

Exam Questions AWS-Solution-Architect-Associate

Amazon AWS Certified Solutions Architect - Associate

<https://www.2passeasy.com/dumps/AWS-Solution-Architect-Associate/>



NEW QUESTION 1

- (Topic 4)

A company has a business-critical application that runs on Amazon EC2 instances. The application stores data in an Amazon DynamoDB table. The company must be able to revert the table to any point within the last 24 hours.

Which solution meets these requirements with the LEAST operational overhead?

- A. Configure point-in-time recovery for the table.
- B. Use AWS Backup for the table.
- C. Use an AWS Lambda function to make an on-demand backup of the table every hour.
- D. Turn on streams on the table to capture a log of all changes to the table in the last 24 hours. Store a copy of the stream in an Amazon S3 bucket.

Answer: A

Explanation:

Point-in-time recovery (PITR) for DynamoDB is a feature that enables you to restore your table data to any point in time during the last 35 days. PITR helps protect your table from accidental write or delete operations, such as a test script writing to a production table or a user issuing a wrong command. PITR is easy to use, fully managed, fast, and scalable. You can enable PITR with a single click in the DynamoDB console or with a simple API call. You can restore a table to a new table using the console, the AWS CLI, or the DynamoDB API. PITR does not consume any provisioned table capacity and has no impact on the performance or availability of your production applications. PITR meets the requirements of the company with the least operational overhead, as it does not require any manual backup creation, scheduling, or maintenance. It also provides per-second granularity for restoring the table to any point within the last 24 hours.

References:

- ? Point-in-time recovery for DynamoDB - Amazon DynamoDB
- ? Amazon DynamoDB point-in-time recovery (PITR)
- ? Enable Point-in-Time Recovery (PITR) for Dynamodb global tables
- ? Restoring a DynamoDB table to a point in time - Amazon DynamoDB
- ? Point-in-time recovery: How it works - Amazon DynamoDB

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 company has a multi-tier payment processing application that is based on virtual machines (VMs). The communication between the tiers occurs asynchronously.

through a third-party middleware solution that guarantees exactly-once delivery.

The company needs a solution that requires the least amount of infrastructure management. The solution must guarantee exactly-once delivery for application messaging

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

- A. Use AWS Lambda for the compute layers in the architecture.
- B. Use Amazon EC2 instances for the compute layers in the architecture.
- C. Use Amazon Simple Notification Service (Amazon SNS) as the messaging component between the compute layers.
- D. Use Amazon Simple Queue Service (Amazon SQS) FIFO queues as the messaging component between the compute layers.
- E. Use containers that are based on Amazon Elastic Kubernetes Service (Amazon EKS) for the compute layers in the architecture.

Answer: AD

Explanation:

This solution meets the requirements because it requires the least amount of infrastructure management and guarantees exactly-once delivery for application messaging. AWS Lambda is a serverless compute service that lets you run code without provisioning or managing servers. You only pay for the compute time you consume. Lambda scales automatically with the size of your workload. Amazon SQS FIFO queues are designed to ensure that messages are processed exactly once, in the exact order that they are sent. FIFO queues have high availability and deliver messages in a strict first-in, first-out order. You can use Amazon SQS to decouple and scale microservices, distributed systems, and serverless applications. References: AWS Lambda, Amazon SQS FIFO queues

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 tasks¹.

NEW QUESTION 7

- (Topic 4)

An image hosting company uploads its large assets to Amazon S3 Standard buckets. The company uses multipart upload in parallel by using S3 APIs and overwrites if the same object is uploaded again. For the first 30 days after upload, the objects will be accessed frequently. The objects will be used less frequently after 30 days, but the access patterns for each object will be inconsistent. The company must optimize its S3 storage costs while maintaining high availability and resiliency of stored assets.

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

- A. Move assets to S3 Intelligent-Tiering after 30 days.
- B. Configure an S3 Lifecycle policy to clean up incomplete multipart uploads.
- C. Configure an S3 Lifecycle policy to clean up expired object delete markers.
- D. Move assets to S3 Standard-Infrequent Access (S3 Standard-IA) after 30 days.
- E. Move assets to S3 One Zone-Infrequent Access (S3 One Zone-IA) after 30 days.

Answer: AB

Explanation:

S3 Intelligent-Tiering is a storage class that automatically moves data to the most cost-effective access tier based on access frequency, without performance impact, retrieval fees, or operational overhead¹. It is ideal for data with unknown or changing access patterns, such as the company's assets. By moving assets to S3 Intelligent-Tiering after 30 days, the company can optimize its storage costs while maintaining high availability and resilience of stored assets.

S3 Lifecycle is a feature that enables you to manage your objects so that they are stored cost effectively throughout their lifecycle². You can create lifecycle rules to define actions that Amazon S3 applies to a group of objects. One of the actions is to abort incomplete multipart uploads that can occur when an upload is interrupted. By configuring an S3 Lifecycle policy to clean up incomplete multipart uploads, the company can reduce its storage costs and avoid paying for parts that are not used.

Option C is incorrect because expired object delete markers are automatically deleted by Amazon S3 and do not incur any storage costs³. Therefore, configuring an S3 Lifecycle policy to clean up expired object delete markers will not have any effect on the company's storage costs.

Option D is incorrect because S3 Standard-IA is a storage class for data that is accessed less frequently, but requires rapid access when needed¹. It has a lower storage cost than S3 Standard, but it has a higher retrieval cost and a minimum storage duration charge of 30 days. Therefore, moving assets to S3 Standard-IA after 30 days may not optimize the company's storage costs if the assets are still accessed occasionally.

Option E is incorrect because S3 One Zone-IA is a storage class for data that is accessed less frequently, but requires rapid access when needed¹. It has a lower storage cost than S3 Standard-IA, but it stores data in only one Availability Zone and has less resilience than other storage classes. It also has a higher retrieval cost and a minimum storage duration charge of 30 days. Therefore, moving assets to S3 One Zone-IA after 30 days may not optimize the company's storage costs if the assets are still accessed occasionally or require high availability. Reference URL: 1: <https://docs.aws.amazon.com/AmazonS3/latest/userguide/storage-class-intro.html> 2:

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/object-lifecycle-mgmt.html> 3: <https://docs.aws.amazon.com/AmazonS3/latest/userguide/delete-or-empty-bucket.html#delete-bucket-considerations> : <https://docs.aws.amazon.com/AmazonS3/latest/userguide/mpuoverview.html> :

<https://aws.amazon.com/certification/certified-solutions-architect-associate/>

NEW QUESTION 8

- (Topic 4)

A company runs a highly available SFTP service. The SFTP service uses two Amazon EC2

Linux instances that run with elastic IP addresses to accept traffic from trusted IP sources on the internet. The SFTP service is backed by shared storage that is attached to the instances. User accounts are created and managed as Linux users in the SFTP servers.

The company wants a serverless option that provides high IOPS performance and highly configurable security. The company also wants to maintain control over user permissions.

Which solution will meet these requirements?

- A. Create an encrypted Amazon Elastic Block Store (Amazon EBS) volume
- B. Create an AWS Transfer Family SFTP service with a public endpoint that allows only trusted IP addresses
- C. Attach the EBS volume to the SFTP service endpoint
- D. Grant users access to the SFTP service.
- E. Create an encrypted Amazon Elastic File System (Amazon EFS) volume
- F. Create an AWS Transfer Family SFTP service with elastic IP addresses and a VPC endpoint that has internet-facing access
- G. Attach a security group to the endpoint that allows only trusted IP addresses
- H. Attach the EFS volume to the SFTP service endpoint
- I. Grant users access to the SFTP service.
- J. Create an Amazon S3 bucket with default encryption enabled
- K. Create an AWS Transfer Family SFTP service with a public endpoint that allows only trusted IP addresses
- L. Attach the S3 bucket to the SFTP service endpoint
- M. Grant users access to the SFTP service.
- N. Create an Amazon S3 bucket with default encryption enabled
- O. Create an AWS Transfer Family SFTP service with a VPC endpoint that has internal access in a private subnet
- P. Attach a security group that allows only trusted IP addresses
- Q. Attach the S3 bucket to the SFTP service endpoint
- R. Grant users access to the SFTP service.

Answer: C

Explanation:

AWS Transfer Family is a secure transfer service that enables you to transfer files into and out of AWS storage services using SFTP, FTPS, FTP, and AS2 protocols. You can use AWS Transfer Family to create an SFTP-enabled server with a public endpoint that allows only trusted IP addresses. You can also attach an Amazon S3 bucket with default encryption enabled to the SFTP service endpoint, which will provide high IOPS performance and highly configurable security for your data at rest. You can also maintain control over user permissions by granting users access to the SFTP service using IAM roles or service-managed identities. References: <https://docs.aws.amazon.com/transfer/latest/userguide/what-is-aws-transfer-family.html>

<https://docs.aws.amazon.com/transfer/latest/userguide/create-server-s3.html>

NEW QUESTION 9

- (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": "Allow",
      "Action": ["ec2:DescribeInstances"],
      "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 10

- (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 10

- (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 11

- (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 ahas/aws/ebs alias 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 13

- (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 subne
- 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 17

- (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 19

- (Topic 4)

A retail company uses a regional Amazon API Gateway API for its public REST APIs. The API Gateway endpoint is a custom domain name that points to an Amazon Route 53 alias record. A solutions architect needs to create a solution that has minimal effects on customers and minimal data loss to release the new version of APIs.

Which solution will meet these requirements?

