Amazon

Exam Questions AWS-Certified-Solutions-Architect-Professional

Amazon AWS Certified Solutions Architect Professional

NEW QUESTION 1
- (Exam Topic 2)
A company has set up its entire infrastructure on AWS. The company uses Amazon EC2 instances to host its ecommerce website and uses Amazon S3 to store static data. Three engineers at the company handle the cloud administration and development through one AWS account. Occasionally, an engineer alters an EC2 security group configuration of another engineer and causes noncompliance issues in the environment.
A solutions architect must set up a system that tracks changes that the engineers make. The system must send alerts when the engineers make noncompliant changes to the security settings for the EC2 instances.
What is the FASTEST way for the solutions architect to meet these requirements?

A. Set up AWS Organizations for the compan
B. Apply SCPs to govern and track noncompliant security group changes that are made to the AWS account.
C. Enable AWS CloudTrail to capture the changes to EC2 security group
D. Enable Amazon CtoudWatch rules to provide alerts when noncompliant security settings are detected.
E. Enable SCPs on the AWS account to provide alerts when noncompliant security group changes are made to the environment.
F. Enable AWS Config on the EC2 security groups to track any noncompliant changes Send the changes as alerts through an Amazon Simple Notification Service (Amazon SNS) topic.

Answer: D

Explanation:
https://aws.amazon.com/es/blogs/industries/how-to-monitor-alert-and-remediate-non-compliant-hipaa-findings

NEW QUESTION 2
- (Exam Topic 2)
A company has deployed its database on an Amazon RDS for MySQL DB instance in the us-east-1 Region. The company needs to make its data available to customers in Europe. The customers in Europe must have access to the same data as customers in the United States (US) and will not tolerate high application latency or stale data. The customers in Europe and the customers in the US need to write to the database. Both groups of customers need to see updates from the other group in real time.
Which solution will meet these requirements?

A. Create an Amazon Aurora MySQL replica of the RDS for MySQL DB instanc
B. Pause application writes to the RDS DB instanc
C. Promote the Aurora Replica to a standalone DB cluste
D. Reconfigure the application to use the Aurora database and resume write
E. Add eu-west-1 as a secondary Region to the 06 cluste
F. Enable write forwarding on the DB cluste
G. Deploy the application in eu-west-1. Configure the application to use the Aurora MySQL endpoint in eu- west-1.
H. Add a cross-Region replica in eu-west-1 for the RDS for MySQL DB instanc
I. Configure the replica to replicate write queries back to the primary DB instanc
J. Deploy the application in eu-west-1. Configure the application to use the RDS for MySQL endpoint in eu-west-1.
K. Copy the most recent snapshot from the RDS for MySQL DB instance to eu-west-1. Create a new RDS for MySQL DB instance in eu-west-1 from the snapsho
L. Configure MySQL logical replication fromus-east-1 to eu-west-1. Enable write forwarding on the DB cluste
M. Deploy the application in eu-west-1. Configure the application to use the RDS for MySQL endpoint in eu-west-1.
N. Convert the RDS for MySQL DB instance to an Amazon Aurora MySQL DB cluste
O. Add eu-west-1 as a secondary Region to the DB cluste
P. Enable write forwarding on the DB cluste
Q. Deploy the application in eu-west-1. Configure the application to use the Aurora MySQL endpoint in eu-west-1.

Answer: D

Explanation:
The company should use AWS Amplify to create a static website for uploads of media files. The company should use Amplify Hosting to serve the website through Amazon CloudFront. The company should use Amazon S3 to store the uploaded media files. The company should use Amazon Cognito to authenticate users.
This solution will meet the requirements with the least operational overhead because AWS Amplify is a complete solution that lets frontend web and mobile developers easily build, ship, and host full-stack applications on AWS, with the flexibility to leverage the breadth of AWS services as use cases evolve. No cloud expertise needed1. By using AWS Amplify, the company can refactor the application to a serverless architecture that reduces operational complexity and costs.
AWS Amplify offers the following features and benefits:

- Amplify Studio: A visual interface that enables you to build and deploy a full-stack app quickly, including frontend UI and backend.
- Amplify CLI: A local toolchain that enables you to configure and manage an app backend with just a few commands.
- Amplify Libraries: Open-source client libraries that enable you to build cloud-powered mobile and web apps.
- Amplify UI Components: Open-source design system with cloud-connected components for building feature-rich apps fast.
- Amplify Hosting: Fully managed CI/CD and hosting for fast, secure, and reliable static and server-side rendered apps.

By using AWS Amplify to create a static website for uploads of media files, the company can leverage Amplify Studio to visually build a pixel-perfect UI and connect it to a cloud backend in clicks. By using Amplify Hosting to serve the website through Amazon CloudFront, the company can easily deploy its web app or website to the fast, secure, and reliable AWS content delivery network (CDN), with hundreds of points of presence globally. By using Amazon S3 to store the uploaded media files, the company can benefit from a highly scalable, durable, and cost-effective object storage service that can handle any amount of data2. By using Amazon Cognito to authenticate users, the company can add user sign-up, sign-in, and access control to its web app with a fully managed service that scales to support millions of users3.
The other options are not correct because:

- Using AWS Application Migration Service to migrate the application server to Amazon EC2 instances would not refactor the application or accelerate development. AWS Application Migration Service (AWS MGN) is a service that enables you to migrate physical servers, virtual machines (VMs), or cloud servers from any source infrastructure to AWS without requiring agents or specialized tools. However, this would not address the challenges of overutilization and data uploads failures. It would also not reduce operational overhead or costs compared to a serverless architecture.
- Creating a static website for uploads of media files and using AWS AppSync to create an API would not be as simple or fast as using AWS Amplify. AWS AppSync is a service that enables you to create flexible APIs for securely accessing, manipulating, and combining data from one or more data sources. However, this would require more configuration and management than using Amplify Studio and Amplify Hosting. It would also not provide authentication features like Amazon Cognito.

- Setting up AWS IAM Identity Center (AWS Single Sign-On) to give users the ability to sign in to the application would not be as suitable as using Amazon Cognito. AWS Single Sign-On (AWS SSO) is a
service that enables you to centrally manage SSO access and user permissions across multiple AWS accounts and business applications. However, this service is designed for enterprise customers who need to manage access for employees or partners across multiple resources. It is not intended for authenticating end users of web or mobile apps.

References:

- https://aws.amazon.com/amplify/
- https://aws.amazon.com/s3/
- https://aws.amazon.com/cognito/
- https://aws.amazon.com/mgn/
- https://aws.amazon.com/appsync/
- https://aws.amazon.com/single-sign-on/

NEW QUESTION 3
- (Exam Topic 2)
A company needs to optimize the cost of backups for Amazon Elastic File System (Amazon EFS). A solutions architect has already configured a backup plan in AWS Backup for the EFS backups. The backup plan contains a rule with a lifecycle configuration to transition EFS backups to cold storage after 7 days and to keep the backups for an additional 90 days.
After I month, the company reviews its EFS storage costs and notices an increase in the EFS backup costs. The EFS backup cold storage produces almost double the cost of the EFS warm backup storage.
What should the solutions architect do to optimize the cost?

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

Answer: A

Explanation:
The cost of EFS backup cold storage is $0.01 per GB-month, whereas the cost of EFS backup warm storage is $0.05 per GB-month1. Therefore, moving the backups to cold storage as soon as possible will reduce the storage cost. However, cold storage backups must be retained for a minimum of 90 day2s, otherwise they incur a pro-rated charge equal to the storage charge for the remaining days1. Therefore, setting the backup retention period to 30 days will incur a penalty of 60 days of cold storage cost for each backup deleted. This penalty will still be lower than keeping the backups in warm storage for 7 days and then in cold storage for 83 days, which is the current configuration. Therefore, option A is the most cost-effective solution.

NEW QUESTION 4
- (Exam Topic 2)
A company is running a two-tier web-based application in an on-premises data center. The application layer consists of a single server running a stateful application. The application connects to a PostgreSQL database running on a separate server. The application's user base is expected to grow significantly, so the company is migrating the application and database to AWS. The solution will use Amazon Aurora PostgreSQL, Amazon EC2 Auto Scaling, and Elastic Load Balancing.
Which solution will provide a consistent user experience that will allow the application and database tiers to scale?

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

Answer: C

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

NEW QUESTION 5
- (Exam Topic 2)
A company uses a load balancer to distribute traffic to Amazon EC2 instances in a single Availability Zone. The company is concerned about security and wants a solutions architect to re-architect the solution to meet the following requirements:
• Inbound requests must be filtered for common vulnerability attacks.
• Rejected requests must be sent to a third-party auditing application.
• All resources should be highly available. Which solution meets these requirements?

