SAA-C03 Q&As

AWS Certified Solutions Architect - Associate (SAA-C03)

Pass Amazon SAA-C03 Exam with 100% Guarantee

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

https://www.pass4itsure.com/saa-c03.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

An application runs on Amazon EC2 instances across multiple Availability Zones The instances run in an Amazon EC2 Auto Scaling group behind an Application Load Balancer The application performs best when the CPU utilization of the EC2 instances is at or near 40%.

What should a solutions architect do to maintain the desired performance across all instances in the group?

A. Use a simple scaling policy to dynamically scale the Auto Scaling group

B. Use a target tracking policy to dynamically scale the Auto Scaling group

C. Use an AWS Lambda function to update the desired Auto Scaling group capacity.

D. Use scheduled scaling actions to scale up and scale down the Auto Scaling group

Correct Answer: B

https://docs.aws.amazon.com/autoscaling/application/userguide/application-auto-scaling-target-tracking.html

QUESTION 2

A company wants to relocate its on-premises MySQL database to AWS. The database accepts regular imports from a client-facing application, which causes a high volume of write operations. The company is concerned that the amount of traffic might be causing performance issues within the application.

How should a solutions architect design the architecture on AWS?

A. Provision an Amazon RDS for MySQL DB instance with Provisioned IOPS SSD storage. Monitor write operation metrics by using Amazon CloudWatch. Adjust the provisioned IOPS if necessary.

B. Provision an Amazon RDS for MySQL DB instance with General Purpose SSD storage. Place an Amazon ElastiCache cluster in front of the DB instance. Configure the application to query ElastiCache instead.

C. Provision an Amazon DocumentDB (with MongoDB compatibility) instance with a memory optimized instance type. Monitor Amazon CloudWatch for performance-related issues. Change the instance class if necessary.

D. Provision an Amazon Elastic File System (Amazon EFS) file system in General Purpose performance mode. Monitor Amazon CloudWatch for IOPS bottlenecks. Change to Provisioned Throughput performance mode if necessary.

Correct Answer: A

QUESTION 3

A company has released a new version of its production application. The company\\'s workload uses Amazon EC2, AWS Lambda, AWS Fargate, and Amazon SageMaker. The company wants to cost optimize the workload now that usage is at a steady state. The company wants to cover the most services with the fewest savings plans. Which combination of savings plans will meet these requirements? (Choose two.)

A. Purchase an EC2 Instance Savings Plan for Amazon EC2 and SageMaker.

B. Purchase a Compute Savings Plan for Amazon EC2, Lambda, and SageMaker.

C. Purchase a SageMaker Savings Plan.

D. Purchase a Compute Savings Plan for Lambda, Fargate, and Amazon EC2.

E. Purchase an EC2 Instance Savings Plan for Amazon EC2 and Fargate.

Correct Answer: CD

https://aws.amazon.com/savingsplans/ml-pricing/ https://aws.amazon.com/savingsplans/compute-pricing/

QUESTION 4

A company hosts an online shopping application that stores all orders in an Amazon RDS for PostgreSQL Singfe-AZ DB instance. Management wants to eliminate single points of C^ilure and has asked a solutions architect to recommend an approach to minimize database downtime without requiring any changes to the application code.

Which solution meets these requirements?

A. Convert the existing database instance to a Multi-AZ deployment by modifying the database instance and specifying the Multi-AZ option.

B. Create a new RDS Multi-AZ deployment. Take a snapshot of the current RDS instance and restore the new Multi-AZ deployment with the snapshot.

C. Create a read-only replica of the PostgreSQL database in another Availability Zone. Use Amazon Route 53 weighted record sets to distribute requests across the databases.

D. Place the RDS for PostgreSQL database in an Amazon EC2 Auto Scaling group with a minimum group size of two. Use Amazon Route 53 weighted record sets to distribute requests across instances.

Correct Answer: A

https://aws.amazon.com/rds/features/multi-az/ To convert an existing Single-AZ DB Instance to a Multi-AZ deployment, use the "Modify" option corresponding to your DB Instance in the AWS Management Console.

QUESTION 5

A solutions architect is designing the cloud architecture for a new application being deployed on AWS. The process should run in parallel while adding and removing application nodes as needed based on the number of jobs to be processed. The processor application is stateless. The solutions architect must ensure that the application is loosely coupled and the job items are durably stored.

