

Amazon-Web-Services

Exam Questions SAA-C03

AWS Certified Solutions Architect - Associate (SAA-C03)



NEW QUESTION 1

- (Topic 1)

A company observes an increase in Amazon EC2 costs in its most recent bill. The billing team notices unwanted vertical scaling of instance types for a couple of EC2 instances. A solutions architect needs to create a graph comparing the last 2 months of EC2 costs and perform an in-depth analysis to identify the root cause of the vertical scaling.

How should the solutions architect generate the information with the LEAST operational overhead?

- A. Use AWS Budgets to create a budget report and compare EC2 costs based on instance types.
- B. Use Cost Explorer's granular filtering feature to perform an in-depth analysis of EC2 costs based on instance types.
- C. Use graphs from the AWS Billing and Cost Management dashboard to compare EC2 costs based on instance types for the last 2 months.
- D. Use AWS Cost and Usage Reports to create a report and send it to an Amazon S3 bucket. Use Amazon QuickSight with Amazon S3 as a source to generate an interactive graph based on instance types.

Answer: B

Explanation:

AWS Cost Explorer is a tool that enables you to view and analyze your costs and usage. You can explore your usage and costs using the main graph, the Cost Explorer cost and usage reports, or the Cost Explorer RI reports. You can view data for up to the last 12 months, forecast how much you're likely to spend for the next 12 months, and get recommendations for what Reserved Instances to purchase. You can use Cost Explorer to identify areas that need further inquiry and see trends that you can use to understand your

costs. <https://docs.aws.amazon.com/cost-management/latest/userguide/ce-what-is.html>

NEW QUESTION 2

- (Topic 1)

A company hosts more than 300 global websites and applications. The company requires a platform to analyze more than 30 TB of clickstream data each day. What should a solutions architect do to transmit and process the clickstream data?

- A. Design an AWS Data Pipeline to archive the data to an Amazon S3 bucket and run an Amazon EMR cluster with the data to generate analytics.
- B. Create an Auto Scaling group of Amazon EC2 instances to process the data and send it to an Amazon S3 data lake for Amazon Redshift to use for analysis.
- C. Cache the data to Amazon CloudFront. Store the data in an Amazon S3 bucket. When an object is added to the S3 bucket, run an AWS Lambda function to process the data for analysis.
- D. Collect the data from Amazon Kinesis Data Stream.
- E. Use Amazon Kinesis Data Firehose to transmit the data to an Amazon S3 data lake. Load the data in Amazon Redshift for analysis.

Answer: D

Explanation:

<https://aws.amazon.com/es/blogs/big-data/real-time-analytics-with-amazon-redshift-streaming-ingestion/>

NEW QUESTION 3

- (Topic 1)

A company's website uses an Amazon EC2 instance store for its catalog of items. The company wants to make sure that the catalog is highly available and that the catalog is stored in a durable location.

What should a solutions architect do to meet these requirements?

- A. Move the catalog to Amazon ElastiCache for Redis.
- B. Deploy a larger EC2 instance with a larger instance store.
- C. Move the catalog from the instance store to Amazon S3 Glacier Deep Archive.
- D. Move the catalog to an Amazon Elastic File System (Amazon EFS) file system.

Answer: D

Explanation:

Moving the catalog to an Amazon Elastic File System (Amazon EFS) file system provides both high availability and durability. Amazon EFS is a fully-managed, highly-available, and durable file system that is built to scale on demand. With Amazon EFS, the catalog data can be stored and accessed from multiple EC2 instances in different availability zones, ensuring high availability. Also, Amazon EFS automatically stores files redundantly within and across multiple availability zones, making it a durable storage option.

NEW QUESTION 4

- (Topic 1)

A company has a three-tier web application that is deployed on AWS. The web servers are deployed in a public subnet in a VPC. The application servers and database servers are deployed in private subnets in the same VPC. The company has deployed a third-party virtual firewall appliance from AWS Marketplace in an inspection VPC. The appliance is configured with an IP interface that can accept IP packets. A solutions architect needs to integrate the web application with the appliance to inspect all traffic to the application before the traffic reaches the web server. Which solution will meet these requirements with the LEAST operational overhead?

- A. Create a Network Load Balancer in the public subnet of the application's VPC to route the traffic to the appliance for packet inspection.
- B. Create an Application Load Balancer in the public subnet of the application's VPC to route the traffic to the appliance for packet inspection.
- C. Deploy a transit gateway in the inspection VPC. Configure route tables to route the incoming packets through the transit gateway.
- D. Deploy a Gateway Load Balancer in the inspection VPC. Create a Gateway Load Balancer endpoint to receive the incoming packets and forward the packets to the appliance.

Answer: D

Explanation:

<https://aws.amazon.com/blogs/networking-and-content-delivery/scaling-network-traffic-inspection-using-aws-gateway-load-balancer/>

NEW QUESTION 5

- (Topic 1)

A company uses NFS to store large video files in on-premises network attached storage. Each video file ranges in size from 1MB to 500 GB. The total storage is 70 TB and is no longer growing. The company decides to migrate the video files to Amazon S3. The company must migrate the video files as soon as possible while using the least possible network bandwidth.

Which solution will meet these requirements?

- A. Create an S3 bucket Create an IAM role that has permissions to write to the S3 bucket
- B. Use the AWS CLI to copy all files locally to the S3 bucket.
- C. Create an AWS Snowball Edge job
- D. Receive a Snowball Edge device on premise
- E. Use the Snowball Edge client to transfer data to the device
- F. Return the device so that AWS can import the data into Amazon S3.
- G. Deploy an S3 File Gateway on premise
- H. Create a public service endpoint to connect to the S3 File Gateway Create an S3 bucket Create a new NFS file share on the S3 File Gateway Point the new file share to the S3 bucket
- I. Transfer the data from the existing NFS file share to the S3 File Gateway.
- J. Set up an AWS Direct Connect connection between the on-premises network and AWS
- K. Deploy an S3 File Gateway on premise
- L. Create a public virtual interface (VIF) to connect to the S3 File Gateway
- M. Create an S3 bucket
- N. Create a new NFS file share on the S3 File Gateway
- O. Point the new file share to the S3 bucket
- P. Transfer the data from the existing NFS file share to the S3 File Gateway.

Answer: B

Explanation:

The basic difference between Snowball and Snowball Edge is the capacity they provide. Snowball provides a total of 50 TB or 80 TB, out of which 42 TB or 72 TB is available, while Amazon Snowball Edge provides 100 TB, out of which 83 TB is available.

NEW QUESTION 6

- (Topic 1)

A company has a production workload that runs on 1,000 Amazon EC2 Linux instances. The workload is powered by third-party software. The company needs to patch the third-party software on all EC2 instances as quickly as possible to remediate a critical security vulnerability.

What should a solutions architect do to meet these requirements?

- A. Create an AWS Lambda function to apply the patch to all EC2 instances.
- B. Configure AWS Systems Manager Patch Manager to apply the patch to all EC2 instances.
- C. Schedule an AWS Systems Manager maintenance window to apply the patch to all EC2 instances.
- D. Use AWS Systems Manager Run Command to run a custom command that applies the patch to all EC2 instances.

Answer: B

Explanation:

<https://docs.aws.amazon.com/systems-manager/latest/userguide/about-windows-app-patching.html>

NEW QUESTION 7

- (Topic 1)

A company uses 50 TB of data for reporting. The company wants to move this data from on premises to AWS. A custom application in the company's data center runs a weekly data transformation job. The company plans to pause the application until the data transfer is complete and needs to begin the transfer process as soon as possible.

The data center does not have any available network bandwidth for additional workloads. A solutions architect must transfer the data and must configure the transformation job to continue to run in the AWS Cloud.

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

- A. Use AWS DataSync to move the data. Create a custom transformation job by using AWS Glue.
- B. Order an AWS Snowcone device to move the data. Deploy the transformation application to the device.
- C. Order an AWS Snowball Edge Storage Optimized device.
- D. Copy the data to the device.
- E. Create a custom transformation job by using AWS Glue.
- F. Order an AWS Snowball Edge Storage Optimized device that includes Amazon EC2 compute. Copy the data to the device. Create a new EC2 instance on AWS to run the transformation application.
- G. Snowball Edge Storage Optimized device that includes Amazon EC2 compute. Copy the data to the device. Create a new EC2 instance on AWS to run the transformation application.

Answer: D

Explanation:

AWS Snowball Edge is a type of Snowball device with on-board storage and compute power for select AWS capabilities. Snowball Edge can do local processing and edge-computing workloads in addition to transferring data between your local environment and the AWS Cloud¹. Users can order an AWS Snowball Edge Storage Optimized device that includes Amazon EC2 compute to move 50 TB of data from on premises to AWS. The Storage Optimized device has 80 TB of usable storage and 40 vCPUs of compute power². Users can copy the data to the device using the AWS OpsHub graphical user interface or the Snowball client command line tool³. Users can also create and run Amazon EC2 instances on the device using Amazon Machine Images (AMIs) that are compatible with the sb1 instance type. Users can use the Snowball Edge device to transfer the data and run the transformation job locally without using any network bandwidth. Users can also create a new EC2 instance on AWS to run the transformation application after the data transfer is complete. Amazon EC2 is a web service that provides secure, resizable compute capacity in the cloud. Users can launch an EC2 instance in the same AWS Region where they send their Snowball Edge device and choose an AMI that matches their application requirements. Users can use the EC2 instance to continue running the transformation job in the AWS Cloud.

NEW QUESTION 8

- (Topic 1)

A company wants to migrate an on-premises data center to AWS. The data center hosts an SFTP server that stores its data on an NFS-based file system. The server holds 200 GB of data that needs to be transferred. The server must be hosted on an Amazon EC2 instance that uses an Amazon Elastic File System (Amazon EFS) file system

When combination of steps should a solutions architect take to automate this task? (Select TWO)

- A. Launch the EC2 instance into the same Availability Zone as the EFS file system
- B. Install an AWS DataSync agent in the on-premises data center
- C. Create a secondary Amazon Elastic Block Store (Amazon EBS) volume on the EC2 instance for the data
- D. Manually use an operating system copy command to push the data to the EC2 instance
- E. Use AWS DataSync to create a suitable location configuration for the on-premises SFTP server

Answer: BE

Explanation:

AWS DataSync is an online data movement and discovery service that simplifies data migration and helps users quickly, easily, and securely move their file or object data to, from, and between AWS storage services¹. Users can use AWS DataSync to transfer data between on-premises and AWS storage services. To use AWS DataSync, users need to install an AWS DataSync agent in the on-premises data center. The agent is a software appliance that connects to the source or destination storage system and handles the data transfer to or from AWS over the network². Users also need to use AWS DataSync to create a suitable location configuration for the on-premises SFTP server. A location is a logical representation of a storage system that contains files or objects that users want to transfer using DataSync. Users can create locations for NFS shares, SMB shares, HDFS file systems, self-managed object storage, Amazon S3 buckets, Amazon EFS file systems, Amazon FSx for Windows File Server file systems, Amazon FSx for Lustre file systems, Amazon FSx for OpenZFS file systems, Amazon FSx for NetApp ONTAP file systems, and AWS Snowcone devices³.

NEW QUESTION 9

- (Topic 1)

A company is implementing a shared storage solution for a media application that is hosted in the AWS Cloud. The company needs the ability to use SMB clients to access data. The solution must be fully managed.

Which AWS solution meets these requirements?

- A. Create an AWS Storage Gateway volume gateway
- B. Create a file share that uses the required client protocol. Connect the application server to the file share.
- C. Create an AWS Storage Gateway tape gateway. Configure it to use Amazon S3. Connect the application server to the tape gateway.
- D. Create an Amazon EC2 Windows instance. Install and configure a Windows file share role on the instance.
- E. Connect the application server to the file share.
- F. Create an Amazon FSx for Windows File Server file system. Attach the file system to the origin server.
- G. Connect the application server to the file system.

Answer: D

Explanation:

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

Amazon FSx has native support for Windows file system features and for the industry-standard Server Message Block (SMB) protocol to access file storage over a network. <https://docs.aws.amazon.com/fsx/latest/WindowsGuide/what-is.html>

NEW QUESTION 10

- (Topic 1)

A company has registered its domain name with Amazon Route 53. The company uses Amazon API Gateway in the us-east-1 Region as a public interface for its backend microservice APIs. Third-party services consume the APIs securely. The company wants to design its API Gateway URL with the company's domain name and corresponding certificate so that the third-party services can use HTTPS.

Which solution will meet these requirements?

- A. Create stage variables in API Gateway with Name="Endpoint-URL" and Value="Company Domain Name" to overwrite the default URL
- B. Import the public certificate associated with the company's domain name into AWS Certificate Manager (ACM).
- C. Create Route 53 DNS records with the company's domain name
- D. Point the alias record to the Regional API Gateway stage endpoint
- E. Import the public certificate associated with the company's domain name into AWS Certificate Manager (ACM) in the us-east-1 Region.
- F. Create a Regional API Gateway endpoint
- G. Associate the API Gateway endpoint with the company's domain name
- H. Import the public certificate associated with the company's domain name into AWS Certificate Manager (ACM) in the same Region
- I. Attach the certificate to the API Gateway endpoint
- J. Configure Route 53 to route traffic to the API Gateway endpoint.
- K. Create a Regional API Gateway endpoint
- L. Associate the API Gateway endpoint with the company's domain name
- M. Import the public certificate associated with the company's domain name into AWS Certificate Manager (ACM) in the us-east-1 Region
- N. Attach the certificate to the API Gateway API
- O. Create Route 53 DNS records with the company's domain name
- P. Point an A record to the company's domain name.

