

[June-2021] Braindump2go SAA-C02 VCE and PDF Dumps SAA-C02 669Q for 100% Passing SAA-C02 Exam [Q654-Q669]

June/2021 Latest Braindump2go SAA-C02 Exam Dumps with PDF and VCE Free Updated Following are some new SAA-C02

Real Exam Questions! QUESTION 654A company is building a web-based application running on Amazon EC2 instances in multiple Availability Zones. The web application will provide access to a repository of text documents totaling about 900 TB in size. The company anticipates that the web application will experience periods of high demand. A solutions architect must ensure that the storage component for the text documents can scale to meet the demand of the application at all times. The company is concerned about the overall cost of the solution. Which storage solution meets these requirements MOST cost-effectively? A. Amazon Elastic Block Store (Amazon EBS) B. Amazon Elastic File System (Amazon EFS) C. Amazon Elasticsearch Service (Amazon ES) D. Amazon S3

Answer: C QUESTION 655A company has an application that uses overnight digital images of products on store shelves to analyze inventory data. The application runs on Amazon EC2 instances behind an Application Load Balancer (ALB) and obtains the images from an Amazon S3 bucket for its metadata to be processed by worker nodes for analysis. A solutions architect needs to ensure that every image is processed by the worker nodes. What should the solutions architect do to meet this requirement in the MOST cost-efficient way? A. Send the image metadata from the application directly to a second ALB for the worker nodes that use an Auto Scaling group of EC2 Spot Instances as the target group. B. Process the image metadata by sending it directly to EC2 Reserved Instances in an Auto Scaling group. With a dynamic scaling policy, use an Amazon CloudWatch metric for average CPU utilization of the Auto Scaling group as soon as the front-end application obtains the images. C. Write messages to Amazon Simple Queue Service (Amazon SQS) when the front-end application obtains an image. Process the images with EC2 On-Demand instances in an Auto Scaling group with instance scale-in protection and a fixed number of instances with periodic health checks. D. Write messages to Amazon Simple Queue Service (Amazon SQS) when the application obtains an image. Process the images with EC2 Spot Instances in an Auto Scaling group with instance scale-in protection and a dynamic scaling policy using a custom Amazon CloudWatch metric for the current number of messages in the queue.

Answer: B QUESTION 656A company is using a centralized AWS account to store log data in various Amazon S3 buckets. A solutions architect needs to ensure that the data is encrypted at rest before the data is uploaded to the S3 buckets. The data also must be encrypted in transit. Which solution meets these requirements? A. Use client-side encryption to encrypt the data that is being uploaded to the S3 buckets. B. Use server-side encryption to encrypt the data that is being uploaded to the S3 buckets. C. Create bucket policies that require the use of server-side encryption with S3 managed encryption keys (SSE-S3) for S3 uploads. D. Enable the security option to encrypt the S3 buckets through the use of a default AWS Key Management Service (AWS KMS) key.

Answer: B QUESTION 657A company's HTTP application is behind a Network Load Balancer (NLB). The NLB's target group is configured to use an Amazon EC2 Auto Scaling group with multiple EC2 instances that run the web service. The company notices that the NLB is not detecting HTTP errors for the application. These errors require a manual restart of the EC2 instances that run the web service. The company needs to improve the application's availability without writing custom scripts or code. What should a solutions architect do to meet these requirements? A. Enable HTTP health checks on the NLB, supplying the URL of the company's application. B. Add a cron job to the EC2 instances to check the local application's logs once each minute. If HTTP errors are detected, the application will restart. C. Replace the NLB with an Application Load Balancer. Enable HTTP health checks by supplying the URL of the company's application. Configure an Auto Scaling action to replace unhealthy instances. D. Create an Amazon CloudWatch alarm that monitors the UnhealthyHostCount metric for the NLB. Configure an Auto Scaling action to replace unhealthy instances when the alarm is in the ALARM state.

