



Amazon

Exam Questions AWS-Certified-Solutions-Architect-Professional

Amazon AWS Certified Solutions Architect Professional

About ExamBible

Your Partner of IT Exam

Found in 1998

ExamBible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, ExamBible has its unique advantages that other companies could not achieve.

Our Advances

* 99.9% Uptime

All examinations will be up to date.

* 24/7 Quality Support

We will provide service round the clock.

* 100% Pass Rate

Our guarantee that you will pass the exam.

* Unique Gurantee

If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

NEW QUESTION 1

- (Exam Topic 2)

A retail company needs to provide a series of data files to another company, which is its business partner. These files are saved in an Amazon S3 bucket under Account A, which belongs to the retail company. The business partner company wants one of its IAM users, User_DataProcessor, to access the files from its own AWS account (Account B).

Which combination of steps must the companies take so that User_DataProcessor can access the S3 bucket successfully? (Select TWO.)

- A. Turn on the cross-origin resource sharing (CORS) feature for the S3 bucket in Account A.
- B. In Account A, set the S3 bucket policy to the following:

```
{
  "Effect": "Allow",
  "Action": [
    "s3:GetObject",
    "s3:ListBucket"
  ],
  "Resource": "arn:aws:s3:::AccountABucketName/*"
}
```

- D. In Account A, set the S3 bucket policy to the following:

```
{
  "Effect": "Allow",
  "Principal": {
    "AWS": "arn:aws:iam::AccountB:user/User_DataProcessor"
  },
  "Action": [
    "s3:GetObject",
    "s3:ListBucket"
  ],
  "Resource": [
    "arn:aws:s3:::AccountABucketName/*"
  ]
}
```

- F. In Account B, set the permissions of User_DataProcessor to the following:

```
{
  "Effect": "Allow",
  "Action": [
    "s3:GetObject",
    "s3:ListBucket"
  ],
  "Resource": "arn:aws:s3:::AccountABucketName/*"
}
```

- H. In Account B, set the permissions of User_DataProcessor to the following:

```
{
  "Effect": "Allow",
  "Principal": {
    "AWS": "arn:aws:iam::AccountB:user/User_DataProcessor"
  },
  "Action": [
    "s3:GetObject",
    "s3:ListBucket"
  ],
  "Resource": [
    "arn:aws:s3:::AccountABucketName/*"
  ]
}
```

Answer: CD

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/cross-account-access-s3/>

NEW QUESTION 2

- (Exam Topic 2)

A company is running a two-tier web-based application in an on-premises data center. The application layer consists of a single server running a stateful application. The application connects to a PostgreSQL database running on a separate server. The application's user base is expected to grow significantly, so the company is migrating the application and database to AWS. The solution will use Amazon Aurora PostgreSQL, Amazon EC2 Auto Scaling, and Elastic Load Balancing.

Which solution will provide a consistent user experience that will allow the application and database tiers to scale?

- A. Enable Aurora Auto Scaling for Aurora Replica

- B. Use a Network Load Balancer with the least outstanding requests routing algorithm and sticky sessions enabled.
- C. Enable Aurora Auto Scaling for Aurora writer
- D. Use an Application Load Balancer with the round robin routing algorithm and sticky sessions enabled.
- E. Enable Aurora Auto Scaling for Aurora Replica
- F. Use an Application Load Balancer with the round robin routing and sticky sessions enabled.
- G. Enable Aurora Scaling for Aurora writer
- H. Use a Network Load Balancer with the least outstanding requests routing algorithm and sticky sessions enabled.

Answer: C

Explanation:

Aurora Auto Scaling enables your Aurora DB cluster to handle sudden increases in connectivity or workload. When the connectivity or workload decreases, Aurora Auto Scaling removes unnecessary Aurora Replicas so that you don't pay for unused provisioned DB instances

NEW QUESTION 3

- (Exam Topic 2)

A company needs to optimize the cost of an AWS environment that contains multiple accounts in an organization in AWS Organizations. The company conducted cost optimization activities 3 years ago and purchased Amazon EC2 Standard Reserved Instances that recently expired. The company needs EC2 instances for 3 more years. Additionally, the company has deployed a new serverless workload. Which strategy will provide the company with the MOST cost savings?

- A. Purchase the same Reserved Instances for an additional 3-year term with All Upfront payment
- B. Purchase a 3-year Compute Savings Plan with All Upfront payment in the management account to cover any additional compute costs.
- C. Purchase a 1-year Compute Savings Plan with No Upfront payment in each member account
- D. Use the Savings Plans recommendations in the AWS Cost Management console to choose the Compute Savings Plan.
- E. Purchase a 3-year EC2 Instance Savings Plan with No Upfront payment in the management account to cover EC2 costs in each AWS Region
- F. Purchase a 3-year Compute Savings Plan with No Upfront payment in the management account to cover any additional compute costs.
- G. Purchase a 3-year EC2 Instance Savings Plan with All Upfront payment in each member account
- H. Use the Savings Plans recommendations in the AWS Cost Management console to choose the EC2 Instance Savings Plan.

Answer: A

Explanation:

The company should purchase the same Reserved Instances for an additional 3-year term with All Upfront payment. The company should purchase a 3-year Compute Savings Plan with All Upfront payment in the management account to cover any additional compute costs. This solution will provide the company with the most cost savings because Reserved Instances and Savings Plans are both pricing models that offer significant discounts compared to On-Demand pricing. Reserved Instances are commitments to use a specific instance type and size in a single Region for a one- or three-year term. You can choose between three payment options:

No Upfront, Partial Upfront, or All Upfront. The more you pay upfront, the greater the discount. Savings Plans are flexible pricing models that offer low prices on EC2 instances, Fargate, and Lambda usage, in exchange for a commitment to a consistent amount of usage (measured in \$/hour) for a one- or three-year term. You can choose between two types of Savings Plans: Compute Savings Plans and EC2 Instance Savings Plans. Compute Savings Plans apply to any EC2 instance regardless of Region, instance family, operating system, or tenancy, including those that are part of EMR, ECS, or EKS clusters, or launched by Fargate or Lambda. EC2 Instance Savings Plans apply to a specific instance family within a Region and provide the most savings. By purchasing the same Reserved Instances for an additional 3-year term with All Upfront payment, the company can lock in the lowest possible price for its EC2 instances that run continuously for 3 years. By purchasing a 3-year Compute Savings Plan with All Upfront payment in the management account, the company can benefit from additional discounts on any other compute usage across its member accounts.

The other options are not correct because:

➤ Purchasing a 1-year Compute Savings Plan with No Upfront payment in each member account would not provide as much cost savings as purchasing a 3-year Compute Savings Plan with All Upfront payment in the management account. A 1-year term offers lower discounts than a 3-year term, and a No Upfront payment option offers lower discounts than an All Upfront payment option. Also, purchasing a Savings Plan in each member account would not allow the company to share the benefits of unused Savings Plan discounts across its organization.

➤ Purchasing a 3-year EC2 Instance Savings Plan with No Upfront payment in the management account to cover EC2 costs in each AWS Region would not provide as much cost savings as purchasing Reserved Instances for an additional 3-year term with All Upfront payment. An EC2 Instance Savings Plan offers lower discounts than Reserved Instances for the same instance family and Region. Also, a No Upfront payment option offers lower discounts than an All Upfront payment option.

➤ Purchasing a 3-year EC2 Instance Savings Plan with All Upfront payment in each member account would not provide as much flexibility or cost savings as purchasing a 3-year Compute Savings Plan with All Upfront payment in the management account. An EC2 Instance Savings Plan applies only to a specific instance family within a Region and does not cover Fargate or Lambda usage. Also, purchasing a Savings Plan in each member account would not allow the company to share the benefits of unused Savings Plan discounts across its organization.

References:

➤ <https://aws.amazon.com/ec2/pricing/reserved-instances/>

➤ <https://aws.amazon.com/savingsplans/>

NEW QUESTION 4

- (Exam Topic 2)

A solutions architect is redesigning a three-tier application that a company hosts on premises. The application provides personalized recommendations based on user profiles. The company already has an AWS account and has configured a VPC to host the application.

The frontend is a Java-based application that runs in on-premises VMs. The company hosts a personalization model on a physical application server and uses TensorFlow to implement the model. The personalization model uses artificial intelligence and machine learning (AI/ML). The company stores user information in a Microsoft SQL Server database. The web application calls the personalization model, which reads the user profiles from the database and provides recommendations.

The company wants to migrate the redesigned application to AWS.

Which solution will meet this requirement with the LEAST operational overhead?

- A. Use AWS Server Migration Service (AWS SMS) to migrate the on-premises physical application server and the web application VMs to AWS
- B. Use AWS Database Migration Service (AWS DMS) to migrate the SQL Server database to Amazon RDS for SQL Server.
- C. Export the personalization model
- D. Store the model artifacts in Amazon S3. Deploy the model to Amazon SageMaker and create an endpoint
- E. Host the Java application in AWS Elastic Beanstalk

- F. Use AWS Database Migration Service (AWS DMS) to migrate the SQL Server database to Amazon RDS for SQL Server.
- G. Use AWS Application Migration Service to migrate the on-premises personalization model and VMs to Amazon EC2 instances in Auto Scaling group
- H. Use AWS Database Migration Service (AWS DMS) to migrate the SQL Server database to an EC2 instance.
- I. Containerize the personalization model and the Java application
- J. Use Amazon Elastic Kubernetes Service (Amazon EKS) managed node groups to deploy the model and the application to Amazon EKS Host the node groups in a VP
- K. Use AWS Database Migration Service (AWS DMS) to migrate the SQL Server database to Amazon RDS for SQL Server.

Answer: B

Explanation:

Amazon SageMaker is a fully managed machine learning service that allows users to build, train, and deploy machine learning models quickly and easily¹. Users can export their existing TensorFlow models and store the model artifacts in Amazon S3, a highly scalable and durable object storage service². Users can then deploy the model to Amazon SageMaker and create an endpoint that can be invoked by the web application to provide recommendations³. This way, the solution can leverage the AI/ML capabilities of Amazon SageMaker without having to rewrite the personalization model.

AWS Elastic Beanstalk is a service that allows users to deploy and manage web applications without worrying about the infrastructure that runs those applications. Users can host their Java application in AWS Elastic Beanstalk and configure it to communicate with the Amazon SageMaker endpoint. This way, the solution can reduce the operational overhead of managing servers, load balancers, scaling, and application health monitoring.

AWS Database Migration Service (AWS DMS) is a service that helps users migrate databases to AWS quickly and securely. Users can use AWS DMS to migrate their SQL Server database to Amazon RDS for SQL Server, a fully managed relational database service that offers high availability, scalability, security, and compatibility. This way, the solution can reduce the operational overhead of managing database servers, backups, patches, and upgrades.

Option A is incorrect because using AWS Server Migration Service (AWS SMS) to migrate the on-premises physical application server and the web application VMs to AWS is not cost-effective or scalable. AWS SMS is a service that helps users migrate on-premises workloads to AWS. However, for this use case, migrating the physical application server and the web application VMs to AWS will not take advantage of the AI/ML capabilities of Amazon SageMaker or the managed services of AWS Elastic Beanstalk and Amazon RDS.

Option C is incorrect because using AWS Application Migration Service to migrate the on-premises personalization model and VMs to Amazon EC2 instances in Auto Scaling groups is not cost-effective or scalable. AWS Application Migration Service is a service that helps users migrate applications from on-premises or other clouds to AWS without making any changes to their applications. However, for this use case, migrating the personalization model and VMs to EC2 instances will not take advantage of the AI/ML capabilities of Amazon SageMaker or the managed services of AWS Elastic Beanstalk and Amazon RDS.

Option D is incorrect because containerizing the personalization model and the Java application and using Amazon Elastic Kubernetes Service (Amazon EKS) managed node groups to deploy them to Amazon EKS is not necessary or cost-effective. Amazon EKS is a service that allows users to run Kubernetes on AWS without needing to install, operate, and maintain their own Kubernetes control plane or nodes. However, for this use case, containerizing and deploying the personalization model and the Java application will not take advantage of the AI/ML capabilities of Amazon SageMaker or the managed services of AWS Elastic Beanstalk. Moreover, using S3 Glacier Deep Archive as a storage class for images will incur a high retrieval fee and latency for accessing them.

NEW QUESTION 5

- (Exam Topic 2)

A company is using AWS CloudFormation to deploy its infrastructure. The company is concerned that, if a production CloudFormation stack is deleted, important data stored in Amazon RDS databases or Amazon EBS volumes might also be deleted.

How can the company prevent users from accidentally deleting data in this way?

