

Amazon-Web-Services

Exam Questions SAP-C02

AWS Certified Solutions Architect - Professional



NEW QUESTION 1

- (Exam Topic 1)

A company wants to change its internal cloud billing strategy for each of its business units. Currently, the cloud governance team shares reports for overall cloud spending with the head of each business unit. The company uses AWS Organizations to manage the separate AWS accounts for each business unit. The existing tagging standard in Organizations includes the application, environment, and owner. The cloud governance team wants a centralized solution so each business unit receives monthly reports on its cloud spending. The solution should also send notifications for any cloud spending that exceeds a set threshold. Which solution is the MOST cost-effective way to meet these requirements?

- A. Configure AWS Budgets in each account and configure budget alerts that are grouped by application, environment, and owner
- B. Add each business unit to an Amazon SNS topic for each alert
- C. Use Cost Explorer in each account to create monthly reports for each business unit.
- D. Configure AWS Budgets in the organization's master account and configure budget alerts that are grouped by application, environment, and owner
- E. Add each business unit to an Amazon SNS topic for each alert
- F. Use Cost Explorer in the organization's master account to create monthly reports for each business unit.
- G. Configure AWS Budgets in each account and configure budget alerts that are grouped by application, environment, and owner
- H. Add each business unit to an Amazon SNS topic for each alert
- I. Use the AWS Billing and Cost Management dashboard in each account to create monthly reports for each business unit.
- J. Enable AWS Cost and Usage Reports in the organization's master account and configure reports grouped by application, environment, and owner
- K. Create an AWS Lambda function that processes AWS Cost and Usage Reports, sends budget alerts, and sends monthly reports to each business unit's email list.

Answer: B

Explanation:

Configure AWS Budgets in the organization's master account and configure budget alerts that are grouped by application, environment, and owner. Add each business unit to an Amazon SNS topic for each alert. Use Cost Explorer in the organization's master account to create monthly reports for each business unit.
<https://aws.amazon.com/about-aws/whats-new/2019/07/introducing-aws-budgets-reports/#:~:text=AWS%20Bud>

NEW QUESTION 2

- (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 instance
- H. Migrate the RDS DB instance to the Aurora DB instance
- 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 3

- (Exam Topic 1)

A company is running a critical application that uses an Amazon RDS for MySQL database to store data. The RDS DB instance is deployed in Multi-AZ mode. A recent RDS database failover test caused a 40-second outage to the application. A solutions architect needs to design a solution to reduce the outage time to less than 20 seconds.

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

- A. Use Amazon ElastiCache for Memcached in front of the database
- B. Use Amazon ElastiCache for Redis in front of the database.
- C. Use RDS Proxy in front of the database
- D. Migrate the database to Amazon Aurora MySQL
- E. Create an Amazon Aurora Replica
- F. Create an RDS for MySQL read replica

Answer: CDE

Explanation:

Migrate the database to Amazon Aurora MySQL. - Create an Amazon Aurora Replica. - Use RDS Proxy in front of the database. - These options are correct because they address the requirement of reducing the failover time to less than 20 seconds. Migrating to Amazon Aurora MySQL and creating an Aurora replica can reduce the failover time to less than 20 seconds. Aurora has a built-in, fault-tolerant storage system that can automatically detect and repair failures. Additionally, Aurora has a feature called "Aurora Global Database" which allows you to create read-only replicas across multiple AWS regions which can further

help to reduce the failover time. Creating an Aurora replica can also help to reduce the failover time as it can take over as the primary DB instance in case of a failure. Using RDS proxy can also help to reduce the failover time as it can route the queries to the healthy DB instance, it also helps to balance the load across multiple DB instances.

NEW QUESTION 4

- (Exam Topic 1)

A company is running a traditional web application on Amazon EC2 instances. The company needs to refactor the application as microservices that run on containers. Separate versions of the application exist in two distinct environments: production and testing. Load for the application is variable, but the minimum load and the maximum load are known. A solutions architect needs to design the updated application with a serverless architecture that minimizes operational complexity.

Which solution will meet these requirements MOST cost-effectively?

- A. Upload the container images to AWS Lambda as function
- B. Configure a concurrency limit for the associated Lambda functions to handle the expected peak load
- C. Configure two separate Lambda integrations within Amazon API Gateway: one for production and one for testing.
- D. Upload the container images to Amazon Elastic Container Registry (Amazon ECR). Configure two auto scaled Amazon Elastic Container Service (Amazon ECS) clusters with the Fargate launch type to handle the expected load
- E. Deploy tasks from the ECR image
- F. Configure two separate Application Load Balancers to direct traffic to the ECS clusters.
- G. Upload the container images to Amazon Elastic Container Registry (Amazon ECR). Configure two auto scaled Amazon Elastic Kubernetes Service (Amazon EKS) clusters with the Fargate launch type to handle the expected load
- H. Deploy tasks from the ECR image
- I. Configure two separate Application Load Balancers to direct traffic to the EKS clusters.
- J. Upload the container images to AWS Elastic Beanstalk
- K. In Elastic Beanstalk, create separate environments and deployments for production and testing
- L. Configure two separate Application Load Balancers to direct traffic to the Elastic Beanstalk deployments.

Answer: B

Explanation:

minimizes operational + microservices that run on containers = AWS Elastic Beanstalk

NEW QUESTION 5

- (Exam Topic 1)

A company has an organization in AWS Organizations that has a large number of AWS accounts. One of the AWS accounts is designated as a transit account and has a transit gateway that is shared with all of the other AWS accounts. AWS Site-to-Site VPN connections are configured between all of the company's global offices and the transit account. The company has AWS Config enabled on all of its accounts.

The company's networking team needs to centrally manage a list of internal IP address ranges that belong to the global offices. Developers will reference this list to gain access to applications securely.

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

- A. Create a JSON file that is hosted in Amazon S3 and that lists all of the internal IP address ranges. Configure an Amazon Simple Notification Service (Amazon SNS) topic in each of the accounts that can be involved when the JSON file is updated.
- B. Subscribe an AWS Lambda function to the SNS topic to update all relevant security group rules with the updated IP address ranges.
- C. Create a new AWS Config managed rule that contains all of the internal IP address ranges. Use the rule to check the security groups in each of the accounts to ensure compliance with the list of IP address ranges.
- D. Configure the rule to automatically remediate any noncompliant security group that is detected.
- E. In the transit account, create a VPC prefix list with all of the internal IP address ranges.
- F. Use AWS Resource Access Manager to share the prefix list with all of the other accounts.
- G. Use the shared prefix list to configure security group rules in the other accounts.
- H. In the transit account, create a security group with all of the internal IP address ranges.
- I. Configure the security groups in the other accounts to reference the transit account's security group by using a nested security group reference of `*<transit-account-id>/.sg-1a2b3c4d`.

Answer: C

Explanation:

Customer-managed prefix lists — Sets of IP address ranges that you define and manage. You can share your prefix list with other AWS accounts, enabling those accounts to reference the prefix list in their own resources. <https://docs.aws.amazon.com/vpc/latest/userguide/managed-prefix-lists.html>

A VPC prefix list is created in the transit account with all of the internal IP address ranges, and then shared to all of the other accounts using AWS Resource Access Manager. This allows for central management of the IP address ranges, and eliminates the need for manual updates to security group rules in each account. This solution also allows for compliance checks to be run using AWS Config and for any non-compliant security groups to be automatically remediated.

NEW QUESTION 6

- (Exam Topic 1)

A solutions architect has developed a web application that uses an Amazon API Gateway Regional endpoint and an AWS Lambda function. The consumers of the web application are all close to the AWS Region where the application will be deployed. The Lambda function only queries an Amazon Aurora MySQL database. The solutions architect has configured the database to have three read replicas.

During testing, the application does not meet performance requirements. Under high load, the application opens a large number of database connections. The solutions architect must improve the application's performance.

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

- A. Use the cluster endpoint of the Aurora database.
- B. Use RDS Proxy to set up a connection pool to the reader endpoint of the Aurora database.
- C. Use the Lambda Provisioned Concurrency feature.
- D. Move the code for opening the database connection in the Lambda function outside of the event handler.
- E. Change the API Gateway endpoint to an edge-optimized endpoint.

Answer: BD

Explanation:

Connect to RDS outside of Lambda handler method to improve performance <https://awstut.com/en/2022/04/30/connect-to-rds-outside-of-lambda-handler-method-to-improve-performance-en>

Using RDS Proxy, you can handle unpredictable surges in database traffic. Otherwise, these surges might cause issues due to oversubscribing connections or creating new connections at a fast rate. RDS Proxy establishes a database connection pool and reuses connections in this pool. This approach avoids the memory and CPU overhead of opening a new database connection each time. To protect the database against oversubscription, you can control the number of database connections that are created. <https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/rds-proxy.html>

NEW QUESTION 7

- (Exam Topic 1)

A company is planning to migrate its business-critical applications from an on-premises data center to AWS. The company has an on-premises installation of a Microsoft SQL Server Always On cluster. The company wants to migrate to an AWS managed database service. A solutions architect must design a heterogeneous database migration on AWS.

Which solution will meet these requirements?

- A. Migrate the SQL Server databases to Amazon RDS for MySQL by using backup and restore utilities.
- B. Use an AWS Snowball Edge Storage Optimized device to transfer data to Amazon S3. Set up Amazon RDS for MySQL
- C. Use S3 integration with SQL Server features, such as BULK INSERT.
- D. Use the AWS Schema Conversion Tool to translate the database schema to Amazon RDS for MySQL
- E. Then use AWS Database Migration Service (AWS DMS) to migrate the data from on-premises databases to Amazon RDS.
- F. Use AWS DataSync to migrate data over the network between on-premises storage and Amazon S3. Set up Amazon RDS for MySQL
- G. Use S3 integration with SQL Server features, such as BULK INSERT.

