

Amazon

Exam Questions AWS-Solution-Architect-Associate

Amazon AWS Certified Solutions Architect - Associate



NEW QUESTION 1

- (Topic 4)

A company has two VPCs named Management and Production. The Management VPC uses VPNs through a customer gateway to connect to a single device in the data center. The Production VPC uses a virtual private gateway AWS Direct Connect connections. The Management and Production VPCs both use a single VPC peering connection to allow communication between the

What should a solutions architect do to mitigate any single point of failure in this architecture?

- A. Add a set of VPNs between the Management and Production VPCs.
- B. Add a second virtual private gateway and attach it to the Management VPC.
- C. Add a second set of VPNs to the Management VPC from a second customer gateway device.
- D. Add a second VPC peering connection between the Management VPC and the Production VPC.

Answer: C

Explanation:

This answer is correct because it provides redundancy for the VPN connection between the Management VPC and the data center. If one customer gateway device or one VPN tunnel becomes unavailable, the traffic can still flow over the second customer gateway device and the second VPN tunnel. This way, the single point of failure in the VPN connection is mitigated.

References:

? <https://docs.aws.amazon.com/vpn/latest/s2svpn/vpn-redundant-connection.html>

? <https://www.trendmicro.com/cloudoneconformity/knowledge-base/aws/VPC/vpn-tunnel-redundancy.html>

NEW QUESTION 2

- (Topic 4)

A company wants to use an AWS CloudFormation stack for its application in a test environment. The company stores the CloudFormation template in an Amazon S3 bucket that blocks public access. The company wants to grant CloudFormation access to the template in the S3 bucket based on specific user requests to create the test environment. The solution must follow security best practices.

Which solution will meet these requirements?

- A. Create a gateway VPC endpoint for Amazon S3. Configure the CloudFormation stack to use the S3 object URL.
- B. Create an Amazon API Gateway REST API that has the S3 bucket as the target.
- C. Configure the CloudFormation stack to use the API Gateway URL.
- D. Create a presigned URL for the template object. Configure the CloudFormation stack to use the presigned URL.
- E. Allow public access to the template object in the S3 bucket.
- F. Block the public access after the test environment is created.

Answer: C

Explanation:

It allows CloudFormation to access the template in the S3 bucket without granting public access or creating additional resources. A presigned URL is a URL that is signed with the access key of an IAM user or role that has permission to access the object. The presigned URL can be used by anyone who receives it, but it expires after a specified time. By creating a presigned URL for the template object and configuring the CloudFormation stack to use it, the company can grant CloudFormation access to the template based on specific user requests and follow security best practices. References:

? [Using Amazon S3 Presigned URLs](#)

? [Using Amazon S3 Buckets](#)

NEW QUESTION 3

- (Topic 4)

A company is moving its data and applications to AWS during a multiyear migration project. The company wants to securely access data on Amazon S3 from the company's AWS Region and from the company's on-premises location. The data must not traverse the internet. The company has established an AWS Direct Connect connection between its Region and its on-premises location.

Which solution will meet these requirements?

- A. Create gateway endpoints for Amazon S3. Use the gateway endpoints to securely access the data from the Region and the on-premises location.
- B. Create a gateway in AWS Transit Gateway to access Amazon S3 securely from the Region and the on-premises location.
- C. Create interface endpoints for Amazon S3. Use the interface endpoints to securely access the data from the Region and the on-premises location.
- D. Use an AWS Key Management Service (AWS KMS) key to access the data securely from the Region and the on-premises location.

Answer: B

Explanation:

A gateway endpoint is a gateway that is a target for a specified route in your route table, used for traffic destined to a supported AWS service¹. Amazon S3 does not support gateway endpoints, only interface endpoints². Therefore, option A is incorrect.

An interface endpoint is an elastic network interface with a private IP address that serves as an entry point for traffic destined to a supported service¹. An interface endpoint can provide secure access to Amazon S3 from within the Region, but not from the on-premises location. Therefore, option C is incorrect.

AWS Key Management Service (AWS KMS) is a service that allows you to create and manage encryption keys to protect your data³. AWS KMS does not provide a way to access data on Amazon S3 without traversing the internet. Therefore, option D is incorrect. AWS Transit Gateway is a service that enables you to connect your Amazon Virtual Private Clouds (VPCs) and your on-premises networks to a single gateway. You can create a gateway in AWS Transit Gateway to access Amazon S3 securely from both the Region and the on-premises location using AWS Direct Connect. Therefore, option B is correct.

NEW QUESTION 4

- (Topic 4)

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.)

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

Answer: AB

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}. References:

- ? Identity-based policies and resource-based policies
- ? AWS Organizations
- ? Amazon ECS task role
- ? Amazon EC2 instance profile

NEW QUESTION 5

- (Topic 4)

A company is building a shopping application on AWS. The application offers a catalog that changes once each month and needs to scale with traffic volume. The company wants the lowest possible latency from the application. Data from each user's shopping cart needs to be highly available. User session data must be available even if the user is disconnected and reconnects.

What should a solutions architect do to ensure that the shopping cart data is preserved at all times?

- A. Configure an Application Load Balancer to enable the sticky sessions feature (session affinity) for access to the catalog in Amazon Aurora.
- B. Configure Amazon ElastiCache for Redis to cache catalog data from Amazon DynamoDB and shopping cart data from the user's session.
- C. Configure Amazon OpenSearch Service to cache catalog data from Amazon DynamoDB and shopping cart data from the user's session.
- D. Configure an Amazon EC2 instance with Amazon Elastic Block Store (Amazon EBS) storage for the catalog and shopping cart.
- E. Configure automated snapshots.

Answer: B

Explanation:

To ensure that the shopping cart data is preserved at all times, a solutions architect should configure Amazon ElastiCache for Redis to cache catalog data from Amazon DynamoDB and shopping cart data from the user's session. This solution has the following benefits:

- ? It offers the lowest possible latency from the application, as ElastiCache for Redis is a blazing fast in-memory data store that provides sub-millisecond latency to power internet-scale real-time applications¹.
- ? It scales with traffic volume, as ElastiCache for Redis supports horizontal scaling by adding more nodes or shards to the cluster, and vertical scaling by changing the node type².
- ? It is highly available, as ElastiCache for Redis supports replication across multiple Availability Zones and automatic failover in case of a primary node failure³.
- ? It preserves user session data even if the user is disconnected and reconnects, as ElastiCache for Redis can store session data, such as user login information and shopping cart contents, in a persistent and durable manner using snapshots or AOF (append-only file) persistence⁴.

References:

- ? 1: <https://aws.amazon.com/elasticache/redis/>
- ? 2: <https://docs.aws.amazon.com/AmazonElastiCache/latest/red-ug/Scaling.html>
- ? 3: <https://docs.aws.amazon.com/AmazonElastiCache/latest/red-ug/Replication.html>
- ? 4: <https://docs.aws.amazon.com/AmazonElastiCache/latest/red-ug/backups.html>

NEW QUESTION 6

- (Topic 4)

A social media company runs its application on Amazon EC2 instances behind an Application Load Balancer (ALB). The ALB is the origin for an Amazon CloudFront distribution. The application has more than a billion images stored in an Amazon S3 bucket and processes thousands of images each second. The company wants to resize the images dynamically and serve appropriate formats to clients.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Install an external image management library on an EC2 instance
- B. Use the image management library to process the images.
- C. Create a CloudFront origin request policy
- D. Use the policy to automatically resize images and to serve the appropriate format based on the User-Agent HTTP header in the request.
- E. Use a Lambda@Edge function with an external image management library
- F. Associate the Lambda@Edge function with the CloudFront behaviors that serve the images.
- G. Create a CloudFront response headers policy
- H. Use the policy to automatically resize images and to serve the appropriate format based on the User-Agent HTTP header in the request.

Answer: C

Explanation:

Lambda@Edge is a service that allows you to run Lambda functions at CloudFront edge locations. It can be used to modify requests and responses that flow through CloudFront. CloudFront origin request policy is a policy that controls the values (URL query strings, HTTP headers, and cookies) that are included in requests that CloudFront sends to the origin. It can be used to collect additional information at the origin or to customize the origin response. CloudFront response headers policy is a policy that specifies the HTTP headers that CloudFront removes or adds in responses that it sends to viewers. It can be used to add security or custom headers to responses.

Based on these definitions, the solution that will meet the requirements with the least operational overhead is:

* C. Use a Lambda@Edge function with an external image management library. Associate the Lambda@Edge function with the CloudFront behaviors that serve the images.

This solution would allow the application to use a Lambda@Edge function to resize the images dynamically and serve appropriate formats to clients based on the User-Agent HTTP header in the request. The Lambda@Edge function would run at the edge locations, reducing latency and load on the origin. The application code would only need to include an external image management library that can perform image manipulation tasks.

NEW QUESTION 7

- (Topic 4)

A solutions architect is designing a highly available Amazon ElastiCache for Redis based solution. The solutions architect needs to ensure that failures do not result in performance degradation or loss of data locally and within an AWS Region. The solution needs to provide high availability at the node level and at the Region level.

Which solution will meet these requirements?

- A. Use Multi-AZ Redis replication groups with shards that contain multiple nodes.
- B. Use Redis shards that contain multiple nodes with Redis append only files (AOF) turned on.
- C. Use a Multi-AZ Redis cluster with more than one read replica in the replication group.
- D. Use Redis shards that contain multiple nodes with Auto Scaling turned on.

Answer: A

Explanation:

This answer is correct because it provides high availability at the node level and at the Region level for the ElastiCache for Redis solution. A Multi-AZ Redis replication group consists of a primary cluster and up to five read replica clusters, each in a different Availability Zone. If the primary cluster fails, one of the read replicas is automatically promoted to be the new primary cluster. A Redis replication group with shards enables partitioning of the data across multiple nodes, which increases the scalability and performance of the solution. Each shard can have one or more replicas to provide redundancy and read scaling.

References:

? <https://docs.aws.amazon.com/AmazonElastiCache/latest/red-ug/AutoFailover.html>

? <https://docs.aws.amazon.com/AmazonElastiCache/latest/red-ug/Shards.html>

NEW QUESTION 8

- (Topic 4)

A financial company needs to handle highly sensitive data. The company will store the data in an Amazon S3 bucket. The company needs to ensure that the data is encrypted in transit and at rest. The company must manage the encryption keys outside the AWS Cloud.

Which solution will meet these requirements?

- A. Encrypt the data in the S3 bucket with server-side encryption (SSE) that uses an AWS Key Management Service (AWS KMS) customer managed key.
- B. Encrypt the data in the S3 bucket with server-side encryption (SSE) that uses an AWS Key Management Service (AWS KMS) AWS managed key.
- C. Encrypt the data in the S3 bucket with the default server-side encryption (SSE).
- D. Encrypt the data at the company's data center before storing the data in the S3 bucket.

Answer: D

Explanation:

This option is the only solution that meets the requirements because it allows the company to encrypt the data with its own encryption keys and tools outside the AWS Cloud. By encrypting the data at the company's data center before storing the data in the S3 bucket, the company can ensure that the data is encrypted in transit and at rest, and that the company has full control over the encryption keys and processes. This option also avoids the need to use any AWS encryption services or features, which may not be compatible with the company's security policies or compliance standards.

* A. Encrypt the data in the S3 bucket with server-side encryption (SSE) that uses an AWS Key Management Service (AWS KMS) customer managed key. This option does not meet the requirements because it does not allow the company to manage the encryption keys outside the AWS Cloud. Although the company can create and use its own customer managed key in AWS KMS, the key is still stored and managed by AWS KMS, which is a service within the AWS Cloud.

Moreover, the company still needs to use the AWS encryption features and APIs to encrypt and decrypt the data in the S3 bucket, which may not be compatible with the company's security policies or compliance standards.

* B. Encrypt the data in the S3 bucket with server-side encryption (SSE) that uses an AWS Key Management Service (AWS KMS) AWS managed key. This option does not meet the requirements because it does not allow the company to manage the encryption keys outside the AWS Cloud. In this option, the company uses

the default AWS managed key in AWS KMS, which is created and managed by AWS on behalf of the company. The company has no control over the key rotation, deletion, or recovery policies. Moreover, the company still needs to use the AWS encryption features and APIs to encrypt and decrypt the data in the S3 bucket, which may not be compatible with the company's security policies or compliance standards.

