

AWS Solution Architect Associate Certification Example Questions

1) A company runs a public-facing three-tier web application in a VPC across multiple Availability Zones. Amazon EC2 instances for the application tier running in private subnets need to download software patches from the internet. However, the EC2 instances cannot be directly accessible from the internet. Which actions should be taken to allow the EC2 instances to download the needed patches? (Select two choices)

A) Configure a NAT gateway in a public subnet.

B) Define a custom route table with a route to the NAT gateway for internet traffic and associate it with the private subnets for the application tier.

C) Assign Elastic IP addresses to the EC2 instances.

D) Define a custom route table with a route to the internet gateway for internet traffic and associate it with the private subnets for the application tier.

E) Configure a NAT instance in a private subnet.

2) A solutions architect wants to design a solution to save costs for Amazon EC2 instances that do not need to run during a 2-week company shutdown. The applications running on the EC2 instances store data in instance memory that must be present when the instances resume operation. Which approach should the solutions architect recommend to shut down and resume the EC2 instances?

A) Modify the application to store the data on instance store volumes. Reattach the volumes while restarting them.

B) Snapshot the EC2 instances before stopping them. Restore the snapshot after restarting the instances.

C) Run the applications on EC2 instances enabled for hibernation. Hibernate the instances before the 2-week company shutdown.

D) Note the Availability Zone for each EC2 instance before stopping it. Restart the instances in the same Availability Zones after the 2-week company shutdown.

3) A company plans to run a monitoring application on an Amazon EC2 instance in a VPC. Connections are made to the EC2 instance using the instance's private IPv4 address. A solutions architect needs to design a solution that will allow traffic to be quickly directed to a standby EC2 instance if the application fails and becomes unreachable. Which approach will meet these requirements? (Select two choices)

A) Deploy an Application Load Balancer configured with a listener for the private IP address and register the primary EC2 instance with the load balancer. Upon failure, de-register the instance and register the standby EC2 instance.

B) Configure a custom DHCP option set. Configure DHCP to assign the same private IP address to the standby EC2 instance when the primary EC2 instance fails.

C) Attach a secondary elastic network interface to the EC2 instance configured with the private IP address. Move the network interface to the standby EC2 instance if the primary EC2 instance becomes unreachable.

D) Associate an Elastic IP address with the network interface of the primary EC2 instance. Disassociate the Elastic IP from the primary instance upon failure and associate it with a standby EC2 instance.

4) An analytics company is planning to offer a web analytics service to its users. The service will require that the users' webpages include a JavaScript script that makes authenticated GET requests to the company's Amazon S3 bucket. What must a solutions architect do to ensure that the script will successfully execute?

A) Enable cross-origin resource sharing (CORS) on the S3 bucket.

B) Enable S3 Versioning on the S3 bucket.

C) Provide the users with a signed URL for the script.

D) Configure an S3 bucket policy to allow public execute privileges.

5) A company's security team requires that all data stored in the cloud be encrypted at rest at all times using encryption keys stored on premises. Which encryption options meet these requirements? (Select two choices)

A) Use server-side encryption with Amazon S3 managed encryption keys (SSE-S3).

B) Use server-side encryption with AWS KMS managed encryption keys (SSE-KMS).

C) Use server-side encryption with customer-provided encryption keys (SSE-C).

D) Use client-side encryption to provide at-rest encryption.

E) Use an AWS Lambda function invoked by Amazon S3 events to encrypt the data using the customer's keys.

6) A company uses Amazon EC2 Reserved Instances to run its data processing workload. The nightly job typically takes 7 hours to run and must finish within a 10-hour time window. The company anticipates temporary increases in demand at the end of each month that will cause the job to run over the time limit with the capacity of the current resources. Once started, the processing job cannot be interrupted before completion. The company wants to implement a solution that would provide increased resource capacity as cost-effectively as possible. What should a solutions architect do to accomplish this?

A) Deploy On-Demand Instances during periods of high demand.

B) Create a second EC2 reservation for additional instances.

C) Deploy Spot Instances during periods of high demand.

D) Increase the EC2 instance size in the EC2 reservation to support the increased workload.

7) A company runs an online voting system for a weekly live television program. During broadcasts, users submit hundreds of thousands of votes within minutes to a front-end fleet of Amazon EC2 instances that run in an Auto Scaling group. The EC2 instances write the votes to an Amazon RDS database. However, the database is unable to keep up with the requests that come from the EC2 instances. A solutions architect must design a solution that processes the votes in the most efficient manner and without downtime.

Which solution meets these requirements?

A) Migrate the front-end application to AWS Lambda. Use Amazon API Gateway to route user requests to the Lambda functions.

B) Scale the database horizontally by converting it to a Multi-AZ deployment. Configure the front-end application to write to both the primary and secondary DB instances.

C) Configure the front-end application to send votes to an Amazon Simple Queue Service (Amazon SQS) queue. Provision worker instances to read the SQS queue and write the vote information to the database.

D) Use Amazon EventBridge (Amazon CloudWatch Events) to create a scheduled event to re-provision the database with larger, memory-optimized instances during voting periods. When voting ends, re-provision the database to use smaller instances.

8) A company has a two-tier application architecture that runs in public and private subnets. Amazon EC2 instances running the web application are in the public subnet and an EC2 instance for the database runs on the private subnet. The web application instances and the database are running in a single Availability Zone (AZ).

Which combination of steps should a solutions architect take to provide high availability for this architecture? (Select two choices)

A) Create new public and private subnets in the same AZ.

B) Create an Amazon EC2 Auto Scaling group and Application Load Balancer spanning multiple AZs for the web application instances.

C) Add the existing web application instances to an Auto Scaling group behind an Application Load Balancer.

D) Create new public and private subnets in a new AZ. Create a database using an EC2 instance in the public subnet in the new AZ. Migrate the old database contents to the new database.

E) Create new public and private subnets in the same VPC, each in a new AZ. Create an Amazon RDS Multi-AZ DB instance in the private subnets. Migrate the old database contents to the new DB instance.

9) A website runs a custom web application that receives a burst of traffic each day at noon. The users upload new pictures and content daily, but have been complaining of timeouts. The architecture uses Amazon EC2 Auto Scaling groups, and the application consistently takes 1 minute to initiate upon boot up before responding to user requests.

How should a solutions architect redesign the architecture to better respond to changing traffic?

A) Configure a Network Load Balancer with a slow start configuration.

B) Configure Amazon ElastiCache for Redis to offload direct requests from the EC2 instances.

C) Configure an Auto Scaling step scaling policy with an EC2 instance warmup condition.

D) Configure Amazon CloudFront to use an Application Load Balancer as the origin.

10) An application running on AWS uses an Amazon Aurora Multi-AZ DB cluster deployment for its database. When evaluating performance metrics, a solutions architect discovered that the database reads are causing high I/O and adding latency to the write requests against the database. What should the solutions architect do to separate the read requests from the write requests?

A) Enable read-through caching on the Aurora database.

B) Update the application to read from the Multi-AZ standby instance.

C) Create an Aurora replica and modify the application to use the appropriate endpoints.

D) Create a second Aurora database and link it to the primary database as a read replica.

Answers

- 1) A-B
- 2) C
- 3) C-A
- 4) A
- 5) C-D
- 6) A
- 7) C
- 8) B-E
- 9) C
- 10) C