aws training and certification

AWS Certified Solutions Architect – Associate
Official Practice Question Set

QUESTIONS

#1

A solutions architect is designing a solution to run a containerized web application by using Amazon Elastic Container Service (Amazon ECS). The solutions architect wants to minimize cost by running multiple copies of a task on each container instance. The number of task copies must scale as the load increases and decreases.

Which routing solution distributes the load to the multiple tasks?

A. Configure an Application Load Balancer to distribute the requests by using path-based routing.
B. Configure an Application Load Balancer to distribute the requests by using dynamic host port mapping.
C. Configure an Amazon Route 53 alias record set to distribute the requests with a failover routing policy.
D. Configure an Amazon Route 53 alias record set to distribute the requests with a weighted routing policy.

AWS Certified Solutions Architect – Associate
Official Practice Question Set

#2

A company has strict data protection requirements. A solutions architect must configure security for a VPC to ensure that backend Amazon RDS DB instances cannot be accessed from the internet. The solutions architect must ensure that the DB instances are accessible from the application tier over a specified port only.

Which actions should the solutions architect take to meet these requirements? (Select TWO.)

A. Specify a DB subnet group that contains only private subnets for the DB instances.
B. Attach an elastic network interface with a private IPv4 address to each DB instance.
C. Configure AWS Shield with the VPC. Update the route tables for the subnets that the DB instances use.
D. Configure an AWS Direct Connect connection on the database port between the application tier and the backend.
E. Add an inbound rule to the database security group that allows requests from the security group of the application tier over the database port. Remove other inbound rules.

#3

An application runs on two Amazon EC2 instances behind a Network Load Balancer. The EC2 instances are in a single Availability Zone.

What should a solutions architect do to make this architecture more highly available?

A. Create a new VPC with two new EC2 instances in the same Availability Zone as the original EC2 instances. Create a VPC peering connection between the two VPCs
B. Replace the Network Load Balancer with an Application Load Balancer that is configured with the EC2 instances in an Auto Scaling group.
C. Configure Amazon Route 53 to perform health checks on the EC2 instances behind the Network Load Balancer. Add a failover routing policy.
D. Place the EC2 instances in an Auto Scaling group that extends across multiple Availability Zones. Designate the Auto Scaling group as the target of the Network Load Balancer.

AWS Certified Solutions Architect – Associate
Official Practice Question Set

#4

A reporting application runs on Amazon EC2 instances behind an Application Load Balancer. The instances run in an Amazon EC2 Auto Scaling group across multiple Availability Zones. For complex reports, the application can take up to 15 minutes to respond to a request. A solutions architect is concerned that users will receive HTTP 5xx errors if a report request is in process during a scale-in event.

What should the solutions architect do to ensure that user requests will be completed before instances are terminated?

A. Enable sticky sessions (session affinity) for the target group of the instances.
B. Increase the instance size in the Application Load Balancer target group.
C. Increase the cooldown period for the Auto Scaling group to a greater amount of time than the time required for the longest running responses.
D. Increase the deregistration delay timeout for the target group of the instances to greater than 900 seconds.

#5

A solutions architect is planning a company's migration to the AWS Cloud. A key component of the company's environment is an application server that sends email notifications to customers. As part of this migration, the solutions architect must use only managed AWS services.

Which solution meets these requirements?

A. Configure Amazon Simple Queue Service (Amazon SQS) with a standard queue.
B. Deploy an Amazon EC2 instance to host the application server. Send email notifications.
C. Configure Amazon Simple Queue Service (Amazon SQS) with a FIFO queue.
D. Configure Amazon Simple Notification Service (Amazon SNS) with the protocol of email.

#6

A company used Amazon EC2 Spot Instances for a demonstration that is now complete. A solutions architect must remove the Spot Instances to stop them from incurring cost.

What should the solutions architect do to meet this requirement?

A. Cancel the Spot request only.
B. Terminate the Spot Instances only.
C. Cancel the Spot request. Terminate the Spot Instances.
D. Terminate the Spot Instances. Cancel the Spot request.

#7

A company needs to look up configuration details about how a Linux-based Amazon EC2 instance was launched.

Which command should a solutions architect run on the EC2 instance to gather the system metadata?

A. curl http://169.254.169.254/latest/meta-data
B. curl http://localhost/latest/meta-data
C. curl http://254.169.254.169/latest/meta-data
D. curl http://192.168.0.1/latest/meta-data

#8

A company has an on-premises application that exports log files about users of a website. These log files range from 20 GB to 30 GB in size. A solutions architect has created an Amazon S3 bucket to store these files. The files will be uploaded directly from the application. The network connection experiences intermittent failures, and the upload sometimes fails.