- A. Modify the CloudFormation templates to add a DeletionPolicy attribute to RDS and EBS resources.
- B. Configure a stack policy that disallows the deletion of RDS and EBS resources.
- C. Modify IAM policies to deny deleting RDS and EBS resources that are tagged with an "awscloudformation: stack-name" tag.
- D. Use AWS Config rules to prevent deleting RDS and EBS resources.

Answer: A

Explanation:

With the DeletionPolicy attribute you can preserve or (in some cases) backup a resource when its stack is deleted. You specify a DeletionPolicy attribute for each resource that you want to control. If a resource has no DeletionPolicy attribute, AWS CloudFormation deletes the resource by default. To keep a resource when its stack is deleted, specify Retain for that resource. You can use retain for any resource. For example, you can retain a nested stack, Amazon S3 bucket, or EC2 instance so that you can continue to use or modify those resources after you delete their stacks.

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-attribute-deletionpolicy.html>

NEW QUESTION 6

- (Exam Topic 2)

A company has an application that runs on Amazon EC2 instances in an Amazon EC2 Auto Scaling group. The company uses AWS CodePipeline to deploy the application. The instances that run in the Auto Scaling group are constantly changing because of scaling events.

When the company deploys new application code versions, the company installs the AWS CodeDeploy agent on any new target EC2 instances and associates the instances with the CodeDeploy deployment group. The application is set to go live within the next 24 hours.

What should a solutions architect recommend to automate the application deployment process with the LEAST amount of operational overhead?

- A. Configure Amazon EventBridge to invoke an AWS Lambda function when a new EC2 instance is launched into the Auto Scaling group
- B. Code the Lambda function to associate the EC2 instances with the CodeDeploy deployment group.
- C. Write a script to suspend Amazon EC2 Auto Scaling operations before the deployment of new code. When the deployment is complete, create a new AMI and configure the Auto Scaling group's launch template to use the new AMI for new launches
- D. Resume Amazon EC2 Auto Scaling operations.
- E. Create a new AWS CodeBuild project that creates a new AMI that contains the new code. Configure CodeBuild to update the Auto Scaling group's launch template to the new AMI
- F. Run an Amazon EC2 Auto Scaling instance refresh operation.
- G. Create a new AMI that has the CodeDeploy agent installed
- H. Configure the Auto Scaling group's launch template to use the new AMI
- I. Associate the CodeDeploy deployment group with the Auto Scaling group instead of the EC2 instances.

Answer: D

Explanation:

<https://docs.aws.amazon.com/codedeploy/latest/userguide/integrations-aws-auto-scaling.html>

NEW QUESTION 7

- (Exam Topic 2)

A company has an application in the AWS Cloud. The application runs on a fleet of 20 Amazon EC2 instances. The EC2 instances are persistent and store data on multiple attached Amazon Elastic Block Store (Amazon EBS) volumes.

The company must maintain backups in a separate AWS Region. The company must be able to recover the EC2 instances and their configuration within 1 business day, with loss of no more than 1 day's worth of data. The company has limited staff and needs a backup solution that optimizes operational efficiency and cost. The company already has created an AWS CloudFormation template that can deploy the required network configuration in a secondary Region.

Which solution will meet these requirements?

- A. Create a second CloudFormation template that can recreate the EC2 instances in the secondary Region. Run daily multivolume snapshots by using AWS Systems Manager Automation runbook
- B. Copy the snapshots to the secondary Region
- C. In the event of a failure, launch the CloudFormation templates, restore the EBS volumes from snapshots, and transfer usage to the secondary Region.
- D. Use Amazon Data Lifecycle Manager (Amazon DLM) to create daily multivolume snapshots of the EBS volume
- E. In the event of a failure, launch the CloudFormation template and use Amazon DLM to restore the EBS volumes and transfer usage to the secondary Region.
- F. Use AWS Backup to create a scheduled daily backup plan for the EC2 instance
- G. Configure the backup task to copy the backups to a vault in the secondary Region
- H. In the event of a failure, launch the CloudFormation template, restore the instance volumes and configurations from the backup vault, and transfer usage to the secondary Region.
- I. Deploy EC2 instances of the same size and configuration to the secondary Region
- J. Configure AWS DataSync daily to copy data from the primary Region to the secondary Region
- K. In the event of a failure, launch the CloudFormation template and transfer usage to the secondary Region.

Answer: C

Explanation:

Using AWS Backup to create a scheduled daily backup plan for the EC2 instances will enable taking snapshots of the EC2 instances and their attached EBS volumes. Configuring the backup task to copy the backups to a vault in the secondary Region will enable maintaining backups in a separate Region. In the event of a failure, launching the CloudFormation template will enable deploying the network configuration in the secondary Region. Restoring the instance volumes and configurations from the backup vault will enable recovering the EC2 instances and their data. Transferring usage to the secondary Region will enable resuming operations.

NEW QUESTION 8

- (Exam Topic 2)

A solutions architect needs to define a reference architecture for a solution for three-tier applications with web, application, and NoSQL data layers. The reference architecture must meet the following requirements:

- High availability within an AWS Region
- Able to fail over in 1 minute to another AWS Region for disaster recovery
- Provide the most efficient solution while minimizing the impact on the user experience Which combination of steps will meet these requirements? (Select THREE.)

- A. Use an Amazon Route 53 weighted routing policy set to 100/0 across the two selected Region
- B. Set Time to Live (TTL) to 1 hour.
- C. Use an Amazon Route 53 failover routing policy for failover from the primary Region to the disaster recovery Region
- D. Set Time to Live (TTL) to 30 seconds.
- E. Use a global table within Amazon DynamoDB so data can be accessed in the two selected Regions.
- F. Back up data from an Amazon DynamoDB table in the primary Region every 60 minutes and then write the data to Amazon S3. Use S3 Cross-Region replication to copy the data from the primary Region to the disaster recovery Region
- G. Have a script import the data into DynamoDB in a disaster recovery scenario.
- H. Implement a hot standby model using Auto Scaling groups for the web and application layers across multiple Availability Zones in the Region
- I. Use zonal Reserved Instances for the minimum number of servers and On-Demand Instances for any additional resources.
- J. Use Auto Scaling groups for the web and application layers across multiple Availability Zones in the Region
- K. Use Spot Instances for the required resources.

Answer: BCE

Explanation:

The requirements can be achieved by using an Amazon DynamoDB database with a global table. DynamoDB is a NoSQL database so it fits the requirements. A global table also allows both reads and writes to occur in both Regions. For the web and application tiers Auto Scaling groups should be configured. Due to the 1-minute RTO these must be configured in an active/passive state. The best pricing model to lower price but ensure resources are available when needed is to use a combination of zonal reserved instances and on-demand instances. To failover between the Regions, a Route 53 failover routing policy can be configured with a TTL configured on the record of 30 seconds. This will mean clients must resolve against Route 53 every 30 seconds to get the latest record. In a failover scenario the clients would be redirected to the secondary site if the primary site is unhealthy.

NEW QUESTION 9

- (Exam Topic 2)

A company runs an IoT application in the AWS Cloud. The company has millions of sensors that collect data from houses in the United States. The sensors use the MQTT protocol to connect and send data to a custom MQTT broker. The MQTT broker stores the data on a single Amazon EC2 instance. The sensors connect to the broker through the domain named `iot.example.com`. The company uses Amazon Route 53 as its DNS service. The company stores the data in Amazon DynamoDB.

On several occasions, the amount of data has overloaded the MQTT broker and has resulted in lost sensor data. The company must improve the reliability of the solution.

Which solution will meet these requirements?

- A. Create an Application Load Balancer (ALB) and an Auto Scaling group for the MQTT broker
- B. Use the Auto Scaling group as the target for the ALB
- C. Update the DNS record in Route 53 to an alias record
- D. Point the alias record to the ALB
- E. Use the MQTT broker to store the data.
- F. Set up AWS IoT Core to receive the sensor data
- G. Create and configure a custom domain to connect to AWS IoT Core

- H. Update the DNS record in Route 53 to point to the AWS IoT Core Data-ATS endpoint
- I. Configure an AWS IoT rule to store the data.
- J. Create a Network Load Balancer (NLB). Set the MQTT broker as the target
- K. Create an AWS Global Accelerator accelerator
- L. Set the NLB as the endpoint for the accelerator
- M. Update the DNS record in Route 53 to a multivalued answer record
- N. Set the Global Accelerator IP addresses as values
- O. Use the MQTT broker to store the data.
- P. Set up AWS IoT Greengrass to receive the sensor data
- Q. Update the DNS record in Route 53 to point to the AWS IoT Greengrass endpoint
- R. Configure an AWS IoT rule to invoke an AWS Lambda function to store the data.

Answer: A

Explanation:

It describes a solution that uses an Application Load Balancer (ALB) and an Auto Scaling group for the MQTT broker. The ALB distributes incoming traffic across the instances in the Auto Scaling group and allows for automatic scaling based on incoming traffic. The use of an alias record in Route 53 allows for easy updates to the DNS record without changing the IP address. This solution improves the reliability of the MQTT broker by allowing it to automatically scale based on incoming traffic, reducing the likelihood of lost data due to broker overload.

Reference: <https://aws.amazon.com/elasticloadbalancing/applicationloadbalancer/> <https://aws.amazon.com/autoscaling/> <https://aws.amazon.com/route53/>

NEW QUESTION 10

- (Exam Topic 2)

A financial services company loaded millions of historical stock trades into an Amazon DynamoDB table. The table uses on-demand capacity mode. Once each day at midnight, a few million new records are loaded into the table. Application read activity against the table happens in bursts throughout the day, and a limited set of keys are repeatedly looked up. The company needs to reduce costs associated with DynamoDB.

Which strategy should a solutions architect recommend to meet this requirement?

- A. Deploy an Amazon ElastiCache cluster in front of the DynamoDB table.
- B. Deploy DynamoDB Accelerator (DAX). Configure DynamoDB auto scaling
- C. Purchase Savings Plans in Cost Explorer
- D. Use provisioned capacity mode
- E. Purchase Savings Plans in Cost Explorer.
- F. Deploy DynamoDB Accelerator (DAX). Use provisioned capacity mode
- G. Configure DynamoDB auto scaling.

Answer: D

Explanation:

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/HowItWorks.ReadWriteCapacityMode.html>

NEW QUESTION 10

- (Exam Topic 2)

A company has developed a hybrid solution between its data center and AWS. The company uses Amazon VPC and Amazon EC2 instances that send application logs to Amazon CloudWatch. The EC2 instances read data from multiple relational databases that are hosted on premises.

The company wants to monitor which EC2 instances are connected to the databases in near-real time. The company already has a monitoring solution that uses Splunk on premises. A solutions architect needs to determine how to send networking traffic to Splunk.

How should the solutions architect meet these requirements?

- A. Enable VPC flows logs, and send them to CloudWatch
- B. Create an AWS Lambda function to periodically export the CloudWatch logs to an Amazon S3 bucket by using the pre-defined export function
- C. Generate ACCESS_KEY and SECRET_KEY AWS credentials
- D. Configure Splunk to pull the logs from the S3 bucket by using those credentials.
- E. Create an Amazon Kinesis Data Firehose delivery stream with Splunk as the destination
- F. Configure a pre-processing AWS Lambda function with a Kinesis Data Firehose stream processor that extracts individual log events from records sent by CloudWatch Logs subscription filter
- G. Enable VPC flows logs, and send them to CloudWatch
- H. Create a CloudWatch Logs subscription that sends log events to the Kinesis Data Firehose delivery stream.
- I. Ask the company to log every request that is made to the databases along with the EC2 instance IP address
- J. Export the CloudWatch logs to an Amazon S3 bucket
- K. Use Amazon Athena to query the logs grouped by database name
- L. Export Athena results to another S3 bucket
- M. Invoke an AWS Lambda function to automatically send any new file that is put in the S3 bucket to Splunk.
- N. Send the CloudWatch logs to an Amazon Kinesis data stream with Amazon Kinesis Data Analytics for SQL Application
- O. Configure a 1-minute sliding window to collect the event
- P. Create a SQL query that uses the anomaly detection template to monitor any networking traffic anomalies in near-real time
- Q. Send the result to an Amazon Kinesis Data Firehose delivery stream with Splunk as the destination.

Answer: B

Explanation:

<https://docs.aws.amazon.com/firehose/latest/dev/creating-the-stream-to-splunk.html>

NEW QUESTION 12

- (Exam Topic 2)

A company is running an application in the AWS Cloud. The application collects and stores a large amount of unstructured data in an Amazon S3 bucket. The S3 bucket contains several terabytes of data and uses the S3 Standard storage class. The data increases in size by several gigabytes every day.

The company needs to query and analyze the data. The company does not access data that is more than 1 year old. However, the company must retain all the data indefinitely for compliance reasons.

Which solution will meet these requirements MOST cost-effectively?

