Amazon

Exam Questions AWS-Certified-Solutions-Architect-Professional

Amazon AWS Certified Solutions Architect Professional

NEW QUESTION 1
- (Exam Topic 1)
A company has a three-tier application running on AWS with a web server, an application server, and an Amazon RDS MySQL DB instance. A solutions architect is designing a disaster recovery (OR) solution with an RPO of 5 minutes.
Which solution will meet the company's requirements?

A. Configure AWS Backup to perform cross-Region backups of all servers every 5 minute
B. Reprovision the three tiers in the DR Region from the backups using AWS CloudFormation in the event of a disaster.
C. Maintain another running copy of the web and application server stack in the DR Region using AWS CloudFormation drill detectio
D. Configure cross-Region snapshots ol the DB instance to the DR Region every 5 minute
E. In the event of a disaster, restore the DB instance using the snapshot in the DR Region.
F. Use Amazon EC2 Image Builder to create and copy AMIs of the web and application server to both the primary and DR Region
G. Create a cross-Region read replica of the DB instance in the DR Regio
H. In the event of a disaster, promote the read replica to become the master and reprovision the servers with AWS CloudFormation using the AMIs.
I. Create AMts of the web and application servers in the DR Regio
J. Use scheduled AWS Glue jobs to synchronize the DB instance with another DB instance in the DR Regio
K. In the event of a disaster, switch to the DB instance in the DR Region and reprovision the servers with AWS CloudFormation using the AMIs.

Answer: C

Explanation:
deploying a brand new RDS instance will take >30 minutes. You will use EC2 Image builder to put the AMIs into the new region, but not use image builder to LAUNCH them.

NEW QUESTION 2
- (Exam Topic 1)
A company is running an application distributed over several Amazon EC2 instances in an Auto Seating group behind an Application Load Balancer The security team requires that all application access attempts be made available for analysis information about the client IP address, connection type, and user agent must be included
Which solution will meet these requirements?

A. Enable EC2 detailed monitoring, and include network log
B. Send all logs through Amazon Kinesis Data Firehose to an Amazon Elasticsearch Service (Amazon ES) cluster that the security team uses for analysis.
C. Enable VPC Flow Logs for all EC2 instance network interfaces Publish VPC Flow Logs to an Amazon S3 bucket Have the security team use Amazon Athena to query and analyze the logs.
D. Enable access logs for the Application Load Balancer, and publish the logs to an Amazon S3 bucket.Have the security team use Amazon Athena to query and analyze the logs
E. Enable Traffic Mirroring and specify all EC2 instance network interfaces as the sourc
F. Send all traffic information through Amazon Kinesis Data Firehose to an Amazon Elasticsearch Service (Amazon ES) cluster that the security team uses for analysis.

Answer: C

Explanation:
https://docs.aws.amazon.com/elasticloadbalancing/latest/application/load-balancer-access-logs.html https://docs.aws.amazon.com/vpc/latest/mirroring/what-is-traffic-mirroring.html

NEW QUESTION 3
- (Exam Topic 1)
A solutions architect at a largo company needs to set up network security for outbound traffic to the internet from all AWS accounts within an organization m AWS Organizations The organization has more than 100 AWS accounts, and the accounts route to each other by using a centralized AWS Transit Gateway. Each account has both an internet gateway and a NAT gateway for outbound traffic to the interne) 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. Creates a new VPC for outbound traffic to the internet Connect the existing transit gateway to the new VPC Configure a new NAT gateway Create an Auto Scaling group of Amazon EC2 Instances that run an open-source internet proxy for rule-based filtering across all Availability Zones in the Region Modify all default routes to point to the proxy's Auto Scaling group
B. Create a new VPC for outbound traffic to the internet Connect the existing transit gateway to the new VPC Configure a new NAT gateway Use an AWS Network Firewall firewall for rule-based filtering Create Network Firewall endpoints In each Availability Zone Modify all default routes to point to the Network Firewall endpoints
C. Create an AWS Network Firewall firewall for rule-based filtering in each AWS account Modify all default routes to point to the Network Firewall firewalls in each account.
D. In each AWS account, create an Auto Scaling group of network-optimized Amazon EC2 instances that run an open-source internet proxy for rule-based filtering 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/ https://aws.amazon.com/blogs/networking-and-content-delivery/deploy-centralized-traffic-filtering-using-aws-n

NEW QUESTION 4
- (Exam Topic 1)
A large company in Europe plans to migrate its applications to the AWS Cloud. The company uses multiple AWS accounts for various business groups. A data privacy law requires the company to restrict developers' access to AWS European Regions only.
What should the solutions architect do to meet this requirement with the LEAST amount of management overhead^

A. Create IAM users and IAM groups in each accoun
B. Create IAM policies to limit access to non-European Regions Attach the IAM policies to the IAM groups
C. Enable AWS Organizations, attach the AWS accounts, and create OUs for European Regions andnon-European Region
D. Create SCPs to limit access to non-European Regions and attach the policies to the OUs.
E. Set up AWS Single Sign-On and attach AWS account
F. Create permission sets with policies to restrict access to non-European Regions Create IAM users and IAM groups in each account.
G. Enable AWS Organizations, attach the AWS accounts, and create OUs for European Regions andnon-European Region
H. Create permission sets with policies to restrict access to non-European Region
I. Create IAM users and IAM groups in the primary account.

Answer: B

Explanation:
"This policy uses the Deny effect to deny access to all requests for operations that don't target one of the two approved regions (eu-central-1 and eu-west-1)."
https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps_examples_general.htm
https://docs.aws.amazon.com/IAM/latest/UserGuide/reference_policies_elements_condition.html

NEW QUESTION 5
- (Exam Topic 1)
A company has many services running in its on-premises data center. The data center is connected to AWS using AWS Direct Connect (DX) and an IPSec VPN.
The service data is sensitive and connectivity cannot traverse the internet. The company wants to expand into a new market segment and begin offering its
services to other companies that are using AWS.
Which solution will meet these requirements?

A. Create a VPC Endpoint Service that accepts TCP traffic, host it behind a Network Load Balancer, and make the service available over DX.
B. Create a VPC Endpoint Service that accepts HTTP or HTTPS traffic, host it behind an Application Load Balancer, and make the service available over DX.
C. Attach an internet gateway to the VP
D. and ensure that network access control and security group rules allow the relevant inbound and outbound traffic.
E. Attach a NAT gateway to the VP
F. and ensure that network access control and security group rules allow the relevant inbound and outbound traffic.

Answer: A

NEW QUESTION 6
- (Exam Topic 1)
A company maintains a restaurant review website. The website is a single-page application where files are stored in Amazon S3 and delivered using Amazon
CloudFront. The company receives several fake postings every day that are manually removed.
The security team has identified that most of the fake posts are from bots with IP addresses that have a bad reputation within the same global region. The team
needs to create a solution to help restrict the bots from accessing the website.
Which strategy should a solutions architect use?

A. Use AWS Firewall Manager to control the CloudFront distribution security setting
B. Create a geographical block rule and associate it with Firewall Manager.
C. Associate an AWS WAF web ACL with the CloudFront distributio
D. Select the managed Amazon IP reputation rule group for the web ACL with a deny action.
E. Use AWS Firewall Manager to control the CloudFront distribution security setting
F. Select the managed Amazon IP reputation rule group and associate it with Firewall Manager with a deny action.
G. Associate an AWS WAF web ACL with the CloudFront distributio
H. Create a rule group for the web ACL with a geographical match statement with a deny action.

Answer: B

Explanation:
IP reputation rule groups allow you to block requests based on their source. Choose one or more of these rule groups if you want to reduce your exposure to
BOTS!!!! traffic or exploitation attempts
The Amazon IP reputation list rule group contains rules that are based on Amazon internal threat intelligence. This is useful if you would like to block IP addresses
typically associated with bots or other threats. Inspects for a list of IP addresses that have been identified as bots by Amazon threat intelligence.

NEW QUESTION 7
- (Exam Topic 1)
A company is running an application on Amazon EC2 instances in three environments; development, testing, and production. The company uses AMIs to deploy
the EC2 instances. The company builds the AMIs by using custom deployment scripts and infrastructure orchestration tools for each release in each environment.
The company is receiving errors in its deployment process. Errors appear during operating system package downloads and during application code installation
from a third-party Git hosting service. The company needs deployments to become more reliable across all environments.
Which combination of steps will meet these requirements? (Select THREE).

A. Mirror the application code to an AWS CodeCommit Git repositor
B. Use the repository to build EC2 AMIs.
C. Produce multiple EC2 AMI
D. one for each environment, for each release.
E. Produce one EC2 AMI for each release for use across all environments.
F. Mirror the application code to a third-party Git repository that uses Amazon S3 storag
G. Use the repository for deployment.
H. Replace the custom scripts and tools with AWS CodeBuil
I. Update the infrastructure deployment process to use EC2 Image Builder.

Answer: ACE

NEW QUESTION 8

- (Exam Topic 1)

A company has a complex web application that leverages Amazon CloudFront for global scalability and performance. Over time, users report that the web application is slowing down.

The company's operations team reports that the CloudFront cache hit ratio has been dropping steadily. The cache metrics report indicates that query strings on some URLs are inconsistently ordered and are specified sometimes in mixed-case letters and sometimes in lowercase letters.

Which set of actions should the solutions architect take to increase the cache hit ratio as quickly as possible?

A. Deploy a Lambda@Edge function to sort parameters by name and force them to be lowercas
B. Select the CloudFront viewer request trigger to invoke the function.
C. Update the CloudFront distribution to disable caching based on query string parameters.
D. Deploy a reverse proxy after the load balancer to post-process the emitted URLs in the application to force the URL strings to be lowercase.
E. Update the CloudFront distribution to specify casing-insensitive query string processing.

Answer: A

Explanation:
https://docs.amazonaws.cn/en_us/AmazonCloudFront/latest/DeveloperGuide/lambda-examples.html#lambda-ex Before CloudFront serves content from the cache it will trigger any Lambda function associated with the
Viewer Request, in which we can normalize parameters.
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/lambda-examples.html#lambda-examp

NEW QUESTION 9

- (Exam Topic 1)

A company wants to host a new global website that consists of static content. A solutions architect is working on a solution that uses Amazon CloudFront with an origin access identity <OAI) to access website content that is stored in a private Amazon S3 bucket.

During testing, the solutions architect receives 404 errors from the S3 bucket. Error messages appear only for attempts to access paths that end with a forward slash. such as example.com/path/. These requests should return the existing S3 object path/index.html. Any potential solution must not prevent CloudFront from caching the content.

What should the solutions architect do to resolve this problem?

A. Change the CloudFront origin to an Amazon API Gateway proxy endpoin
B. Rewrite the S3 request URL by using an AWS Lambda function.
C. Change the CloudFront origin to an Amazon API Gateway endpoin
D. Rewrite the S3 request URL in an AWS service integration.
E. Change the CloudFront configuration to use an AWS Lambda@Edge function that is invoked by a viewer request event to rewrite the S3 request URL.
F. Change the CloudFront configuration to use an AWS Lambda@Edge function that is invoked by an origin request event to rewrite the S3 request URL.

Answer: C

NEW QUESTION 10

- (Exam Topic 1)

A media company uses Amazon DynamoDB to store metadata for its catalog of movies that are available to

stream. Each media item Contains user-facing content that concludes a description of the media, a list of search tags, and similar data. In addition, media items include a list of Amazon S3 key names that relate to movie files. The company stores these movie files in a single S3 bucket that has versioning enable. The company uses Amazon CloudFront to serve these movie files.

The company has 100.000 media items, and each media item can have many different S3 objects that represent different encodings of the same media S3 objects that belong to the same media item are grouped together under the same key prefix, which is a random unique ID

Because of an expiring contract with a media provider, the company must remove 2.000 media Items. The company must completely delete all DynamoDB keys and movie files on Amazon S3 that are related to these media items within 36 hours The company must ensure that the content cannot be recovered.

