Amazon Amazon AWS Certified Solutions Architect Associate-SAA-C01 PDF

Amazon Amazon AWS Certified Solutions Architect Associate-SAA-C01 PDF Questions Available Here at:

https://www.certification-questions.com//amazon-dumps/aws-certified-solutions-architect-associate-saa-c01.html

Enrolling now you will get access to 268 questions in a unique set of Amazon AWS Certified Solutions Architect Associate-SAA-C01

Question 1

A company is developing several critical long-running applications hosted on Docker.
How should a Solutions Architect design a solution to meet the scalability and orchestration requirements on AWS?

Options:

A. Use Amazon ECS and Service Auto Scaling.

B. Use Spot Instances for orchestration and for scaling containers on existing Amazon EC2 instances.

C. Use AWS OpsWorks to launch containers in new Amazon EC2 instances.

D. Use Auto Scaling groups to launch containers on existing Amazon EC2 instances.

Answer: A

Explanation:
https://docs.aws.amazon.com/AmazonECS/latest/developerguide/service-auto-scaling.html

Question 2

An application is running on an Amazon EC2 instance in a private subnet. The application needs to read and write data onto Amazon Kinesis Data Streams, and corporate policy requires that this traffic should not go to the internet.
How can these requirements be met?

Options:

A. Configure a NAT gateway in a public subnet and route all traffic to Amazon Kinesis through the NAT gateway.

B. Configure a gateway VPC endpoint for Kinesis and route all traffic to Kinesis through the gateway
VPC endpoint.

C. Configure an interface VPC endpoint for Kinesis and route all traffic to Kinesis through the gateway
VPC endpoint.

D. Configure an AWS Direct Connect private virtual interface for Kinesis and route all traffic to Kinesis
through the virtual interface.

Answer: C

Explanation:
https://docs.aws.amazon.com/vpc/latest/userguide/vpc-endpoints.html An interface endpoint is an elastic network interface with a private IP address from the IP address range of your subnet that serves as an entry point for traffic destined to a supported service,

Question 3

A Solutions Architect is designing the storage layer for a production relational database. The database will run on Amazon EC2. The database is accessed by an application that performs intensive reads and writes, so the database requires the LOWEST random I/O latency.
Which data storage method fulfills the above requirements?

Options:

A. Store data in a filesystem backed by Amazon Elastic File System (EFS).

B. Store data in Amazon S3 and use a third-party solution to expose Amazon S3 as a filesystem to the
database server.

C. Store data in Amazon Dynamo DB and emulate relational database semantics.

D. Stripe data across multiple Amazon EBS volumes using RAID 0.

Answer: D

Explanation:
When we perform the RAID 0 Striping of multiple volumes, IOPS are distributed among the volumes of a stripe. If you add another volume to RAID 0, you get the straight addition of IOPS throughput of

that volume and additional volume size. Reference:
https://cloudacademy.com/blog/amazon-aws-raid-0-configuration-on-ebs-volumes/

Question 4

A news organization plans to migrate their 20 TB video archive to AWS. The files are rarely accessed, but when they are, a request is made in advance and a 3 to 5-hour retrieval time frame is acceptable. However, when there is a breaking news story, the editors require access to archived footage within minutes.
Which storage solution meets the needs of this organization while providing the LOWEST cost of storage?

Options:

A. Store the archive in Amazon S3 Reduced Redundancy Storage.

B. Store the archive in Amazon Glacier and use standard retrieval for all content.

C. Store the archive in Amazon Glacier and pay the additional charge for expedited retrieval when needed.

D. Store the archive in Amazon S3 with a lifecycle policy to move this to S3 Infrequent Access after 30
days.

Answer: C

Explanation:
Expedited - Expedited retrievals allow you to quickly access your data when occasional urgent requests for a subset of archives are required. For all but the largest archives (250 MB+), data accessed using Expedited retrievals are typically made available within 1-5 minutes. Provisioned Capacity ensures that retrieval capacity for Expedited retrievals is available when you need it. For more information, see Provisioned Capacity.
Standard - Standard retrievals allow you to access any of your archives within several hours. Standard retrievals typically complete within 3-5 hours. This is the default option for retrieval requests that do not specify the retrieval option.
Bulk - Bulk retrievals are Glacier's lowest-cost retrieval option, which you can use to retrieve large amounts, even petabytes, of data inexpensively in a day. Bulk retrievals typically complete within 5-12 hours.