- A. Use S3 Select to query the data
- B. Create an S3 Lifecycle policy to transition data that is more than 1 year old to S3 Glacier Deep Archive.
- C. Use Amazon Redshift Spectrum to query the data
- D. Create an S3 Lifecycle policy to transition data that is more than 1 year old to S3 Glacier Deep Archive.
- E. Use an AWS Glue Data Catalog and Amazon Athena to query the data
- F. Create an S3 Lifecycle policy to transition data that is more than 1 year old to S3 Glacier Deep Archive.
- G. Use Amazon Redshift Spectrum to query the data
- H. Create an S3 Lifecycle policy to transition data that is more than 1 year old to S3 Intelligent-Tiering.

Answer: C

Explanation:

Generally, unstructured data should be converted structured data before querying them. AWS Glue can do that.

<https://docs.aws.amazon.com/glue/latest/dg/schema-relationalize.html> <https://docs.aws.amazon.com/athena/latest/ug/glue-athena.html>

NEW QUESTION 17

- (Exam Topic 2)

A company is running an application that uses an Amazon ElastiCache for Redis cluster as a caching layer. A recent security audit revealed that the company has configured encryption at rest for ElastiCache. However, the company did not configure ElastiCache to use encryption in transit. Additionally, users can access the cache without authentication.

A solutions architect must make changes to require user authentication and to ensure that the company is using end-to-end encryption.

Which solution will meet these requirements?

- A. Create an AUTH token. Store the token in AWS System Manager Parameter Store, as an encrypted parameter. Create a new cluster with AUTH and configure encryption in transit. Update the application to retrieve the AUTH token from Parameter Store when necessary and to use the AUTH token for authentication.
- B. Create an AUTH token. Store the token in AWS Secrets Manager. Configure the existing cluster to use the AUTH token and configure encryption in transit. Update the application to retrieve the AUTH token from Secrets Manager when necessary and to use the AUTH token for authentication.
- C. Create an SSL certificate. Store the certificate in AWS Secrets Manager. Create a new cluster and configure encryption in transit. Update the application to retrieve the SSL certificate from Secrets Manager when necessary and to use the certificate for authentication.
- D. Create an SSL certificate. Store the certificate in AWS Systems Manager Parameter Store, as an encrypted advanced parameter. Update the existing cluster to configure encryption in transit. Update the application to retrieve the SSL certificate from Parameter Store when necessary and to use the certificate for authentication.

Answer: B

Explanation:

Creating an AUTH token and storing it in AWS Secrets Manager and configuring the existing cluster to use the AUTH token and configure encryption in transit, and updating the application to retrieve the AUTH token from Secrets Manager when necessary and to use the AUTH token for authentication, would meet the requirements for user authentication and end-to-end encryption.

AWS Secrets Manager is a service that enables you to easily rotate, manage, and retrieve database credentials, API keys, and other secrets throughout their lifecycle. Secrets Manager also enables you to encrypt the data and ensure that only authorized users and applications can access it.

By configuring the existing cluster to use the AUTH token and encryption in transit, all data will be encrypted as it is sent over the network, providing additional security for the data stored in ElastiCache.

Additionally, by updating the application to retrieve the AUTH token from Secrets Manager when necessary and to use the AUTH token for authentication, it ensures that only authorized users and applications can access the cache.

Reference:

AWS Secrets Manager documentation: <https://aws.amazon.com/secrets-manager/> Encryption in transit for ElastiCache:

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

Authentication and Authorization for ElastiCache: <https://docs.aws.amazon.com/AmazonElastiCache/latest/red-ug/accessing-elasticache.html>

NEW QUESTION 18

- (Exam Topic 2)

A company needs to migrate its customer transactions database from on premises to AWS. The database resides on an Oracle DB instance that runs on a Linux server. According to a new security requirement, the company must rotate the database password each year.

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

- A. Convert the database to Amazon DynamoDB by using the AWS Schema Conversion Tool (AWS SCT). Store the password in AWS Systems Manager Parameter Store.
- B. Create an Amazon CloudWatch alarm to invoke an AWS Lambda function for yearly password rotation.
- C. Migrate the database to Amazon RDS for Oracle.
- D. Store the password in AWS Secrets Manager.
- E. Turn on automatic rotation.
- F. Configure a yearly rotation schedule.
- G. Migrate the database to an Amazon EC2 instance.
- H. Use AWS Systems Manager Parameter Store to keep and rotate the connection string by using an AWS Lambda function on a yearly schedule.
- I. Migrate the database to Amazon Neptune by using the AWS Schema Conversion Tool (AWS SCT). Create an Amazon CloudWatch alarm to invoke an AWS Lambda function for yearly password rotation.

Answer: B

NEW QUESTION 19

- (Exam Topic 2)

A company has a few AWS accounts for development and wants to move its production application to AWS. The company needs to enforce Amazon Elastic Block Store (Amazon EBS) encryption at rest current production accounts and future production accounts only. The company needs a solution that includes built-in blueprints and guardrails.

Which combination of steps will meet these requirements? (Choose three.)

- A. Use AWS CloudFormation StackSets to deploy AWS Config rules on production accounts.
- B. Create a new AWS Control Tower landing zone in an existing developer account.
- C. Create OUs for account.

- D. Add production and development accounts to production and development OUs, respectively.
- E. Create a new AWS Control Tower landing zone in the company's management account
- F. Add production and development accounts to production and development OU
- G. respectively.
- H. Invite existing accounts to join the organization in AWS Organization
- I. Create SCPs to ensure compliance.
- J. Create a guardrail from the management account to detect EBS encryption.
- K. Create a guardrail for the production OU to detect EBS encryption.

Answer: CDF

Explanation:

<https://docs.aws.amazon.com/controltower/latest/userguide/controls.html> <https://docs.aws.amazon.com/controltower/latest/userguide/strongly-recommended-controls.html#ebs-enable-en> AWS is now transitioning the previous term 'guardrail' new term 'control'.

NEW QUESTION 22

- (Exam Topic 2)

A solutions architect must create a business case for migration of a company's on-premises data center to the AWS Cloud. The solutions architect will use a configuration management database (CMDB) export of all the company's servers to create the case. Which solution will meet these requirements MOST cost-effectively?

- A. Use AWS Well-Architected Tool to import the CMDB data to perform an analysis and generate recommendations.
- B. Use Migration Evaluator to perform an analysis
- C. Use the data import template to upload the data from the CMDB export.
- D. Implement resource matching rule
- E. Use the CMDB export and the AWS Price List Bulk API to query CMDB data against AWS services in bulk.
- F. Use AWS Application Discovery Service to import the CMDB data to perform an analysis.

Answer: B

Explanation:

<https://aws.amazon.com/blogs/architecture/accelerating-your-migration-to-aws/> Build a business case with AWS Migration Evaluator The foundation for a successful migration starts with a defined business objective (for example, growth or new offerings). In order to enable the business drivers, the established business case must then be aligned to a technical capability (increased security and elasticity). AWS Migration Evaluator (formerly known as TSO Logic) can help you meet these objectives. To get started, you can choose to upload exports from third-party tools such as Configuration Management Database (CMDB) or install a collector agent to monitor. You will receive an assessment after data collection, which includes a projected cost estimate and savings of running your on-premises workloads in the AWS Cloud. This estimate will provide a summary of the projected costs to re-host on AWS based on usage patterns. It will show the breakdown of costs by infrastructure and software licenses. With this information, you can make the business case and plan next steps.

NEW QUESTION 23

- (Exam Topic 2)

A solutions architect needs to review the design of an Amazon EMR cluster that is using the EMR File System (EMRFS). The cluster performs tasks that are critical to business needs. The cluster is running Amazon EC2 On-Demand Instances at all times for all task, primary, and core nodes. The EMR tasks run each morning, starting at 1 :00 AM. and take 6 hours to finish running. The amount of time to complete the processing is not a priority because the data is not referenced until late in the day.

The solutions architect must review the architecture and suggest a solution to minimize the compute costs. Which solution should the solutions architect recommend to meet these requirements?

- A. Launch all task, primary, and core nodes on Spot Instances in an instance fleet
- B. Terminate the cluster, including all instances, when the processing is completed.
- C. Launch the primary and core nodes on On-Demand Instance
- D. Launch the task nodes on Spot Instances in an instance fleet
- E. Terminate the cluster, including all instances, when the processing is complete
- F. Purchase Compute Savings Plans to cover the On-Demand Instance usage.
- G. Continue to launch all nodes on On-Demand Instance
- H. Terminate the cluster, including all instances, when the processing is complete
- I. Purchase Compute Savings Plans to cover the On-Demand Instance usage
- J. Launch the primary and core nodes on On-Demand Instance
- K. Launch the task nodes on Spot Instances in an instance fleet
- L. Terminate only the task node instances when the processing is complete
- M. Purchase Compute Savings Plans to cover the On-Demand Instance usage.

Answer: A

Explanation:

Amazon EC2 Spot Instances offer spare compute capacity at steep discounts compared to On-Demand prices. Spot Instances can be interrupted by EC2 with two minutes of notification when EC2 needs the capacity back. Amazon EMR can handle Spot interruptions gracefully by decommissioning the nodes and redistributing the tasks to other nodes. By launching all nodes on Spot Instances in an instance fleet, the solutions architect can minimize the compute costs of the EMR cluster. An instance fleet is a collection of EC2 instances with different types and sizes that EMR automatically provisions to meet a defined target capacity. By terminating the cluster when the processing is completed, the solutions architect can avoid paying for idle resources. References:

- > <https://docs.aws.amazon.com/emr/latest/ManagementGuide/emr-managed-scaling.html>
- > <https://docs.aws.amazon.com/emr/latest/ManagementGuide/emr-instance-fleet.html>
- > <https://aws.amazon.com/blogs/big-data/optimizing-amazon-emr-for-resilience-and-cost-with-capacity-opt>

NEW QUESTION 27

- (Exam Topic 2)

A company manufactures smart vehicles. The company uses a custom application to collect vehicle data. The vehicles use the MQTT protocol to connect to the application.

The company processes the data in 5-minute intervals. The company then copies vehicle telematics data to on-premises storage. Custom applications analyze this

data to detect anomalies.

The number of vehicles that send data grows constantly. Newer vehicles generate high volumes of data. The on-premises storage solution is not able to scale for peak traffic, which results in data loss. The company must modernize the solution and migrate the solution to AWS to resolve the scaling challenges.

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

- A. Use AWS IoT Greengrass to send the vehicle data to Amazon Managed Streaming for Apache Kafka (Amazon MSK). Create an Apache Kafka application to store the data in Amazon S3. Use a pretrained model in Amazon SageMaker to detect anomalies.
- B. Use AWS IoT Core to receive the vehicle data
- C. Configure rules to route data to an Amazon Kinesis Data Firehose delivery stream that stores the data in Amazon S3. Create an Amazon Kinesis Data Analytics application that reads from the delivery stream to detect anomalies.
- D. Use AWS IoT FleetWise to collect the vehicle data
- E. Send the data to an Amazon Kinesis data stream. Use an Amazon Kinesis Data Firehose delivery stream to store the data in Amazon S3. Use the built-in machine learning transforms in AWS Glue to detect anomalies.
- F. Use Amazon MQ for RabbitMQ to collect the vehicle data
- G. Send the data to an Amazon Kinesis Data Firehose delivery stream to store the data in Amazon S3. Use Amazon Lookout for Metrics to detect anomalies.

Answer: B

Explanation:

Using AWS IoT Core to receive the vehicle data will enable connecting the smart vehicles to the cloud using the MQTT protocol¹. AWS IoT Core is a platform that enables you to connect devices to AWS Services and other devices, secure data and interactions, process and act upon device data, and enable applications to interact with devices even when they are offline². Configuring rules to route data to an Amazon Kinesis Data Firehose delivery stream that stores the data in Amazon S3 will enable processing and storing the vehicle data in a scalable and reliable way³. Amazon Kinesis Data Firehose is a fully managed service that delivers real-time streaming data to destinations such as Amazon S3. Creating an Amazon Kinesis Data Analytics application that reads from the delivery stream to detect anomalies will enable analyzing the vehicle data using SQL queries or Apache Flink applications. Amazon Kinesis Data Analytics is a fully managed service that enables you to process and analyze streaming data using SQL or Java.

NEW QUESTION 32

- (Exam Topic 2)

A company wants to optimize AWS data-transfer costs and compute costs across developer accounts within the company's organization in AWS Organizations. Developers can configure VPCs and launch Amazon EC2 instances in a single AWS Region. The EC2 instances retrieve approximately 1 TB of data each day from Amazon S3.

The developer activity leads to excessive monthly data-transfer charges and NAT gateway processing charges between EC2 instances and S3 buckets, along with high compute costs. The company wants to proactively enforce approved architectural patterns for any EC2 instance and VPC infrastructure that developers deploy within the AWS accounts. The company does not want this enforcement to negatively affect the speed at which the developers can perform their tasks. Which solution will meet these requirements MOST cost-effectively?