* C. Encrypt the data in the S3 bucket with the default server-side encryption (SSE). This option does not meet the requirements because it does not allow the company to manage the encryption keys outside the AWS Cloud. In this option, the company uses the default server-side encryption with Amazon S3 managed keys (SSE-S3), which is applied to every bucket in Amazon S3. The company has no visibility or control over the encryption keys, which are managed by Amazon S3. Moreover, the company still needs to use the AWS encryption features and APIs to encrypt and decrypt the data in the S3 bucket, which may not be compatible with the company's security policies or compliance standards. References:

- ? 1 Protecting data with encryption - Amazon Simple Storage Service
- ? 2 Protecting data with server-side encryption - Amazon Simple Storage Service
- ? 3 Protecting data by using client-side encryption - Amazon Simple Storage Service
- ? 4 AWS Key Management Service Concepts - AWS Key Management Service

NEW QUESTION 9

- (Topic 4)

A company has users all around the world accessing its HTTP-based application deployed on Amazon EC2 instances in multiple AWS Regions. The company wants to improve the availability and performance of the application. The company also wants to protect the application against common web exploits that may affect availability, compromise security, or consume excessive resources. Static IP addresses are required.

What should a solutions architect recommend to accomplish this?

- A. Put the EC2 instances behind Network Load Balancers (NLBs) in each Region
- B. Deploy AWS WAF on the NLB
- C. Create an accelerator using AWS Global Accelerator and register the NLBs as endpoints.
- D. Put the EC2 instances behind Application Load Balancers (ALBs) in each Region
- E. Deploy AWS WAF on the ALB
- F. Create an accelerator using AWS Global Accelerator and register the ALBs as endpoints.
- G. Put the EC2 instances behind Network Load Balancers (NLBs) in each Region
- H. Deploy AWS WAF on the NLB
- I. Create an Amazon CloudFront distribution with an origin that uses Amazon Route 53 latency-based routing to route requests to the NLBs.
- J. Put the EC2 instances behind Application Load Balancers (ALBs) in each Region
- K. Create an Amazon CloudFront distribution with an origin that uses Amazon Route 53 latency-based routing to route requests to the ALB
- L. Deploy AWS WAF on the CloudFront distribution.

Answer: A

Explanation:

The company wants to improve the availability and performance of the application, as well as protect it against common web exploits. The company also needs static IP addresses for the application. To meet these requirements, a solutions architect should recommend the following solution:

? Put the EC2 instances behind Network Load Balancers (NLBs) in each Region.

NLBs are designed to handle millions of requests per second while maintaining high throughput at ultra-low latency. NLBs also support static IP addresses for each Availability Zone, which can be useful for whitelisting or firewalling purposes.

? Deploy AWS WAF on the NLBs. AWS WAF is a web application firewall that helps

protect web applications from common web exploits that could affect availability, security, or performance. AWS WAF lets you define customizable web security rules that control which traffic to allow or block to your web applications.

? Create an accelerator using AWS Global Accelerator and register the NLBs as

endpoints. AWS Global Accelerator is a service that improves the availability and performance of your applications with local or global users. It provides static IP addresses that act as a fixed entry point to your application endpoints in any AWS Region. It uses the AWS global network to optimize the path from your users to your applications, improving the performance of your TCP and UDP traffic.

This solution will provide high availability across Availability Zones and Regions, improve performance by routing traffic over the AWS global network, protect the application from common web attacks, and provide static IP addresses for the application.

References:

- ? Network Load Balancer
- ? AWS WAF
- ? AWS Global Accelerator

NEW QUESTION 10

- (Topic 4)

A company hosts a multi-tier web application on Amazon Linux Amazon EC2 instances behind an Application Load Balancer. The instances run in an Auto Scaling group across multiple Availability Zones. The company observes that the Auto Scaling group launches more On-Demand Instances when the application's end users access high volumes of static web content. The company wants to optimize cost.

What should a solutions architect do to redesign the application MOST cost-effectively?

- A. Update the Auto Scaling group to use Reserved Instances instead of On-Demand Instances.
- B. Update the Auto Scaling group to scale by launching Spot Instances instead of On-Demand Instances.
- C. Create an Amazon CloudFront distribution to host the static web contents from an Amazon S3 bucket.
- D. Create an AWS Lambda function behind an Amazon API Gateway API to host the static website contents.

Answer: C

Explanation:

This answer is correct because it meets the requirements of optimizing cost and reducing the workload on the database. Amazon CloudFront is a content delivery network (CDN) service that speeds up distribution of your static and dynamic web content, such as .html, .css, .js, and image files, to your users. CloudFront delivers your content through a worldwide network of data centers called edge locations. When a user requests content that you're serving with CloudFront, the request is routed to the edge location that provides the lowest latency (time delay), so that content is delivered with the best possible performance. You can create an Amazon CloudFront distribution to host the static web contents from an Amazon S3 bucket, which is an origin that you define for CloudFront. This way, you can offload the requests for static web content from your EC2 instances to CloudFront, which can improve the performance and availability of your website, and reduce the cost of running your EC2 instances.

References:

- ? <https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/Introduction.html>
- ? <https://docs.aws.amazon.com/AmazonS3/latest/userguide/WebsiteHosting.html>

NEW QUESTION 10

- (Topic 4)

A company runs an application on AWS. The application receives inconsistent amounts of usage. The application uses AWS Direct Connect to connect to an on-premises MySQL-compatible database. The on-premises database consistently uses a minimum of 2 GiB of memory. The company wants to migrate the on-premises database to a managed AWS service. The company wants to use auto scaling capabilities to manage unexpected workload increases.

Which solution will meet these requirements with the LEAST administrative overhead?

- A. Provision an Amazon DynamoDB database with default read and write capacity settings.
- B. Provision an Amazon Aurora database with a minimum capacity of 1 Aurora capacity unit (ACU).
- C. Provision an Amazon Aurora Serverless v2 database with a minimum capacity of 1 Aurora capacity unit (ACU).
- D. Provision an Amazon RDS for MySQL database with 2 GiB of memory.

Answer: C

Explanation:

it allows the company to migrate the on-premises database to a managed AWS service that supports auto scaling capabilities and has the least administrative overhead. Amazon Aurora Serverless v2 is a configuration of Amazon Aurora that automatically scales compute capacity based on workload demand. It can scale from hundreds to hundreds of thousands of transactions in a fraction of a second. Amazon Aurora Serverless v2 also supports MySQL-compatible databases and AWS Direct Connect connectivity. References:

? Amazon Aurora Serverless v2

? Connecting to an Amazon Aurora DB Cluster

NEW QUESTION 15

- (Topic 4)

A company is designing a new web service that will run on Amazon EC2 instances behind an Elastic Load Balancing (ELB) load balancer. However, many of the web service clients can only reach IP addresses authorized on their firewalls. What should a solutions architect recommend to meet the clients' needs?

- A. A Network Load Balancer with an associated Elastic IP address.
- B. An Application Load Balancer with an associated Elastic IP address.
- C. An A record in an Amazon Route 53 hosted zone pointing to an Elastic IP address.
- D. An EC2 instance with a public IP address running as a proxy in front of the load balancer.

Answer: A

Explanation:

A Network Load Balancer can be assigned one Elastic IP address for each Availability Zone it uses¹. This allows the clients to reach the load balancer using a static IP address that can be authorized on their firewalls. An Application Load Balancer cannot be assigned an Elastic IP address². An A record in an Amazon Route 53 hosted zone pointing to an Elastic IP address would not work because the load balancer would still use its own IP address as the source of the forwarded requests to the web service. An EC2 instance with a public IP address running as a proxy in front of the load balancer would add unnecessary complexity and cost, and would not provide the same scalability and availability as a Network Load Balancer. References: ¹: Network Load Balancers - Elastic Load Balancing³, IP address type section²: How to assign Elastic IP to Application Load Balancer in AWS⁴, answer section.

NEW QUESTION 17

- (Topic 4)

A company offers a food delivery service that is growing rapidly. Because of the growth, the company's order processing system is experiencing scaling problems during peak traffic hours. The current architecture includes the following:

- A group of Amazon EC2 instances that run in an Amazon EC2 Auto Scaling group to collect orders from the application
- Another group of EC2 instances that run in an Amazon EC2 Auto Scaling group to fulfill orders

The order collection process occurs quickly, but the order fulfillment process can take longer. Data must not be lost because of a scaling event.

A solutions architect must ensure that the order collection process and the order fulfillment process can both scale properly during peak traffic hours. The solution must optimize

utilization of the company's AWS resources. Which solution meets these requirements?

- A. Use Amazon CloudWatch metrics to monitor the CPU of each instance in the Auto Scaling group
- B. Configure each Auto Scaling group's minimum capacity according to peak workload values.
- C. Use Amazon CloudWatch metrics to monitor the CPU of each instance in the Auto Scaling group
- D. Configure a CloudWatch alarm to invoke an Amazon Simple Notification Service (Amazon SNS) topic that creates additional Auto Scaling groups on demand.
- E. Provision two Amazon Simple Queue Service (Amazon SQS) queues: one for order collection and another for order fulfillment
- F. Configure the EC2 instances to poll their respective queue
- G. Scale the Auto Scaling groups based on notifications that the queues send.
- H. Provision two Amazon Simple Queue Service (Amazon SQS) queues: one for order collection and another for order fulfillment
- I. Configure the EC2 instances to poll their respective queue
- J. Create a metric based on a backlog per instance calculation
- K. Scale the Auto Scaling groups based on this metric.

Answer: D

Explanation:

The number of instances in your Auto Scaling group can be driven by how long it takes to process a message and the acceptable amount of latency (queue delay). The solution is to use a backlog per instance metric with the target value being the acceptable backlog per instance to maintain.

NEW QUESTION 20

- (Topic 4)

A company has established a new AWS account. The account is newly provisioned and no changes have been made to the default settings. The company is concerned about the security of the AWS account root user.

What should be done to secure the root user?

- A. Create IAM users for daily administrative task
- B. Disable the root user.
- C. Create IAM users for daily administrative task
- D. Enable multi-factor authentication on the root user.
- E. Generate an access key for the root user Use the access key for daily administration tasks instead of the AWS Management Console.
- F. Provide the root user credentials to the most senior solutions architect
- G. Have the solutions architect use the root user for daily administration tasks.

Answer: B

Explanation:

This answer is the most secure and recommended option for securing the root user of a new AWS account. The root user is the identity that has complete access to all AWS services and resources in the account. It is accessed by signing in with the email address and password that were used to create the account. To protect the root user credentials from unauthorized use, AWS advises the following best practices:

- ? Create IAM users for daily administrative tasks. IAM users are identities that you create in your account that have specific permissions to access AWS resources. You can create individual IAM users for yourself and for others who need access to your account. You can also assign IAM users to IAM groups that have a set of policies that grant permissions to perform common tasks. By using IAM users instead of the root user, you can follow the principle of least privilege and reduce the risk of compromising your account.
- ? Enable multi-factor authentication (MFA) on the root user. MFA is a security feature that requires users to prove their identity by providing two pieces of information: their password and a code from a device that only they have access to. By enabling MFA on the root user, you can add an extra layer of protection to your account and prevent unauthorized access even if your password is compromised.
- ? Limit the tasks you perform with the root user account. You should use the root user only for tasks that require root user credentials, such as changing your account settings, closing your account, or managing consolidated billing. For a complete list of tasks that require root user credentials, see Tasks that require root user credentials. For all other tasks, you should use IAM users or roles that have the appropriate permissions.

References:

- ? AWS account root user
- ? Root user best practices for your AWS account
- ? Tasks that require root user credentials

NEW QUESTION 21

- (Topic 4)

A manufacturing company runs its report generation application on AWS. The application generates each report in about 20 minutes. The application is built as a monolith that runs on a single Amazon EC2 instance. The application requires frequent updates to its tightly coupled modules. The application becomes complex to maintain as the company adds new features.

Each time the company patches a software module, the application experiences downtime. Report generation must restart from the beginning after any interruptions. The company wants to redesign the application so that the application can be flexible, scalable, and gradually improved. The company wants to minimize application downtime.

Which solution will meet these requirements?

- A. Run the application on AWS Lambda as a single function with maximum provisioned concurrency.
- B. Run the application on Amazon EC2 Spot Instances as microservices with a Spot Fleet default allocation strategy.
- C. Run the application on Amazon Elastic Container Service (Amazon ECS) as microservices with service auto scaling.
- D. Run the application on AWS Elastic Beanstalk as a single application environment with an all-at-once deployment strategy.

Answer: C

Explanation:

The solution that will meet the requirements is to run the application on Amazon Elastic Container Service (Amazon ECS) as microservices with service auto scaling. This solution will allow the application to be flexible, scalable, and gradually improved, as well as minimize application downtime. By breaking down the monolithic application into microservices, the company can decouple the modules and update them independently, without affecting the whole application. By running the microservices on Amazon ECS, the company can leverage the benefits of containerization, such as portability, efficiency, and isolation. By enabling service auto scaling, the company can adjust the number of containers running for each microservice based on demand, ensuring optimal performance and cost. Amazon ECS also supports various deployment strategies, such as rolling update or blue/green deployment, that can reduce or eliminate downtime during updates.