Question 5

A Solutions Architect must select the most cost-efficient architecture for a service that responds to web requests. These web requests are small and query a DynamoDB table.
The request rate ranges from zero to several hundred each second, without any predictable patterns.

What is the MOST cost-efficient architecture for this service?

Options:

A. Network Load Balancer/Amazon EC2

B. Application Load Balancer/Amazon ECS

C. API Gateway/AWS Lambda

D. AWS Elastic Beanstalk/AWS Lambda

Answer: C

Explanation:
Initial Concurrency Burst Limits
3000 - US West (Oregon), US East (N. Virginia), EU (Ireland).
1000 - Asia Pacific (Tokyo), EU (Frankfurt).
500 - Other regions.

Question 6

An application uses an Amazon SQS queue as a transport mechanism to deliver data to a group of EC2 instances for processing. The application owner wants to add a mechanism to archive the incoming data without modifying application code on the EC2 instances.
How can this application be re-architected to archive the data without modifying the processing instances?

Options:

A. Trigger a Lambda function by using Amazon CloudWatch Events to retrieve messages from the SQS queue and archive to Amazon S3.

B. Use an Amazon SNS topic to fan out the data to the SQS queue in addition to a Lambda function
that records the data to an S3 bucket.

C. Set up an Amazon Kinesis Data Stream so that multiple instances can receive data.
Add a separate EC2 instance that is configured to archive all data it receives.

D. Write the data to an S3 bucket, and use an SQS queue for S3 event notifications to tell the instances where to retrieve the data.

Answer: B

Explanation:
https://docs.aws.amazon.com/sns/latest/dg/sns-common-scenarios.html

Question 7

The application tier for a workload runs on EC2 instances that are unevenly distributed across two Availability Zones. The instances are behind a Network Load Balancer and are accessed through layer 4 TCP connections.
The instances in the lesser populated Availability Zone are failing as the result of high CPU utilization.
Which configuration change can help mitigate the issue?

Options:

A. Modify the Network Load Balancer to enable sticky sessions

B. Modify the Network Load Balancer to enable cross-zone load balancing.

C. Switch to using an Application Load Balancer and enable sticky sessions.

D. Switch to using an Application Load Balancer and enable cross-zone load balancing.

Answer: D

Question 8

A company wants to migrate a three-tier web application to AWS. The company wants to control the placement of the instances and have visibility into underlying sockets and cores for licensing purposes.
Which compute model should a Solutions Architect choose to accomplish this task?

Options:

A. EC2 Reserved Instances

B. EC2 Spot Instances

C. EC2 Dedicated Hosts

D. EC2 Placement Groups

Answer: C

Question 9

A Solutions Architect is trying to bring a data warehouse workload to an Amazon EC2 instance. The data will reside in Amazon EBS volumes and full table scans will be executed frequently.
What type of Amazon EBS volume would be most suitable in this scenario?

Options:

A. Throughput Optimized HDD (st1)

B. Provisioned IOPS SSD (io1)

C. General Purpose SSD (gp2)

D. Cold HDD (sc1)

Answer: A

Explanation:
https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSVolumeTypes.html

Question 10

A company has an application that uses Amazon CloudFront for content that is hosted on an Amazon S3 bucket. After an unexpected refresh, the users are still seeing old content.

Which step should the Solutions Architect take to ensure that new content is displayed?

Options:

A. Perform a cache refresh on the CloudFront distribution that is serving the content.

B. Perform an invalidation on the CloudFront distribution that is serving the content.

C. Create a new cache behavior path with the updated content.

D. Change the TTL value for removing the old objects

Answer: B

Explanation:
https://aws.amazon.com/premiumsupport/knowledge-center/cloudfront-serving-outdated-content-s3/