- A. Create SCPs to prevent developers from launching unapproved EC2 instance types. Provide the developers with an AWS CloudFormation template to deploy an approved VPC configuration with S3 interface endpoints. Scope the developers' IAM permissions so that the developers can launch VPC resources only with CloudFormation.
- B. Create a daily forecasted budget with AWS Budgets to monitor EC2 compute costs and S3 data-transfer costs across the developer accounts. When the forecasted cost is 75% of the actual budget cost, send an alert to the developer teams. If the actual budget cost is 100%, create a budget action to terminate the developers' EC2 instances and VPC infrastructure.
- C. Create an AWS Service Catalog portfolio that users can use to create an approved VPC configuration with S3 gateway endpoints and approved EC2 instances. Share the portfolio with the developer accounts. Configure an AWS Service Catalog launch constraint to use an approved IAM role. Scope the developers' IAM permissions to allow access only to AWS Service Catalog.
- D. Create and deploy AWS Config rules to monitor the compliance of EC2 and VPC resources in the developer AWS accounts. If developers launch unapproved EC2 instances or if developers create VPCs without S3 gateway endpoints, perform a remediation action to terminate the unapproved resources.

Answer: C

Explanation:

This solution allows developers to quickly launch resources using pre-approved configurations and instance types, while also ensuring that the resources launched comply with the company's architectural patterns. This can help reduce data transfer and compute costs associated with the resources. Using AWS Service Catalog also allows the company to control access to the approved configurations and resources through the use of IAM roles, while also allowing developers to quickly provision resources without negatively affecting their ability to perform their tasks.

Reference:

AWS Service Catalog: <https://aws.amazon.com/service-catalog/> AWS Service Catalog Constraints:

<https://docs.aws.amazon.com/servicecatalog/latest/adminguide/constraints.html>

AWS Service Catalog Launch Constraints: <https://docs.aws.amazon.com/servicecatalog/latest/adminguide/launch-constraints.html>

NEW QUESTION 33

- (Exam Topic 2)

A company uses an AWS CodeCommit repository. The company must store a backup copy of the data that is in the repository in a second AWS Region. Which solution will meet these requirements?

- A. Configure AWS Elastic Disaster Recovery to replicate the CodeCommit repository data to the second Region.
- B. Use AWS Backup to back up the CodeCommit repository on an hourly schedule. Create a cross-Region copy in the second Region.
- C. Create an Amazon EventBridge rule to invoke AWS CodeBuild when the company pushes code to the repository. Use CodeBuild to clone the repository. Create a zip file of the content. Copy the file to an S3 bucket in the second Region.
- D. Create an AWS Step Functions workflow on an hourly schedule to take a snapshot of the CodeCommit repository. Configure the workflow to copy the snapshot to an S3 bucket in the second Region.

Answer: B

Explanation:

AWS Backup is a fully managed service that makes it easy to centralize and automate the creation, retention, and restoration of backups across AWS services. It provides a way to schedule automatic backups for CodeCommit repositories on an hourly basis. Additionally, it also supports cross-Region replication, which allows you to copy the backups to a second Region for disaster recovery.

By using AWS Backup, the company can set up an automatic and regular backup schedule for the CodeCommit repository, ensuring that the data is regularly backed up and stored in a second Region. This can provide a way to recover quickly from any disaster event that might occur.

Reference:

AWS Backup documentation: <https://aws.amazon.com/backup/> AWS Backup for AWS CodeCommit documentation: https://aws.amazon.com/about-aws/whats-new/2020/07/aws-backup-now-supports-aws-codecommit-repository

NEW QUESTION 38

- (Exam Topic 2)

A company has several AWS accounts. A development team is building an automation framework for cloud governance and remediation processes. The automation framework uses AWS Lambda functions in a centralized account. A solutions architect must implement a least privilege permissions policy that allows the Lambda functions to run in each of the company's AWS accounts.

Which combination of steps will meet these requirements? (Choose two.)

- A. In the centralized account, create an IAM role that has the Lambda service as a trusted entity
- B. Add an inline policy to assume the roles of the other AWS accounts.
- C. In the other AWS accounts, create an IAM role that has minimal permission
- D. Add the centralized account's Lambda IAM role as a trusted entity.
- E. In the centralized account, create an IAM role that has roles of the other accounts as trusted entities. Provide minimal permissions.
- F. In the other AWS accounts, create an IAM role that has permissions to assume the role of the centralized account
- G. Add the Lambda service as a trusted entity.
- H. In the other AWS accounts, create an IAM role that has minimal permission
- I. Add the Lambda service as a trusted entity.

Answer: AB

Explanation:

<https://medium.com/@it.melnichenko/invoke-a-lambda-across-multiple-aws-accounts-8c094b2e70be>

NEW QUESTION 41

- (Exam Topic 2)

A company's public API runs as tasks on Amazon Elastic Container Service (Amazon ECS). The tasks run on AWS Fargate behind an Application Load Balancer (ALB) and are configured with Service Auto Scaling for the tasks based on CPU utilization. This service has been running well for several months.

Recently, API performance slowed down and made the application unusable. The company discovered that a significant number of SQL injection attacks had occurred against the API and that the API service had scaled to its maximum amount.

A solutions architect needs to implement a solution that prevents SQL injection attacks from reaching the ECS API service. The solution must allow legitimate traffic through and must maximize operational efficiency. Which solution meets these requirements?

- A. Create a new AWS WAF web ACL to monitor the HTTP requests and HTTPS requests that are forwarded to the ALB in front of the ECS tasks.
- B. Create a new AWS WAF Bot Control implementation
- C. Add a rule in the AWS WAF Bot Control managed rule group to monitor traffic and allow only legitimate traffic to the ALB in front of the ECS tasks.
- D. Create a new AWS WAF web ACL
- E. Add a new rule that blocks requests that match the SQL database rule group
- F. Set the web ACL to allow all other traffic that does not match those rules
- G. Attach the web ACL to the ALB in front of the ECS tasks.
- H. Create a new AWS WAF web ACL
- I. Create a new empty IP set in AWS WAF
- J. Add a new rule to the web ACL to block requests that originate from IP addresses in the new IP set
- K. Create an AWS Lambda function that scrapes the API logs for IP addresses that send SQL injection attacks, and add those IP addresses to the IP set
- L. Attach the web ACL to the ALB in front of the ECS tasks.

Answer: C

Explanation:

The company should create a new AWS WAF web ACL. The company should add a new rule that blocks requests that match the SQL database rule group. The company should set the web ACL to allow all other traffic that does not match those rules. The company should attach the web ACL to the ALB in front of the ECS tasks. This solution will meet the requirements because AWS WAF is a web application firewall that lets you monitor and control web requests that are forwarded to your web applications. You can use AWS WAF to define customizable web security rules that control which traffic can access your web applications and which traffic should be blocked¹. By creating a new AWS WAF web ACL, the company can create a collection of rules that define the conditions for allowing or blocking web requests. By adding a new rule that blocks requests that match the SQL database rule group, the company can prevent SQL injection attacks from reaching the ECS API service. The SQL database rule group is a managed rule group provided by AWS that contains rules to protect against common SQL injection attack patterns². By setting the web ACL to allow all other traffic that does not match those rules, the company can ensure that legitimate traffic can access the API service. By attaching the web ACL to the ALB in front of the ECS tasks, the company can apply the web security rules to all requests that are forwarded by the load balancer.

The other options are not correct because:

- Creating a new AWS WAF Bot Control implementation would not prevent SQL injection attacks from reaching the ECS API service. AWS WAF Bot Control is a feature that gives you visibility and control over common and pervasive bot traffic that can consume excess resources, skew metrics, cause downtime, or perform other undesired activities. However, it does not protect against SQL injection attacks, which are malicious attempts to execute unauthorized SQL statements against your database³.
- Creating a new AWS WAF web ACL to monitor the HTTP requests and HTTPS requests that are forwarded to the ALB in front of the ECS tasks would not prevent SQL injection attacks from reaching the ECS API service. Monitoring mode is a feature that enables you to evaluate how your rules would perform without actually blocking any requests. However, this mode does not provide any protection against attacks, as it only logs and counts requests that match your rules⁴.
- Creating a new AWS WAF web ACL and creating a new empty IP set in AWS WAF would not prevent SQL injection attacks from reaching the ECS API service. An IP set is a feature that enables you to specify a list of IP addresses or CIDR blocks that you want to allow or block based on their source IP address. However, this approach would not be effective or efficient against SQL injection attacks, as it would require constantly updating the IP set with new IP addresses of attackers, and it would not block attackers who use proxies or VPNs.

References:

- <https://aws.amazon.com/waf/>
- <https://docs.aws.amazon.com/waf/latest/developerguide/waf-bot-control.html>
- <https://docs.aws.amazon.com/waf/latest/developerguide/web-acl-monitoring-mode.html>
-

<https://docs.aws.amazon.com/waf/latest/developerguide/waf-ip-sets.html>

NEW QUESTION 44

- (Exam Topic 2)

A company is running a containerized application in the AWS Cloud. The application is running by using Amazon Elastic Container Service (Amazon ECS) on a set of Amazon EC2 instances. The EC2 instances run in an Auto Scaling group.

The company uses Amazon Elastic Container Registry (Amazon ECR) to store its container images. When a new image version is uploaded, the new image version receives a unique tag.

The company needs a solution that inspects new image versions for common vulnerabilities and exposures. The solution must automatically delete new image tags that have Critical or High severity findings. The solution also must notify the development team when such a deletion occurs.

Which solution meets these requirements?

- A. Configure scan on push on the repository Use Amazon EventBridge to invoke an AWS Step Functions state machine when a scan is complete for images that have Critical or High severity finding
- B. Use the Step Functions state machine to delete the image tag for those images and to notify the development team through Amazon Simple Notification Service (Amazon SNS).
- C. Configure scan on push on the repository Configure scan results to be pushed to an Amazon Simple Queue Service (Amazon SQS) queue
- D. Invoke an AWS Lambda function when a new message is added to the SQS queue
- E. Use the Lambda function to delete the image tag for images that have Critical or High severity finding
- F. Notify the development team by using Amazon Simple Email Service (Amazon SES).
- G. Schedule an AWS Lambda function to start a manual image scan every hour
- H. Configure Amazon EventBridge to invoke another Lambda function when a scan is complete
- I. Use the second Lambda function to delete the image tag for images that have Critical or High severity finding
- J. Notify the development team by using Amazon Simple Notification Service (Amazon SNS).
- K. Configure periodic image scan on the repository
- L. Configure scan results to be added to an Amazon Simple Queue Service (Amazon SQS) queue
- M. Invoke an AWS Step Functions state machine when a new message is added to the SQS queue
- N. Use the Step Functions state machine to delete the image tag for images that have Critical or High severity finding
- O. Notify the development team by using Amazon Simple Email Service (Amazon SES).

Answer: A

Explanation:

<https://docs.aws.amazon.com/AmazonECR/latest/userguide/ecr-eventbridge.html> "Activating an AWS Step Functions state machine"

<https://docs.aws.amazon.com/step-functions/latest/dg/tutorial-creating-lambda-state-machine.html>

NEW QUESTION 48

- (Exam Topic 2)

A company has a critical application in which the data tier is deployed in a single AWS Region. The data tier uses an Amazon DynamoDB table and an Amazon Aurora MySQL DB cluster. The current Aurora MySQL engine version supports a global database. The application tier is already deployed in two Regions.

Company policy states that critical applications must have application tier components and data tier components deployed across two Regions. The RTO and RPO must be no more than a few minutes each. A solutions architect must recommend a solution to make the data tier compliant with company policy.

Which combination of steps will meet these requirements? (Choose two.)

- A. Add another Region to the Aurora MySQL DB cluster
- B. Add another Region to each table in the Aurora MySQL DB cluster
- C. Set up scheduled cross-Region backups for the DynamoDB table and the Aurora MySQL DB cluster
- D. Convert the existing DynamoDB table to a global table by adding another Region to its configuration
- E. Use Amazon Route 53 Application Recovery Controller to automate database backup and recovery to the secondary Region

Answer: AD

Explanation:

The company should use Amazon Aurora global database and Amazon DynamoDB global table to deploy the data tier components across two Regions. Amazon Aurora global database is a feature that allows a single

Aurora database to span multiple AWS Regions, enabling low-latency global reads and fast recovery from Region-wide outages¹. Amazon DynamoDB global table is a feature that allows a single DynamoDB table to span multiple AWS Regions, enabling low-latency global reads and writes and fast recovery from Region-wide outages².