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

A. Configure the dynamoDB table with a TTL fiel
B. Create and invoke an AWS Lambda function to perform a conditional update Set the TTL field to the time of the contract's expiration on every affected media item.
C. Configure an S3 Lifecycle object expiration rule that is based on the contract's expiration date
D. Write a script to perform a conditional delete on all the affected DynamoDB records
E. Temporarily suspend versioning on the S3 bucke
F. Create and invoke an AWS Lambda function that deletes affected objects Reactivate versioning when the operation is complete
G. Write a script to delete objects from Amazon S3 Specify in each request a NoncurrentVersionExpiration property with a NoncurrentDays attribute set to 0.

Answer: CE

NEW QUESTION 10

- (Exam Topic 1)

A company uses AWS Transit Gateway for a hub-and-spoke model to manage network traffic between many VPCs. The company is developing a new service that must be able to send data at 100 Gbps. The company needs a faster connection to other VPCs in the same AWS Region.

Which solution will meet these requirements?

A. Establish VPC peering between the necessary VPC
B. Ensure that all route tables are updated as required.
C. Attach an additional transit gateway to the VPC
D. Update the route tables accordingly.
E. Create AWS Site-to-Site VPN connections that use equal-cost multi-path (ECMP) routing between the necessary VPCs.
F. Create an additional attachment from the necessary VPCs to the existing transit gateway.

Answer: D

NEW QUESTION 14

- (Exam Topic 1)

A company hosts a web application that tuns on a group of Amazon EC2 instances that ate behind an Application Load Balancer (ALB) in a VPC. The company wants to analyze the network payloads lo reverse-engineer a sophisticated attack of the application.
Which approach should the company take to achieve this goal?

A. Enable VPC Flow Log
B. Store the flow logs in an Amazon S3 bucket for analysis.
C. Enable Traffic Mirroring on the network interface of the EC2 instance
D. Send the mirrored traffic lo a target for storage and analysis.
E. Create an AWS WAF web AC
F. and associate it with the AL
G. Configure AWS WAF logging.
H. Enable logging for the AL
I. Store the logs in an Amazon S3 bucket for analysis.

Answer: A

NEW QUESTION 18
- (Exam Topic 1)
A company needs to implement a patching process for its servers. The on-premises servers and Amazon EC2 instances use a variety of tools to perform patching. Management requires a single report showing the patch status of all the servers and instances.
Which set of actions should a solutions architect take to meet these requirements?

A. Use AWS Systems Manager to manage patches on the on-premises servers and EC2 instance
B. Use Systems Manager to generate patch compliance reports.
C. Use AWS OpsWorks to manage patches on the on-premises servers and EC2 instance
D. Use Amazon OuickSight integration with OpsWorks to generate patch compliance reports.
E. Use an Amazon EventBridge (Amazon CloudWatch Events) rule to apply patches by scheduling an AWS Systems Manager patch remediation jo
F. Use Amazon Inspector to generate patch compliance reports.
G. Use AWS OpsWorks to manage patches on the on-premises servers and EC2 instance
H. Use AWS X-Ray to post the patch status to AWS Systems Manager OpsCenter to generate patch compliance reports.

Answer: A

Explanation:
https://docs.aws.amazon.com/systems-manager/latest/userguide/systems-manager-patch.html

NEW QUESTION 20
- (Exam Topic 1)
A company has registered 10 new domain names. The company uses the domains for online marketing. The company needs a solution that will redirect online visitors to a specific URL for each domain. All domains and target URLs are defined in a JSON document. All DNS records are managed by Amazon Route 53.
A solutions architect must implement a redirect service that accepts HTTP and HTTPS requests.
Which combination of steps should the solutions architect take to meet these requirements with the LEAST amount of operational effort? (Choose three.)

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

Answer: CEF

Explanation:
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/lambda-edge-how-it-works-tutorial.ht

NEW QUESTION 21
- (Exam Topic 1)
A company is providing weather data over a REST-based API to several customers. The API is hosted by Amazon API Gateway and is integrated with different AWS Lambda functions for each API operation. The company uses Amazon Route 53 for DNS and has created a resource record of weather.example.com. The company stores data for the API in Amazon DynamoDB tables. The company needs a solution that will give the API the ability to fail over to a different AWS Region.
Which solution will meet these requirements?

A. Deploy a new set of Lambda functions in a new Regio
B. Update the API Gateway API to use an edge-optimized API endpoint with Lambda functions from both Regions as target
C. Convert the DynamoDB tables to global tables.
D. Deploy a new API Gateway API and Lambda functions in another Regio
E. Change the Route 53 DNS record to a multivalue answe
F. Add both API Gateway APIs to the answe
G. Enable target health monitorin
H. Convert the DynamoDB tables to global tables.
I. Deploy a new API Gateway API and Lambda functions in another Regio
J. Change the Route 53 DNS record to a failover recor
K. Enable target health monitorin
L. Convert the DynamoDB tables to global tables.
M. Deploy a new API Gateway API in a new Regio
N. Change the Lambda functions to global functions.Change the Route 53 DNS record to a multivalue answe
O. Add both API Gateway APIs to the answe

P. Enable target health monitorin
Q. Convert the DynamoDB tables to global tables.

Answer: C

Explanation:
https://docs.aws.amazon.com/apigateway/latest/developerguide/dns-failover.html

NEW QUESTION 22
- (Exam Topic 1)
A company is running a web application with On-Demand Amazon EC2 instances in Auto Scaling groups that scale dynamically based on custom metrics After extensive testing, the company determines that the m5.2xlarge instance size is optimal for the workload Application data is stored in db.r4.4xlarge Amazon RDS instances that are confirmed to be optimal. The traffic to the web application spikes randomly during the day.
What other cost-optimization methods should the company implement to further reduce costs without impacting the reliability of the application?

A. Double the instance count in the Auto Scaling groups and reduce the instance size to m5.large
B. Reserve capacity for the RDS database and the minimum number of EC2 instances that are constantly running.
C. Reduce the RDS instance size to db.r4.xlarge and add five equivalent^ sized read replicas to provide reliability.
D. Reserve capacity for all EC2 instances and leverage Spot Instance pricing for the RDS database.

Answer: B

Explanation:
People are being confused by the term 'reserve capacity'. This is not the same as an on-demand capacity reservation. This article by AWS clearly states that by 'reserving capacity' you are reserving the instances and reducing your costs. See https://aws.amazon.com/aws-cost-management/aws-cost-optimization/reserved-instances/

NEW QUESTION 26
- (Exam Topic 1)
A company wants to retire its Oracle Solaris NFS storage arrays. The company requires rapid data migration over its internet network connection to a combination of destinations for Amazon S3. Amazon Elastic File System (Amazon EFS), and Amazon FSx lor Windows File Server. The company also requires a full initial copy, as well as incremental transfers of changes until the retirement of the storage arrays. All data must be encrypted and checked for integrity.
What should a solutions architect recommend to meet these requirements?

A. Configure CloudEndur
B. Create a project and deploy the CloudEndure agent and token to the storage arra
C. Run the migration plan to start the transfer.
D. Configure AWS DataSyn
E. Configure the DataSync agent and deploy it to the local networ
F. Create a transfer task and start the transfer.
G. Configure the aws S3 sync comman
H. Configure the AWS client on the client side with credential
I. Run the sync command to start the transfer.
J. Configure AWS Transfer (or FT
K. Configure the FTP client with credential
L. Script the client to connect and sync to start the transfer.

Answer: B

NEW QUESTION 29
- (Exam Topic 1)
A company is planning on hosting its ecommerce platform on AWS using a multi-tier web application designed for a NoSQL database. The company plans to use the us-west-2 Region as its primary Region. The company want to ensure that copies of the application and data are available in a second Region, us-west-1, for disaster recovery. The company wants to keep the time to fail over as low as possible. Failing back to the primary Region should be possible without administrative interaction after the primary service is restored.
Which design should the solutions architect use?

A. Use AWS Cloud Formation StackSets lo create the stacks in both Regions with Auto Scaling groups for the web and application tier
B. Asynchronously replicate static content between Regions using Amazon S3 cross-Region replicatio
C. Use an Amazon Route 53 DNS failover routing policy to direct users to the secondary site in us-west-1 in the event of an outag
D. Use Amazon DynamoDB global tables for the database tier.
E. Use AWS Cloud Formation StackSets to create the stacks in both Regions with Auto Scaling groups for the web and application tier
F. Asynchronously replicate static content between Regions using Amazon S3 cross-Region replicatio
G. Use an Amazon Route 53 DNS failover routing policy to direct users to the secondary site in us-west-1 in the event of an outag
H. Deploy an Amazon Aurora global database for the database tier.
I. Use AWS Service Catalog to deploy the web and application servers in both Region
J. Asynchronously replicate static content between the two Regions using Amazon S3 cross-Region replicatio
K. Use Amazon Route 53 health checks to identify a primary Region failure and update the public DNS entry listing to the secondary Region in the event of an outag
L. Use Amazon RDS for MySQL withcross-Region replication for the database tier.
M. Use AWS CloudFormation StackSets to create the stacks in both Regions using Auto Scaling groups for the web and application tier
N. Asynchronously replicate static content between Regions using Amazon S3 cross-Region replicatio
O. Use Amazon CloudFront with static files in Amazon S3, and multi-Region origins for the front-end web tie
P. Use Amazon DynamoD8 tables in each Region with scheduled backups to Amazon S3.

Answer: A

NEW QUESTION 33
- (Exam Topic 1)

A company has a data lake in Amazon S3 that needs to be accessed by hundreds of applications across many AWS accounts. The company's information security policy states that the S3 bucket must not be accessed over the public internet and that each application should have the minimum permissions necessary to function.

To meet these requirements, a solutions architect plans to use an S3 access point that is restricted to specific VPCs tor each application.

Which combination of steps should the solutions architect take to implement this solution? (Select TWO.)

A. Create an S3 access point for each application in the AWS account that owns the S3 bucke
B. Configure each access point to be accessible only from the application's VP
C. Update the bucket policy to require access from an access point.
D. Create an interface endpoint for Amazon S3 in each application's VP
E. Configure the endpoint policy to allow access to an S3 access poin
F. Create a VPC gateway attachment for the S3 endpoint.
G. Create a gateway endpoint lor Amazon S3 in each application's VP
H. Configure the endpoint policy to allow access to an S3 access poin
I. Specify the route table that is used to access the access point.
J. Create an S3 access point for each application in each AWS account and attach the access points to the S3 bucke
K. Configure each access point to be accessible only from the application's VP
L. Update the bucket policy to require access from an access point.
M. Create a gateway endpoint for Amazon S3 in the data lake's VP
N. Attach an endpoint policy to allow access to the S3 bucke
O. Specify the route table that is used to access the bucket.

Answer: AC

Explanation:
https://joe.blog.freemansoft.com/2020/04/protect-data-in-cloud-with-s3-access.html https://aws.amazon.com/s3/features/access-points/
https://aws.amazon.com/s3/features/access-points/
&
https://aws.amazon.com/blogs/storage/managing-amazon-s3-access-with-vpc-endpoints-and-s3-access-points/

NEW QUESTION 36
- (Exam Topic 1)
A company stores sales transaction data in Amazon DynamoDB tables. To detect anomalous behaviors and respond quickly, all changes lo the items stored in the DynamoDB tables must be logged within 30 minutes.

Which solution meets the requirements?