A solutions architect must design a solution that resolves this problem. The solution must minimize operational overhead.

Which solution will meet these requirements?

A. Enable S3 Transfer Acceleration.
B. Copy the files to an Amazon EC2 instance in the closest AWS Region. Use S3 Lifecycle policies to copy the log files to Amazon S3.
C. Use multipart upload to Amazon S3.
D. Upload the files to two AWS Regions simultaneously. Enable two-way Cross-Region Replication between the two Regions.

#9

A company is deploying a new database on a new Amazon EC2 instance. The workload of this database requires a single Amazon Elastic Block Store (Amazon EBS) volume that can support up to 20,000 IOPS.

Which type of EBS volume meets this requirement?

A. Throughput Optimized HDD
B. Provisioned IOPS SSD
C. General Purpose SSD
D. Cold HDD

AWS Certified Solutions Architect – Associate
Official Practice Question Set

#10

The usage of a company's image-processing application is increasing suddenly with no set pattern. The application's processing time grows linearly with the size of the image. The processing can take up to 20 minutes for large image files.

The architecture consists of a web tier, an Amazon Simple Queue Service (Amazon SQS) standard queue, and message consumers that process the images on Amazon EC2 instances. When a high volume of requests occurs, the message backlog in Amazon SQS increases. Users are reporting the delays in processing. A solutions architect must improve the performance of the application in compliance with cloud best practices.

Which solution will meet these requirements?

A. Purchase enough Dedicated Instances to meet the peak demand. Deploy the instances for the consumers.
B. Convert the existing SQS standard queue to an SQS FIFO queue. Increase the visibility timeout.
C. Configure a scalable AWS Lambda function as the consumer of the SQS messages.
D. Create a message consumer that is an Auto Scaling group of instances. Configure the Auto Scaling group to scale based upon the ApproximateNumberOfMessages Amazon CloudWatch metric.

#11

A company is developing a chat application that will be deployed on AWS. The application stores the messages by using a key-value data model. Groups of users typically read the messages multiple times. A solutions architect must select a database solution that will scale for a high rate of reads and will deliver messages with microsecond latency.

Which database solution will meet these requirements?

A. Amazon Aurora with Aurora Replicas
B. Amazon DynamoDB with DynamoDB Accelerator (DAX)
C. Amazon Aurora with Amazon ElastiCache for Memcached
D. Amazon Neptune with Amazon ElastiCache for Memcached

AWS Certified Solutions Architect – Associate
Official Practice Question Set

#12

A company uses one AWS account to run production workloads. The company has a separate AWS account for its security team. During periodic audits, the security team needs to view specific account settings and resource configurations in the AWS account that runs production workloads. A solutions architect must provide the required access to the security team by designing a solution that follows AWS security best practices.

Which solution will meet these requirements?

A. Create an IAM user for each security team member in the production account. Attach a permissions policy that provides the permissions required by the security team to each user.
B. Create an IAM role in the production account. Attach a permissions policy that provides the permissions required by the security team. Add the security team account to the trust policy.
C. Create a new IAM user in the production account. Assign administrative privileges to the user. Allow the security team to use this account to log in to the systems that need to be accessed.
D. Create an IAM user for each security team member in the production account. Attach a permissions policy that provides the permissions required by the security team to a new IAM group. Assign the security team members to the group.

AWS Certified Solutions Architect – Associate
Official Practice Question Set

#13

A company is deploying a new application that will consist of an application layer and an online transaction processing (OLTP) relational database. The application must be available at all times. However, the application will have periods of inactivity. The company wants to pay the minimum for compute costs during these idle periods.

Which solution meets these requirements MOST cost-effectively?

A. Run the application in containers with Amazon Elastic Container Service (Amazon ECS) on AWS Fargate. Use Amazon Aurora Serverless for the database.
B. Run the application on Amazon EC2 instances by using a burstable instance type. Use Amazon Redshift for the database.
C. Deploy the application and a MySQL database to Amazon EC2 instances by using AWS CloudFormation. Delete the stack at the beginning of the idle periods.
D. Deploy the application on Amazon EC2 instances in an Auto Scaling group behind an Application Load Balancer. Use Amazon RDS for MySQL for the database.

#14

A company needs to maintain data records for a minimum of 5 years. The data is rarely accessed after it is stored. The data must be accessible within 2 hours.

Which solution will meet these requirements MOST cost-effectively?