- A. Create a canary release deployment stage for API Gatewa
- B. Deploy the latest API versio
- C. Point an appropriate percentage of traffic to the canary stag
- D. After API verification, promote the canary stage to the production stage.
- E. Create a new API Gateway endpoint with a new version of the API in OpenAPI YAMLfile forma
- F. Use the import-to-update operation in merge mode into the API in API Gatewa
- G. Deploy the new version of the API to the production stage.
- H. Create a new API Gateway endpoint with a new version of the API in OpenAPI JSON file forma
- I. Use the import-to-update operation in overwrite mode into the API in API Gatewa
- J. Deploy the new version of the API to the production stage.
- K. Create a new API Gateway endpoint with new versions of the API definition
- L. Create a custom domain name for the new API Gateway AP
- M. Point the Route 53 alias record to the new API Gateway API custom domain name.

Answer: A

Explanation:

This answer is correct because it meets the requirements of releasing the new version of APIs with minimal effects on customers and minimal data loss. A canary release deployment is a software development strategy in which a new version of an API is deployed for testing purposes, and the base version remains deployed as a production release for normal operations on the same stage. In a canary release deployment, total API traffic is separated at random into a production release and a canary release with a pre- configured ratio. Typically, the canary release receives a small percentage of API traffic and the production release takes up the rest. The updated API features are only visible to API traffic through the canary. You can adjust the canary traffic percentage to optimize test coverage or performance. By keeping canary traffic small and the selection random, most users are not adversely affected at any time by potential bugs in the new version, and no single user is adversely affected all the time. After the test metrics pass your requirements, you can promote the canary release to the production release and disable the canary from the deployment. This makes the new features available in the production stage. References:

? <https://docs.aws.amazon.com/apigateway/latest/developerguide/canary-release.html>

NEW QUESTION 20

- (Topic 4)

A solutions architect is designing a REST API in Amazon API Gateway for a cash payback service The application requires 1 GB of memory and 2 GB of storage for its computation resources. The application will require that the data is in a relational format.

Which additional combination of AWS services will meet these requirements with the LEAST administrative effort? {Select TWO.}

- A. Amazon EC2
- B. AWS Lambda
- C. Amazon RDS
- D. Amazon DynamoDB
- E. Amazon Elastic Kubernetes Services (Amazon EKS)

Answer: BC

Explanation:

AWS Lambda is a service that lets users run code without provisioning or managing servers. It automatically scales and manages the underlying compute resources for the code. It supports multiple languages, such as Java, Python, Node.js, and G1o. By using AWS Lambda for the REST API, the solution can meet the requirements of 1 GB of memory and minimal administrative effort.

Amazon RDS is a service that makes it easy to set up, operate, and scale a relational database in the cloud. It provides cost-efficient and resizable capacity while automating time-consuming administration tasks such as hardware provisioning, database setup, patching and backups. It supports multiple database engines,

such as MySQL, PostgreSQL, Oracle, and SQL Server². By using Amazon RDS for the data store, the solution can meet the requirements of 2 GB of storage and a relational format.

* A. Amazon EC2. This solution will not meet the requirement of minimal administrative effort, as Amazon EC2 is a service that provides virtual servers in the cloud that users have to configure and manage themselves. It requires users to choose an instance type, an operating system, a security group, and other options³.

* D. Amazon DynamoDB. This solution will not meet the requirement of a relational format, as Amazon DynamoDB is a service that provides a key-value and document database that delivers single-digit millisecond performance at any scale. It is a non-relational or NoSQL database that does not support joins or transactions.

* E. Amazon Elastic Kubernetes Services (Amazon EKS). This solution will not meet the requirement of minimal administrative effort, as Amazon EKS is a service that provides a fully managed Kubernetes service that users have to configure and manage themselves. It requires users to create clusters, nodes groups, pods, services, and other Kubernetes resources.

Reference URL: <https://aws.amazon.com/lambda/>

NEW QUESTION 21

- (Topic 4)

A solutions architect has created two IAM policies: Policy1 and Policy2. Both policies are attached to an IAM group.

Policy 1

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "iam:Get*",
        "iam:List*",
        "kms:List*",
        "ec2:*",
        "ds:*",
        "logs:Get*",
        "logs:Describe*"
      ],
      "Resource": "*"
    }
  ]
}
```

Policy 2

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Deny",
      "Action": "ds:Delete*",
      "Resource": "*"
    }
  ]
}
```

A cloud engineer is added as an IAM user to the IAM group. Which action will the cloud engineer be able to perform?

- A. Deleting IAM users
- B. Deleting directories
- C. Deleting Amazon EC2 instances
- D. Deleting logs from Amazon CloudWatch Logs

Answer: C

Explanation:

<https://awscli.amazonaws.com/v2/documentation/api/latest/reference/ds/index.html>

NEW QUESTION 26

- (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 31

- (Topic 4)

A company built an application with Docker containers and needs to run the application in the AWS Cloud. The company wants to use a managed service to host the application.

The solution must scale in and out appropriately according to demand on the individual container services. The solution also must not result in additional operational overhead or infrastructure to manage.

Which solutions will meet these requirements? (Select TWO)

- A. Use Amazon Elastic Container Service (Amazon ECS) with AWS Fargate.
- B. Use Amazon Elastic Kubernetes Service (Amazon EKS) with AWS Fargate.
- C. Provision an Amazon API Gateway. API Connect the API to AWS Lambda to run the containers.
- D. Use Amazon Elastic Container Service (Amazon ECS) with Amazon EC2 worker nodes.
- E. Use Amazon Elastic Kubernetes Service (Amazon EKS) with Amazon EC2 worker nodes.

Answer: AB

Explanation:

These options are the best solutions because they allow the company to run the application with Docker containers in the AWS Cloud using a managed service that scales automatically and does not require any infrastructure to manage. By using AWS Fargate, the company can launch and run containers without having to provision, configure, or scale clusters of EC2 instances. Fargate allocates the right amount of compute resources for each container and scales them up or down as needed. By using Amazon ECS or Amazon EKS, the company can choose the container orchestration platform that suits its needs. Amazon ECS is a fully managed service that integrates with other AWS services and simplifies the deployment and management of containers. Amazon EKS is a managed service that runs Kubernetes on AWS and provides compatibility with existing Kubernetes tools and plugins.

* C. Provision an Amazon API Gateway. API Connect the API to AWS Lambda to run the containers. This option is not feasible because AWS Lambda does not support running Docker containers directly. Lambda functions are executed in a sandboxed environment that is isolated from other functions and resources. To run Docker containers on Lambda, the company would need to use a custom runtime or a wrapper library that emulates the Docker API, which can introduce additional complexity and overhead.

* D. Use Amazon Elastic Container Service (Amazon ECS) with Amazon EC2 worker nodes. This option is not optimal because it requires the company to manage the EC2 instances that host the containers. The company would need to provision, configure, scale, patch, and monitor the EC2 instances, which can increase the operational overhead and infrastructure costs.

* E. Use Amazon Elastic Kubernetes Service (Amazon EKS) with Amazon EC2 worker nodes. This option is not ideal because it requires the company to manage the EC2 instances that host the containers. The company would need to provision, configure, scale, patch, and monitor the EC2 instances, which can increase the operational overhead and infrastructure costs.

References:

? 1 AWS Fargate - Amazon Web Services

? 2 Amazon Elastic Container Service - Amazon Web Services

? 3 Amazon Elastic Kubernetes Service - Amazon Web Services

? 4 AWS Lambda FAQs - Amazon Web Services

NEW QUESTION 33

- (Topic 4)

A company runs a three-tier application in two AWS Regions. The web tier, the application tier, and the database tier run on Amazon EC2 instances. The company uses Amazon RDS for Microsoft SQL Server Enterprise for the database tier. The database tier is experiencing high load when weekly and monthly reports are run. The company wants to reduce the load on the database tier.

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

- A. Create read replica
- B. Configure the reports to use the new read replicas.
- C. Convert the RDS database to Amazon DynamoDB. Configure the reports to use DynamoDB
- D. Modify the existing RDS DB instances by selecting a larger instance size.
- E. Modify the existing RDS DB instances and put the instances into an Auto Scaling group.

Answer: A

Explanation:

it allows the company to create read replicas of its RDS database and reduce the load on the database tier. By creating read replicas, the company can offload read traffic from the primary database instance to one or more replicas. By configuring the reports to use the new read replicas, the company can improve performance and availability of its database tier. References:

? Working with Read Replicas

? Read Replicas for Amazon RDS for SQL Server

NEW QUESTION 34

- (Topic 4)

A company designed a stateless two-tier application that uses Amazon EC2 in a single Availability Zone and an Amazon RDS Multi-AZ DB instance. New company management wants to ensure the application is highly available.

What should a solutions architect do to meet this requirement?

- A. Configure the application to use Multi-AZ EC2 Auto Scaling and create an Application Load Balancer
- B. Configure the application to take snapshots of the EC2 instances and send them to a different AWS Region.
- C. Configure the application to use Amazon Route 53 latency-based routing to feed requests to the application.
- D. Configure Amazon Route 53 rules to handle incoming requests and create a Multi-AZ Application Load Balancer

Answer: A

Explanation:

<https://docs.aws.amazon.com/autoscaling/ec2/userguide/as-add-availability-zone.html>

NEW QUESTION 38

- (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 40

- (Topic 4)

A company has created a multi-tier application for its ecommerce website. The website uses an Application Load Balancer that resides in the public subnets, a web tier in the public subnets, and a MySQL cluster hosted on Amazon EC2 instances in the private subnets. The MySQL database needs to retrieve product catalog and pricing information that is hosted on the internet by a third-party provider. A solutions architect must devise a strategy that maximizes security without increasing operational overhead. What should the solutions architect do to meet these requirements?