Which design should the solutions architect use?

A. Create an Amazon SNS topic to send the jobs that need to be processed Create an Amazon Machine Image (AMI) that consists of the processor application Create a launch configuration that uses the AMI Create an Auto Scaling group using the launch configuration Set the scaling policy for the Auto Scaling group to add and remove nodes based on CPU usage

B. Create an Amazon SQS queue to hold the jobs that need to be processed Create an Amazon Machine image (AMI) that consists of the processor application Create a launch configuration that uses the AM\\' Create an Auto Scaling

group using the launch configuration Set the scaling policy for the Auto Scaling group to add and remove nodes based on network usage

C. Create an Amazon SQS queue to hold the jobs that needs to be processed Create an Amazon Machine image (AMI) that consists of the processor application Create a launch template that uses the AMI Create an Auto Scaling group using the launch template Set the scaling policy for the Auto Scaling group to add and remove nodes based on the number of items in the SQS queue

D. Create an Amazon SNS topic to send the jobs that need to be processed Create an Amazon Machine Image (AMI) that consists of the processor application Create a launch template that uses the AMI Create an Auto Scaling group using the launch template Set the scaling policy for the Auto Scaling group to add and remove nodes based on the number of messages published to the SNS topic

Correct Answer: C

"Create an Amazon SQS queue to hold the jobs that needs to be processed. Create an Amazon EC2 Auto Scaling group for the compute application. Set the scaling policy for the Auto Scaling group to add and remove nodes based on the number of items in the SQS queue"

In this case we need to find a durable and loosely coupled solution for storing jobs. Amazon SQS is ideal for this use case and can be configured to use dynamic scaling based on the number of jobs waiting in the queue.To configure this scaling you can use the backlog per instance metric with the target value being the acceptable backlog per instance to maintain. You can calculate these numbers as follows: Backlog per instance: To calculate your backlog per instance, start with the ApproximateNumberOfMessages queue attribute to determine the length of the SQS queue

QUESTION 6

A company uses Amazon S3 to store its confidential audit documents. The S3 bucket uses bucket policies to restrict access to audit team IAM user credentials according to the principle of least privilege. Company managers are worried about accidental deletion of documents in the S3 bucket and want a more secure solution.

What should a solutions architect do to secure the audit documents?

A. Enable the versioning and MFA Delete features on the S3 bucket.

B. Enable multi-factor authentication (MFA) on the IAM user credentials for each audit team IAM user account.

C. Add an S3 Lifecycle policy to the audit team\\'s IAM user accounts to deny the s3:DeleteObject action during audit dates.

D. Use AWS Key Management Service (AWS KMS) to encrypt the S3 bucket and restrict audit team IAM user accounts from accessing the KMS key.

Correct Answer: A

QUESTION 7

A rapidly growing ecommerce company is running its workloads in a single AWS Region. A solutions architect must create a disaster recovery (DR) strategy that includes a different AWS Region The company wants its database to be up to date in the DR Region with the least possible latency The remaining infrastructure in the DR Region needs to run at reduced capacity and must be able to scale up it necessary

Which solution will meet these requirements with the LOWEST recovery time objective (RTO)?

A. Use an Amazon Aurora global database with a pilot light deployment

B. Use an Amazon Aurora global database with a warm standby deployment

C. Use an Amazon RDS Multi-AZ DB instance with a pilot light deployment

D. Use an Amazon RDS Multi-AZ DB instance with a warm standby deployment

Correct Answer: B

Note: The difference between pilot light and warm standby can sometimes be difficult to understand. Both include an environment in your DR Region with copies of your primary Region assets. The distinction is that pilot light cannot process requests without additional action taken first, whereas warm standby can handle traffic (at reduced capacity levels) immediately. The pilot light approach requires you to "turn on" servers, possibly deploy additional (non-core) infrastructure, and scale up, whereas warm standby only requires you to scale up (everything is already deployed and running). Use your RTO and RPO needs to help you choose between these approaches.

https://docs.aws.amazon.com/whitepapers/latest/disaster-recovery-workloads-on-aws/disaster-recovery-options-in-the-cloud.html

QUESTION 8

A company is subscribed to the AWS Business Support plan. Compliance rules require the company to check on AWS infrastructure health before deployments can proceed. The company needs a programmatic and automated way to check on infrastructure health at the beginning of new deployments.