A. Store the data in an Amazon Elastic File System (Amazon EFS) file system. Access the data by using AWS Direct Connect.
B. Store the data in an Amazon S3 bucket. Use an S3 Lifecycle policy to move the data to S3 Glacier.
C. Store the data in an Amazon Elastic Block Store (Amazon EBS) volume. Create snapshots. Store the snapshots in an Amazon S3 bucket.
D. Store the data in an Amazon S3 bucket. Use an S3 Lifecycle policy to move the data to S3 Standard-Infrequent Access (S3 Standard-IA).

AWS Certified Solutions Architect – Associate
Official Practice Question Set

#15

A company is using an Amazon S3 bucket to store legal documents. The company frequently revises the documents and re-uploads them with the same object key to the S3 bucket. The company needs the ability to download older copies of the documents. The company also needs to protect the documents from accidental deletion.

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

A. Enable S3 Versioning on the S3 bucket.
B. Enable multi-factor authentication (MFA) delete on the S3 bucket.
C. Configure S3 Cross-Region Replication from the S3 bucket to an S3 bucket in a different AWS Region.
D. Configure an S3 Lifecycle policy to archive the documents to S3 Glacier after 30 days.

#16

A company runs its website on Amazon EC2 instances behind an Application Load Balancer that is configured as the origin for an Amazon CloudFront distribution. The company wants to protect against cross-site scripting and SQL injection attacks.

Which approach should a solutions architect recommend to meet these requirements?

A. Enable AWS Shield Advanced. List the CloudFront distribution as a protected resource.
B. Define an AWS Shield Advanced policy in AWS Firewall Manager to block cross-site scripting and SQL injection attacks.
C. Set up AWS WAF on the CloudFront distribution. Use conditions and rules that block cross-site scripting and SQL injection attacks.
D. Deploy AWS Firewall Manager on the EC2 instances. Create conditions and rules that block cross-site scripting and SQL injection attacks.

AWS Certified Solutions Architect – Associate
Official Practice Question Set

#17

Which components are required to build a site-to-site VPN connection on AWS? (Select TWO.)

A. An internet gateway
B. A NAT gateway
C. A customer gateway
D. Amazon API Gateway
E. A virtual private gateway

#18

A media company is designing a new solution for graphic rendering. The application requires up to 400 GB of storage for temporary data that is discarded after the frames are rendered. The application requires approximately 40,000 random IOPS to perform the rendering.

What is the MOST cost-effective storage option for this rendering application?

A. A storage optimized Amazon EC2 instance with instance store storage
B. A storage optimized Amazon EC2 instance with a Provisioned IOPS SSD (io1 or io2) Amazon Elastic Block Store (Amazon EBS) volume
C. A burstable Amazon EC2 instance with a Throughput Optimized HDD (st1) Amazon Elastic Block Store (Amazon EBS) volume
D. A burstable Amazon EC2 instance with Amazon S3 storage over a VPC endpoint

AWS Certified Solutions Architect – Associate
Official Practice Question Set

#19

A company is developing a new mobile version of its popular web application in the AWS Cloud. The mobile app must be accessible to internal and external users. The mobile app must handle authorization, authentication, and user management from one central source.

Which solution meets these requirements?

A. IAM roles
B. IAM users and groups
C. Amazon Cognito user pools
D. AWS Security Token Service (AWS STS)

#20

A company that processes satellite images has an application that runs on AWS. The company stores the images in an Amazon S3 bucket. For compliance reasons, the company must replicate all data once a month to an on-premises location. The average amount of data that the company needs to transfer is 60 TB.

What is the MOST cost-effective way to transfer this data?

A. Export the data monthly from the existing S3 bucket to an AWS Snowball Edge Storage Optimized device. Ship the device to the on-premises location. Transfer the data. Return the device a week later.
B. Use S3 bucket replication to copy all objects to a new S3 bucket that uses S3 Standard-Infrequent Access (S3 Standard-IA) storage. Use an AWS Storage Gateway File Gateway to transfer the data from the new S3 bucket to the on-premises location. Delete the images from the new S3 bucket after the transfer of the data.
C. Use S3 bucket replication to copy all objects to a new S3 bucket that uses S3 Standard-Infrequent Access (S3 Standard-IA) storage. Use Amazon S3 to transfer the data from the new S3 bucket to the on-premises location. Delete the images from the new S3 bucket after the transfer of the data.
D. Create an Amazon CloudFront distribution for the objects in the existing S3 bucket. Download the objects from CloudFront to the on-premises location every month.

AWS Certified Solutions Architect – Associate
Official Practice Question Set

ANSWERS

#1

Domain & Topic: 1.1 Design a multi-tier architecture solution
Correct Answer: B

A. Incorrect. With path-based routing, multiple services can use the same listener port on a single Application Load Balancer (ALB). The ALB forwards requests to specific target groups based on the URL path. However, this solution does not help with load distribution between different tasks of the same service.

