Free Questions for SAP-C02 by certsinside

Shared by Burgess on 15-04-2024

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

A public retail web application uses an Application Load Balancer (ALB) in front of Amazon EC2 instances running across multiple Availability Zones (AZs) in a Region backed by an Amazon RDS MySQL Multi-AZ deployment. Target group health checks are configured to use HTTP and pointed at the product catalog page. Auto Scaling is configured to maintain the web fleet size based on the ALB health check.

Recently, the application experienced an outage. Auto Scaling continuously replaced the instances during the outage. A subsequent investigation determined that the web server metrics were within the normal range, but the database tier was experiencing high toad, resulting in severely elevated query response times.

Which of the following changes together would remediate these issues while improving monitoring capabilities for the availability and functionality of the entire application stack for future growth? (Select TWO.)

Options:

A- Configure read replicas for Amazon RDS MySQL and use the single reader endpoint in the web application to reduce the load on the backend database tier.

B- Configure the target group health check to point at a simple HTML page instead of a product catalog page and the Amazon Route 53 health check against the product page to evaluate full application functionality. Configure Ama7on CloudWatch alarms to notify administrators when the site fails.

C- Configure the target group health check to use a TCP check of the Amazon EC2 web server and the Amazon Route S3 health check against the product page to evaluate full application functionality. Configure Amazon CloudWatch alarms to notify administrators when the site fails.

D- Configure an Amazon CtoudWatch alarm for Amazon RDS with an action to recover a high-load, impaired RDS instance in the database tier.

E- Configure an Amazon Elastic ache cluster and place it between the web application and RDS MySQL instances to reduce the load on the backend database tier.

Answer:

A, E

Explanation:

Configuring read replicas for Amazon RDS MySQL and using the single reader endpoint in the web application can significantly reduce the load on the backend database tier, improving overall application performance. Additionally, implementing an Amazon ElastiCache cluster between the web application and RDS MySQL instances can further reduce database load by caching frequently accessed data, thereby enhancing the application's resilience and scalability. These changes address the root cause of the outage by alleviating the database tier's high load and preventing similar issues in the future.

Question 2

Question Type: MultipleChoice

A company deploys workloads in multiple AWS accounts. Each account has a VPC with VPC flow logs published in text log format to a centralized Amazon S3 bucket. Each log file is compressed with gzjp compression. The company must retain the log files indefinitely.

A security engineer occasionally analyzes the togs by using Amazon Athena to query the VPC flow logs. The query performance is degrading over time as the number of ingested togs is growing. A solutions architect: must improve the performance of the tog analysis and reduce the storage space that the VPC flow logs use.

Which solution will meet these requirements with the LARGEST performance improvement?

Options:

A- Create an AWS Lambda function to decompress the gzip flies and to compress the tiles with bzip2 compression. Subscribe the Lambda function to an s3: ObiectCrealed;Put S3 event notification for the S3 bucket.

B- Enable S3 Transfer Acceleration for the S3 bucket. Create an S3 Lifecycle configuration to move files to the S3 Intelligent-Tiering storage class as soon as the ties are uploaded

C- Update the VPC flow log configuration to store the files in Apache Parquet format. Specify Hourly partitions for the log files.

D- Create a new Athena workgroup without data usage control limits. Use Athena engine version 2.

Answer:

C

Explanation:

Converting VPC flow logs to store in Apache Parquet format and specifying hourly partitions significantly improves query performance and reduces storage space usage. Apache Parquet is a columnar storage file format optimized for analytical queries, allowing Athena to scan less data and improve query performance. Partitioning logs by hour further enhances query efficiency by limiting the amount of data scanned during queries, addressing the issue of degrading performance over time due to the growing volume of ingested logs.

Question 3

Question Type: MultipleChoice

A software development company has multiple engineers who ate working remotely. The company is running Active Directory Domain Services (AD DS) on an Amazon EC2 instance. The company's security policy states that al internal, nonpublic services that are deployed in a VPC must be accessible through a VPN. Multi-factor authentication (MFA) must be used for access to a VPN.

What should a solutions architect do to meet these requirements?

Options:

A- Create an AWS Sire-to-Site VPN connection. Configure Integration between a VPN and AD DS. Use an Amazon Workspaces client with MFA support enabled to establish a VPN connection.

B- Create an AWS Client VPN endpoint Create an AD Connector directory tor integration with AD DS. Enable MFA tor AD Connector. Use AWS Client VPN to establish a VPN connection.