Answer: C

Explanation:

<https://aws.amazon.com/dms/schema-conversion-tool/>

AWS Schema Conversion Tool (SCT) can automatically convert the database schema from Microsoft SQL Server to Amazon RDS for MySQL. This allows for a smooth transition of the database schema without any manual intervention. AWS DMS can then be used to migrate the data from the on-premises databases to the newly created Amazon RDS for MySQL instance. This service can perform a one-time migration of the data or can set up ongoing replication of data changes to keep the on-premises and AWS databases in sync.

NEW QUESTION 8

- (Exam Topic 1)

A company developed a pilot application by using AWS Elastic Beanstalk and Java. To save costs during development, the company's development team deployed the application into a single-instance environment. Recent tests indicate that the application consumes more CPU than expected. CPU utilization is regularly greater than 85%, which causes some performance bottlenecks.

A solutions architect must mitigate the performance issues before the company launches the application to production.

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

- A. Create a new Elastic Beanstalk application
- B. Select a load-balanced environment type
- C. Select all Availability Zones
- D. Add a scale-out rule that will run if the maximum CPU utilization is over 85% for 5 minutes.
- E. Create a second Elastic Beanstalk environment
- F. Apply the traffic-splitting deployment policy
- G. Specify a percentage of incoming traffic to direct to the new environment in the average CPU utilization is over 85% for 5 minutes.
- H. Modify the existing environment's capacity configuration to use a load-balanced environment type. Select all Availability Zones
- I. Add a scale-out rule that will run if the average CPU utilization is over 85% for 5 minutes.
- J. Select the Rebuild environment action with the load balancing option. Select an Availability Zone. Add a scale-out rule that will run if the sum CPU utilization is over 85% for 5 minutes.

Answer: C

Explanation:

This solution will meet the requirements with the least operational overhead because it allows the company to modify the existing environment's capacity configuration, so it becomes a load-balanced environment type. By selecting all availability zones, the company can ensure that the application is running in multiple availability zones, which can help to improve the availability and scalability of the application. The company can also add a scale-out rule that will run if the average CPU utilization is over 85% for 5 minutes, which can help to mitigate the performance issues. This solution does not require creating new Elastic Beanstalk environments or rebuilding the existing one, which reduces the operational overhead.

You can refer to the AWS Elastic Beanstalk documentation for more information on how to use this service: <https://aws.amazon.com/elasticbeanstalk/> You can refer to the AWS documentation for more information on how to use autoscaling: <https://aws.amazon.com/autoscaling/>

NEW QUESTION 9

- (Exam Topic 1)

A company is planning to store a large number of archived documents and make the documents available to employees through the corporate intranet. Employees will access the system by connecting through a client VPN service that is attached to a VPC. The data must not be accessible to the public.

The documents that the company is storing are copies of data that is held on physical media elsewhere. The number of requests will be low. Availability and speed of retrieval are not concerns of the company.

Which solution will meet these requirements at the LOWEST cost?

- A. Create an Amazon S3 bucket
- B. Configure the S3 bucket to use the S3 One Zone-Infrequent Access (S3 One Zone-IA) storage class as default
- C. Configure the S3 bucket for website hosting
- D. Create an S3 interface endpoint
- E. Configure the S3 bucket to allow access only through that endpoint.
- F. Launch an Amazon EC2 instance that runs a web server
- G. Attach an Amazon Elastic File System (Amazon EFS) file system to store the archived data in the EFS One Zone-Infrequent Access (EFS One Zone-IA) storage class. Configure the instance security groups to allow access only from private networks.
- H. Launch an Amazon EC2 instance that runs a web server. Attach an Amazon Elastic Block Store (Amazon EBS) volume to store the archived data

- I. Use the Cold HDD (sc1) volume typ
- J. Configure the instance security groups to allow access only from private networks.
- K. Create an Amazon S3 bucket
- L. Configure the S3 bucket to use the S3 Glacier Deep Archive storage class as default
- M. Configure the S3 bucket for website hosting
- N. Create an S3 interface endpoint
- O. Configure the S3 bucket to allow access only through that endpoint.

Answer: D

Explanation:

The S3 Glacier Deep Archive storage class is the lowest-cost storage class offered by Amazon S3, and it is designed for archival data that is accessed infrequently and for which retrieval time of several hours is acceptable. S3 interface endpoint for the VPC ensures that access to the bucket is only from resources within the VPC and this will meet the requirement of not being accessible to the public. And also, S3 bucket can be configured for website hosting, and this will allow employees to access the documents through the corporate intranet. Using an EC2 instance and a file system or block store would be more expensive and unnecessary because the number of requests to the data will be low and availability and speed of retrieval are not concerns. Additionally, using Amazon S3 bucket will provide durability, scalability and availability of data.

NEW QUESTION 10

- (Exam Topic 1)

A company runs a Java application that has complex dependencies on VMs that are in the company's data center. The application is stable. but the company wants to modernize the technology stack. The company wants to migrate the application to AWS and minimize the administrative overhead to maintain the servers.

Which solution will meet these requirements with the LEAST code changes?

- A. Migrate the application to Amazon Elastic Container Service (Amazon ECS) on AWS Fargate by using AWS App2Container
- B. Store container images in Amazon Elastic Container Registry (Amazon ECR). Grant the ECS task execution role permission to access the ECR image repository
- C. Configure Amazon ECS to use an Application Load Balancer (ALB). Use the ALB to interact with the application.
- D. Migrate the application code to a container that runs in AWS Lambda
- E. Build an Amazon API Gateway REST API with Lambda integration
- F. Use API Gateway to interact with the application.
- G. Migrate the application to Amazon Elastic Kubernetes Service (Amazon EKS) on EKS managed node groups by using AWS App2Container
- H. Store container images in Amazon Elastic Container Registry (Amazon ECR). Give the EKS nodes permission to access the ECR image repository
- I. Use Amazon API Gateway to interact with the application.
- J. Migrate the application code to a container that runs in AWS Lambda
- K. Configure Lambda to use an Application Load Balancer (ALB). Use the ALB to interact with the application.

Answer: A

Explanation:

According to the AWS documentation¹, AWS App2Container (A2C) is a command line tool for migrating and modernizing Java and .NET web applications into container format. AWS A2C analyzes and builds an inventory of applications running in bare metal, virtual machines, Amazon Elastic Compute Cloud (EC2) instances, or in the cloud. You can use AWS A2C to generate container images for your applications and deploy them on Amazon ECS or Amazon EKS. Option A meets the requirements of the scenario because it allows you to migrate your existing Java application to AWS and minimize the administrative overhead to maintain the servers. You can use AWS A2C to analyze your application dependencies, extract application artifacts, and generate a Dockerfile. You can then store your container images in Amazon ECR, which is a fully managed container registry service. You can use AWS Fargate as the launch type for your Amazon ECS cluster, which is a serverless compute engine that eliminates the need to provision and manage servers for your containers. You can grant the ECS task execution role permission to access the ECR image repository, which allows your tasks to pull images from ECR. You can configure Amazon ECS to use an ALB, which is a load balancer that distributes traffic across multiple targets in multiple Availability Zones using HTTP or HTTPS protocols. You can use the ALB to interact with your application.

NEW QUESTION 10

- (Exam Topic 1)

A financial services company receives a regular data feed from its credit card servicing partner. Approximately 5.1 records are sent every 15 minutes in plaintext, delivered over HTTPS directly into an Amazon S3 bucket with server-side encryption. This feed contains sensitive credit card primary account number (PAN) data. The company needs to automatically mask the PAN before sending the data to another S3 bucket for additional internal processing. The company also needs to remove and merge specific fields, and then transform the record into JSON format. Additionally, extra feeds are likely to be added in the future, so any design needs to be easily expandable.

Which solutions will meet these requirements?

- A. Trigger an AWS Lambda function on file delivery that extracts each record and writes it to an Amazon SQS queue
- B. Trigger another Lambda function when new messages arrive in the SQS queue to process the records, writing the results to a temporary location in Amazon S3. Trigger a final Lambda function once the SQS queue is empty to transform the records into JSON format and send the results to another S3 bucket for internal processing.
- C. Trigger an AWS Lambda function on file delivery that extracts each record and writes it to an Amazon SQS queue
- D. Configure an AWS Fargate container application to automatically scale to a single instance when the SQS queue contains a message
- E. Have the application process each record, and transform the record into JSON format
- F. When the queue is empty, send the results to another S3 bucket for internal processing and scale down the AWS Fargate instance.
- G. Create an AWS Glue crawler and custom classifier based on the data feed formats and build a table definition to match
- H. Trigger an AWS Lambda function on file delivery to start an AWS Glue ETL job to transform the entire record according to the processing and transformation requirements
- I. Define the output format as JSON
- J. Once complete, have the ETL job send the results to another S3 bucket for internal processing.
- K. Create an AWS Glue crawler and custom classifier based upon the data feed formats and build a table definition to match
- L. Perform an Amazon Athena query on file delivery to start an Amazon EMR ETL job to transform the entire record according to the processing and transformation requirements
- M. Define the output format as JSON
- N. Once complete, send the results to another S3 bucket for internal processing and scale down the EMR cluster.

Answer: C

Explanation:

You can use a Glue crawler to populate the AWS Glue Data Catalog with tables. The Lambda function can be triggered using S3 event notifications when object create events occur. The Lambda function will then trigger the Glue ETL job to transform the records masking the sensitive data and modifying the output format to JSON. This solution meets all requirements.

NEW QUESTION 12

- (Exam Topic 1)

A company is creating a sequel for a popular online game. A large number of users from all over the world will play the game within the first week after launch. Currently, the game consists of the following components deployed in a single AWS Region:

- Amazon S3 bucket that stores game assets
- Amazon DynamoDB table that stores player scores

A solutions architect needs to design a multi-Region solution that will reduce latency improve reliability, and require the least effort to implement

What should the solutions architect do to meet these requirements?

