



Free Questions for SAA-C03 by dumpssheet

Shared by Sellers on 12-12-2023

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

A company has a three-tier environment on AWS that ingests sensor data from its users' devices. The traffic flows through a Network Load Balancer (NLB) then to Amazon EC2 instances for the web tier and finally to EC2 instances for the application tier that makes database calls.

What should a solutions architect do to improve the security of data in transit to the web tier?

Options:

- A-** Configure a TLS listener and add the server certificate on the NLB
- B-** Configure AWS Shield Advanced and enable AWS WAF on the NLB
- C-** Change the load balancer to an Application Load Balancer and attach AWS WAF to it
- D-** Encrypt the Amazon Elastic Block Store (Amazon EBS) volume on the EC2 instances using AWS Key Management Service (AWS KMS)

Answer:

A

Explanation:

A: How do you protect your data in transit?

Best Practices:

Implement secure key and certificate management: Store encryption keys and certificates securely and rotate them at appropriate time intervals while applying strict access control; for example, by using a certificate management service, such as AWS Certificate Manager (ACM).

Enforce encryption in transit: Enforce your defined encryption requirements based on appropriate standards and recommendations to help you meet your organizational, legal, and compliance requirements.

Automate detection of unintended data access: Use tools such as GuardDuty to automatically detect attempts to move data outside of defined boundaries based on data classification level, for example, to detect a trojan that is copying data to an unknown or untrusted network using the DNS protocol.

Authenticate network communications: Verify the identity of communications by using protocols that support authentication, such as Transport Layer Security (TLS) or IPsec.

https://wa.aws.amazon.com/wat.question.SEC_9.en.html

Question 2

Question Type: MultipleChoice

A solutions architect wants to use the following JSON text as an identity-based policy to grant specific permissions:

```
{  "Statement": [
    {
      "Action": [
        "ssm:ListDocuments",
        "ssm:GetDocument"
      ],
      "Effect": "Allow",
      "Resource": "*",
      "Sid": ""
    }
  ],
  "Version": "2012-10-17"
}
```

Which IAM principals can the solutions architect attach this policy to? (Select TWO.)

Options:

A- Role

- B-** Group
- C-** Organization
- D-** Amazon Elastic Container Service (Amazon ECS) resource
- E-** Amazon EC2 resource

Answer:

A, B

Explanation:

This JSON text is an identity-based policy that grants specific permissions. The IAM principals that the solutions architect can attach this policy to are Role and Group. This is because the policy is written in JSON and is an identity-based policy, which can be attached to IAM principals such as users, groups, and roles. Identity-based policies are permissions policies that you attach to IAM identities (users, groups, or roles) and explicitly state what that identity is allowed (or denied) to do¹. Identity-based policies are different from resource-based policies, which define the permissions around the specific resource¹. Resource-based policies are attached to a resource, such as an Amazon S3 bucket or an Amazon EC2 instance¹. Resource-based policies can also specify a principal, which is the entity that is allowed or denied access to the resource¹. Organization is not an IAM principal, but a feature of AWS Organizations that allows you to manage multiple AWS accounts centrally². Amazon ECS resource and Amazon EC2 resource are not IAM principals, but AWS resources that can have resource-based policies attached to them^{3,4}.

Identity-based policies and resource-based policies

[AWS Organizations](#)

[Amazon ECS task role](#)

[Amazon EC2 instance profile](#)

Question 3

Question Type: MultipleChoice

A company is running a legacy system on an Amazon EC2 instance. The application code cannot be modified, and the system cannot run on more than one instance. A solutions architect must design a resilient solution that can improve the recovery time for the system.

What should the solutions architect recommend to meet these requirements?

Options:

A- Enable termination protection for the EC2 instance.

B- Configure the EC2 instance for Multi-AZ deployment.

C- Create an Amazon CloudWatch alarm to recover the EC2 instance in case of failure.

D- Launch the EC2 instance with two Amazon Elastic Block Store (Amazon EBS) volumes that use RAID configurations for storage redundancy.