C- Create multiple AWS Site-to-Site VPN connections by using AWS VPN CloudHub. Configure integration between AWS VPN CloudHub and AD DS. Use AWS Copilot to establish a VPN connection.

D- Create an Amazon WorkLink endpoint. Configure integration between Amazon WorkLink and AD DS. Enable MFA in Amazon WorkLink. Use AWS Client VPN to establish a VPN connection.

Answer:

B

Explanation:

Setting up an AWS Client VPN endpoint and integrating it with Active Directory Domain Services (AD DS) using an AD Connector directory enables secure remote access to internal services deployed in a VPC. Enabling multi-factor authentication (MFA) for AD Connector enhances security by adding an additional layer of authentication. This solution meets the company's requirements for secure remote access through a VPN with MFA, ensuring that the security policy is adhered to while providing a seamless experience for the remote engineers.

Question 4

Question Type: MultipleChoice

A company needs to improve the reliability ticketing application. The application runs on an Amazon Elastic Container Service (Amazon ECS) cluster. The company uses Amazon CloudFront to servo the application. A single ECS service of the ECS cluster is the CloudFront distribution's origin.

The application allows only a specific number of active users to enter a ticket purchasing flow. These users are identified by an encrypted attribute in their JSON Web Token (JWT). All other users are redirected to a waiting room module until there is available capacity for purchasing.

The application is experiencing high loads. The waiting room modulo is working as designed, but load on the waiting room is disrupting the application's availability. This disruption is negatively affecting the application's ticket sale Transactions.

Which solution will provide the MOST reliability for ticket sale transactions during periods of high load? '

Options:

A- Create a separate service in the ECS cluster for the waiting room. Use a separate scaling configuration. Ensure that the ticketing service uses the JWT info-nation and appropriately forwards requests to the waring room service.

B- Move the application to an Amazon Elastic Kubernetes Service (Amazon EKS) cluster. Split the wailing room module into a pod that is separate from the ticketing pod. Make the ticketing pod part of a StatefuISeL Ensure that the ticketing pod uses the JWT information and appropriately forwards requests to the waiting room pod.

C- Create a separate service in the ECS cluster for the waiting room. Use a separate scaling configuration. Create a CloudFront function That inspects the JWT information and appropriately forwards requests to the ticketing service or the waiting room service

D- Move the application to an Amazon Elastic Kubernetes Service (Amazon EKS) cluster. Split the wailing room module into a pod that is separate from the ticketing pod. Use AWS App Mesh by provisioning the App Mesh controller for Kubermetes. Enable mTLS authentication and service-to-service authentication for communication between the ticketing pod and the waiting room pod. Ensure that the ticketing pod uses The JWT information and appropriately forwards requests to the waiting room pod.

Answer:

C

Explanation:

Implementing a CloudFront function that inspects the JWT information and appropriately forwards requests either to the ticketing service or the waiting room service within the Amazon ECS cluster enhances reliability during high load periods. This solution segregates the load between the main application and the waiting room, ensuring that the ticketing service remains unaffected by the high load on the waiting room. Using CloudFront functions for request routing based on JWT attributes allows for efficient distribution of user traffic, thereby maintaining the application's availability and performance during peak times.

Question 5

Question Type: MultipleChoice

A flood monitoring agency has deployed more than 10.000 water-level monitoring sensors. Sensors send continuous data updates, and each update is less than 1 MB in size. The agency has a fleet of on-premises application servers. These servers receive upda.es 'on the sensors, convert the raw data into a human readable format, and write the results loan on-premises relational database server. Data analysts then use simple SOL queries to monitor the data.

The agency wants to increase overall application availability and reduce the effort that is required to perform maintenance tasks These maintenance tasks, which include updates and patches to the application servers, cause downtime. While an application server is down, data is lost from sensors because the remaining servers cannot handle the entire workload.

The agency wants a solution that optimizes operational overhead and costs. A solutions architect recommends the use of AWS IoT Core to collect the sensor data.

What else should the solutions architect recommend to meet these requirements?

Options:

A- Send the sensor data to Amazon Kinesis Data Firehose. Use an AWS Lambda function to read the Kinesis Data Firehose data, convert it to .csv format, and insert it into an Amazon Aurora MySQL DB instance. Instruct the data analysts to query the data directly from the DB instance.

B- Send the sensor data to Amazon Kinesis Data Firehose. Use an AWS Lambda function to read the Kinesis Data Firehose data, convert it to Apache Parquet format and save it to an Amazon S3 bucket. Instruct the data analysts to query the data by using Amazon Athena.