- A. Create an Amazon CloudFront distribution to serve assets from the S3 bucket Configure S3Cross-Region Replication Create a new DynamoDB able in a new Region Use the new table as a replica target tor DynamoDB global tables.
- B. Create an Amazon CloudFront distribution to serve assets from the S3 bucke
- C. Configure S3Same-Region Replicatio
- D. Create a new DynamoDB able m a new Regio
- E. Configure asynchronous replication between the DynamoDB tables by using AWS Database Migration Service (AWS DMS) with change data capture (CDC)
- F. Create another S3 bucket in a new Region and configure S3 Cross-Region Replication between the buckets Create an Amazon CloudFront distribution and configure origin failover with two origins accessing the S3 buckets in each Regio
- G. Configure DynamoDB global tables by enabling Amazon DynamoDB Streams, and add a replica table in a new Region.
- H. Create another S3 bucket in the same Region, and configure S3 Same-Region Replication between the buckets- Create an Amazon CloudFront distribution and configure origin failover with two origin accessing the S3 buckets Create a new DynamoDB table m a new Region Use the new table as a replica target for DynamoDB global tables.

Answer: C

Explanation:

https://aws.amazon.com/premiumsupport/knowledge-center/dynamodb-global-table-stream-lambda/?nc1=h_ls

NEW QUESTION 15

- (Exam Topic 1)

A company is using AWS Organizations lo manage multiple AWS accounts For security purposes, the company requires the creation of an Amazon Simple Notification Service (Amazon SNS) topic that enables integration with a third-party alerting system in all the Organizations member accounts

A solutions architect used an AWS CloudFormation template to create the SNS topic and stack sets to automate the deployment of CloudFormation stacks Trusted access has been enabled in Organizations

What should the solutions architect do to deploy the CloudFormation StackSets in all AWS accounts?

- A. Create a stack set in the Organizations member account
- B. Use service-managed permission
- C. Set deployment options to deploy to an organizatio
- D. Use CloudFormation StackSets drift detection.
- E. Create stacks in the Organizations member account
- F. Use self-service permission
- G. Set deployment options to deploy to an organizatio
- H. Enable the CloudFormation StackSets automatic deployment.
- I. Create a stack set in the Organizations management account Use service-managed permission
- J. Set deployment options to deploy to the organizatio
- K. Enable CloudFormation StackSets automatic deployment.
- L. Create stacks in the Organizations management accoun
- M. Use service-managed permission
- N. Set deployment options to deploy to the organizatio
- O. Enable CloudFormation StackSets drift detection.

Answer: C

Explanation:

<https://aws.amazon.com/blogs/aws/use-cloudformation-stacksets-to-provision-resources-across-multiple-aws-ac>

NEW QUESTION 19

- (Exam Topic 1)

A company is hosting a critical application on a single Amazon EC2 instance. The application uses an Amazon ElastiCache for Redis single-node cluster for an in-memory data store. The application uses an Amazon RDS for MariaDB DB instance for a relational database. For the application to function, each piece of the infrastructure must be healthy and must be in an active state.

A solutions architect needs to improve the application's architecture so that the infrastructure can automatically recover from failure with the least possible downtime.

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

- A. Use an Elastic Load Balancer to distribute traffic across multiple EC2 instance
- B. Ensure that the EC2 instances are part of an Auto Scaling group that has a minimum capacity of two instances.
- C. Use an Elastic Load Balancer to distribute traffic across multiple EC2 instances Ensure that the EC2 instances are configured in unlimited mode.
- D. Modify the DB instance to create a read replica in the same Availability Zon
- E. Promote the read replica to be the primary DB instance in failure scenarios.
- F. Modify the DB instance to create a Multi-AZ deployment that extends across two Availability Zones.
- G. Create a replication group for the ElastiCache for Redis cluste
- H. Configure the cluster to use an Auto Scaling group that has a minimum capacity of two instances.
- I. Create a replication group for the ElastiCache for Redis cluste
- J. Enable Multi-AZ on the cluster.

Answer: ADF

Explanation:

- Option A is correct because using an Elastic Load Balancer and an Auto Scaling group with a minimum capacity of two instances can improve the availability and scalability of the EC2 instances that host the application. The load balancer can distribute traffic across multiple instances and the Auto Scaling group can replace any unhealthy instances automatically1
 - Option D is correct because modifying the DB instance to create a Multi-AZ deployment that extends across two Availability Zones can improve the availability and durability of the RDS for MariaDB database. Multi-AZ deployments provide enhanced data protection and minimize downtime by automatically failing over to a standby replica in another Availability Zone in case of a planned or unplanned outage4
 - Option F is correct because creating a replication group for the ElastiCache for Redis cluster and enabling Multi-AZ on the cluster can improve the availability and fault tolerance of the in-memory data store. A replication group consists of a primary node and up to five read-only replica nodes that are synchronized with the primary node using asynchronous replication. Multi-AZ allows automatic failover to one of the replicas if the primary node fails or becomes unreachable6
- References: 1: <https://docs.aws.amazon.com/elasticloadbalancing/latest/userguide/how-elastic-load-balancing-works.html> 2: <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/burstable-performance-instances-unlimited-mode.htm> 3: https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_ReadRepl.html 4: <https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Concepts.MultiAZ.html> 5: <https://docs.aws.amazon.com/AmazonElastiCache/latest/red-ug/AutoScaling.html> 6: <https://docs.aws.amazon.com/AmazonElastiCache/latest/red-ug/Replication.Redis.Groups.html>

NEW QUESTION 22

- (Exam Topic 1)

A financial services company in North America plans to release a new online web application to its customers on AWS . The company will launch the application in the us-east-1 Region on Amazon EC2 instances. The application must be highly available and must dynamically scale to meet user traffic. The company also wants to implement a disaster recovery environment for the application in the us-west-1 Region by using active-passive failover. Which solution will meet these requirements?

- A. Create a VPC in us-east-1 and a VPC in us-west-1 Configure VPC peering In the us-east-1 VP
- B. create an Application Load Balancer (ALB) that extends across multiple Availability Zones in both VPCs Create an Auto Scaling group that deploys the EC2 instances across the multiple Availability Zones in both VPCs Place the Auto Scaling group behind the ALB.
- C. Create a VPC in us-east-1 and a VPC in us-west-1. In the us-east-1 VP
- D. create an Application Load Balancer (ALB) that extends across multiple Availability Zones in that VP
- E. Create an Auto Scaling group that deploys the EC2 instances across the multiple Availability Zones in the us-east-1 VPC Place the Auto Scaling group behind the ALB Set up the same configuration in the us-west-1 VP
- F. Create an Amazon Route 53 hosted zone Create separate records for each ALB Enable health checks to ensure high availability between Regions.
- G. Create a VPC in us-east-1 and a VPC in us-west-1 In the us-east-1 VP
- H. create an Application Load Balancer (ALB) that extends across multiple Availability Zones in that VPC Create an Auto Scaling group that deploys the EC2 instances across the multiple Availability Zones in the us-east-1 VPC Place the Auto Scaling group behind the ALB Set up the same configuration in the us-west-1 VPC Create an Amazon Route 53 hosted zon
- I. Create separate records for each ALB Enable health checks and configure a failover routing policy for each record.
- J. Create a VPC in us-east-1 and a VPC in us-west-1 Configure VPC peering In the us-east-1 VP
- K. create an Application Load Balancer (ALB) that extends across multiple Availability Zones in Create an Auto Scaling group that deploys the EC2 instances across the multiple Availability Zones in both VPCs Place the Auto Scaling group behind the ALB Create an Amazon Route 53 host.. Create a record for the ALB.

Answer: C

Explanation:

it's the one that handles failover while B (the one shown as the answer today) it almost the same but does not handle failover.

NEW QUESTION 25

- (Exam Topic 1)

A company is using multiple AWS accounts The DNS records are stored in a private hosted zone for Amazon Route 53 in Account A The company's applications and databases are running in Account B.

A solutions architect win deploy a two-net application In a new VPC To simplify the configuration, the db.example com CNAME record set for the Amazon RDS endpoint was created in a private hosted zone for Amazon Route 53.

During deployment, the application failed to start. Troubleshooting revealed that db.example com is not resolvable on the Amazon EC2 instance The solutions architect confirmed that the record set was created correctly in Route 53.

Which combination of steps should the solutions architect take to resolve this issue? (Select TWO)

- A. Deploy the database on a separate EC2 instance in the new VPC Create a record set for the instance's private IP in the private hosted zone
- B. Use SSH to connect to the application tier EC2 instance Add an RDS endpoint IP address to the/etc/resolv.conf file
- C. Create an authorization lo associate the private hosted zone in Account A with the new VPC In Account B
- D. Create a private hosted zone for the example.com domain m Account B Configure Route 53 replication between AWS accounts
- E. Associate a new VPC in Account B with a hosted zone in Account
- F. Delete the association authorization In Account A.

Answer: CE

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/private-hosted-zone-different-account/>

NEW QUESTION 29

- (Exam Topic 1)

A company's solutions architect is reviewing a web application that runs on AWS. The application references static assets in an Amazon S3 bucket in the us-east-1 Region. The company needs resiliency across multiple AWS Regions. The company already has created an S3 bucket in a second Region. Which solution will meet these requirements with the LEAST operational overhead?

- A. Configure the application to write each object to both S3 bucket

- B. Set up an Amazon Route 53 public hosted zone with a record set by using a weighted routing policy for each S3 bucket.
- C. Configure the application to reference the objects by using the Route 53 DNS name.
- D. Create an AWS Lambda function to copy objects from the S3 bucket in us-east-1 to the S3 bucket in the second Region.
- E. Invoke the Lambda function each time an object is written to the S3 bucket in us-east-1. Set up an Amazon CloudFront distribution with an origin group that contains the two S3 buckets as origins.
- F. Configure replication on the S3 bucket in us-east-1 to replicate objects to the S3 bucket in the second Region. Set up an Amazon CloudFront distribution with an origin group that contains the two S3 buckets as origins.
- G. Configure replication on the S3 bucket in us-east-1 to replicate objects to the S3 bucket in the second Region.
- H. If failover is required, update the application code to load S3 objects from the S3 bucket in the second Region.