Answer: C

Explanation:

To design the API Gateway URL with the company's domain name and corresponding certificate, the company needs to do the following: 1. Create a Regional API Gateway endpoint: This will allow the company to create an endpoint that is specific to a region. 2. Associate the API Gateway endpoint with the company's domain name: This will allow the company to use its own domain name for the API Gateway URL. 3. Import the public certificate associated with the company's domain name into AWS Certificate Manager (ACM) in the same Region: This will allow the company to use HTTPS for secure communication with its APIs. 4. Attach the certificate to the API Gateway endpoint: This will allow the company to use the certificate for securing the API Gateway URL. 5. Configure Route 53 to route traffic to the API Gateway endpoint: This will allow the company to use Route 53 to route traffic to the API Gateway URL using the company's domain name.

NEW QUESTION 10

- (Topic 1)

A company needs to review its AWS Cloud deployment to ensure that its Amazon S3 buckets do not have unauthorized configuration changes. What should a solutions architect do to accomplish this goal?

- A. Turn on AWS Config with the appropriate rules.
- B. Turn on AWS Trusted Advisor with the appropriate checks.
- C. Turn on Amazon Inspector with the appropriate assessment template.
- D. Turn on Amazon S3 server access logging
- E. Configure Amazon EventBridge (Amazon Cloud Watch Events).

Answer: A

Explanation:

To ensure that Amazon S3 buckets do not have unauthorized configuration changes, a solutions architect should turn on AWS Config with the appropriate rules. AWS Config is a service that allows users to audit and assess their AWS resource configurations for compliance with industry standards and internal policies. It provides a detailed view of the resources and their configurations, including information on how the resources are related to each other. By turning on AWS Config with the appropriate rules, users can identify and remediate unauthorized configuration changes to their Amazon S3 buckets.

NEW QUESTION 12

- (Topic 1)

A company is storing sensitive user information in an Amazon S3 bucket. The company wants to provide secure access to this bucket from the application tier running on Amazon EC2 instances inside a VPC.

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

- A. Configure a VPC gateway endpoint for Amazon S3 within the VPC
- B. Create a bucket policy to make the objects in the S3 bucket public
- C. Create a bucket policy that limits access to only the application tier running in the VPC
- D. Create an IAM user with an S3 access policy and copy the IAM credentials to the EC2 instance
- E. Create a NAT instance and have the EC2 instances use the NAT instance to access the S3 bucket

Answer: AC

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/s3-private-connection-no-authentication/>

NEW QUESTION 15

- (Topic 1)

An Amazon EC2 administrator created the following policy associated with an IAM group containing several users