- A. Deploy a NAT instance in the VP
- B. Route all the internet-based traffic through the NAT instance.
- C. Deploy a NAT gateway in the public subnet
- D. Modify the private subnet route table to direct all internet-bound traffic to the NAT gateway.
- E. Configure an internet gateway and attach it to the VP

- F. Modify the private subnet route table to direct internet-bound traffic to the internet gateway.
- G. Configure a virtual private gateway and attach it to the VP
- H. Modify the private subnet route table to direct internet-bound traffic to the virtual private gateway.

Answer: B

Explanation:

To allow the MySQL database in the private subnets to access the internet without exposing it to the public, a NAT gateway is a suitable solution. A NAT gateway enables instances in a private subnet to connect to the internet or other AWS services, but prevents the internet from initiating a connection with those instances. A NAT gateway resides in the public subnets and can handle high throughput of traffic with low latency. A NAT gateway is also a managed service that does not require any operational overhead. References:

? NAT Gateways

? NAT Gateway Pricing

NEW QUESTION 44

- (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 45

- (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 50

- (Topic 4)

A company is moving its on-premises Oracle database to Amazon Aurora PostgreSQL. The database has several applications that write to the same tables. The applications need to be migrated one by one with a month in between each migration. Management has expressed concerns that the database has a high number of reads and writes. The data must be kept in sync across both databases throughout the migration.

What should a solutions architect recommend?

- A. Use AWS DataSync for the initial migratio
- B. Use AWS Database Migration Service (AWS DMS) to create a change data capture (CDC) replication task and a table mapping to select all tables.
- C. Use AWS DataSync for the initial migratio
- D. Use AWS Database Migration Service (AWS DMS) to create a full load plus change data capture (CDC) replication task and a table mapping to select all tables.
- E. Use the AWS Schema Conversion Tool with AWS Database Migration Service (AWS DMS) using a memory optimized replication instanc
- F. Create a full load plus change data capture (CDC) replication task and a table mapping to select all tables.
- G. Use the AWS Schema Conversion Tool with AWS Database Migration Service (AWS DMS) using a compute optimized replication instanc
- H. Create a full load plus change data capture (CDC) replication task and a table mapping to select the largest tables.

Answer: C

Explanation:

<https://aws.amazon.com/ko/premiumsupport/knowledge-center/dms-memory-optimization/>

NEW QUESTION 55

- (Topic 4)

A company has a popular gaming platform running on AWS. The application is sensitive to latency because latency can impact the user experience and introduce unfair advantages to some players. The application is deployed in every AWS Region. It runs on Amazon EC2 instances that are part of Auto Scaling groups configured behind Application Load Balancers (ALBs). A solutions architect needs to implement a mechanism to monitor the health of the application and redirect traffic to healthy endpoints.

Which solution meets these requirements?

- A. Configure an accelerator in AWS Global Accelerator
- B. Add a listener for the port that the application listens on, and attach it to a Regional endpoint in each Region
- C. Add the ALB as the endpoint.
- D. Create an Amazon CloudFront distribution and specify the ALB as the origin server
- E. Configure the cache behavior to use origin cache header
- F. Use AWS Lambda functions to optimize the traffic.
- G. Create an Amazon CloudFront distribution and specify Amazon S3 as the origin server
- H. Configure the cache behavior to use origin cache header
- I. Use AWS Lambda functions to optimize the traffic.
- J. Configure an Amazon DynamoDB database to serve as the data store for the application
- K. Create a DynamoDB Accelerator (DAX) cluster to act as the in-memory cache for DynamoDB hosting the application data.

Answer: A

Explanation:

AWS Global Accelerator directs traffic to the optimal healthy endpoint based on health checks, it can also route traffic to the closest healthy endpoint based on geographic location of the client. By configuring an accelerator and attaching it to a Regional endpoint in each Region, and adding the ALB as the endpoint, the solution will redirect traffic to healthy endpoints, improving the user experience by reducing latency and ensuring that the application is running optimally. This solution will ensure that traffic is directed to the closest healthy endpoint and will help to improve the overall user experience.

NEW QUESTION 56

- (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 58

- (Topic 4)

A company wants to rearchitect a large-scale web application to a serverless microservices architecture. The application uses Amazon EC2 instances and is written in Python.

The company selected one component of the web application to test as a microservice. The component supports hundreds of requests each second. The company wants to create and test the microservice on an AWS solution that supports Python. The solution must also scale automatically and require minimal infrastructure and minimal operational support.

Which solution will meet these requirements?

- A. Use a Spot Fleet with auto scaling of EC2 instances that run the most recent Amazon Linux operating system.
- B. Use an AWS Elastic Beanstalk web server environment that has high availability configured.
- C. Use Amazon Elastic Kubernetes Service (Amazon EKS). Launch Auto Scaling groups of self-managed EC2 instances.
- D. Use an AWS Lambda function that runs custom developed code.

Answer: D

Explanation:

AWS Lambda is a serverless compute service that lets you run code without provisioning or managing servers. You can use Lambda to create and test microservices that are written in Python or other supported languages. Lambda scales automatically to handle the number of requests per second. You only pay for the compute time you consume. Lambda also integrates with other AWS services, such as Amazon API Gateway, Amazon S3, Amazon DynamoDB, and Amazon SQS, to enable event-driven architectures. Lambda has minimal infrastructure and operational overhead, as you do not need to manage servers, operating systems, patches, or scaling policies.

The other options are not serverless solutions and require more infrastructure and operational support. They also do not scale automatically to handle the number of requests per second. A Spot Fleet is a collection of EC2 instances that run on spare capacity at low prices. However, Spot Instances can be interrupted by AWS at any time, which can affect the availability and performance of your microservice. AWS Elastic Beanstalk is a service that automates the deployment and management of web applications on EC2 instances. However, you still need to provision, configure, and monitor the underlying EC2 instances and load balancers. Amazon EKS is a service that runs Kubernetes on AWS. However, you still need to create, configure, and manage the EC2 instances that form the Kubernetes

cluster and nodes. You also need to install and update the Kubernetes software and tools. References:

- ? What is AWS Lambda?
- ? Building Lambda functions with Python
- ? Create a layer for a Lambda Python function
- ? AWS Lambda – Function in Python
- ? How do I call my AWS Lambda function from a local python script?

NEW QUESTION 59

- (Topic 4)

A company has hired a solutions architect to design a reliable architecture for its application. The application consists of one Amazon RDS DB instance and two manually provisioned Amazon EC2 instances that run web servers. The EC2 instances are located in a single Availability Zone. An employee recently deleted the DB instance, and the application was unavailable for 24 hours as a result. The company is concerned with the overall reliability of its environment. What should the solutions architect do to maximize reliability of the application's infrastructure?

- A. Delete one EC2 instance and enable termination protection on the other EC2 instance
- B. Update the DB instance to be Multi-AZ, and enable deletion protection.
- C. Update the DB instance to be Multi-AZ, and enable deletion protection
- D. Place the EC2 instances behind an Application Load Balancer, and run them in an EC2 Auto Scaling group across multiple Availability Zones.
- E. Create an additional DB instance along with an Amazon API Gateway and an AWS Lambda function
- F. Configure the application to invoke the Lambda function through API Gateway
- G. Have the Lambda function write the data to the two DB instances.
- H. Place the EC2 instances in an EC2 Auto Scaling group that has multiple subnets located in multiple Availability Zones
- I. Use Spot Instances instead of On-Demand Instance
- J. Set up Amazon CloudWatch alarms to monitor the health of the instance
- K. Update the DB instance to be Multi-AZ, and enable deletion protection.

Answer: B

Explanation:

This answer is correct because it meets the requirements of maximizing the reliability of the application's infrastructure. You can update the DB instance to be Multi-AZ, which means that Amazon RDS automatically provisions and maintains a synchronous standby replica in a different Availability Zone. The primary DB instance is synchronously replicated across Availability Zones to a standby replica to provide data redundancy and minimize latency spikes during system backups. Running a DB instance with high availability can enhance availability during planned system maintenance. It can also help protect your databases against DB instance failure and Availability Zone disruption. You can also enable deletion protection on the DB instance, which prevents the DB instance from being deleted by any user. You can place the EC2 instances behind an Application Load Balancer, which distributes incoming application traffic across multiple targets, such as EC2 instances, in multiple Availability Zones. This increases the availability and fault tolerance of your applications. You can run the EC2 instances in an EC2 Auto Scaling group across multiple Availability Zones, which ensures that you have the correct number of EC2 instances available to handle the load for your application. You can use scaling policies to adjust the number of instances in your Auto Scaling group in response to changing demand.

References:

- ? <https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Concepts.MultiAZSingleStandby.html>
- ? https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_DeleteInstance.html#USER_DeleteInstance.DeletionProtection
- ? <https://docs.aws.amazon.com/elasticloadbalancing/latest/application/introduction.html>
- ? <https://docs.aws.amazon.com/autoscaling/ec2/userguide/AutoScalingGroup.html>

NEW QUESTION 64

- (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:

To resize images dynamically and serve appropriate formats to clients, a Lambda@Edge function with an external image management library can be used. Lambda@Edge allows running custom code at the edge locations of CloudFront, which can process the images on the fly and optimize them for different devices and browsers. An external image management library can provide various image manipulation and optimization features. References:

- ? Lambda@Edge
- ? Resizing Images with Amazon CloudFront & Lambda@Edge

NEW QUESTION 65

- (Topic 4)

A company manages AWS accounts in AWS Organizations. AWS IAM Identity Center (AWS Single Sign-On) and AWS Control Tower are configured for the accounts. The company wants to manage multiple user permissions across all the accounts. The permissions will be used by multiple IAM users and must be split between the developer and administrator teams. Each team requires different permissions. The company wants a solution that includes new users that are hired on both teams. Which solution will meet these requirements with the LEAST operational overhead?

