SAA-C02^Q&As

AWS Certified Solutions Architect - Associate (SAA-C02)

Pass Amazon SAA-C02 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

https://www.leads4pass.com/saa-c02.html

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Amazon Official Exam Center

- **Instant Download** After Purchase
- **100% Money Back** Guarantee
- **365 Days** Free Update
- **800,000+** Satisfied Customers

QUESTION 1

A company is experiencing growth as demand for its product has increased The company\\'s existing purchasing application is slow when traffic spikes The application is a monolithic three tier application that uses synchronous transactions

and sometimes sees bottlenecks in the application tier A solutions architect needs to design a solution that can meet required application response times while accounting for traffic volume spikes.

Which solution will meet these requirements?

A. Vertically scale the application instance using a larger Amazon EC2 instance size.

B. Scale the application\\'s persistence layer horizontally by introducing Oracle RAC on AWS

C. Scale the web and application tiers horizontally using Auto Scaling groups and an Application Load Balancer

D. Decouple the application and data tiers using Amazon Simple Queue Service (Amazon SQS) with asynchronous AWS Lambda calls.

Correct Answer: C

QUESTION 2

A company has created a VPC with multiple private subnets in multiple Availability Zones (AZs) and one public subnet in one of the AZs. The public subnet is used to launch a NAT gateway. There are instance in the private subnet that use a

NAT gateway to connect to the internet. In case of an AZ failure, the company wants to ensure that the instance are not all experiencing internet connectivity issues and that there is a backup plan ready.

Which solution should a solutions architect recommend that is MOST highly available?

A. Create a new public subnet with a NAT gateway in the same AZ Distribute the traffic between the two NAT gateways

B. Create an Amazon EC2 NAT instance in a now public subnet Distribute the traffic between the NAT gateway and the NAT instance

C. Create public subnets In each AZ and launch a NAT gateway in each subnet Configure the traffic from the private subnets In each A2 to the respective NAT gateway

D. Create an Amazon EC2 NAT instance in the same public subnet Replace the NAT gateway with the NAT instance and associate the instance with an Auto Scaling group with an appropriate scaling policy.

Correct Answer: C

QUESTION 3

A product team is creating a new application that will store a large amount of data The data will be analyzed hourly and modified by multiple Amazon EC2 Linux instances The application team believes the amount of space needed will continue to grow for the next 6 months Which set of actions should a solutions architect take to support these needs\\'?

A. Store the data in an Amazon EBS volume Mount the EBS volume on the application instances

B. Store the data in an Amazon EFS file system Mount the file system on the application instances

C. Store the data in Amazon S3 Glacier Update the vault policy to allow access to the application instances

D. Store the data in Amazon S3 Standard-Infrequent Access (S3 Standard-IA) Update the bucket policy to allow access to the application instances

Correct Answer: B

Amazon Elastic File System Amazon Elastic File System (Amazon EFS) provides a simple, scalable, fully managed elastic NFS file system for use with AWS Cloud services and on-premises resources. It is built to scale on demand to petabytes without disrupting applications, growing and shrinking automatically as you add and remove files, eliminating the need to provision and manage capacity to accommodate growth. Amazon EFS is designed to provide massively parallel shared access to thousands of Amazon EC2 instances, enabling your applications to achieve high levels of aggregate throughput and IOPS with consistent low latencies.

Amazon EFS is well suited to support a broad spectrum of use cases from home directories to business-critical applications. Customers can use EFS to lift-and-shift existing enterprise applications to the AWS Cloud. Other use cases include: big data analytics, web serving and content management, application development and testing, media and entertainment workflows, database backups, and container storage. Amazon EFS is a regional service storing data within and across multiple Availability Zones (AZs) for high availability and durability. Amazon EC2 instances can access your file system across AZs, regions, and VPCs, while on-premises servers can access using AWS Direct Connect or AWS VPN. https://aws.amazon.com/efs/

QUESTION 4

A company is hosting its website by using Amazon EC2 instance behind an Elastic Load Balancer across multiple Availability Zones. The instance run in an EC2 Auto Scaling group. The website uses Amazon Elastic Block Store (Amazon EBS) volumes to store product manuals for users to download. The company updates the product content often, so new instance launched by the Auto Scaling group often have old data It can take up to 30 minutes for the new instances to receive all the updates. The updates also requires the EBS volumes to be resized during business hours. The company wants to ensure that the product manuals are always up to data on all that the architecture adjusts quickly to increased user demand. A solutions architect needs to meet these requirements without causing the company to update its application code or adjust its website. What should the solution architect do to accomplish this goal?

A. Store the product manuals in an EBS volume. Mount that volume to the EC2 instances.

B. Store the product manuals in an Amazon S3 bucket. Redirect the downloads to this bucket.

C. Store the product manual in an Amazon Elastic File System (Amazon EFS) volume Mount that volume to the EC2 instances.

D. Store the product manual in an Amazon S3 Standard-infrequent Access (S3 Standard-IA) bucket Redirect the downloads to this bucket.

Correct Answer: D

QUESTION 5

A company has an application that collects data from loT sensors on automobiles. The data is streamed and stored in Amazon S3 through Amazon Kinesis Date Firehose The data produces trillions of S3 objects each year. Each morning,

the company uses the data from the previous 30 days to retrain a suite of machine learning (ML) models.

Four times each year, the company uses the data from the previous 12 months to perform analysis and train other ML models The data must be available with minimal delay for up to 1 year. After 1 year, the data must be retained for archival purposes.

Which storage solution meets these requirements MOST cost-effectively?

A. Use the S3 Intelligent-Tiering storage class. Create an S3 Lifecycle policy to transition objects to S3 Glacier Deep Archive after 1 year

B. Use the S3 Intelligent-Tiering storage class. Configure S3 Intelligent-Tiering to automatically move objects to S3 Glacier Deep Archive after 1 year.

C. Use the S3 Standard-Infrequent Access (S3 Standard-IA) storage class. Create an S3 Lifecycle policy to transition objects to S3 Glacier Deep Archive after 1 year.

D. Use the S3 Standard storage class. Create an S3 Lifecycle policy to transition objects to S3 Standard-Infrequent Access (S3 Standard-IA) after 30 days, and then to S3 Glacier Deep Archive after 1 year.

Correct Answer: B

Reference: https://aws.amazon.com/blogs/aws/s3-intelligent-tiering-adds-archive-access-tiers/