References:

> <https://aws.amazon.com/rds/aurora/global-database/>

> https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/globaltables_HowItWorks.html

> <https://aws.amazon.com/route53/application-recovery-controller/>

NEW QUESTION 50

- (Exam Topic 2)

A company runs an application on AWS. The company curates data from several different sources. The company uses proprietary algorithms to perform data transformations and aggregations. After the company performs ETL processes, the company stores the results in Amazon Redshift tables. The company sells this data to other companies. The company downloads the data as files from the Amazon Redshift tables and transmits the files to several data customers by using FTP. The number of data customers has grown significantly. Management of the data customers has become difficult.

The company will use AWS Data Exchange to create a data product that the company can use to share data with customers. The company wants to confirm the identities of the customers before the company shares data.

The customers also need access to the most recent data when the company publishes the data. Which solution will meet these requirements with the LEAST operational overhead?

- A. Use AWS Data Exchange for APIs to share data with customer
- B. Configure subscription verification In the AWS account of the company that produces the data, create an Amazon API Gateway Data API service integration with Amazon Redshift
- C. Require the data customers to subscribe to the data product In the AWS account of the company that produces the data, create an AWS Data Exchange datashare by connecting AWS Data Exchange to the Redshift

- D. cluste
- E. Configure subscription verificatio
- F. Require the data customers to subscribe to the data product.
- G. Download the data from the Amazon Redshift tables to an Amazon S3 bucket periodicall
- H. Use AWS Data Exchange for S3 to share data with customers.
- I. Configure subscription verificatio
- J. Require the data customers to subscribe to the data product Publish the Amazon Redshift data to an Open Data on AWS Data Exchange
- K. Require the customers to subscribe to the data product in AWS Data Exchange
- L. In the AWS account of the company that produces the data, attach IAM resource-based policies to the Amazon Redshift tables to allow access only to verified AWS accounts.

Answer: C

Explanation:

The company should download the data from the Amazon Redshift tables to an Amazon S3 bucket periodically and use AWS Data Exchange for S3 to share data with customers. The company should configure subscription verification and require the data customers to subscribe to the data product. This solution will meet the requirements with the least operational overhead because AWS Data Exchange for S3 is a feature that enables data subscribers to access third-party data files directly from data providers' Amazon S3 buckets. Subscribers can easily use these files for their data analysis with AWS services without needing to create or manage data copies. Data providers can easily set up AWS Data Exchange for S3 on top of their existing S3 buckets to share direct access to an entire S3 bucket or specific prefixes and S3 objects. AWS Data Exchange automatically manages subscriptions, entitlements, billing, and payment¹.

The other options are not correct because:

- Using AWS Data Exchange for APIs to share data with customers would not work because AWS Data Exchange for APIs is a feature that enables data subscribers to access third-party APIs directly from data providers' AWS accounts. Subscribers can easily use these APIs for their data analysis with AWS services without needing to manage API keys or tokens. Data providers can easily set up AWS Data Exchange for APIs on top of their existing API Gateway resources to share direct access to an entire API or specific routes and stages². However, this feature is not suitable for sharing data from Amazon Redshift tables, which are not exposed as APIs.
- Creating an Amazon API Gateway Data API service integration with Amazon Redshift would not work because the Data API is a feature that enables you to query your Amazon Redshift cluster using HTTP requests, without needing a persistent connection or a SQL client³. It is useful for building applications that interact with Amazon Redshift, but not for sharing data files with customers.
- Creating an AWS Data Exchange datashare by connecting AWS Data Exchange to the Redshift cluster would not work because AWS Data Exchange does not support datashares for Amazon Redshift clusters. A datashare is a feature that enables you to share live and secure access to your Amazon Redshift data across your accounts or with third parties without copying or moving the underlying data⁴. It is useful for sharing query results and views with other users, but not for sharing data files with customers.
- Publishing the Amazon Redshift data to an Open Data on AWS Data Exchange would not work because Open Data on AWS Data Exchange is a feature that enables you to find and use free and public datasets from AWS customers and partners. It is useful for accessing open and free data, but not for confirming the identities of the customers or charging them for the data.

References:

- <https://aws.amazon.com/data-exchange/why-aws-data-exchange/s3/>
- <https://aws.amazon.com/data-exchange/why-aws-data-exchange/api/>
- <https://docs.aws.amazon.com/redshift/latest/mgmt/data-api.html>
- <https://docs.aws.amazon.com/redshift/latest/dg/datashare-overview.html>
- <https://aws.amazon.com/data-exchange/open-data/>

NEW QUESTION 51

- (Exam Topic 2)

A company is designing an AWS Organizations structure. The company wants to standardize a process to apply tags across the entire organization. The company will require tags with specific values when a user creates a new resource. Each of the company's OUs will have unique tag values.

Which solution will meet these requirements?

- A. Use an SCP to deny the creation of resources that do not have the required tag
- B. Create a tag policy that Includes the tag values that the company has assigned to each O
- C. Attach the tag policies to the OUs.
- D. Use an SCP to deny the creation of resources that do not have the required tag
- E. Create a tag policy that includes the tag values that the company has assigned to each O
- F. Attach the tag policies to the organization's management account.
- G. Use an SCP to allow the creation of resources only when the resources have the required tag
- H. Create a tag policy that includes the tag values that the company has assigned to each O
- I. Attach the tag policies to the OUs.
- J. Use an SCP to deny the creation of resources that do not have the required tag
- K. Define the list of tags.Attach the SCP to the OUs

Answer: A

Explanation:

<https://aws.amazon.com/blogs/mt/implement-aws-resource-tagging-strategy-using-aws-tag-policies-and-service>

NEW QUESTION 54

- (Exam Topic 2)

A company is deploying a new web-based application and needs a storage solution for the Linux application servers. The company wants to create a single location for updates to application data for all instances. The active dataset will be up to 100 GB in size. A solutions architect has determined that peak operations will occur for 3 hours daily and will require a total of 225 MiBps of read throughput.

The solutions architect must design a Multi-AZ solution that makes a copy of the data available in another AWS Region for disaster recovery (DR). The DR copy has an RPO of less than 1 hour.

Which solution will meet these requirements?

- A. Deploy a new Amazon Elastic File System (Amazon EFS) Multi-AZ file syste
- B. Configure the file system for 75 MiBps of provisioned throughpu
- C. Implement replication to a file system in the DR Region.

- D. Deploy a new Amazon FSx for Lustre file system
- E. Configure Bursting Throughput mode for the file system
- F. Use AWS Backup to back up the file system to the DR Region.
- G. Deploy a General Purpose SSD (gp3) Amazon Elastic Block Store (Amazon EBS) volume with 225 MiBps of throughput
- H. Enable Multi-Attach for the EBS volume
- I. Use AWS Elastic Disaster Recovery to replicate the EBS volume to the DR Region.
- J. Deploy an Amazon FSx for OpenZFS file system in both the production Region and the DR Region. Create an AWS DataSync scheduled task to replicate the data from the production file system to the DR file system every 10 minutes.

Answer: A

Explanation:

The company should deploy a new Amazon Elastic File System (Amazon EFS) Multi-AZ file system. The company should configure the file system for 75 MiBps of provisioned throughput. The company should implement replication to a file system in the DR Region. This solution will meet the requirements because Amazon EFS is a serverless, fully elastic file storage service that lets you share file data without provisioning or managing storage capacity and performance. Amazon EFS is built to scale on demand to petabytes without disrupting applications, growing and shrinking automatically as you add and remove files¹. By deploying a new Amazon EFS Multi-AZ file system, the company can create a single location for updates to application data for all instances. A Multi-AZ file system replicates data across multiple Availability Zones (AZs) within a Region, providing high availability and durability². By configuring the file system for 75 MiBps of provisioned throughput, the company can ensure that it meets the peak operations requirement of 225 MiBps of read throughput. Provisioned throughput is a feature that enables you to specify a level of throughput that the file system can drive independent of the file system's size or burst credit balance³. By implementing replication to a file system in the DR Region, the company can make a copy of the data available in another AWS Region for disaster recovery. Replication is a feature that enables you to replicate data from one EFS file system to another EFS file system across AWS Regions. The replication process has an RPO of less than 1 hour.

The other options are not correct because:

- Deploying a new Amazon FSx for Lustre file system would not provide a single location for updates to application data for all instances. Amazon FSx for Lustre is a fully managed service that provides cost-effective, high-performance storage for compute workloads. However, it does not support concurrent write access from multiple instances. Using AWS Backup to back up the file system to the DR Region would not provide real-time replication of data. AWS Backup is a service that enables you to centralize and automate data protection across AWS services. However, it does not support continuous data replication or cross-Region disaster recovery.
- Deploying a General Purpose SSD (gp3) Amazon Elastic Block Store (Amazon EBS) volume with 225 MiBps of throughput would not provide a single location for updates to application data for all instances. Amazon EBS is a service that provides persistent block storage volumes for use with Amazon EC2 instances. However, it does not support concurrent access from multiple instances, unless Multi-Attach is enabled. Enabling Multi-Attach for the EBS volume would not provide Multi-AZ resilience or cross-Region replication. Multi-Attach is a feature that enables you to attach an EBS volume to multiple EC2 instances within the same Availability Zone. Using AWS Elastic Disaster Recovery to replicate the EBS volume to the DR Region would not provide real-time replication of data. AWS Elastic Disaster Recovery (AWS DRS) is a service that enables you to orchestrate and automate disaster recovery workflows across AWS Regions. However, it does not support continuous data replication or sub-hour RPOs.
- Deploying an Amazon FSx for OpenZFS file system in both the production Region and the DR Region would not be as simple or cost-effective as using Amazon EFS. Amazon FSx for OpenZFS is a fully managed service that provides high-performance storage with strong data consistency and advanced data management features for Linux workloads. However, it requires more configuration and management than Amazon EFS, which is serverless and fully elastic. Creating an AWS DataSync scheduled task to replicate the data from the production file system to the DR file system every 10 minutes would not provide real-time replication of data. AWS DataSync is a service that enables you to transfer data between on-premises storage and AWS services, or between AWS services. However, it does not support continuous data replication or sub-minute RPOs.

References:

- <https://aws.amazon.com/efs/>
- <https://docs.aws.amazon.com/efs/latest/ug/how-it-works.html#how-it-works-azs>
- <https://docs.aws.amazon.com/efs/latest/ug/performance.html#provisioned-throughput>
- <https://docs.aws.amazon.com/efs/latest/ug/replication.html>
- <https://aws.amazon.com/fsx/lustre/>
- <https://aws.amazon.com/backup/>
- <https://aws.amazon.com/ebs/>
- <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-volumes-multi.html>

NEW QUESTION 55

- (Exam Topic 2)

A company is building a hybrid environment that includes servers in an on-premises data center and in the AWS Cloud. The company has deployed Amazon EC2 instances in three VPCs. Each VPC is in a different AWS Region. The company has established an AWS Direct Connect connection to the data center from the Region that is closest to the data center.

The company needs the servers in the on-premises data center to have access to the EC2 instances in all three VPCs. The servers in the on-premises data center also must have access to AWS public services.

Which combination of steps will meet these requirements with the LEAST cost? (Select TWO.)

- A. Create a Direct Connect gateway in the Region that is closest to the data center
- B. Attach the Direct Connect connection to the Direct Connect gateway
- C. Use the Direct Connect gateway to connect the VPCs in the other two Regions.
- D. Direct Connect gateway to connect the VPCs in the other two Regions.
- E. Set up additional Direct Connect connections from the on-premises data center to the other two Regions.
- F. Create a private VIF
- G. Establish an AWS Site-to-Site VPN connection over the private VIF to the VPCs in the other two Regions.
- H. Create a public VIF
- I. Establish an AWS Site-to-Site VPN connection over the public VIF to the VPCs in the other two Regions.
- J. Use VPC peering to establish a connection between the VPCs across the Region
- K. Create a private VIF with the existing Direct Connect connection to connect to the peered VPCs.

Answer: AE

Explanation:

A Direct Connect gateway allows you to connect multiple VPCs across different Regions to a Direct Connect connection¹. A public VIF allows you to access AWS public services such as EC2¹. A Site-to-Site VPN connection over the public VIF provides encryption and redundancy for the traffic between the on-premises data

center and the VPCs2. This solution is cheaper than setting up additional Direct Connect connections or using a private VIF with VPC peering.

NEW QUESTION 57

- (Exam Topic 2)

A company has many separate AWS accounts and uses no central billing or management. Each AWS account hosts services for different departments in the company. The company has a Microsoft Azure Active Directory that is deployed.

A solution architect needs to centralize billing and management of the company's AWS accounts. The company wants to start using identify federation instead of manual user management. The company also wants to use temporary credentials instead of long-lived access keys.