Answer: C QUESTION 658A company has two VPCs that are located in the us-west-2 Region within the same AWS account. The company needs to allow network traffic between these VPCs. Approximately 500 GB of data transfer will occur between the VPCs each month. What is the MOST cost-effective solution to connect these VPCs? A. Implement AWS Transit Gateway to connect the VPCs. Update the route tables of each VPC to use the transit gateway for inter-VPC communication. B. Implement an AWS Site-to-Site VPN tunnel between the VPCs. Update the route tables of each VPC to use the VPN tunnel for inter-VPC communication. C. Set up a VPC peering connection between the VPCs. Update the route tables of each VPC to use the VPC peering connection for inter-VPC communication. D. Set up a 1 GB AWS Direct Connect connection between the VPCs. Update the route tables of each VPC to use the Direct Connect connection for inter-VPC communication.

Answer: D QUESTION 659A company is running a global application. The application's users submit multiple videos that are then merged into a single video file. The application uses a single Amazon S3 bucket in the us-east-1 Region to receive uploads from users. The same S3 bucket provides the download location of the single video file that is produced. The final video file output has an average size of 250 GB. The company needs to develop a solution that delivers faster uploads and downloads of the video files that are stored in Amazon S2. The

QUESTION 659A company is running a global application. The application's users submit multiple videos that are then merged into a single video file. The application uses a single Amazon S3 bucket in the us-east-1 Region to receive uploads from users. The same S3 bucket provides the download location of the single video file that is produced. The final video file output has an average size of 250 GB. The company needs to develop a solution that delivers faster uploads and downloads of the video files that are stored in Amazon S2. The

QUESTION 659A company is running a global application. The application's users submit multiple videos that are then merged into a single video file. The application uses a single Amazon S3 bucket in the us-east-1 Region to receive uploads from users. The same S3 bucket provides the download location of the single video file that is produced. The final video file output has an average size of 250 GB. The company needs to develop a solution that delivers faster uploads and downloads of the video files that are stored in Amazon S2. The

QUESTION 659A company is running a global application. The application's users submit multiple videos that are then merged into a single video file. The application uses a single Amazon S3 bucket in the us-east-1 Region to receive uploads from users. The same S3 bucket provides the download location of the single video file that is produced. The final video file output has an average size of 250 GB. The company needs to develop a solution that delivers faster uploads and downloads of the video files that are stored in Amazon S2. The

company will offer the solution as a subscription to users who want to pay for the increased speed. What should a solutions architect do to meet these requirements?

A. Enable AWS Global Accelerator for the S3 endpoint. Adjust the application's upload and download links to use the Global Accelerator S3 endpoint for users who have a subscription.

B. Enable S3 Cross-Region Replication to S3 buckets in all other AWS Regions. Use an Amazon Route 53 geolocation routing policy to route S3 requests based on the location of users who have a subscription.

C. Create an Amazon CloudFront distribution and use the S3 bucket in us-east-1 as an origin. Adjust the application to use the CloudFront URL as the upload and download links for users who have a subscription.

D. Enable S3 Transfer Acceleration for the S3 bucket in us-east-1. Configure the application to use the bucket's S3-accelerate endpoint domain name for the upload and download links for users who have a subscription.

Answer: C

QUESTION 660A company is using AWS Organizations with two AWS accounts: Logistics and Sales. The Logistics account operates an Amazon Redshift cluster. The Sales account includes Amazon EC2 instances. The Sales account needs to access the Logistics account's Amazon Redshift cluster. What should a solutions architect recommend to meet this requirement MOST cost-effectively?

A. Set up VPC sharing with the Logistics account as the owner and the Sales account as the participant to transfer the data.

B. Create an AWS Lambda function in the Logistics account to transfer data to the Amazon EC2 instances in the Sales account.

C. Create a snapshot of the Amazon Redshift cluster, and share the snapshot with the Sales account. In the Sales account, restore the cluster by using the snapshot ID that is shared by the Logistics account.

D. Run COPY commands to load data from Amazon Redshift into Amazon S3 buckets in the Logistics account. Grant permissions to the Sales account to access the S3 buckets of the Logistics account.

Answer: C

QUESTION 661A company is using Amazon Redshift for analytics and to generate customer reports. The company recently acquired 50 TB of additional customer demographic data. The data is stored in .csv files in Amazon S3. The company needs a solution that joins the data and visualizes the results with the least possible cost and effort. What should a solutions architect recommend to meet these requirements?