A. Copy the DynamoDB tables into Apache Hive tables on Amazon EMR every hour and analyze them (or anomalous behavior
B. Send Amazon SNS notifications when anomalous behaviors are detected.
C. Use AWS CloudTrail to capture all the APIs that change the DynamoDB table
D. Send SNS notifications when anomalous behaviors are detected using CloudTrail event filtering.
E. Use Amazon DynamoDB Streams to capture and send updates to AWS Lambd
F. Create a Lambda function to output records lo Amazon Kinesis Data Stream
G. Analyze any anomalies with Amazon Kinesis Data Analytic
H. Send SNS notifications when anomalous behaviors are detected.
I. Use event patterns in Amazon CloudWatch Events to capture DynamoDB API call events with an AWS Lambda (unction as a target to analyze behavio
J. Send SNS notifications when anomalous behaviors are detected.

Answer: C

Explanation:
https://aws.amazon.com/blogs/database/dynamodb-streams-use-cases-and-design-patterns/#:~:text=DynamoDB DynamoDb Stream to capture DynamoDB
update. And Kinesis Data Analytics for anomaly detection (it uses
AWS proprietary Random Cut Forest Algorithm)

NEW QUESTION 39
- (Exam Topic 1)
A company has a website that enables users to upload videos. Company policy states the uploaded videos must be analyzed for restricted content. An uploaded video is placed in Amazon S3, and a message is pushed to an Amazon SOS queue with the video's location. A backend application pulls this location from Amazon SOS and analyzes the video.

The video analysis is compute-intensive and occurs sporadically during the day The website scales with demand. The video analysis application runs on a fixed number of instances. Peak demand occurs during the holidays, so the company must add instances to the application dunng this time. All instances used are currently on-demand Amazon EC2 T2 instances. The company wants to reduce the cost of the current solution.

Which of the following solutions is MOST cost-effective?

A. Keep the website on T2 instance
B. Determine the minimum number of website instances required during off-peak times and use Spot Instances to cover them while using Reserved Instances to cover peak deman
C. Use Amazon EC2 R4 and Amazon EC2 R5 Reserved Instances in an Auto Scaling group for the video analysis application
D. Keep the website on T2 instance
E. Determine the minimum number of website instances required during off-peak times and use Reserved Instances to cover them while using On-Demand Instances to cover peak deman
F. Use Spot Fleet for the video analysis application comprised of Amazon EC2 C4 and Amazon EC2 C5 Spot Instances.
G. Migrate the website to AWS Elastic Beanstalk and Amazon EC2 C4 instance
H. Determine the minimum number of website instances required during off-peak times and use On-Demand Instances to cover them while using Spot capacity to cover peak demand Use Spot Fleet for the video anarysis application comprised of C4 and Amazon EC2 C5 instances.
I. Migrate the website to AWS Elastic Beanstalk and Amazon EC2 R4 instance
J. Determine the minimum number of website instances required during off-peak times and use Reserved Instances to cover them while using On-Demand Instances to cover peak demand Use Spot Fleet for the video analysis application comprised of R4 and Amazon EC2 R5 instances

Answer: B

NEW QUESTION 43
- (Exam Topic 1)
A company has an Amazon VPC that is divided into a public subnet and a pnvate subnet. A web application runs in Amazon VPC. and each subnet has its own NACL The public subnet has a CIDR of 10.0.0 0/24 An Application Load Balancer is deployed to the public subnet The private subnet has a CIDR of 10.0.1.0/24. Amazon EC2 instances that run a web server on port 80 are launched into the private subnet

Onty network traffic that is required for the Application Load Balancer to access the web application can be allowed to travel between the public and private subnets

What collection of rules should be written to ensure that the private subnet's NACL meets the requirement? (Select TWO.)

A. An inbound rule for port 80 from source 0.0 0.0/0
B. An inbound rule for port 80 from source 10.0 0 0/24
C. An outbound rule for port 80 to destination 0.0.0.0/0
D. An outbound rule for port 80 to destination 10.0.0.0/24
E. An outbound rule for ports 1024 through 65535 to destination 10.0.0.0/24

Answer: BE

Explanation:
Ephemeral ports are not covered in the syllabus so be careful that you don't confuse day to day best practise with what is required for the exam. Link to an explanation on Ephemeral ports here. https://acloud.guru/forums/aws-certified-solutions-architect-associate/discussion/-KUbcwo4IXefMl7janaK/netw

NEW QUESTION 45
- (Exam Topic 1)
A company runs an e-commerce platform with front-end and e-commerce tiers. Both tiers run on LAMP stacks with the front-end instances running behind a load balancing appliance that has a virtual offering on AWS Current*/, the operations team uses SSH to log in to the instances to maintain patches and address other concerns. The platform has recently been the target of multiple attacks, including.
• A DDoS attack.
• An SOL injection attack
• Several successful dictionary attacks on SSH accounts on the web servers
The company wants to improve the security of the e-commerce platform by migrating to AWS. The company's solutions architects have decided to use the following approach;
• Code review the existing application and fix any SQL injection issues.
• Migrate the web application to AWS and leverage the latest AWS Linux AMI to address initial security patching.
• Install AWS Systems Manager to manage patching and allow the system administrators to run commands on all instances, as needed.
What additional steps will address all of the identified attack types while providing high availability and minimizing risk?

A. Enable SSH access to the Amazon EC2 instances using a security group that limits access to specific IP
B. Migrate on-premises MySQL to Amazon RDS Multi-AZ Install the third-party load balancer from the AWS Marketplace and migrate the existing rules to the load balancer's AWS instances Enable AWS Shield Standard for DDoS protection
C. Disable SSH access to the Amazon EC2 instance
D. Migrate on-premises MySQL to Amazon RDS Multi-AZ Leverage an Elastic Load Balancer to spread the load and enable AWS Shield Advanced for protectio
E. Add an Amazon CloudFront distribution in front of the website Enable AWS WAF on the distribution to manage the rules.
F. Enable SSH access to the Amazon EC2 instances through a bastion host secured by limiting access to specific IP addresse
G. Migrate on-premises MySQL to a self-managed EC2 instanc
H. Leverage an AWS Elastic Load Balancer to spread the load, and enable AWS Shield Standard for DDoS protection Add anAmazon CloudFront distribution in front of the website.
I. Disable SSH access to the EC2 instance
J. Migrate on-premises MySQL to Amazon RDS Single-A
K. Leverage an AWS Elastic Load Balancer to spread the load Add an Amazon CloudFront distribution in front of the website Enable AWS WAF on the distribution to manage the rules.

Answer: B

NEW QUESTION 47
- (Exam Topic 1)
A solutions architect works for a government agency that has strict disaster recovery requirements All Amazon Elastic Block Store (Amazon EBS) snapshots are required to be saved in at least two additional AWS Regions. The agency also is required to maintain the lowest possible operational overhead.
Which solution meets these requirements?

A. Configure a policy in Amazon Data Lifecycle Manager (Amazon DLMJ to run once daily to copy the EBS snapshots to the additional Regions.
B. Use Amazon EventBridge (Amazon CloudWatch Events) to schedule an AWS Lambda function to copy the EBS snapshots to the additional Regions.
C. Set up AWS Backup to create the EBS snapshot
D. Configure Amazon S3 cross-Region replication to copy the EBS snapshots to the additional Regions.
E. Schedule Amazon EC2 Image Builder to run once daily to create an AMI and copy the AMI to the additional Regions.

Answer: B

NEW QUESTION 48
- (Exam Topic 1)
A group of research institutions and hospitals are in a partnership to study 2 PBs of genomic data. The institute that owns the data stores it in an Amazon S3 bucket and updates it regularly. The institute would like to give all of the organizations in the partnership read access to the data. All members of the partnership are extremety cost-conscious, and the institute that owns the account with the S3 bucket is concerned about covering the costs tor requests and data transfers from Amazon S3.
Which solution allows for secure datasharing without causing the institute that owns the bucket to assume all the costs for S3 requests and data transfers'?

A. Ensure that all organizations in the partnership have AWS account
B. In the account with the S3 bucket, create a cross-account role for each account in the partnership that allows read access to the dat
C. Have the organizations assume and use that read role when accessing the data.
D. Ensure that all organizations in the partnership have AWS account
E. Create a bucket policy on the bucket that owns the data The policy should allow the accounts in the partnership read access to the bucke

F. Enable Requester Pays on the bucke
G. Have the organizations use their AWS credentials when accessing the data.
H. Ensure that all organizations in the partnership have AWS account
I. Configure buckets in each of the accounts with a bucket policy that allows the institute that owns the data the ability to write to the bucket Periodically sync the data from the institute's account to the other organization
J. Have the organizations use their AWS credentials when accessing the data using their accounts
K. Ensure that all organizations in the partnership have AWS account
L. In the account with the S3 bucket, create a cross-account role for each account in the partnership that allows read access to the dat
M. Enable Requester Pays on the bucke
N. Have the organizations assume and use that read role when accessing the data.

Answer: B

Explanation:
In general, bucket owners pay for all Amazon S3 storage and data transfer costs associated with their bucket. A bucket owner, however, can configure a bucket to be a Requester Pays bucket. With Requester Pays buckets, the requester instead of the bucket owner pays the cost of the request and the data download from the bucket. The bucket owner always pays the cost of storing data. If you enable Requester Pays on a bucket, anonymous access to that bucket is not allowed.
https://docs.aws.amazon.com/AmazonS3/latest/userguide/RequesterPaysExamples.html

NEW QUESTION 53
- (Exam Topic 1)
A company is developing and hosting several projects in the AWS Cloud. The projects are developed across multiple AWS accounts under the same organization in AWS Organizations. The company requires the cost lor cloud infrastructure to be allocated to the owning project. The team responsible for all of the AWS accounts has discovered that several Amazon EC2 instances are lacking the Project tag used for cost allocation.
Which actions should a solutions architect take to resolve the problem and prevent it from happening in the future? (Select THREE.)

A. Create an AWS Config rule in each account to find resources with missing tags.
B. Create an SCP in the organization with a deny action for ec2:RunInstances if the Project tag is missing.
C. Use Amazon Inspector in the organization to find resources with missing tags.
D. Create an IAM policy in each account with a deny action for ec2:RunInstances if the Project tag is missing.
E. Create an AWS Config aggregator for the organization to collect a list of EC2 instances with the missing Project tag.
F. Use AWS Security Hub to aggregate a list of EC2 instances with the missing Project tag.

Answer: CDE

NEW QUESTION 57
- (Exam Topic 1)
A fitness tracking company serves users around the world, with its primary markets in North America and Asia. The company needs to design an infrastructure for its read-heavy user authorization application with the following requirements:
• Be resilient to problems with the application in any Region.
• Write to a database in a single Region.
• Read from multiple Regions.
• Support resiliency across application tiers in each Region.
• Support the relational database semantics reflected in the application. Which combination of steps should a solutions architect take? (Select TWO.)

A. Use an Amazon Route 53 geoproximity routing policy combined with a multivalue answer routing policy.
B. Deploy we
C. application, and MySQL database servers to Amazon EC2 instances in each Regio
D. Set up the application so that reads and writes are local to the Regio
E. Create snapshots of the web, application, and database servers and store the snapshots in an Amazon S3 bucket in both Region
F. Set upcross-Region replication for the database layer.
G. Use an Amazon Route 53 geolocation routing policy combined with a failover routing policy.
H. Set up web, application, and Amazon RDS for MySQL instances in each Regio
I. Set up the application so that reads are local and writes are partitioned based on the use
J. Set up a Multi-AZ failover for the web, application, and database server
K. Set up cross-Region replication for the database layer.
L. Set up active-active web and application servers in each Regio
M. Deploy an Amazon Aurora global database with clusters in each Regio
N. Set up the application to use the in-Region Aurora database endpoint
O. Create snapshots of the web and application servers and store them in an Amazon S3 bucket in both Regions.