Answer: C

Explanation:

https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/high_availability_origin_failover.html

NEW QUESTION 34

- (Exam Topic 1)

A financial company is planning to migrate its web application from on premises to AWS. The company uses a third-party security tool to monitor the inbound traffic to the application. The company has used the security tool for the last 15 years, and the tool has no cloud solutions available from its vendor. The company's security team is concerned about how to integrate the security tool with AWS technology.

The company plans to deploy the application migration to AWS on Amazon EC2 instances. The EC2 instances will run in an Auto Scaling group in a dedicated VPC. The company needs to use the security tool to inspect all packets that come in and out of the VPC. This inspection must occur in real time and must not affect the application's performance. A solutions architect must design a target architecture on AWS that is highly available within an AWS Region.

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

- A. Deploy the security tool on EC2 instances in a new Auto Scaling group in the existing VPC.
- B. Deploy the web application behind a Network Load Balancer.
- C. Deploy an Application Load Balancer in front of the security tool instances.
- D. Provision a Gateway Load Balancer for each Availability Zone to redirect the traffic to the security tool.
- E. Provision a transit gateway to facilitate communication between VPCs.

Answer: AD

Explanation:

Option A, Deploy the security tool on EC2 instances in a new Auto Scaling group in the existing VPC, allows the company to use its existing security tool while still running it within the AWS environment. This ensures that all packets coming in and out of the VPC are inspected by the security tool in real time. Option D, Provision a Gateway Load Balancer for each Availability Zone to redirect the traffic to the security tool, allows for high availability within an AWS Region. By provisioning a Gateway Load Balancer for each Availability Zone, the traffic is redirected to the security tool in the event of any failures or outages. This ensures that the security tool is always available to inspect the traffic, even in the event of a failure.

NEW QUESTION 35

- (Exam Topic 1)

A company has registered 10 new domain names. The company uses the domains for online marketing. The company needs a solution that will redirect online visitors to a specific URL for each domain. All domains and target URLs are defined in a JSON document. All DNS records are managed by Amazon Route 53. A solutions architect must implement a redirect service that accepts HTTP and HTTPS requests.

Which combination of steps should the solutions architect take to meet these requirements with the LEAST amount of operational effort? (Choose three.)

- A. Create a dynamic webpage that runs on an Amazon EC2 instance.
- B. Configure the webpage to use the JSON document in combination with the event message to look up and respond with a redirect URL.
- C. Create an Application Load Balancer that includes HTTP and HTTPS listeners.
- D. Create an AWS Lambda function that uses the JSON document in combination with the event message to look up and respond with a redirect URL.
- E. Use an Amazon API Gateway API with a custom domain to publish an AWS Lambda function.
- F. Create an Amazon CloudFront distribution.
- G. Deploy a Lambda@Edge function.
- H. Create an SSL certificate by using AWS Certificate Manager (ACM). Include the domains as Subject Alternative Names.

Answer: CEF

Explanation:

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/lambda-edge-how-it-works-tutorial.html>

NEW QUESTION 38

- (Exam Topic 1)

A weather service provides high-resolution weather maps from a web application hosted on AWS in the eu-west-1 Region. The weather maps are updated frequently and stored in Amazon S3 along with static HTML content. The web application is fronted by Amazon CloudFront.

The company recently expanded to serve users in the us-east-1 Region, and these new users report that viewing their respective weather maps is slow from time to time.

Which combination of steps will resolve the us-east-1 performance issues? (Choose two.)

- A. Configure the AWS Global Accelerator endpoint for the S3 bucket in eu-west-1. Configure endpoint groups for TCP ports 80 and 443 in us-east-1.
- B. Create a new S3 bucket in us-east-1. Configure S3 cross-Region replication to synchronize from the S3 bucket in eu-west-1.
- C. Use Lambda@Edge to modify requests from North America to use the S3 Transfer Acceleration endpoint in us-east-1.
- D. Use Lambda@Edge to modify requests from North America to use the S3 bucket in us-east-1.
- E. Configure the AWS Global Accelerator endpoint for us-east-1 as an origin on the CloudFront distribution.
- F. Use Lambda@Edge to modify requests from North America to use the new origin.

Answer: BD

Explanation:

<https://aws.amazon.com/about-aws/whats-new/2016/04/transfer-files-into-amazon-s3-up-to-300-percent-faster/>

NEW QUESTION 43

- (Exam Topic 1)

A team collects and routes behavioral data for an entire company. The company runs a Multi-AZ VPC environment with public subnets, private subnets, and an internet gateway. Each public subnet also contains a NAT gateway. Most of the company's applications read from and write to Amazon Kinesis Data Streams. Most of the workloads are in private subnets. A solutions architect must review the infrastructure. The solutions architect needs to reduce costs and maintain the function of the applications. The solutions architect uses Cost Explorer and notices that the cost in the EC2-Other category is consistently high. A further review shows that NatGateway-Bytes charges are increasing the cost in the EC2-Other category. What should the solutions architect do to meet these requirements?

- A. Enable VPC Flow Log
- B. Use Amazon Athena to analyze the logs for traffic that can be removed
- C. Ensure that security groups are blocking traffic that is responsible for high costs.
- D. Add an interface VPC endpoint for Kinesis Data Streams to the VPC
- E. Ensure that applications have the correct IAM permissions to use the interface VPC endpoint.
- F. Enable VPC Flow Logs and Amazon Detective. Review Detective findings for traffic that is not related to Kinesis Data Streams. Configure security groups to block that traffic.
- G. Add an interface VPC endpoint for Kinesis Data Streams to the VPC
- H. Ensure that the VPC endpoint policy allows traffic from the applications.

Answer: D

Explanation:

<https://docs.aws.amazon.com/vpc/latest/privatelink/vpc-endpoints-access.html> <https://aws.amazon.com/premiumsupport/knowledge-center/vpc-reduce-nat-gateway-transfer-costs/>

VPC endpoint policies enable you to control access by either attaching a policy to a VPC endpoint or by using additional fields in a policy that is attached to an IAM user, group, or role to restrict access to only occur via the specified VPC endpoint.

NEW QUESTION 46

- (Exam Topic 1)

A company is storing data on-premises on a Windows file server. The company produces 5 GB of new data daily. The company migrated part of its Windows-based workload to AWS and needs the data to be available on a file system in the cloud. The company already has established an AWS Direct Connect connection between the on-premises network and AWS. Which data migration strategy should the company use?

- A. Use the file gateway option in AWS Storage Gateway to replace the existing Windows file server, and point the existing file share to the new file gateway.
- B. Use AWS DataSync to schedule a daily task to replicate data between the on-premises Windows file server and Amazon FSx.
- C. Use AWS Data Pipeline to schedule a daily task to replicate data between the on-premises Windows file server and Amazon Elastic File System (Amazon EFS).
- D. Use AWS DataSync to schedule a daily task to replicate data between the on-premises Windows file server and Amazon Elastic File System (Amazon EFS).

Answer: B

Explanation:

<https://aws.amazon.com/storagegateway/file/>

<https://docs.aws.amazon.com/fsx/latest/WindowsGuide/migrate-files-to-fsx-datasync.html> <https://docs.aws.amazon.com/systems-manager/latest/userguide/prereqs-operating-systems.html#prereqs-os-win>

NEW QUESTION 48

- (Exam Topic 1)

A company has its cloud infrastructure on AWS. A solutions architect needs to define the infrastructure as code. The infrastructure is currently deployed in one AWS Region. The company's business expansion plan includes deployments in multiple Regions across multiple AWS accounts. What should the solutions architect do to meet these requirements?

- A. Use AWS CloudFormation templates. Add IAM policies to control the various accounts. Deploy the templates across the multiple Regions.
- B. Use AWS Organizations. Deploy AWS CloudFormation templates from the management account. Use AWS Control Tower to manage deployments across accounts.
- C. Use AWS Organizations and AWS CloudFormation StackSets. Deploy a CloudFormation template from an account that has the necessary IAM permissions.
- D. Use nested stacks with AWS CloudFormation templates. Change the Region by using nested stacks.

Answer: C

Explanation:

<https://aws.amazon.com/blogs/aws/new-use-aws-cloudformation-stacksets-for-multiple-accounts-in-an-aws-org/> AWS Organizations allows the management of multiple AWS accounts as a single entity and AWS

CloudFormation StackSets allows creating, updating, and deleting stacks across multiple accounts and regions in an organization. This solution allows creating a single CloudFormation template that can be deployed across multiple accounts and regions, and also allows for the management of access and permissions for the different accounts through the use of IAM roles and policies in the management account.

NEW QUESTION 52

- (Exam Topic 1)

A company is storing data in several Amazon DynamoDB tables. A solutions architect must use a serverless architecture to make the data accessible publicly through a simple API over HTTPS. The solution must scale automatically in response to demand. Which solutions meet these requirements? (Choose two.)

- A. Create an Amazon API Gateway REST API
- B. Configure this API with direct integrations to DynamoDB by using API Gateway's AWS integration type.

- C. Create an Amazon API Gateway HTTP AP
- D. Configure this API with direct integrations to Dynamo DB by using API Gateway's AWS integration type.
- E. Create an Amazon API Gateway HTTP AP
- F. Configure this API with integrations to AWS Lambda functions that return data from the DynamoDB tables.
- G. Create an accelerator in AWS Global Accelerator
- H. Configure this accelerator with AWS Lambda@Edge function integrations that return data from the DynamoDB tables.
- I. Create a Network Load Balance
- J. Configure listener rules to forward requests to the appropriate AWS Lambda functions

Answer: AC

Explanation:

<https://docs.aws.amazon.com/apigateway/latest/developerguide/api-gateway-overview-developer-experience.htm>

NEW QUESTION 54

- (Exam Topic 1)

A health insurance company stores personally identifiable information (PII) in an Amazon S3 bucket. The company uses server-side encryption with S3 managed encryption keys (SSE-S3) to encrypt the objects. According to a new requirement, all current and future objects in the S3 bucket must be encrypted by keys that the company's security team manages. The S3 bucket does not have versioning enabled. Which solution will meet these requirements?