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

- A. Create a new AWS account to serve as a management account
- B. Deploy an organization in AWS Organization
- C. Invite each existing AWS account to join the organization
- D. Ensure that each account accepts the invitation.
- E. Configure each AWS Account's email address to be aws+<account id>@example.com so that account management email messages and invoices are sent to the same place.
- F. Deploy AWS IAM Identity Center (AWS Single Sign-On) in the management account
- G. Connect IAM Identity Center to the Azure Active Director
- H. Configure IAM Identity Center for automatic synchronization of users and groups.
- I. Deploy an AWS Managed Microsoft AD directory in the management account
- J. Share the directory with all other accounts in the organization by using AWS Resource Access Manager (AWS RAM).
- K. Create AWS IAM Identity Center (AWS Single Sign-On) permission set
- L. Attach the permission sets to the appropriate IAM Identity Center groups and AWS accounts.
- M. Configure AWS Identity and Access Management (IAM) in each AWS account to use AWS Managed Microsoft AD for authentication and authorization.

Answer: ACE

NEW QUESTION 61

- (Exam Topic 2)

A company has millions of objects in an Amazon S3 bucket. The objects are in the S3 Standard storage class. All the S3 objects are accessed frequently. The number of users and applications that access the objects is increasing rapidly. The objects are encrypted with server-side encryption with AWS KMS Keys (SSE-KMS).

A solutions architect reviews the company's monthly AWS invoice and notices that AWS KMS costs are increasing because of the high number of requests from Amazon S3. The solutions architect needs to optimize costs with minimal changes to the application.

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

- A. Create a new S3 bucket that has server-side encryption with customer-provided keys (SSE-C) as the encryption type
- B. Copy the existing objects to the new S3 bucket
- C. Specify SSE-C.
- D. Create a new S3 bucket that has server-side encryption with Amazon S3 managed keys (SSE-S3) as the encryption type
- E. Use S3 Batch Operations to copy the existing objects to the new S3 bucket
- F. Specify SSE-S3.
- G. Use AWS CloudHSM to store the encryption key
- H. Create a new S3 bucket
- I. Use S3 Batch Operations to copy the existing objects to the new S3 bucket
- J. Encrypt the objects by using the keys from CloudHSM.
- K. Use the S3 Intelligent-Tiering storage class for the S3 bucket
- L. Create an S3 Intelligent-Tiering archive configuration to transition objects that are not accessed for 90 days to S3 Glacier Deep Archive.

Answer: B

Explanation:

To reduce the volume of Amazon S3 calls to AWS KMS, use Amazon S3 bucket keys, which are protected encryption keys that are reused for a limited time in Amazon S3. Bucket keys can reduce costs for AWS KMS requests by up to 99%. You can configure a bucket key for all objects in an Amazon S3 bucket, or for a specific object in an Amazon S3 bucket. https://docs.aws.amazon.com/fr_fr/kms/latest/developerguide/services-s3.html

NEW QUESTION 63

- (Exam Topic 2)

A company runs an application on a fleet of Amazon EC2 instances that are in private subnets behind an

internet-facing Application Load Balancer (ALB). The ALB is the origin for an Amazon CloudFront distribution. An AWS WAF web ACL that contains various AWS managed rules is associated with the CloudFront distribution.

The company needs a solution that will prevent internet traffic from directly accessing the ALB. Which solution will meet these requirements with the LEAST operational overhead?

- A. Create a new web ACL that contains the same rules that the existing web ACL contain
- B. Associate the new web ACL with the ALB.
- C. Associate the existing web ACL with the ALB.
- D. Add a security group rule to the ALB to allow traffic from the AWS managed prefix list for CloudFront only.
- E. Add a security group rule to the ALB to allow only the various CloudFront IP address ranges.

Answer: C

Explanation:

<https://aws.amazon.com/about-aws/whats-new/2022/02/amazon-cloudfront-managed-prefix-list/>

NEW QUESTION 65

- (Exam Topic 2)

A company wants to run a custom network analysis software package to inspect traffic as traffic leaves and enters a VPC. The company has deployed the solution by using AWS Cloud Formation on three Amazon EC2 instances in an Auto Scaling group. All network routing has been established to direct traffic to the EC2

instances.

Whenever the analysis software stops working, the Auto Scaling group replaces an instance. The network routes are not updated when the instance replacement occurs.

Which combination of steps will resolve this issue? (Select THREE.)

- A. Create alarms based on EC2 status check metrics that will cause the Auto Scaling group to replace the failed instance.
- B. Update the Cloud Formation template to install the Amazon CloudWatch agent on the EC2 instances. Configure the CloudWatch agent to send process metrics for the application.
- C. Update the Cloud Formation template to install AWS Systems Manager Agent on the EC2 instances. Configure Systems Manager Agent to send process metrics for the application.
- D. Create an alarm for the custom metric in Amazon CloudWatch for the failure scenario
- E. Configure the alarm to publish a message to an Amazon Simple Notification Service (Amazon SNS) topic.
- F. Create an AWS Lambda function that responds to the Amazon Simple Notification Service (Amazon SNS) message to take the instance out of service.
- G. Update the network routes to point to the replacement instance.
- H. In the Cloud Formation template, write a condition that updates the network routes when a replacement instance is launched.

Answer: BDE

NEW QUESTION 68

- (Exam Topic 2)

A solutions architect wants to cost-optimize and appropriately size Amazon EC2 instances in a single AWS account. The solutions architect wants to ensure that the instances are optimized based on CPU, memory, and network metrics.

Which combination of steps should the solutions architect take to meet these requirements? (Choose two.)

- A. Purchase AWS Business Support or AWS Enterprise Support for the account.
- B. Turn on AWS Trusted Advisor and review any "Low Utilization Amazon EC2 Instances" recommendations.
- C. Install the Amazon CloudWatch agent and configure memory metric collection on the EC2 instances.
- D. Configure AWS Compute Optimizer in the AWS account to receive findings and optimization recommendations.
- E. Create an EC2 Instance Savings Plan for the AWS Regions, instance families, and operating systems of interest.

Answer: BD

Explanation:

AWS Trusted Advisor is a service that provides real-time guidance to help users provision their resources following AWS best practices¹. One of the Trusted Advisor checks is "Low Utilization Amazon EC2 Instances", which identifies EC2 instances that appear to be underutilized based on CPU, network I/O, and disk I/O metrics¹. This check can help users optimize the cost and size of their EC2 instances by recommending smaller or more appropriate instance types.

AWS Compute Optimizer is a service that analyzes the configuration and utilization metrics of AWS resources and generates optimization recommendations to reduce the cost and improve the performance of workloads². Compute Optimizer supports four types of AWS resources: EC2 instances, EBS volumes, ECS services on AWS Fargate, and Lambda functions². For EC2 instances, Compute Optimizer evaluates the vCPUs, memory, storage, and other specifications, as well as the CPU utilization, network in and out, disk read and write, and other utilization metrics of currently running instances³. It then recommends optimal instance types based on price-performance trade-offs.

Option A is incorrect because purchasing AWS Business Support or AWS Enterprise Support for the account will not directly help with cost-optimization and sizing of EC2 instances. However, these support plans do provide access to more Trusted Advisor checks than the basic support plan¹.

Option C is incorrect because installing the Amazon CloudWatch agent and configuring memory metric collection on the EC2 instances will not provide any optimization recommendations by itself. However, memory metrics can be used by Compute Optimizer to enhance its recommendations if enabled³.

Option E is incorrect because creating an EC2 Instance Savings Plan for the AWS Regions, instance families, and operating systems of interest will not help with cost-optimization and sizing of EC2 instances. Savings Plans are a flexible pricing model that offer lower prices on Amazon EC2 usage in exchange for a commitment to a consistent amount of usage for a 1- or 3-year term⁴. Savings Plans do not affect the configuration or utilization of EC2 instances.

NEW QUESTION 69

- (Exam Topic 2)

A company has built a high performance computing (HPC) cluster in AWS for a tightly coupled workload that generates a large number of shared files stored in Amazon EFS. The cluster was performing well when the number of Amazon EC2 instances in the cluster was 100. However, when the company increased the cluster size to 1,000 EC2 instances, overall performance was well below expectations.

Which collection of design choices should a solutions architect make to achieve the maximum performance from the HPC cluster? (Select THREE.)

- A. Ensure the HPC cluster is launched within a single Availability Zone.
- B. Launch the EC2 instances and attach elastic network interfaces in multiples of four.
- C. Select EC2 Instance types with an Elastic Fabric Adapter (EFA) enabled.
- D. Ensure the cluster is launched across multiple Availability Zones.
- E. Replace Amazon EFS with multiple Amazon EBS volumes in a RAID array.
- F. Replace Amazon EFS with Amazon FSx for Lustre.

Answer: ACF

Explanation:

* A. High performance computing (HPC) workload cluster should be in a single AZ.

* C. Elastic Fabric Adapter (EFA) is a network device that you can attach to your Amazon EC2 instances to accelerate High Performance Computing (HPC)

* F. Amazon FSx for Lustre - Use it for workloads where speed matters, such as machine learning, high performance computing (HPC), video processing, and financial modeling.

Cluster – packs instances close together inside an Availability Zone. This strategy enables workloads to achieve the low-latency network performance necessary for tightly-coupled node-to-node communication that is typical of HPC applications.

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/placement-groups.html>

NEW QUESTION 73

- (Exam Topic 2)

A telecommunications company is running an application on AWS. The company has set up an AWS Direct Connect connection between the company's on-premises data center and AWS. The company deployed the application on Amazon EC2 instances in multiple Availability Zones behind an internal Application Load Balancer (ALB). The company's clients connect from the on-premises network by using HTTPS. The TLS terminates in the ALB. The company has multiple target groups and uses path-based routing to forward requests based on the URL path.

The company is planning to deploy an on-premises firewall appliance with an allow list that is based on IP address. A solutions architect must develop a solution to allow traffic flow to AWS from the on-premises network so that the clients can continue to access the application.

Which solution will meet these requirements?

- A. Configure the existing ALB to use static IP addresses
- B. Assign IP addresses in multiple Availability Zones to the ALB
- C. Add the ALB IP addresses to the firewall appliance.
- D. Create a Network Load Balancer (NLB). Associate the NLB with one static IP addresses in multiple Availability Zone
- E. Create an ALB-type target group for the NLB and add the existing ALB. Add the NLB IP addresses to the firewall appliance
- F. Update the clients to connect to the NLB.
- G. Create a Network Load Balancer (NLB). Associate the NLB with one static IP addresses in multiple Availability Zone
- H. Add the existing target groups to the NLB
- I. Update the clients to connect to the NLB
- J. Delete the ALB. Add the NLB IP addresses to the firewall appliance.
- K. Create a Gateway Load Balancer (GWLB). Assign static IP addresses to the GWLB in multiple Availability Zone
- L. Create an ALB-type target group for the GWLB and add the existing ALB
- M. Add the GWLB IP addresses to the firewall appliance
- N. Update the clients to connect to the GWLB.

Answer: B

Explanation:

The company should create a Network Load Balancer (NLB) and associate it with one static IP address in multiple Availability Zones. The company should also create an ALB-type target group for the NLB and add the existing ALB. The company should add the NLB IP addresses to the firewall appliance and update the clients to connect to the NLB. This solution will allow traffic flow to AWS from the on-premises network by using static IP addresses that can be added to the firewall appliance's allow list. The NLB will forward requests to the ALB, which will use path-based routing to forward requests to the target groups.

NEW QUESTION 78

- (Exam Topic 2)

A company has a website that runs on Amazon EC2 instances behind an Application Load Balancer (ALB). The instances are in an Auto Scaling group. The ALB is associated with an AWS WAF web ACL.

The website often encounters attacks in the application layer. The attacks produce sudden and significant increases in traffic on the application server. The access logs show that each attack originates from different IP addresses. A solutions architect needs to implement a solution to mitigate these attacks.

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

- A. Create an Amazon CloudWatch alarm that monitors server access
- B. Set a threshold based on access by IP address
- C. Configure an alarm action that adds the IP address to the web ACL's deny list.
- D. Deploy AWS Shield Advanced in addition to AWS WAF
- E. Add the ALB as a protected resource.
- F. Create an Amazon CloudWatch alarm that monitors user IP addresses
- G. Set a threshold based on access by IP address
- H. Configure the alarm to invoke an AWS Lambda function to add a deny rule in the application server's subnet route table for any IP addresses that activate the alarm.
- I. Inspect access logs to find a pattern of IP addresses that launched the attack
- J. Use an Amazon Route 53 geolocation routing policy to deny traffic from the countries that host those IP addresses.

Answer: C

Explanation:

"The AWS WAF API supports security automation such as blacklisting IP addresses that exceed request limits, which can be useful for mitigating HTTP flood attacks." >

<https://aws.amazon.com/blogs/security/how-to-protect-dynamic-web-applications-against-ddos-attacks-by-using>

NEW QUESTION 79

- (Exam Topic 2)

A company has five development teams that have each created five AWS accounts to develop and host applications. To track spending, the development teams log in to each account every month, record the current cost from the AWS Billing and Cost Management console, and provide the information to the company's finance team.