Answer: CE

Explanation:
https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-policy.html
Geoproximity routing policy is good to control the user traffic to specific regions. However, a multivalue answer routing policy may cause the users to be randomly sent to other healthy regions that may be far away from the user's location. You can use geolocation routing policy to direct the North American users to your servers on the North America region and configure failover routing to the Asia region in case the North America region fails. You can configure the same for the Asian users pointed to the Asia region servers and have the North America region as its backup.

NEW QUESTION 58
- (Exam Topic 1)
An online retail company hosts its stateful web-based application and MySQL database in an on-premises data center on a single server. The company wants to increase its customer base by conducting more marketing campaigns and promotions. In preparation, the company wants to migrate its application and database to AWS to increase the reliability of its architecture.
Which solution should provide the HIGHEST level of reliability?

A. Migrate the database to an Amazon RDS MySQL Multi-AZ DB instanc
B. Deploy the application in an Auto Scaling group on Amazon EC2 instances behind an Application Load Balance
C. Store sessions in Amazon Neptune.

D. Migrate the database to Amazon Aurora MySQ
E. Deploy the application in an Auto Scaling group on Amazon EC2 instances behind an Application Load Balance
F. Store sessions in an Amazon ElastiCache for Redis replication group.
G. Migrate the database to Amazon DocumentDB (with MongoDB compatibility). Deploy the application in an Auto Scaling group on Amazon EC2 instances behind a Network Load Balance
H. Store sessions in Amazon Kinesis Data Firehose.
I. Migrate the database to an Amazon RDS MariaDB Multi-AZ DB instanc
J. Deploy the application in an Auto Scaling group on Amazon EC2 instances behind an Application Load Balance
K. Store sessions in Amazon ElastiCache for Memcached.

Answer: B

NEW QUESTION 62
- (Exam Topic 1)
A financial services company logs personally identifiable information 10 its application logs stored in Amazon S3. Due to regulatory compliance requirements, the log files must be encrypted at rest. The security team has mandated that the company's on-premises hardware security modules (HSMs) be used to generate the CMK material.
Which steps should the solutions architect take to meet these requirements?

A. Create an AWS CloudHSM cluste
B. Create a new CMK in AWS KMS using AWS_CloudHSM as the source (or the key material and an origin of AWS_CLOUDHS
C. Enable automatic key rotation on the CMK with a duration of 1 yea
D. Configure a bucket policy on the togging bucket thai disallows uploads of unencrypted data and requires that the encryption source be AWS KMS.
E. Provision an AWS Direct Connect connection, ensuring there is no overlap of the RFC 1918 address space between on-premises hardware and the VPC
F. Configure an AWS bucket policy on the logging bucket that requires all objects to be encrypte
G. Configure the logging application to query theon-premises HSMs from the AWS environment for the encryption key material, and create a unique CMK for each logging event.
H. Create a CMK in AWS KMS with no key material and an origin of EXTERNA
I. Import the key material generated from the on-premises HSMs into the CMK using the public key and import token provided by AW
J. Configure a bucket policy on the logging bucket that disallows uploads ofnon-encrypted data and requires that the encryption source be AWS KMS.
K. Create a new CMK in AWS KMS with AWS-provided key material and an origin of AWS_KMS.Disable this CM
L. and overwrite the key material with the key material from the on-premises HSM using the public key and import token provided by AW
M. Re-enable the CM
N. Enable automatic key rotation on the CMK with a duration of 1 yea
O. Configure a bucket policy on the logging bucket that disallows uploads of non-encrypted data and requires that the encryption source be AWS KMS.

Answer: C

Explanation:
https://aws.amazon.com/blogs/security/how-to-byok-bring-your-own-key-to-aws-kms-for-less-than-15-00-a-yea
https://docs.aws.amazon.com/kms/latest/developerguide/importing-keys-create-cmk.html

NEW QUESTION 66
- (Exam Topic 1)
A company needs to run a software package that has a license that must be run on the same physical host for the duration of Its use. The software package is only going to be used for 90 days The company requires patching and restarting of all instances every 30 days
How can these requirements be met using AWS?

A. Run a dedicated instance with auto-placement disabled.
B. Run the instance on a dedicated host with Host Affinity set to Host.
C. Run an On-Demand Instance with a Reserved Instance to ensure consistent placement.
D. Run the instance on a licensed host with termination set for 90 days.

Answer: B

Explanation:
Host Affinity is configured at the instance level. It establishes a launch relationship between an instance and a Dedicated Host. (This set which host the instance can run on) Auto-placement allows you to manage whether instances that you launch are launched onto a specific host, or onto any available host that has matching configurations. Auto-placement must be configured at the host level. (This sets which instance the host can run.) When affinity is set to Host, an instance launched onto a specific host always restarts on the same host if stopped. This applies to both targeted and untargeted launches.
https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/how-dedicated-hosts-work.html
When affinity is set to Off, and you stop and restart the instance, it can be restarted on any available host. However, it tries to launch back onto the last Dedicated Host on which it ran (on a best-effort basis).

NEW QUESTION 70
- (Exam Topic 1)
A company wants to change its internal cloud billing strategy for each of its business units. Currently, the cloud governance team shares reports for overall cloud spending with the head of each business unit. The company uses AWS Organizations lo manage the separate AWS accounts for each business unit. The existing tagging standard in Organizations includes the application, environment, and owner. The cloud governance team wants a centralized solution so each business unit receives monthly reports on its cloud spending. The solution should also send notifications for any cloud spending that exceeds a set threshold.
Which solution is the MOST cost-effective way to meet these requirements?

A. Configure AWS Budgets in each account and configure budget alerts that are grouped by application, environment, and owne
B. Add each business unit to an Amazon SNS topic for each aler
C. Use Cost Explorer in each account to create monthly reports for each business unit.
D. Configure AWS Budgets in the organization's master account and configure budget alerts that are grouped by application, environment, and owne
E. Add each business unit to an Amazon SNS topic for each aler
F. Use Cost Explorer in the organization's master account to create monthly reports for each business unit.
G. Configure AWS Budgets in each account and configure budget alerts lhat are grouped by application, environment, and owne
H. Add each business unit to an Amazon SNS topic for each aler

I. Use the AWS Billing and Cost Management dashboard in each account to create monthly reports for each business unit.
J. Enable AWS Cost and Usage Reports in the organization's master account and configure reports grouped by application, environment, and owne
K. Create an AWS Lambda function that processes AWS Cost and Usage Reports, sends budget alerts, and sends monthly reports to each business unit's email list.

Answer: B

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

NEW QUESTION 71
- (Exam Topic 1)
A web application is hosted in a dedicated VPC that is connected to a company's on-premises data center over a Site-to-Site VPN connection. The application is accessible from the company network only. This is a temporary non-production application that is used during business hours. The workload is generally low with occasional surges.
The application has an Amazon Aurora MySQL provisioned database cluster on the backend. The VPC has an internet gateway and a NAT gateways attached. The web servers are in private subnets in an Auto Scaling group behind an Elastic Load Balancer. The web servers also upload data to an Amazon S3 bucket through the internet.
A solutions architect needs to reduce operational costs and simplify the architecture. Which strategy should the solutions architect use?

A. Review the Auto Scaling group settings and ensure the scheduled actions are specified to operate the Amazon EC2 instances during business hours onl
B. Use 3-year scheduled Reserved Instances for the web server EC2 instance
C. Detach the internet gateway and remove the NAT gateways from the VP
D. Use an Aurora Servertess database and set up a VPC endpoint for the S3 bucket.
E. Review the Auto Scaling group settings and ensure the scheduled actions are specified to operate the Amazon EC2 instances during business hours onl
F. Detach the internet gateway and remove the NAT gateways from the VP
G. Use an Aurora Servertess database and set up a VPC endpoint for the S3 bucket, then update the network routing and security rules and policies related to the changes.
H. Review the Auto Scaling group settings and ensure the scheduled actions are specified to operate the Amazon EC2 instances during business hours onl
I. Detach the internet gateway from the VPC, and use an Aurora Servertess databas
J. Set up a VPC endpoint for the S3 bucket, then update the network routing and security rules and policies related to the changes.
K. Use 3-year scheduled Reserved Instances for the web server Amazon EC2 instance
L. Remove the NAT gateways from the VPC, and set up a VPC endpoint for the S3 bucke
M. Use Amazon
N. CloudWatch and AWS Lambda to stop and start the Aurora DB cluster so it operates during business hours onl
O. Update the network routing and security rules and policies related to the changes.

Answer: B

Explanation:
The application is accessible from the company network only remove NAT and IGW, application - S3 with VPC endpoint. Non-Production application no need to go for Reserved instances
To build site-to-site vpn, you don't need internet gateway. Instead, customer gateway is needed.
https://docs.aws.amazon.com/vpn/latest/s2svpn/SetUpVPNConnections.html#vpn-create-cgw

NEW QUESTION 76
- (Exam Topic 1)
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 to block clients "hat submit more than five requests per da
C. Associate the web ACL with the CloudFront distributio
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 execute 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 distributio
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. 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

NEW QUESTION 80
- (Exam Topic 2)
A company is running an application in the AWS Cloud. The application uses AWS Lambda functions and Amazon Elastic Container Service (Amazon ECS) containers that run with AWS Fargate technology as its primary compute. The load on the application is irregular. The application experiences long periods of no usage, followed by sudden and significant increases and decreases in traffic. The application is write-heavy and stores data in an Amazon Aurora MySQL database. The database runs on an Amazon RDS memory optimized D8 instance that is not able to handle the load.
What is the MOST cost-effective way for the company to handle the sudden and significant changes in traffic?

A. Add additional read replicas to the databas
B. Purchase Instance Savings Plans and RDS Reserved Instances.
C. Migrate the database to an Aurora multi-master DB cluste
D. Purchase Instance Savings Plans.
E. Migrate the database to an Aurora global database Purchase Compute Savings Plans and RDS Reserved Instances
F. Migrate the database to Aurora Serverless v1. Purchase Compute Savings Plans

Answer: D

NEW QUESTION 85
- (Exam Topic 2)
A company is planning to migrate an application from on premises to the AWS Cloud. The company will begin the migration by moving the application's underlying data storage to AWS The application data is stored on a shared tie system on premises, and the application servers connect to the shared We system through SMB.
A solutions architect must implement a solution that uses an Amazon S3 bucket tor shared storage Until the application Is fully migrated and code is rewritten to use native Amazon S3 APIs, the application must continue to have access to the data through SMB The solutions architect must migrate the application data to AWS to its new location while still allowing the on-premises application to access the data.
Which solution will meet these requirements?

A. Create a new Amazon FSx for Windows File Server fie system Configure AWS DataSync with onelocation tor the on-premises file share and one location for the new Amazon FSx file system Create a new DataSync task to copy the data from the on-premises file share location to the Amazon FSx file system
B. Create an S3 bucket for the applicatio
C. Copy the data from the on-premises storage to the S3 bucket
D. Deploy an AWS Server Migration Service (AWS SMS) VM to the on-premises environmen
E. Use AWS SMS to migrate the file storage server from on premises to an Amazon EC2 instance
F. Create an S3 bucket for the applicatio
G. Deploy a new AWS Storage Gateway Me gateway on anon-premises V
H. Create a new file share that stores data in the S3 bucket and is associated with the tie gatewa
I. Copy the data from the on-premises storage to the new file gateway endpoint.

Answer: A

NEW QUESTION 89
- (Exam Topic 2)
A company is deploying a distributed in-memory database on a fleet of Amazon EC2 instances. The fleet consists of a primary node and eight worker nodes. The primary node is responsible for monitoring cluster health, accepting user requests, distributing user requests to worker nodes and sending an aggregate response back to a client. Worker nodes communicate with each other to replicate data partitions.
The company requires the lowest possible networking latency to achieve maximum performance. Which solution will meet these requirements?