   For more information about load balancing, see Service load balancing.

B. Correct. With dynamic host port mapping, multiple tasks from the same service are allowed for each container instance.

   For more information about load balancing, see Service load balancing.

C. Incorrect. You can use failover routing policies to route traffic to backup instances, in case a primary instance fails. You cannot use failover routing policies to manage multiple tasks on a single container.

   For more information about routing policies, see Choosing a routing policy.

D. Incorrect. You can use weighted routing policies to route traffic to instances at proportions that you specify. You cannot use weighted routing policies to manage multiple tasks on a single container.

   For more information about routing policies, see Choosing a routing policy.

AWS Certified Solutions Architect – Associate
Official Practice Question Set

#2

Domain & Topic: 3.2 Design secure application tiers
Correct Answers: A & E

A. Correct. A private subnet is one component to use to secure the database tier. Internet traffic is not routed to a private subnet. When you place DB instances in a private subnet, you add a layer of security.

   For more information about VPCs with public subnets and private subnets, see Routing.

B. Incorrect. An elastic network interface is a logical networking component in a VPC that represents a virtual network card. The use of an elastic network interface would not meet the requirements in the scenario.

   For more information about elastic network interfaces, see Elastic network interfaces.

C. Incorrect. Shield provides protection against DDoS attacks. Shield cannot be a target of routes in a route table. The use of Shield would not meet the requirements in the scenario.

   For more information about Shield, see AWS Shield.

D. Incorrect. Direct Connect provides a dedicated connection to your AWS environment. The use of Direct Connect would not meet the requirements in the scenario.

   For more information about Direct Connect, see AWS Direct Connect features.

E. Correct. Security groups can restrict access to the DB instances. Security groups provide access from only the application tier on only a specific port.

   For more information about security groups, see Security group basics.

AWS Certified Solutions Architect – Associate
Official Practice Question Set

#3

Domain & Topic: 1.2 Design highly available and/or fault-tolerant architectures
Correct Answer: D

A. Incorrect. VPC peering will provide connectivity to the other Availability Zone. However, VPC peering does not ensure high availability because the EC2 instances are still in one Availability Zone.

   For more information about VPC peering, see What is VPC Peering?

B. Incorrect. The replacement of the Network Load Balancer with an Application Load Balancer provides no additional availability. Both load balancers are inherently highly available. However, the EC2 instances would be highly available only if they extended across two Availability Zones.

   For more information about Elastic Load Balancing, see How Elastic Load Balancing works.

C. Incorrect. Failover routing requires a primary destination and a secondary (failover) destination. No failover destination is specified in this solution. In addition, this approach does not ensure high availability because the EC2 instances are still in one Availability Zone.

   For more information about DNS failover, see Configuring DNS failover.

D. Correct. This solution extends the EC2 instances across multiple Availability Zones and automatically adds capacity when additional capacity is needed.

   For more information about Amazon EC2 Auto Scaling, see Amazon EC2 Auto Scaling benefits.

AWS Certified Solutions Architect – Associate
Official Practice Question Set

#4

Domain & Topic: 2.1 Identify elastic and scalable compute solutions for a workload
Correct Answer: D

A. Incorrect. If an EC2 instance were removed from the target group during a scale-in process, the EC2 instance would fail (or would be unhealthy if it were checked). An Application Load Balancer would stop routing requests to that target and would choose a new healthy target.

   For more information about sticky sessions, see Sticky sessions for your Application Load Balancer.

B. Incorrect. An increase of the instance size likely would increase the speed of processing. However, this solution does not directly ensure that instances that process a request are unaffected by scale-in actions.

   For more information about deregistration delay, see Deregistration delay.

C. Incorrect. Amazon EC2 Auto Scaling cooldown periods help you prevent Auto Scaling groups from launching or terminating additional instances before the effects of previous activities are apparent.

   For more information about cooldown periods, see Scaling cooldowns for Amazon EC2 Auto Scaling.

D. Correct. By default, Elastic Load Balancing waits 300 seconds before the completion of the deregistration process, which can help in-flight requests to the target become complete. To change the amount of time that Elastic Load Balancing waits, update the deregistration delay value.

   For more information about deregistration delay, see Deregistration Delay.

#5

Domain & Topic: 4.1 Identify cost-effective storage solutions
Correct Answer: D

A. Incorrect. Amazon SQS is a fully managed message queuing service that sends messages between software components. However, Amazon SQS cannot push messages to customers.

   For information about Amazon SQS, see Amazon Simple Queue Service.