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "ec2:TerminateInstances",
      "Resource": "*",
      "Condition": {
        "IpAddress": {
          "aws:SourceIp": "10.100.100.0/24"
        }
      }
    },
    {
      "Effect": "Deny",
      "Action": "ec2:*",
      "Resource": "*",
      "Condition": {
        "StringNotEquals": {
          "ec2:Region": "us-east-1"
        }
      }
    }
  ]
}
```

What is the effect of this policy?

- A. Users can terminate an EC2 instance in any AWS Region except us-east-1.
- B. Users can terminate an EC2 instance with the IP address 10.100.100.1 in the us-east-1 Region
- C. Users can terminate an EC2 instance in the us-east-1 Region when the user's source IP is 10.100.100.254.
- D. Users cannot terminate an EC2 instance in the us-east-1 Region when the user's source IP is 10.100.100.254

Answer: C

Explanation:

as the policy prevents anyone from doing any EC2 action on any region except us-east-1 and allows only users with source ip 10.100.100.0/24 to terminate instances. So user with source ip 10.100.100.254 can terminate instances in us-east-1 region.

NEW QUESTION 17

- (Topic 1)

A company's dynamic website is hosted using on-premises servers in the United States. The company is launching its product in Europe, and it wants to optimize site loading times for new European users. The site's backend must remain in the United States. The product is being launched in a few days, and an immediate solution is needed.

What should the solutions architect recommend?

- A. Launch an Amazon EC2 instance in us-east-1 and migrate the site to it.
- B. Move the website to Amazon S3. Use cross-Region replication between Regions.
- C. Use Amazon CloudFront with a custom origin pointing to the on-premises servers.
- D. Use an Amazon Route 53 geo-proximity routing policy pointing to on-premises servers.

Answer: C

Explanation:

<https://aws.amazon.com/pt/blogs/aws/amazon-cloudfront-support-for-custom-origins/>

You can now create a CloudFront distribution using a custom origin. Each distribution will can point to an S3 or to a custom origin. This could be another storage service, or it could be something more interesting and more dynamic, such as an EC2 instance or even an Elastic Load Balancer

NEW QUESTION 19

- (Topic 1)

A company recently signed a contract with an AWS Managed Service Provider (MSP) Partner for help with an application migration initiative. A solutions architect needs to share an Amazon Machine Image (AMI) from an existing AWS account with the MSP Partner's AWS account. The AMI is backed by Amazon Elastic Block Store (Amazon EBS) and uses a customer managed customer master key (CMK) to encrypt EBS volume snapshots.

What is the MOST secure way for the solutions architect to share the AMI with the MSP Partner's AWS account?

- A. Make the encrypted AMI and snapshots publicly availabl
- B. Modify the CMK's key policy to allow the MSP Partner's AWS account to use the key
- C. Modify the launchPermission property of the AM
- D. Share the AMI with the MSP Partner's AWS account onl
- E. Modify the CMK's key policy to allow the MSP Partner's AWS account to use the key.
- F. Modify the launchPermission property of the AMI Share the AMI with the MSP Partner's AWS account onl
- G. Modify the CMK's key policy to trust a new CMK that is owned by the MSP Partner for encryption.
- H. Export the AMI from the source account to an Amazon S3 bucket in the MSP Partner's AWS accoun
- I. Encrypt the S3 bucket with a CMK that is owned by the MSP Partner Copy and launch the AMI in the MSP Partner's AWS account.

Answer: B

Explanation:

Share the existing KMS key with the MSP external account because it has already been used to encrypt the AMI snapshot.

<https://docs.aws.amazon.com/kms/latest/developerguide/key-policy-modifying-external-accounts.html>

NEW QUESTION 21

- (Topic 1)

A company is storing backup files by using Amazon S3 Standard storage. The files are accessed frequently for 1 month. However, the files are not accessed after 1 month. The company must keep the files indefinitely.

Which storage solution will meet these requirements MOST cost-effectively?

- A. Configure S3 Intelligent-Tiering to automatically migrate objects.
- B. Create an S3 Lifecycle configuration to transition objects from S3 Standard to S3 Glacier Deep Archive after 1 month.
- C. Create an S3 Lifecycle configuration to transition objects from S3 Standard to S3 Standard-Infrequent Access (S3 Standard-IA) after 1 month.
- D. Create an S3 Lifecycle configuration to transition objects from S3 Standard to S3 One Zone-Infrequent Access (S3 One Zone-IA) after 1 month.

Answer: B

Explanation:

The storage solution that will meet these requirements most cost-effectively is B: Create an S3 Lifecycle configuration to transition objects from S3 Standard to S3 Glacier Deep Archive after 1 month. Amazon S3 Glacier Deep Archive is a secure, durable, and extremely low-cost Amazon S3 storage class for long-term retention of data that is rarely accessed and for which retrieval times of several hours are acceptable. It is the lowest-cost storage option in Amazon S3, making it a cost-effective choice for storing backup files that are not accessed after 1 month. You can use an S3 Lifecycle configuration to automatically transition objects from S3 Standard to S3 Glacier Deep Archive after 1 month. This will minimize the storage costs for the backup files that are not accessed frequently.

NEW QUESTION 23

- (Topic 1)

A company runs a shopping application that uses Amazon DynamoDB to store customer information. In case of data corruption, a solutions architect needs to design a solution that meets a recovery point objective (RPO) of 15 minutes and a recovery time objective (RTO) of 1 hour.

What should the solutions architect recommend to meet these requirements?

- A. Configure DynamoDB global table
- B. For RPO recovery, point the application to a different AWS Region.
- C. Configure DynamoDB point-in-time recover
- D. For RPO recovery, restore to the desired point in time.
- E. Export the DynamoDB data to Amazon S3 Glacier on a daily basi
- F. For RPO recovery, import the data from S3 Glacier to DynamoDB.
- G. Schedule Amazon Elastic Block Store (Amazon EBS) snapshots for the DynamoDB table every 15 minute

H. For RPO recovery, restore the DynamoDB table by using the EBS snapshot.

Answer: B

Explanation:

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

NEW QUESTION 26

- (Topic 1)

A solutions architect is designing a VPC with public and private subnets. The VPC and subnets use IPv4 CIDR blocks. There is one public subnet and one private subnet in each of three Availability Zones (AZs) for high availability. An internet gateway is used to provide internet access for the public subnets. The private subnets require access to the internet to allow Amazon EC2 instances to download software updates.

What should the solutions architect do to enable Internet access for the private subnets?

- A. Create three NAT gateways, one for each public subnet in each AZ
- B. Create a private route table for each AZ that forwards non-VPC traffic to the NAT gateway in its AZ.
- C. Create three NAT instances, one for each private subnet in each AZ
- D. Create a private route table for each AZ that forwards non-VPC traffic to the NAT instance in its AZ.
- E. Create a second internet gateway on one of the private subnets
- F. Update the route table for the private subnets that forward non-VPC traffic to the private internet gateway.
- G. Create an egress-only internet gateway on one of the public subnets
- H. Update the route table for the private subnets that forward non-VPC traffic to the egress-only internet gateway.

Answer: A

Explanation:

<https://aws.amazon.com/about-aws/whats-new/2018/03/introducing-amazon-vpc-nat-gateway-in-the-aws-govcloud-us-region/#:~:text=NAT%20Gateway%20is%20a%20highly,instances%20in%20a%20private%20subnet.>

<https://docs.aws.amazon.com/vpc/latest/userguide/vpc-nat-comparison.html>

NEW QUESTION 31

- (Topic 1)

An application runs on an Amazon EC2 instance in a VPC. The application processes logs that are stored in an Amazon S3 bucket. The EC2 instance needs to access the S3 bucket without connectivity to the internet.

Which solution will provide private network connectivity to Amazon S3?

- A. Create a gateway VPC endpoint to the S3 bucket.
- B. Stream the logs to Amazon CloudWatch Log
- C. Export the logs to the S3 bucket.
- D. Create an instance profile on Amazon EC2 to allow S3 access.
- E. Create an Amazon API Gateway API with a private link to access the S3 endpoint.

Answer: A

Explanation:

VPC endpoint allows you to connect to AWS services using a private network instead of using the public Internet

NEW QUESTION 36

- (Topic 1)

A company recently launched a variety of new workloads on Amazon EC2 instances in its AWS account. The company needs to create a strategy to access and administer the instances remotely and securely. The company needs to implement a repeatable process that works with native AWS services and follows the AWS Well-Architected Framework.

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

- A. Use the EC2 serial console to directly access the terminal interface of each instance for administration.
- B. Attach the appropriate IAM role to each existing instance and new instances
- C. Use AWS Systems Manager Session Manager to establish a remote SSH session.
- D. Create an administrative SSH key pair
- E. Load the public key into each EC2 instance
- F. Deploy a bastion host in a public subnet to provide a tunnel for administration of each instance.
- G. Establish an AWS Site-to-Site VPN connection
- H. Instruct administrators to use their local on-premises machines to connect directly to the instances by using SSH keys across the VPN tunnel.

Answer: B

Explanation:

<https://docs.aws.amazon.com/systems-manager/latest/userguide/setup-launch-managed-instance.html>

NEW QUESTION 41

- (Topic 1)

A company wants to reduce the cost of its existing three-tier web architecture. The web, application, and database servers are running on Amazon EC2 instances for the development, test, and production environments. The EC2 instances average 30% CPU utilization during peak hours and 10% CPU utilization during non-peak hours.

The production EC2 instances run 24 hours a day. The development and test EC2 instances run for at least 8 hours each day. The company plans to implement automation to stop the development and test EC2 instances when they are not in use.

Which EC2 instance purchasing solution will meet the company's requirements MOST cost-effectively?

- A. Use Spot Instances for the production EC2 instances

- B. Use Reserved Instances for the development and test EC2 instances.
- C. Use Reserved Instances for the production EC2 instance
- D. Use On-Demand Instances for the development and test EC2 instances.
- E. Use Spot blocks for the production EC2 instance
- F. Use Reserved Instances for the development and test EC2 instances.
- G. Use On-Demand Instances for the production EC2 instance
- H. Use Spot blocks for the development and test EC2 instances.

Answer: B

NEW QUESTION 44

- (Topic 2)

A company hosts a two-tier application on Amazon EC2 instances and Amazon RDS. The application's demand varies based on the time of day. The load is minimal after work hours and on weekends. The EC2 instances run in an EC2 Auto Scaling group that is configured with a minimum of two instances and a maximum of five instances. The application must be available at all times, but the company is concerned about overall cost. Which solution meets the availability requirement MOST cost-effectively?

- A. Use all EC2 Spot Instance
- B. Stop the RDS database when it is not in use.
- C. Purchase EC2 Instance Savings Plans to cover five EC2 instance
- D. Purchase an RDS Reserved DB Instance
- E. Purchase two EC2 Reserved Instances Use up to three additional EC2 Spot Instances as needed
- F. Stop the RDS database when it is not in use.
- G. Purchase EC2 Instance Savings Plans to cover two EC2 instance
- H. Use up to three additional EC2 On-Demand Instances as needed
- I. Purchase an RDS Reserved DB Instance.

Answer: C

Explanation:

This solution meets the requirements of a two-tier application that has a variable demand based on the time of day and must be available at all times, while minimizing the overall cost. EC2 Reserved Instances can provide significant savings compared to On-Demand Instances for the baseline level of usage, and they can guarantee capacity reservation when needed. EC2 Spot Instances can provide up to 90% savings compared to On-Demand Instances for any additional capacity that the application needs during peak hours. Spot Instances are suitable for stateless applications that can tolerate interruptions and can be replaced by other instances. Stopping the RDS database when it is not in use can reduce the cost of running the database tier.

Option A is incorrect because using all EC2 Spot Instances can affect the availability of the application if there are not enough spare capacity or if the Spot price exceeds the maximum price. Stopping the RDS database when it is not in use can reduce the cost of running the database tier, but it can also affect the availability of the application. Option B is incorrect because purchasing EC2 Instance Savings Plans to cover five EC2 instances can lock in a fixed amount of compute usage per hour, which may not match the actual usage pattern of the application. Purchasing an RDS Reserved DB Instance can provide savings for the database tier, but it does not allow stopping the database when it is not in use. Option D is incorrect because purchasing EC2 Instance Savings Plans to cover two EC2 instances can lock in a fixed amount of compute usage per hour, which may not match the actual usage pattern of the application. Using up to three additional EC2 On-Demand Instances as needed can incur higher costs than using Spot Instances.

References:

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

? <https://aws.amazon.com/ec2/spot/>

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

NEW QUESTION 49

- (Topic 2)

A company has a Windows-based application that must be migrated to AWS. The application requires the use of a shared Windows file system attached to multiple Amazon EC2 Windows instances that are deployed across multiple Availability Zones.

What should a solutions architect do to meet this requirement?

- A. Configure AWS Storage Gateway in volume gateway mode
- B. Mount the volume to each Windows instance.
- C. Configure Amazon FSx for Windows File Server
- D. Mount the Amazon FSx file system to each Windows instance.
- E. Configure a file system by using Amazon Elastic File System (Amazon EFS). Mount the EFS file system to each Windows instance.
- F. Configure an Amazon Elastic Block Store (Amazon EBS) volume with the required size
- G. Attach each EC2 instance to the volume
- H. Mount the file system within the volume to each Windows instance.

Answer: B

Explanation:

This solution meets the requirement of migrating a Windows-based application that requires the use of a shared Windows file system attached to multiple Amazon EC2 Windows instances that are deployed across multiple Availability Zones. Amazon FSx for Windows File Server provides fully managed shared storage built on Windows Server, and delivers a wide range of data access, data management, and administrative capabilities. It supports the Server Message Block (SMB) protocol and can be mounted to EC2 Windows instances across multiple Availability Zones.

Option A is incorrect because AWS Storage Gateway in volume gateway mode provides cloud-backed storage volumes that can be mounted as iSCSI devices from on-premises application servers, but it does not support SMB protocol or EC2 Windows instances. Option C is incorrect because Amazon Elastic File System (Amazon EFS) provides a scalable and elastic NFS file system for Linux-based workloads, but it does not support SMB protocol or EC2 Windows instances.

Option D is incorrect because Amazon Elastic Block Store (Amazon EBS) provides persistent block storage volumes for use with EC2 instances, but it does not support SMB protocol or attaching multiple instances to the same volume.

References:

? <https://aws.amazon.com/fsx/windows/>

? <https://docs.aws.amazon.com/fsx/latest/WindowsGuide/using-file-shares.html>

NEW QUESTION 51

- (Topic 2)

A solutions architect needs to help a company optimize the cost of running an application on AWS. The application will use Amazon EC2 instances, AWS Fargate, and AWS Lambda for compute within the architecture.

The EC2 instances will run the data ingestion layer of the application. EC2 usage will be sporadic and unpredictable. Workloads that run on EC2 instances can be interrupted at any time. The application front end will run on Fargate, and Lambda will serve the API layer. The front-end utilization and API layer utilization will be predictable over the course of the next year.

Which combination of purchasing options will provide the MOST cost-effective solution for hosting this application? (Choose two.)

- A. Use Spot Instances for the data ingestion layer
- B. Use On-Demand Instances for the data ingestion layer
- C. Purchase a 1-year Compute Savings Plan for the front end and API layer.
- D. Purchase 1-year All Upfront Reserved instances for the data ingestion layer.
- E. Purchase a 1-year EC2 instance Savings Plan for the front end and API layer.

Answer: AC

Explanation:

EC2 instance Savings Plan saves 72% while Compute Savings Plans saves 66%. But according to link, it says "Compute Savings Plans provide the most flexibility and help to reduce your costs by up to 66%. These plans automatically apply to EC2 instance usage regardless of instance family, size, AZ, region, OS or tenancy, and also apply to Fargate and Lambda usage." EC2 instance Savings Plans are not applied to Fargate or Lambda

