

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.)

- 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?

- 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.)

- 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.)

- 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 – A [NAT gateway](#) forwards traffic from the EC2 instances in the private subnet to the internet or other AWS services, and then sends the response back to the instances. After a NAT gateway is created, the route tables for private subnets must be updated to point internet traffic to the NAT gateway.
- 2) C – [Hibernating](#) EC2 instances save the contents of instance memory to an Amazon Elastic Block Store (Amazon EBS) root volume. When the instances restart, the instance memory contents are reloaded.
- 3) C – A [secondary elastic network interface](#) can be added to an EC2 instance. While primary network interfaces cannot be detached from an instance, secondary network interfaces can be detached and attached to a different EC2 instance.
- 4) A – Web browsers will block running a script that originates from a server with a domain name that is different from the webpage. [Amazon S3 can be configured with CORS](#) to send HTTP headers that allow the script to run.
- 5) C, D – [Server-side encryption with customer-provided keys \(SSE-C\)](#) enables Amazon S3 to encrypt objects on the server side using an encryption key provided in the PUT request. The same key must be provided in the GET requests for Amazon S3 to decrypt the object. Customers also have the option to encrypt data on the client side before uploading it to Amazon S3, and then they can decrypt the data after downloading it. AWS software development kits (SDKs) provide an S3 encryption client that streamlines the process.
- 6) A – While [Spot Instances](#) would be the least costly option, they are not suitable for jobs that cannot be interrupted or must complete within a certain time period. [On-Demand Instances](#) would be billed for the number of seconds they are running.
- 7) C – [Decouple](#) the ingestion of votes from the database to allow the voting system to continue processing votes without waiting for the database writes. Add dedicated workers to read from the [SQS queue](#) to allow votes to be entered into the database at a controllable rate. The votes will be added to the database as fast as the database can process them, but no votes will be lost.
- 8) B, E – Create new subnets in a new Availability Zone (AZ) to provide a redundant network. Create an [Auto Scaling group with instances in two AZs behind the load balancer](#) to ensure high availability of the web application and redistribution of web traffic between the two public AZs. Create an RDS DB instance in the two private subnets to make the [database tier highly available](#) too.
- 9) C – The current configuration puts new EC2 instances into service before they are able to respond to transactions. This could also cause the instances to overscale. With a [step scaling policy](#), you can specify the number of seconds that it takes for a newly launched instance to [warm up](#). Until its specified warm-up time has expired, an EC2 instance is not counted toward the aggregated metrics of the Auto Scaling group. While scaling out, the Auto Scaling logic does not consider EC2 instances that are warming up as part of the current capacity of the Auto Scaling group. Therefore, multiple alarm breaches that fall in the range of the same step adjustment result in a single scaling activity. This ensures that you do not add more instances than you need.
- 10) C – [Aurora Replicas](#) provide a way to offload read traffic. Aurora Replicas share the same [underlying storage](#) as the main database, so lag time is generally very low. Aurora Replicas have their own endpoints, so the application will need to be configured to direct read traffic to the new endpoints.