Which solution will meet these requirements?

A. Use the AWS Trusted Advisor API at the start of each deployment. Pause all new deployments if the API returns any issues.

B. Use the AWS Health API at the start of each deployment. Pause all new deployments if the API returns any issues.

C. Query the AWS Support API at the start of each deployment. Pause all new deployments if the API returns any open issues.

D. Send an API call to each workload ahead of deployment. Pause the deployments if the API call fails.

Correct Answer: B

The AWS Health API provides programmatic access to the AWS Health information that is presented in the AWS Personal Health Dashboard. You can use the API operations to get information about AWS Health events that affect your AWS services and resources. You can also use the API to enable or disable health-based insights for your organization. You can use the AWS Health API at the start of each deployment to check on AWS infrastructure health and pause all new deployments if the API returns any issues.

References: https://docs.aws.amazon.com/health/latest/APIReference/Welcome.html

QUESTION 9

A company is building a new furniture inventory application. The company has deployed the application on a fleet ofAmazon EC2 instances across multiple Availability Zones. The EC2 instances run behind an Application Load Balancer (ALB) in their VPC. A solutions architect has observed that incoming traffic seems to favor one EC2 instance,

resulting in latency for some requests. What should the solutions architect do to resolve this issue?

A. Disable session affinity (sticky sessions) on the ALB

B. Replace the ALB with a Network Load Balancer

C. Increase the number of EC2 instances in each Availability Zone

D. Adjust the frequency of the health checks on the ALB\\'s target group

Correct Answer: A

QUESTION 10

A serverless application uses Amazon API Gateway, AWS Lambda, and Amazon DynamoDB. The Lambda function needs permissions to read and write to the DynamoDB table.

Which solution will give the Lambda function access to the DynamoDB table MOST securely?

A. Create an IAM user with programmatic access to the Lambda function. Attach a policy to the user that allows read and write access to the DynamoDB table. Store the access_key_id and secret_access_key parameters as part of the Lambda environment variables. Ensure that other AWS users do not have read and write access to the Lambda function configuration.

B. Create an IAM role that includes Lambda as a trusted service. Attach a policy to the role that allows read and write access to the DynamoDB table. Update the configuration of the Lambda function to use the new role as the execution role.

C. Create an IAM user with programmatic access to the Lambda function. Attach a policy to the user that allows read and write access to the DynamoDB table. Store the access_key_id and secret_access_key parameters in AWS Systems Manager Parameter Store as secure string parameters. Update the Lambda function code to retrieve the secure string parameters before connecting to the DynamoDB table.

D. Create an IAM role that includes DynamoDB as a trusted service. Attach a policy to the role that allows read and write access from the Lambda function. Update the code of the Lambda function to attach to the new role as an execution role.

Correct Answer: B

Role key word and trusted service lambda

QUESTION 11

A company has thousands of edge devices that collectively generate 1 TB of status alerts each day. Each alert is approximately 2 KB in size. A solutions architect needs to implement a solution to ingest and store the alerts for future analysis.

The company wants a highly available solution. However, the company needs to minimize costs and does not want to manage additional infrastructure. Ad ditionally, the company wants to keep 14 days of data available for immediate analysis and archive any data older than 14 days.

What is the MOST operationally efficient solution that meets these requirements?

A. Create an Amazon Kinesis Data Firehose delivery stream to ingest the alerts Configure the Kinesis Data Firehose stream to deliver the alerts to an Amazon S3 bucket Set up an S3 Lifecycle configuration to transition data to Amazon S3 Glacier after 14 days

B. Launch Amazon EC2 instances across two Availability Zones and place them behind an Elastic Load Balancer to ingest the alerts Create a script on the EC2 instances that will store tne alerts m an Amazon S3 bucket Set up an S3 Lifecycle configuration to transition data to Amazon S3 Glacier after 14 days

C. Create an Amazon Kinesis Data Firehose delivery stream to ingest the alerts Configure the Kinesis Data Firehose stream to deliver the alerts to an Amazon Elasticsearch Service (Amazon ES) duster Set up the Amazon ES cluster to take manual snapshots every day and delete data from the duster that is older than 14 days

D. Create an Amazon Simple Queue Service (Amazon SQS i standard queue to ingest the alerts and set the message retention period to 14 days Configure consumers to poll the SQS queue check the age of the message and analyze the message data as needed If the message is 14 days old the consumer should copy the message to an Amazon S3 bucket and delete the message from the SQS queue