- A. Create individual users in IAM Identity Center for each account
- B. Create separate developer and administrator groups in IAM Identity Center

- C. Assign the users to the appropriate groups Create a custom 1AM policy for each group to set fine-grained permissions.
- D. Create individual users in 1AM Identity Center for each account
- E. Create separate developer and administrator groups in 1AM Identity Center
- F. Assign the users to the appropriate group
- G. Attach AWS managed 1AM policies to each user as needed for fine-grained permissions.
- H. Create individual users in 1AM Identity Center Create new developer and administrator groups in 1AM Identity Center
- I. Create new permission sets that include the appropriate 1AM policies for each group
- J. Assign the new groups to the appropriate accounts Assign the new permission sets to the new groups When new users are hired, add them to the appropriate group.
- K. Create individual users in 1AM Identity Center
- L. Create new permission sets that include the appropriate 1AM policies for each user
- M. Assign the users to the appropriate account
- N. Grant additional 1AM permissions to the users from within specific account
- O. When new users are hired, add them to 1AM Identity Center and assign them to the accounts.

Answer: C

Explanation:

This solution meets the requirements with the least operational overhead because it leverages the features of IAM Identity Center and AWS Control Tower to centrally manage multiple user permissions across all the accounts. By creating new groups and permission sets, the company can assign fine-grained permissions to the developer and administrator teams based on their roles and responsibilities. The permission sets are applied to the groups at the organization level, so they are automatically inherited by all the accounts in the organization. When new users are hired, the company only needs to add them to the appropriate group in IAM Identity Center, and they will automatically get the permissions assigned to that group. This simplifies the user management and reduces the manual effort of assigning permissions to each user individually.

References:

- ? Managing access to AWS accounts and applications
- ? Managing permission sets
- ? Managing groups

NEW QUESTION 70

- (Topic 4)

A company has a nightly batch processing routine that analyzes report files that an on-premises file system receives daily through SFTP. The company wants to move the solution to the AWS Cloud. The solution must be highly available and resilient. The solution also must minimize operational effort.

Which solution meets these requirements?

- A. Deploy AWS Transfer for SFTP and an Amazon Elastic File System (Amazon EFS) file system for storage
- B. Use an Amazon EC2 instance in an Auto Scaling group with a scheduled scaling policy to run the batch operation.
- C. Deploy an Amazon EC2 instance that runs Linux and an SFTP service
- D. Use an Amazon Elastic Block Store (Amazon EBS) volume for storage
- E. Use an Auto Scaling group with the minimum number of instances and desired number of instances set to 1.
- F. Deploy an Amazon EC2 instance that runs Linux and an SFTP service
- G. Use an Amazon Elastic File System (Amazon EFS) file system for storage
- H. Use an Auto Scaling group with the minimum number of instances and desired number of instances set to 1.
- I. Deploy AWS Transfer for SFTP and an Amazon S3 bucket for storage
- J. Modify the application to pull the batch files from Amazon S3 to an Amazon EC2 instance for processing
- K. Use an EC2 instance in an Auto Scaling group with a scheduled scaling policy to run the batch operation.

Answer: D

Explanation:

The solution that meets the requirements of high availability, performance, security, and static IP addresses is to use Amazon CloudFront, Application Load Balancers (ALBs), Amazon Route 53, and AWS WAF. This solution allows the company to distribute its HTTP-based application globally using CloudFront, which is a content delivery network (CDN) service that caches content at edge locations and provides static IP addresses for each edge location. The company can also use Route 53 latency-based routing to route requests to the closest ALB in each Region, which balances the load across the EC2 instances. The company can also deploy AWS WAF on the CloudFront distribution to protect the application against common web exploits by creating rules that allow, block, or count web requests based on conditions that are defined. The other solutions do not meet all the requirements because they either use Network Load Balancers (NLBs), which do not support HTTP-based applications, or they do not use CloudFront, which provides better performance and security than AWS Global Accelerator.

References :=

- ? Amazon CloudFront
- ? Application Load Balancer
- ? Amazon Route 53
- ? AWS WAF

NEW QUESTION 71

- (Topic 4)

A company wants to run its payment application on AWS. The application receives payment notifications from mobile devices. Payment notifications require a basic validation before they are sent for further processing.

The backend processing application is long running and requires compute and memory to be adjusted. The company does not want to manage the infrastructure. Which solution will meet these requirements with the LEAST operational overhead?

- A. Create an Amazon Simple Queue Service (Amazon SQS) queue. Integrate the queue with an Amazon EventBridge rule to receive payment notifications from mobile devices. Configure the rule to validate payment notifications and send the notifications to the backend application. Deploy the backend application on Amazon Elastic Kubernetes Service (Amazon EKS). Anywhere. Create a standalone cluster.
- B. Create an Amazon API Gateway API. Integrate the API with an AWS Step Functions state machine to receive payment notifications from mobile devices. Invoke the state machine to validate payment notifications and send the notifications to the backend application. Deploy the backend application on Amazon Elastic Kubernetes Service (Amazon EKS). Configure an EKS cluster with self-managed nodes.
- C. Create an Amazon Simple Queue Service (Amazon SQS) queue. Integrate the queue with an Amazon EventBridge rule to receive payment notifications from mobile devices. Configure the rule to validate payment notifications and send the notifications to the backend application. Deploy the backend application on Amazon EC2 Spot Instances. Configure a Spot Fleet with a default allocation strategy.
- D. Create an Amazon API Gateway API. Integrate the API with AWS Lambda to receive payment notifications from mobile devices. Invoke a Lambda function to validate payment notifications and send the notifications to the backend application. Deploy the backend application on Amazon Elastic Container Service (Amazon ECS).

ECS). Configure Amazon ECS with an AWS Fargate launch type.

Answer: D

Explanation:

This option is the best solution because it allows the company to run its payment application on AWS with minimal operational overhead and infrastructure management. By using Amazon API Gateway, the company can create a secure and scalable API to receive payment notifications from mobile devices. By using AWS Lambda, the company can run a serverless function to validate the payment notifications and send them to the backend application. Lambda handles the provisioning, scaling, and security of the function, reducing the operational complexity and cost. By using Amazon ECS with AWS Fargate, the company can run the backend application on a fully managed container service that scales the compute resources automatically and does not require any EC2 instances to manage. Fargate allocates the right amount of CPU and memory for each container and adjusts them as needed.

* A. Create an Amazon Simple Queue Service (Amazon SQS) queue Integrate the queue with an Amazon EventBridge rule to receive payment notifications from mobile devices Configure the rule to validate payment notifications and send the notifications to the backend application Deploy the backend application on Amazon Elastic Kubernetes Service (Amazon EKS) Anywhere Create a standalone cluster. This option is not optimal because it requires the company to manage the Kubernetes cluster that runs the backend application. Amazon EKS Anywhere is a deployment option that allows the company to create and operate Kubernetes clusters on-premises or in other environments outside AWS. The company would need to provision, configure, scale, patch, and monitor the cluster nodes, which can increase the operational overhead and complexity. Moreover, the company would need to ensure the connectivity and security between the AWS services and the EKS Anywhere cluster, which can also add challenges and risks.

* B. Create an Amazon API Gateway API Integrate the API with an AWS Step Functions state machine to receive payment notifications from mobile devices Invoke the state machine to validate payment notifications and send the notifications to the backend application Deploy the backend application on Amazon Elastic Kubernetes Service (Amazon EKS). Configure an EKS cluster with self-managed nodes. This option is not ideal because it requires the company to manage the EC2 instances that host the Kubernetes cluster that runs the backend application. Amazon EKS is a fully managed service that runs Kubernetes on AWS, but it still requires the company to manage the worker nodes that run the containers. The company would need to provision, configure, scale, patch, and monitor the EC2 instances, which can increase the operational overhead and infrastructure costs. Moreover, using AWS Step Functions to validate the payment notifications may be unnecessary and complex, as the validation logic can be implemented in a simpler way with Lambda or other services.

* C. Create an Amazon Simple Queue Service (Amazon SQS) queue Integrate the queue with an Amazon EventBridge rule to receive payment notifications from mobile devices Configure the rule to validate payment notifications and send the notifications to the backend application Deploy the backend application on Amazon EC2 Spot Instances Configure a Spot Fleet with a default allocation strategy. This option is not cost-effective because it requires the company to manage the EC2 instances that run the backend application. The company would need to provision, configure, scale, patch, and monitor the EC2 instances, which can increase the operational overhead and infrastructure costs. Moreover, using Spot Instances can introduce the risk of interruptions, as Spot Instances are reclaimed by AWS when the demand for On-Demand Instances increases. The company would need to handle the interruptions gracefully and ensure the availability and reliability of the backend application.

References:

? 1 Amazon API Gateway - Amazon Web Services

? 2 AWS Lambda - Amazon Web Services

? 3 Amazon Elastic Container Service - Amazon Web Services

? 4 AWS Fargate - Amazon Web Services

NEW QUESTION 74

- (Topic 4)

A company has multiple AWS accounts with applications deployed in the us-west-2 Region Application logs are stored within Amazon S3 buckets in each account

The company wants

to build a centralized log analysis solution that uses a single S3 bucket Logs must not leave us-west-2, and the company wants to incur minimal operational overhead

Which solution meets these requirements and is MOST cost-effective?

A. Create an S3 Lifecycle policy that copies the objects from one of the application S3 buckets to the centralized S3 bucket

B. Use S3 Same-Region Replication to replicate logs from the S3 buckets to another S3 bucket in us-west-2 Use this S3 bucket for log analysis.

C. Write a script that uses the PutObject API operation every day to copy the entire contents of the buckets to another S3 bucket in us-west-2 Use this S3 bucket for log analysis.