NEW QUESTION 56

- (Topic 2)

An application runs on Amazon EC2 instances across multiple Availability Zones The instances run in an Amazon EC2 Auto Scaling group behind an Application Load Balancer The application performs best when the CPU utilization of the EC2 instances is at or near 40%.

What should a solutions architect do to maintain the desired performance across all instances in the group?

- A. Use a simple scaling policy to dynamically scale the Auto Scaling group
- B. Use a target tracking policy to dynamically scale the Auto Scaling group
- C. Use an AWS Lambda function to update the desired Auto Scaling group capacity.
- D. Use scheduled scaling actions to scale up and scale down the Auto Scaling group

Answer: B

Explanation:

<https://docs.aws.amazon.com/autoscaling/application/userguide/application-auto-scaling-target-tracking.html>

NEW QUESTION 59

- (Topic 2)

A company is migrating its on-premises PostgreSQL database to Amazon Aurora PostgreSQL. The on-premises database must remain online and accessible during the migration. The Aurora database must remain synchronized with the on-premises database.

Which combination of actions must a solutions architect take to meet these requirements? (Choose two.)

- A. Create an ongoing replication task.
- B. Create a database backup of the on-premises database
- C. Create an AWS Database Migration Service (AWS DMS) replication server
- D. Convert the database schema by using the AWS Schema Conversion Tool (AWS SCT).
- E. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to monitor the database synchronization

Answer: AC

Explanation:

AWS Database Migration Service supports homogeneous migrations such as Oracle to Oracle, as well as heterogeneous migrations between different database platforms, such as Oracle or Microsoft SQL Server to Amazon Aurora. With AWS Database Migration Service, you can also continuously replicate data with low latency from any supported source to any supported target. For example, you can replicate from multiple sources to Amazon Simple Storage Service (Amazon S3) to build a highly available and scalable data lake solution. You can also consolidate databases into a petabyte-scale data warehouse by streaming data to Amazon Redshift. Learn more about the supported source and target databases. <https://aws.amazon.com/dms/>

NEW QUESTION 63

- (Topic 2)

An online retail company has more than 50 million active customers and receives more than 25,000 orders each day. The company collects purchase data for customers and stores this data in Amazon S3. Additional customer data is stored in Amazon RDS.

The company wants to make all the data available to various teams so that the teams can perform analytics. The solution must provide the ability to manage fine-grained permissions for the data and must minimize operational overhead.

Which solution will meet these requirements?

- A. Migrate the purchase data to write directly to Amazon RD
- B. Use RDS access controls to limit access.
- C. Schedule an AWS Lambda function to periodically copy data from Amazon RDS to Amazon S3. Create an AWS Glue crawle
- D. Use Amazon Athena to query the dat
- E. Use S3 policies to limit access.
- F. Create a data lake by using AWS Lake Formatio
- G. Create an AWS Glue JDBC connection to Amazon RD
- H. Register (he S3 bucket in Lake Formatio
- I. Use Lake Formation access controls to limit access.
- J. Create an Amazon Redshift cluste
- K. Schedule an AWS Lambda function to periodically copy data from Amazon S3 and Amazon RDS to Amazon Redshif
- L. Use Amazon Redshift access controls to limit access.

Answer: C

Explanation:

To make all the data available to various teams and minimize operational overhead, the company can create a data lake by using AWS Lake Formation. This will allow the company to centralize all the data in one place and use fine-grained access controls to manage access to the data. To meet the requirements of the company, the solutions architect can create a data lake by using AWS Lake Formation, create an AWS Glue JDBC connection to Amazon RDS, and register the S3 bucket in Lake Formation. The solutions architect can then use Lake Formation access controls to limit access to the data. This solution will provide the ability to manage fine-grained permissions for the data and minimize operational overhead.

NEW QUESTION 65

- (Topic 2)

A company wants to measure the effectiveness of its recent marketing campaigns. The company performs batch processing on csv files of sales data and stores the results in an Amazon S3 bucket once every hour. The S3 bucket contains petabytes of objects. The company runs one-time queries in Amazon Athena to determine which products are most popular on a particular date for a particular region. Queries sometimes fail or take longer than expected to finish.

Which actions should a solutions architect take to improve the query performance and reliability? (Select TWO.)

- A. Reduce the S3 object sizes to less than 126 MB
- B. Partition the data by date and region in Amazon S3
- C. Store the files as large, single objects in Amazon S3.
- D. Use Amazon Kinesis Data Analytics to run the Queries as part of the batch processing operation
- E. Use an AWS data extract, transform, and load (ETL) process to convert the csv files into Apache Parquet format.

Answer: BE

Explanation:

<https://aws.amazon.com/blogs/big-data/top-10-performance-tuning-tips-for-amazon-athena/>

This solution meets the requirements of measuring the effectiveness of marketing campaigns by performing batch processing on csv files of sales data and storing the results in an Amazon S3 bucket once every hour. An AWS data ETL process can use services such as AWS Glue or AWS Data Pipeline to extract data from S3, transform it into a more efficient format such as Apache Parquet, and load it back into S3. Apache Parquet is a columnar storage format that can improve the query performance and reliability of Athena by reducing the amount of data scanned, improving compression ratio, and enabling predicate pushdown.

NEW QUESTION 66

- (Topic 2)

A company has an ecommerce checkout workflow that writes an order to a database and calls a service to process the payment. Users are experiencing timeouts during the checkout process. When users resubmit the checkout form, multiple unique orders are created for the same desired transaction.

How should a solutions architect refactor this workflow to prevent the creation of multiple orders?

- A. Configure the web application to send an order message to Amazon Kinesis Data Firehose
- B. Set the payment service to retrieve the message from Kinesis Data Firehose and process the order.
- C. Create a rule in AWS CloudTrail to invoke an AWS Lambda function based on the logged application path request. Use Lambda to query the database, call the payment service, and pass in the order information.
- D. Store the order in the database
- E. Send a message that includes the order number to Amazon Simple Notification Service (Amazon SNS). Set the payment service to poll Amazon SNS
- F. retrieve the message, and process the order.
- G. Store the order in the database
- H. Send a message that includes the order number to an Amazon Simple Queue Service (Amazon SQS) FIFO queue
- I. Set the payment service to retrieve the message and process the order
- J. Delete the message from the queue.

Answer: D

Explanation:

This approach ensures that the order creation and payment processing steps are separate and atomic. By sending the order information to an SQS FIFO queue, the payment service can process the order one at a time and in the order they were received. If the payment service is unable to process an order, it can be retried later, preventing the creation of multiple orders. The deletion of the message from the queue after it is processed will prevent the same message from being processed multiple times.

NEW QUESTION 71

- (Topic 2)

A company wants to move its application to a serverless solution. The serverless solution needs to analyze existing and new data by using SQL. The company stores the data in an Amazon S3 bucket. The data requires encryption and must be replicated to a different AWS Region.

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

- A. Create a new S3 bucket
- B. Load the data into the new S3 bucket
- C. Use S3 Cross-Region Replication (CRR) to replicate encrypted objects to an S3 bucket in another Region
- D. Use server-side encryption with AWS KMS multi-Region keys (SSE-KMS). Use Amazon Athena to query the data.
- E. Create a new S3 bucket
- F. Load the data into the new S3 bucket
- G. Use S3 Cross-Region Replication (CRR) to replicate encrypted objects to an S3 bucket in another Region
- H. Use server-side encryption with AWS KMS multi-Region keys (SSE-KMS). Use Amazon RDS to query the data.
- I. Load the data into the existing S3 bucket
- J. Use S3 Cross-Region Replication (CRR) to replicate encrypted objects to an S3 bucket in another Region
- K. Use server-side encryption with Amazon S3 managed encryption keys (SSE-S3). Use Amazon Athena to query the data.
- L. Load the data into the existing S3 bucket
- M. Use S3 Cross-Region Replication (CRR) to replicate encrypted objects to an S3 bucket in another Region
- N. Use server-side encryption with Amazon S3 managed encryption keys (SSE-S3). Use Amazon RDS to query the data.

Answer: A

Explanation:

This solution meets the requirements of a serverless solution, encryption, replication, and SQL analysis with the least operational overhead. Amazon Athena is a

serverless interactive query service that can analyze data in S3 using standard SQL. S3 Cross-Region Replication (CRR) can replicate encrypted objects to an S3 bucket in another Region automatically. Server-side encryption with AWS KMS multi-Region keys (SSE-KMS) can encrypt the data at rest using keys that are replicated across multiple Regions. Creating a new S3 bucket can avoid potential conflicts with existing data or configurations.

Option B is incorrect because Amazon RDS is not a serverless solution and it cannot query data in S3 directly. Option C is incorrect because server-side encryption with Amazon S3 managed encryption keys (SSE-S3) does not use KMS keys and it does not support multi-Region replication. Option D is incorrect because Amazon RDS is not a serverless solution and it cannot query data in S3 directly. It is also incorrect for the same reason as option C. References:

? <https://docs.aws.amazon.com/AmazonS3/latest/userguide/replication-walkthrough-4.html>

? <https://aws.amazon.com/blogs/storage/considering-four-different-replication-options-for-data-in-amazon-s3/>

? <https://docs.aws.amazon.com/AmazonS3/latest/userguide/UsingEncryption.html>

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

NEW QUESTION 74

- (Topic 2)

A media company is evaluating the possibility of moving its systems to the AWS Cloud. The company needs at least 10 TB of storage with the maximum possible I/O performance for video processing, 300 TB of very durable storage for storing media content, and 900 TB of storage to meet requirements for archival media that is not in use anymore.

Which set of services should a solutions architect recommend to meet these requirements?

- A. Amazon EBS for maximum performance, Amazon S3 for durable data storage, and Amazon S3 Glacier for archival storage
- B. Amazon EBS for maximum performance, Amazon EFS for durable data storage and Amazon S3 Glacier for archival storage
- C. Amazon EC2 instance store for maximum performance
- D. Amazon EFS for durable data storage and Amazon S3 for archival storage
- E. Amazon EC2 Instance store for maximum performance
- F. Amazon S3 for durable data storage, and Amazon S3 Glacier for archival storage

Answer: A

Explanation:

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/InstanceStorage.html>

NEW QUESTION 78

- (Topic 2)

A solutions architect is creating a new Amazon CloudFront distribution for an application. Some of the information submitted by users is sensitive. The application uses HTTPS but needs another layer of security. The sensitive information should be protected throughout the entire application stack, and access to the information should be restricted to certain applications.

Which action should the solutions architect take?

- A. Configure a CloudFront signed URL.
- B. Configure a CloudFront signed cookie.
- C. Configure a CloudFront field-level encryption profile.
- D. Configure CloudFront and set the Origin Protocol Policy setting to HTTPS Only for the Viewer Protocol Policy.

Answer: C

Explanation:

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/field-level-encryption.html>

"With Amazon CloudFront, you can enforce secure end-to-end connections to origin servers by using HTTPS. Field-level encryption adds an additional layer of security that lets you protect specific data throughout system processing so that only certain applications can see it."

NEW QUESTION 81

- (Topic 2)

A company runs workloads on AWS. The company needs to connect to a service from an external provider. The service is hosted in the provider's VPC. According to the company's security team, the connectivity must be private and must be restricted to the target service. The connection must be initiated only from the company's VPC.

Which solution will meet these requirements?

- A. Create a VPC peering connection between the company's VPC and the provider's VPC
- B. Update the route table to connect to the target service.
- C. Ask the provider to create a virtual private gateway in its VPC
- D. Use AWS PrivateLink to connect to the target service.
- E. Create a NAT gateway in a public subnet of the company's VPC
- F. Update the route table to connect to the target service.
- G. Ask the provider to create a VPC endpoint for the target service
- H. Use AWS PrivateLink to connect to the target service.

Answer: D

Explanation:

AWS PrivateLink provides private connectivity between VPCs, AWS services, and your on-premises networks, without exposing your traffic to the public internet. AWS PrivateLink makes it easy to connect services across different accounts and VPCs to significantly simplify your network architecture. Interface

VPC endpoints, powered by AWS PrivateLink, connect you to services hosted by AWS Partners and supported solutions available in AWS Marketplace.

<https://aws.amazon.com/privatelink/>

NEW QUESTION 83

- (Topic 2)

A company runs a web-based portal that provides users with global breaking news, local alerts, and weather updates. The portal delivers each user a personalized view by using a mixture of static and dynamic content. Content is served over HTTPS through an API server running on an Amazon EC2 instance behind an Application Load Balancer (ALB). The company wants the portal to provide this content to its users across the world as quickly as possible.

How should a solutions architect design the application to ensure the LEAST amount of latency for all users?

- A. Deploy the application stack in a single AWS Region
- B. Use Amazon CloudFront to serve all static and dynamic content by specifying the ALB as an origin.
- C. Deploy the application stack in two AWS Region
- D. Use an Amazon Route 53 latency routing policy to serve all content from the ALB in the closest Region.
- E. Deploy the application stack in a single AWS Region
- F. Use Amazon CloudFront to serve the static content
- G. Serve the dynamic content directly from the ALB.
- H. Deploy the application stack in two AWS Region
- I. Use an Amazon Route 53 geolocation routing policy to serve all content from the ALB in the closest Region.

Answer: A

Explanation:

<https://aws.amazon.com/blogs/networking-and-content-delivery/deliver-your-apps-dynamic-content-using-amazon-cloudfront-getting-started-template/>

NEW QUESTION 85

- (Topic 2)

A company runs an Oracle database on premises. As part of the company's migration to AWS, the company wants to upgrade the database to the most recent available version. The company also wants to set up disaster recovery (DR) for the database. The company needs to minimize the operational overhead for normal operations and DR setup. The company also needs to maintain access to the database's underlying operating system. Which solution will meet these requirements?

- A. Migrate the Oracle database to an Amazon EC2 instance
- B. Set up database replication to a different AWS Region.
- C. Migrate the Oracle database to Amazon RDS for Oracle
- D. Activate Cross-Region automated backups to replicate the snapshots to another AWS Region.
- E. Migrate the Oracle database to Amazon RDS Custom for Oracle
- F. Create a read replica for the database in another AWS Region.
- G. Migrate the Oracle database to Amazon RDS for Oracle
- H. Create a standby database in another Availability Zone.

Answer: C

Explanation:

<https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/rds-custom.html> and <https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/working-with-custom-oracle.html>

NEW QUESTION 86

- (Topic 2)

A solutions architect must design a solution that uses Amazon CloudFront with an Amazon S3 origin to store a static website. The company's security policy requires that all website traffic be inspected by AWS WAF. How should the solutions architect comply with these requirements?

- A. Configure an S3 bucket policy to accept requests coming from the AWS WAF Amazon Resource Name (ARN) only.
- B. Configure Amazon CloudFront to forward all incoming requests to AWS WAF before requesting content from the S3 origin.
- C. Configure a security group that allows Amazon CloudFront IP addresses to access Amazon S3 only.
- D. Associate AWS WAF to CloudFront.
- E. Configure Amazon CloudFront and Amazon S3 to use an origin access identity (OAI) to restrict access to the S3 bucket.