The company has strict compliance requirements and needs to ensure that resources are created only in AWS Regions in the United States. However, some resources have been created in other Regions.

A solutions architect needs to implement a solution that gives the finance team the ability to track and consolidate expenditures for all the accounts. The solution also must ensure that the company can create resources only in Regions in the United States.

Which combination of steps will meet these requirements in the MOST operationally efficient way? (Select THREE.)

- A. Create a new account to serve as a management account
- B. Create an Amazon S3 bucket for the finance team. Use AWS Cost and Usage Reports to create monthly reports and to store the data in the finance team's S3 bucket.
- C. Create a new account to serve as a management account
- D. Deploy an organization in AWS Organizations with all features enabled
- E. Invite all the existing accounts to the organization
- F. Ensure that each account accepts the invitation.
- G. Create an OU that includes all the development teams
- H. Create an SCP that allows the creation of resources only in Regions that are in the United States
- I. Apply the SCP to the OU.
- J. Create an OU that includes all the development teams
- K. Create an SCP that denies the creation of resources in Regions that are outside the United States
- L. Apply the SCP to the OU.
- M. Create an IAM role in the management account. Attach a policy that includes permissions to view the Billing and Cost Management console
- N. Allow the finance team users to assume the role
- O. Use AWS Cost Explorer and the Billing and Cost Management console to analyze cost.

- P. Create an IAM role in each AWS account
- Q. Attach a policy that includes permissions to view the Billing and Cost Management console
- R. Allow the finance team users to assume the role.

Answer: BCE

Explanation:

AWS Organizations is a service that enables you to consolidate multiple AWS accounts into an organization that you create and centrally manage. By creating a management account and inviting all the existing accounts to join the organization, the solutions architect can track and consolidate expenditures for all the accounts using AWS Cost Management tools such as AWS Cost Explorer and AWS Budgets. An organizational unit (OU) is a group of accounts within an organization that can be used to apply policies and simplify management. A service control policy (SCP) is a type of policy that you can use to manage permissions in your organization. By creating an OU that includes all the development teams and applying an SCP that allows the creation of resources only in Regions that are in the United States, the solutions architect can ensure that the company meets its compliance requirements and avoids unwanted charges from other Regions. An IAM role is an identity with permission policies that determine what the identity can and cannot do in AWS. By creating an IAM role in the management account and allowing the finance team users to assume it, the solutions architect can give them access to view the Billing and Cost Management console without sharing credentials or creating additional users. References:

- > https://docs.aws.amazon.com/organizations/latest/userguide/orgs_introduction.html
- > https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scp.html
- > https://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles.html
- > <https://docs.aws.amazon.com/aws-cost-management/latest/userguide/what-is-costmanagement.html>

NEW QUESTION 80

- (Exam Topic 2)

A company processes environment data. The has a set up sensors to provide a continuous stream of data from different areas in a city. The data is available in JSON format.

The company wants to use an AWS solution to send the data to a database that does not require fixed schemas for storage. The data must be send in real time. Which solution will meet these requirements?

- A. Use Amazon Kinesis Data Firehose to send the data to Amazon Redshift.
- B. Use Amazon Kinesis Data streams to send the data to Amazon DynamoDB.
- C. Use Amazon Managed Streaming for Apache Kafka (Amazon MSK) to send the data to Amazon Aurora.
- D. Use Amazon Kinesis Data firehose to send the data to Amazon Keyspaces (for Apache Cassandra).

Answer: B

Explanation:

Amazon Kinesis Data Streams is a service that enables real-time data ingestion and processing. Amazon DynamoDB is a NoSQL database that does not require fixed schemas for storage. By using Kinesis Data Streams and DynamoDB, the company can send the JSON data to a database that can handle schemaless data in real time. References:

- > <https://docs.aws.amazon.com/streams/latest/dev/introduction.html>
- > <https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Introduction.html>

NEW QUESTION 83

- (Exam Topic 2)

A company's solutions architect is analyzing costs of a multi-application environment. The environment is deployed across multiple Availability Zones in a single AWS Region. After a recent acquisition, the company manages two organizations in AWS Organizations. The company has created multiple service provider applications as AWS PrivateLink-powered VPC endpoint services in one organization. The company has created multiple service consumer applications in the other organization.

Data transfer charges are much higher than the company expected, and the solutions architect needs to reduce the costs. The solutions architect must recommend guidelines for developers to follow when they deploy services. These guidelines must minimize data transfer charges for the whole environment. Which guidelines meet these requirements? (Select TWO.)

- A. Use AWS Resource Access Manager to share the subnets that host the service provider applications with other accounts in the organization.
- B. Place the service provider applications and the service consumer applications in AWS accounts in the same organization.
- C. Turn off cross-zone load balancing for the Network Load Balancer in all service provider application deployments.
- D. Ensure that service consumer compute resources use the Availability Zone-specific endpoint service by using the endpoint's local DNS name.
- E. Create a Savings Plan that provides adequate coverage for the organization's planned inter-Availability Zone data transfer usage.

Answer: CD

Explanation:

Cross-zone load balancing enables traffic to be distributed evenly across all registered instances in all enabled Availability Zones. However, this also increases data transfer charges between Availability Zones. By turning off cross-zone load balancing, the service provider applications can reduce inter-Availability Zone data transfer costs. Similarly, by using the Availability Zone-specific endpoint service, the service consumer applications can ensure that they connect to the nearest service provider application in the same Availability Zone, avoiding cross-Availability Zone data transfer charges. References:

- > <https://docs.aws.amazon.com/vpc/latest/userguide/vpce-interface.html#vpce-interface-dns>

NEW QUESTION 85

- (Exam Topic 1)

A company wants to use a third-party software-as-a-service (SaaS) application. The third-party SaaS application is consumed through several API calls. The third-party SaaS application also runs on AWS inside a VPC.

The company will consume the third-party SaaS application from inside a VPC. The company has internal security policies that mandate the use of private connectivity that does not traverse the internet. No resources that run in the company VPC are allowed to be accessed from outside the company's VPC. All permissions must conform to the principles of least privilege.

Which solution meets these requirements?

- A. Create an AWS PrivateLink interface VPC endpoint

- B. Connect this endpoint to the endpoint service that the third-party SaaS application provide
- C. Create a security group to limit the access to the endpoint
- D. Associate the security group with the endpoint.
- E. Create an AWS Site-to-Site VPN connection between the third-party SaaS application and the company VP
- F. Configure network ACLs to limit access across the VPN tunnels.
- G. Create a VPC peering connection between the third-party SaaS application and the company VPC update route tables by adding the needed routes for the peering connection.
- H. Create an AWS PrivateLink endpoint service
- I. Ask the third-party SaaS provider to create an interface VPC endpoint for this endpoint service
- J. Grant permissions for the endpoint service to the specific account of the third-party SaaS provider.

Answer: A

Explanation:

Reference architecture - <https://docs.aws.amazon.com/vpc/latest/privatelink/privatelink-access-saas.html> Note from documentation that Interface Endpoint is at client side

NEW QUESTION 87

- (Exam Topic 1)

A company is running an event ticketing platform on AWS and wants to optimize the platform's cost-effectiveness. The platform is deployed on Amazon Elastic Kubernetes Service (Amazon EKS) with Amazon EC2 and is backed by an Amazon RDS for MySQL DB instance. The company is developing new application features to run on Amazon EKS with AWS Fargate. The platform experiences infrequent high peaks in demand. The surges in demand depend on event dates. Which solution will provide the MOST cost-effective setup for the platform?

- A. Purchase Standard Reserved Instances for the EC2 instances that the EKS cluster uses in its baseline load
- B. Scale the cluster with Spot Instances to handle peak
- C. Purchase 1-year All Upfront Reserved Instances for the database to meet predicted peak load for the year.
- D. Purchase Compute Savings Plans for the predicted medium load of the EKS cluster
- E. Scale the cluster with On-Demand Capacity Reservations based on event dates for peak
- F. Purchase 1-year No Upfront Reserved Instances for the database to meet the predicted base load
- G. Temporarily scale out database read replicas during peaks.
- H. Purchase EC2 Instance Savings Plans for the predicted base load of the EKS cluster
- I. Scale the cluster with Spot Instances to handle peak
- J. Purchase 1-year All Upfront Reserved Instances for the database to meet the predicted base load
- K. Temporarily scale up the DB instance manually during peaks.
- L. Purchase Compute Savings Plans for the predicted base load of the EKS cluster
- M. Scale the cluster with Spot Instances to handle peak
- N. Purchase 1-year All Upfront Reserved Instances for the database to meet the predicted base load
- O. Temporarily scale up the DB instance manually during peaks.

Answer: B

Explanation:

They all mention using spot instances and EKS based on EC2. A spot instance is not appropriate for a production server and the company is developing new application designed for AWS Fargate, which means we must plan the future cost improvement including AWS Fargate.
<https://aws.amazon.com/savingsplans/compute-pricing/>

NEW QUESTION 92

- (Exam Topic 1)

A company is hosting an image-processing service on AWS in a VPC. The VPC extends across two Availability Zones. Each Availability Zone contains one public subnet and one private subnet.

The service runs on Amazon EC2 instances in the private subnets. An Application Load Balancer in the public subnets is in front of the service. The service needs to communicate with the internet and does so through two NAT gateways. The service uses Amazon S3 for image storage. The EC2 instances retrieve approximately 1 GB of data from an S3 bucket each day.

The company has promoted the service as highly secure. A solutions architect must reduce cloud expenditures as much as possible without compromising the service's security posture or increasing the time spent on ongoing operations.

Which solution will meet these requirements?

- A. Replace the NAT gateways with NAT instance
- B. In the VPC route table, create a route from the private subnets to the NAT instances.
- C. Move the EC2 instances to the public subnet
- D. Remove the NAT gateways.
- E. Set up an S3 gateway VPC endpoint in the VPC
- F. Attach an endpoint policy to the endpoint to allow the required actions on the S3 bucket.
- G. Attach an Amazon Elastic File System (Amazon EFS) volume to the EC2 instance
- H. Host the image on the EFS volume.

Answer: C

Explanation:

Create Amazon S3 gateway endpoint in the VPC and add a VPC endpoint policy. This VPC endpoint policy will have a statement that allows S3 access only via access points owned by the organization.

NEW QUESTION 96

- (Exam Topic 1)

A company is developing a new service that will be accessed using TCP on a static port. A solutions architect must ensure that the service is highly available, has redundancy across Availability Zones, and is accessible using the DNS name `myservice.com`, which is publicly accessible. The service must use fixed address assignments so other companies can add the addresses to their allow lists.

Assuming that resources are deployed in multiple Availability Zones in a single Region, which solution will meet these requirements?

- A. Create Amazon EC2 instances with an Elastic IP address for each instance Create a Network Load Balancer (NLB) and expose the static TCP port Register EC2 instances with the NLB Create a new name server record set named my service com, and assign the Elastic IP addresses of the EC2 instances to the record set Provide the Elastic IP addresses of the EC2 instances to the other companies to add to their allow lists
- B. Create an Amazon ECS cluster and a service definition for the application Create and assign public IP addresses for the ECS cluster Create a Network Load Balancer (NLB) and expose the TCP port Create a target group and assign the ECS cluster name to the NLB Create a new A record set named my service com and assign the public IP addresses of the ECS cluster to the record set Provide the public IP addresses of the ECS cluster to the other companies to add to their allow lists
- C. Create Amazon EC2 instances for the service Create one Elastic IP address for each Availability Zone Create a Network Load Balancer (NLB) and expose the assigned TCP port Assign the Elastic IP addresses to the NLB for each Availability Zone Create a target group and register the EC2 instances with the NLB Create a new A (alias) record set named my service com, and assign the NLB DNS name to the record set.
- D. Create an Amazon ECS cluster and a service definition for the application Create and assign public IP address for each host in the cluster Create an Application Load Balancer (ALB) and expose the static TCP port Create a target group and assign the ECS service definition name to the ALB Create a new CNAME record set and associate the public IP addresses to the record set Provide the Elastic IP addresses of the Amazon EC2 instances to the other companies to add to their allow lists

Answer: C

Explanation:

<https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-to-elb-load-balancer.html>

Create a Network Load Balancer (NLB) and expose the assigned TCP port. Assign the Elastic IP addresses to the NLB for each Availability Zone. Create a target group and register the EC2 instances with the NLB. Create a new A (alias) record set named my.service.com, and assign the NLB DNS name to the record set. As it uses the NLB as the resource in the A-record, traffic will be routed through the NLB, and it will automatically route the traffic to the healthy instances based on the health checks and also it provides the fixed address assignments as the other companies can add the NLB's Elastic IP addresses to their allow lists.