B. Incorrect. The deployment of an EC2 instance gives the company the ability to run its application. However, this solution does not use only managed services.

   For more information about Amazon EC2, see Amazon EC2.

C. Incorrect. Amazon SQS is a fully managed message queuing service that sends messages between software components. However, Amazon SQS cannot push messages to customers.

   For information about Amazon SQS, see Amazon Simple Queue Service.

D. Correct. Amazon SNS is a fully managed messaging service for application-to-application communication and application-to-person communication.

   For more information about Amazon SNS, see Amazon SNS FAQs.

#6

Domain & Topic: 2.1 Identify elastic and scalable compute solutions for a workload
Correct Answer: C

A. Incorrect. If the only action you take is to cancel the Spot request, the running instances will not be terminated automatically. These instances will continue to run and will incur additional cost.

   For more information about the termination of Spot Instances, see Terminate a Spot Instance.

B. Incorrect. When Spot Instances are terminated, new instances will launch until the Spot request is canceled.

   For more information about the termination of Spot Instances, see Terminate a Spot Instance.

C. Correct. To remove the Spot Instances, the appropriate steps are to cancel the Spot request and then to terminate the Spot Instances.

   For more information about the termination of Spot Instances, see Terminate a Spot Instance.

D. Incorrect. When Spot Instances are terminated, new instances will launch until the Spot request is canceled.

   For more information about the termination of Spot Instances, see Terminate a Spot Instance.

AWS Certified Solutions Architect – Associate
Official Practice Question Set

#7

Domain & Topic: 2.1 Identify elastic and scalable compute solutions for a workload
Correct Answer: A

A. Correct. The only way to retrieve instance metadata is to use the link-local address, which is 169.254.169.254.
For more information about instance metadata, see Retrieve instance metadata.

B. Incorrect. The use of localhost will not work because this solution checks an IP address of 127.0.0.1. Metadata is not available through the use of the localhost name.
For more information about instance metadata, see Retrieve instance metadata.

C. Incorrect. The format for the link-local address is 169.254.169.254.
For more information about instance metadata, see Retrieve instance metadata.

D. Incorrect. The 192.168.x.x. IP address range is a public block. Instance metadata is not available through a public block.
For more information about instance metadata, see Retrieve instance metadata.

#8

Domain & Topic: 2.2 Select high-performing and scalable storage solutions for a workload
Correct Answer: C

A. Incorrect. S3 Transfer Acceleration facilitates quicker uploads by using edge locations to copy data into Amazon S3. S3 Transfer Acceleration does not solve the problem of the file size limitation (5 GB) for a single PUT operation.

   For more information about S3 Transfer Acceleration, see S3 Transfer Acceleration.

B. Incorrect. This solution does not solve the problem of the file size limitation (5 GB) for a single PUT operation. S3 Lifecycle policies cannot transfer files from EC2 block storage to Amazon S3. This solution also adds unnecessary services and operational overhead to the environment.

   For more information about Amazon EC2, see What is Amazon EC2?

C. Correct. With a single PUT operation, you can upload a single object that is up to 5 GB in size. You can use a multipart upload to upload larger files, such as the files in this scenario.

   For more information about multipart uploads, see Uploading and copying objects using multipart upload.

D. Incorrect. This solution does not solve the problem of the file size limitation (5 GB) for a single PUT operation. Each destination Region would have the same problem as a single Region. This solution also adds operational overhead.

   For more information about configuring replication, see Configuring replication for source and destination buckets owned by the same account.

#9

Domain & Topic: 2.2 Select high-performing and scalable storage solutions for a workload
Correct Answer: B

A. Incorrect. A Throughput Optimized HDD EBS volume is an HDD-backed storage device that is limited to 500 IOPS for each volume.

   For more information about Throughput Optimized HDD EBS volumes, see Hard disk drives (HDD).

B. Correct. A Provisioned IOPS SSD EBS volume provides up to 64,000 IOPS for each volume.

   For more information about Provisioned IOPS SSD EBS volumes, see Solid state drives (SSD).

C. Incorrect. A General Purpose SSD EBS volume is limited to 16,000 IOPS for each volume.

   For more information about General Purpose SSD EBS volumes, see Solid state drives (SSD).

D. Incorrect. A Cold HDD volume provides low-cost magnetic storage that defines performance in terms of throughput rather than IOPS. Cold HDD volumes are a good fit for large, sequential cold-data workloads.

   For more information about Cold HDD EBS volumes, see Amazon EBS Cold HDD Volumes.