Answer:

C

Explanation:

To design a resilient solution that can improve the recovery time for the system, a solutions architect should recommend creating an Amazon CloudWatch alarm to recover the EC2 instance in case of failure. This solution has the following benefits:

It allows the EC2 instance to be automatically recovered when a system status check failure occurs, such as loss of network connectivity, loss of system power, software issues on the physical host, or hardware issues on the physical host that impact network reachability¹.

It preserves the instance ID, private IP addresses, Elastic IP addresses, and all instance metadata of the original instance. A recovered instance is identical to the original instance, except for any data that is in-memory, which is lost during the recovery process¹.

It does not require any modification of the application code or the EC2 instance configuration. The solutions architect can create a CloudWatch alarm using the AWS Management Console, the AWS CLI, or the CloudWatch API².

1: <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-instance-recover.html>

2: <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-instance-recover.html#ec2-instance-recover-create-alarm>

Question 4

Question Type: MultipleChoice

A gaming company wants to launch a new internet-facing application in multiple AWS Regions. The application will use the TCP and UDP protocols for communication. The company needs to provide high availability and minimum latency for global users.

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

Options:

- A- Create internal Network Load Balancers in front of the application in each Region.
- B- Create external Application Load Balancers in front of the application in each Region.
- C- Create an AWS Global Accelerator accelerator to route traffic to the load balancers in each Region.
- D- Configure Amazon Route 53 to use a geolocation routing policy to distribute the traffic.
- E- Configure Amazon CloudFront to handle the traffic and route requests to the application in each Region.

Answer:

B, C

Explanation:

This combination of actions will provide high availability and minimum latency for global users by using AWS Global Accelerator and Application Load Balancers. AWS Global Accelerator is a networking service that helps you improve the availability, performance, and security of your internet-facing applications by using the AWS global network. It provides two global static public IPs that act as a fixed entry point to your application endpoints, such as Application Load Balancers, in multiple Regions¹. Global Accelerator uses the AWS backbone network to route traffic to the optimal regional endpoint based on health, client location, and policies that you configure. It also offers TCP and UDP support, traffic encryption, and DDoS protection². Application Load Balancers are external load balancers that distribute incoming application traffic across multiple targets, such as EC2 instances, in multiple Availability Zones. They support both HTTP and HTTPS (SSL/TLS) protocols, and offer advanced features such as content-based routing, health checks, and integration with other AWS services³. By creating external Application Load Balancers in front of the application in each Region, you can ensure that the application can handle varying load patterns and scale on demand. By creating an AWS Global Accelerator accelerator to route traffic to the load balancers in each Region, you can leverage the performance, security, and availability of the AWS global network to deliver the best possible user experience.

Question 5

Question Type: MultipleChoice

A company is designing a tightly coupled high performance computing (HPC) environment in the AWS Cloud. The company needs to include features that will optimize the HPC environment for networking and storage.

Which combination of solutions will meet these requirements? (Select TWO)

Options:

- A- Create an accelerator in AWS Global Accelerator. Configure custom routing for the accelerator.
- B- Create an Amazon FSx for Lustre file system. Configure the file system with scratch storage.
- C- Create an Amazon CloudFront distribution. Configure the viewer protocol policy to be HTTP and HTTPS.
- D- Launch Amazon EC2 instances. Attach an Elastic Fabric Adapter (EFA) to the instances.
- E- Create an AWS Elastic Beanstalk deployment to manage the environment.

Answer:

B, D

Explanation:

These two solutions will optimize the HPC environment for networking and storage. Amazon FSx for Lustre is a fully managed service that provides cost-effective, high-performance, scalable storage for compute workloads. It is built on the world's most popular high-performance file system, Lustre, which is designed for applications that require fast storage, such as HPC and machine learning. By configuring the file system with scratch storage, you can achieve sub-millisecond latencies, up to hundreds of GBs/s of throughput, and millions of IOPS. Scratch file systems are ideal for temporary storage and shorter-term processing of data. Data is not replicated and does not persist if a file server fails. For more information, see [Amazon FSx for Lustre](#).