NEW QUESTION 99

- (Exam Topic 1)

A company is developing a new serverless API by using Amazon API Gateway and AWS Lambda. The company integrated the Lambda functions with API Gateway to use several shared libraries and custom classes.

A solutions architect needs to simplify the deployment of the solution and optimize for code reuse. Which solution will meet these requirements?

- A. Deploy the shared libraries and custom classes into a Docker image
- B. Store the image in an S3 bucket. Create a Lambda layer that uses the Docker image as the source
- C. Deploy the API's Lambda functions as Zip package
- D. Configure the packages to use the Lambda layer.
- E. Deploy the shared libraries and custom classes to a Docker image
- F. Upload the image to Amazon Elastic Container Registry (Amazon ECR). Create a Lambda layer that uses the Docker image as the source
- G. Deploy the API's Lambda functions as Zip package
- H. Configure the packages to use the Lambda layer.
- I. Deploy the shared libraries and custom classes to a Docker container in Amazon Elastic Container Service (Amazon ECS) by using the AWS Fargate launch type
- J. Deploy the API's Lambda functions as Zip package
- K. Configure the packages to use the deployed container as a Lambda layer.
- L. Deploy the shared libraries, custom classes, and code for the API's Lambda functions to a Docker image
- M. Upload the image to Amazon Elastic Container Registry (Amazon ECR). Configure the API's Lambda functions to use the Docker image as the deployment package.

Answer: B

Explanation:

Deploying the shared libraries and custom classes to a Docker image and uploading the image to Amazon Elastic Container Registry (Amazon ECR) and creating a Lambda layer that uses the Docker image as the source. Then, deploying the API's Lambda functions as Zip packages and configuring the packages to use the Lambda layer would meet the requirements for simplifying the deployment and optimizing for code reuse.

A Lambda layer is a distribution mechanism for libraries, custom runtimes, and other function dependencies. It allows you to manage your in-development function code separately from your dependencies, this way you can easily update your dependencies without having to update your entire function code.

By deploying the shared libraries and custom classes to a Docker image and uploading the image to Amazon Elastic Container Registry (ECR), it makes it easy to manage and version the dependencies. This way, the company can use the same version of the dependencies across different Lambda functions.

By creating a Lambda layer that uses the Docker image as the source, the company can configure the API's Lambda functions to use the layer, reducing the need to include the dependencies in each function package, and making it easy to update the dependencies across all functions at once.

Reference:

AWS Lambda Layers documentation: <https://docs.aws.amazon.com/lambda/latest/dg/configuration-layers.html>

AWS Elastic Container Registry (ECR) documentation: <https://aws.amazon.com/ecr/> Building Lambda Layers with Docker documentation:

<https://aws.amazon.com/blogs/compute/building-lambda-layers-with-docker/>

NEW QUESTION 102

- (Exam Topic 1)

A company wants to migrate its data analytics environment from on premises to AWS The environment consists of two simple Node.js applications One of the applications collects sensor data and loads it into a MySQL database The other application aggregates the data into reports When the aggregation jobs run, some of the load jobs fail to run correctly

The company must resolve the data loading issue The company also needs the migration to occur without interruptions or changes for the company's customers What should a solutions architect do to meet these requirements?

- A. Set up an Amazon Aurora MySQL database as a replication target for the on-premises database Create an Aurora Replica for the Aurora MySQL database, and move the aggregation jobs to run against the Aurora Replica Set up collection endpoints as AWS Lambda functions behind a Network Load Balancer (NLB). and use Amazon RDS Proxy to write to the Aurora MySQL database When the databases are synced disable the replication job and restart the Aurora Replica as the primary instance
- B. Point the collector DNS record to the NLB.
- C. Set up an Amazon Aurora MySQL database Use AWS Database Migration Service (AWS DMS) to perform continuous data replication from the on-premises database to Aurora Move the aggregation jobs to run against the Aurora MySQL database Set up collection endpoints behind an Application Load Balancer (ALB) as Amazon EC2 instances in an Auto Scaling group When the databases are synced, point the collector DNS record to the ALB Disable the AWS DMS sync task after the cutover from on premises to AWS
- D. Set up an Amazon Aurora MySQL database Use AWS Database Migration Service (AWS DMS) to perform continuous data replication from the on-premises database to Aurora Create an Aurora Replica for the Aurora MySQL database and move the aggregation jobs to run against the Aurora Replica Set up collection

endpoints as AWS Lambda functions behind an Application Load Balancer (ALB) and use Amazon RDS Proxy to write to the Aurora MySQL database. When the databases are synced, point the collector DNS record to the ALB. Disable the AWS DMS sync task after the cutover from on-premises to AWS.

E. Set up an Amazon Aurora MySQL database. Create an Aurora Replica for the Aurora MySQL database and move the aggregation jobs to run against the Aurora Replica. Set up collection endpoints as an Amazon Kinesis data stream. Use Amazon Kinesis Data Firehose to replicate the data to the Aurora MySQL database. When the databases are synced, disable the replication job and restart the Aurora Replica as the primary instance. Point the collector DNS record to the Kinesis data stream.

Answer: C

Explanation:

Set up an Amazon Aurora MySQL database. Use AWS Database Migration Service (AWS DMS) to perform continuous data replication from the on-premises database to Aurora. Create an Aurora Replica for the Aurora MySQL database, and move the aggregation jobs to run against the Aurora Replica. Set up collection endpoints as AWS Lambda functions behind an Application Load Balancer (ALB), and use Amazon RDS Proxy to write to the Aurora MySQL database. When the databases are synced, point the collector DNS record to the ALB. Disable the AWS DMS sync task after the cutover from on-premises to AWS. Amazon RDS Proxy allows applications to pool and share connections established with the database, improving database efficiency and application scalability. With RDS Proxy, failover times for Aurora and RDS databases are reduced by up to 66%.

NEW QUESTION 103

- (Exam Topic 1)

An application is using an Amazon RDS for MySQL Multi-AZ DB instance in the us-east-1 Region. After a failover test, the application lost the connections to the database and could not re-establish the connections. After a restart of the application, the application re-established the connections.

A solutions architect must implement a solution so that the application can re-establish connections to the database without requiring a restart.

Which solution will meet these requirements?

- A. Create an Amazon Aurora MySQL Serverless v1 DB instance
- B. Migrate the RDS DB instance to the Aurora Serverless v1 DB instance
- C. Update the connection settings in the application to point to the Aurora reader endpoint.
- D. Create an RDS proxy
- E. Configure the existing RDS endpoint as a target
- F. Update the connection settings in the application to point to the RDS proxy endpoint.
- G. Create a two-node Amazon Aurora MySQL DB cluster
- H. Migrate the RDS DB instance to the Aurora DB cluster
- I. Create an RDS proxy
- J. Configure the existing RDS endpoint as a target
- K. Update the connection settings in the application to point to the RDS proxy endpoint.
- L. Create an Amazon S3 bucket
- M. Export the database to Amazon S3 by using AWS Database Migration Service (AWS DMS). Configure Amazon Athena to use the S3 bucket as a data store
- N. Install the latest Open Database Connectivity (ODBC) driver for the application
- O. Update the connection settings in the application to point to the Athena endpoint

Answer: B

Explanation:

Amazon RDS Proxy is a fully managed database proxy service for Amazon Relational Database Service (RDS) that makes applications more scalable, resilient, and secure. It allows applications to pool and share connections to an RDS database, which can help reduce database connection overhead, improve scalability, and provide automatic failover and high availability.

NEW QUESTION 107

- (Exam Topic 1)

A company runs an IoT platform on AWS IoT sensors in various locations send data to the company's Node.js API servers on Amazon EC2 instances running behind an Application Load Balancer. The data is stored in an Amazon RDS MySQL DB instance that uses a 4 TB General Purpose SSD volume.

The number of sensors the company has deployed in the field has increased over time and is expected to grow significantly. The API servers are consistently overloaded and RDS metrics show high write latency.

Which of the following steps together will resolve the issues permanently and enable growth as new sensors are provisioned, while keeping this platform cost-efficient? (Select TWO.)

- A. Resize the MySQL General Purpose SSD storage to 6 TB to improve the volume's IOPS
- B. Re-architect the database tier to use Amazon Aurora instead of an RDS MySQL DB instance and add read replicas
- C. Leverage Amazon Kinesis Data Streams and AWS Lambda to ingest and process the raw data
- D. Use AWS X-Ray to analyze and debug application issues and add more API servers to match the load
- E. Re-architect the database tier to use Amazon DynamoDB instead of an RDS MySQL DB instance

Answer: CE

Explanation:

➤ Option C is correct because leveraging Amazon Kinesis Data Streams and AWS Lambda to ingest and process the raw data resolves the issues permanently and enable growth as new sensors are provisioned. Amazon Kinesis Data Streams is a serverless streaming data service that simplifies the capture, processing, and storage of data streams at any scale. Kinesis Data Streams can handle any amount of streaming data and process data from hundreds of thousands of sources with very low latency. AWS Lambda is a serverless compute service that lets you run code without provisioning or managing servers. Lambda can be triggered by Kinesis Data Streams events and process the data records in real time. Lambda can also scale automatically based on the incoming data volume. By using Kinesis Data Streams and Lambda, the company can reduce the load on the API servers and improve the performance and scalability of the data ingestion and processing layer.

➤ Option E is correct because re-architecting the database tier to use Amazon DynamoDB instead of an RDS MySQL DB instance resolves the issues permanently and enable growth as new sensors are provisioned. Amazon DynamoDB is a fully managed key-value and document database that delivers single-digit millisecond performance at any scale. DynamoDB supports auto scaling, which automatically adjusts read and write capacity based on actual traffic patterns. DynamoDB also supports on-demand capacity mode, which instantly accommodates up to double the previous peak traffic on a table. By using DynamoDB instead of RDS MySQL DB instance, the company can eliminate high write latency and improve scalability and performance of the database tier.

References: 1: <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-volume-types.html> 2:

https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/CHAP_AuroraOverview.html 3:

<https://docs.aws.amazon.com/streams/latest/dev/introduction.html> : <https://docs.aws.amazon.com/lambda/latest/dg/welcome.html> :

<https://docs.aws.amazon.com/xray/latest/devguide/aws-xray.html> : <https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Introduction.html> :

NEW QUESTION 111

- (Exam Topic 1)

A software as a service (SaaS) based company provides a case management solution to customers A3 part of the solution. The company uses a standalone Simple Mail Transfer Protocol (SMTP) server to send email messages from an application. The application also stores an email template for acknowledgement email messages that populate customer data before the application sends the email message to the customer.

The company plans to migrate this messaging functionality to the AWS Cloud and needs to minimize operational overhead.

Which solution will meet these requirements MOST cost-effectively?

- A. Set up an SMTP server on Amazon EC2 instances by using an AMI from the AWS Marketplace
- B. Store the email template in an Amazon S3 bucket
- C. Create an AWS Lambda function to retrieve the template from the S3 bucket and to merge the customer data from the application with the template
- D. Use an SDK in the Lambda function to send the email message.
- E. Set up Amazon Simple Email Service (Amazon SES) to send email message
- F. Store the email template in an Amazon S3 bucket
- G. Create an AWS Lambda function to retrieve the template from the S3 bucket and to merge the customer data from the application with the template
- H. Use an SDK in the Lambda function to send the email message.
- I. Set up an SMTP server on Amazon EC2 instances by using an AMI from the AWS Marketplace
- J. Store the email template in Amazon Simple Email Service (Amazon SES) with parameters for the customer data
- K. Create an AWS Lambda function to call the SES SendTemplatedEmail API operation and to pass customer data to replace the parameters
- L. Use the AWS Marketplace SMTP server to send the email message.
- M. Set up Amazon Simple Email Service (Amazon SES) to send email message
- N. Store the email template on Amazon SES with parameters for the customer data
- O. Create an AWS Lambda function to call the SendTemplatedEmail API operation and to pass customer data to replace the parameters and the email destination.

Answer: D

Explanation:

In this solution, the company can use Amazon SES to send email messages, which will minimize operational overhead as SES is a fully managed service that handles sending and receiving email messages. The company can store the email template on Amazon SES with parameters for the customer data and use an AWS Lambda function to call the SendTemplatedEmail API operation, passing in the customer data to replace the parameters and the email destination. This solution eliminates the need to set up and manage an SMTP server on EC2 instances, which can be costly and time-consuming.

NEW QUESTION 112

.....

Relate Links

100% Pass Your AWS-Certified-Solutions-Architect-Professional Exam with ExamBible Prep Materials

<https://www.exambible.com/AWS-Certified-Solutions-Architect-Professional-exam/>

Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - <https://www.exambible.com/>