A. Configure a Multi-AZ Auto Scaling group using the application's AM
B. Create an Application Load Balancer (ALB) and select the previously created Auto Scaling group as the targe
C. Use Amazon Inspector to monitor traffic to the ALB and EC2 instance
D. Create a web ACL in WA
E. Create an AWS WAF using the web ACL and AL
F. Use an AWS Lambda function to frequently push the Amazon Inspector report to the third-party auditing application.
G. Configure an Application Load Balancer (ALB) and add the EC2 instances as targets Create a web ACL in WA
H. Create an AWS WAF using the web ACL and ALB name and enable logging with Amazon CloudWatch Log
I. Use an AWS Lambda function to frequently push the logs to the third-party auditing application.
J. Configure an Application Load Balancer (ALB) along with a target group adding the EC2 instances as target

K. Create an Amazon Kinesis Data Firehose with the destination of the third-party auditing applicatio
L. Create a web ACL in WA
M. Create an AWS WAF using the web ACL and ALB then enable logging by selecting the Kinesis Data Firehose as the destinatio
N. Subscribe to AWS Managed Rules in AWS Marketplace, choosing the WAF as the subscriber.
O. Configure a Multi-AZ Auto Scaling group using the application's AM
P. Create an Application Load Balancer (ALB) and select the previously created Auto Scaling group as the targe
Q. Create an Amazon Kinesis Data Firehose with a destination of the third-party auditing applicatio
R. Create a web ACL inWA
S. Create an AWS WAF using the WebACL and ALB then enable logging by selecting the Kinesis Data Firehose as the destinatio
T. Subscribe to AWS Managed Rules in AWS Marketplace, choosing the WAF as the subscriber.

Answer: D

Explanation:
https://docs.aws.amazon.com/waf/latest/developerguide/marketplace-managed-rule-groups.html

NEW QUESTION 6
- (Exam Topic 2)
A company built an application based on AWS Lambda deployed in an AWS CloudFormation stack. The last production release of the web application introduced an issue that resulted in an outage lasting several minutes. A solutions architect must adjust the deployment process to support a canary release.
Which solution will meet these requirements?

A. Create an alias for every new deployed version of the Lambda functio
B. Use the AWS CLI update-alias command with the routing-config parameter to distribute the load.
C. Deploy the application into a new CloudFormation stac
D. Use an Amazon Route 53 weighted routing policy to distribute the load.
E. Create a version for every new deployed Lambda functio
F. Use the AWS CLIupdate-function-contiguration command with the routing-config parameter to distribute the load.
G. Configure AWS CodeDeploy and use CodeDeployDefault.OneAtATime in the Deployment configuration to distribute the load.

Answer: A

Explanation:
https://aws.amazon.com/blogs/compute/implementing-canary-deployments-of-aws-lambda-functions-with-alias

NEW QUESTION 7
- (Exam Topic 2)
A company has loT sensors that monitor traffic patterns throughout a large city. The company wants to read and collect data from the sensors and perform aggregations on the data.
A solutions architect designs a solution in which the loT devices are streaming to Amazon Kinesis Data Streams. Several applications are reading from the stream. However, several consumers are experiencing throttling and are periodically and are periodically encountering a RealProvisioned Throughput Exceeded error.
Which actions should the solution architect take to resolve this issue? (Select THREE.)

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

Answer: ACE

Explanation:
https://repost.aws/knowledge-center/kinesis-readprovisionedthroughputexceeded Follow Data Streams best practices
To mitigate ReadProvisionedThroughputExceeded exceptions, apply these best practices:
• Reshard your stream to increase the number of shards in the stream.
• Use consumers with enhanced fan-out. For more information about enhanced fan-out, see Developing custom consumers with dedicated throughput (enhanced fan-out).
• Use an error retry and exponential backoff mechanism in the consumer logic if ReadProvisionedThroughputExceeded exceptions are encountered. For consumer applications that use an AWS SDK, the requests are retried by default.

NEW QUESTION 8
- (Exam Topic 2)
A company wants to refactor its retail ordering web application that currently has a load-balanced Amazon EC2 instance fleet for web hosting, database API services, and business logic. The company needs to create a decoupled, scalable architecture with a mechanism for retaining failed orders while also minimizing operational costs.
Which solution will meet these requirements?

A. Use Amazon S3 for web hosting with Amazon API Gateway for database API service
B. Use Amazon Simple Queue Service (Amazon SQS) for order queuin
C. Use Amazon Elastic Container Service (Amazon ECS) for business logic with Amazon SQS long polling for retaining failed orders.
D. Use AWS Elastic Beanstalk for web hosting with Amazon API Gateway for database API service
E. Use Amazon MQ for order queuin
F. Use AWS Step Functionsfor business logic with Amazon S3 Glacier Deep Archive for retaining failed orders.
G. Use Amazon S3 for web hosting with AWS AppSync for database API service
H. Use Amazon Simple Queue Service (Amazon SQS) for order queuin
I. Use AWS Lambda for business logic with an Amazon SQS dead-letter queue for retaining failed orders.
J. Use Amazon Lightsail for web hosting with AWS AppSync for database API service
K. Use Amazon Simple Email Service (Amazon SES) for order queuin
L. UseAmazon Elastic Kubernetes Service (Amazon EKS) for business logic with Amazon OpenSearch Service for retaining failed orders.

Answer: C

Explanation:
•Use Amazon S3 for web hosting with AWS AppSync for database API services. Use Amazon Simple Queue Service (Amazon SQS) for order queuing. Use AWS Lambda for business logic with an Amazon SQS dead-letter queue for retaining failed orders.
This solution will allow you to:
•Host a static website on Amazon S3 without provisioning or managing servers1.
•Use AWS AppSync to create a scalable GraphQL API that connects to your database and other data sources1.
•Use Amazon SQS to decouple and scale your order processing microservices1.
•Use AWS Lambda to run code for your business logic without provisioning or managing servers1.
•Use an Amazon SQS dead-letter queue to retain messages that can't be processed by your Lambda function1.

NEW QUESTION 9
- (Exam Topic 2)
A company wants to use AWS for disaster recovery for an on-premises application. The company has hundreds of Windows-based servers that run the application. All the servers mount a common share.
The company has an RTO of 15 minutes and an RPO of 5 minutes. The solution must support native failover and fallback capabilities.
Which solution will meet these requirements MOST cost-effectively?

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

Answer: D

NEW QUESTION 10
- (Exam Topic 2)
A company is using AWS CloudFormation to deploy its infrastructure. The company is concerned that, if a production CloudFormation stack is deleted, important data stored in Amazon RDS databases or Amazon EBS volumes might also be deleted.
How can the company prevent users from accidentally deleting data in this way?

A. Modify the CloudFormation templates to add a DeletionPolicy attribute to RDS and EBS resources.
B. Configure a stack policy that disallows the deletion of RDS and EBS resources.
C. Modify 1AM policies to deny deleting RDS and EBS resources that are tagged with an "awsrcloudformation: stack-name" tag.
D. Use AWS Config rules to prevent deleting RDS and EBS resources.

Answer: A

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

NEW QUESTION 10
- (Exam Topic 2)
A company has an application that runs on Amazon EC2 instances in an Amazon EC2 Auto Scaling group. The company uses AWS CodePipeline to deploy the application. The instances that run in the Auto Scaling group are constantly changing because of scaling events.
When the company deploys new application code versions, the company installs the AWS CodeDeploy agent on any new target EC2 instances and associates the instances with the CodeDeploy deployment group. The application is set to go live within the next 24 hours.
What should a solutions architect recommend to automate the application deployment process with the LEAST amount of operational overhead?

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

Answer: D

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

NEW QUESTION 13
- (Exam Topic 2)
A company is storing sensitive data in an Amazon S3 bucket. The company must log all activities for objects in the S3 bucket and must keep the logs for 5 years. The company's security team also must receive an email notification every time there is an attempt to delete data in the S3 bucket.
Which combination of steps will meet these requirements MOST cost-effectively? (Select THREE.)

A. Configure AWS CloudTrail to log S3 data events.
B. Configure S3 server access logging for the S3 bucket.
C. Configure Amazon S3 to send object deletion events to Amazon Simple Email Service (Amazon SES).
D. Configure Amazon S3 to send object deletion events to an Amazon EventBridge event bus that publishes to an Amazon Simple Notification Service (Amazon SNS) topic.
E. Configure Amazon S3 to send the logs to Amazon Timestream with data storage tiering.
F. Configure a new S3 bucket to store the logs with an S3 Lifecycle policy.

Answer: ADF

Explanation:
Configuring AWS CloudTrail to log S3 data events will enable logging all activities for objects in the S3 bucket1. Data events are object-level API operations such as GetObject, DeleteObject, and PutObject1. Configuring Amazon S3 to send object deletion events to an Amazon EventBridge event bus that publishes to an Amazon Simple Notification Service (Amazon SNS) topic will enable sending email notifications every time there is an attempt to delete data in the S3 bucket2. EventBridge can route events from S3 to SNS, which can send emails to subscribers2. Configuring a new S3 bucket to store the logs with an S3 Lifecycle policy will enable keeping the logs for 5 years in a cost-effective way3. A lifecycle policy can transition the logs to a cheaper storage class such as Glacier or delete them after a specified period of time3.

NEW QUESTION 17
- (Exam Topic 2)
A company is building a call center by using Amazon Connect. The company’s operations team is defining a disaster recovery (DR) strategy across AWS Regions. The contact center has dozens of contact flows, hundreds of users, and dozens of claimed phone numbers.
Which solution will provide DR with the LOWEST RTO?