The other solutions are not as effective as the first one because they either do not meet the requirements or introduce new challenges. Running the application on AWS Lambda as a single function with maximum provisioned concurrency will not meet the requirements, as it will not break down the monolith into microservices, nor will it reduce the complexity of maintenance. Lambda functions are also limited by execution time (15 minutes), memory size (10 GB), and concurrency quotas, which may not be sufficient for the report generation application. Running the application on Amazon EC2 Spot Instances as microservices with a Spot Fleet default allocation strategy will not meet the requirements, as it will introduce the risk of interruptions due to spot price fluctuations. Spot Instances are not guaranteed to be available or stable, and may be reclaimed by AWS at any time with a two-minute warning. This may cause report generation to fail or restart from scratch. Running the application on AWS Elastic Beanstalk as a single application environment with an all-at-once deployment strategy will not meet the requirements, as it will not break down the monolith into microservices, nor will it minimize application downtime. The all-at-once deployment strategy will deploy updates to all instances simultaneously, causing a brief outage for the application.

References:

- ? Amazon Elastic Container Service
- ? Microservices on AWS
- ? Service Auto Scaling - Amazon Elastic Container Service
- ? AWS Lambda
- ? Amazon EC2 Spot Instances
- ? [AWS Elastic Beanstalk]

NEW QUESTION 22

- (Topic 4)

A company wants to move from many standalone AWS accounts to a consolidated, multi-account architecture. The company plans to create many new AWS accounts for different business units. The company needs to authenticate access to these AWS accounts by using a centralized corporate directory service. Which combination of actions should a solutions architect recommend to meet these requirements? (Select TWO.)

- A. Create a new organization in AWS Organizations with all features turned on
- B. Create the new AWS accounts in the organization.
- C. Set up an Amazon Cognito identity pool

- D. Configure AWS IAM Identity Center (AWS Single Sign-On) to accept Amazon Cognito authentication.
- E. Configure a service control policy (SCP) to manage the AWS account
- F. Add AWS IAM Identity Center (AWS Single Sign-On) to AWS Directory Service.
- G. Create a new organization in AWS Organization
- H. Configure the organization's authentication mechanism to use AWS Directory Service directly.
- I. Set up AWS IAM Identity Center (AWS Single Sign-On) in the organization
- J. Configure IAM Identity Center, and integrate it with the company's corporate directory service.

Answer: AE

Explanation:

AWS Organizations is a service that helps users centrally manage and govern multiple AWS accounts. It allows users to create organizational units (OUs) to group accounts based on business needs or other criteria. It also allows users to define and attach service control policies (SCPs) to OUs or accounts to restrict the actions that can be performed by the accounts¹. By creating a new organization in AWS Organizations with all features turned on, the solution can consolidate and manage the new AWS accounts for different business units.

AWS IAM Identity Center (formerly known as AWS Single Sign-On) is a service that provides single sign-on access for all of your AWS accounts and cloud applications. It connects with Microsoft Active Directory through AWS Directory Service to allow users in that directory to sign in to a personalized AWS access portal using their existing Active Directory user names and passwords. From the AWS access portal, users have access to all the AWS accounts and cloud applications that they have permissions for². By setting up IAM Identity Center in the organization and integrating it with the company's corporate directory service, the solution can authenticate access to these AWS accounts using a centralized corporate directory service.

* B. Set up an Amazon Cognito identity pool. Configure AWS IAM Identity Center (AWS Single Sign-On) to accept Amazon Cognito authentication. This solution will not meet the requirement of authenticating access to these AWS accounts by using a centralized corporate directory service, as Amazon Cognito is a service that provides user sign-up, sign-in, and access control for web and mobile applications, not for corporate directory services³.

* C. Configure a service control policy (SCP) to manage the AWS accounts. Add AWS IAM Identity Center (AWS Single Sign-On) to AWS Directory Service. This solution will not work, as SCPs are used to restrict the actions that can be performed by the accounts in an organization, not to manage the accounts themselves¹. Also, IAM Identity Center cannot be added to AWS Directory Service, as it is a separate service that connects with Microsoft Active Directory through AWS Directory Service².

* D. Create a new organization in AWS Organizations. Configure the organization's authentication mechanism to use AWS Directory Service directly. This solution will not work, as AWS Organizations does not have an authentication mechanism that can use AWS Directory Service directly. AWS Organizations relies on IAM Identity Center to provide single sign-on access for the accounts in an organization.

Reference URL: https://docs.aws.amazon.com/organizations/latest/userguide/orgs_integrate_services.html

NEW QUESTION 27

- (Topic 4)

An online video game company must maintain ultra-low latency for its game servers. The game servers run on Amazon EC2 instances. The company needs a solution that can

handle millions of UDP internet traffic requests each second.

Which solution will meet these requirements MOST cost-effectively?

- A. Configure an Application Load Balancer with the required protocol and ports for the internet traffic
- B. Specify the EC2 instances as the targets.
- C. Configure a Gateway Load Balancer for the internet traffic
- D. Specify the EC2 instances as the targets.
- E. Configure a Network Load Balancer with the required protocol and ports for the internet traffic
- F. Specify the EC2 instances as the targets.
- G. Launch an identical set of game servers on EC2 instances in separate AWS Region
- H. Route internet traffic to both sets of EC2 instances.

Answer: C

Explanation:

The most cost-effective solution for the online video game company is to configure a Network Load Balancer with the required protocol and ports for the internet traffic and specify the EC2 instances as the targets. This solution will enable the company to handle millions of UDP requests per second with ultra-low latency and high performance. A Network Load Balancer is a type of Elastic Load Balancing that operates at the connection level (Layer 4) and routes traffic to targets (EC2 instances, microservices, or containers) within Amazon VPC based on IP protocol data. A Network Load Balancer is ideal for load balancing of both TCP and UDP traffic, as it is capable of handling millions of requests per second while maintaining high throughput at ultra-low latency. A Network Load Balancer also preserves the source IP address of the clients to the back-end applications, which can be useful for logging or security purposes¹.

NEW QUESTION 32

- (Topic 4)

A company is using an Application Load Balancer (ALB) to present its application to the internet. The company finds abnormal traffic access patterns across the application. A solutions architect needs to improve visibility into the infrastructure to help the company understand these abnormalities better.

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

- A. Create a table in Amazon Athena for AWS CloudTrail log
- B. Create a query for the relevant information.
- C. Enable ALB access logging to Amazon S3. Create a table in Amazon Athena, and query the logs.
- D. Enable ALB access logging to Amazon S3. Open each file in a text editor, and search each line for the relevant information
- E. Use Amazon EMR on a dedicated Amazon EC2 instance to directly query the ALB to acquire traffic access log information.

Answer: B

Explanation:

This solution meets the requirements because it allows the company to improve visibility into the infrastructure by using ALB access logging and Amazon Athena. ALB access logging is a feature that captures detailed information about requests sent to the load balancer, such as the client's IP address, request path, response code, and latency. By enabling ALB access logging to Amazon S3, the company can store the access logs in an S3 bucket as compressed files. Amazon Athena is an interactive query service that makes it easy to analyze data in Amazon S3 using standard SQL. By creating a table in Amazon Athena for the access logs, the company can query the logs and get results in seconds. This way, the company can better understand the abnormal traffic access patterns across the application.

References:

? Access logs for your Application Load Balancer

? Querying Application Load Balancer Logs

NEW QUESTION 37

- (Topic 4)

A company uses Amazon EC2 instances to host its internal systems. As part of a deployment operation, an administrator tries to use the AWS CLI to terminate an EC2 instance. However, the administrator receives a 403 (Access Denied) error message.

