Exam Questions AWS-Certified-Solutions-Architect-Professional

Amazon AWS Certified Solutions Architect Professional

https://www.2passeasy.com/dumps/AWS-Certified-Solutions-Architect-Professional/

NEW QUESTION 1
- (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 elastic network 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 2
- (Exam Topic 2)
A company built an application based on AWS Lambda deployed in an AWS CloudFormation stack. The last production release of the web application introduced an issue that resulted in an outage lasting several minutes. A solutions architect must adjust the deployment process to support a canary release.
Which solution will meet these requirements?

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

Answer: A

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

NEW QUESTION 3
- (Exam Topic 2)
A solutions architect is designing an AWS account structure for a company that consists of multiple teams. All the teams will work in the same AWS Region. The company needs a VPC that is connected to the on-premises network. The company expects less than 50 Mbps of total traffic to and from the on-premises network.
Which combination of steps will meet these requirements MOST cost-effectively? (Select TWO.)

A. Create an AWS Cloud Formation template that provisions a VPC and the required subnet
B. Deploy the template to each AWS account.
C. Create an AWS Cloud Formation template that provisions a VPC and the required subnet
D. Deploy the template to a shared services account Share the subnets by using AWS Resource Access Manager.
E. Use AWS Transit Gateway along with an AWS Site-to-Site VPN for connectivity to the on-premises networ
F. Share the transit gateway by using AWS Resource Access Manager.
G. Use AWS Site-to-Site VPN for connectivity to the on-premises network.
H. Use AWS Direct Connect for connectivity to the on-premises network.

Answer: BD

NEW QUESTION 4
- (Exam Topic 2)
A company wants to refactor its retail ordering web application that currently has a load-balanced Amazon EC2 instance fleet for web hosting, database API services, and business logic. The company needs to create a decoupled, scalable architecture with a mechanism for retaining failed orders while also minimizing operational costs.
Which solution will meet these requirements?

A. Use Amazon S3 for web hosting with Amazon API Gateway for database API service
B. Use Amazon Simple Queue Service (Amazon SQS) for order queuin
C. Use Amazon Elastic Container Service (Amazon ECS) for business logic with Amazon SQS long polling for retaining failed orders.
D. Use AWS Elastic Beanstalk for web hosting with Amazon API Gateway for database API service
E. Use Amazon MQ for order queuin

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

Answer: C

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

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

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

Answer: D

NEW QUESTION 6
- (Exam Topic 2)
A solutions architect needs to improve an application that is hosted in the AWS Cloud. The application uses an Amazon Aurora MySQL DB instance that is experiencing overloaded connections. Most of the application's operations insert records into the database. The application currently stores credentials in a text-based configuration file.
The solutions architect needs to implement a solution so that the application can handle the current connection load. The solution must keep the credentials secure and must provide the ability to rotate the credentials automatically on a regular basis.
Which solution will meet these requirements?

A. Deploy an Amazon RDS Proxy layer in front of the DB instanc
B. Store the connection credentials as a secret in AWS Secrets Manager.
C. Deploy an Amazon RDS Proxy layer in front of the DB instanc
D. Store the connection credentials in AWS Systems Manager Parameter Store.
E. Create an Aurora Replic
F. Store the connection credentials as a secret in AWS Secrets Manager.
G. Create an Aurora Replic
H. Store the connection credentials in AWS Systems Manager Parameter Store.

Answer: A

Explanation:
https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/rds-proxy.html

NEW QUESTION 7
- (Exam Topic 2)
A company has an application that runs on Amazon EC2 instances in an Amazon EC2 Auto Scaling group. The company uses AWS CodePipeline to deploy the application. The instances that run in the Auto Scaling group are constantly changing because of scaling events.
When the company deploys new application code versions, the company installs the AWS CodeDeploy agent on any new target EC2 instances and associates the instances with the CodeDeploy deployment group. The application is set to go live within the next 24 hours.
What should a solutions architect recommend to automate the application deployment process with the LEAST amount of operational overhead?

A. Configure Amazon EventBridge to invoke an AWS Lambda function when a new EC2 instance is launched into the Auto Scaling grou

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

Answer: D

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

NEW QUESTION 8
- (Exam Topic 2)
A company has an application in the AWS Cloud. The application runs on a fleet of 20 Amazon EC2 instances. The EC2 instances are persistent and store data on multiple attached Amazon Elastic Block Store (Amazon EBS) volumes.
The company must maintain backups in a separate AWS Region. The company must be able to recover the EC2 instances and their configuration within I business day, with loss of no more than I day's worth of data. The company has limited staff and needs a backup solution that optimizes operational efficiency and cost. The company already has created an AWS CloudFormation template that can deploy the required network configuration in a secondary Region.
Which solution will meet these requirements?
A. Create a second CloudFormation template that can recreate the EC2 instances in the secondary Region.Run daily multivolume snapshots by using AWS Systems Manager Automation runbook
B. Copy the snapshots to the secondary Regio
C. In the event of a failure, launch the CloudFormation templates, restore the EBS volumes from snapshots, and transfer usage to the secondary Region.
D. Use Amazon Data Lifecycle Manager (Amazon DLM) to create daily multivolume snapshots of the EBS volume
E. In the event of a failure, launch theCloudFormation template and use Amazon DLM to restore the EBS volumes and transfer usage to the secondary Region.
F. Use AWS Backup to create a scheduled daily backup plan for the EC2 instance
G. Configure the backup task to copy the backups to a vault in the secondary Regio
H. In the event of a failure, launch the CloudFormation template, restore the instance volumes and configurations from the backup vault, and transfer usage to the secondary Region.
I. Deploy EC2 instances of the same size and configuration to the secondary Regio
J. Configure AWS DataSync daily to copy data from the primary Region to the secondary Regio
K. In the event of a failure, launch the CloudFormation template and transfer usage to the secondary Region.

