Free Questions for SAP-C02 by go4braindumps

Shared by Espinoza on 09-08-2024

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

A solutions architect is responsible (or redesigning a legacy Java application to improve its availability, data durability, and scalability. Currently, the application runs on a single high-memory Amazon EC2 instance. It accepts HTTP requests from upstream clients, adds them to an in-memory queue, and responds with a 200 status. A separate application thread reads items from the queue, processes them, and persists the results to an Amazon RDS MySQL instance. The processing time for each item takes 90 seconds on average, most of which is spent waiting on external service calls, but the application is written to process multiple items in parallel.

Traffic to this service is unpredictable. During periods of high load, items may sit in the internal queue for over an hour while the application processes the backlog. In addition, the current system has issues with availability and data loss if the single application node fails.

Clients that access this service cannot be modified. They expect to receive a response to each HTTP request they send within 10 seconds before they will time out and retry the request.

Which approach would improve the availability and durability of (he system while decreasing the processing latency and minimizing costs?

Options:

A) Create an Amazon API Gateway REST API that uses Lambda proxy integration to pass requests to an AWS Lambda function. Migrate the core processing code to a Lambda function and write a wrapper class that provides a handler method that converts the

proxy events to the internal application data model and invokes the processing module.

B) Create an Amazon API Gateway REST API that uses a service proxy to put items in an Amazon SOS queue. Extract the core processing code from the existing application and update it to pull items from Amazon SOS instead of an in-memory queue. Deploy the new processing application to smaller EC2 instances within an Auto Scaling group that scales dynamically based on the approximate number of messages in the Amazon SOS queue.

C) Modify the application to use Amazon DynamoDB instead of Amazon RDS. Configure Auto Scaling for the DynamoDB table. Deploy the application within an Auto Scaling group with a scaling policy based on CPU utilization. Back the in-memory queue with a memory-mapped file to an instance store volume and periodically write that file to Amazon S3.

D) Update the application to use a Redis task queue instead of the in-memory queue. 8uild a Docker container image for the application. Create an Amazon ECS task definition that includes the application container and a separate container to host Redis. Deploy the new task definition as an ECS service using AWS Fargate, and enable Auto Scaling.

Answer:

B

Explanation:

The obvious challenges here are long workloads, scalability based on queue load, and reliability. Almost always the defacto answer to queue related workload is SQS. Since the workloads are very long (90 minutes) Lambdas cannot be used (15 mins max timeout). So, autoscaled smaller EC2 nodes that wait on external services to complete the task makes more sense. If the task fails, the message is returned to the queue and retried.

Question 2

Question Type: MultipleChoice

A company has a three-tier application running on AWS with a web server, an application server, and an Amazon RDS MySQL DB instance. A solutions architect is designing a disaster recovery (OR) solution with an RPO of 5 minutes.

Which solution will meet the company's requirements?

Options:

A) Configure AWS Backup to perform cross-Region backups of all servers every 5 minutes. Reprovision the three tiers in the DR Region from the backups using AWS CloudFormation in the event of a disaster.

B) Maintain another running copy of the web and application server stack in the DR Region using AWS CloudFormation drill detection. Configure cross-Region snapshots ol the DB instance to the DR Region every 5 minutes. In the event of a disaster, restore the DB instance using the snapshot in the DR Region.

C) Use Amazon EC2 Image Builder to create and copy AMIs of the web and application server to both the primary and DR Regions. Create a cross-Region read replica of the DB instance in the DR Region. In the event of a disaster, promote the read replica to become the master and reprovision the servers with AWS CloudFormation using the AMIs.

D) Create AMts of the web and application servers in the DR Region. Use scheduled AWS Glue jobs to synchronize the DB instance with another DB instance in the DR Region. In the event of a disaster, switch to the DB instance in the DR Region and reprovision the servers with AWS CloudFormation using the AMIs.

Answer:

C

Explanation:

deploying a brand new RDS instance will take >30 minutes. You will use EC2 Image builder to put the AMIs into the new region, but not use image builder to LAUNCH them.