The administrator is using an IAM role that has the following IAM policy attached:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": ["ec2:TerminateInstances"],
      "Resource": ["*"]
    },
    {
      "Effect": "Deny",
      "Action": ["ec2:TerminateInstances"],
      "Condition": {
        "NotIpAddress": {
          "aws:SourceIp": [
            "192.0.2.0/24",
            "203.0.113.0/24"
          ]
        }
      }
    },
    {
      "Effect": "Deny",
      "Action": ["ec2:TerminateInstances"],
      "Resource": ["*"]
    }
  ]
}
```

What is the cause of the unsuccessful request?

- A. The EC2 instance has a resource-based policy with a Deny statement.
- B. The principal has not been specified in the policy statement
- C. The "Action" field does not grant the actions that are required to terminate the EC2 instance.
- D. The request to terminate the EC2 instance does not originate from the CIDR blocks 192.0.2.0/24 or 203.0.113.0/24

Answer: D

NEW QUESTION 39

- (Topic 4)

A company wants to migrate its three-tier application from on premises to AWS. The web tier and the application tier are running on third-party virtual machines (VMs). The database tier is running on MySQL.

The company needs to migrate the application by making the fewest possible changes to the architecture. The company also needs a database solution that can restore data to a specific point in time.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Migrate the web tier and the application tier to Amazon EC2 instances in private subnet
- B. Migrate the database tier to Amazon RDS for MySQL in private subnets.
- C. Migrate the web tier to Amazon EC2 instances in public subnet
- D. Migrate the application tier to EC2 instances in private subnet
- E. Migrate the database tier to Amazon Aurora MySQL in private subnets.
- F. Migrate the web tier to Amazon EC2 instances in public subnet
- G. Migrate the application tier to EC2 instances in private subnet
- H. Migrate the database tier to Amazon RDS for MySQL in private subnets.
- I. Migrate the web tier and the application tier to Amazon EC2 instances in public subnet
- J. Migrate the database tier to Amazon Aurora MySQL in public subnets.

Answer: C

Explanation:

The solution that meets the requirements with the least operational overhead is to migrate the web tier to Amazon EC2 instances in public subnets, migrate the application tier to EC2 instances in private subnets, and migrate the database tier to Amazon RDS for MySQL in private subnets. This solution allows the company to migrate its three-tier application to AWS by making minimal changes to the architecture, as it preserves the same web, application, and database tiers and uses the same MySQL database engine. The solution also provides a database solution that can restore data to a specific point in time, as Amazon RDS for MySQL supports automated backups and point-in-time recovery. This solution also reduces the operational overhead by using managed services such as Amazon EC2 and Amazon RDS, which handle tasks such as provisioning, patching, scaling, and monitoring.

The other solutions do not meet the requirements as well as the first one because they either involve more changes to the architecture, do not provide point-in-time recovery, or do not follow best practices for security and availability. Migrating the database tier to Amazon Aurora MySQL would require changing the database

engine and potentially modifying the application code to ensure compatibility. Migrating the web tier and the application tier to public subnets would expose them to more security risks and reduce their availability in case of a subnet failure. Migrating the database tier to public subnets would also compromise its security and performance. References:

- ? Migrate Your Application Database to Amazon RDS
- ? Amazon RDS for MySQL
- ? Amazon Aurora MySQL
- ? Amazon VPC

NEW QUESTION 40

- (Topic 4)

A solutions architect is implementing a document review application using an Amazon S3 bucket for storage. The solution must prevent accidental deletion of the documents and ensure that all versions of the documents are available. Users must be able to download, modify, and upload documents.

Which combination of actions should be taken to meet these requirements? (Choose two.)

- A. Enable a read-only bucket ACL.
- B. Enable versioning on the bucket.
- C. Attach an IAM policy to the bucket.
- D. Enable MFA Delete on the bucket.
- E. Encrypt the bucket using AWS KMS.

Answer: BD

Explanation:

Versioning is a feature of Amazon S3 that allows users to keep multiple versions of the same object in a bucket. It can help prevent accidental deletion of the documents and ensure that all versions of the documents are available¹. MFA Delete is a feature of Amazon S3 that adds an extra layer of security by requiring two forms of authentication to delete a version or change the versioning state of a bucket. It can help prevent unauthorized or accidental deletion of the documents². By enabling both versioning and MFA Delete on the bucket, the solution can meet the requirements.

- * A. Enable a read-only bucket ACL. This solution will not meet the requirement of allowing users to download, modify, and upload documents, as a read-only bucket ACL will prevent write access to the bucket³.
 - * C. Attach an IAM policy to the bucket. This solution will not meet the requirement of preventing accidental deletion of the documents and ensuring that all versions of the documents are available, as an IAM policy is used to grant or deny permissions to users or roles, not to enable versioning or MFA Delete⁴.
 - * E. Encrypt the bucket using AWS KMS. This solution will not meet the requirement of preventing accidental deletion of the documents and ensuring that all versions of the documents are available, as encrypting the bucket using AWS KMS is a method of protecting data at rest, not enabling versioning or MFA Delete.
- Reference URL: <https://docs.aws.amazon.com/AmazonS3/latest/userguide/Versioning.html>

NEW QUESTION 42

- (Topic 4)

A company maintains about 300 TB in Amazon S3 Standard storage month after month. The S3 objects are each typically around 50 GB in size and are frequently replaced with multipart uploads by their global application. The number and size of S3 objects remain constant but the company's S3 storage costs are increasing each month.

How should a solutions architect reduce costs in this situation?

- A. Switch from multipart uploads to Amazon S3 Transfer Acceleration.
- B. Enable an S3 Lifecycle policy that deletes incomplete multipart uploads.
- C. Configure S3 inventory to prevent objects from being archived too quickly.
- D. Configure Amazon CloudFront to reduce the number of objects stored in Amazon S3.

Answer: B

Explanation:

This option is the most cost-effective way to reduce the S3 storage costs in this situation. Incomplete multipart uploads are parts of objects that are not completed or aborted by the application. They consume storage space and incur charges until they are deleted. By enabling an S3 Lifecycle policy that deletes incomplete multipart uploads, you can automatically remove them after a specified period of time (such as one day) and free up the storage space. This will reduce the S3 storage costs and also improve the performance of the application by avoiding unnecessary retries or errors.

Option A is not correct because switching from multipart uploads to Amazon S3 Transfer Acceleration will not reduce the S3 storage costs. Amazon S3 Transfer Acceleration is a feature that enables faster data transfers to and from S3 by using the AWS edge network. It is useful for improving the upload speed of large objects over long distances, but it does not affect the storage space or charges. In fact, it may increase the costs by adding a data transfer fee for using the feature.

Option C is not correct because configuring S3 inventory to prevent objects from being archived too quickly will not reduce the S3 storage costs. Amazon S3 Inventory is a feature that provides a report of the objects and their metadata in an S3 bucket. It is useful for managing and auditing the S3 objects, but it does not affect the storage space or charges. In fact, it may increase the costs by generating additional S3 objects for the inventory reports.

Option D is not correct because configuring Amazon CloudFront to reduce the number of objects stored in Amazon S3 will not reduce the S3 storage costs. Amazon CloudFront is a content delivery network (CDN) that distributes the S3 objects to edge locations for faster and lower latency access. It is useful for improving the download speed and availability of the S3 objects, but it does not affect the storage space or charges. In fact, it may increase the costs by adding a data transfer fee for using the service. References:

- ? Managing your storage lifecycle
- ? Using multipart upload
- ? Amazon S3 Transfer Acceleration
- ? Amazon S3 Inventory
- ? What Is Amazon CloudFront?

NEW QUESTION 46

- (Topic 4)

A company runs an SMB file server in its data center. The file server stores large files that the company frequently accesses for up to 7 days after the file creation date. After 7 days, the company needs to be able to access the files with a maximum retrieval time of 24 hours.

Which solution will meet these requirements?

- A. Use AWS DataSync to copy data that is older than 7 days from the SMB file server to AWS.
- B. Create an Amazon S3 File Gateway to increase the company's storage space.
- C. Create an S3 Lifecycle policy to transition the data to S3 Glacier Deep Archive after 7 days.

- D. Create an Amazon FSx File Gateway to increase the company's storage space.
- E. Create an Amazon S3 Lifecycle policy to transition the data after 7 days.
- F. Configure access to Amazon S3 for each use.
- G. Create an S3 Lifecycle policy to transition the data to S3 Glacier Flexible Retrieval after 7 days.

Answer: B

Explanation:

Amazon S3 File Gateway is a service that provides a file-based interface to Amazon S3, which appears as a network file share. It enables you to store and retrieve Amazon S3 objects through standard file storage protocols such as SMB. S3 File Gateway can also cache frequently accessed data locally for low-latency access. S3 Lifecycle policy is a feature that allows you to define rules that automate the management of your objects throughout their lifecycle. You can use S3 Lifecycle policy to transition objects to different storage classes based on their age and access patterns. S3 Glacier Deep Archive is a storage class that offers the lowest cost for long-term data archiving, with a retrieval time of 12 hours or 48 hours. This solution will meet the requirements, as it allows the company to store large files in S3 with SMB file access, and to move the files to S3 Glacier Deep Archive after 7 days for cost savings and compliance.

References:

- ? 1 provides an overview of Amazon S3 File Gateway and its benefits.
- ? 2 explains how to use S3 Lifecycle policy to manage object storage lifecycle.
- ? 3 describes the features and use cases of S3 Glacier Deep Archive storage class.

NEW QUESTION 50

- (Topic 4)

A company is deploying an application that processes large quantities of data in parallel. The company plans to use Amazon EC2 instances for the workload. The network architecture must be configurable to prevent groups of nodes from sharing the same underlying hardware. Which networking solution meets these requirements?

- A. Run the EC2 instances in a spread placement group.
- B. Group the EC2 instances in separate accounts.
- C. Configure the EC2 instances with dedicated tenancy.
- D. Configure the EC2 instances with shared tenancy.

Answer: A

Explanation:

It allows the company to deploy an application that processes large quantities of data in parallel and prevent groups of nodes from sharing the same underlying hardware. By running the EC2 instances in a spread placement group, the company can launch a small number of instances across distinct underlying hardware to reduce correlated failures. A spread placement group ensures that each instance is isolated from each other at the rack level. References:

- ? Placement Groups
- ? Spread Placement Groups

NEW QUESTION 53

- (Topic 4)

A company is building an Amazon Elastic Kubernetes Service (Amazon EKS) cluster for its workloads. All secrets that are stored in Amazon EKS must be encrypted in the Kubernetes etcd key-value store. Which solution will meet these requirements?

- A. Create a new AWS Key Management Service (AWS KMS) key. Use AWS Secrets Manager to manage rotate, and store all secrets in Amazon EKS.
- B. Create a new AWS Key Management Service (AWS KMS) key. Enable Amazon EKS KMS secrets encryption on the Amazon EKS cluster.
- C. Create the Amazon EKS cluster with default options. Use the Amazon Elastic Block Store (Amazon EBS) Container Storage Interface (CSI) driver as an add-on.
- D. Create a new AWS Key Management Service (AWS KMS) key with the alias `aws/ebs` and enable default Amazon Elastic Block Store (Amazon EBS) volume encryption for the account.

Answer: B

Explanation:

This option is the most secure and simple way to encrypt the secrets that are stored in Amazon EKS. AWS Key Management Service (AWS KMS) is a service that allows you to create and manage encryption keys that can be used to encrypt your data. Amazon EKS KMS secrets encryption is a feature that enables you to use a KMS key to encrypt the secrets that are stored in the Kubernetes etcd key-value store. This provides an additional layer of protection for your sensitive data, such as passwords, tokens, and keys. You can create a new KMS key or use an existing one, and then enable the Amazon EKS KMS secrets encryption on the Amazon EKS cluster. You can also use IAM policies to control who can access or use the KMS key.

Option A is not correct because using AWS Secrets Manager to manage, rotate, and store all secrets in Amazon EKS is not necessary or efficient. AWS Secrets Manager is a service that helps you securely store, retrieve, and rotate your secrets, such as database credentials, API keys, and passwords. You can use it to manage secrets that are used by your applications or services outside of Amazon EKS, but it is not designed to encrypt the secrets that are stored in the Kubernetes etcd key-value store. Moreover, using AWS Secrets Manager would incur additional costs and complexity, and it would not leverage the native Kubernetes secrets management capabilities.

Option C is not correct because using the Amazon EBS Container Storage Interface (CSI) driver as an add-on does not encrypt the secrets that are stored in Amazon EKS. The Amazon EBS CSI driver is a plugin that allows you to use Amazon EBS volumes as persistent storage for your Kubernetes pods. It is useful for providing durable and scalable storage for your applications, but it does not affect the encryption of the secrets that are stored in the Kubernetes etcd key-value store. Moreover, using the Amazon EBS CSI driver would require additional configuration and resources, and it would not provide the same level of security as using a KMS key.

Option D is not correct because creating a new AWS KMS key with the alias `aws/ebs` and enabling default Amazon EBS volume encryption for the account does not encrypt the secrets that are stored in Amazon EKS. The alias `aws/ebs` is a reserved alias that is used by AWS to create a default KMS key for your account. This key is used to encrypt the Amazon EBS volumes that are created in your account, unless you specify a different KMS key. Enabling default Amazon EBS volume encryption for the account is a setting that ensures that all new Amazon EBS volumes are encrypted by default. However, these features do not affect the encryption of the secrets that are stored in the Kubernetes etcd key-value store. Moreover, using the default KMS key or the default encryption setting would not provide the same level of control and security as using a custom KMS key and enabling the Amazon EKS KMS secrets encryption feature. References:

- ? Encrypting secrets used in Amazon EKS
- ? What Is AWS Key Management Service?