A. Launch memory optimized EC2 instances in a partition placement group
B. Launch compute optimized EC2 instances in a partition placement group
C. Launch memory optimized EC2 instances in a cluster placement group
D. Launch compute optimized EC2 instances in a spread placement group.

Answer: B

NEW QUESTION 94
- (Exam Topic 2)
A company that runs applications on AWS recently subscribed to a new software-as-a-service (SaaS) data vendor. The vendor provides the data by way of a REST API that the vendor hosts in its AWS environment The vendor offers multiple options for connectivity to the API and Is working with the company to find the best way to connect.
The company's AWS account does not allow outbound internet access from Its AWS environment The vendor's services run on AWS in the same AWS Region as the company's applications
A solutions architect must Implement connectivity to the vendor's API so that the API is highly available In the company's VPC.
Which solution will meet these requirements?

A. Connect to the vendor's public API address for the data service.
B. Connect to the vendor by way of a VPC peering connection between the vendor's VPC and the company's VPC
C. Connect to the vendor by way of a VPC endpoint service that uses AWS PrivateLink
D. Connect to a public bastion host that the vendor provides Tunnel the API traffic.

Answer: C

NEW QUESTION 97
- (Exam Topic 2)
A company's interactive web application uses an Amazon CloudFront distribution to serve images from an Amazon S3 bucket. Occasionally, third-party tools ingest corrupted images into the S3 bucket. This image corruption causes a poor user experience in the application later. The company has successfully implemented and tested Python logic to detect corrupt images.
A solutions architect must recommend a solution to integrate the detection logic with minimal latency between the ingestion and serving.
Which solution will meet these requirements?

A. Use a Lambda@Edge function that is invoked by a viewer-response event.
B. Use a Lambda@Edge function that is invoked by an origin-response event.
C. Use an S3 event notification that invokes an AWS Lambda function.
D. Use an S3 event notification that invokes an AWS Step Functions state machine.

Answer: A

NEW QUESTION 101
- (Exam Topic 2)
A video processing company has an application that downloads images from an Amazon S3 bucket, processes the images, stores a transformed image in a second S3 bucket, and updates metadata about the image in an Amazon DynamoDB table. The application is written in Node.js and runs by using an AWS Lambda function. The Lambda function is invoked when a new image is uploaded to Amazon S3.
The application ran without incident for a while. However, the size of the images has grown significantly. The Lambda function is now failing frequently with timeout errors. The function timeout is set to its maximum value. A solutions architect needs to refactor the application's architecture to prevent invocation failures. The company does not want to manage the underlying infrastructure.
Which combination of steps should the solutions architect take to meet these requirements? (Choose two.)

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

Answer: DE

NEW QUESTION 102
- (Exam Topic 2)
A company runs a content management application on a single Windows Amazon EC2 instance in a development environment. The application reads and writes static content to a 2 TB Amazon Elastic Block Store (Amazon EBS) volume that is attached to the instance as the root device. The company plans to deploy this application in production as a highly available and fault-tolerant solution that runs on at least three EC2 instances across multiple Availability Zones.
A solutions architect must design a solution that joins all the instances that run the application to an Active Directory domain. The solution also must implement Windows ACLs to control access to file contents. The application always must maintain exactly the same content on all running instances at any given point in time.
Which solution will meet these requirements with the LEAST management overhead?

A. Create an Amazon Elastic File System (Amazon EFS) file shar
B. Create an Auto Scaling group that extends across three Availability Zones and maintains a minimum size of three instance
C. Implement a user data script to install the application, join the instance to the AD domain, and mount the EFS file share.
D. Create a new AMI from the current EC2 instance that is runnin
E. Create an Amazon FSx for Lustre file syste
F. Create an Auto Scaling group that extends across three Availability Zones and maintains a minimum size of three instance
G. Implement a user data script to join the instance to the AD domain and mount the FSx for Lustre file system.
H. Create an Amazon FSx for Windows File Server file syste
I. Create an Auto Scaling group that extends across three Availability Zones and maintains a minimum size of three instance
J. Implement a user data script to install the application and mount the FSx for Windows File Server file syste
K. Perform a seamless domain join to join the instance to the AD domain.
L. Create a new AMI from the current EC2 instance that is runnin
M. Create an Amazon Elastic File System (Amazon EFS) file syste
N. Create an Auto Scaling group that extends across three Availability Zones and maintains a minimum size of three instance
O. Perform a seamless domain join to join the instance to the AD domain.

Answer: B

NEW QUESTION 105
- (Exam Topic 2)
A company has a web application that allows users to upload short videos. The videos are stored on Amazon EBS volumes and analyzed by custom recognition software for categorization.
The website contains stat c content that has variable traffic with peaks in certain months. The architecture consists of Amazon EC2 instances running in an Auto Scaling group for the web application and EC2 instances running in an Auto Scaling group to process an Amazon SQS queue The company wants to re-architect the application to reduce operational overhead using AWS managed services where possible and remove dependencies on third-party software.
Which solution meets these requirements?

A. Use Amazon ECS containers for the web application and Spot Instances for the Auto Scaling group that processes the SQS queu
B. Replace the custom software with Amazon Recognition to categorize the videos.
C. Store the uploaded videos n Amazon EFS and mount the file system to the EC2 instances for Te web applicatio
D. Process the SOS queue with an AWS Lambda function that calls the Amazon Rekognition API to categorize the videos.
E. Host the web application in Amazon S3. Store the uploaded videos in Amazon S3. Use S3 event notifications to publish events to the SQS queue Process the SQS queue with an AWS Lambda function that calls the Amazon Rekognition API to categorize the videos.
F. Use AWS Elastic Beanstalk to launch EC2 instances in an Auto Scaling group for the web application and launch a worker environment to process the SQS queue Replace the custom software with Amazon Rekognition to categorize the videos.

Answer: C

NEW QUESTION 106
- (Exam Topic 2)
A company is using an Amazon CloudFront distribution to distribute both static and dynamic content from a web application running behind an Application Load Balancer The web application requires user authorization and session tracking tor dynamic content The CloudFront distribution has a single cache behavior configured to forward the Authorization, Host, and Agent HTTP allow list headers and a session cookie to the origin All other cache behavior settings are set to their default value

A valid ACM certificate is applied to the CloudFront distribution with a matching CNAME in the distribution settings The ACM certificate is also applied to the HTTPS listener for the Application Load Balancer The CloudFront origin protocol policy is set to HTTPS only Analysis of the cache statistics report shows that the miss rate for this distribution is very high

What can the solutions architect do to improve the cache hit rate for this distribution without causing the SSL/TLS handshake between CloudFront and the Application Load Balancer to fail?

A. Create two cache behaviors for static and dynamic content Remove the user-Agent and Host HTTP headers from the allow list headers section on both of the cache behaviors Remove the session cookie from the allow list cookies section and the Authorization HTTP header from the allow list headers section for cache behavior configured for static content
B. Remove the user-Agent and Authorization HTTP headers from the allow list headers section of the cache behaviou
C. Then update the cache behaviour to use resigned cookies for authorization
D. Remove the Host HTTP header from the allow list headers section and remove the session cookie from the allow list cookies section for the default cache behaviour Enable automatic object compression and use Lambda@Edge viewer request events for user authorization
E. Create two cache behaviours for static and dynamic content Remove the User-Agent HTTP header from the allow list headers section on both of the cache behavioursRemove the session cookie from the allow list cookies section and the Authorization HTTP header from the allow list headers section for cache behaviour configured for static content

Answer: D

Explanation:
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/understanding-the-cache-key.html Removing the host header will result in failed flow between CloudFront and ALB, because they have same certificate.

NEW QUESTION 108
- (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/JpostId) 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 making the comments appear 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/commentsV{commentId} 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

NEW QUESTION 111
- (Exam Topic 2)
During an audit, a security team discovered that a development team was putting IAM user secret access keys in their code and then committing it to an AWS CodeCommit repository . The security team wants to automatically find and remediate instances of this security vulnerability
Which solution will ensure that the credentials are appropriately secured automatically?

A. Run a script nightly using AWS Systems Manager Run Command to search for credentials on the development instances If found use AWS Secrets Manager to rotate the credentials.
B. Use a scheduled AWS Lambda function to download and scan the application code from CodeCommit If credentials are found, generate new credentials and store them in AWS KMS
C. Configure Amazon Macie to scan for credentials in CodeCommit repositories If credentials are found, trigger an AWS Lambda function to disable the credentials and notify the user
D. Configure a CodeCommit trigger to invoke an AWS Lambda function to scan new code submissions for credentials If credentials are found, disable them in AWS IAM and notify the user.

Answer: A

NEW QUESTION 112
- (Exam Topic 2)
A company is planning to migrate an Amazon RDS for Oracle database to an RDS for PostgreSQL DB instance in another AWS account A solutions architect needs to design a migration strategy that will require no downtime and that will minimize the amount of time necessary to complete the migration The migration strategy must replicate all existing data and any new data that is created during the migration The target database must be identical to the source database at completion of the migration process
All applications currently use an Amazon Route 53 CNAME record as their endpoint for communication with the RDS for Oracle DB instance The RDS for Oracle DB instance is in a private subnet
Which combination of steps should the solutions architect take to meet these requirements? (Select THREE)

A. Create a new RDS for PostgreSQL DB instance in the target account Use the AWS Schema Conversion Tool (AWS SCT) to migrate the database schema from the source database to the target database.
B. Use the AWS Schema Conversion Tool (AWS SCT) to create a new RDS for PostgreSQL DB instance in the target account with the schema and initial data from the source database
C. Configure VPC peering between the VPCs in the two AWS accounts to provide connectivity to both DB instances from the target accoun
D. Configure the security groups that are attached to each DB instance to allow traffic on the database port from the VPC in the target account
E. Temporarily allow the source DB instance to be publicly accessible to provide connectivity from the VPC in the target account Configure the security groups that

are attached to each DB instance to allow traffic on the database port from the VPC in the target account.
F. Use AWS Database Migration Service (AWS DMS) in the target account to perform a full load plus change data capture (CDC) migration from the source database to the target database When the migration is complete, change the CNAME record to point to the target DB instance endpoint
G. Use AWS Database Migration Service (AWS DMS) in the target account to perform a change data capture (CDC) migration from the source database to the target database When the migration is complete change the CNAME record to point to the target DB instance endpoint

Answer: BCE

NEW QUESTION 114
- (Exam Topic 2)
A solutions architect is working with a company that is extremely sensitive to its IT costs and wishes to implement controls that will result in a predictable AWS spend each month Which combination ot steps can help the company control and monitor its monthly AWS usage to achieve a cost that is as close as possible to the target amount? (Select THREE.)

A. Implement an IAM policy that requires users to specify a 'workload' tag for cost allocation when launching Amazon EC2 instances
B. Contact AWS Support and ask that they apply limits to the account so that users are not able to launch more than a certain number of instance types
C. Purchase all upfront Reserved Instances that cover 100% of the account's expected Amazon EC2 usage
D. Place conditions in the users' IAM policies that limit the number of instances they are able to launch
E. Define 'workload' as a cost allocation tag in the AWS Billing and Cost Management console
F. Set up AWS Budgets to alert and notify when a given workload is expected to exceed a defined cost

Answer: AEF

NEW QUESTION 116
- (Exam Topic 2)
An external audit of a company's serverless application reveals IAM policies that grant too many permissions. These policies are attached to the company's AWS Lambda execution roles. Hundreds of the company's Lambda functions have broad access permissions, such as full access to Amazon S3 buckets and Amazon DynamoDB tables. The company wants each function to have only the minimum permissions that the function needs to complete its task.
A solutions architect must determine which permissions each Lambda function needs.
What should the solutions architect do to meet this requirement with the LEAST amount of effort?