Question 3

Question Type: MultipleChoice

A company is running a web application on Amazon EC2 instances in a production AWS account. The company requires all logs generated from the web application to be copied to a central AWS account for analysis and archiving. The company's AWS accounts are currently managed independently. Logging agents are configured on the EC2 instances to upload the tog files to an Amazon S3 bucket in the central AWS account.

A solutions architect needs to provide access for a solution that will allow the production account to store log files in the central account. The central account also needs to have read access to the tog files.

What should the solutions architect do to meet these requirements?

Options:

A) Create a cross-account role in the central account. Assume the role from the production account when the logs are being copied.

B) Create a policy on the S3 bucket with the production account ID as the principal. Allow S3 access from a delegated user.

C) Create a policy on the S3 bucket with access from only the CIDR range of the EC2 instances in the production account. Use the production account ID as the principal.

D) Create a cross-account role in the production account. Assume the role from the production account when the logs are being copied.

Answer:

A

Explanation:

Cross-account roles are used to grant access to resources in one AWS account, from another AWS account. In this case, a cross-account role should be created in the central account, which will grant access to the S3 bucket from the production account. The production account will then assume this role when the logs are being copied to the S3 bucket, allowing the log files to be stored in the central account.

For more information, please refer to the AWS documentation on Cross-Account Access and IAM Roles in AWS Organizations.

https://docs.aws.amazon.com/IAM/latest/UserGuide/tutorial_cross-account-with-roles.html

Question 4

Question Type: MultipleChoice

A company has many services running in its on-premises data center. The data center is connected to AWS using AWS Direct Connect (DX) and an IPSec VPN. The service data is sensitive and connectivity cannot traverse the internet. The company wants to expand into a new market segment and begin offering its services to other companies that are using AWS.

Which solution will meet these requirements?

Options:

A) Create a VPC Endpoint Service that accepts TCP traffic, host it behind a Network Load Balancer, and make the service available over DX.

B) Create a VPC Endpoint Service that accepts HTTP or HTTPS traffic, host it behind an Application Load Balancer, and make the service available over DX.

C) Attach an internet gateway to the VPC. and ensure that network access control and security group rules allow the relevant inbound and outbound traffic.

D) Attach a NAT gateway to the VPC. and ensure that network access control and security group rules allow the relevant inbound and outbound traffic.

Answer:

A

Explanation:

https://docs.aws.amazon.com/vpc/latest/privatelink/create-endpoint-service.html Endpoint services require either a Network Load Balancer or a Gateway Load Balancer.

Question 5

Question Type: MultipleChoice

A company runs an application on AWS. An AWS Lambda function uses credentials to authenticate to an Amazon RDS tor MySQL DB instance. A security risk assessment identified that these credentials are not frequently rotated. Also, encryption at rest is not enabled for the DB instance. The security team requires that both of these issues be resolved.

Which strategy should a solutions architect recommend to remediate these security risks?

Options:

A) Configure the Lambda function to store and retrieve the database credentials in AWS Secrets Manager and enable rotation of the credentials. Take a snapshot ol the DB instance and encrypt a copy of that snapshot. Replace the DB instance with a new DB instance that is based on the encrypted snapshot.

B) Enable IAM DB authentication on the DB instance. Grant the Lambda execution role access to the DB instance. Modify the DB instance and enable encryption.

C) Enable IAM DB authentication on the DB instance. Grant the Lambda execution role access to the DB instance. Create an encrypted read replica of the DB instance. Promote Ihe encrypted read replica to be the new primary node.

D) Configure the Lambda function to store and retrieve the database credentials as encrypted AWS Systems Manager Parameter Store parameters. Create another Lambda function to automatically rotate the credentials. Create an encrypted read replica of the DB instance. Promote the encrypted read replica to be the new primary node.