D. Write AWS Lambda functions in these accounts that are triggered every time logs are delivered to the S3 buckets (s3 ObjectCreated a event) Copy the logs to another S3 bucket in us-west-2. Use this S3 bucket for log analysis.

Answer: B

Explanation:

This solution meets the following requirements:

? It is cost-effective, as it only charges for the storage and data transfer of the replicated objects, and does not require any additional AWS services or custom scripts. S3 Same-Region Replication (SRR) is a feature that automatically replicates objects across S3 buckets within the same AWS Region. SRR can help you aggregate logs from multiple sources to a single destination for analysis and auditing. SRR also preserves the metadata, encryption, and access control of the source objects.

? It is operationally efficient, as it does not require any manual intervention or scheduling. SRR replicates objects as soon as they are uploaded to the source bucket, ensuring that the destination bucket always has the latest log data. SRR also handles any updates or deletions of the source objects, keeping the destination bucket in sync. SRR can be enabled with a few clicks in the S3 console or with a simple API call.

? It is secure, as it does not allow the logs to leave the us-west-2 Region. SRR only replicates objects within the same AWS Region, ensuring that the data sovereignty and compliance requirements are met. SRR also supports encryption of the source and destination objects, using either server-side encryption with AWS KMS or S3-managed keys, or client-side encryption.

References:

? Same-Region Replication - Amazon Simple Storage Service

? How do I replicate objects across S3 buckets in the same AWS Region?

? Centralized Logging on AWS | AWS Solutions | AWS Solutions Library

NEW QUESTION 79

- (Topic 4)

A company wants to run its experimental workloads in the AWS Cloud. The company has a budget for cloud spending. The company's CFO is concerned about cloud spending accountability for each department. The CFO wants to receive notification when the spending threshold reaches 60% of the budget.

Which solution will meet these requirements?

A. Use cost allocation tags on AWS resources to label owner

B. Create usage budgets in AWS Budget

C. Add an alert threshold to receive notification when spending exceeds 60% of the budget.

- D. Use AWS Cost Explorer forecasts to determine resource owner
- E. Use AWS Cost Anomaly Detection to create alert threshold notifications when spending exceeds 60% of the budget.
- F. Use cost allocation tags on AWS resources to label owner
- G. Use AWS Support API on AWS Trusted Advisor to create alert threshold notifications when spending exceeds 60% of the budget
- H. Use AWS Cost Explorer forecasts to determine resource owner
- I. Create usage budgets in AWS Budget
- J. Add an alert threshold to receive notification when spending exceeds 60% of the budget.

Answer: A

Explanation:

This solution meets the requirements because it allows the company to track and manage its cloud spending by using cost allocation tags to assign costs to different departments, creating usage budgets to set spending limits, and adding alert thresholds to receive notifications when the spending reaches a certain percentage of the budget. This way, the company can monitor its experimental workloads and avoid overspending on the cloud.

References:

- ? Using Cost Allocation Tags
- ? Creating an AWS Budget
- ? Creating an Alert for an AWS Budget

NEW QUESTION 84

- (Topic 4)

A company is building a solution that will report Amazon EC2 Auto Scaling events across all the applications in an AWS account. The company needs to use a serverless solution to store the EC2 Auto Scaling status data in Amazon S3. The company then will use the data in Amazon S3 to provide near-real-time updates in a dashboard. The solution must not affect the speed of EC2 instance launches.

How should the company move the data to Amazon S3 to meet these requirements?

- A. Use an Amazon CloudWatch metric stream to send the EC2 Auto Scaling status data to Amazon Kinesis Data Firehose
- B. Store the data in Amazon S3.
- C. Launch an Amazon EMR cluster to collect the EC2 Auto Scaling status data and send the data to Amazon Kinesis Data Firehose
- D. Store the data in Amazon S3.
- E. Create an Amazon EventBridge rule to invoke an AWS Lambda function on a schedule
- F. Configure the Lambda function to send the EC2 Auto Scaling status data directly to Amazon S3.
- G. Use a bootstrap script during the launch of an EC2 instance to install Amazon Kinesis Agent
- H. Configure Kinesis Agent to collect the EC2 Auto Scaling status data and send the data to Amazon Kinesis Data Firehose
- I. Store the data in Amazon S3.

Answer: A

Explanation:

You can use metric streams to continually stream CloudWatch metrics to a destination of your choice, with near-real-time delivery and low latency. One of the use cases is Data Lake: create a metric stream and direct it to an Amazon Kinesis Data Firehose delivery stream that delivers your CloudWatch metrics to a data lake such as Amazon S3. <https://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/CloudWatch-Metric-Streams.html>

NEW QUESTION 86

- (Topic 4)

A company's web application that is hosted in the AWS Cloud recently increased in popularity. The web application currently exists on a single Amazon EC2 instance in a single public subnet. The web application has not been able to meet the demand of the increased web traffic.

The company needs a solution that will provide high availability and scalability to meet the increased user demand without rewriting the web application.

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

- A. Replace the EC2 instance with a larger compute optimized instance.
- B. Configure Amazon EC2 Auto Scaling with multiple Availability Zones in private subnets.
- C. Configure a NAT gateway in a public subnet to handle web requests.
- D. Replace the EC2 instance with a larger memory optimized instance.
- E. Configure an Application Load Balancer in a public subnet to distribute web traffic

Answer: BE

Explanation:

These two steps will meet the requirements because they will provide high availability and scalability for the web application without rewriting it. Amazon EC2 Auto Scaling allows you to automatically adjust the number of EC2 instances in response to changes in demand. By configuring Auto Scaling with multiple Availability Zones in private subnets, you can ensure that your web application is distributed across isolated and fault-tolerant locations, and that your instances are not directly exposed to the internet. An Application Load Balancer operates at the application layer and distributes incoming web traffic across multiple targets, such as EC2 instances, containers, or Lambda functions. By configuring an Application Load Balancer in a public subnet, you can enable your web application to handle requests from the internet and route them to the appropriate targets in the private subnets.

References:

- ? What is Amazon EC2 Auto Scaling?
- ? What is an Application Load Balancer?

NEW QUESTION 87

- (Topic 4)

A company wants to migrate its on-premises Microsoft SQL Server Enterprise edition database to AWS. The company's online application uses the database to process transactions. The data analysis team uses the same production database to run reports for analytical processing. The company wants to reduce operational overhead by moving to managed services wherever possible.

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

- A. Migrate to Amazon RDS for Microsoft SQL Server
- B. Use read replicas for reporting purposes.
- C. Migrate to Microsoft SQL Server on Amazon EC2. Use Always On read replicas for reporting purposes.
- D. Migrate to Amazon DynamoDB

- E. Use DynamoDB on-demand replicas for reporting purposes.
- F. Migrate to Amazon Aurora MySQL
- G. Use Aurora read replicas for reporting purposes.

Answer: A

Explanation:

Amazon RDS for Microsoft SQL Server is a fully managed service that offers SQL Server 2014, 2016, 2017, and 2019 editions while offloading database administration tasks such as backups, patching, and scaling. Amazon RDS supports read replicas, which are read-only copies of the primary database that can be used for reporting purposes without affecting the performance of the online application. This solution will meet the requirements with the least operational overhead, as it does not require any code changes or manual intervention.

References:

? 1 provides an overview of Amazon RDS for Microsoft SQL Server and its benefits.

? 2 explains how to create and use read replicas with Amazon RDS.

NEW QUESTION 91

- (Topic 4)

A company stores text files in Amazon S3. The text files include customer chat messages, date and time information, and customer personally identifiable information (PII).

The company needs a solution to provide samples of the conversations to an external service provider for quality control. The external service provider needs to randomly pick sample conversations up to the most recent conversation. The company must not share the customer PII with the external service provider. The solution must scale when the number of customer conversations increases.

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

- A. Create an Object Lambda Access Point
- B. Create an AWS Lambda function that redacts the PII when the function reads the file
- C. Instruct the external service provider to access the Object Lambda Access Point.
- D. Create a batch process on an Amazon EC2 instance that regularly reads all new files, redacts the PII from the files, and writes the redacted files to a different S3 bucket
- E. Instruct the external service provider to access the bucket that does not contain the PII.
- F. Create a web application on an Amazon EC2 instance that presents a list of the files, redacts the PII from the files, and allows the external service provider to download new versions of the files that have the PII redacted.
- G. Create an Amazon DynamoDB table
- H. Create an AWS Lambda function that reads only the data in the files that does not contain PII
- I. Configure the Lambda function to store the non-PII data in the DynamoDB table when a new file is written to Amazon S3. Grant the external service provider access to the DynamoDB table.

Answer: A

Explanation:

The correct solution is to create an Object Lambda Access Point and an AWS Lambda function that redacts the PII when the function reads the file. This way, the company can use the S3 Object Lambda feature to modify the S3 object content on the fly, without creating a copy or changing the original object. The external service provider can access the Object Lambda Access Point and get the redacted version of the file. This solution has the least operational overhead because it does not require any additional storage, processing, or synchronization. The solution also scales automatically with the number of customer conversations and the demand from the external service provider. The other options are incorrect because:

? Option B is using a batch process on an EC2 instance to read, redact, and write the files to a different S3 bucket. This solution has more operational overhead because it requires managing the EC2 instance, the batch process, and the additional S3 bucket. It also introduces latency and inconsistency between the original and the redacted files.

? Option C is using a web application on an EC2 instance to present, redact, and download the files. This solution has more operational overhead because it requires managing the EC2 instance, the web application, and the download process. It also exposes the original files to the web application, which increases the risk of leaking the PII.

? Option D is using a DynamoDB table and a Lambda function to store the non-PII data from the files. This solution has more operational overhead because it requires managing the DynamoDB table, the Lambda function, and the data transformation. It also changes the format and the structure of the original files, which may affect the quality control process.