- A. In the S3 bucket properties, change the default encryption to SSE-S3 with a customer managed key
- B. Use the AWS CLI to re-upload all objects in the S3 bucket
- C. Set an S3 bucket policy to deny unencrypted PutObject requests.
- D. In the S3 bucket properties, change the default encryption to server-side encryption with AWS KMS managed encryption keys (SSE-KMS). Set an S3 bucket policy to deny unencrypted PutObject request
- E. Use the AWS CLI to re-upload all objects in the S3 bucket.
- F. In the S3 bucket properties, change the default encryption to server-side encryption with AWS KMS managed encryption keys (SSE-KMS). Set an S3 bucket policy to automatically encrypt objects on GetObject and PutObject requests.
- G. In the S3 bucket properties, change the default encryption to AES-256 with a customer managed key. Attach a policy to deny unencrypted PutObject requests to any entities that access the S3 bucket
- H. Use the AWS CLI to re-upload all objects in the S3 bucket.

Answer: D

Explanation:

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/ServerSideEncryptionCustomerKeys.html> Clearly says we need following header for SSE-C x-amz-server-side-encryption-customer-algorithm Use this header to specify the encryption algorithm. The header value must be AES256.

NEW QUESTION 55

- (Exam Topic 1)

A video processing company has an application that downloads images from an Amazon S3 bucket, processes the images, stores a transformed image in a second S3 bucket, and updates metadata about the image in an Amazon DynamoDB table. The application is written in Node.js and runs by using an AWS Lambda function. The Lambda function is invoked when a new image is uploaded to Amazon S3.

The application ran without incident for a while. However, the size of the images has grown significantly. The Lambda function is now failing frequently with timeout errors. The function timeout is set to its maximum value. A solutions architect needs to refactor the application's architecture to prevent invocation failures. The company does not want to manage the underlying infrastructure.

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

- A. Modify the application deployment by building a Docker image that contains the application code. Publish the image to Amazon Elastic Container Registry (Amazon ECR).
- B. Create a new Amazon Elastic Container Service (Amazon ECS) task definition with a compatibility type of AWS Fargate
- C. Configure the task definition to use the new image in Amazon Elastic Container Registry (Amazon ECR). Adjust the Lambda function to invoke an ECS task by using the ECS task definition when a new file arrives in Amazon S3.
- D. Create an AWS Step Functions state machine with a Parallel state to invoke the Lambda function. Increase the provisioned concurrency of the Lambda function.
- E. Create a new Amazon Elastic Container Service (Amazon ECS) task definition with a compatibility type of Amazon EC2. Configure the task definition to use the new image in Amazon Elastic Container Registry (Amazon ECR). Adjust the Lambda function to invoke an ECS task by using the ECS task definition when a new file arrives in Amazon S3.
- F. Modify the application to store images on Amazon Elastic File System (Amazon EFS) and to store metadata on an Amazon RDS DB instance
- G. Adjust the Lambda function to mount the EFS file share.

Answer: AB

Explanation:

A. Modify the application deployment by building a Docker image that contains the application code. Publish the image to Amazon Elastic Container Registry (Amazon ECR). - This step is necessary to package the application code in a container and make it available for running on ECS. B. Create a new Amazon Elastic Container Service (Amazon ECS) task definition with a compatibility type of AWS Fargate. Configure the task definition to use the new image in Amazon Elastic Container Registry (Amazon ECR). Adjust the Lambda function to invoke an ECS task by using the ECS task definition when a new file arrives in Amazon S3.

NEW QUESTION 60

- (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 myservice.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 myservice.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 64

- (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 VP
- 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 67

- (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 loa
- 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 cluste
- 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 loa
- G. Temporarily scale out database read replicas during peaks.
- H. Purchase EC2 Instance Savings Plans for the predicted base load of the EKS cluste
- 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 loa
- K. Temporarily scale up the DB instance manually during peaks.
- L. Purchase Compute Savings Plans for the predicted base load of the EKS cluste
- 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 loa
- 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 72

- (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 73

- (Exam Topic 2)

A solutions architect is designing a solution to process events. The solution must have the ability to scale in and out based on the number of events that the solution receives. If a processing error occurs, the event must move into a separate queue for review. Which solution will meet these requirements?

- A. Send event details to an Amazon Simple Notification Service (Amazon SNS) topic
- B. Configure an AWS Lambda function as a subscriber to the SNS topic to process the event
- C. Add an on-failure destination to the function
- D. Set an Amazon Simple Queue Service (Amazon SQS) queue as the target.
- E. Publish events to an Amazon Simple Queue Service (Amazon SQS) queue
- F. Create an Amazon EC2 Auto Scaling group
- G. Configure the Auto Scaling group to scale in and out based on the ApproximateAgeOfOldestMessage metric of the queue
- H. Configure the application to write failed messages to a dead-letter queue.
- I. Write events to an Amazon DynamoDB table
- J. Configure a DynamoDB stream for the table
- K. Configure the stream to invoke an AWS Lambda function
- L. Configure the Lambda function to process the events.
- M. Publish events to an Amazon EventBridge event bus
- N. Create and run an application on an Amazon EC2 instance with an Auto Scaling group that is behind an Application Load Balancer (ALB). Set the ALB as the event bus target
- O. Configure the event bus to retry event
- P. Write messages to a dead-letter queue if the application cannot process the messages.

Answer: A

Explanation:

Amazon Simple Notification Service (Amazon SNS) is a fully managed pub/sub messaging service that enables users to send messages to multiple subscribers¹. Users can send event details to an Amazon SNS topic and configure an AWS Lambda function as a subscriber to the SNS topic to process the events. Lambda is a serverless compute service that runs code in response to events and automatically manages the underlying compute resources². Users can add an on-failure destination to the function and set an Amazon Simple Queue

Service (Amazon SQS) queue as the target. Amazon SQS is a fully managed message queuing service that enables users to decouple and scale microservices, distributed systems, and serverless applications³. This way, if a processing error occurs, the event will move into the separate queue for review.

Option B is incorrect because publishing events to an Amazon SQS queue and creating an Amazon EC2 Auto Scaling group will not have the ability to scale in and out based on the number of events that the solution receives. Amazon EC2 is a web service that provides secure, resizable compute capacity in the cloud. Auto Scaling is a feature that helps users maintain application availability and allows them to scale their EC2 capacity up or down automatically according to conditions they define. However, for this use case, using SQS and EC2 will not take advantage of the serverless capabilities of Lambda and SNS.

Option C is incorrect because writing events to an Amazon DynamoDB table and configuring a DynamoDB stream for the table will not have the ability to move events into a separate queue for review if a processing error occurs. Amazon DynamoDB is a fully managed key-value and document database that delivers single-digit millisecond performance at any scale. DynamoDB Streams is a feature that captures data modification events in DynamoDB tables. Users can configure the stream to invoke a Lambda function, but they cannot configure an on-failure destination for the function.

Option D is incorrect because publishing events to an Amazon EventBridge event bus and setting an Application Load Balancer (ALB) as the event bus target will not have the ability to move events into a separate queue for review if a processing error occurs. Amazon EventBridge is a serverless event bus service that makes it easy to connect applications with data from a variety of sources. An ALB is a load balancer that distributes incoming application traffic across multiple targets, such as EC2 instances, containers, IP addresses, Lambda functions, and virtual appliances. Users can configure EventBridge to retry events, but they cannot configure an on-failure destination for the ALB.

NEW QUESTION 74

- (Exam Topic 2)

A company runs a customer service center that accepts calls and automatically sends all customers a managed, interactive, two-way experience survey by text message.

The applications that support the customer service center run on machines that the company hosts in an on-premises data center. The hardware that the company uses is old, and the company is experiencing downtime with the system. The company wants to migrate the system to AWS to improve reliability.

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

- A. Use Amazon Connect to replace the old call center hardware
- B. Use Amazon Pinpoint to send text message surveys to customers.
- C. Use Amazon Connect to replace the old call center hardware

- D. Use Amazon Simple Notification Service (Amazon SNS) to send text message surveys to customers.
- E. Migrate the call center software to Amazon EC2 instances that are in an Auto Scaling group.
- F. Use the EC2 instances to send text message surveys to customers.
- G. Use Amazon Pinpoint to replace the old call center hardware and to send text message surveys to customers.

Answer: A

Explanation:

Amazon Connect is a cloud-based contact center service that allows you to set up a virtual call center for your business. It provides an easy-to-use interface for managing customer interactions through voice and chat. Amazon Connect integrates with other AWS services, such as Amazon S3 and Amazon Kinesis, to help you collect, store, and analyze customer data for insights into customer behavior and trends. On the other hand, Amazon Pinpoint is a marketing automation and analytics service that allows you to engage with your customers across different channels, such as email, SMS, push notifications, and voice. It helps you create personalized campaigns based on user behavior and enables you to track user engagement and retention. While both services allow you to communicate with your customers, they serve different purposes. Amazon Connect is focused on customer support and service, while Amazon Pinpoint is focused on marketing and engagement.

NEW QUESTION 78

- (Exam Topic 2)

A company uses AWS Organizations for a multi-account setup in the AWS Cloud. The company's finance team has a data processing application that uses AWS Lambda and Amazon DynamoDB. The company's marketing team wants to access the data that is stored in the DynamoDB table. The DynamoDB table contains confidential data. The marketing team can have access to only specific attributes of data in the DynamoDB table. The finance team and the marketing team have separate AWS accounts.

What should a solutions architect do to provide the marketing team with the appropriate access to the DynamoDB table?