A. Use Amazon Redshift Spectrum to query the data in Amazon S3 directly and join that data with the existing data in Amazon Redshift. Use Amazon QuickSight to build the visualizations.

B. Use Amazon Athena to query the data in Amazon S3. Use Amazon QuickSight to join the data from Athena with the existing data in Amazon Redshift and to build the visualizations.

C. Increase the size of the Amazon Redshift cluster, and load the data from Amazon S3. Use Amazon EMR Notebooks to query the data and build the visualizations in Amazon Redshift.

D. Export the data from the Amazon Redshift cluster into Apache Parquet files in Amazon S3. Use Amazon Elasticsearch Service (Amazon ES) to query the data. Use Kibana to visualize the results.

Answer: A

QUESTION 662A company's database is hosted on an Amazon Aurora MySQL DB cluster in the us-east-1 Region. The database is 4 TB in size. The company needs to expand its disaster recovery strategy to the us-west-2 Region. The company must have the ability to fail over to us-west-2 with a recovery time objective (RTO) of 15 minutes. What should a solutions architect recommend to meet these requirements?

A. Create a Multi-Region Aurora MySQL DB cluster in us-east-1 and us-west-2. Use an Amazon Route 53 health check to monitor us-east-1 and fail over to us-west-2 upon failure.

B. Take a snapshot of the DB cluster in us-east-1. Configure an Amazon EventBridge (Amazon CloudWatch Events) rule that invokes an AWS Lambda function upon receipt of resource events. Configure the Lambda function to copy the snapshot to us-west-2 and restore the snapshot in us-west-2 when failure is detected.

C. Create an AWS CloudFormation script to create another Aurora MySQL DB cluster in us-west-2 in case of failure. Configure an Amazon EventBridge (Amazon CloudWatch Events) rule that invokes an AWS Lambda function upon receipt of resource events. Configure the Lambda function to deploy the AWS CloudFormation stack in us-west-2 when failure is detected.

D. Recreate the database as an Aurora global database with the primary DB cluster in us-east-1 and a secondary DB cluster in us-west-2. Configure an Amazon EventBridge (Amazon CloudWatch Events) rule that invokes an AWS Lambda function upon receipt of resource events. Configure the Lambda function to promote the DB cluster in us-west-2 when failure is detected.

Answer: B

QUESTION 663A company is migrating its applications to AWS. Currently, applications that run on premises generate hundreds of terabytes of data that is stored on a shared file system. The company is running an analytics application in the cloud that runs hourly to generate insights from this data. The company needs a solution to handle the ongoing data transfer between the on-premises shared file system and Amazon S3. The solution also must be able to handle occasional interruptions in internet connectivity. Which solutions should the company use for the data transfer to meet these requirements?

A. AWS DataSync

B. AWS Migration Hub

C. AWS Snowball Edge Storage Optimized

D. AWS Transfer for SFTP

Answer: A

QUESTION 664A solutions architect is designing the architecture for a new web application. The application will run on AWS Fargate containers with an Application Load Balancer (ALB) and an Amazon Aurora PostgreSQL database. The web application will perform primarily read queries against the database. What should the solutions architect do to ensure that the website can scale with increasing traffic? (Choose two.)

A. Enable auto scaling on the ALB to scale the load balancer horizontally.

B. Configure Aurora Auto Scaling to adjust the number of Aurora Replicas in the Aurora cluster dynamically.

C. Enable cross-zone load balancing on the ALB to distribute the load evenly across containers in all Availability Zones.

D. Configure an Amazon Elastic

Container Service (Amazon ECS) cluster in each Availability Zone to distribute the load across multiple Availability Zones.E. Configure Amazon Elastic Container Service (Amazon ECS) Service Auto Scaling with a target tracking scaling policy that is based on CPU utilization.

Answer: BQUESTION 665A company captures ordered clickstream data from multiple websites and uses batch processing to analyze the data. The company receives 100 million event records, all approximately 1 KB in size, each day. The company loads the data into Amazon Redshift each night, and business analysts consume the data. The company wants to move toward near-real-time data processing for timely insights. The solution should process the streaming data while requiring the least possible operational overhead. Which combination of AWS services will meet these requirements MOST cost-effectively? (Choose two.)