Elastic Fabric Adapter (EFA) is a network interface for Amazon EC2 instances that enables customers to run applications requiring high levels of inter-node communications at scale on AWS. Its custom-built operating system (OS) bypass hardware interface enhances the

performance of inter-instance communications, which is critical to scaling HPC and machine learning applications. EFA provides a low-latency, low-jitter channel for inter-instance communications, enabling your tightly-coupled HPC or distributed machine learning applications to scale to thousands of cores. EFA uses libfabric interface and libfabric APIs for communications, which are supported by most HPC programming models. For more information, see [Elastic Fabric Adapter](#).

The other solutions are not suitable for optimizing the HPC environment for networking and storage. AWS Global Accelerator is a networking service that helps you improve the availability, performance, and security of your public applications by using the AWS global network. It provides two global static public IPs, deterministic routing, fast failover, and TCP termination at the edge for your application endpoints. However, it does not support OS-bypass capabilities or high-performance file systems that are required for HPC and machine learning applications. For more information, see [AWS Global Accelerator](#).

Amazon CloudFront is a content delivery network (CDN) service that securely delivers data, videos, applications, and APIs to customers globally with low latency, high transfer speeds, all within a developer-friendly environment. CloudFront is integrated with AWS services such as Amazon S3, Amazon EC2, AWS Elemental Media Services, AWS Shield, AWS WAF, and AWS Lambda@Edge. However, CloudFront is not designed for HPC and machine learning applications that require high levels of inter-node communications and fast storage. For more information, see [\[Amazon CloudFront\]](#).

AWS Elastic Beanstalk is an easy-to-use service for deploying and scaling web applications and services developed with Java, .NET, PHP, Node.js, Python, Ruby, Go, and Docker on familiar servers such as Apache, Nginx, Passenger, and IIS. You can simply upload your code and Elastic Beanstalk automatically handles the deployment, from capacity provisioning, load balancing, auto-scaling to application health monitoring. However, Elastic Beanstalk is not optimized for HPC and machine learning applications that require OS-bypass capabilities and high-performance file systems. For more information, see [\[AWS Elastic Beanstalk\]](#).

Question 6

Question Type: MultipleChoice

An analytics company uses Amazon VPC to run its multi-tier services. The company wants to use RESTful APIs to offer a web analytics service to millions of users. Users must be verified by using an authentication service to access the APIs.

Which solution will meet these requirements with the MOST operational efficiency?

Options:

- A-** Configure an Amazon Cognito user pool for user authentication. Implement Amazon API Gateway REST APIs with a Cognito authorizer.
- B-** Configure an Amazon Cognito identity pool for user authentication. Implement Amazon API Gateway HTTP APIs with a Cognito authorizer.
- C-** Configure an AWS Lambda function to handle user authentication. Implement Amazon API Gateway REST APIs with a Lambda authorizer.
- D-** Configure an IAM user to handle user authentication. Implement Amazon API Gateway HTTP APIs with an IAM authorizer.

Answer:

A

Explanation:

This solution will meet the requirements with the most operational efficiency because:

Amazon Cognito user pools provide a secure and scalable user directory that can store and manage user profiles, and handle user sign-up, sign-in, and access control. User pools can also integrate with social identity providers and enterprise identity providers via SAML or OIDC. User pools can issue JSON Web Tokens (JWTs) that can be used to authenticate users and authorize API requests.

Amazon API Gateway REST APIs enable you to create and deploy APIs that expose your backend services to your clients. REST APIs support multiple authorization mechanisms, including Cognito user pools, IAM, Lambda, and custom authorizers. A Cognito authorizer is a type of Lambda authorizer that uses a Cognito user pool as the identity source. When a client makes a request to a REST API method that is configured with a Cognito authorizer, API Gateway verifies the JWTs that are issued by the user pool and grants access based on the token's claims and the authorizer's configuration.

By using Cognito user pools and API Gateway REST APIs with a Cognito authorizer, you can achieve a high level of security, scalability, and performance for your web analytics service. You can also leverage the built-in features of Cognito and API Gateway, such as user management, token validation, caching, throttling, and monitoring, without having to implement them yourself. This reduces the operational overhead and complexity of your solution.