AWS Certified Solutions Architect – Associate
Official Practice Question Set

#10

Domain & Topic: 1.1 Design a multi-tier architecture solution
Correct Answer: D

A. Incorrect. With Dedicated Instances, you can reduce your costs by using existing server-bound software licenses. However, server-bound licenses are not mentioned in this scenario. One of the benefits of the AWS Cloud is that you do not have to purchase for peak consumption. The AWS Cloud can scale on demand.

   For more information about Dedicated Hosts, see Amazon EC2 Dedicated Hosts Pricing.
   For more information about demand-based scaling, see Dynamic Supply.

B. Incorrect. FIFO queues will solve problems that occur when messages are processed out of order. FIFO queues will not improve performance during sudden volume increases. Additionally, you cannot convert SQS queues after you create them.

   For more information about FIFO queues, see Amazon SQS FIFO (First-In-First-Out) queues.
   For more information about SQS queues, see Editing an Amazon SQS queue (console).

C. Incorrect. Some files in this scenario can take up to 20 minutes to process. Lambda has a 15-minute operational limit.

   For more information about Lambda constraints, see Lambda quotas.

D. Correct. With Amazon EC2 Auto Scaling, the processing capacity can keep up with the demand.

   For more information about the use of Amazon EC2 Auto Scaling with Amazon SQS, see Scaling based on Amazon SQS.

#11

Domain & Topic: 2.4 Choose high-performing database solutions for a workload
Correct Answer: B

A. Incorrect. Aurora is a relational database (not a key-value database). Aurora is not likely to achieve microsecond latency consistently.
For more information about Aurora, see What is Amazon Aurora?

B. Correct. DynamoDB is a NoSQL database that supports key-value records. DAX delivers response times in microseconds.
For more information about DynamoDB, see What is Amazon DynamoDB?
For more information about DAX, see In-Memory Acceleration with DynamoDB A (DAX).

C. Incorrect. Aurora is a relational database (not a key-value database). Aurora is not likely to achieve microsecond latency consistently, even with ElastiCache.
For more information about Aurora, see What is Amazon Aurora?
For more information about ElastiCache for Memcached, see What is Amazon ElastiCache for Memcached?

D. Incorrect. Neptune is a graph database that is optimized for working with highly connected data. Neptune is not optimized for simple key-value data.
For more information about Neptune, see What Is Amazon Neptune?

#12

Domain & Topic: 3.1 Design secure access to AWS resources
Correct Answer: B

A. Incorrect. This solution does not follow the security best practice of using roles to delegate permissions.

   For more information about how to use roles to delegate permissions, see Use roles to delegate permissions.

B. Correct. This solution follows security best practices by using a role to delegate permissions that consist of least privilege access.

   For more information about how to use roles to delegate permissions, see Use roles to delegate permissions.

C. Incorrect. The assignment of administrative privileges to a user violates security best practices and the principle of least privilege.

   For more information about how to grant least privilege in roles, see Grant least privilege.

D. Incorrect. This solution does not follow the security best practice of using roles to delegate permissions.

   For more information about how to use roles to delegate permissions, see Use roles to delegate permissions.

AWS Certified Solutions Architect – Associate
Official Practice Question Set

#13

Domain & Topic: 4.2 Identify cost-effective compute and database services
Correct Answer: A

A. Correct. When Amazon ECS uses Fargate for compute, it incurs no costs when the application is idle. Aurora Serverless also incurs no compute costs when it is idle.

   For more information about Aurora Serverless, see Amazon Aurora Serverless.

B. Incorrect. EC2 burstable instances offer burstable capability without scaling. However, this solution does not minimize cost during the periods of inactivity and is not the most cost-effective option. In addition, an Amazon Redshift database is not ideal for OLTP. Amazon Redshift is specifically designed for online analytic processing (OLAP).

   For more information about Amazon Redshift, see What is Amazon Redshift?

C. Incorrect. Although infrastructure as code (IaC) helps with availability, this solution does not meet the requirement of being always available. In addition, this solution offers no way to keep the database data after the database is recreated.

   For more information about CloudFormation, see What is AWS CloudFormation?

D. Incorrect. With this solution, at least one instance and a database will run during the periods of inactivity. This solution does not minimize cost during the periods of inactivity. This solution is not the most cost-effective option.

   For more information about Auto Scaling groups, see What is Amazon EC2 Auto Scaling?

AWS Certified Solutions Architect – Associate
Official Practice Question Set

#14

Domain & Topic: 4.1 Identify cost-effective storage solutions
Correct Answer: B

A. Incorrect. This solution is not the most cost-effective solution. Amazon S3 is a more cost-effective solution if there is not a requirement for a file system.

   For more information about Amazon EFS and Direct Connect, see How Amazon EFS works with AWS Direct Connect and AWS Managed VPN.