- A. Create an SCP to grant the marketing team's AWS account access to the specific attributes of the DynamoDB table.
- B. Attach the SCP to the OU of the finance team.
- C. Create an IAM role in the finance team's account by using IAM policy conditions for specific DynamoDB attributes (fine-grained access control). Establish trust with the marketing team's account.
- D. In the marketing team's account, create an IAM role that has permissions to assume the IAM role in the finance team's account.
- E. Create a resource-based IAM policy that includes conditions for specific DynamoDB attributes (fine-grained access control). Attach the policy to the DynamoDB table.
- F. In the marketing team's account, create an IAM role that has permissions to access the DynamoDB table in the finance team's account.
- G. Create an IAM role in the finance team's account to access the DynamoDB table.
- H. Use an IAM permissions boundary to limit the access to the specific attribute.
- I. In the marketing team's account, create an IAM role that has permissions to assume the IAM role in the finance team's account.

Answer: C

Explanation:

The company should create a resource-based IAM policy that includes conditions for specific DynamoDB attributes (fine-grained access control). The company should attach the policy to the DynamoDB table. In the marketing team's account, the company should create an IAM role that has permissions to access the DynamoDB table in the finance team's account. This solution will meet the requirements because a resource-based IAM policy is a policy that you attach to an AWS resource (such as a DynamoDB table) to control who can access that resource and what actions they can perform on it. You can use IAM policy conditions to specify fine-grained access control for DynamoDB items and attributes. For example, you can allow or deny access to specific attributes of all items in a table by matching on attribute names¹. By creating a resource-based policy that allows access to only specific attributes of the DynamoDB table and attaching it to the table, the company can restrict access to confidential data. By creating an IAM role in the marketing team's account that has permissions to access the DynamoDB table in the finance team's account, the company can enable cross-account access. The other options are not correct because:

- Creating an SCP to grant the marketing team's AWS account access to the specific attributes of the DynamoDB table would not work because SCPs are policies that you can use with AWS Organizations to manage permissions in your organization's accounts. SCPs do not grant permissions; instead, they specify the maximum permissions that identities in an account can have². SCPs cannot be used to specify fine-grained access control for DynamoDB items and attributes.
- Creating an IAM role in the finance team's account by using IAM policy conditions for specific DynamoDB attributes and establishing trust with the marketing team's account would not work because IAM roles are identities that you can create in your account that have specific permissions. You can use an IAM role to delegate access to users, applications, or services that don't normally have access to your AWS resources³. However, creating an IAM role in the finance team's account would not restrict access to specific attributes of the DynamoDB table; it would only allow cross-account access. The company would still need a resource-based policy attached to the table to enforce fine-grained access control.
- Creating an IAM role in the finance team's account to access the DynamoDB table and using an IAM permissions boundary to limit the access to the specific attributes would not work because IAM permissions boundaries are policies that you use to delegate permissions management to other users. You can use permissions boundaries to limit the maximum permissions that an identity-based policy can grant to an IAM entity (user or role)⁴. Permissions boundaries cannot be used to specify fine-grained access control for DynamoDB items and attributes.

References:

- <https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/specifying-conditions.html>
- https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps.html
- https://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles.html
- https://docs.aws.amazon.com/IAM/latest/UserGuide/access_policies_boundaries.html

NEW QUESTION 80

- (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 85

- (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 88

- (Exam Topic 2)

A company has an application that runs as a ReplicaSet of multiple pods in an Amazon Elastic Kubernetes Service (Amazon EKS) cluster. The EKS cluster has nodes in multiple Availability Zones. The application generates many small files that must be accessible across all running instances of the application. The company needs to back up the files and retain the backups for 1 year.

Which solution will meet these requirements while providing the FASTEST storage performance?

- A. Create an Amazon Elastic File System (Amazon EFS) file system and a mount target for each subnet that contains nodes in the EKS cluster
- B. Configure the ReplicaSet to mount the file system
- C. Direct the application to store files in the file system
- D. Configure AWS Backup to back up and retain copies of the data for 1 year.
- E. Create an Amazon Elastic Block Store (Amazon EBS) volume
- F. Enable the EBS Multi-Attach feature. Configure the ReplicaSet to mount the EBS volume
- G. Direct the application to store files in the EBS volume
- H. Configure AWS Backup to back up and retain copies of the data for 1 year.
- I. Create an Amazon S3 bucket
- J. Configure the ReplicaSet to mount the S3 bucket
- K. Direct the application to store files in the S3 bucket
- L. Configure S3 Versioning to retain copies of the data
- M. Configure an S3 Lifecycle policy to delete objects after 1 year.
- N. Configure the ReplicaSet to use the storage available on each of the running application pods to store the files locally
- O. Use a third-party tool to back up the EKS cluster for 1 year.

Answer: A

Explanation:

In the past, EBS can be attached only to one EC2 instance but not anymore but there are limitations like - it works only on io1/io2 instance types and many others as described here. <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-volumes-multi.html> EFS has shareable storage

In terms of performance, Amazon EFS is optimized for workloads that require high levels of aggregate throughput and IOPS, whereas EBS is optimized for low-latency, random access I/O operations. Amazon EFS is designed to scale throughput and capacity automatically as your storage needs grow, while EBS volumes can be resized on demand.

NEW QUESTION 91

- (Exam Topic 2)

A company is migrating a document processing workload to AWS. The company has updated many applications to natively use the Amazon S3 API to store, retrieve, and modify documents that a processing server generates at a rate of approximately 5 documents every second. After the document processing is

finished, customers can download the documents directly from Amazon S3.

During the migration, the company discovered that it could not immediately update the processing server that generates many documents to support the S3 API. The server runs on Linux and requires fast local access to the files that the server generates and modifies. When the server finishes processing, the files must be available to the public for download within 30 minutes.

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

- A. Migrate the application to an AWS Lambda function
- B. Use the AWS SDK for Java to generate, modify, and access the files that the company stores directly in Amazon S3.
- C. Set up an Amazon S3 File Gateway and configure a file share that is linked to the document store. Mount the file share on an Amazon EC2 instance by using NFS
- D. When changes occur in Amazon S3, initiate a RefreshCache API call to update the S3 File Gateway.
- E. Configure Amazon FSx for Lustre with an import and export policy
- F. Link the new file system to an S3 bucket
- G. Install the Lustre client and mount the document store to an Amazon EC2 instance by using NFS.
- H. Configure AWS DataSync to connect to an Amazon EC2 instance
- I. Configure a task to synchronize the generated files to and from Amazon S3.

Answer: C

Explanation:

Amazon FSx for Lustre is a fully managed service that provides cost-effective, high-performance, scalable storage for compute workloads. Powered by Lustre, the world's most popular high-performance file system, FSx for Lustre offers shared storage with sub-ms latencies, up to terabytes per second of throughput, and millions of IOPS. FSx for Lustre file systems can also be linked to Amazon Simple Storage Service (S3) buckets, allowing you to access and process data concurrently from both a high-performance file system and from the S3 API.

NEW QUESTION 96

- (Exam Topic 2)

A company operates an on-premises software-as-a-service (SaaS) solution that ingests several files daily. The company provides multiple public SFTP endpoints to its customers to facilitate the file transfers. The customers add the SFTP endpoint IP addresses to their firewall allow list for outbound traffic. Changes to the SFTP endpoint IP addresses are not permitted.

The company wants to migrate the SaaS solution to AWS and decrease the operational overhead of the file transfer service.

Which solution meets these requirements?

- A. Register the customer-owned block of IP addresses in the company's AWS account
- B. Create Elastic IP addresses from the address pool and assign them to an Amazon S3 endpoint
- C. Use Amazon S3 Transfer Acceleration to store the files in Amazon S3.
- D. Add a subnet containing the customer-owned block of IP addresses to a VPC. Create Elastic IP addresses from the address pool and assign them to an Application Load Balancer (ALB). Launch EC2 instances hosting SFTP services in an Auto Scaling group behind the ALB.
- E. Store the files in attached Amazon Elastic Block Store (Amazon EBS) volumes.
- F. Register the customer-owned block of IP addresses with Amazon Route 53. Create alias records in Route 53 that point to a Network Load Balancer (NLB). Launch EC2 instances hosting SFTP services in an Auto Scaling group behind the NLB.
- G. Store the files in Amazon S3.
- H. Register the customer-owned block of IP addresses in the company's AWS account
- I. Create Elastic IP addresses from the address pool and assign them to an Amazon S3 VPC endpoint
- J. Enable SFTP support on the S3 bucket.

Answer: A

Explanation:

Bring your own IP addresses (BYOIP) You can bring part or all of your publicly routable IPv4 or IPv6 address range from your on-premises network to your AWS account. You continue to own the address range, but AWS advertises it on the internet by default. After you bring the address range to AWS, it appears in your AWS account as an address pool. <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-byoip.html> Amazon S3 Transfer Acceleration enables you to easily move your file transfer workloads that use the Secure Shell File Transfer Protocol (SFTP) to AWS without needing to modify your applications or manage any SFTP servers. <https://aws.amazon.com/about-aws/whats-new/2018/11/aws-transfer-for-sftp-fully-managed-sftp-for-s3/>

NEW QUESTION 97

- (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 101

- (Exam Topic 2)

A company is creating a centralized logging service running on Amazon EC2 that will receive and analyze logs from hundreds of AWS accounts. AWS PrivateLink is being used to provide connectivity between the client services and the logging service.

In each AWS account with a client, an interface endpoint has been created for the logging service and is available. The logging service running on EC2 instances with a Network Load Balancer (NLB) are deployed in different subnets. The clients are unable to submit logs using the VPC endpoint.

Which combination of steps should a solutions architect take to resolve this issue? (Select TWO.)