[Amazon Cognito User Pools](#)

[Amazon API Gateway REST APIs](#)

[Use API Gateway Lambda authorizers](#)

Question 7

Question Type: MultipleChoice

A company has an application that delivers on-demand training videos to students around the world. The application also allows authorized content developers to upload videos. The data is stored in an Amazon S3 bucket in the us-east-2 Region.

The company has created an S3 bucket in the eu-west-2 Region and an S3 bucket in the ap-southeast-1 Region. The company wants to replicate the data to the new S3 buckets. The company needs to minimize latency for developers who upload videos and students who stream videos near eu-west-2 and ap-southeast-1.

Which combination of steps will meet these requirements with the FEWEST changes to the application? (Select TWO.)

Options:

- A-** Configure one-way replication from the us-east-2 S3 bucket to the eu-west-2 S3 bucket. Configure one-way replication from the us-east-2 S3 bucket to the ap-southeast-1 S3 bucket.
- B-** Configure one-way replication from the us-east-2 S3 bucket to the eu-west-2 S3 bucket. Configure one-way replication from the eu-west-2 S3 bucket to the ap-southeast-1 S3 bucket.
- C-** Configure two-way (bidirectional) replication among the S3 buckets that are in all three Regions.
- D-** Create an S3 Multi-Region Access Point. Modify the application to use the Amazon Resource Name (ARN) of the Multi-Region Access Point for video streaming. Do not modify the application for video uploads.
- E-** Create an S3 Multi-Region Access Point. Modify the application to use the Amazon Resource Name (ARN) of the Multi-Region Access

Point for video streaming and uploads.

Answer:

A, E

Explanation:

These two steps will meet the requirements with the fewest changes to the application because they will enable the company to replicate the data to the new S3 buckets and minimize latency for both video streaming and uploads. One-way replication from the us-east-2 S3 bucket to the other two S3 buckets will ensure that the data is synchronized across all three regions. The company can use S3 Cross-Region Replication (CRR) to automatically copy objects across buckets in different AWS Regions. CRR can help the company achieve lower latency and compliance requirements by keeping copies of their data in different regions. Creating an S3 Multi-Region Access Point and modifying the application to use its ARN will allow the company to access the data through a single global endpoint. An S3 Multi-Region Access Point is a globally unique name that can be used to access objects stored in S3 buckets across multiple regions. It automatically routes requests to the closest S3 bucket with the lowest latency. By using an S3 Multi-Region Access Point, the company can simplify the application architecture and improve the performance and reliability of the application.

[Replicating objects](#)

[Multi-Region Access Points in Amazon S3](#)

Question 8

Question Type: MultipleChoice

A company is deploying an application in three AWS Regions using an Application Load Balancer. Amazon Route 53 will be used to distribute traffic between these Regions. Which Route 53 configuration should a solutions architect use to provide the MOST high-performing experience?

Options:

- A-** Create an A record with a latency policy.
- B-** Create an A record with a geolocation policy.
- C-** Create a CNAME record with a failover policy.
- D-** Create a CNAME record with a geoproximity policy.

Answer:

A

Explanation:

To provide the most high-performing experience for the users of the application, a solutions architect should use a latency routing policy for the Route 53 A record. This policy allows Route 53 to route traffic to the AWS Region that provides the lowest possible latency for the users. A latency routing policy can also improve the availability of the application, as Route 53 can automatically route traffic to another

Region if the primary Region becomes unavailable2.

1: <https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-policy.html#routing-policy-latency>

2: https://aws.amazon.com/route53/faqs/#Latency_Based_Routing

Question 9

Question Type: MultipleChoice

A company is running a photo hosting service in the us-east-1 Region. The service enables users across multiple countries to upload and view photos. Some photos are heavily viewed for months, and others are viewed for less than a week. The application allows uploads of up to 20 MB for each photo. The service uses the photo metadata to determine which photos to display to each user.

Which solution provides the appropriate user access MOST cost-effectively?