- ? What Is AWS Secrets Manager?
- ? Amazon EBS CSI driver
- ? Encryption at rest

NEW QUESTION 56

- (Topic 4)

A company stores multiple Amazon Machine Images (AMIs) in an AWS account to launch its Amazon EC2 instances. The AMIs contain critical data and configurations that are necessary for the company's operations. The company wants to implement a solution that will recover accidentally deleted AMIs quickly and efficiently.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Create Amazon Elastic Block Store (Amazon EBS) snapshots of the AMI
- B. Store the snapshots in a separate AWS account.
- C. Copy all AMIs to another AWS account periodically.
- D. Create a retention rule in Recycle Bin.
- E. Upload the AMIs to an Amazon S3 bucket that has Cross-Region Replication.

Answer: C

Explanation:

Recycle Bin is a data recovery feature that enables you to restore accidentally deleted Amazon EBS snapshots and EBS-backed AMIs. When using Recycle Bin, if your resources are deleted, they are retained in the Recycle Bin for a time period that you specify before being permanently deleted. You can restore a resource from the Recycle Bin at any time before its retention period expires. This solution has the least operational overhead, as you do not need to create, copy, or upload any additional resources. You can also manage tags and permissions for AMIs in the Recycle Bin. AMIs in the Recycle Bin do not incur any additional charges.

References:

- ? Recover AMIs from the Recycle Bin
- ? Recover an accidentally deleted Linux AMI

NEW QUESTION 57

- (Topic 4)

A solutions architect is designing a new API using Amazon API Gateway that will receive requests from users. The volume of requests is highly variable; several hours can pass without receiving a single request. The data processing will take place asynchronously, but should be completed within a few seconds after a request is made.

Which compute service should the solutions architect have the API invoke to deliver the requirements at the lowest cost?

- A. An AWS Glue job
- B. An AWS Lambda function
- C. A containerized service hosted in Amazon Elastic Kubernetes Service (Amazon EKS)
- D. A containerized service hosted in Amazon ECS with Amazon EC2

Answer: B

Explanation:

API Gateway + Lambda is the perfect solution for modern applications with serverless architecture.

NEW QUESTION 61

- (Topic 4)

A company operates a two-tier application for image processing. The application uses two Availability Zones, each with one public subnet and one private subnet. An Application Load Balancer (ALB) for the web tier uses the public subnets. Amazon EC2 instances for the application tier use the private subnets.

Users report that the application is running more slowly than expected. A security audit of the web server log files shows that the application is receiving millions of illegitimate requests from a small number of IP addresses. A solutions architect needs to resolve the immediate performance problem while the company investigates a more permanent solution.

What should the solutions architect recommend to meet this requirement?

- A. Modify the inbound security group for the web tier
- B. Add a deny rule for the IP addresses that are consuming resources.
- C. Modify the network ACL for the web tier subnet
- D. Add an inbound deny rule for the IP addresses that are consuming resources
- E. Modify the inbound security group for the application tier
- F. Add a deny rule for the IP addresses that are consuming resources.
- G. Modify the network ACL for the application tier subnet
- H. Add an inbound deny rule for the IP addresses that are consuming resources

Answer: B

Explanation:

Deny the request from the first entry at the public subnet, don't allow it to cross and get to the private subnet.

In this scenario, the security audit reveals that the application is receiving millions of illegitimate requests from a small number of IP addresses. To address this issue, it is recommended to modify the network ACL (Access Control List) for the web tier subnets. By adding an inbound deny rule specifically targeting the IP addresses that are consuming resources, the network ACL can block the illegitimate traffic at the subnet level before it reaches the web servers. This will help alleviate the excessive load on the web tier and improve the application's performance.

NEW QUESTION 64

- (Topic 4)

A company runs a web application that is deployed on Amazon EC2 instances in the private subnet of a VPC. An Application Load Balancer (ALB) that extends across the public subnets directs web traffic to the EC2 instances. The company wants to implement new security measures to restrict inbound traffic from the ALB to the EC2 instances while preventing access from any other source inside or outside the private subnet of the EC2 instances.

Which solution will meet these requirements?

- A. Configure a route in a route table to direct traffic from the internet to the private IP addresses of the EC2 instances.
- B. Configure the security group for the EC2 instances to only allow traffic that comes from the security group for the ALB.
- C. Move the EC2 instances into the public subnet
- D. Give the EC2 instances a set of Elastic IP addresses.
- E. Configure the security group for the ALB to allow any TCP traffic on any port.

Answer: B

Explanation:

To restrict inbound traffic from the ALB to the EC2 instances, the security group for the EC2 instances should only allow traffic that comes from the security group for the ALB. This way, the EC2 instances can only receive requests from the ALB and not from any other source inside or outside the private subnet.

References:

- ? Security Groups for Your Application Load Balancers
- ? Security Groups for Your VPC

NEW QUESTION 67

- (Topic 4)

A company runs demonstration environments for its customers on Amazon EC2 instances. Each environment is isolated in its own VPC. The company's operations team needs to be notified when RDP or SSH access to an environment has been established.

- A. Configure Amazon CloudWatch Application Insights to create AWS Systems Manager OpsItems when RDP or SSH access is detected.
- B. Configure the EC2 instances with an IAM instance profile that has an IAM role with the AmazonSSMManagedInstanceCore policy attached.
- C. Publish VPC flow logs to Amazon CloudWatch Log
- D. Create required metric filter
- E. Create an Amazon CloudWatch metric alarm with a notification action for when the alarm is in the ALARM state.
- F. Configure an Amazon EventBridge rule to listen for events of type EC2 Instance State- change Notificatio
- G. Configure an Amazon Simple Notification Service (Amazon SNS) topic as a target
- H. Subscribe the operations team to the topic.

Answer: C

Explanation:

<https://aws.amazon.com/blogs/security/how-to-monitor-and-visualize-failed-ssh-access-attempts-to-amazon-ec2-linux-instances/>

NEW QUESTION 71

- (Topic 4)

A company is concerned that two NAT instances in use will no longer be able to support the traffic needed for the company's application. A solutions architect wants to implement a solution that is highly available, fault tolerant, and automatically scalable. What should the solutions architect recommend?

- A. Remove the two NAT instances and replace them with two NAT gateways in the same Availability Zone.
- B. Use Auto Scaling groups with Network Load Balancers for the NAT instances in different Availability Zones.
- C. Remove the two NAT instances and replace them with two NAT gateways in different Availability Zones.
- D. Replace the two NAT instances with Spot Instances in different Availability Zones and deploy a Network Load Balancer.

Answer: C

Explanation:

If you have resources in multiple Availability Zones and they share one NAT gateway, and if the NAT gateway's Availability Zone is down, resources in the other Availability Zones lose internet access. To create an Availability Zone-independent architecture, create a NAT gateway in each Availability Zone and configure your routing to ensure that resources use the NAT gateway in the same Availability Zone. <https://docs.aws.amazon.com/vpc/latest/userguide/vpc-nat-gateway.html#nat-gateway-basics>

NEW QUESTION 72

- (Topic 4)

A company is conducting an internal audit. The company wants to ensure that the data in an Amazon S3 bucket that is associated with the company's AWS Lake Formation data lake does not contain sensitive customer or employee data. The company wants to discover personally identifiable information (PII) or financial information, including passport numbers and credit card numbers.

Which solution will meet these requirements?

- A. Configure AWS Audit Manager on the account
- B. Select the Payment Card Industry Data Security Standards (PCI DSS) for auditing.
- C. Configure Amazon S3 Inventory on the S3 bucket
- D. Configure Amazon Athena to query the inventory.
- E. Configure Amazon Macie to run a data discovery job that uses managed identifiers for the required data types.
- F. Use Amazon S3 Select to run a report across the S3 bucket.

Answer: C

Explanation:

Amazon Macie is a fully managed data security and data privacy service that uses machine learning and pattern matching to discover and protect your sensitive data in AWS. Macie can run data discovery jobs that use managed identifiers for various types of PII or financial information, such as passport numbers and credit card numbers. Macie can also generate findings that alert you to potential issues or risks with your data. References:

<https://docs.aws.amazon.com/macie/latest/userguide/macie-identifiers.html>

NEW QUESTION 77

- (Topic 4)

A company wants to host a scalable web application on AWS. The application will be accessed by users from different geographic regions of the world. Application users will be able to download and upload unique data up to gigabytes in size. The development team wants a cost-effective solution to minimize upload and download latency and maximize performance.

What should a solutions architect do to accomplish this?

- A. Use Amazon S3 with Transfer Acceleration to host the application.
- B. Use Amazon S3 with CacheControl headers to host the application.
- C. Use Amazon EC2 with Auto Scaling and Amazon CloudFront to host the application.

D. Use Amazon EC2 with Auto Scaling and Amazon ElastiCache to host the application.

Answer: C

Explanation:

This answer is correct because it meets the requirements of hosting a scalable web application that can handle large data transfers from different geographic regions. Amazon EC2 provides scalable compute capacity for hosting web applications. Auto Scaling can automatically adjust the number of EC2 instances based on the demand and traffic patterns. Amazon CloudFront is a content delivery network (CDN) that can cache static and dynamic content at edge locations closer to the users, reducing latency and improving performance. CloudFront can also use S3 Transfer Acceleration to speed up the transfers between S3 buckets and CloudFront edge locations.

References:

? <https://docs.aws.amazon.com/autoscaling/ec2/userguide/what-is-amazon-ec2-auto-scaling.html>

? <https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/Introduction.html>

? <https://aws.amazon.com/s3/transfer-acceleration/>

NEW QUESTION 82

- (Topic 4)

A retail company has several businesses. The IT team for each business manages its own AWS account. Each team account is part of an organization in AWS Organizations. Each team monitors its product inventory levels in an Amazon DynamoDB table in the team's own AWS account.

The company is deploying a central inventory reporting application into a shared AWS account. The application must be able to read items from all the teams' DynamoDB tables.

Which authentication option will meet these requirements MOST securely?

- A. Integrate DynamoDB with AWS Secrets Manager in the inventory application account
- B. Configure the application to use the correct secret from Secrets Manager to authenticate and read the DynamoDB table
- C. Schedule secret rotation for every 30 days.
- D. In every business account, create an IAM user that has programmatic access
- E. Configure the application to use the correct IAM user access key ID and secret access key to authenticate and read the DynamoDB table
- F. Manually rotate IAM access keys every 30 days.
- G. In every business account, create an IAM role named BU_ROLE with a policy that gives the role access to the DynamoDB table and a trust policy to trust a specific role in the inventory application account
- H. In the inventory account, create a role named APP_ROLE that allows access to the STS AssumeRole API operation
- I. Configure the application to use APP_ROLE and assume the cross-account role BU_ROLE to read the DynamoDB table.
- J. Integrate DynamoDB with AWS Certificate Manager (ACM). Generate identity certificates to authenticate DynamoDB
- K. Configure the application to use the correct certificate to authenticate and read the DynamoDB table.

Answer: C

Explanation:

This solution meets the requirements most securely because it uses IAM roles and the STS AssumeRole API operation to authenticate and authorize the inventory application to access the DynamoDB tables in different accounts. IAM roles are more secure than IAM users or certificates because they do not require long-term credentials or passwords. Instead, IAM roles provide temporary security credentials that are automatically rotated and can be configured with a limited duration. The STS AssumeRole API operation enables you to request temporary credentials for a role that you are allowed to assume. By using this operation, you can delegate access to resources that are in different AWS accounts that you own or that are owned by third parties. The trust policy of the role defines which entities can assume the role, and the permissions policy of the role defines which actions can be performed on the resources. By using this solution, you can avoid hard-coding credentials or certificates in the inventory application, and you can also avoid storing them in Secrets Manager or ACM. You can also leverage the built-in security features of IAM and STS, such as MFA, access logging, and policy conditions.

References:

? IAM Roles

? STS AssumeRole

? Tutorial: Delegate Access Across AWS Accounts Using IAM Roles

NEW QUESTION 87

- (Topic 4)

A company wants to use high-performance computing and artificial intelligence to improve its fraud prevention and detection technology. The company requires distributed processing to complete a single workload as quickly as possible.

Which solution will meet these requirements?

- A. Use Amazon Elastic Kubernetes Service (Amazon EKS) and multiple containers.
- B. Use AWS ParallelCluster and the Message Passing Interface (MPI) libraries.
- C. Use an Application Load Balancer and Amazon EC2 instances.
- D. Use AWS Lambda functions.

Answer: B

Explanation:

AWS ParallelCluster is a service that allows you to create and manage high-performance computing (HPC) clusters on AWS. It supports multiple schedulers, including AWS Batch, which can run distributed workloads across multiple EC2 instances¹.

MPI is a standard for message passing between processes in parallel computing. It provides functions for sending and receiving data, synchronizing processes, and managing communication groups².

By using AWS ParallelCluster and MPI libraries, you can take advantage of the following benefits:

? You can easily create and configure HPC clusters that meet your specific requirements, such as instance type, number of nodes, network configuration, and storage options¹.

? You can leverage the scalability and elasticity of AWS to run large-scale parallel workloads without worrying about provisioning or managing servers¹.

? You can use MPI libraries to optimize the performance and efficiency of your parallel applications by enabling inter-process communication and data exchange².

? You can choose from a variety of MPI implementations that are compatible with AWS ParallelCluster, such as Open MPI, Intel MPI, and MPICH3.

NEW QUESTION 89

- (Topic 4)