Answer:

A

Explanation:

Parameter store can store DB credentials as secure string but CANNOT rotate secrets, hence, go with A + Cannot enable encryption on existing MySQL RDS instance, must create a new encrypted one from unencrypted snapshot.

https://aws.amazon.com/blogs/security/rotate-amazon-rds-database-credentials-automatically-with-aws-secrets-manager/#:~:text=Secrets%20Manager%20offers%20built%2Din%20integrations%20for%20rotating%20credentials%20for,rotate%20other%20type

Encrypting a unencrypted instance of DB or creating a encrypted replica of an un encrypted DB instance are not possible Hence A is the only solution possible.
https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Overview.Encryption.html#Overview.Encryption.Limitations

Question 6

Question Type: MultipleChoice

A fitness tracking company serves users around the world, with its primary markets in North America and Asi

a. The company needs to design an infrastructure for its read-heavy user authorization application with the following requirements:

* Be resilient to problems with the application in any Region.

* Write to a database in a single Region.

* Read from multiple Regions.

* Support resiliency across application tiers in each Region.

* Support the relational database semantics reflected in the application.

Which combination of steps should a solutions architect take? (Select TWO.)

Options:

A) Use an Amazon Route 53 geoproximity routing policy combined with a multivalue answer routing policy.

B) Deploy web. application, and MySQL database servers to Amazon EC2 instances in each Region. Set up the application so that reads and writes are local to the Region. Create snapshots of the web, application, and database servers and store the snapshots in an Amazon S3 bucket in both Regions. Set up cross-Region replication for the database layer.

C) Use an Amazon Route 53 geolocation routing policy combined with a failover routing policy.

D) Set up web, application, and Amazon RDS for MySQL instances in each Region. Set up the application so that reads are local and writes are partitioned based on the user. Set up a Multi-AZ failover for the web, application, and database servers. Set up cross-Region replication for the database layer.

E) Set up active-active web and application servers in each Region. Deploy an Amazon Aurora global database with clusters in each Region. Set up the application to use the in-Region Aurora database endpoints. Create snapshots of the web and application servers and store them in an Amazon S3 bucket in both Regions.

Answer:

C, E

Explanation:

https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-policy.html

Geoproximity routing policy is good to control the user traffic to specific regions. However, a multivalue answer routing policy may cause the users to be randomly sent to other healthy regions that may be far away from the user's location. You can use geolocation routing

policy to direct the North American users to your servers on the North America region and configure failover routing to the Asia region in case the North America region fails. You can configure the same for the Asian users pointed to the Asia region servers and have the North America region as its backup.

Question 7

Question Type: MultipleChoice

An ecommerce website running on AWS uses an Amazon RDS for MySQL DB instance with General Purpose SSD storage. The developers chose an appropriate instance type based on demand, and configured 100 GB of storage with a sufficient amount of free space.

The website was running smoothly for a few weeks until a marketing campaign launched. On the second day of the campaign, users reported long wait times and time outs. Amazon CloudWatch metrics indicated that both reads and writes to the DB instance were experiencing long response times. The CloudWatch metrics show 40% to 50% CPU and memory utilization, and sufficient free storage space is still available. The application server logs show no evidence of database connectivity issues.

What could be the root cause of the issue with the marketing campaign?

Options:

A) It exhausted the I/O credit balance due to provisioning low disk storage during the setup phase.

B) It caused the data in the tables to change frequently, requiring indexes to be rebuilt to optimize queries.

C) It exhausted the maximum number of allowed connections to the database instance.

D) It exhausted the network bandwidth available to the RDS for MySQL DB instance.

Answer:

A

Explanation:

'When using General Purpose SSD storage, your DB instance receives an initial I/O credit balance of 5.4 million I/O credits. This initial credit balance is enough to sustain a burst performance of 3,000 IOPS for 30 minutes.'

https://aws.amazon.com/blogs/database/how-to-use-cloudwatch-metrics-to-decide-between-general-purpose-or-provisioned-iops-for-your-rds-database/

Question 8

Question Type: MultipleChoice

A company that tracks medical devices in hospitals wants to migrate its existing storage solution to the AWS Cloud. The company equips all of its devices with sensors that collect location and usage information. This sensor data is sent in unpredictable patterns with large spikes. The data is stored in a MySQL database running on premises at each hospital. The company wants the cloud storage solution to scale with usage.