B. Correct. The storage of the data in an S3 bucket provides a cost-effective initial location for the data. S3 Glacier is the most cost-effective archival storage solution that meets the requirement of a 2-hour retrieval time.

   For more information about how to move data between S3 storage classes automatically, see Managing your storage lifecycle.
   For more information about S3 storage classes, see Using Amazon S3 storage classes.

C. Incorrect. This solution is not the most cost-effective solution because it requires the use of Amazon EBS in addition to Amazon S3.

   For more information about EBS snapshots, see Amazon EBS snapshots.

D. Incorrect. The storage of the data in an S3 bucket provides a cost-effective initial location for the data. However, S3 Standard-IA is not the most cost-effective storage class to meet the requirements in the scenario. The S3 Glacier storage class is designed for low-cost data archiving.

   For more information about S3 storage classes, see Using Amazon S3 storage classes.

AWS Certified Solutions Architect – Associate
Official Practice Question Set

#15

Domain & Topic: 1.4 Choose appropriate resilient storage
Correct Answer: A

A. Correct. With S3 Versioning, you can keep multiple variants of an object in the same S3 bucket. You can use S3 Versioning to preserve, retrieve, and restore every version of every object that is stored in your S3 bucket. By using S3 Versioning, you can recover from unintended user actions and application failures.

   For more information about S3 Versioning, see How S3 Versioning works.

B. Incorrect. MFA delete provides an additional layer of object security. MFA delete ensures that the entity that deletes the objects possesses an authorized MFA token. However, MFA delete does not meet the document versioning requirements of this question.

   For more information about MFA delete, see Configuring MFA delete.

C. Incorrect. S3 Cross-Region Replication requires S3 Versioning as a prerequisite. However, S3 Versioning alone meets the requirements of this question and is the more operationally efficient solution.

   For information about S3 Cross-Region Replication, see Replicating objects.

D. Incorrect. You can use S3 Lifecycle rules to store objects cost-effectively throughout their lifecycle. S3 Lifecycle rules do not meet the document versioning requirements of this question.

   For more information about S3 Lifecycle rules, see Managing your storage lifecycle.

AWS Certified Solutions Architect – Associate
Official Practice Question Set

#16

Domain & Topic: 3.3 Select appropriate data security options
Correct Answer: C

A. Incorrect. Shield Advanced protects against DDoS attacks. Shield Advanced does not protect against cross-site scripting or SQL injection.

   For more information about Shield Advanced, see AWS Shield.

B. Incorrect. With Firewall Manager, you can manage AWS WAF, Shield Advanced, and other AWS services. Shield Advanced protects against DDoS attacks. Shield Advanced does not protect against cross-site scripting or SQL injection.

   For more information about AWS WAF, AWS Shield, and Firewall Manager, see What are AWS WAF, AWS Shield, and AWS Firewall Manager?

C. Correct. AWS WAF can detect the presence of SQL code that is likely to be malicious (known as SQL injection). AWS WAF also can detect the presence of a script that is likely to be malicious (known as cross-site scripting).

   For more information about AWS WAF, see AWS WAF.

D. Incorrect. With Firewall Manager, you can manage AWS WAF, AWS Shield Advanced, and other AWS services. Firewall Manager is a managed service that is not installed on EC2 instances.

   For more information about Firewall Manager, see Firewall Manager.

AWS Certified Solutions Architect – Associate
Official Practice Question Set

#17

Domain & Topic: 2.3 Select high-performing networking solutions for a workload
Correct Answers: C & E

A. Incorrect. An internet gateway is attached to a VPC to allow traffic from the internet to flow into or out of the VPC. A VPN connection does not flow through an internet gateway. The internet gateway is designed to allow traffic from the open internet, not an encrypted VPN connection.
For more information about internet gateways, see Internet gateways.

B. Incorrect. A NAT gateway provides a way for private Amazon EC2 instances to send requests to the internet. A NAT gateway does not give you the ability to create an encrypted site-to-site VPN connection.
For more information about NAT gateways, see NAT gateways.
For more information about VPN connections to AWS, see What is AWS Site-to-Site VPN?

C. Correct. A customer gateway is required for the VPN connection to be established.A customer gateway device is set up and configured in the customer's data center.
For more information about customer gateways and VPN connections to AWS, see What is an AWS Site-to-Site VPN?

D. Incorrect. API Gateway is a fully managed service for developers to create, publish, maintain, monitor, and secure APIs at any scale. APIs act as the front door for applications to use to access data, business logic, or functionality from backend services. However, API Gateway is not necessary for the implementation of a VPN connection.
For more information about API Gateway, see What is Amazon API Gateway?