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

Answer: B

Explanation:
IAM Access Analyzer helps you identify the resources in your organization and accounts, such as Amazon S3 buckets or IAM roles, shared with an external entity. This lets you identify unintended access to your resources and data, which is a security risk. IAM Access Analyzer identifies resources shared with external principals by using logic-based reasoning to analyze the resource-based policies in your AWS environment.
https://docs.aws.amazon.com/IAM/latest/UserGuide/what-is-access-analyzer.html

---

NEW QUESTION 117
- (Exam Topic 2)
A company has VPC flow logs enabled for its NAT gateway. The company is seeing Action = ACCEPT for inbound traffic that comes from public IP address 198.51.100.2 destined for a private Amazon EC2 instance.
A solutions architect must determine whether the traffic represents unsolicited inbound connections from the internet. The first two octets of the VPC CIDR block are 203.0.
Which set of steps should the solutions architect take to meet these requirements?

A. Open the AWS CloudTrail consol
B. Select the log group that contains the NAT gateway's elastic network interface and the private instance's elastic network interfac
C. Run a query to filter with the destination address set as "like 203.0" and the source address set as "like 198.51.100.2". Run the stats command to filter the sum of bytes transferred by the source address and the destination address.
D. Open the Amazon CloudWatch consol
E. Select the log group that contains the NAT gateway's elastic network interface and the private instance's elastic network interfac
F. Run a query to filter with the destination address set as "like 203.0" and the source address set as "like 198.51.100.2". Run the stats command to filter the sum of bytes transferred by the source address and the destination address.
G. Open the AWS CloudTrail consol
H. Select the log group that contains the NAT gateway's elastic network interface and the private instance's elastic network interfac
I. Run a query to filter with the destination address set as "like 198.51.100.2" and the source address set as "like 203.0". Run the stats command to filter the sum of bytes transferred by the source address and the destination address.
J. Open the Amazon CloudWatch consol
K. Select the log group that contains the NAT gateway's elasticnetwork interface and the private instance's elastic network interfac
L. Run a query to filter with the destination address set as "like 198.51.100.2" and the source address set as "like 203.0". Run the stats command to filter the sum

of bytes transferred by the source address and the destination address.

Answer: D

Explanation:
https://aws.amazon.com/premiumsupport/knowledge-center/vpc-analyze-inbound-traffic-nat-gateway/ by Cloudxie says "select appropriate log"

NEW QUESTION 119
- (Exam Topic 2)
A company has its cloud infrastructure on AWS A solutions architect needs to define the infrastructure as code. The infrastructure is currently deployed in one AWS Region. The company's business expansion plan includes deployments in multiple Regions across multiple AWS accounts
What should the solutions architect do to meet these requirements?

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

Answer: C

Explanation:
https://aws.amazon.com/blogs/aws/new-use-aws-cloudformation-stacksets-for-multiple-accounts-in-an-aws-orga

NEW QUESTION 123
- (Exam Topic 2)
A company is developing a gene reporting device that will collect genomic information to assist researchers with collecting large samples of data from a diverse population. The device will push 8 KB of genomic data every second to a data platform that will need to process and analyze the data and provide information back to researchers The data platform must meet the following requirements:
• Provide near-real-time analytics of the inbound genomic data
• Ensure the data is flexible, parallel, and durable
• Deliver results of processing to a data warehouse
Which strategy should a solutions architect use to meet these requirements?

A. Use Amazon Kinesis Data Firehose to collect the inbound sensor data analyze the data with Kinesis client
B. and save the results to an Amazon RDS instance
C. Use Amazon Kinesis Data Streams to collect the inbound sensor data analyze the data with Kinesis clients and save the results to an Amazon Redshift duster using Amazon EMR
D. Use Amazon S3 to collect the inbound device data analyze the data from Amazon SOS with Kinesis and save the results to an Amazon Redshift duster
E. Use an Amazon API Gateway to put requests into an Amazon SQS queue analyze the data with an AWS Lambda function and save the results » an Amazon Redshift duster using Amazon EMR

Answer: A

NEW QUESTION 126
- (Exam Topic 2)
A digital marketing company has multiple AWS accounts that belong to various teams. The creative team uses an Amazon S3 bucket in its AWS account to securely store images and media files that are used as content for the company's marketing campaigns. The creative team wants to share the S3 bucket with the strategy team so that the strategy team can view the objects.
A solutions architect has created an IAM role that is named strategy_reviewer in the Strategy account. The solutions architect also has set up a custom AWS Key Management Service (AWS KMS) key in the Creative account and has associated the key with the S3 bucket. However, when users from the Strategy account assume the IAM role and try to access objects in the S3 bucket, they receive an Account.
The solutions architect must ensure that users in the Strategy account can access the S3 bucket. The solution must provide these users with only the minimum permissions that they need.
Which combination of steps should the solutions architect take to meet these requirements? (Select THREE.)

A. Create a bucket policy that includes read permissions for the S3 bucke
B. Set the principal of the bucket policy to the account ID of the Strategy account
C. Update the strategy_reviewer IAM role to grant full permissions for the S3 bucket and to grant decrypt permissions for the custom KMS key.
D. Update the custom KMS key policy in the Creative account to grant decrypt permissions to the strategy_reviewer IAM role.
E. Create a bucket policy that includes read permissions for the S3 bucke
F. Set the principal of the bucket policy to an anonymous user.
G. Update the custom KMS key policy in the Creative account to grant encrypt permissions to the strategy_reviewer IAM role.
H. Update the strategy_reviewer IAM role to grant read permissions for the S3 bucket and to grant decrypt permissions for the custom KMS key

Answer: ACF

Explanation:
https://aws.amazon.com/premiumsupport/knowledge-center/cross-account-access-denied-error-s3/

NEW QUESTION 131
- (Exam Topic 2)
A company has automated the nightly retraining ot its machine learning models by using AWS Step Functions. The workflow consists of multiple steps that use AWS Lambda. Each step can fail for various reasons, and any failure causes a failure of the overall workflow.
A review reveals that the retraining has failed multiple nights in a row without the company noticing the failure. A solutions architect needs to improve the workflow so that notifications are sent for all types of failures in the retraining process.
Which combination of steps should the solutions architect take to meet these requirements? (Select THREE.)

A. Create an Amazon Simple Notification Service {Amazon SNS) topic with a subscription of type "Email" that targets the team's mailing list.

B. Create a task named "Email" that forwards the input arguments to the SNS topic
C. Add a Catch field to all Tas
D. Ma
E. and Parallel states that have a statement of "ErrorEquals": ["states.all"] and "Next": "Email".
F. Add a new email address to Amazon Simple Email Service (Amazon SES). Verify the email address.
G. Create a task named "Email" that forwards the input arguments to the SES email address
H. Add a Catch field to all Task, Map, and Parallel states that have a statement of "ErrorEquals": ["states.Bun time"] and "Next": "Email".

Answer: BCD

NEW QUESTION 132
- (Exam Topic 2)
A company is deploying a third-party firewall appliance solution from AWS Marketplace to monitor and protect traffic that leaves the company's AWS environments.
The company wants to deploy this appliance into a shared services VPC and route all outbound internet-bound traffic through the appliances.
A solutions architect needs to recommend a deployment method that prioritizes reliability and minimizes failover time between firewall appliances within a single
AWS Region. The company has set up routing from the shared services VPC to other VPCs.
Which steps should the solutions architect recommend to meet these requirements? (Select THREE)

A. Deploy two firewall appliances into the shared services VP
B. each in a separate Availability Zone
C. Create a new Network Load Balancer in the shared services VPC Create a new target group, and attach it to the new Network Load Balancer Add each of the firewall appliance instances to the target group.
D. Create a new Gateway Load Balancer in the shared services VPC Create a new target group, and attach it to the new Gateway Load Balancer Add each of the firewall appliance instances to the target group
E. Create a VPC interface endpoint Add a route to the route table in the shared services VP
F. Designate the new endpoint as the next hop for traffic that enters the shared services VPC from other VPCs.
G. Deploy two firewall appliances into the shared services VP
H. each in the same Availability Zone

Answer: AC

NEW QUESTION 136
- (Exam Topic 2)
An online magazine will launch Its latest edition this month. This edition will be the first to be distributed globally. The magazine's dynamic website currently uses
an Application Load Balancer in front of the web tier a fleet of Amazon EC2 instances for web and application servers, and Amazon Aurora MySQL. Portions of the
website include static content and almost all traffic is read-only
The magazine is expecting a significant spike m internet traffic when the new edition is launched Optimal performance is a top priority for the week following the
launch
Which combination of steps should a solutions architect take to reduce system response antes for a global audience? (Select TWO)

A. Use logical cross-Region replication to replicate the Aurora MySQL database to a secondary Region Replace the web servers with Amazon S3 Deploy S3
buckets in cross-Region replication mode
B. Ensure the web and application tiers are each m Auto Scaling group
C. Introduce an AWS Direct Connect connection Deploy the web and application tiers in Regions across the world
D. Migrate the database from Amazon Aurora to Amazon RDS for MySQ
E. Ensure all three of the application tiers—we
F. application, and database—are in private subnets.
G. Use an Aurora global database for physical cross-Region replicatio
H. Use Amazon S3 with cross-Region replication for static content and resource
I. Deploy the web and application tiers in Regions across the world
J. Introduce Amazon Route 53 with latency-based routing and Amazon CloudFront distribution
K. Ensure me web and application tiers are each in Auto Scaling groups

Answer: DE

NEW QUESTION 140
- (Exam Topic 2)
A retail company has an on-premises data center in Europe. The company also has a multi-Region AWS presence that includes the eu-west-1 and us-east-1
Regions. The company wants to be able to route network traffic from its on-premises infrastructure into VPCs in either of those Regions. The company also needs
to support traffic that is routed directly between VPCs in those Regions. No single points of failure can exist on the network.
The company already has created two 1 Gbps AWS Direct Connect connections from its on-premises data center. Each connection goes into a separate Direct
Connect location in Europe for high availability. These two locations are named DX-A and DX-B, respectively. Each Region has a single AWS Transit Gateway that
is configured to route all inter-VPC traffic within that Region.
Which solution will meet these requirements?

A. Create a private VIF from the DX-A connection into a Direct Connect gatewa
B. Create a private VIF from the DX-B connection into the same Direct Connect gateway for high availabilit
C. Associate both the eu-west-1 and us-east-1 transit gateways with the Direct Connect gatewa
D. Peer the transit gateways with each other to support cross-Region routing.
E. Create a transit VIF from the DX-A connection into a Direct Connect gatewa
F. Associate the eu-west-1 transit gateway with this Direct Connect gatewa
G. Create a transit VIF from the DX-B connection into a separate Direct Connect gatewa
H. Associate the us-east-1 transit gateway with this separate Direct Connect gatewa
I. Peer the Direct Connect gateways with each other to support high availability and cross-Region routing.
J. Create a transit VIF from the DX-A connection into a Direct Connect gatewa
K. Create a transit VIF from the DX-B connection into the same Direct Connect gateway for high availabilit
L. Associate both the eu-west-1 and us-east-1 transit gateways with this Direct Connect gatewa
M. Configure the Direct Connect gateway to route traffic between the transit gateways.
N. Create a transit VIF from the DX-A connection into a Direct Connect gatewa
O. Create a transit VIF from the DX-B connection into the same Direct Connect gateway for high availabilit

P. Associate both the eu-west-1 and us-east-1 transit gateways with this Direct Connect gatewa
Q. Peer the transit gateways with each other to support cross-Region routing.