References:

? S3 Object Lambda

? Object Lambda Access Point

? Lambda function

NEW QUESTION 95

- (Topic 4)

A company has an AWS Direct Connect connection from its on-premises location to an AWS account. The AWS account has 30 different VPCs in the same AWS Region. The VPCs use private virtual interfaces (VIFs). Each VPC has a CIDR block that does not overlap with other networks under the company's control.

The company wants to centrally manage the networking architecture while still allowing each VPC to communicate with all other VPCs and on-premises networks.

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

- A. Create a transit gateway and associate the Direct Connect connection with a new transit VIF. Turn on the transit gateway's route propagation feature.
- B. Create a Direct Connect gateway. Recreate the private VIFs to use the new gateway. Associate each VPC by creating new virtual private gateways.
- C. Create a transit VPC. Connect the Direct Connect connection to the transit VPC. Create a peering connection between all other VPCs in the Region. Update the route tables.
- D. Create AWS Site-to-Site VPN connections from on-premises to each VPC. Ensure that both VPN tunnels are UP for each connection. Turn on the route propagation feature.

Answer: A

Explanation:

This solution meets the following requirements:

? It is operationally efficient, as it only requires one transit gateway and one transit VIF to connect the Direct Connect connection to all the VPCs in the same AWS Region. The transit gateway acts as a regional network hub that simplifies the network management and reduces the number of VIFs and gateways needed.

? It is scalable, as it can support up to 5000 attachments per transit gateway, which can include VPCs, VPNs, Direct Connect gateways, and peering connections. The transit gateway can also be connected to other transit gateways in different Regions or accounts using peering connections, enabling cross-Region and cross-

account connectivity.

? It is flexible, as it allows each VPC to communicate with all other VPCs and on-premises networks using dynamic routing protocols such as Border Gateway Protocol (BGP). The transit gateway's route propagation feature automatically propagates the routes from the attached VPCs and VPNs to the transit gateway route table, eliminating the need to manually update the route tables.

References:

? Transit Gateways - Amazon Virtual Private Cloud

? Working with transit gateways - AWS Direct Connect

? Amazon VPC-to-Amazon VPC connectivity options - Amazon Virtual Private Cloud Connectivity Options

NEW QUESTION 99

- (Topic 4)

A company is using a centralized AWS account to store log data in various Amazon S3 buckets. A solutions architect needs to ensure that the data is encrypted at rest before the data is uploaded to the S3 buckets. The data also must be encrypted in transit.

Which solution meets these requirements?

- A. Use client-side encryption to encrypt the data that is being uploaded to the S3 buckets.
- B. Use server-side encryption to encrypt the data that is being uploaded to the S3 buckets.
- C. Create bucket policies that require the use of server-side encryption with S3 managed encryption keys (SSE-S3) for S3 uploads.
- D. Enable the security option to encrypt the S3 buckets through the use of a default AWS Key Management Service (AWS KMS) key.

Answer: A

Explanation:

Client-side encryption is a method of encrypting data before uploading it to Amazon S3. It allows users to manage the encryption process, encryption keys, and related tools¹. By using client-side encryption, the solution can ensure that the data is encrypted at rest and in transit, as Amazon S3 will not have access to the encryption keys or the unencrypted data².

NEW QUESTION 101

- (Topic 4)

A company needs to extract the names of ingredients from recipe records that are stored as text files in an Amazon S3 bucket. A web application will use the ingredient names to query an Amazon DynamoDB table and determine a nutrition score.

The application can handle non-food records and errors. The company does not have any employees who have machine learning knowledge to develop this solution.

Which solution will meet these requirements MOST cost-effectively?

- A. Use S3 Event Notifications to invoke an AWS Lambda function when PutObject requests occur. Program the Lambda function to analyze the object and extract the ingredient names by using Amazon Comprehend. Store the Amazon Comprehend output in the DynamoDB table.
- B. Use an Amazon EventBridge rule to invoke an AWS Lambda function when PutObject requests occur.
- C. Program the Lambda function to analyze the object by using Amazon Forecast to extract the ingredient names. Store the Forecast output in the DynamoDB table.
- D. Use S3 Event Notifications to invoke an AWS Lambda function when PutObject requests occur. Use Amazon Polly to create audio recordings of the recipe record.
- E. Save the audio files in the S3 bucket. Use Amazon Simple Notification Service (Amazon SNS) to send a URL as a message to employees. Instruct the employees to listen to the audio files and calculate the nutrition score. Store the ingredient names in the DynamoDB table.
- F. Use an Amazon EventBridge rule to invoke an AWS Lambda function when a PutObject request occurs. Program the Lambda function to analyze the object and extract the ingredient names by using Amazon SageMaker. Store the inference output from the SageMaker endpoint in the DynamoDB table.

Answer: A

Explanation:

This solution meets the following requirements:

- ? It is cost-effective, as it only uses serverless components that are charged based on usage and do not require any upfront provisioning or maintenance.
- ? It is scalable, as it can handle any number of recipe records that are uploaded to the S3 bucket without any performance degradation or manual intervention.
- ? It is easy to implement, as it does not require any machine learning knowledge or complex data processing logic. Amazon Comprehend is a natural language processing service that can automatically extract entities such as ingredients from text files. The Lambda function can simply invoke the Comprehend API and store the results in the DynamoDB table.
- ? It is reliable, as it can handle non-food records and errors gracefully. Amazon Comprehend can detect the language and domain of the text files and return an appropriate response. The Lambda function can also implement error handling and logging mechanisms to ensure the data quality and integrity.

References:

? Using AWS Lambda with Amazon S3 - AWS Lambda

? What Is Amazon Comprehend? - Amazon Comprehend

? Working with Tables - Amazon DynamoDB

NEW QUESTION 106

- (Topic 4)

A company wants to share accounting data with an external auditor. The data is stored in an Amazon RDS DB instance that resides in a private subnet. The auditor has its own AWS account and requires its own copy of the database.

What is the MOST secure way for the company to share the database with the auditor?

- A. Create a read replica of the database.
- B. Configure IAM standard database authentication to grant the auditor access.
- C. Export the database contents to a text file.
- D. Store the files in an Amazon S3 bucket.
- E. Create a new IAM user for the auditor.
- F. Grant the user access to the S3 bucket.
- G. Copy a snapshot of the database to an Amazon S3 bucket.
- H. Create an IAM user.
- I. Share the user's keys with the auditor to grant access to the object in the S3 bucket.
- J. Create an encrypted snapshot of the database.
- K. Share the snapshot with the auditor.

L. Allow access to the AWS Key Management Service (AWS KMS) encryption key.

Answer: D

Explanation:

This answer is correct because it meets the requirements of sharing the database with the auditor in a secure way. You can create an encrypted snapshot of the database by using AWS Key Management Service (AWS KMS) to encrypt the snapshot with a customer managed key. You can share the snapshot with the auditor by modifying the permissions of the snapshot and specifying the AWS account ID of the auditor. You can also allow access to the AWS KMS encryption key by adding a key policy statement that grants permissions to the auditor's account. This way, you can ensure that only the auditor can access and restore the snapshot in their own AWS account.

References:

? https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_ShareSnapshot.html

? <https://docs.aws.amazon.com/kms/latest/developerguide/key-policies.html#key-policy-default-allow-root-enable-iam>

NEW QUESTION 108

- (Topic 4)

An ecommerce application uses a PostgreSQL database that runs on an Amazon EC2 instance. During a monthly sales event, database usage increases and causes database connection issues for the application. The traffic is unpredictable for subsequent monthly sales events, which impacts the sales forecast. The company needs to maintain performance when there is an unpredictable increase in traffic.

Which solution resolves this issue in the MOST cost-effective way?

- A. Migrate the PostgreSQL database to Amazon Aurora Serverless v2.
- B. Enable auto scaling for the PostgreSQL database on the EC2 instance to accommodate increased usage.
- C. Migrate the PostgreSQL database to Amazon RDS for PostgreSQL with a larger instance type
- D. Migrate the PostgreSQL database to Amazon Redshift to accommodate increased usage

Answer: A

Explanation:

Amazon Aurora Serverless v2 is a cost-effective solution that can automatically scale the database capacity up and down based on the application's needs. It can handle unpredictable traffic spikes without requiring any provisioning or management of database instances. It is compatible with PostgreSQL and offers high performance, availability, and durability¹. References: 1: AWS Ramp-Up Guide: Architect², page 312: AWS Certified Solutions Architect - Associate exam guide³, page 9.

NEW QUESTION 109

- (Topic 4)

A company has an application that processes customer orders. The company hosts the application on an Amazon EC2 instance that saves the orders to an Amazon Aurora database. Occasionally when traffic is high: the workload does not process orders fast enough.

What should a solutions architect do to write the orders reliably to the database as quickly as possible?

- A. Increase the instance size of the EC2 instance when traffic is high
- B. Write orders to Amazon Simple Notification Service (Amazon SNS). Subscribe the database endpoint to the SNS topic.
- C. Write orders to an Amazon Simple Queue Service (Amazon SQS) queue
- D. Use EC2 instances in an Auto Scaling group behind an Application Load Balancer to read from the SQS queue and process orders into the database.
- E. Write orders to Amazon Simple Notification Service (Amazon SNS). Subscribe the database endpoint to the SNS topic. Use EC2 instances in an Auto Scaling group behind an Application Load Balancer to read from the SNS topic.
- F. Write orders to an Amazon Simple Queue Service (Amazon SQS) queue when the EC2 instance reaches CPU threshold limit
- G. Use scheduled scaling of EC2 instances in an Auto Scaling group behind an Application Load Balancer to read from the SQS queue and process orders into the database

Answer: B

Explanation:

Amazon SQS is a fully managed message queuing service that can decouple and scale microservices, distributed systems, and serverless applications. By writing orders to an SQS queue, the application can handle spikes in traffic without losing any orders. The EC2 instances in an Auto Scaling group can read from the SQS queue and process orders into the database at a steady pace. The Application Load Balancer can distribute the load across the EC2 instances and provide health checks. This solution meets all the requirements of the question, while the other options do not. References:

? <https://docs.aws.amazon.com/wellarchitected/latest/reliability-pillar/welcome.html>

? <https://aws.amazon.com/architecture/serverless/>

? <https://aws.amazon.com/sqs/>

NEW QUESTION 113

- (Topic 4)

A company website hosted on Amazon EC2 instances processes classified data stored in The application writes data to Amazon Elastic Block Store (Amazon EBS) volumes The company needs to ensure that all data that is written to the EBS volumes is encrypted at rest.