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

Answer: D

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

NEW QUESTION 22
- (Exam Topic 2)
A financial services company loaded millions of historical stock trades into an Amazon DynamoDB table. The table uses on-demand capacity mode. Once each day at midnight, a few million new records are loaded into the table. Application read activity against the table happens in bursts throughout the day. and a limited set of keys are repeatedly looked up. The company needs to reduce costs associated with DynamoDB.
Which strategy should a solutions architect recommend to meet this requirement?

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

Answer: D

Explanation:
https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/HowItWorks.ReadWriteCapacityMode.h

NEW QUESTION 27
- (Exam Topic 2)
A company has developed a hybrid solution between its data center and AWS. The company uses Amazon VPC and Amazon EC2 instances that send application togs to Amazon CloudWatch. The EC2 instances read data from multiple relational databases that are hosted on premises.
The company wants to monitor which EC2 instances are connected to the databases in near-real time. The company already has a monitoring solution that uses Splunk on premises. A solutions architect needs to determine how to send networking traffic to Splunk.
How should the solutions architect meet these requirements?

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

Answer: B

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

NEW QUESTION 32
- (Exam Topic 2)
A company uses AWS Organizations with a single OU named Production to manage multiple accounts All accounts are members of the Production OU Administrators use deny list SCPs in the root of the organization to manage access to restricted services.
The company recently acquired a new business unit and invited the new unit's existing AWS account to the organization Once onboarded the administrators of the new business unit discovered that they are not able to update existing AWS Config rules to meet the company's policies.
Which option will allow administrators to make changes and continue to enforce the current policies without introducing additional long-term maintenance?

A. Remove the organization's root SCPs that limit access to AWS Config Create AWS Service Catalog products for the company's standard AWS Config rules and deploy them throughout the organization, including the new account.
B. Create a temporary OU named Onboarding for the new account Apply an SCP to the Onboarding OU to allow AWS Config actions Move the new account to the Production OU when adjustments to AWS Config are complete
C. Convert the organization's root SCPs from deny list SCPs to allow list SCPs to allow the required services only Temporarily apply an SCP to the organization's root that allows AWS Config actions forprincipals only in the new account.
D. Create a temporary OU named Onboarding for the new account Apply an SCP to the Onboarding OU to allow AWS Config action
E. Move the organization's root SCP to the Production O
F. Move the new account to the Production OU when adjustments to AWS Config are complete.

Answer: D

Explanation:
An SCP at a lower level can't add a permission after it is blocked by an SCP at a higher level. SCPs can only filter; they never add permissions. SO you need to create a new OU for the new account assign an SCP, and move the root SCP to Production OU. Then move the new account to production OU when AWS config is done.

NEW QUESTION 35
- (Exam Topic 2)
A company is running an application in the AWS Cloud. The application collects and stores a large amount of unstructured data in an Amazon S3 bucket. The S3 bucket contains several terabytes of data and uses the S3 Standard storage class. The data increases in size by several gigabytes every day.
The company needs to query and analyze the data. The company does not access data that is more than 1 year old. However, the company must retain all the data indefinitely for compliance reasons.
Which solution will meet these requirements MOST cost-effectively?

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

Answer: C

Explanation:
Generally, unstructured data should be converted structured data before querying them. AWS Glue can do that.
https://docs.aws.amazon.com/glue/latest/dg/schema-relationalize.html https://docs.aws.amazon.com/athena/latest/ug/glue-athena.html

NEW QUESTION 37
- (Exam Topic 2)
A solutions architect at a large company needs to set up network security tor outbound traffic to the internet from all AWS accounts within an organization in AWS Organizations. The organization has more than 100 AWS accounts, and the accounts route to each other by using a centralized AWS Transit Gateway. Each account has both an internet gateway and a NAT gateway tor outbound traffic to the internet The company deploys resources only into a single AWS Region. The company needs the ability to add centrally managed rule-based filtering on all outbound traffic to the internet for all AWS accounts in the organization. The peak load of outbound traffic will not exceed 25 Gbps in each Availability Zone.
Which solution meets these requirements?

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

Answer: B

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

NEW QUESTION 39
- (Exam Topic 2)
A large company runs workloads in VPCs that are deployed across hundreds of AWS accounts. Each VPC consists to public subnets and private subnets that span across multiple Availability Zones. NAT gateways are deployed in the public subnets and allow outbound connectivity to the internet from the private subnets. A solutions architect is working on a hub-and-spoke design. All private subnets in the spoke VPCs must route traffic to the internet through an egress VPC. The solutions architect already has deployed a NAT gateway in an egress VPC in a central AWS account.
Which set of additional steps should the solutions architect take to meet these requirements?

A. Create peering connections between the egress VPC and the spoke VPC
B. Configure the required routing to allow access to the internet.
C. Create a transit gateway, and share it with the existing AWS account
D. Attach existing VPCs to the transit gateway Configure the required routing to allow access to the internet.
E. Create a transit gateway in every accoun
F. Attach the NAT gateway to the transit gateway
G. Configure the required routing to allow access to the internet.
H. Create an AWS PrivateLink connection between the egress VPC and the spoke VPC
I. Configure the required routing to allow access to the internet

Answer: B

Explanation:
https://d1.awsstatic.com/architecture-diagrams/ArchitectureDiagrams/NAT-gateway-centralized-egress-ra.pdf?d

NEW QUESTION 44
- (Exam Topic 2)
A company needs to migrate its customer transactions database from on premises to AWS. The database resides on an Oracle DB instance that runs on a Linux server. According to a new security requirement, the company must rotate the database password each year.
Which solution will meet these requirements with the LEAST operational overhead?

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

Answer: B

NEW QUESTION 49
- (Exam Topic 2)

A company is implementing a serverless architecture by using AWS Lambda functions that need to access a Microsoft SQL Server DB instance on Amazon RDS.
The company has separate environments for development and production, including a clone of the database system.
The company's developers are allowed to access the credentials for the development database. However, the credentials for the production database must be encrypted with a key that only members of the IT security team's IAM user group can access. This key must be rotated on a regular basis.
What should a solutions architect do in the production environment to meet these requirements?

A. Store the database credentials in AWS Systems Manager Parameter Store by using a SecureString parameter that is encrypted by an AWS Key Management Service (AWS KMS) customer managed ke
B. Attach a role to each Lambda function to provide access to the SecureString paramete
C. Restrict access to the Securestring parameter and the customer managed key so that only the IT security team can access the parameter and the key.
D. Encrypt the database credentials by using the AWS Key Management Service (AWS KMS) default Lambda ke
E. Store the credentials in the environment variables of each Lambda functio
F. Load the credentials from the environment variables in the Lambda cod
G. Restrict access to the KMS key o that only the IT security team can access the key.
H. Store the database credentials in the environment variables of each Lambda functio
I. Encrypt the environment variables by using an AWS Key Management Service (AWS KMS) customer managed ke
J. Restrict access to the customer managed key so that only the IT security team can access the key.
K. Store the database credentials in AWS Secrets Manager as a secret that is associated with an AWS Key Management Service (AWS KMS) customermanaged ke
L. Attach a role to each Lambda function to provide access to the secre
M. Restrict access to the secret and the customer managed key so that only the IT security team can access the secret and the key.

Answer: D

Explanation:
Storing the database credentials in AWS Secrets Manager as a secret that is associated with an AWS Key Management Service (AWS KMS) customer managed key will enable encrypting and managing the credentials securely1. AWS Secrets Manager helps you to securely encrypt, store, and retrieve credentials for your databases and other services2. Attaching a role to each Lambda function to provide access to the secret will enable retrieving the credentials programmatically1. Restricting access to the secret and the customer managed key so that only members of the IT security team's IAM user group can access them will enable meeting the security requirements1.

NEW QUESTION 52
- (Exam Topic 2)
A company has a few AWS accounts for development and wants to move its production application to AWS. The company needs to enforce Amazon Elastic Block Store (Amazon EBS) encryption at rest current production accounts and future production accounts only. The company needs a solution that includes built-in blueprints and guardrails.
Which combination of steps will meet these requirements? (Choose three.)

A. Use AWS CloudFormation StackSets to deploy AWS Config rules on production accounts.
B. Create a new AWS Control Tower landing zone in an existing developer accoun
C. Create OUs for account
D. Add production and development accounts to production and development OUs, respectively.
E. Create a new AWS Control Tower landing zone in the company's management accoun
F. Add production and development accounts to production and development OU
G. respectively.
H. Invite existing accounts to join the organization in AWS Organization
I. Create SCPs to ensure compliance.
J. Create a guardrail from the management account to detect EBS encryption.
K. Create a guardrail for the production OU to detect EBS encryption.

Answer: CDF

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

NEW QUESTION 55
- (Exam Topic 2)
A company is using AWS Organizations to manage multiple AWS accounts. For security purposes, the company requires the creation of an Amazon Simple Notification Service (Amazon SNS) topic that enables integration with a third-party alerting system in all the Organizations member accounts.
A solutions architect used an AWS CloudFormation template to create the SNS topic and stack sets to automate the deployment of Cloud Formation stacks.
Trusted access has been enabled in Organizations.
What should the solutions architect do to deploy the CloudFormation StackSets in all AWS accounts?