The company's analytics team uses the sensor data to calculate usage by device type and hospital. The team needs to keep analysis tools running locally while fetching data from the cloud. The team also needs to use existing Java application and SQL queries with as few changes as possible.

How should a solutions architect meet these requirements while ensuring the sensor data is secure?

Options:

A) Store the data in an Amazon Aurora Serverless database. Serve the data through a Network Load Balancer (NLB). Authenticate users using the NLB with credentials stored in AWS Secrets Manager.

B) Store the data in an Amazon S3 bucket. Serve the data through Amazon QuickSight using an IAM user authorized with AWS Identity and Access Management (IAM) with the S3 bucket as the data source.

C) Store the data in an Amazon Aurora Serverless database. Serve the data through the Aurora Data API using an IAM user authorized with AWS Identity and Access Management (IAM) and the AWS Secrets Manager ARN.

D) Store the data in an Amazon S3 bucket. Serve the data through Amazon Athena using AWS PrivateLink to secure the data in transit.

Answer:

C

Explanation:

https://aws.amazon.com/blogs/aws/new-data-api-for-amazon-aurora-serverless/

https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/data-api.html

https://aws.amazon.com/blogs/aws/aws-privatelink-for-amazon-s3-now-available/

https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/data-api.html#data-api.access

The data is currently stored in a MySQL database running on-prem. Storing MySQL data in S3 doesn't sound good so B & D are out. Aurora Data API 'enables the SQL HTTP endpoint, a connectionless Web Service API for running SQL queries against this database. When the SQL HTTP endpoint is enabled, you can also query your database from inside the RDS console (these features are free to use).'

Question 9

Question Type: MultipleChoice

A company hosts a web application that tuns on a group of Amazon EC2 instances that ate behind an Application Load Balancer (ALB) in a VPC. The company wants to analyze the network payloads lo reverse-engineer a sophisticated attack of the application.

Which approach should the company take to achieve this goal?

Options:

A) Enable VPC Flow Logs. Store the flow logs in an Amazon S3 bucket for analysis.

B) Enable Traffic Mirroring on the network interface of the EC2 instances. Send the mirrored traffic lo a target for storage and analysis.

C) Create an AWS WAF web ACL. and associate it with the ALB. Configure AWS WAF logging.

D) Enable logging for the ALB. Store the logs in an Amazon S3 bucket for analysis.

Answer:

B

Explanation:

Traffic Mirroring allows to copy network traffic from a network interface to a destination network interface, Amazon EC2 instance or Amazon S3 bucket. The company can use Traffic Mirroring to analyze network payloads, detect sophisticated attacks and reverse-engineer the same.

Question 10

Question Type: MultipleChoice

A solutions architect is designing an application to accept timesheet entries from employees on their mobile devices. Timesheets will be submitted weekly, with most of the submissions occurring on Friday. The data must be stored in a format that allows payroll administrators to run monthly reports. The infrastructure must be highly available and scale to match the rate of incoming data and reporting requests.

Which combination of steps meets these requirements while minimizing operational overhead? (Select TWO.)

Options:

A) Deploy the application to Amazon EC2 On-Demand Instances With load balancing across multiple Availability Zones. Use scheduled Amazon EC2 Auto Scaling to add capacity before the high volume of submissions on Fridays.

B) Deploy the application in a container using Amazon Elastic Container Service (Amazon ECS) with load balancing across multiple Availability Zones. Use scheduled Service Auto Scaling to add capacity before the high volume of submissions on Fridays.

C) Deploy the application front end to an Amazon S3 bucket served by Amazon CloudFront. Deploy the application backend using Amazon API Gateway with an AWS Lambda proxy integration.

D) Store the timesheet submission data in Amazon Redshift. Use Amazon OuickSight to generate the reports using Amazon Redshift as the data source.

E) Store the timesheet submission data in Amazon S3. Use Amazon Athena and Amazon OuickSight to generate the reports using Amazon S3 as the data source.

Answer:

A, E