A company operates an ecommerce website on Amazon EC2 instances behind an Application Load Balancer (ALB) in an Auto Scaling group. The site is experiencing performance issues related to a high request rate from illegitimate external systems with changing IP addresses. The security team is worried about potential DDoS attacks against the website. The company must block the illegitimate incoming requests in a way that has a minimal impact on legitimate users. What should a solutions architect recommend?

- A. Deploy Amazon Inspector and associate it with the ALB.
- B. Deploy AWS WAF, associate it with the ALB, and configure a rate-limiting rule.
- C. Deploy rules to the network ACLs associated with the ALB to block the incoming traffic.
- D. Deploy Amazon GuardDuty and enable rate-limiting protection when configuring GuardDuty.

Answer: B

Explanation:

This answer is correct because it meets the requirements of blocking the illegitimate incoming requests in a way that has a minimal impact on legitimate users. AWS WAF is a web application firewall that helps protect your web applications or APIs against common web exploits that may affect availability, compromise security, or consume excessive resources. AWS WAF gives you control over how traffic reaches your applications by enabling you to create security rules that block common attack patterns, such as SQL injection or cross-site scripting, and rules that filter out specific traffic patterns you define. You can associate AWS WAF with an ALB to protect the web application from malicious requests. You can configure a rate-limiting rule in AWS WAF to track the rate of requests for each originating IP address and block requests from an IP address that exceeds a certain limit within a five-minute period. This way, you can mitigate potential DDoS attacks and improve the performance of your website.

References:

? <https://docs.aws.amazon.com/waf/latest/developerguide/what-is-aws-waf.html>

? <https://docs.aws.amazon.com/waf/latest/developerguide/waf-rule-statement-type-rate-based.html>

NEW QUESTION 94

- (Topic 4)

A company runs a container application by using Amazon Elastic Kubernetes Service (Amazon EKS). The application includes microservices that manage customers and place orders. The company needs to route incoming requests to the appropriate microservices. Which solution will meet this requirement MOST cost-effectively?

- A. Use the AWS Load Balancer Controller to provision a Network Load Balancer.
- B. Use the AWS Load Balancer Controller to provision an Application Load Balancer.
- C. Use an AWS Lambda function to connect the requests to Amazon EKS.
- D. Use Amazon API Gateway to connect the requests to Amazon EKS.

Answer: B

Explanation:

An Application Load Balancer is a type of Elastic Load Balancer that operates at the application layer (layer 7) of the OSI model. It can distribute incoming traffic across multiple targets, such as Amazon EC2 instances, containers, IP addresses, and Lambda functions. It can also route requests based on the content of the request, such as the host name, path, or query parameters¹.

The AWS Load Balancer Controller is a controller that helps you manage Elastic Load Balancers for your Kubernetes cluster. It can provision Application Load Balancers or Network Load Balancers when you create Kubernetes Ingress or Service resources².

By using the AWS Load Balancer Controller to provision an Application Load Balancer for your Amazon EKS cluster, you can achieve the following benefits:

? You can route incoming requests to the appropriate microservices based on the

rules you define in your Ingress resource. For example, you can route requests with different host names or paths to different microservices that handle customers and orders².

? You can improve the performance and availability of your container applications by distributing the load across multiple targets and enabling health checks and automatic scaling¹.

? You can reduce the cost and complexity of managing your load balancers by using

a single controller that integrates with Amazon EKS and Kubernetes. You do not need to manually create or configure load balancers or update them when your cluster changes².

NEW QUESTION 96

- (Topic 4)

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.)

- 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: BC

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.

References: 1: What is AWS Global Accelerator? - AWS Global Accelerator⁴, Overview section²: Network Acceleration Service - AWS Global Accelerator - AWS⁵, Why AWS Global Accelerator? section. 3: What is an Application Load Balancer? - Elastic Load Balancing⁶, Overview section.

NEW QUESTION 100

- (Topic 4)

A company runs multiple workloads in its on-premises data center. The company's data center cannot scale fast enough to meet the company's expanding business needs. The company wants to collect usage and configuration data about the on-premises servers and workloads to plan a migration to AWS. Which solution will meet these requirements?

- A. Set the home AWS Region in AWS Migration Hub
- B. Use AWS Systems Manager to collect data about the on-premises servers.
- C. Set the home AWS Region in AWS Migration Hub
- D. Use AWS Application Discovery Service to collect data about the on-premises servers.
- E. Use the AWS Schema Conversion Tool (AWS SCT) to create the relevant template
- F. Use AWS Trusted Advisor to collect data about the on-premises servers.
- G. Use the AWS Schema Conversion Tool (AWS SCT) to create the relevant templates. Use AWS Database Migration Service (AWS DMS) to collect data about the on-premises servers.

Answer: B

Explanation:

The most suitable solution for the company's requirements is to set the home AWS Region in AWS Migration Hub and use AWS Application Discovery Service to collect data about the on-premises servers. This solution will enable the company to gather usage and configuration data of its on-premises servers and workloads, and plan a migration to AWS.

AWS Migration Hub is a service that simplifies and accelerates migration tracking by aggregating migration status information into a single console. Users can view the discovered servers, group them into applications, and track the migration status of each application from the Migration Hub console in their home Region. The home Region is the AWS Region where users store their migration data, regardless of which Regions they migrate into¹.

AWS Application Discovery Service is a service that helps users plan their migration to AWS by collecting usage and configuration data about their on-premises servers and databases. Application Discovery Service is integrated with AWS Migration Hub and supports two methods of performing discovery: agentless discovery and agent-based discovery. Agentless discovery can be performed by deploying the Application Discovery Service Agentless Collector through VMware vCenter, which collects static configuration data and utilization data for virtual machines (VMs) and databases. Agent-based discovery can be performed by deploying the AWS Application Discovery Agent on each of the VMs and physical servers, which collects static configuration data, detailed time-series system-performance information, inbound and outbound network connections, and processes that are running².

The other options are not correct because they do not meet the requirements or are not relevant for the use case. Using the AWS Schema Conversion Tool (AWS SCT) to create the relevant templates and using AWS Trusted Advisor to collect data about the on-premises servers is not correct because this solution is not suitable for collecting usage and configuration data of on-premises servers and workloads. AWS SCT is a tool that helps users convert database schemas and code objects from one database engine to another, such as from Oracle to PostgreSQL³. AWS Trusted Advisor is a service that provides best practice recommendations for cost optimization, performance, security, fault tolerance, and service limits⁴. Using the AWS Schema Conversion Tool (AWS SCT) to create the relevant templates and using AWS Database Migration Service (AWS DMS) to collect data about the on-premises servers is not correct because this solution is not suitable for collecting usage and configuration data of on-premises servers and workloads. As mentioned above, AWS SCT is a tool that helps users convert database schemas and code objects from one database engine to another. AWS DMS is a service that helps users migrate relational databases, non-relational databases, and other types of data stores to

AWS with minimal downtime⁵. References:

? Home Region - AWS Migration Hub

? What is AWS Application Discovery Service? - AWS Application Discovery Service

? AWS Schema Conversion Tool - Amazon Web Services

? What Is Trusted Advisor? - Trusted Advisor

? What Is AWS Database Migration Service? - AWS Database Migration Service

NEW QUESTION 102

- (Topic 4)

A media company collects and analyzes user activity data on premises. The company wants to migrate this capability to AWS. The user activity data store will continue to grow and will be petabytes in size. The company needs to build a highly available data ingestion solution that facilitates on-demand analytics of existing data and new data with SQL.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Send activity data to an Amazon Kinesis data stream
- B. Configure the stream to deliver the data to an Amazon S3 bucket.
- C. Send activity data to an Amazon Kinesis Data Firehose delivery stream
- D. Configure the stream to deliver the data to an Amazon Redshift cluster.
- E. Place activity data in an Amazon S3 bucket
- F. Configure Amazon S3 to run an AWS Lambda function on the data as the data arrives in the S3 bucket.
- G. Create an ingestion service on Amazon EC2 instances that are spread across multiple Availability Zones
- H. Configure the service to forward data to an Amazon RDS Multi-AZ database.

Answer: B

Explanation:

Amazon Redshift is a fully managed, petabyte-scale data warehouse service in the cloud. You can start with just a few hundred gigabytes of data and scale to a petabyte or more. This allows you to use your data to gain new insights for your business and customers. The first step to create a data warehouse is to launch a set of nodes, called an Amazon Redshift cluster. After you provision your cluster, you can upload your data set and then perform data analysis queries. Regardless of the size of the data set, Amazon Redshift offers fast query performance using the same SQL-based tools and business intelligence applications that you use today.

NEW QUESTION 107

- (Topic 4)

A company sends AWS CloudTrail logs from multiple AWS accounts to an Amazon S3 bucket in a centralized account. The company must keep the CloudTrail logs. The company must also be able to query the CloudTrail logs at any time. Which solution will meet these requirements?

- A. Use the CloudTrail event history in the centralized account to create an Amazon Athena table
- B. Query the CloudTrail logs from Athena.
- C. Configure an Amazon Neptune instance to manage the CloudTrail logs

- D. Query the CloudTrail logs from Neptune.
- E. Configure CloudTrail to send the logs to an Amazon DynamoDB tabl
- F. Create a dashboard in Amazon QuIcKsight to query the logs in the table.
- G. use Amazon Athena to create an Athena noteboo
- H. Configure CloudTrail to send the logs to the noteboo
- I. Run queries from Athena.

Answer: A

Explanation:

it allows the company to keep the CloudTrail logs and query them at any time. By using the CloudTrail event history in the centralized account, the company can view, filter, and download recent API activity across multiple AWS accounts. By creating an Amazon Athena table from the CloudTrail event history, the company can use a serverless interactive query service that makes it easy to analyze data in S3 using standard SQL. By querying the CloudTrail logs from Athena, the company can gain insights into user activity and resource changes. References:

- ? Viewing Events with CloudTrail Event History
- ? Querying AWS CloudTrail Logs
- ? Amazon Athena

NEW QUESTION 109

- (Topic 4)

A company is deploying an application that processes streaming data in near-real time The company plans to use Amazon EC2 instances for the workload The network architecture must be configurable to provide the lowest possible latency between nodes Which combination of network solutions will meet these requirements? (Select TWO)

- A. Enable and configure enhanced networking on each EC2 instance
- B. Group the EC2 instances in separate accounts
- C. Run the EC2 instances in a cluster placement group
- D. Attach multiple elastic network interfaces to each EC2 instance
- E. Use Amazon Elastic Block Store (Amazon EBS) optimized instance types.

Answer: AC

Explanation:

These options are the most suitable ways to configure the network architecture to provide the lowest possible latency between nodes. Option A enables and configures enhanced networking on each EC2 instance, which is a feature that improves the network performance of the instance by providing higher bandwidth, lower latency, and lower jitter. Enhanced networking uses single root I/O virtualization (SR-IOV) or Elastic Fabric Adapter (EFA) to provide direct access to the network hardware. You can enable and configure enhanced networking by choosing a supported instance type and a compatible operating system, and installing the required drivers. Option C runs the EC2 instances in a cluster placement group, which is a logical grouping of instances within a single Availability Zone that are placed close together on the same underlying hardware. Cluster placement groups provide the lowest network latency and the highest network throughput among the placement group options. You can run the EC2 instances in a cluster placement group by creating a placement group and launching the instances into it. Option B is not suitable because grouping the EC2 instances in separate accounts does not provide the lowest possible latency between nodes. Separate accounts are used to isolate and organize resources for different purposes, such as security, billing, or compliance. However, they do not affect the network performance or proximity of the instances. Moreover, grouping the EC2 instances in separate accounts would incur additional costs and complexity, and it would require setting up cross-account networking and permissions.

Option D is not suitable because attaching multiple elastic network interfaces to each EC2 instance does not provide the lowest possible latency between nodes. Elastic network interfaces are virtual network interfaces that can be attached to EC2 instances to provide additional network capabilities, such as multiple IP addresses, multiple subnets, or enhanced security. However, they do not affect the network performance or proximity of the instances. Moreover, attaching multiple elastic network interfaces to each EC2 instance would consume additional resources and limit the instance type choices.

Option E is not suitable because using Amazon EBS optimized instance types does not provide the lowest possible latency between nodes. Amazon EBS optimized instance types are instances that provide dedicated bandwidth for Amazon EBS volumes, which are block storage volumes that can be attached to EC2 instances. EBS optimized instance types improve the performance and consistency of the EBS volumes, but they do not affect the network performance or proximity of the instances. Moreover, using EBS optimized instance types would incur additional costs and may not be necessary for the streaming data workload. References:

- ? Enhanced networking on Linux
- ? Placement groups
- ? Elastic network interfaces
- ? Amazon EBS-optimized instances

NEW QUESTION 111

- (Topic 4)

A company wants to manage Amazon Machine Images (AMIs). The company currently copies AMIs to the same AWS Region where the AMIs were created. The company needs to design an application that captures AWS API calls and sends alerts whenever the Amazon EC2 CreateImage API operation is called within the company's account.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Create an AWS Lambda function to query AWS CloudTrail logs and to send an alert when a CreateImage API call is detected.