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

Answer: C

Explanation:
https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/stacksets-orgs-manage-auto-deployment.h

NEW QUESTION 58
- (Exam Topic 2)
A software-as-a-service (SaaS) provider exposes APIs through an Application Load Balancer (ALB). The ALB connects to an Amazon Elastic Kubernetes Service (Amazon EKS) cluster that is deployed in the
us-east-I Region. The exposed APIs contain usage of a few non-standard REST methods: LINK, UNLINK, LOCK, and UNLOCK.
Users outside the United States are reporting long and inconsistent response times for these APIs. A solutions architect needs to resolve this problem with a solution that minimizes operational overhead.
Which solution meets these requirements?

A. Add an Amazon CloudFront distributio
B. Configure the ALB as the origin.
C. Add an Amazon API Gateway edge-optimized API endpoint to expose the API
D. Configure the ALB as the target.
E. Add an accelerator in AWS Global Accelerato
F. Configure the ALB as the origin.
G. Deploy the APIs to two additional AWS Regions: eu-west-l and ap-southeast-2. Add latency-based routing records in Amazon Route 53.

Answer: C

Explanation:
Adding an accelerator in AWS Global Accelerator will enable improving the performance of the APIs for local and global users1. AWS Global Accelerator is a service that uses the AWS global network to route traffic to the optimal regional endpoint based on health, client location, and policies1. Configuring the ALB as the origin will enable connecting the accelerator to the ALB that exposes the APIs2. AWS Global Accelerator supports non-standard REST methods such as LINK, UNLINK, LOCK, and UNLOCK3.

NEW QUESTION 60
- (Exam Topic 2)
A solutions architect must create a business case for migration of a company's on-premises data center to the AWS Cloud. The solutions architect will use a configuration management database (CMDB) export of all the company's servers to create the case.
Which solution will meet these requirements MOST cost-effectively?

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

Answer: B

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

NEW QUESTION 64
- (Exam Topic 2)
A company is running a critical stateful web application on two Linux Amazon EC2 instances behind an Application Load Balancer (ALB) with an Amazon RDS for MySQL database The company hosts the DNS records for the application in Amazon Route 53 A solutions architect must recommend a solution to improve the resiliency of the application
The solution must meet the following objectives:
• Application tier RPO of 2 minutes. RTO of 30 minutes
• Database tier RPO of 5 minutes RTO of 30 minutes
The company does not want to make significant changes to the existing application architecture The company must ensure optimal latency after a failover
Which solution will meet these requirements?

A. Configure the EC2 instances to use AWS Elastic Disaster Recovery Create a cross-Region read replica for the RDS DB instance Create an ALB in a second AWS Region Create an AWS Global Accelerator endpoint and associate the endpoint with the ALBs Update DNS records to point to the Global Accelerator endpoint
B. Configure the EC2 instances to use Amazon Data Lifecycle Manager (Amazon DLM) to take snapshots of the EBS volumes Configure RDS automated backups Configure backup replication to a second AWS Region Create an ALB in the second Region Create an AWS Global Accelerator endpoint, and associate the endpoint with the ALBs Update DNS records to point to the Global Accelerator endpoint
C. Create a backup plan in AWS Backup for the EC2 instances and RDS DB instance Configure backup replication to a second AWS Region Create an ALB in the second Region Configure an Amazon CloudFront distribution in front of the ALB Update DNS records to point to CloudFront
D. Configure the EC2 instances to use Amazon Data Lifecycle Manager (Amazon DLM) to take snapshots of the EBS volumes Create a cross-Region read replica for the RDS DB instance Create an ALB in a second AWS Region Create an AWS Global Accelerator endpoint and associate the endpoint with the ALBs

Answer: B

Explanation:
This option meets the RPO and RTO requirements for both the application and database tiers and uses tools like Amazon DLM and RDS automated backups to

create and manage the backups. Additionally, it uses Global Accelerator to ensure low latency after failover by directing traffic to the closest healthy endpoint.

NEW QUESTION 66
- (Exam Topic 2)
A company hosts a blog post application on AWS using Amazon API Gateway, Amazon DynamoDB, and AWS Lambda. The application currently does not use API keys to authorize requests. The API model is as follows: GET/posts/[postid] to get post details GET/users[userid] to get user details GET/comments/[commentid] to get comments details
The company has noticed users are actively discussing topics in the comments section, and the company wants to increase user engagement by marking the comments appears in real time.
Which design should be used to reduce comment latency and improve user experience?

A. Use edge-optimized API with Amazon CloudFront to cache API responses.
B. Modify the blog application code to request GET comment[commented] every 10 seconds.
C. Use AWS AppSync and leverage WebSockets to deliver comments.
D. Change the concurrency limit of the Lambda functions to lower the API response time.

Answer: C

Explanation:
https://docs.aws.amazon.com/appsync/latest/devguide/graphql-overview.html
AWS AppSync is a fully managed GraphQL service that allows applications to securely access, manipulate, and receive data as well as real-time updates from multiple data sources1. AWS AppSync supports GraphQL subscriptions to perform real-time operations and can push data to clients that choose to listen to specific events from the backend1. AWS AppSync uses WebSockets to establish and maintain a secure connection between the clients and the API endpoint2. Therefore, using AWS AppSync and leveraging WebSockets is a suitable design to reduce comment latency and improve user experience.

NEW QUESTION 70
- (Exam Topic 2)
A company is creating a REST API to share information with six of its partners based in the United States. The company has created an Amazon API Gateway Regional endpoint. Each of the six partners will access the API once per day to post daily sales figures.
After initial deployment, the company observes 1.000 requests per second originating from 500 different IP addresses around the world. The company believes this traffic is originating from a botnet and wants to secure its API while minimizing cost.
Which approach should the company take to secure its API?

A. Create an Amazon CloudFront distribution with the API as the origi
B. Create an AWS WAF web ACL with a rule lo block clients thai submit more than fiverequests per da
C. Associate the web ACL with the CloudFront distnbutio
D. Configure CloudFront with an origin access identity (OAI) and associate it with the distributio
E. Configure API Gateway to ensure only the OAI can run the POST method.
F. Create an Amazon CloudFront distribution with the API as the origi
G. Create an AWS WAF web ACL with a rule to block clients that submit more than five requests per da
H. Associate the web ACL with the CloudFront distnbutio
I. Add a custom header to the CloudFront distribution populated with an API ke
J. Configure the API to require an API key on the POST method.
K. Create an AWS WAF web ACL with a rule to allow access to the IP addresses used by the six partners.Associate the web ACL with the AP
L. Create a resource policy with a request limit and associate it with the AP
M. Configure the API to require an API key on the POST method.
N. Create an AWS WAF web ACL with a rule to allow access to the IP addresses used by the six partners.Associate the web ACL with the AP
O. Create a usage plan with a request limit and associate it with the AP
P. Create an API key and add it to the usage plan.

Answer: D

Explanation:
"A usage plan specifies who can access one or more deployed API stages and methods—and also how much and how fast they can access them. The plan uses API keys to identify API clients and meters access to the associated API stages for each key. It also lets you configure throttling limits and quota limits that are enforced on individual client API keys."
https://docs.aws.amazon.com/apigateway/latest/developerguide/api-gateway-api-usage-plans.html
A rate-based rule tracks the rate of requests for each originating IP address, and triggers the rule action on IPs with rates that go over a limit. You set the limit as the number of requests per 5-minute time span..... The following caveats apply to AWS WAF rate-based rules: The minimum rate that you can set is 100. AWS WAF checks the rate of requests every 30 seconds, and counts requests for the prior five minutes each time. Because of this, it's possible for an IP address to send requests at too high a rate for 30 seconds before AWS WAF detects and blocks it. AWS WAF can block up to 10,000 IP addresses. If more than 10,000 IP addresses send high rates of requests at the same time, AWS WAF will only block 10,000 of them. " https://docs.aws.amazon.com/waf/latest/developerguide/waf-rule-statement-type-rate-based.html

NEW QUESTION 72
- (Exam Topic 2)
A company has multiple business units that each have separate accounts on AWS. Each business unit manages its own network with several VPCs that have CIDR ranges that overlap. The company’s marketing team has created a new internal application and wants to make the application accessible to all the other business units. The solution must use private IP addresses only.
Which solution will meet these requirements with the LEAST operational overhead?

A. Instruct each business unit to add a unique secondary CIDR range to the business unit's VP
B. Peer the VPCs and use a private NAT gateway in the secondary range to route traffic to the marketing team.
C. Create an Amazon EC2 instance to serve as a virtual appliance in the marketing account's VP
D. Create an AWS Site-to-Site VPN connection between the marketing team and each business unit's VP
E. Perform NAT where necessary.
F. Create an AWS PrivateLink endpoint service to share the marketing applicatio
G. Grant permission to specific AWS accounts to connect to the servic
H. Create interface VPC endpoints in other accounts to access the application by using private IP addresses.
I. Create a Network Load Balancer (NLB) in front of the marketing application in a private subne

J. Create an API Gateway AP
K. Use the Amazon API Gateway private integration to connect the API to the NL
L. Activate IAM authorization for the AP
M. Grant access to the accounts of the other business units.

