCs by using AWS VPN
   Use Remote Desktop Protocol (RDP) or SSH over the VPN connection to access the EC2 instances remotely.
B. Configure the EC2 instances to use an instance profile that trusts AWS Systems Manager
   Use Systems Manager Session Manager to gain console access to the EC2 instances
C. Create a new VPC Purchase and install a virtual router from AWS Marketplace
   Establish a VPN connection from the corporate network to this router.
   Establish another VPN connection from the 'outer to the other three VPCs
   Use Remote Desktop Protocol (RDP) or SSH over the VPN connection to access the EC2 instances remotely.
D. Create a new VPC Establish a VPN connection to the new VPC.
   Configure peering connections between the new VPC and the existing VPCs
   In the new VPC create an EC2 bastion host to serve as a jump box lo EC2 instances in the other VPCs
   Use Remote Desktop Protocol (RDP) or SSH over the VPN connection to the bastion host

Answer: A

QUESTION 773
A solutions architect It 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 gateways used to provide internet access for the public subnets. The private subnets require access to the internet to allow Amazon BC2 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 tie 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 trade to the private Internet gateway.
D. Create an egress-only internet gateway on one of the pubic subnets.
   Update the route table for the private subnets that forward non-VPC traffic to the egress-only internet gateway.

Answer: A

QUESTION 774
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 (or MySQL with Multi-AZ enabled.

Answer: D