- B. Configure AWS CloudTrail with an Amazon Simple Notification Service (Amazon SNS) notification that occurs when updated logs are sent to Amazon S3. Use Amazon Athena to create a new table and to query on CreateImage when an API call is detected.
- C. Create an Amazon EventBridge (Amazon CloudWatch Events) rule for the CreateImage API cal
- D. Configure the target as an Amazon Simple Notification Service (Amazon SNS) topic to send an alert when a CreateImage API call is detected.
- E. Configure an Amazon Simple Queue Service (Amazon SQS) FIFO queue as a target for AWS CloudTrail log
- F. Create an AWS Lambda function to send an alert to an Amazon Simple Notification Service (Amazon SNS) topic when a CreateImage API call is detected.

Answer: C

Explanation:

<https://docs.aws.amazon.com/AWSEC2/latest/WindowsGuide/monitor-ami-events.html#:~:text=For%20example%2C%20you%20can%20create%20an%20EventBridge%20rule%20that%20detects%20when%20the%20AMI%20creation%20process%20has%20completed%20and%20then%20invokes%20an%20Amazon%20SNS%20topic%20to%20send%20an%20email%20notification%20to%20you.>

NEW QUESTION 116

- (Topic 4)

A company previously migrated its data warehouse solution to AWS. The company also has an AWS Direct Connect connection. Corporate office users query the data warehouse using a visualization tool. The average size of a query returned by the data warehouse is 50 MB and each webpage sent by the visualization tool is approximately 500 KB. Result sets returned by the data warehouse are not cached.

Which solution provides the LOWEST data transfer egress cost for the company?

- A. Host the visualization tool on premises and query the data warehouse directly over the internet.
- B. Host the visualization tool in the same AWS Region as the data warehouse
- C. Access it over the internet.
- D. Host the visualization tool on premises and query the data warehouse directly over a Direct Connect connection at a location in the same AWS Region.
- E. Host the visualization tool in the same AWS Region as the data warehouse and access it over a Direct Connect connection at a location in the same Region.

Answer: D

Explanation:

<https://aws.amazon.com/directconnect/pricing/> <https://aws.amazon.com/blogs/aws/aws-data-transfer-prices-reduced/>

NEW QUESTION 118

- (Topic 4)

A company runs a web application on Amazon EC2 instances in an Auto Scaling group that has a target group. The company designed the application to work with session affinity (sticky sessions) for a better user experience.

The application must be available publicly over the internet as an endpoint. A WAF must be applied to the endpoint for additional security. Session affinity (sticky sessions) must be configured on the endpoint

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

- A. Create a public Network Load Balancer Specify the application target group.
- B. Create a Gateway Load Balancer Specify the application target group.
- C. Create a public Application Load Balancer Specify the application target group.
- D. Create a second target group
- E. Add Elastic IP addresses to the EC2 instances
- F. Create a web ACL in AWS WAF Associate the web ACL with the endpoint

Answer: CE

Explanation:

C and E are the correct answers because they allow the company to create a public endpoint for its web application that supports session affinity (sticky sessions) and has a WAF applied for additional security. By creating a public Application Load Balancer, the company can distribute incoming traffic across multiple EC2 instances in an Auto Scaling group and specify the application target group. By creating a web ACL in AWS WAF and associating it with the Application Load Balancer, the company can protect its web application from common web exploits. By enabling session stickiness on the Application Load Balancer, the company can ensure that subsequent requests from a user during a session are routed to the same target. References:

? Application Load Balancers

? AWS WAF

? Target Groups for Your Application Load Balancers

? How Application Load Balancer Works with Sticky Sessions

NEW QUESTION 119

- (Topic 4)

A company runs a website that uses a content management system (CMS) on Amazon EC2. The CMS runs on a single EC2 instance and uses an Amazon Aurora MySQL Multi-AZ DB instance for the data tier. Website images are stored on an Amazon Elastic Block Store (Amazon EBS) volume that is mounted inside the EC2 instance.

Which combination of actions should a solutions architect take to improve the performance and resilience of the website? (Select TWO.)

- A. Move the website images into an Amazon S3 bucket that is mounted on every EC2 instance.
- B. Share the website images by using an NFS share from the primary EC2 instance
- C. Mount this share on the other EC2 instances.
- D. Move the website images onto an Amazon Elastic File System (Amazon EFS) file system that is mounted on every EC2 instance.
- E. Create an Amazon Machine Image (AMI) from the existing EC2 instance Use the AMI to provision new instances behind an Application Load Balancer as part of an Auto Scaling group
- F. Configure the Auto Scaling group to maintain a minimum of two instances
- G. Configure an accelerator in AWS Global Accelerator for the website.
- H. Create an Amazon Machine Image (AMI) from the existing EC2 instance
- I. Use the AMI to provision new instances behind an Application Load Balancer as part of an Auto Scaling group
- J. Configure the Auto Scaling group to maintain a minimum of two instances
- K. Configure an Amazon CloudFront distribution for the website.

Answer: CE

Explanation:

Option C provides moving the website images onto an Amazon EFS file system that is mounted on every EC2 instance. Amazon EFS provides a scalable and fully managed file storage solution that can be accessed concurrently from multiple EC2 instances. This ensures that the website images can be accessed efficiently and consistently by all instances, improving performance. In Option E The Auto Scaling group maintains a minimum of two instances, ensuring resilience by automatically replacing any unhealthy instances. Additionally, configuring an Amazon CloudFront distribution for the website further improves performance by caching content at edge locations closer to the end-users, reducing latency and improving content delivery. Hence combining these actions, the website's performance is improved through efficient image storage and content delivery

NEW QUESTION 122

- (Topic 4)

A company is running a microservices application on Amazon EC2 instances. The company wants to migrate the application to an Amazon Elastic Kubernetes Service (Amazon EKS) cluster for scalability. The company must configure the Amazon EKS control plane with endpoint private access set to true and endpoint

public access set to false to maintain security compliance. The company must also put the data plane in private subnets. However, the company has received error notifications because the node cannot join the cluster.
Which solution will allow the node to join the cluster?

- A. Grant the required permission in AWS Identity and Access Management (IAM) to the AmazonEKSNoderole IAM role.
- B. Create interface VPC endpoints to allow nodes to access the control plane.
- C. Recreate nodes in the public subnet. Restrict security groups for EC2 nodes.
- D. Allow outbound traffic in the security group of the nodes.

Answer: B

Explanation:

Kubernetes API requests within your cluster's VPC (such as node to control plane communication) use the private VPC endpoint.
<https://docs.aws.amazon.com/eks/latest/userguide/cluster-endpoint.html>

NEW QUESTION 125

- (Topic 4)

A company stores critical data in Amazon DynamoDB tables in the company's AWS account. An IT administrator accidentally deleted a DynamoDB table. The deletion caused a significant loss of data and disrupted the company's operations. The company wants to prevent this type of disruption in the future.
Which solution will meet this requirement with the LEAST operational overhead?

- A. Configure a trail in AWS CloudTrail
- B. Create an Amazon EventBridge rule for delete action
- C. Create an AWS Lambda function to automatically restore deleted DynamoDB tables.
- D. Create a backup and restore plan for the DynamoDB table
- E. Recover the DynamoDB tables manually.
- F. Configure deletion protection on the DynamoDB tables.
- G. Enable point-in-time recovery on the DynamoDB tables.

Answer: C

Explanation:

Deletion protection is a feature of DynamoDB that prevents accidental deletion of tables. When deletion protection is enabled, you cannot delete a table unless you explicitly disable it first. This adds an extra layer of security and reduces the risk of data loss and operational disruption. Deletion protection is easy to enable and disable using the AWS Management Console, the AWS CLI, or the DynamoDB API. This solution has the least operational overhead, as you do not need to create, manage, or invoke any additional resources or services. References:

- ? Using deletion protection to protect your table
- ? Preventing Accidental Table Deletion in DynamoDB
- ? Amazon DynamoDB now supports table deletion protection

NEW QUESTION 130

- (Topic 4)

A company has multiple Windows file servers on premises. The company wants to migrate and consolidate its files into an Amazon FSx for Windows File Server file system. File permissions must be preserved to ensure that access rights do not change.
Which solutions will meet these requirements? (Select TWO.)

- A. Deploy AWS DataSync agents on premise
- B. Schedule DataSync tasks to transfer the data to the FSx for Windows File Server file system.
- C. Copy the shares on each file server into Amazon S3 buckets by using the AWS CLI. Schedule AWS DataSync tasks to transfer the data to the FSx for Windows File Server file system.
- D. Remove the drives from each file server. Ship the drives to AWS for import into Amazon S3. Schedule AWS DataSync tasks to transfer the data to the FSx for Windows File Server file system.
- E. Order an AWS Snowcone device
- F. Connect the device to the on-premises network
- G. Launch AWS DataSync agents on the device
- H. Schedule DataSync tasks to transfer the data to the FSx for Windows File Server file system.
- I. Order an AWS Snowball Edge Storage Optimized device
- J. Connect the device to the on-premises network
- K. Copy data to the device by using the AWS CLI
- L. Ship the device back to AWS for import into Amazon S3. Schedule AWS DataSync tasks to transfer the data to the FSx for Windows File Server file system.

Answer: AD

Explanation:

A This option involves deploying DataSync agents on your on-premises file servers and using DataSync to transfer the data directly to the FSx for Windows File Server. DataSync ensures that file permissions are preserved during the migration process. D This option involves using an AWS Snowcone device, a portable data transfer device. You would connect the Snowcone device to your on-premises network, launch DataSync agents on the device, and schedule DataSync tasks to transfer the data to FSx for Windows File Server. DataSync handles the migration process while preserving file permissions.

NEW QUESTION 133

- (Topic 4)

A company has a large workload that runs every Friday evening. The workload runs on Amazon EC2 instances that are in two Availability Zones in the us-east-1 Region. Normally, the company must run no more than two instances at all times. However, the company wants to scale up to six instances each Friday to handle a regularly repeating increased workload.
Which solution will meet these requirements with the LEAST operational overhead?

- A. Create a reminder in Amazon EventBridge to scale the instances.
- B. Create an Auto Scaling group that has a scheduled action.
- C. Create an Auto Scaling group that uses manual scaling.
- D. Create an Auto Scaling group that uses automatic scaling.

Answer: B

Explanation:

An Auto Scaling group is a collection of EC2 instances that share similar characteristics and can be scaled in or out automatically based on demand. An Auto Scaling group can have a scheduled action, which is a configuration that tells the group to scale to a specific size at a specific time. This way, the company can scale up to six instances each Friday evening to handle the increased workload, and scale down to two instances at other times to save costs. This solution meets the requirements with the least operational overhead, as it does not require manual intervention or custom scripts. References:

? 1 explains how to create a scheduled action for an Auto Scaling group.

? 2 describes the concept and benefits of an Auto Scaling group.

NEW QUESTION 137

- (Topic 4)

The customers of a finance company request appointments with financial advisors by sending text messages. A web application that runs on Amazon EC2 instances accepts the appointment requests. The text messages are published to an Amazon Simple Queue Service (Amazon SQS) queue through the web application. Another application that runs on EC2 instances then sends meeting invitations and meeting confirmation email messages to the customers. After successful scheduling, this application stores the meeting information in an Amazon DynamoDB database.

As the company expands, customers report that their meeting invitations are taking longer to arrive.

What should a solutions architect recommend to resolve this issue?

- A. Add a DynamoDB Accelerator (DAX) cluster in front of the DynamoDB database.
- B. Add an Amazon API Gateway API in front of the web application that accepts the appointment requests.
- C. Add an Amazon CloudFront distributio
- D. Set the origin as the web application that accepts the appointment requests.
- E. Add an Auto Scaling group for the application that sends meeting invitation
- F. Configure the Auto Scaling group to scale based on the depth of the SQS queue.

Answer: D

Explanation:

To resolve the issue of longer delivery times for meeting invitations, the solutions architect can recommend adding an Auto Scaling group for the application that sends meeting invitations and configuring the Auto Scaling group to scale based on the depth of the SQS queue. This will allow the application to scale up as the number of appointment requests increases, improving the performance and delivery times of the meeting invitations.

NEW QUESTION 140

- (Topic 4)

A company has an online gaming application that has TCP and UDP multiplayer gaming capabilities. The company uses Amazon Route 53 to point the application traffic to multiple Network Load Balancers (NLBs) in different AWS Regions. The company needs to improve application performance and decrease latency for the online game in preparation for user growth.

Which solution will meet these requirements?

- A. Add an Amazon CloudFront distribution in front of the NLB
- B. Increase the Cache- Control: max-age parameter.
- C. Replace the NLBs with Application Load Balancers (ALBs). Configure Route 53 to use latency-based routing.
- D. Add AWS Global Accelerator in front of the NLB
- E. Configure a Global Accelerator endpoint to use the correct listener ports.
- F. 'Add an Amazon API Gateway endpoint behind the NLB
- G. Enable API cachin
- H. Override method caching for the different stages.

Answer: C

Explanation:

This answer is correct because it improves the application performance and decreases latency for the online game by using AWS Global Accelerator. AWS Global Accelerator is a networking service that helps you improve the availability, performance, and security of your public applications. Global Accelerator provides two global static public IPs that act as a fixed entry point to your application endpoints, such as NLBs, in different AWS Regions. Global Accelerator uses the AWS global network to route traffic to the optimal regional endpoint based on health, client location, and policies that you configure. Global Accelerator also terminates TCP and UDP traffic at the edge locations, which reduces the number of hops and improves the network performance. By adding AWS Global Accelerator in front of the NLBs, you can achieve up to 60% improvement in latency for your online game.