C- Send the sensor data to an Amazon Managed Service for Apache Flink {previously known as Amazon Kinesis Data Analytics) application to convert the data to .csv format and store it in an Amazon S3 bucket. Import the data into an Amazon Aurora MySQL DB instance. Instruct the data analysts to query the data directly from the DB instance.

D- Send the sensor data to an Amazon Managed Service for Apache Flink (previously known as Amazon Kinesis Data Analytics) application to convert the data to Apache Parquet format and store it in an Amazon S3 bucket Instruct the data analysis to query the data by using Amazon Athena.

Answer:

B

Explanation:

To enhance application availability and reduce maintenance-induced downtime, sending sensor data to Amazon Kinesis Data Firehose, processing it with an AWS Lambda function, converting it to Apache Parquet format, and storing it in Amazon S3 is an effective strategy. This approach leverages serverless architectures for scalability and reliability. Data analysts can then query the optimized data using Amazon Athena, a serverless interactive query service, which supports complex queries on data stored in S3 without the need for traditional database servers, optimizing operational overhead and costs.

Question 6

Question Type: **MultipleChoice**

A company needs to gather data from an experiment in a remote location that does not have internet connectivity. During the experiment, sensors that are connected to a total network will generate 6 TB of data in a preprimary formal over the course of 1 week. The sensors can be configured to upload their data files to an FTP server periodically, but the sensors do not have their own FTP server. The sensors also do not support other protocols. The company needs to collect the data centrally and move lie data to object storage in the AWS Cloud as soon. as possible after the experiment.

Which solution will meet these requirements?

Options:

A- Order an AWS Snowball Edge Compute Optimized device. Connect the device to the local network. Configure AWS DataSync with a target bucket name, and unload the data over NFS to the device. After the experiment return the device to AWS so that the data can be loaded into Amazon S3.

B- Order an AWS Snowcone device, including an Amazon Linux 2 AMI. Connect the device to the local network. Launch an Amazon EC2 instance on the device. Create a shell script that periodically downloads data from each sensor. After the experiment, return the device to AWS so that the data can be loaded as an Amazon Elastic Block Store [Amazon EBS) volume.

C- Order an AWS Snowcone device, including an Amazon Linux 2 AMI. Connect the device to the local network. Launch an Amazon EC2 instance on the device. Install and configure an FTP server on the EC2 instance. Configure the sensors to upload data to the EC2 instance. After the experiment, return the device to AWS so that the data can be loaded into Amazon S3.

D- Order an AWS Snowcone device. Connect the device to the local network. Configure the device to use Amazon FSx. Configure the sensors to upload data to the device. Configure AWS DataSync on the device to synchronize the uploaded data with an Amazon S3 bucket Return the device to AWS so that the data can be loaded as an Amazon Elastic Block Store (Amazon EBS) volume.

Answer:

C

Explanation:

For collecting data from remote sensors without internet connectivity, using an AWS Snowcone device with an Amazon EC2 instance running an FTP server presents a practical solution. This setup allows the sensors to upload data to the EC2 instance via FTP, and after the experiment, the Snowcone device can be returned to AWS for data ingestion into Amazon S3. This approach minimizes operational complexity and ensures efficient data transfer to AWS for further processing or storage.

Question 7

Question Type: **MultipleChoice**

A company has a Windows-based desktop application that is packaged and deployed to the users' Windows machines. The company recently acquired another company that has employees who primarily use machines with a Linux operating system. The acquiring company has decided to migrate and rehost the Windows-based desktop application lo AWS.

All employees must be authenticated before they use the application. The acquiring company uses Active Directory on premises but wants a simplified way to manage access to the application on AWS (or all the employees.

Which solution will rehost the application on AWS with the LEAST development effort?

Options:

A- Set up and provision an Amazon Workspaces virtual desktop for every employee. Implement authentication by using Amazon Cognito identity pools. Instruct employees to run the application from their provisioned Workspaces virtual desktops.

B- Create an Auto Scarlet group of Windows-based Ama7on EC2 instances. Join each EC2 instance to the company's Active Directory domain. Implement authentication by using the Active Directory That is running on premises. Instruct employees to run the application by using a Windows remote desktop.

C- Use an Amazon AppStream 2.0 image builder to create an image that includes the application and the required configurations. Provision an AppStream 2.0 On-Demand fleet with dynamic Fleet Auto Scaling process for running the image. Implement authentication by using AppStream 2.0 user pools. Instruct the employees to access the application by starling browse'-based AppStream 2.0 streaming sessions.