Answer: C

Explanation:
Using AWS Backup to create a scheduled daily backup plan for the EC2 instances will enable taking snapshots of the EC2 instances and their attached EBS volumes1. Configuring the backup task to copy the backups to a vault in the secondary Region will enable maintaining backups in a separate Region1. In the event of a failure, launching the CloudFormation template will enable deploying the network configuration in the secondary Region2. Restoring the instance volumes and configurations from the backup vault will enable recovering the EC2 instances and their data1. Transferring usage to the secondary Region will enable resuming operations2.

NEW QUESTION 9
- (Exam Topic 2)
A company runs an loT application in the AWS Cloud. The company has millions of sensors that collect data from houses in the United States. The sensors use the MOTT protocol to connect and send data to a custom MQTT broker. The MQTT broker stores the data on a single Amazon EC2 instance. The sensors connect to the broker through the domain named iot.example.com. The company uses Amazon Route 53 as its DNS service. The company stores the data in Amazon DynamoDB.
On several occasions, the amount of data has overloaded the MOTT broker and has resulted in lost sensor data. The company must improve the reliability of the solution.
Which solution will meet these requirements?

A. Create an Application Load Balancer (ALB) and an Auto Scaling group for the MOTT broke
B. Use the Auto Scaling group as the target for the AL
C. Update the DNS record in Route 53 to an alias recor
D. Point the alias record to the AL
E. Use the MQTT broker to store the data.
F. Set up AWS loT Core to receive the sensor dat
G. Create and configure a custom domain to connect to AWS loT Cor
H. Update the DNS record in Route 53 to point to the AWS loT Core Data-ATS endpoin
I. Configure an AWS loT rule to store the data.
J. Create a Network Load Balancer (NLB). Set the MQTT broker as the targe
K. Create an AWS Global Accelerator accelerato
L. Set the NLB as the endpoint for the accelerato
M. Update the DNS record in Route 53 to a multivalue answer recor
N. Set the Global Accelerator IP addresses as value
O. Use the MQTT broker to store the data.
P. Set up AWS loT Greengrass to receive the sensor dat
Q. Update the DNS record in Route 53 to point to the AWS loT Greengrass endpoin
R. Configure an AWS loT rule to invoke an AWS Lambda function to store the data.

Answer: A

Explanation:
it describes a solution that uses an Application Load Balancer (ALB) and an Auto Scaling group for the MQTT broker. The ALB distributes incoming traffic across

the instances in the Auto Scaling group and allows for automatic scaling based on incoming traffic. The use of an alias record in Route 53 allows for easy updates to the DNS record without changing the IP address. This solution improves the reliability of the MQTT broker by allowing it to automatically scale based on incoming traffic, reducing the likelihood of lost data due to broker overload.
Reference: https://aws.amazon.com/elasticloadbalancing/applicationloadbalancer/ https://aws.amazon.com/autoscaling/ https://aws.amazon.com/route53/

NEW QUESTION 10
- (Exam Topic 2)
A company has developed a hybrid solution between its data center and AWS. The company uses Amazon VPC and Amazon EC2 instances that send application togs to Amazon CloudWatch. The EC2 instances read data from multiple relational databases that are hosted on premises.
The company wants to monitor which EC2 instances are connected to the databases in near-real time. The company already has a monitoring solution that uses Splunk on premises. A solutions architect needs to determine how to send networking traffic to Splunk.
How should the solutions architect meet these requirements?

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

Answer: B

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

NEW QUESTION 10
- (Exam Topic 2)
A company is migrating a document processing workload to AWS. The company has updated many applications to natively use the Amazon S3 API to store, retrieve, and modify documents that a processing server generates at a rate of approximately 5 documents every second. After the document processing is finished, customers can download the documents directly from Amazon S3.
During the migration, the company discovered that it could not immediately update the processing server that generates many documents to support the S3 API. The server runs on Linux and requires fast local access to the files that the server generates and modifies. When the server finishes processing, the files must be available to the public for download within 30 minutes.
Which solution will meet these requirements with the LEAST amount of effort?