References:

? <https://docs.aws.amazon.com/global-accelerator/latest/dg/what-is-global-accelerator.html>

? <https://aws.amazon.com/global-accelerator/>

NEW QUESTION 145

- (Topic 4)

A company hosts a website on Amazon EC2 instances behind an Application Load Balancer (ALB) The website serves static content Website traffic is increasing and the company is concerned about a potential increase in cost.

What should a solutions architect do to reduce the cost of the website?

- A. Create an Amazon CloudFront distribution to cache static files at edge locations.
- B. Create an Amazon ElastiCache cluster Connect the ALB to the ElastiCache cluster to serve cached files.
- C. Create an AWS WAF web ACL and associate it with the AL
- D. Add a rule to the web ACL to cache static files.
- E. Create a second ALB in an alternative AWS Region Route user traffic to the closest Region to minimize data transfer costs

Answer: A

Explanation:

Amazon CloudFront is a content delivery network (CDN) that can improve the performance and reduce the cost of serving static content from a website. CloudFront

can cache static files at edge locations closer to the users, reducing the latency and data transfer costs. CloudFront can also integrate with Amazon S3 as the origin for the static content, eliminating the need for EC2 instances to host the website. CloudFront meets all the requirements of the question, while the other options do not. References:

- ? <https://aws.amazon.com/blogs/architecture/architecting-a-low-cost-web-content-publishing-system/>
- ? <https://nodeployfriday.com/posts/static-website-hosting/>
- ? <https://aws.amazon.com/cloudfront/>

NEW QUESTION 150

- (Topic 4)

A company runs multiple Amazon EC2 Linux instances in a VPC across two Availability Zones. The instances host applications that use a hierarchical directory structure. The applications need to read and write rapidly and concurrently to shared storage.

What should a solutions architect do to meet these requirements?

- A. Create an Amazon S3 bucket
- B. Allow access from all the EC2 instances in the VPC.
- C. Create an Amazon Elastic File System (Amazon EFS) file system
- D. Mount the EFS file system from each EC2 instance.
- E. Create a file system on a Provisioned IOPS SSD (102) Amazon Elastic Block Store (Amazon EBS) volume
- F. Attach the EBS volume to all the EC2 instances.
- G. Create file systems on Amazon Elastic Block Store (Amazon EBS) volumes that are attached to each EC2 instance
- H. Synchronize the EBS volumes across the different EC2 instances.

Answer: B

Explanation:

It allows the EC2 instances to read and write rapidly and concurrently to shared storage across two Availability Zones. Amazon EFS provides a scalable, elastic, and highly available file system that can be mounted from multiple EC2 instances. Amazon EFS supports high levels of throughput and IOPS, and consistent low latencies. Amazon EFS also supports NFSv4 lock upgrading and downgrading, which enables high levels of concurrency. References:

- ? [Amazon EFS Features](#)
- ? [Using Amazon EFS with Amazon EC2](#)

NEW QUESTION 152

- (Topic 4)

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?

- 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 unavailable.

References:

- ? 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

NEW QUESTION 154

- (Topic 4)

A company has two VPCs that are located in the us-west-2 Region within the same AWS account. The company needs to allow network traffic between these VPCs. Approximately 500 GB of data transfer will occur between the VPCs each month.

What is the MOST cost-effective solution to connect these VPCs?

- A. Implement AWS Transit Gateway to connect the VPC
- B. Update the route tables of each VPC to use the transit gateway for inter-VPC communication.
- C. Implement an AWS Site-to-Site VPN tunnel between the VPC
- D. Update the route tables of each VPC to use the VPN tunnel for inter-VPC communication.
- E. Set up a VPC peering connection between the VPC
- F. Update the route tables of each VPC to use the VPC peering connection for inter-VPC communication.
- G. Set up a 1 GB AWS Direct Connect connection between the VPC
- H. Update the route tables of each VPC to use the Direct Connect connection for inter-VPC communication.

Answer: C

Explanation:

To connect two VPCs in the same Region within the same AWS account, VPC peering is the most cost-effective solution. VPC peering allows direct network traffic between the VPCs without requiring a gateway, VPN connection, or AWS Transit Gateway. VPC peering also does not incur any additional charges for data transfer between the VPCs.

References:

- ? [What Is VPC Peering?](#)
- ? [VPC Peering Pricing](#)

NEW QUESTION 156

- (Topic 4)

A manufacturing company has machine sensors that upload .csv files to an Amazon S3 bucket. These .csv files must be converted into images and must be made

available as soon as possible for the automatic generation of graphical reports.

The images become irrelevant after 1 month, but the .csv files must be kept to train machine learning (ML) models twice a year. The ML trainings and audits are planned weeks in advance.

Which combination of steps will meet these requirements MOST cost-effectively? (Select TWO.)

- A. Launch an Amazon EC2 Spot Instance that downloads the .csv files every hour, generates the image files, and uploads the images to the S3 bucket.
- B. Design an AWS Lambda function that converts the .csv files into images and stores the images in the S3 bucket
- C. Invoke the Lambda function when a .csv file is uploaded.
- D. Create S3 Lifecycle rules for .csv files and image files in the S3 bucket
- E. Transition the .csv files from S3 Standard to S3 Glacier 1 day after they are uploaded
- F. Expire the image files after 30 days.
- G. Create S3 Lifecycle rules for .csv files and image files in the S3 bucket
- H. Transition the .csv files from S3 Standard to S3 One Zone-Infrequent Access (S3 One Zone-IA) 1 day after they are uploaded
- I. Expire the image files after 30 days.
- J. Create S3 Lifecycle rules for .csv files and image files in the S3 bucket
- K. Transition the .csv files from S3 Standard to S3 Standard-Infrequent Access (S3 Standard-IA) 1 day after they are uploaded
- L. Keep the image files in Reduced Redundancy Storage (RRS).

Answer: BC

Explanation:

These answers are correct because they meet the requirements of converting the .csv files into images, making them available as soon as possible, and minimizing the storage costs. AWS Lambda is a service that lets you run code without provisioning or managing servers. You can use AWS Lambda to design a function that converts the .csv files into images and stores the images in the S3 bucket. You can invoke the Lambda function when a .csv file is uploaded to the S3 bucket by using an S3 event notification. This way, you can ensure that the images are generated and made available as soon as possible for the graphical reports. S3 Lifecycle is a feature that enables you to manage your objects so that they are stored cost effectively throughout their lifecycle. You can create S3 Lifecycle rules for .csv files and image files in the S3 bucket to transition them to different storage classes or expire them based on your business needs. You can transition the .csv files from S3 Standard to S3 Glacier 1 day after they are uploaded, since they are only needed twice a year for ML trainings and audits that are planned weeks in advance. S3 Glacier is a storage class for data archiving that offers secure, durable, and extremely low-cost storage with retrieval times ranging from minutes to hours. You can expire the image files after 30 days, since they become irrelevant after 1 month. References:

- ? <https://docs.aws.amazon.com/lambda/latest/dg/welcome.html>
- ? <https://docs.aws.amazon.com/AmazonS3/latest/userguide/NotificationHowTo.html>
- ? <https://docs.aws.amazon.com/AmazonS3/latest/userguide/object-lifecycle-mgmt.html>
- ? <https://docs.aws.amazon.com/AmazonS3/latest/userguide/storage-class-intro.html#sc-glacier>

NEW QUESTION 160

- (Topic 4)

A company is running its production and nonproduction environment workloads in multiple AWS accounts. The accounts are in an organization in AWS Organizations. The company needs to design a solution that will prevent the modification of cost usage tags.

Which solution will meet these requirements?

- A. Create a custom AWS Config rule to prevent tag modification except by authorized principals.
- B. Create a custom trail in AWS CloudTrail to prevent tag modification
- C. Create a service control policy (SCP) to prevent tag modification except by authorized principals.
- D. Create custom Amazon CloudWatch logs to prevent tag modification.

Answer: C

Explanation:

This solution meets the requirements because it uses SCPs to restrict the actions that can be performed on cost usage tags in the organization. SCPs are a type of organization policy that you can use to manage permissions in your organization. SCPs specify the maximum permissions for an organization, organizational unit (OU), or account. You can use SCPs to enforce consistent tag policies across your organization and prevent unauthorized or accidental changes to your tags. You can also create exceptions for authorized principals, such as administrators or auditors, who need to modify tags for legitimate purposes.

References:

- ? [Service control policies \(SCPs\) - AWS Organizations](#)
- ? [Tag policies - AWS Organizations](#)

NEW QUESTION 164

- (Topic 4)

An ecommerce company runs applications in AWS accounts that are part of an organization in AWS Organizations. The applications run on Amazon Aurora PostgreSQL databases across all the accounts. The company needs to prevent malicious activity and must identify abnormal failed and incomplete login attempts to the databases.

Which solution will meet these requirements in the MOST operationally efficient way?

- A. Attach service control policies (SCPs) to the root of the organization to identify the failed login attempts
- B. Enable the Amazon RDS Protection feature in Amazon GuardDuty for the member accounts of the organization
- C. Publish the Aurora general logs to a log group in Amazon CloudWatch Logs. Export the log data to a central Amazon S3 bucket
- D. Publish all the Aurora PostgreSQL database events in AWS CloudTrail to a central Amazon S3 bucket

Answer: C

Explanation:

This option is the most operationally efficient way to meet the requirements because it allows the company to monitor and analyze the database login activity across all the accounts in the organization. By publishing the Aurora general logs to a log group in Amazon CloudWatch Logs, the company can enable the logging of the database connections, disconnections, and failed authentication attempts. By exporting the log data to a central Amazon S3 bucket, the company can store the log data in a durable and cost-effective way and use other AWS services or tools to perform further analysis or alerting on the log data. For example, the company can use Amazon Athena to query the log data in Amazon S3, or use Amazon SNS to send notifications based on the log data.

* A. Attach service control policies (SCPs) to the root of the organization to identify the failed login attempts. This option is not effective because SCPs are not designed to identify the failed login attempts, but to restrict the actions that the users and roles can perform in the member accounts of the organization. SCPs are applied to the AWS API calls, not to the database login attempts. Moreover, SCPs do not provide any logging or analysis capabilities for the database activity.

* B. Enable the Amazon RDS Protection feature in Amazon GuardDuty for the member accounts of the organization. This option is not optimal because the Amazon RDS Protection feature in Amazon GuardDuty is not available for Aurora PostgreSQL databases, but only for Amazon RDS for MySQL and Amazon RDS for MariaDB databases. Moreover, the Amazon RDS Protection feature does not monitor the database login attempts, but the network and API activity related to the RDS instances.

* D. Publish all the Aurora PostgreSQL database events in AWS CloudTrail to a central Amazon S3 bucket. This option is not sufficient because AWS CloudTrail does not capture the database login attempts, but only the AWS API calls made by or on behalf of the Aurora PostgreSQL database. For example, AWS CloudTrail can record the events such as creating, modifying, or deleting the database instances, clusters, or snapshots, but not the events such as connecting, disconnecting, or failing to authenticate to the database. References:

? 1 Working with Amazon Aurora PostgreSQL - Amazon Aurora

? 2 Working with log groups and log streams - Amazon CloudWatch Logs

? 3 Exporting Log Data to Amazon S3 - Amazon CloudWatch Logs

? [4] Amazon GuardDuty FAQs

? [5] Logging Amazon RDS API Calls with AWS CloudTrail - Amazon Relational Database Service

NEW QUESTION 168

- (Topic 4)

A company is storing 700 terabytes of data on a large network-attached storage (NAS) system in its corporate data center. The company has a hybrid environment with a 10 Gbps AWS Direct Connect connection.

After an audit from a regulator, the company has 90 days to move the data to the cloud. The company needs to move the data efficiently and without disruption.

The company still needs to be able to access and update the data during the transfer window.

Which solution will meet these requirements?

- A. Create an AWS DataSync agent in the corporate data center.
- B. Create a data transfer task.
- C. Start the transfer to an Amazon S3 bucket.
- D. Back up the data to AWS Snowball Edge Storage Optimized device.
- E. Ship the devices to an AWS data center.
- F. Mount a target Amazon S3 bucket on the on-premises file system.
- G. Use rsync to copy the data directly from local storage to a designated Amazon S3 bucket over the Direct Connect connection.
- H. Back up the data on tape.
- I. Ship the tapes to an AWS data center.
- J. Mount a target Amazon S3 bucket on the on-premises file system.

Answer: A

Explanation:

This answer is correct because it meets the requirements of moving the data efficiently and without disruption, and still being able to access and update the data during the transfer window. AWS DataSync is an online data movement and discovery service that simplifies and accelerates data migrations to AWS and helps you move data quickly and securely between on-premises storage, edge locations, other clouds, and AWS Storage. You can create an AWS DataSync agent in the corporate data center to connect your NAS system to AWS over the Direct Connect connection. You can create a data transfer task to specify the source location, destination location, and options for transferring the data. You can start the transfer to an Amazon S3 bucket and monitor the progress of the task. DataSync automatically encrypts data in transit and verifies data integrity during transfer. DataSync also supports incremental transfers, which means that only files that have changed since the last transfer are copied. This way, you can ensure that your data is synchronized between your NAS system and S3 bucket, and you can access and update the data during the transfer window.

References:

? <https://docs.aws.amazon.com/datasync/latest/userguide/what-is-datasync.html>

? <https://docs.aws.amazon.com/datasync/latest/userguide/how-datasync-works.html>

NEW QUESTION 173

.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

AWS-Solution-Architect-Associate Practice Exam Features:

- * AWS-Solution-Architect-Associate Questions and Answers Updated Frequently
- * AWS-Solution-Architect-Associate Practice Questions Verified by Expert Senior Certified Staff
- * AWS-Solution-Architect-Associate Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- * AWS-Solution-Architect-Associate Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The AWS-Solution-Architect-Associate Practice Test Here](#)