Answer: C

Explanation:
With AWS PrivateLink, the marketing team can create an endpoint service to share their internal application with other accounts securely using private IP addresses. They can grant permission to specific AWS accounts to connect to the service and create interface VPC endpoints in the other accounts to access the application by using private IP addresses. This option does not require any changes to the network of the other business units, and it does not require peering or NATing. This solution is both scalable and secure.
https://aws.amazon.com/blogs/networking-and-content-delivery/connecting-networks-with-overlapping-ip-range

NEW QUESTION 77
- (Exam Topic 2)
A company is running a containerized application in the AWS Cloud. The application is running by using Amazon Elastic Container Service (Amazon ECS) on a set of Amazon EC2 instances. The EC2 instances run in an Auto Scaling group.
The company uses Amazon Elastic Container Registry (Amazon ECR) to store its container images. When a new image version is uploaded, the new image version receives a unique tag.
The company needs a solution that inspects new image versions for common vulnerabilities and exposures. The solution must automatically delete new image tags that have Critical or High severity findings. The solution also must notify the development team when such a deletion occurs.
Which solution meets these requirements?

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

Answer: A

Explanation:
https://docs.aws.amazon.com/AmazonECR/latest/userguide/ecr-eventbridge.html "Activating an AWS Step Functions state machine"
https://docs.aws.amazon.com/step-functions/latest/dg/tutorial-creating-lambda-state-machine.html

NEW QUESTION 78
- (Exam Topic 2)
A company has a critical application in which the data tier is deployed in a single AWS Region. The data tier uses an Amazon DynamoDB table and an Amazon Aurora MySQL DB cluster. The current Aurora MySQL engine version supports a global database. The application tier is already deployed in two Regions.
Company policy states that critical applications must have application tier components and data tier components deployed across two Regions. The RTO and RPO must be no more than a few minutes each. A solutions architect must recommend a solution to make the data tier compliant with company policy.
Which combination of steps will meet these requirements? (Choose two.)

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

Answer: AD

Explanation:
The company should use Amazon Aurora global database and Amazon DynamoDB global table to deploy the data tier components across two Regions. Amazon Aurora global database is a feature that allows a single
Aurora database to span multiple AWS Regions, enabling low-latency global reads and fast recovery from Region-wide outages1. Amazon DynamoDB global table is a feature that allows a single DynamoDB table to span multiple AWS Regions, enabling low-latency global reads and writes and fast recovery from Region-wide outages2.
References:
- https://aws.amazon.com/rds/aurora/global-database/
- https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/globaltables_HowItWorks.html
- https://aws.amazon.com/route53/application-recovery-controller/

NEW QUESTION 81
- (Exam Topic 2)
A company has migrated a legacy application to the AWS Cloud. The application runs on three Amazon EC2 instances that are spread across three Availability Zones. One EC2 instance is in each Availability Zone. The EC2 instances are running in three private subnets of the VPC and are set up as targets for an

Application Load Balancer (ALB) that is associated with three public subnets.
The application needs to communicate with on-premises systems. Only traffic from IP addresses in the company's IP address range are allowed to access the on-premises systems. The company's security team is bringing only one IP address from its internal IP address range to the cloud. The company has added this IP address to the allow list for the company firewall. The company also has created an Elastic IP address for this IP address.
A solutions architect needs to create a solution that gives the application the ability to communicate with the on-premises systems. The solution also must be able to mitigate failures automatically.
Which solution will meet these requirements?

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

Answer: C

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

NEW QUESTION 84
- (Exam Topic 2)
A company operates a proxy server on a fleet of Amazon EC2 instances. Partners in different countries use the proxy server to test the company's functionality. The EC2 instances are running in a VPC. and the instances have access to the internet.
The company's security policy requires that partners can access resources only from domains that the company owns.
Which solution will meet these requirements?

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

Answer: A

Explanation:
The company should create an Amazon Route 53 Resolver DNS Firewall domain list that contains the allowed domains. The company should configure a DNS Firewall rule group with a rule that has a high numeric value that blocks all requests. The company should configure a rule that has a low numeric value that allows requests for domains in the allowed list. The company should associate the rule group with the VPC. This solution will meet the requirements because Amazon Route 53 Resolver DNS Firewall is a feature that enables you to filter and regulate outbound DNS traffic for your VPC. You can create reusable collections of filtering rules in DNS Firewall rule groups and associate them with your VPCs. You can specify lists of domain names to allow or block, and you can customize the responses for the DNS queries that you block1. By creating a domain list with the allowed domains and a rule group with rules to allow or block requests based on the domain list, the company can enforce its security policy and control access to sites.
The other options are not correct because:

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

References:

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

NEW QUESTION 88
- (Exam Topic 2)
A company is running a web application in a VPC. The web application runs on a group of Amazon EC2 instances behind an Application Load Balancer (ALB). The ALB is using AWS WAF.
An external customer needs to connect to the web application. The company must provide IP addresses to all external customers.
Which solution will meet these requirements with the LEAST operational overhead?

A. Replace the ALB with a Network Load Balancer (NLB). Assign an Elastic IP address to the NLB.
B. Allocate an Elastic IP addres
C. Assign the Elastic IP address to the ALProvide the Elastic IP address to the customer.
D. Create an AWS Global Accelerator standard accelerato
E. Specify the ALB as the accelerator's endpoint.Provide the accelerator's IP addresses to the customer.
F. Configure an Amazon CloudFront distributio
G. Set the ALB as the origi
H. Ping the distribution's DNS name to determine the distribution's public IP addres
I. Provide the IP address to the customer.

Answer: C

Explanation:
https://docs.aws.amazon.com/global-accelerator/latest/dg/about-accelerators.alb-accelerator.html Option A is wrong. AWS WAF does not support associating with NLB.
https://docs.aws.amazon.com/waf/latest/developerguide/waf-chapter.html Option B is wrong. An ALB does not support an Elastic IP address.
https://aws.amazon.com/elasticloadbalancing/features/

NEW QUESTION 91
- (Exam Topic 2)
A company uses a Grafana data visualization solution that runs on a single Amazon EC2 instance to monitor the health of the company's AWS workloads. The company has invested time and effort to create dashboards that the company wants to preserve. The dashboards need to be highly available and cannot be down for longer than 10 minutes. The company needs to minimize ongoing maintenance.
Which solution will meet these requirements with the LEAST operational overhead?

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

Answer: C

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

NEW QUESTION 94
- (Exam Topic 2)
A company runs an intranet application on premises. The company wants to configure a cloud backup of the application. The company has selected AWS Elastic Disaster Recovery for this solution.
The company requires that replication traffic does not travel through the public internet. The application also must not be accessible from the internet. The company does not want this solution to consume all available network bandwidth because other applications require bandwidth.
Which combination of steps will meet these requirements? (Select THREE.)

A. Create a VPC that has at least two private subnets, two NAT gateways, and a virtual private gateway.
B. Create a VPC that has at least two public subnets, a virtual private gateway, and an internet gateway.
C. Create an AWS Site-to-Site VPN connection between the on-premises network and the target AWS network.
D. Create an AWS Direct Connect connection and a Direct Connect gateway between the on-premises network and the target AWS network.
E. During configuration of the replication servers, select the option to use private IP addresses for data replication.
F. During configuration of the launch settings for the target servers, select the option to ensure that the Recovery instance's private IP address matches the source server's private IP address.

Answer: BDE

Explanation:
AWS Elastic Disaster Recovery (AWS DRS) is a service that minimizes downtime and data loss with fast, reliable recovery of on-premises and cloud-based applications using affordable storage, minimal compute, and point-in-time recovery1. Users can set up AWS DRS on their source servers to initiate secure data replication to a staging area subnet in their AWS account, in the AWS Region they select. Users can then launch recovery instances on AWS within minutes, using the most up-to-date server state or a previous point in time.
To configure a cloud backup of the application with AWS DRS, users need to create a VPC that has at least
two public subnets, a virtual private gateway, and an internet gateway. A VPC is a logically isolated section of the AWS Cloud where users can launch AWS resources in a virtual network that they define2. A public subnet is a subnet that has a route to an internet gateway3. A virtual private gateway is the VPN concentrator on the Amazon side of the Site-to-Site VPN connection4. An internet gateway is a horizontally scaled, redundant, and highly available VPC component that allows communication between instances in the VPC and the internet. Users need to create at least two public subnets for redundancy and high availability. Users need to create a virtual private gateway and attach it to the VPC to enable VPN connectivity between the on-premises network and the target