E. Correct. A virtual private gateway is attached to a VPC to create a site-to-site VPN connection on AWS. You can accept private encrypted network traffic from an on-premises data center into your VPC without the need to traverse the open public internet.

For more information about virtual private gateways and VPN connections to AWS, see What is AWS Site-to-Site VPN?

AWS Certified Solutions Architect – Associate
Official Practice Question Set

#18

Domain & Topic: 4.1 Identify cost-effective storage solutions
Correct Answer: A

A. Correct. SSD-Backed Storage Optimized (i2) instances provide more than 365,000 random IOPS. The instance store has no additional cost, compared with the regular hourly cost of the instance.

   For more information about i2 storage, see Amazon EC2 Instance Storage.
   For more information about pricing for EC2 instances, see Amazon EC2 pricing.

B. Incorrect. Provisioned IOPS SSD (io1 or io2) EBS volumes can deliver more than the 40,000 IOPS that are required in the scenario. However, this solution is not as cost-effective as an instance store because Amazon EBS adds cost to the hourly instance rate. This solution provides persistence of data beyond the lifecycle of the instance, but persistence is not required in this use case.

   For more information about Provisioned IOPS SSD (io1 or io2) EBS volumes, see Provisioned IOPS SSD volumes.
   For more information about pricing for Amazon EBS, see Amazon EBS pricing.

C. Incorrect. Throughput Optimized HDD (st1) EBS volumes are engineered to maximize the throughput of data that can be sent to and from a volume, not the random IOPS. Consequently, this solution does not meet the IOPS requirement. In addition, Amazon EBS adds cost to the hourly instance rate. This solution provides persistence of data beyond the lifecycle of the instance, but persistence is not required in this use case.

   For more information about Throughput Optimized HDD (st1) EBS volumes, see Throughput Optimized HDD volumes.
   For more information about pricing for Amazon EBS, see Amazon EBS pricing.

D. Incorrect. The rapidly changing data that is required for the scratch volume space makes Amazon S3 (object storage) the wrong storage. Block storage is appropriate for the read/write functionality to work smoothly.

   For more information about usage patterns for Amazon S3, see Usage Patterns.

AWS Certified Solutions Architect – Associate
Official Practice Question Set

#19

Domain & Topic: 3.1 Design secure access to AWS resources
Correct Answer: C

A. Incorrect. An IAM role is an IAM entity that is assumable by an IAM user. An IAM role has permissions policies that define what the entity can and cannot do. However, an IAM role does not control access to an application.

   For more information about IAM roles, see IAM roles.

B. Incorrect. You can use IAM users and groups to control who is authenticated and authorized to use an AWS service. However, users and groups do not control access to an application.

   For more information about IAM users, see IAM users.
   For more information about IAM groups, see IAM groups.

C. Correct. Amazon Cognito provides authentication, authorization, and user management for your web and mobile apps. Users can sign in directly with a user name and password, or through a trusted third party.

   For more information about Amazon Cognito, see What is Amazon Cognito?

D. Incorrect. You can use AWS STS to create and provide trusted users with temporary security credentials that can control access to your AWS resources. However, AWS STS does not control access to an application.

   For more information about temporary security credentials, see Temporary security credentials in IAM.

AWS Certified Solutions Architect – Associate
Official Practice Question Set

#20

Domain & Topic: 4.3 Design cost-optimized network architectures
Correct Answer: A

A. Correct. The base price covers the device and 10 days of usage at the on-premises location. If the company returns the device within a week, the company pays the base price and the price for data transfer out of AWS.

   For more information about Snowball pricing, see AWS Snowball Pricing.

B. Incorrect. There is no cost advantage if the company copies all the data to another S3 bucket that uses S3 Standard-IA storage. The company could transfer the data directly from the original S3 bucket. This solution is not the most cost-effective option because the additional replication increases the cost.

   For more information about data transfer pricing for Storage Gateway, see AWS Storage Gateway pricing.

C. Incorrect. There is no cost advantage if the company copies all the data to another S3 bucket that uses S3 Standard-IA storage. The company could transfer the data directly from the original S3 bucket. This solution is not the most cost-effective option because the additional replication increases the cost.

   For more information about S3 data transfer pricing, see Amazon S3 pricing.

D. Incorrect. Data transfer to CloudFront is free of cost, but data transfer of 60 TB from CloudFront to an on-premises location incurs a cost. This option would cost approximately twice as much as the option to use the AWS Snowball Edge Storage Optimized device.

   For more information about CloudFront data pricing, see Amazon CloudFront Pricing.