- A. Check that the NACL is attached to the logging service subnet to allow communications to and from the NLB subnet
- B. Check that the NACL is attached to the NLB subnet to allow communications to and from the logging service subnets running on EC2 instances.
- C. Check that the NACL is attached to the logging service subnets to allow communications to and from the interface endpoint subnet
- D. Check that the NACL is attached to the interface endpoint subnet to allow communications to and from the logging service subnets running on EC2 instances.
- E. Check the security group for the logging service running on the EC2 instances to ensure it allows Ingress from the NLB subnets.
- F. Check the security group for the logging service running on EC2 instances to ensure it allows ingress from the clients.
- G. Check the security group for the NLB to ensure it allows ingress from the interface endpoint subnets.

Answer: AC

NEW QUESTION 102

- (Exam Topic 2)

A company uses a Grafana data visualization solution that runs on a single Amazon EC2 instance to monitor the health of the company's AWS workloads. The company has invested time and effort to create dashboards that the company wants to preserve. The dashboards need to be highly available and cannot be down for longer than 10 minutes. The company needs to minimize ongoing maintenance.

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

- A. Migrate to Amazon CloudWatch dashboard
- B. Recreate the dashboards to match the existing Grafana dashboard
- C. Use automatic dashboards where possible.
- D. Create an Amazon Managed Grafana workspace
- E. Configure a new Amazon CloudWatch data source. Export dashboards from the existing Grafana instance
- F. Import the dashboards into the new workspace.
- G. Create an AMI that has Grafana pre-installed
- H. Store the existing dashboards in Amazon Elastic File System (Amazon EFS). Create an Auto Scaling group that uses the new AMI
- I. Set the Auto Scaling group's minimum, desired, and maximum number of instances to one
- J. Create an Application Load Balancer that serves at least two Availability Zones.
- K. Configure AWS Backup to back up the EC2 instance that runs Grafana once each hour
- L. Restore the EC2 instance from the most recent snapshot in an alternate Availability Zone when required.

Answer: C

Explanation:

By creating an AMI that has Grafana pre-installed and storing the existing dashboards in Amazon Elastic File System (Amazon EFS) it allows for faster and more efficient scaling, and by creating an Auto Scaling group that uses the new AMI and setting the Auto Scaling group's minimum, desired, and maximum number of instances to one and creating an Application Load Balancer that serves at least two Availability Zones, it ensures high availability and minimized downtime.

NEW QUESTION 107

- (Exam Topic 2)

A company operates a proxy server on a fleet of Amazon EC2 instances. Partners in different countries use the proxy server to test the company's functionality. The EC2 instances are running in a VPC, and the instances have access to the internet.

The company's security policy requires that partners can access resources only from domains that the company owns.

Which solution will meet these requirements?

- A. Create an Amazon Route 53 Resolver DNS Firewall domain list that contains the allowed domains. Configure a DNS Firewall rule group with a rule that has a high numeric value that blocks all requests
- B. Configure a rule that has a low numeric value that allows requests for domains in the allowed list
- C. Associate the rule group with the VPC.
- D. Create an Amazon Route 53 Resolver DNS Firewall domain list that contains the allowed domains. Configure a Route 53 outbound endpoint
- E. Associate the outbound endpoint with the VPC
- F. Associate the domain list with the outbound endpoint.
- G. Create an Amazon Route 53 traffic flow policy to match the allowed domain
- H. Configure the traffic flow policy to forward requests that match to the Route 53 Resolver
- I. Associate the traffic flow policy with the VPC.
- J. Create an Amazon Route 53 outbound endpoint
- K. Associate the outbound endpoint with the VPC
- L. Configure a Route 53 traffic flow policy to forward requests for allowed domains to the outbound endpoint
- M. Associate the traffic flow policy with the VPC.

Answer: A

Explanation:

The company should create an Amazon Route 53 Resolver DNS Firewall domain list that contains the allowed domains. The company should configure a DNS Firewall rule group with a rule that has a high numeric value that blocks all requests. The company should configure a rule that has a low numeric value that allows requests for domains in the allowed list. The company should associate the rule group with the VPC. This solution will meet the requirements because Amazon Route 53 Resolver DNS Firewall is a feature that enables you to filter and regulate outbound DNS traffic for your VPC. You can create reusable collections of filtering rules in DNS Firewall rule groups and associate them with your VPCs. You can specify lists of domain names to allow or block, and you can customize the responses for the DNS queries that you block. By creating a domain list with the allowed domains and a rule group with rules to allow or block requests based on the domain list, the company can enforce its security policy and control access to sites.

The other options are not correct because:

- Configuring a Route 53 outbound endpoint and associating it with the VPC would not help with filtering outbound DNS traffic. A Route 53 outbound endpoint is a resource that enables you to forward DNS queries from your VPC to your network over AWS Direct Connect or VPN connections². It does not provide any filtering capabilities.
- Creating a Route 53 traffic flow policy to match the allowed domains would not help with filtering outbound DNS traffic. A Route 53 traffic flow policy is a resource that enables you to route traffic based on multiple criteria, such as endpoint health, geographic location, and latency³. It does not provide any filtering capabilities.
- Creating a Gateway Load Balancer (GWLB) would not help with filtering outbound DNS traffic. A GWLB is a service that enables you to deploy, scale, and manage third-party virtual appliances such as firewalls, intrusion detection and prevention systems, and deep packet inspection systems in the cloud⁴. It does not provide any filtering capabilities.

References:

- <https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/resolver-dns-firewall.html>
- <https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/resolver-outbound-endpoints.html>
- <https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/traffic-flow.html>
- <https://docs.aws.amazon.com/elasticloadbalancing/latest/gateway/introduction.html>

NEW QUESTION 110

- (Exam Topic 2)

A company has migrated a legacy application to the AWS Cloud. The application runs on three Amazon EC2 instances that are spread across three Availability Zones. One EC2 instance is in each Availability Zone. The EC2 instances are running in three private subnets of the VPC and are set up as targets for an Application Load Balancer (ALB) that is associated with three public subnets.

The application needs to communicate with on-premises systems. Only traffic from IP addresses in the company's IP address range are allowed to access the on-premises systems. The company's security team is bringing only one IP address from its internal IP address range to the cloud. The company has added this IP address to the allow list for the company firewall. The company also has created an Elastic IP address for this IP address.

A solutions architect needs to create a solution that gives the application the ability to communicate with the on-premises systems. The solution also must be able to mitigate failures automatically.

Which solution will meet these requirements?

- A. Deploy three NAT gateways, one in each public subne
- B. Assign the Elastic IP address to the NAT gateway
- C. Turn on health checks for the NAT gateway
- D. If a NAT gateway fails a health check, recreate the NAT gateway and assign the Elastic IP address to the new NAT gateway.
- E. Replace the ALB with a Network Load Balancer (NLB). Assign the Elastic IP address to the NLB Turn on health checks for the NL
- F. In the case of a failed health check, redeploy the NLB in different subnets.
- G. Deploy a single NAT gateway in a public subne
- H. Assign the Elastic IP address to the NAT gateway. Use Amazon CloudWatch with a custom metric to monitor the NAT gatewa
- I. If the NAT gateway is unhealthy, invoke an AWS Lambda function to create a new NAT gateway in a different subne
- J. Assign the Elastic IP address to the new NAT gateway.
- K. Assign the Elastic IP address to the AL
- L. Create an Amazon Route 53 simple record with the Elastic IP address as the valu
- M. Create a Route 53 health chec
- N. In the case of a failed health check, recreate the ALB in different subnets.

Answer: C

Explanation:

to connect out from the private subnet you need an NAT gateway and since only one Elastic IP whitelisted on firewall its one NATGateway at time and if AZ failure happens Lambda creates a new NATGATEWAY in a different AZ using the Same Elastic IP ,dont be tempted to select D since application that needs to connect is on a private subnet whose outbound connections use the NATGateway Elastic IP

NEW QUESTION 115

- (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 analysi
- 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 119

- (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 122

- (Exam Topic 2)

A solutions architect at a large company needs to set up network security for outbound traffic to the internet from all AWS accounts within an organization in AWS Organizations. The organization has more than 100 AWS accounts, and the accounts route to each other by using a centralized AWS Transit Gateway. Each account has both an internet gateway and a NAT gateway for outbound traffic to the internet. The company deploys resources only into a single AWS Region. The company needs the ability to add centrally managed rule-based filtering on all outbound traffic to the internet for all AWS accounts in the organization. The peak load of outbound traffic will not exceed 25 Gbps in each Availability Zone.

Which solution meets these requirements?

- A. Create a new VPC for outbound traffic to the internet
- B. Connect the existing transit gateway to the new VPC
- C. Configure a new NAT gateway
- D. Create an Auto Scaling group of Amazon EC2 instances that run an open-source internet proxy for rule-based filtering across all Availability Zones in the Region
- E. Modify all default routes to point to the proxy's Auto Scaling group.
- F. Create a new VPC for outbound traffic to the internet
- G. Connect the existing transit gateway to the new VPC
- H. Configure a new NAT gateway
- I. Use an AWS Network Firewall firewall for rule-based filtering
- J. Create Network Firewall endpoints in each Availability Zone
- K. Modify all default routes to point to the Network Firewall endpoints.
- L. Create an AWS Network Firewall firewall for rule-based filtering in each AWS account
- M. Modify all default routes to point to the Network Firewall firewalls in each account.
- N. In each AWS account, create an Auto Scaling group of network-optimized Amazon EC2 instances that run an open-source internet proxy for rule-based filtering
- O. Modify all default routes to point to the proxy's Auto Scaling group.

Answer: B

Explanation:

<https://aws.amazon.com/blogs/networking-and-content-delivery/deployment-models-for-aws-network-firewall/>

NEW QUESTION 123

- (Exam Topic 2)

A company is building a call center by using Amazon Connect. The company's operations team is defining a disaster recovery (DR) strategy across AWS Regions. The contact center has dozens of contact flows, hundreds of users, and dozens of claimed phone numbers.

Which solution will provide DR with the LOWEST RTO?