- F. Enable AWS WAF on the distribution.

Answer: D

Explanation:

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/private-content-restricting-access-to-s3.html>
<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/distribution-web-aws-waf.html>

NEW QUESTION 90

- (Topic 2)

A gaming company has a web application that displays scores. The application runs on Amazon EC2 instances behind an Application Load Balancer. The application stores data in an Amazon RDS for MySQL database. Users are starting to experience long delays and interruptions that are caused by database read performance. The company wants to improve the user experience while minimizing changes to the application's architecture. What should a solutions architect do to meet these requirements?

- A. Use Amazon ElastiCache in front of the database.
- B. Use RDS Proxy between the application and the database.
- C. Migrate the application from EC2 instances to AWS Lambda.
- D. Migrate the database from Amazon RDS for MySQL to Amazon DynamoDB.

Answer: A

Explanation:

ElastiCache can help speed up the read performance of the database by caching frequently accessed data, reducing latency and allowing the application to access the data more quickly. This solution requires minimal modifications to the current architecture, as ElastiCache can be used in conjunction with the existing Amazon RDS for MySQL database.

NEW QUESTION 94

- (Topic 2)

A business's backup data totals 700 terabytes (TB) and is kept in network attached storage (NAS) at its data center. This backup data must be available in the event of occasional regulatory inquiries and preserved for a period of seven years. The organization has chosen to relocate its backup data from its on-premises

data center to Amazon Web Services (AWS). Within one month, the migration must be completed. The company's public internet connection provides 500 Mbps of dedicated capacity for data transport.

What should a solutions architect do to ensure that data is migrated and stored at the LOWEST possible cost?

- A. Order AWS Snowball devices to transfer the data
- B. Use a lifecycle policy to transition the files to Amazon S3 Glacier Deep Archive.
- C. Deploy a VPN connection between the data center and Amazon VPC
- D. Use the AWS CLI to copy the data from on premises to Amazon S3 Glacier.
- E. Provision a 500 Mbps AWS Direct Connect connection and transfer the data to Amazon S3. Use a lifecycle policy to transition the files to Amazon S3 Glacier Deep Archive.
- F. Use AWS DataSync to transfer the data and deploy a DataSync agent on premise
- G. Use the DataSync task to copy files from the on-premises NAS storage to Amazon S3 Glacier.

Answer: A

Explanation:

<https://www.omnicalculator.com/other/data-transfer>

NEW QUESTION 96

- (Topic 2)

A global company is using Amazon API Gateway to design REST APIs for its loyalty club users in the us-east-1 Region and the ap-southeast-2 Region. A solutions architect must design a solution to protect these API Gateway managed REST APIs across multiple accounts from SQL injection and cross-site scripting attacks.

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

- A. Set up AWS WAF in both Region
- B. Associate Regional web ACLs with an API stage.
- C. Set up AWS Firewall Manager in both Region
- D. Centrally configure AWS WAF rules.
- E. Set up AWS Shield in both Region
- F. Associate Regional web ACLs with an API stage.
- G. Set up AWS Shield in one of the Region
- H. Associate Regional web ACLs with an API stage.

Answer: A

Explanation:

Using AWS WAF has several benefits. Additional protection against web attacks using criteria that you specify. You can define criteria using characteristics of web requests such as the following: Presence of SQL code that is likely to be malicious (known as SQL injection). Presence of a script that is likely to be malicious (known as cross-site scripting). AWS Firewall Manager simplifies your administration and maintenance tasks across multiple accounts and resources for a variety of protections. <https://docs.aws.amazon.com/waf/latest/developerguide/what-is-aws-waf.html>

NEW QUESTION 101

- (Topic 2)

A company has a dynamic web application hosted on two Amazon EC2 instances. The company has its own SSL certificate, which is on each instance to perform SSL termination.

There has been an increase in traffic recently, and the operations team determined that SSL encryption and decryption is causing the compute capacity of the web servers to reach their maximum limit.

What should a solutions architect do to increase the application's performance?

- A. Create a new SSL certificate using AWS Certificate Manager (ACM) install the ACM certificate on each instance
- B. Create an Amazon S3 bucket Migrate the SSL certificate to the S3 bucket Configure the EC2 instances to reference the bucket for SSL termination
- C. Create another EC2 instance as a proxy server Migrate the SSL certificate to the new instance and configure it to direct connections to the existing EC2 instances
- D. Import the SSL certificate into AWS Certificate Manager (ACM) Create an Application Load Balancer with an HTTPS listener that uses the SSL certificate from ACM

Answer: D

Explanation:

<https://aws.amazon.com/certificate-manager/>:

"With AWS Certificate Manager, you can quickly request a certificate, deploy it on ACM- integrated AWS resources, such as Elastic Load Balancers, Amazon CloudFront distributions, and APIs on API Gateway, and let AWS Certificate Manager handle certificate renewals. It also enables you to create private certificates for your internal resources and manage the certificate lifecycle centrally."

NEW QUESTION 102

- (Topic 2)

A company runs a stateless web application in production on a group of Amazon EC2 On- Demand Instances behind an Application Load Balancer. The application experiences heavy usage during an 8-hour period each business day. Application usage is moderate and steady overnight Application usage is low during weekends.

The company wants to minimize its EC2 costs without affecting the availability of the application.

Which solution will meet these requirements?

- A. Use Spot Instances for the entire workload.
- B. Use Reserved instances for the baseline level of usage Use Spot Instances for any additional capacity that the application needs.
- C. Use On-Demand Instances for the baseline level of usage
- D. Use Spot Instances for any additional capacity that the application needs
- E. Use Dedicated Instances for the baseline level of usage
- F. Use On-Demand Instances for any additional capacity that the application needs

Answer: B

Explanation:

Reserved is cheaper than on demand the company has. And it's meet the availability (HA) requirement as to spot instance that can be disrupted at any time. PRICING BELOW. On- Demand: 0% There's no commitment from you. You pay the most with this option. Reserved : 40%-60% 1-year or 3-year commitment from you. You save money from that commitment. Spot 50%-90% Ridiculously inexpensive because there's no commitment from the AWS side.

NEW QUESTION 105

- (Topic 2)

A company has a legacy data processing application that runs on Amazon EC2 instances. Data is processed sequentially, but the order of results does not matter. The application uses a monolithic architecture. The only way that the company can scale the application to meet increased demand is to increase the size of the instances.

The company's developers have decided to rewrite the application to use a microservices architecture on Amazon Elastic Container Service (Amazon ECS).

What should a solutions architect recommend for communication between the microservices?

- A. Create an Amazon Simple Queue Service (Amazon SQS) queue
- B. Add code to the data producers, and send data to the queue
- C. Add code to the data consumers to process data from the queue.
- D. Create an Amazon Simple Notification Service (Amazon SNS) topic
- E. Add code to the data producers, and publish notifications to the topic
- F. Add code to the data consumers to subscribe to the topic.
- G. Create an AWS Lambda function to pass message
- H. Add code to the data producers to call the Lambda function with a data object
- I. Add code to the data consumers to receive a data object that is passed from the Lambda function.
- J. Create an Amazon DynamoDB table
- K. Enable DynamoDB Stream
- L. Add code to the data producers to insert data into the table
- M. Add code to the data consumers to use the DynamoDB Streams API to detect new table entries and retrieve the data.

Answer: A

Explanation:

Queue has Limited throughput (300 msg/s without batching, 3000 msg/s with batching whereby up-to 10 msg per batch operation; Msg duplicates not allowed in the queue (exactly-once delivery); Msg order is preserved (FIFO); Queue name must end with .fifo

NEW QUESTION 106

- (Topic 2)

A company wants to run applications in containers in the AWS Cloud. These applications are stateless and can tolerate disruptions within the underlying infrastructure. The company needs a solution that minimizes cost and operational overhead.

What should a solutions architect do to meet these requirements?

- A. Use Spot Instances in an Amazon EC2 Auto Scaling group to run the application containers.
- B. Use Spot Instances in an Amazon Elastic Kubernetes Service (Amazon EKS) managed node group.
- C. Use On-Demand Instances in an Amazon EC2 Auto Scaling group to run the application containers.
- D. Use On-Demand Instances in an Amazon Elastic Kubernetes Service (Amazon EKS) managed node group.

Answer: A

Explanation:

<https://aws.amazon.com/cn/blogs/compute/cost-optimization-and-resilience-eks-with-spot-instances/>

NEW QUESTION 110

- (Topic 2)

A company uses a three-tier web application to provide training to new employees. The application is accessed for only 12 hours every day. The company is using an Amazon RDS for MySQL DB instance to store information and wants to minimize costs.

What should a solutions architect do to meet these requirements?

- A. Configure an IAM policy for AWS Systems Manager Session Manager
- B. Create an IAM role for the policy
- C. Update the trust relationship of the role
- D. Set up automatic start and stop for the DB instance.
- E. Create an Amazon ElastiCache for Redis cache cluster that gives users the ability to access the data from the cache when the DB instance is stopped
- F. Invalidate the cache after the DB instance is started.
- G. Launch an Amazon EC2 instance
- H. Create an IAM role that grants access to Amazon RDS
- I. Attach the role to the EC2 instance
- J. Configure a cron job to start and stop the EC2 instance on the desired schedule.
- K. Create AWS Lambda functions to start and stop the DB instance
- L. Create Amazon EventBridge (Amazon CloudWatch Events) scheduled rules to invoke the Lambda function
- M. Configure the Lambda functions as event targets for the rules

Answer: D

Explanation:

In a typical development environment, dev and test databases are mostly utilized for 8 hours a day and sit idle when not in use. However, the databases are billed for the compute and storage costs during this idle time. To reduce the overall cost, Amazon RDS allows instances to be stopped temporarily. While the instance is stopped, you're charged for storage and backups, but not for the DB instance hours. Please note that a stopped instance will automatically be started after 7 days. This post presents a solution using AWS Lambda and Amazon EventBridge that allows you to schedule a Lambda function to stop and start the idle databases with specific tags to save on compute costs. The second post presents a solution that accomplishes stop and start of the idle Amazon RDS databases using AWS Systems Manager.

NEW QUESTION 112

- (Topic 3)

An ecommerce company is experiencing an increase in user traffic. The company's store is deployed on Amazon EC2 instances as a two-tier web application consisting of a web tier and a separate database tier. As traffic increases, the company notices that the architecture is causing significant delays in sending timely marketing and order confirmation email to users. The company wants to reduce the time it spends resolving complex email delivery issues and minimize operational overhead.

What should a solutions architect do to meet these requirements?

- A. Create a separate application tier using EC2 instances dedicated to email processing.
- B. Configure the web instance to send email through Amazon Simple Email Service (Amazon SES).
- C. Configure the web instance to send email through Amazon Simple Notification Service (Amazon SNS)
- D. Create a separate application tier using EC2 instances dedicated to email processing
- E. Place the instances in an Auto Scaling group.

Answer: B

Explanation:

Amazon SES is a cost-effective and scalable email service that enables businesses to send and receive email using their own email addresses and domains. Configuring the web instance to send email through Amazon SES is a simple and effective solution that can reduce the time spent resolving complex email delivery issues and minimize operational overhead.

NEW QUESTION 113

- (Topic 3)

A company is running a multi-tier recommence web application in the AWS Cloud. The application runs on Amazon EC2 instances with an Amazon RDS for MySQL Multi-AZ DB instance. Amazon RDS is configured with the latest generation DB instance with 2,000 GB of storage in a General Purpose SSD (gp3) Amazon Elastic Block Store (Amazon EBS) volume. The database performance affects the application during periods of high demand.

A database administrator analyzes the logs in Amazon CloudWatch Logs and discovers that the application performance always degrades when the number of read and write IOPS is higher than 20,000.

What should a solutions architect do to improve the application performance?

- A. Replace the volume with a magnetic volume.
- B. Increase the number of IOPS on the gp3 volume.
- C. Replace the volume with a Provisioned IOPS SSD (io2) volume.
- D. Replace the 2,000 GB gp3 volume with two 1,000 GB gp3 volumes

Answer: C

Explanation:

<https://aws.amazon.com/ebs/features/> Amazon EBS provides a range of options that allow you to optimize storage performance and cost for your workload. These options are divided into two major categories: SSD-backed storage for transactional workloads, such as databases and boot volumes (performance depends primarily on IOPS), and HDD-backed storage for throughput intensive workloads, such as MapReduce and log processing (performance depends primarily on MB/s).

NEW QUESTION 118

- (Topic 3)

A gaming company is moving its public scoreboard from a data center to the AWS Cloud. The company uses Amazon EC2 Windows Server instances behind an Application Load Balancer to host its dynamic application. The company needs a highly available storage solution for the application. The application consists of static files and dynamic server-side code.

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

- A. Store the static files on Amazon S3. Use Amazon CloudFront to cache objects at the edge.
- B. Store the static files on Amazon S3. Use Amazon ElastiCache to cache objects at the edge.
- C. Store the server-side code on Amazon Elastic File System (Amazon EFS). Mount the EFS volume on each EC2 instance to share the files.
- D. Store the server-side code on Amazon FSx for Windows File Server
- E. Mount the FSx for Windows File Server volume on each EC2 instance to share the files.
- F. Store the server-side code on a General Purpose SSD (gp2) Amazon Elastic Block Store (Amazon EBS) volume
- G. Mount the EBS volume on each EC2 instance to share the files.

Answer: AD

Explanation:

A because ElastiCache, despite being ideal for leaderboards per Amazon, doesn't cache at edge locations. D because FSx has higher performance for low latency needs. <https://www.techtarget.com/searchaws/tip/Amazon-FSx-vs-EFS-Compare-the-AWS-file-services> "FSx is built for high performance and submillisecond latency using solid-state drive storage volumes. This design enables users to select storage capacity and latency independently. Thus, even a subterabyte file system can have 256 Mbps or higher throughput and support volumes up to 64 TB."

Amazon S3 is an object storage service that can store static files such as images, videos, documents, etc. Amazon EFS is a file storage service that can store files in a hierarchical structure and supports NFS protocol. Amazon FSx for Windows File Server is a file storage service that can store files in a hierarchical structure and supports SMB protocol. Amazon EBS is a block storage service that can store data in fixed-size blocks and attach to EC2 instances.

Based on these definitions, the combination of steps that should be taken to meet the requirements are:

* A. Store the static files on Amazon S3. Use Amazon CloudFront to cache objects at the edge. D. Store the server-side code on Amazon FSx for Windows File Server. Mount the FSx for Windows File Server volume on each EC2 instance to share the files.

NEW QUESTION 122

- (Topic 3)

A company wants to use high performance computing (HPC) infrastructure on AWS for financial risk modeling. The company's HPC workloads run on Linux. Each HPC workflow runs on hundreds of Amazon EC2 Spot Instances, is short-lived, and generates thousands of output files that are ultimately stored in persistent storage for analytics and long-term future use.

The company seeks a cloud storage solution that permits the copying of on-premises data to long-term persistent storage to make data available for processing by all EC2 instances. The solution should also be a high performance file system that is integrated with persistent storage to read and write datasets and output files.

Which combination of AWS services meets these requirements?

- A. Amazon FSx for Lustre integrated with Amazon S3