Options:

- A-** Store the photos in Amazon DynamoDB. Turn on DynamoDB Accelerator (DAX) to cache frequently viewed items.
- B-** Store the photos in the Amazon S3 Intelligent-Tiering storage class. Store the photo metadata and its S3 location in DynamoDB.
- C-** Store the photos in the Amazon S3 Standard storage class. Set up an S3 Lifecycle policy to move photos older than 30 days to the

S3 Standard-Infrequent Access (S3 Standard-IA) storage class. Use the object tags to keep track of metadata.

D- Store the photos in the Amazon S3 Glacier storage class. Set up an S3 Lifecycle policy to move photos older than 30 days to the S3 Glacier Deep Archive storage class. Store the photo metadata and its S3 location in Amazon OpenSearch Service.

Answer:

B

Explanation:

This solution provides the appropriate user access most cost-effectively because it uses the Amazon S3 Intelligent-Tiering storage class, which automatically optimizes storage costs by moving data to the most cost-effective access tier when access patterns change, without performance impact or operational overhead¹. This storage class is ideal for data with unknown, changing, or unpredictable access patterns, such as photos that are heavily viewed for months or less than a week. By storing the photo metadata and its S3 location in DynamoDB, the application can quickly query and retrieve the relevant photos for each user. DynamoDB is a fast, scalable, and fully managed NoSQL database service that supports key-value and document data models².

Question 10

Question Type: MultipleChoice

A company copies 200 TB of data from a recent ocean survey onto AWS Snowball Edge Storage Optimized devices. The company has a high performance computing (HPC) cluster that is hosted on AWS to look for oil and gas deposits. A solutions architect must provide the cluster with consistent sub-millisecond latency and high-throughput access to the data on the Snowball Edge Storage Optimized devices. The company is sending the devices back to AWS.

Which solution will meet these requirements?

Options:

- A-** Create an Amazon S3 bucket. Import the data into the S3 bucket. Configure an AWS Storage Gateway file gateway to use the S3 bucket. Access the file gateway from the HPC cluster instances.
- B-** Create an Amazon S3 bucket. Import the data into the S3 bucket. Configure an Amazon FSx for Lustre file system, and integrate it with the S3 bucket. Access the FSx for Lustre file system from the HPC cluster instances.
- C-** Create an Amazon S3 bucket and an Amazon Elastic File System (Amazon EFS) file system. Import the data into the S3 bucket. Copy the data from the S3 bucket to the EFS file system. Access the EFS file system from the HPC cluster instances.
- D-** Create an Amazon FSx for Lustre file system. Import the data directly into the FSx for Lustre file system. Access the FSx for Lustre file system from the HPC cluster instances.

Answer:

B

Explanation:

To provide the HPC cluster with consistent sub-millisecond latency and high-throughput access to the data on the Snowball Edge Storage Optimized devices, a solutions architect should configure an Amazon FSx for Lustre file system, and integrate it with an Amazon S3 bucket. This solution has the following benefits:

It allows the HPC cluster to access the data on the Snowball Edge devices using a POSIX-compliant file system that is optimized for fast processing of large datasets¹.

It enables the data to be imported from the Snowball Edge devices into the S3 bucket using the AWS Snow Family Console or the AWS CLI². The data can then be accessed from the FSx for Lustre file system using the S3 integration feature³.

It supports high availability and durability of the data, as the FSx for Lustre file system can automatically copy the data to and from the S3 bucket³. The data can also be accessed from other AWS services or applications using the S3 API⁴.

1: <https://aws.amazon.com/fsx/lustre/>

2: <https://docs.aws.amazon.com/snowball/latest/developer-guide/using-adapter.html>

3: <https://docs.aws.amazon.com/fsx/latest/LustreGuide/create-fs-linked-data-repo.html>

4: <https://docs.aws.amazon.com/fsx/latest/LustreGuide/export-data-repo.html>

To Get Premium Files for SAA-C03 Visit

<https://www.p2pexams.com/products/saa-c03>

For More Free Questions Visit

<https://www.p2pexams.com/amazon/pdf/saa-c03>