D- Refactor and containerize the application to run as a web-based application. Run the application in Amazon Elastic Container Service (Amazon ECS) on AWS Fargate with step scaling policies Implement authentication by using Amazon Cognito user pools. Instruct the employees to run the application from their browsers.

Answer:

C

Explanation:

Amazon AppStream 2.0 offers a streamlined solution for rehosting a Windows-based desktop application on AWS with minimal development effort. By creating an AppStream 2.0 image that includes the application and using an On-Demand fleet for streaming, the application becomes accessible from any device, including Linux machines. AppStream 2.0 user pools can be used for authentication, simplifying access management without the need for extensive changes to the application or infrastructure.

Question 8

Question Type: MultipleChoice

A company wants to use Amazon Workspaces in combination with thin client devices to replace aging desktops. Employees use the desktops to access applications that work with clinical trial dat

a. Corporate security policy states that access to the applications must be restricted to only company branch office locations. The company is considering adding an additional branch office in the next 6 months.

Which solution meets these requirements with the MOST operational efficiency?

Options:

A- Create an IP access control group rule with the list of public addresses from the branch offices. Associate the IP access control group with the Workspaces directory.

B- Use AWS Firewall Manager to create a web ACL rule with an IPSet with the list to public addresses from the branch office Locations-Associate the web ACL with the Workspaces directory.

C- Use AWS Certificate Manager (ACM) to issue trusted device certificates to the machines deployed in the branch office locations. Enable restricted access on the Workspaces directory.

D- Create a custom Workspace image with Windows Firewall configured to restrict access to the public addresses of the branch offices. Use the image to deploy the Workspaces.

Answer:

A

Explanation:

Utilizing an IP access control group rule with the list of public addresses from branch offices and associating it with the Amazon WorkSpaces directory is the most operationally efficient solution. This method ensures that access to WorkSpaces is restricted to specified locations, aligning with the corporate security policy. This approach offers simplicity and flexibility, especially with the potential addition of a new branch office, as updating the IP access control group is straightforward.

Question 9

Question Type: **MultipleChoice**

A company needs to implement disaster recovery for a critical application that runs in a single AWS Region. The application's users interact with a web frontend that is hosted on Amazon EC2 Instances behind an Application Load Balancer (ALB). The application writes to an Amazon RD5 tor MySQL DB instance. The application also outputs processed documents that are stored in an Amazon S3 bucket

The company's finance team directly queries the database to run reports. During busy periods, these queries consume resources and negatively affect application performance.

A solutions architect must design a solution that will provide resiliency during a disaster. The solution must minimize data loss and must resolve the performance problems that result from the finance team's queries.

Which solution will meet these requirements?

Options:

A- Migrate the database to Amazon DynamoDB and use DynamoDB global tables. Instruct the finance team to query a global table in a separate Region. Create an AWS Lambda function to periodically synchronize the contents of the original S3 bucket to a new S3 bucket in the separate Region. Launch EC2 instances and create an ALB in the separate Region. Configure the application to point to the new S3 bucket.

B- Launch additional EC2 instances that host the application in a separate Region. Add the additional instances to the existing ALB. In the separate Region, create a read replica of the RDS DB instance. Instruct the finance team to run queries ageist the read replica. Use S3 Cross-Region Replication (CRR) from the original S3 bucket to a new S3 Docket in the separate Region. During a disaster, promote the read replace to a standalone DB instance. Configure the application to point to the new S3 bucket and to the newly project read replica.

C- Create a read replica of the RDS DB instance in a separate Region. Instruct the finance team to run queries against the read replica. Create AMIs of the EC2 instances mat host the application frontend- Copy the AMIs to the separate Region. Use S3 Cross-Region Replication (CRR) from the original S3 bucket to a new S3 bucket in the separate Region. During a disaster, promote the read replica to a standalone DB instance. Launch EC2 instances from the AMIs and create an ALB to present the application to end users. Configure the application to point to the new S3 bucket.

D- Create hourly snapshots of the RDS DB instance. Copy the snapshots to a separate Region. Add an Amazon Elastic ache cluster m front of the existing RDS database. Create AMIs of the EC2 instances that host the application frontend Copy the AMIs to the separate Region. Use S3 Cross-Region Replication (CRR) from the original S3 bucket to a new S3 bucket in the separate Region. During a disaster, restore The database from the latest RDS snapshot. Launch EC2 Instances from the AMIs and create an ALB to present the application to end users. Configure the application to point to the new S3 bucket

Answer:

C

Explanation:

Implementing a disaster recovery strategy that minimizes data loss and addresses performance issues involves creating a read replica of the RDS DB instance in a separate region and directing the finance team's queries to this replica. This solution alleviates the performance impact on the primary database. Using Amazon S3 Cross-Region Replication (CRR) ensures that processed documents are available in the disaster recovery region. In the event of a disaster, the read replica can be promoted to a standalone DB instance, and EC2 instances can be launched from pre-created AMIs to serve the web frontend, thereby ensuring resiliency and minimal data loss.