- B. Amazon FSx for Windows File Server integrated with Amazon S3
- C. Amazon S3 Glacier integrated with Amazon Elastic Block Store (Amazon EBS)
- D. Amazon S3 bucket with a VPC endpoint integrated with an Amazon Elastic Block Store (Amazon EBS) General Purpose SSD (gp2) volume

Answer: A

Explanation:

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

Amazon FSx for Lustre is a fully managed service that provides cost-effective, high-performance, scalable storage for compute workloads. Many workloads such as machine learning, high performance computing (HPC), video rendering, and financial simulations depend on compute instances accessing the same set of data through high-performance shared storage.

NEW QUESTION 125

- (Topic 3)

A company hosts a marketing website in an on-premises data center. The website consists of static documents and runs on a single server. An administrator updates the website content infrequently and uses an SFTP client to upload new documents.

The company decides to host its website on AWS and to use Amazon CloudFront. The company's solutions architect creates a CloudFront distribution. The solutions architect

must design the most cost-effective and resilient architecture for website hosting to serve as the CloudFront origin.

Which solution will meet these requirements?

- A. Create a virtual server by using Amazon Lightsail
- B. Configure the web server in the Lightsail instance
- C. Upload website content by using an SFTP client.
- D. Create an AWS Auto Scaling group for Amazon EC2 instance
- E. Use an Application Load Balance
- F. Upload website content by using an SFTP client.
- G. Create a private Amazon S3 bucket
- H. Use an S3 bucket policy to allow access from a CloudFront origin access identity (OAI). Upload website content by using the AWS CLI.
- I. Create a public Amazon S3 bucket
- J. Configure AWS Transfer for SFTP
- K. Configure the S3 bucket for website hosting
- L. Upload website content by using the SFTP client.

Answer: C

Explanation:

<https://docs.aws.amazon.com/cli/latest/reference/transfer/describe-server.html>

NEW QUESTION 127

- (Topic 3)

A company has hundreds of Amazon EC2 Linux-based instances in the AWS Cloud. Systems administrators have used shared SSH keys to manage the instances. After a recent audit, the company's security team is mandating the removal of all shared keys. A solutions architect must design a solution that provides secure access to the EC2 instances.

Which solution will meet this requirement with the LEAST amount of administrative overhead?

- A. Use AWS Systems Manager Session Manager to connect to the EC2 instances.
- B. Use AWS Security Token Service (AWS STS) to generate one-time SSH keys on demand.
- C. Allow shared SSH access to a set of bastion instances
- D. Configure all other instances to allow only SSH access from the bastion instances
- E. Use an Amazon Cognito custom authorizer to authenticate user
- F. Invoke an AWS Lambda function to generate a temporary SSH key.

Answer: A

Explanation:

Session Manager is a fully managed AWS Systems Manager capability. With Session Manager, you can manage your Amazon Elastic Compute Cloud (Amazon EC2) instances, edge devices, on-premises servers, and virtual machines (VMs). You can use either an interactive one-click browser-based shell or the AWS Command Line Interface (AWS CLI). Session Manager provides secure and auditable node management without the need to open inbound ports, maintain bastion hosts, or manage SSH keys. Session Manager also allows you to comply with corporate policies that require controlled access to managed nodes, strict security practices, and fully auditable logs with node access details, while providing end users with simple one-click cross-platform access to your managed nodes.

<https://docs.aws.amazon.com/systems-manager/latest/userguide/session-manager.html>

NEW QUESTION 129

- (Topic 3)

A company's facility has badge readers at every entrance throughout the building. When badges are scanned, the readers send a message over HTTPS to indicate who attempted to access that particular entrance.

A solutions architect must design a system to process these messages from the sensors. The solution must be highly available, and the results must be made available for the company's security team to analyze.

Which system architecture should the solutions architect recommend?

- A. Launch an Amazon EC2 instance to serve as the HTTPS endpoint and to process the messages. Configure the EC2 instance to save the results to an Amazon S3 bucket.
- B. Create an HTTPS endpoint in Amazon API Gateway
- C. Configure the API Gateway endpoint to invoke an AWS Lambda function to process the messages and save the results to an Amazon DynamoDB table.
- D. Use Amazon Route 53 to direct incoming sensor messages to an AWS Lambda function
- E. Configure the Lambda function to process the messages and save the results to an Amazon DynamoDB table.

F. Create a gateway VPC endpoint for Amazon S3. Configure a Site-to-Site VPN connection from the facility network to the VPC so that sensor data can be written directly to an S3 bucket by way of the VPC endpoint.

Answer: B

Explanation:

Deploy Amazon API Gateway as an HTTPS endpoint and AWS Lambda to process and save the messages to an Amazon DynamoDB table. This option provides a highly available and scalable solution that can easily handle large amounts of data. It also integrates with other AWS services, making it easier to analyze and visualize the data for the security team.

NEW QUESTION 132

- (Topic 3)

A company provides an online service for posting video content and transcoding it for use by any mobile platform. The application architecture uses Amazon Elastic File System (Amazon EFS) Standard to collect and store the videos so that multiple Amazon EC2 Linux instances can access the video content for processing. As the popularity of the service has grown over time, the storage costs have become too expensive. Which storage solution is MOST cost-effective?

- A. Use AWS Storage Gateway for files to store and process the video content
- B. Use AWS Storage Gateway for volumes to store and process the video content
- C. Use Amazon EFS for storing the video content. Once processing is complete, transfer the files to Amazon Elastic Block Store (Amazon EBS)
- D. Use Amazon S3 for storing the video content. Move the files temporarily over to an Amazon Elastic Block Store (Amazon EBS) volume attached to the server for processing

Answer: D

Explanation:

• Amazon S3 for large-scale, durable, and inexpensive storage of the video content. S3 storage costs are significantly lower than EFS. • Amazon EBS only temporarily during processing. By mounting an EBS volume only when a video needs to be processed, and unmounting it after, the time the content spends on the higher-cost EBS storage is minimized. • The EBS volume can be sized to match the workload needs for active processing, keeping costs lower. The volume does not need to store the entire video library long-term.

NEW QUESTION 134

- (Topic 3)

A developer has an application that uses an AWS Lambda function to upload files to Amazon S3 and needs the required permissions to perform the task. The developer already has an IAM user with valid IAM credentials required for Amazon S3. What should a solutions architect do to grant the permissions?

- A. Add required IAM permissions in the resource policy of the Lambda function
- B. Create a signed request using the existing IAM credentials in the Lambda function
- C. Create a new IAM user and use the existing IAM credentials in the Lambda function.
- D. Create an IAM execution role with the required permissions and attach the IAM role to the Lambda function

Answer: D

Explanation:

To grant the necessary permissions to an AWS Lambda function to upload files to Amazon S3, a solutions architect should create an IAM execution role with the required permissions and attach the IAM role to the Lambda function. This approach follows the principle of least privilege and ensures that the Lambda function can only access the resources it needs to perform its specific task.

NEW QUESTION 136

- (Topic 3)

A solutions architect must migrate a Windows Internet Information Services (IIS) web application to AWS. The application currently relies on a file share hosted in the user's on-premises network-attached storage (NAS). The solutions architect has proposed migrating the MS web servers to Amazon EC2 instances in multiple Availability Zones that are connected to the storage solution, and configuring an Elastic Load Balancer attached to the instances. Which replacement to the on-premises file share is MOST resilient and durable?

- A. Migrate the file share to Amazon RDS
- B. Migrate the file share to AWS Storage Gateway
- C. Migrate the file share to Amazon FSx for Windows File Server
- D. Migrate the file share to Amazon Elastic File System (Amazon EFS)

Answer: C

Explanation:

This answer is correct because it provides a resilient and durable replacement for the on-premises file share that is compatible with Windows IIS web servers. Amazon FSx for Windows File Server is a fully managed service that provides shared file storage built on Windows Server. It supports the SMB protocol and integrates with Microsoft Active Directory, which enables seamless access and authentication for Windows-based applications. Amazon FSx for Windows File Server also offers the following benefits:

? Resilience: Amazon FSx for Windows File Server can be deployed in multiple

Availability Zones, which provides high availability and failover protection. It also supports automatic backups and restores, as well as self-healing features that detect and correct issues.

? Durability: Amazon FSx for Windows File Server replicates data within and across

Availability Zones, and stores data on highly durable storage devices. It also supports encryption at rest and in transit, as well as file access auditing and data deduplication.

? Performance: Amazon FSx for Windows File Server delivers consistent sub-

millisecond latencies and high throughput for file operations. It also supports SSD storage, native Windows features such as Distributed File System (DFS)

Namespaces and Replication, and user-driven performance scaling.

References:

? Amazon FSx for Windows File Server

? Using Microsoft Windows file shares

NEW QUESTION 137

- (Topic 3)

A transaction processing company has weekly scripted batch jobs that run on Amazon EC2 instances. The EC2 instances are in an Auto Scaling group. The number of transactions can vary but the baseline CPU utilization that is noted on each run is at least 60%. The company needs to provision the capacity 30 minutes before the jobs run.

Currently engineering complete this task by manually modifying the Auto Scaling group parameters. The company does not have the resources to analyze the required capacity trends for the Auto Scaling group counts. The company needs an automated way to modify the Auto Scaling group's capacity.

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

- A. Create a dynamic scaling policy for the Auto Scaling group
- B. Configure the policy to scale based on the CPU utilization metric to 60%.
- C. Create a scheduled scaling policy for the Auto Scaling group
- D. Set the appropriate desired capacity, minimum capacity, and maximum capacity
- E. Set the recurrence to weekly
- F. Set the start time to 30 minutes
- G. Before the batch jobs run.
- H. Create a predictive scaling policy for the Auto Scaling group
- I. Configure the policy to scale based on forecast
- J. Set the scaling metric to CPU utilization
- K. Set the target value for the metric to 60%. In the Policy, set the instances to pre-launch 30 minutes before the jobs run.
- L. Create an Amazon EventBridge event to invoke an AWS Lambda function when the CPU utilization metric value for the Auto Scaling group reaches 60%. Configure the Lambda function to increase the Auto Scaling group's desired capacity and maximum capacity by 20%.

Answer: C

Explanation:

This option is the most efficient because it uses a predictive scaling policy for the Auto Scaling group, which is a type of scaling policy that uses machine learning to predict capacity requirements based on historical data from CloudWatch¹. It also configures the policy to scale based on forecast, which enables the Auto Scaling group to adjust its capacity in advance of traffic changes. It also sets the scaling metric to CPU utilization and the target value for the metric to 60%, which aligns with the baseline CPU utilization that is noted on each run. It also sets the instances to pre-launch 30 minutes before the jobs run, which ensures that enough capacity is provisioned before the weekly scripted batch jobs start. This solution meets the requirement of provisioning the capacity 30 minutes before the jobs run with the least operational overhead. Option A is less efficient because it uses a dynamic scaling policy for the Auto Scaling group, which is a type of scaling policy that adjusts your Auto Scaling group's capacity in response to changing demand². However, this does not provide a way to provision the capacity 30 minutes before the jobs run, as it only reacts to changing traffic. Option B is less efficient because it uses a scheduled scaling policy for the Auto Scaling group, which is a type of scaling policy that lets you scale your Auto Scaling group based on a schedule that you create³. However, this does not provide a way to scale based on forecast or CPU utilization, as it only scales based on predefined metrics and policies. Option D is less efficient because it uses an Amazon EventBridge event to invoke an AWS Lambda function when the CPU utilization metric value for the Auto Scaling group reaches 60%, which is a way to trigger serverless functions based on events. However, this does not provide a way to provision the capacity 30 minutes before the jobs run, as it only reacts to changing traffic.

NEW QUESTION 140

- (Topic 3)

A company hosts its application on AWS. The company uses Amazon Cognito to manage users. When users log in to the application, the application fetches required data from Amazon DynamoDB by using a REST API that is hosted in Amazon API Gateway. The company wants an AWS managed solution that will control access to the REST API to reduce development efforts.

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

- A. Configure an AWS Lambda function to be an authorizer in API Gateway to validate which user made the request
- B. For each user, create and assign an API key that must be sent with each request. Validate the key by using an AWS Lambda function
- C. Send the user's email address in the header with every request. Invoke an AWS Lambda function to validate that the user with that email address has proper access
- D. Configure an Amazon Cognito user pool authorizer in API Gateway to allow Amazon Cognito to validate each request

Answer: D

Explanation:

<https://docs.aws.amazon.com/apigateway/latest/developerguide/apigateway-integrate-with-cognito.html>

To control access to the REST API and reduce development efforts, the company can use an Amazon Cognito user pool authorizer in API Gateway. This will allow Amazon Cognito to validate each request and ensure that only authenticated users can access the API. This solution has the LEAST operational overhead, as it does not require the company to develop and maintain any additional infrastructure or code.

NEW QUESTION 145

- (Topic 3)

A company is deploying a two-tier web application in a VPC. The web tier is using an Amazon EC2 Auto Scaling group with public subnets that span multiple Availability Zones. The database tier consists of an Amazon RDS for MySQL DB instance in separate private subnets. The web tier requires access to the database to retrieve product information.

The web application is not working as intended. The web application reports that it cannot connect to the database. The database is confirmed to be up and running. All configurations for the network ACLs, security groups, and route tables are still in their default states.

What should a solutions architect recommend to fix the application?

- A. Add an explicit rule to the private subnet's network ACL to allow traffic from the web tier's EC2 instances.
- B. Add a route in the VPC route table to allow traffic between the web tier's EC2 instances and the database tier.
- C. Deploy the web tier's EC2 instances and the database tier's RDS instance into two separate VPCs
- D. and configure VPC peering.
- E. Add an inbound rule to the security group of the database tier's RDS instance to allow traffic from the web tier's security group.

Answer: D

Explanation:

This answer is correct because it allows the web tier to access the database tier by using

security groups as a source, which is a recommended best practice for VPC connectivity. Security groups are stateful and can reference other security groups in the same VPC, which simplifies the configuration and maintenance of the firewall rules. By adding an inbound rule to the database tier's security group, the web tier's EC2 instances can connect to the RDS instance on port 3306, regardless of their IP addresses or subnets. References:

? Security groups - Amazon Virtual Private Cloud

? Best practices and reference architectures for VPC design

NEW QUESTION 146

- (Topic 3)

A company must migrate 20 TB of data from a data center to the AWS Cloud within 30 days. The company's network bandwidth is limited to 15 Mbps and cannot exceed 70% utilization. What should a solutions architect do to meet these requirements?

- A. Use AWS Snowball.
- B. Use AWS DataSync.
- C. Use a secure VPN connection.
- D. Use Amazon S3 Transfer Acceleration.

Answer: A

Explanation:

AWS Snowball is a secure data transport solution that accelerates moving large amounts of data into and out of the AWS cloud. It can move up to 80 TB of data at a time, and provides a network bandwidth of up to 50 Mbps, so it is well-suited for the task. Additionally, it is secure and easy to use, making it the ideal solution for this migration.

NEW QUESTION 149

- (Topic 3)

A solutions architect needs to design a system to store client case files. The files are core company assets and are important. The number of files will grow over time.

The files must be simultaneously accessible from multiple application servers that run on Amazon EC2 instances. The solution must have built-in redundancy. Which solution meets these requirements?

- A. Amazon Elastic File System (Amazon EFS)
- B. Amazon Elastic Block Store (Amazon EBS)
- C. Amazon S3 Glacier Deep Archive
- D. AWS Backup

Answer: A

Explanation:

Amazon EFS provides a simple, scalable, fully managed file system that can be simultaneously accessed from multiple EC2 instances and provides built-in redundancy. It is optimized for multiple EC2 instances to access the same files, and it is designed to be highly available, durable, and secure. It can scale up to petabytes of data and can handle thousands of concurrent connections, and is a cost-effective solution for storing and accessing large amounts of data.

NEW QUESTION 150

- (Topic 3)

A company uses a payment processing system that requires messages for a particular payment ID to be received in the same order that they were sent. Otherwise, the payments might be processed incorrectly.

Which actions should a solutions architect take to meet this requirement? (Select TWO.)

- A. Write the messages to an Amazon DynamoDB table with the payment ID as the partition key
- B. Write the messages to an Amazon Kinesis data stream with the payment ID as the partition key.
- C. Write the messages to an Amazon ElastiCache for Memcached cluster with the payment ID as the key
- D. Write the messages to an Amazon Simple Queue Service (Amazon SQS) queue. Set the message attribute to use the payment ID
- E. Write the messages to an Amazon Simple Queue Service (Amazon SQS) FIFO queue
- F. Set the message group to use the payment ID.

Answer: BE

Explanation:

1) SQS FIFO queues guarantee that messages are received in the exact order they are sent. Using the payment ID as the message group ensures all messages for a payment ID are received sequentially. 2) Kinesis data streams can also enforce ordering on a per partition key basis. Using the payment ID as the partition key will ensure strict ordering of messages for each payment ID.

NEW QUESTION 155

- (Topic 3)

A company has an on-premises MySQL database used by the global sales team with infrequent access patterns. The sales team requires the database to have minimal downtime. A database administrator wants to migrate this database to AWS without selecting a particular instance type in anticipation of more users in the future.

Which service should a solutions architect recommend?

- A. Amazon Aurora MySQL
- B. Amazon Aurora Serverless for MySQL
- C. Amazon Redshift Spectrum
- D. Amazon RDS for MySQL

Answer: B

Explanation:

Amazon Aurora Serverless for MySQL is a fully managed, auto-scaling relational database service that scales up or down automatically based on the application demand. This service provides all the capabilities of Amazon Aurora, such as high availability, durability, and security, without requiring the customer to provision

any database instances. With Amazon Aurora Serverless for MySQL, the sales team can enjoy minimal downtime since the database is designed to automatically scale to accommodate the increased traffic. Additionally, the service allows the customer to pay only for the capacity used, making it cost-effective for infrequent access patterns. Amazon RDS for MySQL could also be an option, but it requires the customer to select an instance type, and the database administrator would need to monitor and adjust the instance size manually to accommodate the increasing traffic.

NEW QUESTION 156

- (Topic 3)

A company has an application that collects data from IoT sensors on automobiles. The data is streamed and stored in Amazon S3 through Amazon Kinesis Data Firehose. The data produces trillions of S3 objects each year. Each morning, the company uses the data from the previous 30 days to retrain a suite of machine learning (ML) models.

Four times each year, the company uses the data from the previous 12 months to perform analysis and train other ML models. The data must be available with minimal delay for up to 1 year. After 1 year, the data must be retained for archival purposes.

Which storage solution meets these requirements MOST cost-effectively?

- A. Use the S3 Intelligent-Tiering storage class
- B. Create an S3 Lifecycle policy to transition objects to S3 Glacier Deep Archive after 1 year
- C. Use the S3 Intelligent-Tiering storage class
- D. Configure S3 Intelligent-Tiering to automatically move objects to S3 Glacier Deep Archive after 1 year.
- E. Use the S3 Standard-Infrequent Access (S3 Standard-IA) storage class
- F. Create an S3 Lifecycle policy to transition objects to S3 Glacier Deep Archive after 1 year.
- G. Use the S3 Standard storage class
- H. Create an S3 Lifecycle policy to transition objects to S3 Standard-Infrequent Access (S3 Standard-IA) after 30 days, and then to S3 Glacier Deep Archive after 1 year.

Answer: D

Explanation:

- First 30 days- data access every morning (predictable and frequently) – S3 standard - After 30 days, accessed 4 times a year – S3 infrequently access - Data preserved- S3 Glacier Deep Archive

NEW QUESTION 160

- (Topic 3)

A company plans to use Amazon ElastiCache for its multi-tier web application. A solutions architect creates a Cache VPC for the ElastiCache cluster and an App VPC for the application's Amazon EC2 instances. Both VPCs are in the us-east-1 Region.

The solutions architect must implement a solution to provide the application's EC2 instances with access to the ElastiCache cluster.

Which solution will meet these requirements MOST cost-effectively?

- A. Create a peering connection between the VPCs. Add a route table entry for the peering connection in both VPCs. Configure an inbound rule for the ElastiCache cluster's security group to allow inbound connection from the application's security group.
- B. Create a Transit VPC. Update the VPC route tables in the Cache VPC and the App VPC to route traffic through the Transit VPC. Configure an inbound rule for the ElastiCache cluster's security group to allow inbound connection from the application's security group.
- C. Create a peering connection between the VPCs. Add a route table entry for the peering connection in both VPCs. Configure an inbound rule for the peering connection's security group to allow inbound connection from the application's security group.
- D. Create a Transit VPC. Update the VPC route tables in the Cache VPC and the App VPC to route traffic through the Transit VPC. Configure an inbound rule for the Transit VPC's security group to allow inbound connection from the application's security group.

Answer: A

Explanation:

Creating a peering connection between the two VPCs and configuring an inbound rule for the ElastiCache cluster's security group to allow inbound connection from the application's security group is the most cost-effective solution. Peering connections are free and you only incur the cost of configuring the security group rules. The Transit VPC solution requires additional VPCs and associated resources, which would incur additional costs.

Before Testing | AWS Certification Information and Policies | AWS <https://aws.amazon.com/certification/policies/before-testing/>

NEW QUESTION 161

- (Topic 3)

A company hosts a three-tier web application that includes a PostgreSQL database. The database stores the metadata from documents. The company searches the metadata for key terms to retrieve documents that the company reviews in a report each month. The documents are stored in Amazon S3. The documents are usually written only once, but they are updated frequently. The reporting process takes a few hours with the use of relational queries. The reporting process must not affect any document modifications or the addition of new documents.

What are the MOST operationally efficient solutions that meet these requirements? (Select TWO)

- A. Set up a new Amazon DocumentDB (with MongoDB compatibility) cluster that includes a read replica. Scale the read replica to generate the reports.
- B. Set up a new Amazon RDS for PostgreSQL Reserved Instance and an On-Demand read replica. Scale the read replica to generate the reports.
- C. Set up a new Amazon Aurora PostgreSQL DB cluster that includes a Reserved Instance and an Aurora Replica. Issue queries to the Aurora Replica to generate the reports.
- D. Set up a new Amazon RDS for PostgreSQL Multi-AZ Reserved Instance. Configure the reporting module to query the secondary RDS node so that the reporting module does not affect the primary node.
- E. Set up a new Amazon DynamoDB table to store the documents. Use a fixed write capacity to support new document entries. Automatically scale the read capacity to support the reports.

Answer: BC

Explanation:

These options are operationally efficient because they use Amazon RDS read replicas to offload the reporting workload from the primary DB instance and avoid affecting any document modifications or the addition of new documents. They also use Reserved Instances for the primary DB instance to reduce costs and On-Demand or Aurora Replicas for the read replicas to scale as needed. Option A is less efficient because it uses Amazon S3 Glacier Flexible Retrieval, which is a cold storage class that has higher retrieval costs and longer retrieval times than Amazon S3 Standard. It also uses EventBridge rules to invoke the job nightly, which does not meet the requirement of processing incoming data files as soon as possible. Option D is less efficient because it uses AWS Lambda to process the files, which has a maximum execution time of 15 minutes per invocation, which might not be enough for processing each file that needs 3-8 minutes. It also uses

S3 event notifications to invoke the Lambda function when the files arrive, which could cause concurrency issues if there are thousands of small data files arriving periodically. Option E is less efficient because it uses Amazon DynamoDB, which is a NoSQL database service that does not support relational queries, which are needed for generating the reports. It also uses fixed write capacity, which could cause throttling or underutilization depending on the incoming data files.

NEW QUESTION 166

- (Topic 3)

A company runs an application on a large fleet of Amazon EC2 instances. The application reads and write entries into an Amazon DynamoDB table. The size of the DynamoDB table continuously grows, but the application needs only data from the last 30 days. The company needs a solution that minimizes cost and development effort.

Which solution meets these requirements?

- A. Use an AWS CloudFormation template to deploy the complete solution.
- B. Redeploy the CloudFormation stack every 30 days, and delete the original stack.
- C. Use an EC2 instance that runs a monitoring application from AWS Marketplace.
- D. Configure the monitoring application to use Amazon DynamoDB Streams to store the timestamp when a new item is created in the table.
- E. Use a script that runs on the EC2 instance to delete items that have a timestamp that is older than 30 days.
- F. Configure Amazon DynamoDB Streams to invoke an AWS Lambda function when a new item is created in the table.
- G. Configure the Lambda function to delete items in the table that are older than 30 days.
- H. Extend the application to add an attribute that has a value of the current timestamp plus 30 days to each new item that is created in the table.
- I. Configure DynamoDB to use the attribute as the TTL attribute.

Answer: D

Explanation:

Amazon DynamoDB Time to Live (TTL) allows you to define a per-item timestamp to determine when an item is no longer needed. Shortly after the date and time of the specified timestamp, DynamoDB deletes the item from your table without consuming any write throughput. TTL is provided at no extra cost as a means to reduce stored data volumes by retaining only the items that remain current for your workload's needs.

TTL is useful if you store items that lose relevance after a specific time. The following are example TTL use cases:

Remove user or sensor data after one year of inactivity in an application.

Archive expired items to an Amazon S3 data lake via Amazon DynamoDB Streams and AWS Lambda.

Retain sensitive data for a certain amount of time according to contractual or regulatory obligations.

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

NEW QUESTION 168

- (Topic 3)

A company runs an application that receives data from thousands of geographically dispersed remote devices that use UDP. The application processes the data immediately and sends a message back to the device if necessary. No data is stored.

The company needs a solution that minimizes latency for the data transmission from the devices. The solution also must provide rapid failover to another AWS Region.

Which solution will meet these requirements?

- A. Configure an Amazon Route 53 failover routing policy. Create a Network Load Balancer (NLB) in each of the two Regions. Configure the NLB to invoke an AWS Lambda function to process the data.
- B. Use AWS Global Accelerator. Create a Network Load Balancer (NLB) in each of the two Regions as an endpoint.
- C. Create an Amazon Elastic Container Service (Amazon ECS) cluster with the Fargate launch type. Create an ECS service on the cluster. Set the ECS service as the target for the NLB. Process the data in Amazon ECS.
- D. Use AWS Global Accelerator. Create an Application Load Balancer (ALB) in each of the two Regions as an endpoint. Create an Amazon Elastic Container Service (Amazon ECS) cluster with the Fargate launch type. Create an ECS service on the cluster.
- E. Set the ECS service as the target for the ALB. Process the data in Amazon ECS.
- F. Configure an Amazon Route 53 failover routing policy. Create an Application Load Balancer (ALB) in each of the two Regions. Create an Amazon Elastic Container Service (Amazon ECS) cluster with the Fargate launch type. Create an ECS service on the cluster. Set the ECS service as the target for the ALB. Process the data in Amazon ECS.

Answer: B

Explanation:

To meet the requirements of minimizing latency for data transmission from the devices and providing rapid failover to another AWS Region, the best solution would be to use AWS Global Accelerator in combination with a Network Load Balancer (NLB) and Amazon Elastic Container Service (Amazon ECS). AWS Global Accelerator is a service that improves the availability and performance of applications by using static IP addresses (Anycast) to route traffic to optimal AWS endpoints. With Global Accelerator, you can direct traffic to multiple Regions and endpoints, and provide automatic failover to another AWS Region.

NEW QUESTION 173

- (Topic 3)

An application that is hosted on Amazon EC2 instances needs to access an Amazon S3 bucket. Traffic must not traverse the internet. How should a solutions architect configure access to meet these requirements?

- A. Create a private hosted zone by using Amazon Route 53.
- B. Set up a gateway VPC endpoint for Amazon S3 in the VPC.
- C. Configure the EC2 instances to use a NAT gateway to access the S3 bucket.
- D. Establish an AWS Site-to-Site VPN connection between the VPC and the S3 bucket.

Answer: B

Explanation:

This option is the most efficient because it uses a gateway VPC endpoint for Amazon S3, which provides reliable connectivity to Amazon S3 without requiring an internet gateway or a NAT device for the VPC. A gateway VPC endpoint routes traffic from the VPC to Amazon S3 using a prefix list for the service and does not leave the AWS network. This meets the requirement of not traversing the internet. Option A is less efficient because it uses a private hosted zone by using Amazon Route 53, which is a DNS service that allows you to create custom domain names for your resources within your VPC. However, this does not provide connectivity to Amazon S3 without an internet gateway or a NAT device. Option C is less efficient because it uses a NAT gateway to access the S3 bucket, which is a highly available, managed Network Address Translation (NAT) service that enables instances in a private subnet to connect to the internet or other AWS

services, but prevents the internet from initiating a connection with those instances⁴. However, this does not meet the requirement of not traversing the internet. Option D is less efficient because it uses an AWS Site-to-Site VPN connection between the VPC and the S3 bucket, which is a secure and encrypted network connection between your on-premises network and your VPC. However, this does not meet the requirement of not traversing the internet.

NEW QUESTION 176

- (Topic 3)

A company runs a web application that is backed by Amazon RDS. A new database administrator caused data loss by accidentally editing information in a database table. To help recover from this type of incident, the company wants the ability to restore the database to its state from 5 minutes before any change within the last 30 days.

Which feature should the solutions architect include in the design to meet this requirement?

- A. Read replicas
- B. Manual snapshots
- C. Automated backups
- D. Multi-AZ deployments

Answer: C

Explanation:

<https://aws.amazon.com/rds/features/backup/>