Answer: D

NEW QUESTION 143
- (Exam Topic 2)
A company gives users the ability to upload images from a custom application. The upload process invokes an AWS Lambda function that processes and stores the image in an Amazon S3 bucket. The application invokes the Lambda function by using a specific function version ARN.
The Lambda function accepts image processing parameters by using environment variables. The company often adjusts the environment variables of the Lambda function to achieve optimal image processing output. The company tests different parameters and publishes a new function version with the updated environment variables after validating results. This update process also requires frequent changes to the custom application to invoke the new function version ARN. These changes cause interruptions for users.
A solutions architect needs to simplify this process to minimize disruption to users. Which solution will meet these requirements with the LEAST operational overhead?

A. Directly modify the environment variables of the published Lambda function versio
B. Use the SLATEST version to test image processing parameters.
C. Create an Amazon DynamoDB table to store the image processing parameter
D. Modify the Lambda function to retrieve the image processing parameters from the DynamoDB table.
E. Directly code the image processing parameters within the Lambda function and remove the environment variable
F. Publish a new function version when the company updates the parameters.
G. Create a Lambda function alia
H. Modify the client application to use the function alias AR
I. Reconfigure the Lambda alias to point to new versions of the function when the company finishes testing.

Answer: D

NEW QUESTION 148
- (Exam Topic 2)
A company is using Amazon OpenSearch Service to analyze data. The company loads data into an OpenSearch Service cluster with 10 data nodes from an Amazon S3 bucket that uses S3 Standard storage. The data resides in the cluster for 1 month for read-only analysis. After 1 month, the company deletes the index that contains the data from the cluster. For compliance purposes, the company must retain a copy of all input data.
The company is concerned about ongoing costs and asks a solutions architect to recommend a new solution.
Which solution will meet these requirements MOST cost-effectively?

A. Replace all the data nodes with UltraWarm nodes to handle the expected capacit
B. Transition the input data from S3 Standard to S3 Glacier Deep Archive when the company loads the data into the cluster.
C. Reduce the number of data nodes in the cluster to 2 Add UltraWarm nodes to handle the expected capacit
D. Configure the indexes to transition to UltraWarm when OpenSearch Service ingests the dat
E. Transition the input data to S3 Glacier Deep Archive after 1 month by using an S3 Lifecycle policy.
F. Reduce the number of data nodes in the cluster to 2. Add UltraWarm nodes to handle the expected capacit
G. Configure the indexes to transition to UltraWarm when OpenSearch Service ingests the dat
H. Add cold storage nodes to the cluster Transition the indexes from UltraWarm to cold storag
I. Delete the input data from the S3 bucket after 1 month by using an S3 Lifecycle policy.
J. Reduce the number of data nodes in the cluster to 2. Add instance-backed data nodes to handle the expected capacit
K. Transition the input data from S3 Standard to S3 Glacier Deep Archive when the company loads the data into the cluster.

Answer: B

NEW QUESTION 153
- (Exam Topic 2)
A data analytics company has an Amazon Redshift cluster that consists of several reserved nodes. The duster is experiencing unexpected bursts of usage because a team of employees is compiling a deep audit analysis report The queries to generate the report are complex read queries and are CPU intensive.
Business requirements dictate that the cluster must be able to service read and write queries at at) times A solutions architect must devise a solution that accommodates the bursts of usage
Which solution meets these requirements MOST cost-effectively?

A. Provision an Amazon EMR duster Offload the complex data processing tasks
B. Deploy an AWS Lambda function to add capacity to the Amazon Redshift cluster by using a classic resize operation when the duster's CPU metrics in Amazon CloudWatch reach 80%.
C. Deploy an AWS Lambda function to add capacity to the Amazon Redshift duster by using an elastic resize operation when the duster's CPU metrics in Amazon CloudWatch leach 80%.
D. Turn on the Concurrency Scaling feature for the Amazon Redshift duster

Answer: D

NEW QUESTION 155
- (Exam Topic 2)
A company is planning to migrate its business-critical applications from an on-premises data center to AWS. The company has an on-premises installation of a Microsoft SQL Server Always On cluster. The company wants to migrate to an AWS managed database service. A solutions architect must design a heterogeneous database migration on AWS.
Which solution will meet these requirements?

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

F. Use AWS DataSync to migrate data over the network between on-premises storage and Amazon S3. Set up Amazon RDS for MySQ
G. Use S3 integration with SQL Server features, such as BULK INSERT.

Answer: C

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

NEW QUESTION 160
- (Exam Topic 2)
A company is in the process of implementing AWS Organizations to constrain its developers to use only Amazon EC2. Amazon S3 and Amazon DynamoDB. The developers account resides In a dedicated organizational unit (OU). The solutions architect has implemented the following SCP on the developers account:

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "AllowEC2",
            "Effect": "Allow",
            "Action": "ec2:*",
            "Resource": "*"
        },
        {
            "Sid": "AllowDynamoDB",
            "Effect": "Allow",
            "Action": "dynamodb:*",
            "Resource": "*"
        },
        {
            "Sid": "AllowS3",
            "Effect": "Allow",
            "Action": "s3:*",
            "Resource": "*"
        }
    ]
}
```

When this policy is deployed, IAM users in the developers account are still able to use AWS services that are not listed in the policy. What should the solutions architect do to eliminate the developers' ability to use services outside the scope of this policy?

A. Create an explicit deny statement for each AWS service that should be constrained
B. Remove the Full AWS Access SCP from the developer account's OU
C. Modify the Full AWS Access SCP to explicitly deny all services
D. Add an explicit deny statement using a wildcard to the end of the SCP

Answer: B

NEW QUESTION 161
- (Exam Topic 2)
A large company runs workloads in VPCs that are deployed across hundreds of AWS accounts Each VPC consists of 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 VPCs Configure the required routing to allow access to the internet
B. Create a transit gateway and share it with the existing AWS accounts Attach existing VPCs to the transit gateway Configure the required routing to allow access to the internet
C. Create a transit gateway in every account Attach the NAT gateway to the transit gateways Configure the required routing to allow access to the internet
D. Create an AWS PrivateLink connection between the egress VPC and the spoke VPCs Configure the required routing to allow access to the internet

Answer: B

NEW QUESTION 166
- (Exam Topic 2)
A company has deployed an application to multiple environments in AWS. including production and testing the company has separate accounts for production and testing, and users are allowed to create additional application users for team members or services. as needed. The security team has asked the operations team tor better isolation between production and testing with centralized controls on security credentials and improved management of permissions between environments
Which of the following options would MOST securely accomplish this goal?

A. Create a new AWS account to hold user and service accounts, such as an identity account Create users and groups m the identity accoun
B. Create roles with appropriate permissions in the production and testing accounts Add the identity account to the trust policies for the roles
C. Modify permissions in the production and testing accounts to limit creating new IAM users to members of the operations team Set a strong IAM password policy on each account Create new IAM users and groups in each account to Limit developer access to just the services required to complete their job function.
D. Create a script that runs on each account that checks user accounts For adherence to a security policy.Disable any user or service accounts that do not comply.
E. Create all user accounts in the production account Create roles for access in me production account and testing account

F. Grant cross-account access from the production account to the testing account

Answer: A

NEW QUESTION 169
- (Exam Topic 2)
A company has more than 10.000 sensors that send data to an on-premises Apache Kafka server by using the Message Queuing Telemetry Transport (MQTT) protocol . The on-premises Kafka server transforms the data and then stores the results as objects in an Amazon S3 bucket
Recently, the Kafka server crashed. The company lost sensor data while the server was being restored A
solutions architect must create a new design on AWS that is highly available and scalable to prevent a similar occurrence
Which solution will meet these requirements?

A. Launch two Amazon EC2 instances to host the Kafka server in an active/standby configuration across two Availability Zone
B. Create a domain name in Amazon Route 53 Create a Route 53 failover policy Route the sensors to send the data to the domain name
C. Migrate the on-premises Kafka server to Amazon Managed Streaming for Apache Kafka (Amazon MSK). Create a Network Load Balancer (NLB) that points to the Amazon MSK broke
D. Enable NLB health checks Route the sensors to send the data to the NLB.
E. Deploy AWS IoT Core, and connect it to an Amazon Kinesis Data Firehose delivery stream Use an AWS Lambda function to handle data transformation Route the sensors to send the data to AWS IoT Core
F. Deploy AWS IoT Core, and launch an Amazon EC2 instance to host the Kafka server Configure AWS IoT Core to send the data to the EC2 instance Route the sensors to send the data to AWSIoT Core.

Answer: A

NEW QUESTION 170
- (Exam Topic 2)
A solutions architect has been assigned to migrate a 50 TB Oracle data warehouse that contains sales data from on-premises to Amazon Redshift Major updates to the sales data occur on the final calendar day of the month For the remainder of the month, the data warehouse only receives minor daily updates and is primarily used for reading and reporting Because of this the migration process must start on the first day of the month and must be complete before the next set of updates occur. This provides approximately 30 days to complete the migration and ensure that the minor daily changes have been synchronized with the Amazon Redshift data warehouse Because the migration cannot impact normal business network operations, the bandwidth allocated to the migration for moving data over the internet is 50 Mbps The company wants to keep data migration costs low
Which steps will allow the solutions architect to perform the migration within the specified timeline?

A. Install Oracle database software on an Amazon EC2 instance Configure VPN connectivity between AWS and the company's data center Configure the Oracle database running on Amazon EC2 to join the Oracle Real Application Clusters (RAC) When the Oracle database on Amazon EC2 finishes synchronizing, create an AWS DMS ongoing replication task to migrate the data from the Oracle database on Amazon EC2 to Amazon Redshift Verify the data migration is complete and perform the cut over to Amazon Redshift.
B. Create an AWS Snowball import job Export a backup of the Oracle data warehouse Copy the exported data to the Snowball device Return the Snowball device to AWS Create an Amazon RDS for Oracle database and restore the backup file to that RDS instance Create an AWS DMS task to migrate the data from the RDS for Oracle database to Amazon Redshift Copy daily incremental backups from Oracle in the data center to the RDS for Oracle database over the internet Verify the data migration is complete and perform the cut over to Amazon Redshift.
C. Install Oracle database software on an Amazon EC2 instance To minimize the migration time configure VPN connectivity between AWS and the company's data center by provisioning a 1 Gbps AWS Direct Connect connection Configure the Oracle database running on Amazon EC2 to be a read replica of the data center Oracle database Start the synchronization process between the company's on-premises data center and the Oracle database on Amazon EC2 When the Oracle database on Amazon EC2 is synchronized with the on-premises database create an AWS DMS ongoing replication task from the Oracle database read replica that is running on Amazon EC2 to Amazon Redshift Verify the data migration is complete and perform the cut over to Amazon Redshift.
D. Create an AWS Snowball import jo
E. Configure a server in the company€™s data center with an extraction agen
F. Use AWS SCT to manage the extraction agent and convert the Oracle schema to an Amazon Redshift schem
G. Create a new project in AWS SCT using the registered data extraction agen
H. Create a local task and an AWS DMS task in AWS SCT with replication of ongoing change
I. Copy data to the Snowball device and return the Snowball device to AW
J. Allow AWS DMS to copy data from Amazon S3 to Amazon Redshif
K. Verify that the data migration is complete and perform the cut over to Amazon Redshift.

Answer: D

Explanation:
Create an AWS Snowball import job. Configure a server in the company€™s data center with an extraction agent. Use AWS SCT to manage the extraction agent and convert the Oracle schema to an Amazon Redshift schema. Create a new project in AWS SCT using the registered data extraction agent. Create a local task and an AWS DMS task in AWS SCT with replication of ongoing changes. Copy data to the Snowball device and return the Snowball device to AWS. Allow AWS DMS to copy data from Amazon S3 to Amazon Redshift. Verify that the data migration is complete and perform the cut over to Amazon Redshift.
https://aws.amazon.com/getting-started/hands-on/migrate-oracle-to-amazon-redshift/

NEW QUESTION 175
- (Exam Topic 2)
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: C

NEW QUESTION 177
- (Exam Topic 2)
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 181
- (Exam Topic 2)
A global media company is planning a multi-Region deployment of an application. Amazon DynamoDB global tables will back the deployment to keep the user experience consistent across the two continents where users are concentrated. Each deployment will have a public Application Load Balancer (ALB). The company manages public DNS internally. The company wants to make the application available through an apex domain.
Which solution will meet these requirements with the LEAST effort?

A. Migrate public DNS to Amazon Route 53. Create CNAME records for the apex domain to point to the AL
B. Use a geolocation routing policy to route traffic based on user location.
C. Place a Network Load Balancer (NLB) in front of the AL
D. Migrate public DNS to Amazon Route 53.Create a CNAME record for the apex domain to point to the NLB's static IP addres
E. Use a geolocation routing policy to route traffic based on user location.
F. Create an AWS Global Accelerator accelerator with multiple endpoint groups that target endpoints in appropriate AWS Region
G. Use the accelerator's static IP address to create a record in public DNS for the apex domain.
H. Create an Amazon API Gateway API that is backed by AWS Lambda in one of the AWS Regions.Configure a Lambda function to route traffic to application deployments by using the round robin metho
I. Create CNAME records for the apex domain to point to the API's URL.

Answer: C

NEW QUESTION 182
- (Exam Topic 2)
A company is running several workloads in a single AWS account. A new company policy states that engineers can provision only approved resources and that engineers must use AWS CloudFormation to provision these resources. A solutions architect needs to create a solution to enforce the new restriction on the IAM role that the engineers use for access.
What should the solutions architect do to create the solution?

A. Upload AWS CloudFormation templates that contain approved resources to an Amazon S3 bucket.Update the IAM policy for the engineers' IAM role to only allow access to Amazon S3 and AWS CloudFormatio
B. Use AWS CloudFormation templates to provision resources.
C. Update the IAM policy for the engineers' IAM role with permissions to only allow provisioning of approved resources and AWS CloudFormatio
D. Use AWS CloudFormation templates to create stacks with approved resources.
E. Update the IAM policy for the engineers' IAM role with permissions to only allow AWS CloudFormation action
F. Create a new IAM policy with permission to provision approved resources, and assign the policy to a new IAM service rol
G. Assign the IAM service role to AWS CloudFormation during stack creation.
H. Provision resources in AWS CloudFormation stack
I. Update the IAM policy for the engineers' IAM role to only allow access to their own AWS CloudFormation stack.

Answer: B

NEW QUESTION 184
- (Exam Topic 2)
A company recently started hosting new application workloads in the AWS Cloud. The company is using Amazon EC2 instances. Amazon Elastic File System (Amazon EFS) file systems, and Amazon RDS DB instances.
To meet regulatory and business requirements, the company must make the following changes for data backups:
• Backups must be retained based on custom daily, weekly, and monthly requirements.
• Backups must be replicated to at least one other AWS Region immediately after capture.

• The backup solution must provide a single source of backup status across the AWS environment.
• The backup solution must send immediate notifications upon failure of any resource backup.
Which combination of steps will meet these requirements with the LEAST amount of operational overhead? (Select THREE.)

A. Create an AWS Backup plan with a backup rule for each of the retention requirements.
B. Configure an AWS Backup plan to copy backups to another Region.
C. Create an AWS Lambda function to replicate backups to another Region and send notification if a failure occurs.
D. Add an Amazon Simple Notification Service (Amazon SNS) topic to the backup plan to send a notification for finished jobs that have any status except BACKUP_JOB_COMPLETEO.
E. Create an Amazon Data Lifecycle Manager (Amazon DLM) snapshot lifecycle policy for each of the retention requirements.
F. Set up RDS snapshots on each database.

Answer: BDE

NEW QUESTION 189
- (Exam Topic 2)
A media company has a 30-TB repository of digital news videos These videos are stored on tape in an
on-premises tape library and referenced by a Media Asset Management (MAM) system The company wants to enrich the metadata for these videos in an automated fashion and put them into a searchable catalog by using a MAM feature The company must be able to search based on information in the video such as objects scenery items or people's faces A catalog is available that contains faces of people who have appeared in the videos that include an image of each person The company would like to migrate these videos to AWS
The company has a high-speed AWS Direct Connect connection with AWS and would like to move the MAM solution video content directly from its current file system
How can these requirements be met by using the LEAST amount of ongoing management overhead and causing MINIMAL disruption to the existing system"'

A. Set up an AWS Storage Gateway file gateway appliance on-premise
B. Use the MAM solution to extract the videos from the current archive and push them into the file gateway Use the catalog of faces to build a collection in Amazon Rekognition Build an AWS Lambda function that invokes the Rekognition Javascript SDK to have Rekognition pull the video from the Amazon S3 files backing the file gateway, retrieve the required metadata and push the metadata into the MAM solution
C. Set up an AWS Storage Gateway tape gateway appliance on-premises Use the MAM solution to extract the videos from the current archive and push them into the tape gateway Use the catalog of faces to build a collection in Amazon Rekognition Build an AWS Lambda function that invokes the Rekognition Javascript SDK to have Amazon Rekognition process the video in the tape gateway retrieve the required metadata, and push the metadata into the MAM solution
D. Configure a video ingestion stream by using Amazon Kinesis Video Streams Use the catalog of faces to build a collection in Amazon Rekognition Stream the videos from the MAM solution into Kinesis Video Streams Configure Amazon Rekognition to process the streamed videos Then, use a stream consumer to retrieve the required metadata and push the metadata into the MAM solution Configure the stream to store the videos in Amazon S3
E. Set up an Amazon EC2 instance that runs the OpenCV libranes Copy the videos, images, and face catalog from the on-premises library into an Amazon EBS volumemounted on this EC2 instance Process the videos to retrieve the required metadata, and push the metadata into the MAM solution, while also copying the video files to an Amazon S3 bucket

Answer: C

NEW QUESTION 191
- (Exam Topic 2)
A company has a media metadata extraction pipeline running on AWS. Notifications containing a reference to a file Amazon S3 are sent to an Amazon Simple Notification Service (Amazon SNS) topic The pipeline consists of a number of AWS Lambda functions that are subscribed to the SNS topic The Lambda functions extract the S3 file and write metadata to an Amazon RDS PostgreSQL DB instance.
Users report that updates to the metadata are sometimes stow to appear or are lost. During these times, the CPU utilization on the database is high and the number of failed Lambda invocations increases.
Which combination of actions should a solutions architect take to r-e'p resolve this issue? (Select TWO.)

A. Enable massage delivery status on the SNS topic Configure the SNS topic delivery policy to enable retries with exponential backoff
B. Create an Amazon Simple Queue Service (Amazon SOS) FIFO queue and subscribe the queue to the SNS topic Configure the Lambda functions to consume messages from the SQS queue.
C. Create an RDS proxy for the RDS instance Update the Lambda functions to connect to the RDS instance using the proxy.
D. Enable the RDS Data API for the RDS instanc
E. Update the Lambda functions to connect to the RDS instance using the Data API
F. Create an Amazon Simple Queue Service (Amazon SQS) standard queue for each Lambda function and subscribe the queues to the SNS topi
G. Configure the Lambda functions to consume messages from their respective SQS queue.

Answer: CE

NEW QUESTION 192
- (Exam Topic 2)
A company recently completed the migration from an on-premises data center to the AWS Cloud by using a replatforming strategy. One of the migrated servers is running a legacy Simple Mail Transfer Protocol (SMTP) service that a critical application relies upon. The application sends outbound email messages to the company's customers. The legacy SMTP server does not support TLS encryption and uses TCP port 25. The application can use SMTP only.
The company decides to use Amazon Simple Email Service (Amazon SES) and to decommission the legacy SMTP server. The company has created and validated the SES domain. The company has lifted the SES limits.
What should the company do to modify the application to send email messages from Amazon SES?

A. Configure the application to connect to Amazon SES by using TLS Wrappe
B. Create an IAM role that has ses:SendEmail and ses:SendRawEmail permission
C. Attach the IAM role to an Amazon EC2 instance.
D. Configure the application to connect to Amazon SES by using STARTTL
E. Obtain Amazon SES SMTP credential
F. Use the credentials to authenticate with Amazon SES.
G. Configure the application to use the SES API to send email message
H. Create an IAM role that has ses:SendEmail and ses:SendRawEmail permission
I. Use the IAM role as a service role for Amazon SES.
J. Configure the application to use AWS SDKs to send email message

K. Create an IAM user for Amazon SE
L. Generate API access key
M. Use the access keys to authenticate with Amazon SES.

Answer: A

NEW QUESTION 197
- (Exam Topic 2)
A company wants to deploy an API to AWS. The company plans to run the API on AWS Fargate behind a load balancer. The API requires the use of header-based routing and must be accessible from on-premises networks through an AWS Direct Connect connection and a private VIF.
The company needs to add the client IP addresses that connect to the API to an allow list in AWS. The company also needs to add the IP addresses of the API to the allow list. The company's security team will allow /27 CIDR ranges to be added to the allow list. The solution must minimize complexity and operational overhead.
Which solution will meet these requirements?

A. Create a new Network Load Balancer (NLB) in the same subnets as the Fargate task deployments.Create a security group that includes only the client IP addresses that need access to the AP
B. Attach the new security group to the Fargate task
C. Provide the security team with the NLB's IP addresses for the allow list.
D. Create two new /27 subnet
E. Create a new Application Load Balancer (ALB) that extends across the new subnet
F. Create a security group that includes only the client IP addresses that need access to the AP
G. Attach the security group to the AL
H. Provide the security team with the new subnet IP ranges for the allow list.
I. Create two new '27 subnet
J. Create a new Network Load Balancer (NLB) that extends across the new subnet
K. Create a new Application Load Balancer (ALB) within the new subnet
L. Create a security group that includes only the client IP addresses that need access to the AP
M. Attach the security group to the AL
N. Add the ALB's IP addresses as targets behind the NL
O. Provide the security team with the NLB's IP addresses for the allow list.
P. Create a new Application Load Balancer (ALB) in the same subnets as the Fargate task deployments.Create a security group that includes only the client IP addresses that need access to the AP
Q. Attach the security group to the AL
R. Provide the security team with the ALB's IP addresses for the allow list.

Answer: A

NEW QUESTION 201
- (Exam Topic 2)
A company's security compliance requirements state that all Amazon EC2 images must be scanned for vulnerabilities and must pass a CVE assessment A solutions architect is developing a mechanism to create security-approved AMIs that can be used by developers Any new AMIs should go through an automated assessment process and be marked as approved before developers can use them The approved images must be scanned every 30 days to ensure compliance Which combination of steps should the solutions architect take to meet these requirements while following best practices'? (Select TWO)

A. Use the AWS Systems Manager EC2 agent to run the CVE assessment on the EC2 instances launched from the AMIs that need to be scanned
B. Use AWS Lambda to write automatic approval rules Store the approved AMI list in AWS Systems Manager Parameter Store Use Amazon EventBridge to trigger an AWS Systems Manager Automation document on all EC2 instances every 30 days.
C. Use Amazon Inspector to run the CVE assessment on the EC2 instances launched from the AMIs that need to be scanned
D. Use AWS Lambda to write automatic approval rules Store the approved AMI list in AWS Systems Manager Parameter Store Use a managed AWS Config rule for continuous scanning on all EC2 instances, and use AWS Systems Manager Automation documents for remediation
E. Use AWS CloudTrail to run the CVE assessment on the EC2 instances launched from the AMIs that need to be scanned

Answer: BC

NEW QUESTION 204
......

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

100% Actual & Verified — Instant Download, Please Click
[Order The AWS-Certified-Solutions-Architect-Professional Practice Test Here](https://www.certshared.com)