A. Amazon EC2 B. AWS Batch C. Amazon Simple Queue Service (Amazon SQS) D. Amazon Kinesis Data Firehose E. Amazon Kinesis Data Analytics

Answer: CEQUESTION 666A company has a customer relationship management (CRM) application that stores data in an Amazon RDS DB instance that runs Microsoft SQL Server. The company's IT staff has administrative access to the database. The database contains sensitive data. The company wants to ensure that the data is not accessible to the IT staff and that only authorized personnel can view the data. What should a solutions architect do to secure the data?

A. Use client-side encryption with an Amazon RDS managed key. B. Use client-side encryption with an AWS Key Management Service (AWS KMS) customer managed key. C. Use Amazon RDS encryption with an AWS Key Management Service (AWS KMS) default encryption key. D. Use Amazon RDS encryption with an AWS Key Management Service (AWS KMS) customer managed key.

Answer: CQUESTION 667A company with a single AWS account runs its internet-facing containerized web application on an Amazon Elastic Kubernetes Service (Amazon EKS) cluster. The EKS cluster is placed in a private subnet of a VPC. System administrators access the EKS cluster through a bastion host on a public subnet. A new corporate security policy requires the company to avoid the use of bastion hosts. The company also must not allow internet connectivity to the EKS cluster. Which solution meets these requirements MOST cost-effectively?

A. Set up an AWS Direct Connect connection. B. Create a transit gateway. C. Establish a VPN connection. D. Use AWS Storage Gateway.

Answer: BQUESTION 668A company has deployed a multiplayer game for mobile devices. The game requires live location tracking of players based on latitude and longitude. The data store for the game must support rapid updates and retrieval of locations. The game uses an Amazon RDS for PostgreSQL DB instance with read replicas to store the location data. During peak usage periods, the database is unable to maintain the performance that is needed for reading and writing updates. The game's user base is increasing rapidly. What should a solutions architect do to improve the performance of the data tier?

A. Take a snapshot of the existing DB instance. Restore the snapshot with Multi-AZ enabled. B. Migrate from Amazon RDS to Amazon Elasticsearch Service (Amazon ES) with Kibana. C. Deploy Amazon DynamoDB Accelerator (DAX) in front of the existing DB instance. Modify the game to use DAX. D. Deploy an Amazon ElastiCache for Redis cluster in front of the existing DB instance. Modify the game to use Redis.

Answer: CQUESTION 669A company is migrating a large, mission-critical database to AWS. A solutions architect has decided to use an Amazon RDS for MySQL Multi-AZ DB instance that is deployed with 80,000 Provisioned IOPS for storage. The solutions architect is using AWS Database Migration Service (AWS DMS) to perform the data migration. The migration is taking longer than expected, and the company wants to speed up the process. The company's network team has ruled out bandwidth as a limiting factor. Which actions should the solutions architect take to speed up the migration? (Choose two.)

A. Disable Multi-AZ on the target DB instance. B. Create a new DMS instance that has a larger instance size. C. Turn off logging on the target DB instance until the initial load is complete. D. Restart the DMS task on a new DMS instance with transfer acceleration enabled. E. Change the storage type on the target DB instance to Amazon Elastic Block Store (Amazon EBS) General Purpose SSD (gp2).

Answer: CD

Resources From: 1. 2021 Latest Braindump2go SAA-C02 Exam Dumps (PDF & VCE) Free Share: <https://www.braindump2go.com/saa-c02.html> 2. 2021 Latest Braindump2go SAA-C02 PDF and SAA-C02 VCE Dumps Free Share: https://drive.google.com/drive/folders/1_5IK3H_eM74C6AKwU7sKaLn1rrn8xTfm?usp=sharing 3. 2021 Free Braindump2go SAA-C02 Exam Questions Download: [https://www.braindump2go.com/free-online-pdf/SAA-C02-PDF-Dumps\(654-669\).pdf](https://www.braindump2go.com/free-online-pdf/SAA-C02-PDF-Dumps(654-669).pdf) Free Resources from Braindump2go, We Devoted to Helping You 100% Pass All Exams!