- A. Create an AWS Lambda function to check the availability of the Amazon Connect instance and to send a notification to the operations team in case of unavailability
- B. Create an Amazon EventBridge rule to invoke the Lambda function every 5 minutes
- C. After notification, instruct the operations team to use the AWS Management Console to provision a new Amazon Connect instance in a second Region
- D. Deploy the contact flows, users, and claimed phone numbers by using an AWS CloudFormation template.
- E. Provision a new Amazon Connect instance with all existing users in a second Region
- F. Create an AWS Lambda function to check the availability of the Amazon Connect instance
- G. Create an Amazon EventBridge rule to invoke the Lambda function every 5 minutes
- H. In the event of an issue, configure the Lambda function to deploy an AWS CloudFormation template that provisions contact flows and claimed numbers in the second Region.
- I. Provision a new Amazon Connect instance with all existing contact flows and claimed phone numbers in a second Region
- J. Create an Amazon Route 53 health check for the URL of the Amazon Connect instance
- K. Create an Amazon CloudWatch alarm for failed health check
- L. Create an AWS Lambda function to deploy an AWS CloudFormation template that provisions all users
- M. Configure the alarm to invoke the Lambda function.
- N. Provision a new Amazon Connect instance with all existing users and contact flows in a second Region. Create an Amazon Route 53 health check for the URL of the Amazon Connect instance
- O. Create an Amazon CloudWatch alarm for failed health check
- P. Create an AWS Lambda function to deploy an AWS CloudFormation template that provisions claimed phone numbers
- Q. Configure the alarm to invoke the Lambda function.

Answer: D

Explanation:

Option D provisions a new Amazon Connect instance with all existing users and contact flows in a second Region. It also sets up an Amazon Route 53 health check for the URL of the Amazon Connect instance, an Amazon CloudWatch alarm for failed health checks, and an AWS Lambda function to deploy an AWS CloudFormation template that provisions claimed phone numbers. This option allows for the fastest recovery time because all the necessary components are already provisioned and ready to go in the second Region. In the event of a disaster, the failed health check will trigger the AWS Lambda function to deploy the CloudFormation template to provision the claimed phone numbers, which is the only missing component.

NEW QUESTION 125

- (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 volumes¹. Configuring the backup task to copy the backups to a vault in the secondary Region will enable maintaining backups in a separate Region¹. In the event of a failure, launching the CloudFormation template will enable deploying the network configuration in the secondary Region². Restoring the instance volumes and configurations from the backup vault will enable recovering the EC2 instances and their data¹. Transferring usage to the secondary Region will enable resuming operations².

NEW QUESTION 126

- (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 128

- (Exam Topic 2)

A company wants to use AWS for disaster recovery for an on-premises application. The company has hundreds of Windows-based servers that run the application. All the servers mount a common share.

The company has an RTO of 15 minutes and an RPO of 5 minutes. The solution must support native failover and fallback capabilities.

Which solution will meet these requirements MOST cost-effectively?

- A. Create an AWS Storage Gateway File Gateway
- B. Schedule daily Windows server backup
- C. Save the data to Amazon S3. During a disaster, recover the on-premises servers from the backup
- D. During failback
- E. run the on-premises servers on Amazon EC2 instances.
- F. Create a set of AWS CloudFormation templates to create infrastructure
- G. Replicate all data to Amazon Elastic File System (Amazon EFS) by using AWS DataSync
- H. During a disaster, use AWS CodePipeline to deploy the templates to restore the on-premises server
- I. Fail back the data by using DataSync.
- J. Create an AWS Cloud Development Kit (AWS CDK) pipeline to stand up a multi-site active-active environment on AWS
- K. Replicate data into Amazon S3 by using the s3 sync command
- L. During a disaster, swap DNS endpoints to point to AWS
- M. Fail back the data by using the s3 sync command.
- N. Use AWS Elastic Disaster Recovery to replicate the on-premises server
- O. Replicate data to an Amazon FSx for Windows File Server file system by using AWS DataSync

- P. Mount the file system to AWS server
- Q. During a disaster, fail over the on-premises servers to AW
- R. Fail back to new or existing servers by using Elastic Disaster Recovery.

Answer: D

NEW QUESTION 130

- (Exam Topic 2)

A company has IoT sensors that monitor traffic patterns throughout a large city. The company wants to read and collect data from the sensors and perform aggregations on the data.

A solutions architect designs a solution in which the IoT devices are streaming to Amazon Kinesis Data Streams. Several applications are reading from the stream. However, several consumers are experiencing throttling and are periodically and are periodically encountering a RealProvisioned Throughput Exceeded error. Which actions should the solution architect take to resolve this issue? (Select THREE.)

- A. Reshard the stream to increase the number of shards s in the stream.
- B. Use the Kinesis Producer Library (KPL). Adjust the polling frequency.
- C. Use consumers with the enhanced fan-out feature.
- D. Reshard the stream to reduce the number of shards in the stream.
- E. Use an error retry and exponential backoff mechanism in the consumer logic.
- F. Configure the stream to use dynamic partitioning.

Answer: ACE

Explanation:

<https://repost.aws/knowledge-center/kinesis-readprovisionedthroughputexceeded> Follow Data Streams best practices

To mitigate ReadProvisionedThroughputExceeded exceptions, apply these best practices:

- Reshard your stream to increase the number of shards in the stream.
- Use consumers with enhanced fan-out. For more information about enhanced fan-out, see Developing custom consumers with dedicated throughput (enhanced fan-out).
- Use an error retry and exponential backoff mechanism in the consumer logic if ReadProvisionedThroughputExceeded exceptions are encountered. For consumer applications that use an AWS SDK, the requests are retried by default.

NEW QUESTION 135

- (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 savings². 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 137

- (Exam Topic 2)

A company needs to optimize the cost of backups for Amazon Elastic File System (Amazon EFS). A solutions architect has already configured a backup plan in AWS Backup for the EFS backups. The backup plan contains a rule with a lifecycle configuration to transition EFS backups to cold storage after 7 days and to keep the backups for an additional 90 days.

After 1 month, the company reviews its EFS storage costs and notices an increase in the EFS backup costs. The EFS backup cold storage produces almost double the cost of the EFS warm backup storage.

What should the solutions architect do to optimize the cost?

- A. Modify the backup rule's lifecycle configuration to move the EFS backups to cold storage after 1 day. Set the backup retention period to 30 days.
- B. Modify the backup rule's lifecycle configuration to move the EFS backups to cold storage after 8 days. Set the backup retention period to 30 days.
- C. Modify the backup rule's lifecycle configuration to move the EFS backups to cold storage after 1 day. Set the backup retention period to 90 days.
- D. Modify the backup rule's lifecycle configuration to move the EFS backups to cold storage after 8 days. Set the backup retention period to 98 days.

Answer: A

Explanation:

The cost of EFS backup cold storage is \$0.01 per GB-month, whereas the cost of EFS backup warm storage is \$0.05 per GB-month¹. Therefore, moving the backups to cold storage as soon as possible will reduce the storage cost. However, cold storage backups must be retained for a minimum of 90 days², otherwise they incur a pro-rated charge equal to the storage charge for the remaining days¹. Therefore, setting the backup retention period to 30 days will incur a penalty of 60 days of cold storage cost for each backup deleted. This penalty will still be lower than keeping the backups in warm storage for 7 days and then in cold storage for 83 days, which is the current configuration. Therefore, option A is the most cost-effective solution.

NEW QUESTION 141

- (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
- B. In Account
- C. set the S3 bucket policy to the following:

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

- D. In Account
- E. 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
- G. 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 146

- (Exam Topic 2)

An external audit of a company's serverless application reveals IAM policies that grant too many permissions. These policies are attached to the company's AWS Lambda execution roles. Hundreds of the company's Lambda functions have broad access permissions, such as full access to Amazon S3 buckets and Amazon DynamoDB tables. The company wants each function to have only the minimum permissions that the function needs to complete its task.

A solutions architect must determine which permissions each Lambda function needs.

What should the solutions architect do to meet this requirement with the LEAST amount of effort?

- A. Set up Amazon CodeGuru to profile the Lambda functions and search for AWS API call
- B. Create an inventory of the required API calls and resources for each Lambda function
- C. Create new IAM access policies for each Lambda function
- D. Review the new policies to ensure that they meet the company's business requirements.
- E. Turn on AWS CloudTrail logging for the AWS account
- F. Use AWS Identity and Access Management Access Analyzer to generate IAM access policies based on the activity recorded in the CloudTrail logs
- G. Review the generated policies to ensure that they meet the company's business requirements.
- H. Turn on AWS CloudTrail logging for the AWS account
- I. Create a script to parse the CloudTrail log, search for AWS API calls by Lambda execution role, and create a summary report
- J. Review the report
- K. Create IAM access policies that provide more restrictive permissions for each Lambda function.
- L. Turn on AWS CloudTrail logging for the AWS account
- M. Export the CloudTrail logs to Amazon S3. Use Amazon EMR to process the CloudTrail logs in Amazon S3 and produce a report of API calls and resources used by each execution role
- N. Create a new IAM access policy for each role
- O. Export the generated roles to an S3 bucket
- P. Review the generated policies to ensure that they meet the company's business requirements.

Answer: B

Explanation:

IAM Access Analyzer helps you identify the resources in your organization and accounts, such as Amazon S3 buckets or IAM roles, shared with an external entity. This lets you identify unintended access to your resources and data, which is a security risk. IAM Access Analyzer identifies resources shared with external principals by using logic-based reasoning to analyze the resource-based policies in your AWS environment.

<https://docs.aws.amazon.com/IAM/latest/UserGuide/what-is-access-analyzer.html>

NEW QUESTION 148

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questions and Answers in PDF Format

SAP-C02 Practice Exam Features:

- * SAP-C02 Questions and Answers Updated Frequently
- * SAP-C02 Practice Questions Verified by Expert Senior Certified Staff
- * SAP-C02 Most Realistic Questions that Guarantee you a Pass on Your First Try
- * SAP-C02 Practice Test Questions in Multiple Choice Formats and Updates for 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The SAP-C02 Practice Test Here](#)