Correct Answer: A

https://aws.amazon.com/kinesis/datafirehose/features/?nc=snandloc=2#:~:text=into%20Amazon%20S3%2C%20Amazon%20Red shift%2C%20Amazon%20OpenSearch%20Service%2C%20Kinesis,Delivery%20streams

QUESTION 12

A solutions architect is developing a VPC architecture that includes multiple subnets. The architecture will host applications that use Amazon EC2 instances and Amazon RDS DB instances. The architecture consists of six subnets in two

Availability Zones. Each Availability Zone includes a public subnet, a private subnet, and a dedicated subnet for databases. Only EC2 instances that run in the private subnets can have access to the RDS databases.

Which solution will meet these requirements?

A. Create a new route table that excludes the route to the public subnets\\' CIDR blocks. Associate the route table with the database subnets.

B. Create a security group that denies inbound traffic from the security group that is assigned to instances in the public subnets. Attach the security group to the DB instances.

C. Create a security group that allows inbound traffic from the security group that is assigned to instances in the private subnets. Attach the security group to the DB instances.

D. Create a new peering connection between the public subnets and the private subnets. Create a different peering connection between the private subnets and the database subnets.

Correct Answer: C

The solution that meets the requirements described in the question is option C: Create a security group that allows inbound traffic from the security group that is assigned to instances in the private subnets. Attach the security group to the DB

instances.

In this solution, the security group applied to the DB instances allows inbound traffic from the security group assigned to instances in the private subnets. This ensures that only EC2 instances running in the private subnets can have access

to

the RDS databases.

QUESTION 13

A solutions architect needs to design a system to store client case files. The files are core company assets and are important. The number of files will grow over time.

The files must be simultaneously accessible from multiple application servers that run on Amazon EC2 instances. The solution must have built-in redundancy.

Which solution meets these requirements?

A. Amazon Elastic File System (Amazon EFS)

B. Amazon Elastic Block Store (Amazon EBS)

C. Amazon S3 Glacier Deep Archive

D. AWS Backup

Correct Answer: A

Amazon EFS provides a simple, scalable, fully managed file system that can be simultaneously accessed from multiple EC2 instances and provides built-in redundancy. It is optimized for multiple EC2 instances to access the same files, and it is designed to be highly available, durable, and secure. It can scale up to petabytes of data and can handle thousands of concurrent connections, and is a cost-effective solution for storing and accessing large amounts of data.

QUESTION 14

A company is creating an application that runs on containers in a VPC. The application stores and accesses data m an Amazon S3 bucket. During the development phase the application will store and access 1 TB of data in Amazon S3 each day. The company wants to minimize costs and wants to prevent traffic from traversing the internet whenever possible.

Which solution will meet these requirements?

A. Enable S3 Intelligent-Tiering for the S3 bucket.

B. Enable S3 Transfer Acceleration for the S3 bucket

C. Create a gateway VPC endpoint for Amazon S3 Associate this endpoint with all route tables in the VPC.

D. Create an interlace endpoint for Amazon S3 in the VPC Associate this endpoint with all route tables in the VPC.

Correct Answer: C

QUESTION 15

A marketing company receives a large amount of new clickstream data in Amazon S3 from a marketing campaign. The

company needs to analyze the clickstream data in Amazon S3 quickly. Then the company needs to determine whether to process the data further in the data pipeline.

Which solution will meet these requirements with the LEAST operational overhead?

A. Create external tables in a Spark catalog. Configure jobs in AWS Glue to query the data.

B. Configure an AWS Glue crawler to crawl the data. Configure Amazon Athena to query the data.

C. Create external tables in a Hive metastore. Configure Spark jobs in Amazon EMR to query the data.

D. Configure an AWS Glue crawler to crawl the data. Configure Amazon Kinesis Data Analytics to use SQL to query the data.

Correct Answer: B

AWS Glue with Athena (Option B): AWS Glue is a fully managed extract, transform, and load (ETL) service, and Athena is a serverless query service that allows you to analyze data directly in Amazon S3 using SQL queries. By configuring an AWS Glue crawler to crawl the data, you can create a schema for the data, and then use Athena to query the data directly without the need to load it into a separate database. This minimizes operational overhead.