Which solution will meet this requirement?

- A. Create an IAM role that specifies EBS encryption Attach the role to the EC2 instances
- B. Create the EBS volumes as encrypted volumes Attach the EBS volumes to the EC2 instances
- C. Create an EC2 instance tag that has a key of Encrypt and a value of True Tag all instances that require encryption at the EBS level
- D. Create an AWS Key Management Service (AWS KMS) key policy that enforces EBS encryption in the account Ensure that the key policy is active

Answer: B

Explanation:

The simplest and most effective way to ensure that all data that is written to the EBS volumes is encrypted at rest is to create the EBS volumes as encrypted volumes. You can do this by selecting the encryption option when you create a new EBS volume, or by copying an existing unencrypted volume to a new encrypted volume. You can also specify the AWS KMS key that you want to use for encryption, or use the default AWS-managed key. When you attach the encrypted EBS volumes to the EC2 instances, the data will be automatically encrypted and decrypted by the EC2 host. This solution does not require any additional IAM roles, tags, or policies.

References:

- ? Amazon EBS encryption
- ? Creating an encrypted EBS volume
- ? Encrypting an unencrypted EBS volume

NEW QUESTION 118

- (Topic 4)

A company moved its on-premises PostgreSQL database to an Amazon RDS for PostgreSQL DB instance. The company successfully launched a new product. The workload on the database has increased.

The company wants to accommodate the larger workload without adding infrastructure. Which solution will meet these requirements MOST cost-effectively?

- A. Buy reserved DB instances for the total workload
- B. Make the Amazon RDS for PostgreSQL DB instance larger.
- C. Make the Amazon RDS for PostgreSQL DB instance a Multi-AZ DB instance.
- D. Buy reserved DB instances for the total workload
- E. Add another Amazon RDS for PostgreSQL DB instance.
- F. Make the Amazon RDS for PostgreSQL DB instance an on-demand DB instance.

Answer: A

Explanation:

This answer is correct because it meets the requirements of accommodating the larger workload without adding infrastructure and minimizing the cost. Reserved DB instances are a billing discount applied to the use of certain on-demand DB instances in your account. Reserved DB instances provide you with a significant discount compared to on-demand DB instance pricing. You can buy reserved DB instances for the total workload and choose between three payment options: No Upfront, Partial Upfront, or All Upfront. You can make the Amazon RDS for PostgreSQL DB instance larger by modifying its instance type to a higher performance class. This way, you can increase the CPU, memory, and network capacity of your DB instance and handle the increased workload. References:

? https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_WorkingWithReservedDBInstances.html

? <https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Concepts.DBInstanceClass.html>

NEW QUESTION 122

- (Topic 4)

A company provides an API interface to customers so the customers can retrieve their financial information. The company expects a larger number of requests during peak usage times of the year.

The company requires the API to respond consistently with low latency to ensure customer satisfaction. The company needs to provide a compute host for the API.

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

- A. Use an Application Load Balancer and Amazon Elastic Container Service (Amazon ECS).
- B. Use Amazon API Gateway and AWS Lambda functions with provisioned concurrency.
- C. Use an Application Load Balancer and an Amazon Elastic Kubernetes Service (Amazon EKS) cluster.
- D. Use Amazon API Gateway and AWS Lambda functions with reserved concurrency.

Answer: B

Explanation:

Amazon API Gateway is a fully managed service that makes it easy for developers to create, publish, maintain, monitor, and secure APIs at any scale. AWS Lambda is a serverless compute service that lets you run code without provisioning or managing servers. Lambda scales automatically based on the incoming requests, but it may take some time to initialize new instances of your function if there is a sudden increase in demand. This may result in high latency or cold starts for your API. To avoid this, you can use provisioned concurrency, which ensures that your function is initialized and ready to respond at any time. Provisioned concurrency also helps you achieve consistent low latency for your API by reducing the impact of scaling on performance. References:

<https://docs.aws.amazon.com/apigateway/latest/developerguide/http-api-develop-integrations-lambda.html>

<https://docs.aws.amazon.com/lambda/latest/dg/configuration-concurrency.html>

NEW QUESTION 124

- (Topic 4)

A company hosts multiple production applications. One of the applications consists of resources from Amazon EC2, AWS Lambda, Amazon RDS, Amazon Simple Notification Service (Amazon SNS), and Amazon Simple Queue Service (Amazon SQS) across multiple AWS Regions. All company resources are tagged with a tag name of "application" and a value that corresponds to each application. A solutions architect must provide the quickest solution for identifying all of the tagged components.

Which solution meets these requirements?

- A. Use AWS CloudTrail to generate a list of resources with the application tag.
- B. Use the AWS CLI to query each service across all Regions to report the tagged components.
- C. Run a query in Amazon CloudWatch Logs Insights to report on the components with the application tag.
- D. Run a query with the AWS Resource Groups Tag Editor to report on the resources globally with the application tag.

Answer: D

Explanation:

<https://docs.aws.amazon.com/tag-editor/latest/userguide/tagging.html>

NEW QUESTION 129

- (Topic 4)

A company wants to analyze and troubleshoot Access Denied errors and Unauthorized errors that are related to IAM permissions. The company has AWS CloudTrail turned on. Which solution will meet these requirements with the LEAST effort?

- A. Use AWS Glue and write custom scripts to query CloudTrail logs for the errors
- B. Use AWS Batch and write custom scripts to query CloudTrail logs for the errors

- C. Search CloudTrail logs with Amazon Athena queries to identify the errors
- D. Search CloudTrail logs with Amazon QuickSight
- E. Create a dashboard to identify the errors.

Answer: C

Explanation:

This solution meets the following requirements:

? It is the least effort, as it does not require any additional AWS services, custom scripts, or data processing steps. Amazon Athena is a serverless interactive query service that allows you to analyze data in Amazon S3 using standard SQL. You can use Athena to query CloudTrail logs directly from the S3 bucket where they are stored, without any data loading or transformation. You can also use the AWS Management Console, the AWS CLI, or the Athena API to run and manage your queries.

? It is effective, as it allows you to filter, aggregate, and join CloudTrail log data using SQL syntax. You can use various SQL functions and operators to specify the criteria for identifying Access Denied and Unauthorized errors, such as the error code, the user identity, the event source, the event name, the event time, and the resource ARN. You can also use subqueries, views, and common table expressions to simplify and optimize your queries.

? It is flexible, as it allows you to customize and save your queries for future use.

You can also export the query results to other formats, such as CSV or JSON, or integrate them with other AWS services, such as Amazon QuickSight, for further analysis and visualization.

References:

? Querying AWS CloudTrail Logs - Amazon Athena

? Analyzing Data in S3 using Amazon Athena | AWS Big Data Blog

? Troubleshoot IAM permission access denied or unauthorized errors | AWS re:Post

NEW QUESTION 133

- (Topic 4)

A company hosts its application in the AWS Cloud. The application runs on Amazon EC2 instances behind an Elastic Load Balancer in an Auto Scaling group and with an Amazon DynamoDB table. The company wants to ensure the application can be made available in another AWS Region with minimal downtime.

What should a solutions architect do to meet these requirements with the LEAST amount of downtime?

- A. Create an Auto Scaling group and a load balancer in the disaster recovery Region
- B. Configure the DynamoDB table as a global table
- C. Configure DNS failover to point to the new disaster recovery Region's load balancer.
- D. Create an AWS CloudFormation template to create EC2 instances, load balancers, and DynamoDB tables to be launched when needed
- E. Configure DNS failover to point to the new disaster recovery Region's load balancer.
- F. Create an AWS CloudFormation template to create EC2 instances and a load balancer to be launched when needed
- G. Configure the DynamoDB table as a global table
- H. Configure DNS failover to point to the new disaster recovery Region's load balancer.
- I. Create an Auto Scaling group and load balancer in the disaster recovery Region
- J. Configure the DynamoDB table as a global table
- K. Create an Amazon CloudWatch alarm to trigger an AWS Lambda function that updates Amazon Route 53 pointing to the disaster recovery load balancer.

Answer: A

Explanation:

This answer is correct because it meets the requirements of securely migrating the existing data to AWS and satisfying the new regulation. AWS DataSync is a service that makes it easy to move large amounts of data online between on-premises storage and Amazon S3. DataSync automatically encrypts data in transit and verifies data integrity during transfer. AWS CloudTrail is a service that records AWS API calls for your account and delivers log files to Amazon S3. CloudTrail can log data events, which show the resource operations performed on or within a resource in your AWS account, such as S3 object-level API activity. By using CloudTrail to log data events, you can audit access at all levels of the stored data.

References:

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

? <https://docs.aws.amazon.com/awscloudtrail/latest/userguide/logging-data-events-with-cloudtrail.html>

NEW QUESTION 134

- (Topic 4)

A company has stored 10 TB of log files in Apache Parquet format in an Amazon S3 bucket. The company occasionally needs to use SQL to analyze the log files.