Question 10

Question Type: MultipleChoice

A company is designing an AWS environment tor a manufacturing application. The application has been successful with customers, and the application's user base has increased. The company has connected the AWS environment to the company's on-premises data center through a 1 Gbps AWS Direct Connect connection. The company has configured BGP for the connection.

The company must update the existing network connectivity solution to ensure that the solution is highly available, fault tolerant, and secure.

Which solution win meet these requirements MOST cost-effectively?

Options:

A- Add a dynamic private IP AWS Site-to-Site VPN as a secondary path to secure data in transit and provide resilience for the Direct Conned connection. Configure MACsec to encrypt traffic inside the Direct Connect connection.

B- Provision another Direct Conned connection between the company's on-premises data center and AWS to increase the transfer speed and provide resilience. Configure MACsec to encrypt traffic inside the Dried Conned connection.

C- Configure multiple private VIFs. Load balance data across the VIFs between the on-premises data center and AWS to provide resilience.

D- Add a static AWS Site-to-Site VPN as a secondary path to secure data in transit and to provide resilience for the Direct Connect connection.

Answer:

A

Explanation:

To enhance the network connectivity solution's availability, fault tolerance, and security in a cost-effective manner, adding a dynamic private IP AWS Site-to-Site VPN as a secondary path is a viable option. This VPN serves as a resilient backup for the Direct Connect connection, ensuring continuous data flow even if the primary path fails. Implementing MACsec (Media Access Control Security) on the Direct Connect connection further secures the data in transit by providing encryption, thus addressing the security requirement. This solution strikes a balance between cost and operational efficiency, avoiding the higher expenses associated with provisioning an additional Direct Connect connection.

Question 11

Question Type: MultipleChoice

A company wants to establish a dedicated connection between its on-premises infrastructure and AWS. The company is setting up a 1 Gbps AWS Direct Connect connection to its account VPC. The architecture includes a transit gateway and a Direct Connect gateway to connect multiple VPCs and the on-premises infrastructure.

The company must connect to VPC resources over a transit VIF by using the Direct Connect connection.

Which combination of steps will meet these requirements? (Select TWO.)

Options:

A- Update the 1 Gbps Direct Connect connection to 10 Gbps.

B- Advertise the on-premises network prefixes over the transit VIF.

C- Adverse the VPC prefixes from the Direct Connect gateway to the on-premises network over the transit VIF.

D- Update the Direct Connect connection's MACsec encryption mode attribute to must encrypt.

E- Associate a MACsec Connection Key Name-Connectivity Association Key (CKN/CAK) pair with the Direct Connect connection.

Answer:

B, C

Explanation:

To connect VPC resources over a transit Virtual Interface (VIF) using a Direct Connect connection, the company should advertise the on-premises network prefixes over the transit VIF and advertise the VPC prefixes from the Direct Connect gateway to the on-premises network over the same VIF. This configuration ensures seamless connectivity between the on-premises infrastructure and the AWS VPCs through the transit gateway, facilitating efficient and secure communication across the network.

Question 12

Question Type: **MultipleChoice**

A company's compliance audit reveals that some Amazon Elastic Block Store (Amazon EBS) volumes that were created in an AWS account were not encrypted. A solutions architect must Implement a solution to encrypt all new EBS volumes at rest

Which solution will meet this requirement with the LEAST effort?

Options:

A- Create an Amazon EventBridge rule to detect the creation of unencrypted EBS volumes. Invoke an AWS Lambda function to delete noncompliant volumes.

B- Use AWS Audit Manager with data encryption.

C- Create an AWS Config rule to detect the creation of a new EBS volume. Encrypt the volume by using AWS Systems Manager Automation.

D- Turn in EBS encryption by default in all AWS Regions.

Answer:

D

Explanation:

The most effortless way to ensure that all new Amazon Elastic Block Store (EBS) volumes are encrypted at rest is to enable EBS encryption by default in all AWS Regions. This setting automatically encrypts all new EBS volumes and snapshots created in the account, thereby ensuring compliance with encryption policies without the need for manual intervention or additional monitoring.