AWS network. Users need to create an internet gateway and attach it to the VPC to enable internet access for the replication servers.
To ensure that replication traffic does not travel through the public internet, users need to create an AWS Direct Connect connection and a Direct Connect gateway between the on-premises network and the target AWS network. AWS Direct Connect is a service that establishes a dedicated network connection from an on-premises network to one or more VPCs. A Direct Connect gateway is a globally available resource that allows users to connect multiple VPCs across different Regions to their on-premises networks using one or more Direct Connect connections. Users need to create an AWS Direct Connect connection between their on-premises network and an AWS Region. Users need to create a Direct Connect gateway and associate it with their VPC and their Direct Connect connection.
To ensure that the application is not accessible from the internet, users need to select the option to use private IP addresses for data replication during configuration of the replication servers. This option configures the replication servers with private IP addresses only, without assigning any public IP addresses or Elastic IP addresses. This way, the replication servers can only communicate with other resources within the VPC or through VPN connections.
Option A is incorrect because creating a VPC that has at least two private subnets, two NAT gateways, and a virtual private gateway is not necessary or cost-effective. A private subnet is a subnet that does not have a route to an internet gateway3. A NAT gateway 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 connections with those instances. Users do not need to create private subnets or NAT gateways for this use case, as they can use public subnets with private IP addresses for data replication.
Option C is incorrect because creating an AWS Site-to-Site VPN connection between the on-premises network and the target AWS network will not ensure that replication traffic does not travel through the public
internet. A Site-to-Site VPN connection consists of two VPN tunnels between an on-premises customer
gateway device and a virtual private gateway in your VPC4. The VPN tunnels are encrypted using IPSec protocols, but they still use public IP addresses for communication. Users need to use AWS Direct Connect instead of Site-to-Site VPN for this use case.
Option F is incorrect because selecting the option to ensure that the Recovery instance's private IP address matches the source server's private IP address during configuration of the launch settings for the target servers will not ensure that the application is not accessible from the internet. This option configures the Recovery instance with an identical private IP address as its source server when launched in drills or recovery mode. However, this option does not prevent assigning public IP addresses or Elastic IP addresses to the Recovery instance. Users need to select the option to use private IP addresses for data replication instead.

NEW QUESTION 96
- (Exam Topic 2)
A company is building a hybrid environment that includes servers in an on-premises data center and in the AWS Cloud. The company has deployed Amazon EC2 instances in three VPCs. Each VPC is in a different AWS Region. The company has established an AWS Direct Connect connection to the data center from the Region that is closest to the data center.
The company needs the servers in the on-premises data center to have access to the EC2 instances in all three VPCs. The servers in the on-premises data center also must have access to AWS public services.
Which combination of steps will meet these requirements with the LEAST cost? (Select TWO.)

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

Answer: AE

Explanation:
A Direct Connect gateway allows you to connect multiple VPCs across different Regions to a Direct Connect connection1. A public VIF allows you to access AWS public services such as EC21. A Site-to-Site VPN connection over the public VIF provides encryption and redundancy for the traffic between the on-premises data center and the VPCs2. This solution is cheaper than setting up additional Direct Connect connections or using a private VIF with VPC peering.

NEW QUESTION 100
- (Exam Topic 2)
A company has many separate AWS accounts and uses no central billing or management. Each AWS account hosts services for different departments in the company. The company has a Microsoft Azure Active Directory that is deployed.
A solution architect needs to centralize billing and management of the company's AWS accounts. The company wants to start using identify federation instead of manual user management. The company also wants to use temporary credentials instead of long-lived access keys.
Which combination of steps will meet these requirements? (Select THREE)

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

Answer: ACE

NEW QUESTION 104
- (Exam Topic 2)

A company needs to architect a hybrid DNS solution. This solution will use an Amazon Route 53 private hosted zone for the domain cloud.example.com for the resources stored within VPCs.
The company has the following DNS resolution requirements:
• On-premises systems should be able to resolve and connect to cloud.example.com.
• All VPCs should be able to resolve cloud.example.com.
There is already an AWS Direct Connect connection between the on-premises corporate network and AWS Transit Gateway. Which architecture should the company use to meet these requirements with the HIGHEST performance?

A. Associate the private hosted zone to all the VPC
B. Create a Route 53 inbound resolver in the shared services VP
C. Attach all VPCs to the transit gateway and create forwarding rules in the on-premises DNS server for cloud.example.com that point to the inbound resolver.
D. Associate the private hosted zone to all the VPC
E. Deploy an Amazon EC2 conditional forwarder in the shared services VP
F. Attach all VPCs to the transit gateway and create forwarding rules in theon-premises DNS server for cloud.example.com that point to the conditional forwarder.
G. Associate the private hosted zone to the shared services VP
H. Create a Route 53 outbound resolver in the shared services VP
I. Attach all VPCs to the transit gateway and create forwarding rules in theon-premises DNS server for cloud.example.com that point to the outbound resolver.
J. Associate the private hosted zone to the shared services VP
K. Create a Route 53 inbound resolver in the shared services VP
L. Attach the shared services VPC to the transit gateway and create forwarding rules in the on-premises DNS server for cloud.example.com that point to the inbound resolver.

Answer: A

Explanation:
Amazon Route 53 Resolver is a managed DNS resolver service from Route 53 that helps to create conditional forwarding rules to redirect query traffic1. By associating the private hosted zone to all the VPCs, the solutions architect can enable DNS resolution for cloud.example.com within the VPCs. By creating a Route 53 inbound resolver in the shared services VPC, the solutions architect can enable DNS resolution for cloud.example.com from on-premises systems. By attaching all VPCs to the transit gateway, the solutions architect can enable connectivity between the VPCs and the on-premises network through AWS Direct Connect. By creating forwarding rules in the on-premises DNS server for cloud.example.com that point to the inbound resolver, the solutions architect can direct DNS queries for cloud.example.com to the Route 53 Resolver endpoint in AWS. This solution will provide the highest performance as it leverages Route 53 Resolver's optimized routing and caching capabilities.
References: 1: https://aws.amazon.com/route53/resolver/

NEW QUESTION 107
- (Exam Topic 2)
A company operates an on-premises software-as-a-service (SaaS) solution that ingests several files daily. The company provides multiple public SFTP endpoints to its customers to facilitate the file transfers. The customers add the SFTP endpoint IP addresses to their firewall allow list for outbound traffic. Changes to the SFTP endmost IP addresses are not permitted.
The company wants to migrate the SaaS solution to AWS and decrease the operational overhead of the file transfer service.
Which solution meets these requirements?

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

Answer: A

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

NEW QUESTION 108
- (Exam Topic 2)
A company needs to establish a connection from its on-premises data center to AWS. The company needs to connect all of its VPCs that are located in different AWS Regions with transitive routing capabilities between VPC networks. The company also must reduce network outbound traffic costs, increase bandwidth throughput, and provide a consistent network experience for end users.
Which solution will meet these requirements?

A. Create an AWS Site-to-Site VPN connection between the on-premises data center and a new central VP
B. Create VPC peering connections that initiate from the central VPC to all other VPCs.
C. Create an AWS Direct Connect connection between the on-premises data center and AW
D. Provision a transit VIF, and connect it to a Direct Connect gatewa
E. Connect the Direct Connect gateway to all the other VPCs by using a transit gateway in each Region.
F. Create an AWS Site-to-Site VPN connection between the on-premises data center and a new central VP
G. Use a transit gateway with dynamic routin
H. Connect the transit gateway to all other VPCs.
I. Create an AWS Direct Connect connection between the on-premises data center and AWS Establish an AWS Site-to-Site VPN connection between all VPCs in each Regio

J. Create VPC peering connections that initiate from the central VPC to all other VPCs.

Answer: B

Explanation:
Transit GW + Direct Connect GW + Transit VIF + enabled SiteLink if two different DX locations https://aws.amazon.com/blogs/networking-and-content-delivery/introducing-aws-direct-connect-sitelink/

NEW QUESTION 111
- (Exam Topic 2)
A company has an application that runs as a ReplicaSet of multiple pods in an Amazon Elastic Kubernetes Service (Amazon EKS) cluster. The EKS cluster has nodes in multiple Availability Zones. The application generates many small files that must be accessible across all running instances of the application. The company needs to back up the files and retain the backups for 1 year.
Which solution will meet these requirements while providing the FASTEST storage performance?

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

Answer: A

Explanation:
In the past, EBS can be attached only to one ec2 instance but not anymore but there are limitations like - it works only on io1/io2 instance types and many others as described here. https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-volumes-multi.html EFS has shareable storage
In terms of performance, Amazon EFS is optimized for workloads that require high levels of aggregate throughput and IOPS, whereas EBS is optimized for low-latency, random access I/O operations. Amazon EFS is designed to scale throughput and capacity automatically as your storage needs grow, while EBS volumes can be resized on demand.

NEW QUESTION 115
- (Exam Topic 2)
A solutions architect wants to cost-optimize and appropriately size Amazon EC2 instances in a single AWS account. The solutions architect wants to ensure that the instances are optimized based on CPU, memory, and network metrics.
Which combination of steps should the solutions architect take to meet these requirements? (Choose two.)