Which solution will meet these requirements MOST cost-effectively?

- A. Create an Amazon Aurora MySQL database. Migrate the data from the S3 bucket into Aurora by using AWS Database Migration Service (AWS DMS). Issue SQL statements to the Aurora database.
- B. Create an Amazon Redshift cluster. Use Redshift Spectrum to run SQL statements directly on the data in the S3 bucket.
- C. Create an AWS Glue crawler to store and retrieve table metadata from the S3 bucket. Use Amazon Athena to run SQL statements directly on the data in the S3 bucket.
- D. Create an Amazon EMR cluster. Use Apache Spark SQL to run SQL statements directly on the data in the S3 bucket.

Answer: C

Explanation:

AWS Glue is a serverless data integration service that can crawl, catalog, and prepare data for analysis. AWS Glue can automatically discover the schema and partitioning of the data stored in Apache Parquet format in S3, and create a table in the AWS Glue Data Catalog. Amazon Athena is a serverless interactive query service that can run SQL queries directly on data in S3, without requiring any data loading or transformation. Athena can use the table metadata from the AWS Glue Data Catalog to query the data in S3. By using AWS Glue and Athena, you can analyze the log files in S3 most cost-effectively, as you only pay for the resources consumed by the crawler and the queries, and you do not need to provision or manage any servers or clusters.

References:

? AWS Glue

? Amazon Athena

? Analyzing Data in S3 using Amazon Athena

NEW QUESTION 135

- (Topic 4)

A company stores data in PDF format in an Amazon S3 bucket. The company must follow a legal requirement to retain all new and existing data in Amazon S3 for 7 years.

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

- A. Turn on the S3 Versioning feature for the S3 bucket. Configure S3 Lifecycle to delete the data after 7 years.
- B. Configure multi-factor authentication (MFA) delete for all S3 objects.
- C. Turn on S3 Object Lock with governance retention mode for the S3 bucket. Set the retention period to expire after 7 years.
- D. Recopy all existing objects to bring the existing data into compliance.
- E. Turn on S3 Object Lock with compliance retention mode for the S3 bucket.
- F. Set the retention period to expire after 7 years.
- G. Recopy all existing objects to bring the existing data into compliance.
- H. Turn on S3 Object Lock with compliance retention mode for the S3 bucket.
- I. Set the retention period to expire after 7 years.
- J. Use S3 Batch Operations to bring the existing data into compliance.

Answer: C

Explanation:

S3 Object Lock enables a write-once-read-many (WORM) model for objects stored in Amazon S3. It can help prevent objects from being deleted or overwritten for a fixed amount of time or indefinitely¹. S3 Object Lock has two retention modes: governance mode and compliance mode. Compliance mode provides the highest level of protection and prevents any user, including the root user, from deleting or modifying an object version until the retention period expires. To use S3 Object Lock, a new bucket with Object Lock enabled must be created, and a default retention period can be optionally configured for objects placed in the bucket². To bring existing objects into compliance, they must be recopied into the bucket with a retention period specified.

Option A is incorrect because S3 Versioning and S3 Lifecycle do not provide WORM protection for objects. Moreover, MFA delete only applies to deleting object versions, not modifying them.

Option B is incorrect because governance mode allows users with special permissions to override or remove the retention settings or delete the object if necessary. This does not meet the legal requirement of retaining all data for 7 years.

Option D is incorrect because S3 Batch Operations cannot be used to apply compliance mode retention periods to existing objects. S3 Batch Operations can only apply governance mode retention periods or legal holds. Reference URL: 2: <https://docs.aws.amazon.com/AmazonS3/latest/userguide/object-lock-console.html> 3: <https://docs.aws.amazon.com/AmazonS3/latest/userguide/storage-class-intro.html#sc-dynamic-data-access> 4:

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/transfer-acceleration.html> 1: <https://docs.aws.amazon.com/AmazonS3/latest/userguide/object-lock.html> : <https://docs.aws.amazon.com/AmazonS3/latest/userguide/object-lock-overview.html> : <https://docs.aws.amazon.com/AmazonS3/latest/userguide/object-lock-managing.html> : <https://aws.amazon.com/blogs/storage/managing-amazon-s3-access-with-vpc-endpoints-and-s3-access-points/>

NEW QUESTION 138

- (Topic 4)

A company uses Amazon EC2 instances and Amazon Elastic Block Store (Amazon EBS) volumes to run an application. The company creates one snapshot of each EBS volume every day to meet compliance requirements. The company wants to implement an architecture that prevents the accidental deletion of EBS volume snapshots. The solution must not change the administrative rights of the storage administrator user.

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

- A. Create an IAM role that has permission to delete snapshots.
- B. Attach the role to a new EC2 instance.
- C. Use the AWS CLI from the new EC2 instance to delete snapshots.
- D. Create an IAM policy that denies snapshot deletion.
- E. Attach the policy to the storage administrator user.
- F. Add tags to the snapshot.
- G. Create retention rules in Recycle Bin for EBS snapshots that have the tags.
- H. Lock the EBS snapshots to prevent deletion.

Answer: D

Explanation:

EBS snapshots are point-in-time backups of EBS volumes that can be used to restore data or create new volumes. EBS snapshots can be locked to prevent accidental deletion using a feature called EBS Snapshot Lock. When a snapshot is locked, it cannot be deleted by any user, including the root user, until it is unlocked. The lock policy can also specify a retention period, after which the snapshot can be deleted. This solution will meet the requirements with the least administrative effort, as it does not require any code development or policy changes.

References:

? 1 explains how to lock and unlock EBS snapshots using EBS Snapshot Lock.

? 2 describes the concept and benefits of EBS snapshots.

NEW QUESTION 140

- (Topic 4)

The DNS provider that hosts a company's domain name records is experiencing outages that cause service disruption for a website running on AWS. The company needs to

migrate to a more resilient managed DNS service and wants the service to run on AWS. What should a solutions architect do to rapidly migrate the DNS hosting service?

- A. Create an Amazon Route 53 public hosted zone for the domain name.
- B. Import the zone file containing the domain records hosted by the previous provider.
- C. Create an Amazon Route 53 private hosted zone for the domain name. Import the zone file containing the domain records hosted by the previous provider.
- D. Create a Simple AD directory in AWS.
- E. Enable zone transfer between the DNS provider and AWS Directory Service for Microsoft Active Directory for the domain records.
- F. Create an Amazon Route 53 Resolver inbound endpoint in the VPC.
- G. Specify the IP addresses that the provider's DNS will forward DNS queries to.
- H. Configure the provider's DNS to forward DNS queries for the domain to the IP addresses that are specified in the inbound endpoint.

Answer: A

Explanation:

To migrate the DNS hosting service to a more resilient managed DNS service on AWS, the company should use Amazon Route 53, which is a highly available

and scalable cloud DNS web service. Route 53 can host public DNS records for the company's domain name and provide reliable and secure DNS resolution. To rapidly migrate the DNS hosting service, the company should create a public hosted zone for the domain name in Route 53, which is a container for the domain's DNS records. Then, the company should import the zone file containing the domain records hosted by the previous provider, which is a text file that defines the DNS records for the domain. This way, the company can quickly transfer the existing DNS records to Route 53 without manually creating them. After importing the zone file, the company should update the domain registrar to use the name servers that Route 53 assigns to the hosted zone. This will ensure that DNS queries for the domain name are routed to Route 53 and resolved by the imported records.

NEW QUESTION 145

- (Topic 3)

A company is developing a real-time multiplayer game that uses UDP for communications between the client and servers. In an Auto Scaling group, spikes in demand are anticipated during the day, so the game server platform must adapt accordingly. Developers want to store gamer scores and other non-relational data in a database solution that will scale without intervention.

Which solution should a solutions architect recommend?

- A. Use Amazon Route 53 for traffic distribution and Amazon Aurora Serverless for data storage.
- B. Use a Network Load Balancer for traffic distribution and Amazon DynamoDB on-demand for data storage.
- C. Use a Network Load Balancer for traffic distribution and Amazon Aurora Global Database for data storage.
- D. Use an Application Load Balancer for traffic distribution and Amazon DynamoDB global tables for data storage.

Answer: B

Explanation:

A Network Load Balancer is a type of load balancer that operates at the connection level (Layer 4) and can load balance both TCP and UDP traffic¹. A Network Load Balancer is

suitable for scenarios where high performance and low latency are required, such as real-time multiplayer games¹. A Network Load Balancer can also handle sudden and volatile traffic patterns while using a single static IP address per Availability Zone¹.

To meet the requirements of the scenario, the solutions architect should use a Network Load Balancer for traffic distribution between the EC2 instances in the Auto Scaling

group. The Network Load Balancer can route UDP traffic from the client to the servers on the appropriate port². The Network Load Balancer can also support TLS offloading for secure communications between the client and servers¹.

Amazon DynamoDB is a fully managed NoSQL database service that can store and retrieve any amount of data with consistent performance and low latency³.

Amazon DynamoDB on-demand is a flexible billing option that requires no capacity planning and charges only for the read and write requests that are performed on the tables³. Amazon DynamoDB on-demand is ideal for scenarios where the application traffic is unpredictable or sporadic, such as gaming applications³.

To meet the requirements of the scenario, the solutions architect should use Amazon DynamoDB on-demand for data storage. Amazon DynamoDB on-demand can store gamer scores and other non-relational data without intervention from the developers. Amazon DynamoDB on-demand can also scale automatically to handle any level of request traffic without affecting performance or availability³.

NEW QUESTION 147

.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual AWS-Solution-Architect-Associate Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the AWS-Solution-Architect-Associate Product From:

<https://www.2passeasy.com/dumps/AWS-Solution-Architect-Associate/>

Money Back Guarantee

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