Automated backups, will meet the requirement. Amazon RDS allows you to automatically create backups of your DB instance. Automated backups enable point-in-time recovery (PITR) for your DB instance down to a specific second within the retention period, which can be up to 35 days. By setting the retention period to 30 days, the company can restore the database to its state from up to 5 minutes before any change within the last 30 days.

NEW QUESTION 177

- (Topic 3)

A company is using Amazon CloudFront with this website. The company has enabled logging on the CloudFront distribution, and logs are saved in one of the company's Amazon S3 buckets. The company needs to perform advanced analyses on the logs and build visualizations. What should a solutions architect do to meet these requirements?

- A. Use standard SQL queries in Amazon Athena to analyze the CloudFront logs in the S3 bucket. Visualize the results with AWS Glue.
- B. Use standard SQL queries in Amazon Athena to analyze the CloudFront logs in the S3 bucket. Visualize the results with Amazon QuickSight.
- C. Use standard SQL queries in Amazon DynamoDB to analyze the CloudFront logs in the S3 bucket. Visualize the results with AWS Glue.
- D. Use standard SQL queries in Amazon DynamoDB to analyze the CloudFront logs in the S3 bucket. Visualize the results with Amazon QuickSight.

Answer: B

Explanation:

<https://docs.aws.amazon.com/quicksight/latest/user/welcome.html>

Using Athena to query the CloudFront logs in the S3 bucket and QuickSight to visualize the results is the best solution because it is cost-effective, scalable, and requires no infrastructure setup. It also provides a robust solution that enables the company to perform advanced analysis and build interactive visualizations without the need for a dedicated team of developers.

NEW QUESTION 180

- (Topic 3)

A company is migrating its on-premises workload to the AWS Cloud. The company already uses several Amazon EC2 instances and Amazon RDS DB instances. The company wants a solution that automatically starts and stops the EC2 instances and DB instances outside of business hours. The solution must minimize cost and infrastructure maintenance.

Which solution will meet these requirements?

- A. Scale the EC2 instances by using elastic resize. Scale the DB instances to zero outside of business hours.
- B. Explore AWS Marketplace for partner solutions that will automatically start and stop the EC2 instances and DB instances on a schedule.
- C. Launch another EC2 instance.
- D. Configure a crontab schedule to run shell scripts that will start and stop the existing EC2 instances and DB instances on a schedule.
- E. Create an AWS Lambda function that will start and stop the EC2 instances and DB instances. Configure Amazon EventBridge to invoke the Lambda function on a schedule.

Answer: D

Explanation:

The most efficient solution for automatically starting and stopping EC2 instances and DB instances on a schedule while minimizing cost and infrastructure maintenance is to create an AWS Lambda function and configure Amazon EventBridge to invoke the function on a schedule.

Option A, scaling EC2 instances by using elastic resize and scaling DB instances to zero outside of business hours, is not feasible as DB instances cannot be scaled to zero.

Option B, exploring AWS Marketplace for partner solutions, may be an option, but it may not be the most efficient solution and could potentially add additional costs.

Option C, launching another EC2 instance and configuring a crontab schedule to run shell scripts that will start and stop the existing EC2 instances and DB instances on a schedule, adds unnecessary infrastructure and maintenance.

NEW QUESTION 181

- (Topic 3)

A company is migrating a Linux-based web server group to AWS. The web servers must access files in a shared file store for some content. The company must not make any changes to the application.

What should a solutions architect do to meet these requirements?

- A. Create an Amazon S3 Standard bucket with access to the web servers.
- B. Configure an Amazon CloudFront distribution with an Amazon S3 bucket as the origin.
- C. Create an Amazon Elastic File System (Amazon EFS) file system.
- D. Mount the EFS file system on all web servers.

- E. Configure a General Purpose SSD (gp3) Amazon Elastic Block Store (Amazon EBS) volum
- F. Mount the EBS volume to all web servers.

Answer: C

Explanation:

Create an Amazon Elastic File System (Amazon EFS) file system. Mount the EFS file system on all web servers. To meet the requirements of providing a shared file store for Linux-based web servers without making changes to the application, using an Amazon EFS file system is the best solution. Amazon EFS is a managed NFS file system service that provides shared access to files across multiple Linux-based instances, which makes it suitable for this use case. Amazon S3 is not ideal for this scenario since it is an object storage service and not a file system, and it requires additional tools or libraries to mount the S3 bucket as a file system. Amazon CloudFront can be used to improve content delivery performance but is not necessary for this requirement. Additionally, Amazon EBS volumes can only be mounted to one instance at a time, so it is not suitable for sharing files across multiple instances.

NEW QUESTION 185

- (Topic 3)

A company's order system sends requests from clients to Amazon EC2 instances. The EC2 instances process the orders and then store the orders in a database on Amazon RDS. Users report that they must reprocess orders when the system fails. The company wants a resilient solution that can process orders automatically if a system outage occurs.

What should a solutions architect do to meet these requirements?

- A. Move the EC2 instances into an Auto Scaling group. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to target an Amazon Elastic Container Service (Amazon ECS) task.
- B. Move the EC2 instances into an Auto Scaling group behind an Application Load Balancer (ALB). Update the order system to send messages to the ALB endpoint.
- C. Move the EC2 instances into an Auto Scaling group. Configure the order system to send messages to an Amazon Simple Queue Service (Amazon SQS) queue. Configure the EC2 instances to consume messages from the queue.
- D. Create an Amazon Simple Notification Service (Amazon SNS) topic. Create an AWS Lambda function, and subscribe the function to the SNS topic. Configure the order system to send messages to the SNS topic. Send a command to the EC2 instances to process the messages by using AWS Systems Manager Run Command.

Answer: C

Explanation:

To meet the company's requirements of having a resilient solution that can process orders automatically in case of a system outage, the solutions architect needs to implement a fault-tolerant architecture. Based on the given scenario, a potential solution is to move the EC2 instances into an Auto Scaling group and configure the order system to send messages to an Amazon Simple Queue Service (Amazon SQS) queue. The EC2 instances can then consume messages from the queue.

NEW QUESTION 190

- (Topic 3)

A company recently migrated its entire IT environment to the AWS Cloud. The company discovers that users are provisioning oversized Amazon EC2 instances and modifying security group rules without using the appropriate change control process. A solutions architect must devise a strategy to track and audit these inventory and configuration changes.

Which actions should the solutions architect take to meet these requirements? (Select TWO)

- A. Enable AWS CloudTrail and use it for auditing.
- B. Use data lifecycle policies for the Amazon EC2 instances.
- C. Enable AWS Trusted Advisor and reference the security dashboard.
- D. Enable AWS Config and create rules for auditing and compliance purposes.
- E. Restore previous resource configurations with an AWS CloudFormation template.

Answer: AD

Explanation:

A) Enable AWS CloudTrail and use it for auditing. AWS CloudTrail provides a record of API calls and can be used to audit changes made to EC2 instances and security groups. By analyzing CloudTrail logs, the solutions architect can track who provisioned oversized instances or modified security groups without proper approval. D) Enable AWS Config and create rules for auditing and compliance purposes. AWS Config can record the configuration changes made to resources like EC2 instances and security groups. The solutions architect can create AWS Config rules to monitor for non-compliant changes, like launching certain instance types or opening security group ports without permission. AWS Config would alert on any violations of these rules.

NEW QUESTION 195

- (Topic 3)

A company wants to implement a disaster recovery plan for its primary on-premises file storage volume. The file storage volume is mounted from an Internet Small Computer Systems Interface (iSCSI) device on a local storage server. The file storage volume holds hundreds of terabytes (TB) of data.

The company wants to ensure that end users retain immediate access to all file types from the on-premises systems without experiencing latency.

Which solution will meet these requirements with the LEAST amount of change to the company's existing infrastructure?

- A. Provision an Amazon S3 File Gateway as a virtual machine (VM) that is hosted on-premise.
- B. Set the local cache to 10 TB.
- C. Modify existing applications to access the files through the NFS protocol.
- D. To recover from a disaster, provision an Amazon EC2 instance and mount the S3 bucket that contains the files.
- E. Provision an AWS Storage Gateway tape gateway.
- F. Use a data backup solution to back up all existing data to a virtual tape library.
- G. Configure the data backup solution to run nightly after the initial backup is complete.
- H. To recover from a disaster, provision an Amazon EC2 instance and restore the data to an Amazon Elastic Block Store (Amazon EBS) volume from the volumes in the virtual tape library.
- I. Provision an AWS Storage Gateway Volume Gateway cached volume.
- J. Set the local cache to 10 TB.
- K. Mount the Volume Gateway cached volume to the existing file server by using iSCSI.
- L. and copy all files to the storage volume.

- M. Configure scheduled snapshots of the storage volum
- N. To recover from a disaster, restore a snapshot to an Amazon Elastic Block Store (Amazon EBS) volume and attach the EBS volume to an Amazon EC2 instance.
- O. Provision an AWS Storage Gateway Volume Gateway stored volume with the same amount of disk space as the existing file storage volum
- P. Mount the Volume Gateway stored volume to the existing file server by using iSCSI, and copy all files to the storage volum
- Q. Configure scheduled snapshots of the storage volum
- R. To recover from a disaster, restore a snapshot to an Amazon Elastic Block Store (Amazon EBS) volume and attach the EBS volume to an Amazon EC2 instance.

Answer: D

Explanation:

"The company wants to ensure that end users retain immediate access to all file types from the on-premises systems " - Cached volumes: low latency access to most recent data - Stored volumes: entire dataset is on premise, scheduled backups to S3 Hence Volume Gateway stored volume is the apt choice.

NEW QUESTION 198

- (Topic 3)

A company needs a backup strategy for its three-tier stateless web application The web application runs on Amazon EC2 instances in an Auto Scaling group with a dynamic scaling policy that is configured to respond to scaling events The database tier runs on Amazon RDS for PostgreSQL The web application does not require temporary local storage on the EC2 instances The company's recovery point objective (RPO) is 2 hours

The backup strategy must maximize scalability and optimize resource utilization for this environment

Which solution will meet these requirements?

- A. Take snapshots of Amazon Elastic Block Store (Amazon EBS) volumes of the EC2 instances and database every 2 hours to meet the RPO
- B. Configure a snapshot lifecycle policy to take Amazon Elastic Block Store (Amazon EBS) snapshots Enable automated backups in Amazon RDS to meet the RPO
- C. Retain the latest Amazon Machine Images (AMIs) of the web and application tiers Enable automated backups in Amazon RDS and use point-in-time recovery to meet the RPO
- D. Take snapshots of Amazon Elastic Block Store (Amazon EBS) volumes of the EC2 instances every 2 hours Enable automated backups in Amazon RDS and use point-in-time recovery to meet the RPO

Answer: C

Explanation:

Since the application has no local data on instances, AMIs alone can meet the RPO by restoring instances from the most recent AMI backup. When combined with automated RDS backups for the database, this provides a complete backup solution for this environment. The other options involving EBS snapshots would be unnecessary given the stateless nature of the instances. AMIs provide all the backup needed for the app tier. This uses native, automated AWS backup features that require minimal ongoing management: - AMI automated backups provide point-in-time recovery for the stateless app tier. - RDS automated backups provide point-in-time recovery for the database.

NEW QUESTION 199

- (Topic 3)

A solutions architect must secure a VPC network that hosts Amazon EC2 instances The EC2 instances contain highly sensitive data and run on a private subnet According to company policy the EC2 instances that run in the VPC can access only approved third-party software repositories on the internet for software product updates that use the third party's URL Other internet traffic must be blocked.

Which solution meets these requirements?

- A. Update the route table for the private subnet to route the outbound traffic to an AWS Network Firewall
- B. Configure domain list rule groups
- C. Set up an AWS WAF web ACL
- D. Create a custom set of rules that filter traffic requests based on source and destination IP address range sets.
- E. Implement strict inbound security group rules Configure an outbound rule that allows traffic only to the authorized software repositories on the internet by specifying the URLs
- F. Configure an Application Load Balancer (ALB) in front of the EC2 instance
- G. Direct an outbound traffic to the ALB Use a URL-based rule listener in the ALB's target group for outbound access to the internet

Answer: A

Explanation:

Send the outbound connection from EC2 to Network Firewall. In Network Firewall, create stateful outbound rules to allow certain domains for software patch download and deny all other domains. <https://docs.aws.amazon.com/network-firewall/latest/developerguide/suricata-examples.html#suricata-example-domain-filtering>

NEW QUESTION 202

- (Topic 3)

A company needs to create an Amazon Elastic Kubernetes Service (Amazon EKS) cluster to host a digital media streaming application. The EKS cluster will use a managed node group that is backed by Amazon Elastic Block Store (Amazon EBS) volumes for storage. The company must encrypt all data at rest by using a customer managed key that is stored in AWS Key Management Service (AWS KMS)

Which combination of actions will meet this requirement with the LEAST operational overhead? (Select TWO.)

- A. Use a Kubernetes plugin that uses the customer managed key to perform data encryption.
- B. After creation of the EKS cluster, locate the EBS volume
- C. Enable encryption by using the customer managed key.
- D. Enable EBS encryption by default in the AWS Region where the EKS cluster will be create
- E. Select the customer managed key as the default key.
- F. Create the EKS cluster Create an IAM role that has cwulicy that grants permission to the customer managed ke
- G. Associate the role with the EKS cluster.
- H. Store the customer managed key as a Kubernetes secret in the EKS cluste
- I. Use the customer managed key to encrypt the EBS volumes.

Answer: CD

Explanation:

EBS encryption by default is a feature that enables encryption for all new EBS volumes and snapshots created in a Region¹. EBS encryption by default uses a service managed key or a customer managed key that is stored in AWS KMS¹. EBS encryption by default is suitable for scenarios where data at rest must be encrypted by using a customer managed key, such as the digital media streaming application in the scenario¹.

To meet the requirements of the scenario, the solutions architect should enable EBS encryption by default in the AWS Region where the EKS cluster will be created. The solutions architect should select the customer managed key as the default key for encryption¹. This way, all new EBS volumes and snapshots created in that Region will be encrypted by using the customer managed key.

EKS encryption provider support is a feature that enables envelope encryption of Kubernetes secrets in EKS with a customer managed key that is stored in AWS KMS²

. Envelope encryption means that data is encrypted by data encryption keys (DEKs) using AES-GCM; DEKs are encrypted by key encryption keys (KEKs) according to configuration in AWS KMS³. EKS encryption provider support is suitable for scenarios where secrets must be encrypted by using a customer managed key, such as the digital media streaming application in the scenario².

To meet the requirements of the scenario, the solutions architect should create the EKS cluster and create an IAM role that has a policy that grants permission to the customer managed key. The solutions architect should associate the role with the EKS cluster^{2r}. This way, the EKS cluster can use envelope encryption of Kubernetes secrets with the customer managed key.

NEW QUESTION 203

.....

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

SAA-C03 Practice Exam Features:

- * SAA-C03 Questions and Answers Updated Frequently
- * SAA-C03 Practice Questions Verified by Expert Senior Certified Staff
- * SAA-C03 Most Realistic Questions that Guarantee you a Pass on Your First Try
- * SAA-C03 Practice Test Questions in Multiple Choice Formats and Updates for 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The SAA-C03 Practice Test Here](#)