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

Answer: BD

Explanation:
AWS Trusted Advisor is a service that provides real-time guidance to help users provision their resources following AWS best practices1. One of the Trusted Advisor checks is “Low Utilization Amazon EC2 Instances”, which identifies EC2 instances that appear to be underutilized based on CPU, network I/O, and disk I/O metrics1. This check can help users optimize the cost and size of their EC2 instances by recommending smaller or more appropriate instance types.
AWS Compute Optimizer is a service that analyzes the configuration and utilization metrics of AWS resources and generates optimization recommendations to reduce the cost and improve the performance of workloads2. Compute Optimizer supports four types of AWS resources: EC2 instances, EBS volumes, ECS services on AWS Fargate, and Lambda functions2. For EC2 instances, Compute Optimizer evaluates the vCPUs, memory, storage, and other specifications, as well as the CPU utilization, network in and out, disk read and write, and other utilization metrics of currently running instances3. It then recommends optimal instance types based on price-performance trade-offs.
Option A is incorrect because purchasing AWS Business Support or AWS Enterprise Support for the account will not directly help with cost-optimization and sizing of EC2 instances. However, these support plans do provide access to more Trusted Advisor checks than the basic support plan1.
Option C is incorrect because installing the Amazon CloudWatch agent and configuring memory metric collection on the EC2 instances will not provide any optimization recommendations by itself. However, memory metrics can be used by Compute Optimizer to enhance its recommendations if enabled3.
Option E is incorrect because creating an EC2 Instance Savings Plan for the AWS Regions, instance families, and operating systems of interest will not help with cost-optimization and sizing of EC2 instances. Savings Plans are a flexible pricing model that offer lower prices on Amazon EC2 usage in exchange for a commitment to a consistent amount of usage for a 1- or 3-year term4. Savings Plans do not affect the configuration or utilization of EC2 instances.

NEW QUESTION 120
- (Exam Topic 2)
A solutions architect must provide a secure way for a team of cloud engineers to use the AWS CLI to upload objects into an Amazon S3 bucket Each cloud engineer has an IAM user. IAM access keys and a virtual multi-factor authentication (MFA) device The IAM users for the cloud engineers are in a group that is named S3-access The cloud engineers must use MFA to perform any actions in Amazon S3
Which solution will meet these requirements?

A. Attach a policy to the S3 bucket to prompt the 1AM user for an MFA code when the 1AM user performs actions on the S3 bucket Use 1AM access keys with the AWS CLI tocall Amazon S3

B. Update the trust policy for the S3-access group to require principals to use MFA when principals assume the group Use 1AM access keys with the AWS CLI to call Amazon S3
C. Attach a policy to the S3-access group to deny all S3 actions unless MFA is present Use 1AM accesskeys with the AWS CLI to call Amazon S3
D. Attach a policy to the S3-access group to deny all S3 actions unless MFA is present Request temporary credentials from AWS Security Token Service (AWS STS) Attach the temporary credentials in a profile that Amazon S3 will reference when the user performs actions in Amazon S3

Answer: D

Explanation:
The company should attach a policy to the S3-access group to deny all S3 actions unless MFA is present. The company should request temporary credentials from AWS Security Token Service (AWS STS). The company should attach the temporary credentials in a profile that Amazon S3 will reference when the user performs actions in Amazon S3. This solution will meet the requirements because AWS STS is a service that enables you to request temporary, limited-privilege credentials for IAM users or for users that you authenticate (federated users). You can use MFA with AWS STS to provide an extra layer of security when requesting temporary credentials1. You can use the sts get-session-token AWS CLI command to request temporary credentials that include an MFA token2. You can then use these credentials with the AWS CLI to access Amazon S3 resources. To do this, you need to attach a policy to the IAM group that denies all S3 actions unless MFA is present3. You also need to create a profile in the AWS CLI configuration file that references the temporary credentials.
The other options are not correct because:

- Attaching a policy to the S3 bucket to prompt the IAM user for an MFA code when the IAM user performs actions on the S3 bucket would not work because policies attached to S3 buckets cannot enforce MFA authentication. Policies attached to S3 buckets are resource-based policies that define what actions can be performed on the bucket and by whom. They do not have any logic to prompt for an MFA code or verify it.
- Updating the trust policy for the S3-access group to require principals to use MFA when principals assume the group would not work because trust policies are used for roles, not groups. Trust policies are policies that define which principals can assume a role. They do not apply to groups, which are collections of IAM users that share permissions.
- Creating an Amazon Route 53 Resolver DNS Firewall domain list that contains the allowed domains and configuring a DNS Firewall rule group with rules to allow or block requests based on the domain list would not help with enforcing MFA authentication for Amazon S3 actions. Amazon Route 53 Resolver DNS Firewall is a feature that enables you to filter and regulate outbound DNS traffic for your VPC. You can create reusable collections of filtering rules in DNS Firewall rule groups and associate them with your VPCs. You can specify lists of domain names to allow or block, and you can customize the responses for the DNS queries that you block. This feature is useful for controlling access to sites and blocking DNS-level threats, but not for requiring MFA authentication.

References:

- https://docs.aws.amazon.com/IAM/latest/UserGuide/id_credentials_temp.html
- https://docs.aws.amazon.com/IAM/latest/UserGuide/id_credentials_mfa_enable_cliapi.html
- https://docs.aws.amazon.com/IAM/latest/UserGuide/id_credentials_mfa_sample-policies.html
- https://docs.aws.amazon.com/cli/latest/userguide/cli-configure-profiles.html
- https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/resolver-dns-firewall.html

NEW QUESTION 125
- (Exam Topic 2)
A company wants to containerize a multi-tier web application and move the application from an on-premises data center to AWS. The application includes web. application, and database tiers. The company needs to make the application fault tolerant and scalable. Some frequently accessed data must always be available across application servers. Frontend web servers need session persistence and must scale to meet increases in traffic.
Which solution will meet these requirements with the LEAST ongoing operational overhead?

A. Run the application on Amazon Elastic Container Service (Amazon ECS) on AWS Fargat
B. Use Amazon Elastic File System (Amazon EFS) for data that is frequently accessed between the web and application tier
C. Store the frontend web server session data in Amazon Simple Queue Service (Amazon SOS).
D. Run the application on Amazon Elastic Container Service (Amazon ECS) on Amazon EC2. Use Amazon ElastiCache for Redis to cache frontend web server session dat
E. Use Amazon Elastic Block Store (Amazon EBS) with Multi-Attach on EC2 instances that are distributed across multiple Availability Zones.
F. Run the application on Amazon Elastic Kubernetes Service (Amazon EKS). Configure Amazon EKS to use managed node group
G. Use ReplicaSets to run the web servers and application
H. Create an Amazon Elastic File System (Amazon EFS) Me syste
I. Mount the EFS file system across all EKS pods to store frontend web server session data.
J. Deploy the application on Amazon Elastic Kubernetes Service (Amazon EKS) Configure Amazon EKS to use managed node group
K. Run the web servers and application as Kubernetes deployments in the EKS cluste
L. Store the frontend web server session data in an Amazon DynamoDB tabl
M. Create an Amazon Elastic File System (Amazon EFS) volume that all applications will mount at the time of deployment.

Answer: D

Explanation:
Deploying the application on Amazon EKS with managed node groups simplifies the operational overhead of managing the Kubernetes cluster. Running the web servers and application as Kubernetes deployments ensures that the desired number of pods are always running and can scale up or down as needed. Storing the frontend web server session data in an Amazon DynamoDB table provides a fast, scalable, and durable storage option that can be accessed across multiple Availability Zones. Creating an Amazon EFS volume that all applications will mount at the time of deployment allows the application to share data that is frequently accessed between the web and application tiers. References:

- https://docs.aws.amazon.com/eks/latest/userguide/managed-node-groups.html
- https://docs.aws.amazon.com/eks/latest/userguide/deployments.html
- https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Introduction.html
- https://docs.aws.amazon.com/efs/latest/ug/mounting-fs.html

NEW QUESTION 127
- (Exam Topic 2)
A company runs its sales reporting application in an AWS Region in the United States. The application uses an Amazon API Gateway Regional API and AWS Lambda functions to generate on-demand reports from data in an Amazon RDS for MySQL database. The frontend of the application is hosted on Amazon S3 and is accessed by users through an Amazon CloudFront distribution. The company is using Amazon Route 53 as the DNS service for the domain. Route 53 is configured with a simple routing policy to route traffic to the API Gateway API.
In the next 6 months, the company plans to expand operations to Europe. More than 90% of the database traffic is read-only traffic. The company has already

deployed an API Gateway API and Lambda functions in the new Region.
A solutions architect must design a solution that minimizes latency for users who download reports. Which solution will meet these requirements?

A. Use an AWS Database Migration Service (AWS DMS) task with full load to replicate the primary database in the original Region to the database in the new Regio
B. Change the Route 53 record to latency-based routing to connect to the API Gateway API.
C. Use an AWS Database Migration Service (AWS DMS) task with full load plus change data capture (CDC) to replicate the primary database in the original Region to the database in the new Regio
D. Change the Route 53 record to geolocation routing to connect to the API Gateway API.
E. Configure a cross-Region read replica for the RDS database in the new Regio
F. Change the Route 53 record to latency-based routing to connect to the API Gateway API.
G. Configure a cross-Region read replica for the RDS database in the new Regio
H. Change the Route 53 record to geolocation routing to connect to the API

Answer: C

Explanation:
The company should configure a cross-Region read replica for the RDS database in the new Region. The company should change the Route 53 record to latency-based routing to connect to the API Gateway API. This solution will meet the requirements because a cross-Region read replica is a feature that enables you to create a MariaDB, MySQL, Oracle, PostgreSQL, or SQL Server read replica in a different Region from the source DB instance. You can use cross-Region read replicas to improve availability and disaster recovery, scale out globally, or migrate an existing database to a new Region1. By creating a cross-Region read replica for the RDS database in the new Region, the company can have a standby copy of its primary database that can serve read-only traffic from users in Europe. A latency-based routing policy is a feature that enables you to route traffic based on the latency between your users and your resources. You can use latency-based routing to route traffic to the resource that provides the best latency2. By changing the Route 53 record to latency-based routing, the company can minimize latency for users who download reports by connecting them to the API Gateway API in the Region that provides the best response time.
The other options are not correct because:

- Using AWS Database Migration Service (AWS DMS) to replicate the primary database in the original Region to the database in the new Region would not be as cost-effective or simple as using a
cross-Region read replica. AWS DMS is a service that enables you to migrate relational databases, data
warehouses, NoSQL databases, and other types of data stores. You can use AWS DMS to perform one-time migrations or continuous data replication with high availability and consolidate databases into
a petabyte-scale data warehouse3. However, AWS DMS requires more configuration and management than creating a cross-Region read replica, which is fully managed by Amazon RDS. AWS DMS also incurs additional charges for replication instances and tasks.
- Creating an Amazon API Gateway Data API service integration with Amazon Redshift would not help with disaster recovery or minimizing latency. The Data API is a feature that enables you to query your Amazon Redshift cluster using HTTP requests, without needing a persistent connection or a SQL client. It is useful for building applications that interact with Amazon Redshift, but not for replicating or recovering data from an RDS database.
- Creating an AWS Data Exchange datashare by connecting AWS Data Exchange to the Redshift cluster would not help with disaster recovery or minimizing latency. AWS Data Exchange is a service that makes it easy for AWS customers to exchange data in the cloud. You can use AWS Data Exchange to subscribe to a diverse selection of third-party data products or offer your own data products to other AWS customers. A datashare is a feature that enables you to share live and secure access to your Amazon Redshift data across your accounts or with third parties without copying or moving the underlying data. It is useful for sharing query results and views with other users, but not for replicating or recovering data from an RDS database.

References:

- https://aws.amazon.com/dms/
- https://docs.aws.amazon.com/redshift/latest/mgmt/data-api.html
- https://aws.amazon.com/data-exchange/
- https://docs.aws.amazon.com/redshift/latest/dg/datashare-overview.html

NEW QUESTION 128
- (Exam Topic 2)
A solutions architect is designing a solution to process events. The solution must have the ability to scale in and out based on the number of events that the solution receives. If a processing error occurs, the event must move into a separate queue for review.
Which solution will meet these requirements?

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

Answer: A

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

distributed systems, and serverless applications3. This way, if a processing error occurs, the event will move into the separate queue for review.
Option B is incorrect because publishing events to an Amazon SQS queue and creating an Amazon EC2 Auto Scaling group will not have the ability to scale in and out based on the number of events that the solution receives. Amazon EC2 is a web service that provides secure, resizable compute capacity in the cloud. Auto Scaling is a feature that helps users maintain application availability and allows them to scale their EC2 capacity up or down automatically according to conditions they define. However, for this use case, using SQS and EC2 will not take advantage of the serverless capabilities of Lambda and SNS.
Option C is incorrect because writing events to an Amazon DynamoDB table and configuring a DynamoDB stream for the table will not have the ability to move events into a separate queue for review if a processing error occurs. Amazon DynamoDB is a fully managed key-value and document database that delivers single-digit millisecond performance at any scale. DynamoDB Streams is a feature that captures data
modification events in DynamoDB tables. Users can configure the stream to invoke a Lambda function, but they cannot configure an on-failure destination for the function.
Option D is incorrect because publishing events to an Amazon EventBridge event bus and setting an Application Load Balancer (ALB) as the event bus target will not have the ability to move events into a separate queue for review if a processing error occurs. Amazon EventBridge is a serverless event bus service that makes it easy to connect applications with data from a variety of sources. An ALB is a load balancer that distributes incoming application traffic across multiple targets, such as EC2 instances, containers, IP addresses, Lambda functions, and virtual appliances. Users can configure EventBridge to retry events, but they cannot configure an on-failure destination for the ALB.

NEW QUESTION 133
- (Exam Topic 2)
A company processes environment data. The has a set up sensors to provide a continuous stream of data from different areas in a city. The data is available in JSON format.
The company wants to use an AWS solution to send the data to a database that does not require fixed schemas for storage. The data must be send in real time.
Which solution will meet these requirements?

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

Answer: B

Explanation:
Amazon Kinesis Data Streams is a service that enables real-time data ingestion and processing. Amazon DynamoDB is a NoSQL database that does not require fixed schemas for storage. By using Kinesis Data Streams and DynamoDB, the company can send the JSON data to a database that can handle schemaless data in real time. References:
- https://docs.aws.amazon.com/streams/latest/dev/introduction.html
- https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Introduction.html

NEW QUESTION 138
- (Exam Topic 1)
A solutions architect is investigating an issue in which a company cannot establish new sessions in Amazon Workspaces. An initial analysis indicates that the issue involves user profiles. The Amazon Workspaces environment is configured to use Amazon FSx for Windows File Server as the profile share storage. The FSx for Windows File Server file system is configured with 10 TB of storage.
The solutions architect discovers that the file system has reached its maximum capacity. The solutions architect must ensure that users can regain access. The solution also must prevent the problem from occurring again.
Which solution will meet these requirements?

A. Remove old user profiles to create spac
B. Migrate the user profiles to an Amazon FSx for Lustre file system.
C. Increase capacity by using the update-file-system comman
D. Implement an Amazon CloudWatch metric that monitors free spac
E. Use Amazon EventBridge to invoke an AWS Lambda function to increase capacity as required.
F. Monitor the file system by using the FreeStorageCapacity metric in Amazon CloudWatc
G. Use AWS Step Functions to increase the capacity as required.
H. Remove old user profiles to create spac
I. Create an additional FSx for Windows File Server file system.Update the user profile redirection for 50% of the users to use the new file system.

Answer: B

Explanation:
It can prevent the issue from happening again by monitoring the file system with the FreeStorageCapacity metric in Amazon CloudWatch and using Amazon EventBridge to invoke an AWS Lambda function to increase the capacity as required. This ensures that the file system always has enough free space to store user profiles and avoids reaching maximum capacity.

NEW QUESTION 140
- (Exam Topic 1)
A company is running an event ticketing platform on AWS and wants to optimize the platform's
cost-effectiveness. The platform is deployed on Amazon Elastic Kubernetes Service (Amazon EKS) with Amazon EC2 and is backed by an Amazon RDS for MySQL DB instance. The company is developing new application features to run on Amazon EKS with AWS Fargate.
The platform experiences infrequent high peaks in demand. The surges in demand depend on event dates. Which solution will provide the MOST cost-effective setup for the platform?

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

I. Scale the cluster with Spot Instances to handle peak
J. Purchase 1-year All Upfront Reserved Instances for the database to meet the predicted base loa
K. Temporarily scale up the DB instance manually during peaks.
L. Purchase Compute Savings Plans for the predicted base load of the EKS cluste
M. Scale the cluster with Spot Instances to handle peak
N. Purchase 1-year All Upfront Reserved Instances for the database to meet the predicted base loa
O. Temporarily scale up the DB instance manually during peaks.

Answer: B

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

NEW QUESTION 141
- (Exam Topic 1)
A company has an on-premises monitoring solution using a PostgreSQL database for persistence of events. The database is unable to scale due to heavy ingestion and it frequently runs out of storage.
The company wants to create a hybrid solution and has already set up a VPN connection between its network and AWS. The solution should include the following attributes:
• Managed AWS services to minimize operational complexity
• A buffer that automatically scales to match the throughput of data and requires no on-going administration.
• A visualization toot to create dashboards to observe events in near-real time.
• Support for semi -structured JSON data and dynamic schemas.
Which combination of components will enabled© company to create a monitoring solution that will satisfy these requirements" (Select TWO.)

A. Use Amazon Kinesis Data Firehose to buffer events Create an AWS Lambda function 10 process and transform events
B. Create an Amazon Kinesis data stream to buffer events Create an AWS Lambda function to process and transform evens
C. Configure an Amazon Aurora PostgreSQL DB cluster to receive events Use Amazon Quick Sight to read from the database and create near-real-time visualizations and dashboards
D. Configure Amazon Elasticsearch Service (Amazon ES) to receive events Use the Kibana endpoint deployed with Amazon ES to create near-real-time visualizations and dashboards.
E. Configure an Amazon Neptune 0 DB instance to receive events Use Amazon QuickSight to read from the database and create near-real-time visualizations and dashboards

Answer: AD

Explanation:
https://aws.amazon.com/kinesis/data-firehose/faqs/

NEW QUESTION 142
.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questons and Answers in PDF Format

AWS-Certified-Solutions-Architect-Professional Practice Exam Features:

* AWS-Certified-Solutions-Architect-Professional Questions and Answers Updated Frequently

* AWS-Certified-Solutions-Architect-Professional Practice Questions Verified by Expert Senior Certified Staff

* AWS-Certified-Solutions-Architect-Professional Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* AWS-Certified-Solutions-Architect-Professional Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year