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

Answer: C

Explanation:
The company should configure Amazon FSx for Lustre with an import and export policy. The company should link the new file system to an S3 bucket. The company should install the Lustre client and mount the document store to an Amazon EC2 instance by using NFS. This solution will meet the requirements with the least amount of effort because Amazon FSx for Lustre is a fully managed service that provides a high-performance file system optimized for fast processing of workloads such as machine learning, high performance computing, video processing, financial modeling, and electronic design automation1. Amazon FSx for Lustre can be linked to an S3 bucket and can import data from and export data to the bucket2. The import and export policy can be configured to automatically import new or changed objects from S3 and export new or changed files to S33. This will ensure that the files are available to the public for download within 30 minutes. Amazon FSx for Lustre supports NFS version 3.0 protocol for Linux clients.
The other options are not correct because:

- Migrating the application to an AWS Lambda function would require a lot of effort and may not be feasible for the existing server that generates many documents. Lambda functions have limitations on execution time, memory, disk space, and network bandwidth.
- Setting up an Amazon S3 File Gateway would not work because S3 File Gateway does not support write-back caching, which means that files written to the file share are uploaded to S3 immediately and are not available locally until they are downloaded again. This would not provide fast local access to the files that the server generates and modifies.
- Configuring AWS DataSync to connect to an Amazon EC2 instance would not meet the requirement of making the files available to the public for download within 30 minutes. DataSync is a service that transfers data between on-premises storage systems and AWS storage services over the internet or AWS Direct Connect. DataSync tasks can be scheduled to run at specific times or intervals, but they are not triggered by file changes.

References:

- https://aws.amazon.com/fsx/lustre/

- https://docs.aws.amazon.com/fsx/latest/LustreGuide/create-fs-linked-data-repo.html
- https://docs.aws.amazon.com/fsx/latest/LustreGuide/import-export-data-repositories.html
- https://docs.aws.amazon.com/fsx/latest/LustreGuide/mounting-on-premises.html
- https://docs.aws.amazon.com/lambda/latest/dg/gettingstarted-limits.html
- https://docs.aws.amazon.com/storagegateway/latest/userguide/StorageGatewayConcepts.html
- https://docs.aws.amazon.com/datasync/latest/userguide/what-is-datasync.html

NEW QUESTION 13
- (Exam Topic 2)
A company is running an application in the AWS Cloud. The application collects and stores a large amount of unstructured data in an Amazon S3 bucket. The S3 bucket contains several terabytes of data and uses the S3 Standard storage class. The data increases in size by several gigabytes every day.
The company needs to query and analyze the data. The company does not access data that is more than 1 year old. However, the company must retain all the data indefinitely for compliance reasons.
Which solution will meet these requirements MOST cost-effectively?

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

Answer: C

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

NEW QUESTION 18
- (Exam Topic 2)
A solutions architect at a large company needs to set up network security tor outbound traffic to the internet from all AWS accounts within an organization in AWS Organizations. The organization has more than 100 AWS accounts, and the accounts route to each other by using a centralized AWS Transit Gateway. Each account has both an internet gateway and a NAT gateway tor outbound traffic to the internet The company deploys resources only into a single AWS Region.
The company needs the ability to add centrally managed rule-based filtering on all outbound traffic to the internet for all AWS accounts in the organization. The peak load of outbound traffic will not exceed 25 Gbps in each Availability Zone.
Which solution meets these requirements?

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

Answer: B

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

NEW QUESTION 20
- (Exam Topic 2)
A company has a few AWS accounts for development and wants to move its production application to AWS. The company needs to enforce Amazon Elastic Block Store (Amazon EBS) encryption at rest current production accounts and future production accounts only. The company needs a solution that includes built-in blueprints and guardrails.
Which combination of steps will meet these requirements? (Choose three.)

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

Answer: CDF

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

NEW QUESTION 25
- (Exam Topic 2)
A software-as-a-service (SaaS) provider exposes APIs through an Application Load Balancer (ALB). The ALB connects to an Amazon Elastic Kubernetes Service (Amazon EKS) cluster that is deployed in the
us-east-I Region. The exposed APIs contain usage of a few non-standard REST methods: LINK, UNLINK, LOCK, and UNLOCK.
Users outside the United States are reporting long and inconsistent response times for these APIs. A solutions architect needs to resolve this problem with a solution that minimizes operational overhead.
Which solution meets these requirements?

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

Answer: C

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

NEW QUESTION 29
- (Exam Topic 2)
A company is updating an application that customers use to make online orders. The number of attacks on the application by bad actors has increased recently. The company will host the updated application on an Amazon Elastic Container Service (Amazon ECS) cluster. The company will use Amazon DynamoDB to store application data. A public Application Load Balancer (ALB) will provide end users with access to the application. The company must prevent prevent attacks and ensure business continuity with minimal service interruptions during an ongoing attack.
Which combination of steps will meet these requirements MOST cost-effectively? (Select TWO.)

A. Create an Amazon CloudFront distribution with the ALB as the origi
B. Add a custom header and random value on the CloudFront domai
C. Configure the ALB to conditionally forward traffic if the header and value match.
D. Deploy the application in two AWS Region
E. Configure Amazon Route 53 to route to both Regions with equal weight.
F. Configure auto scaling for Amazon ECS task
G. Create a DynamoDB Accelerator (DAX) cluster.
H. Configure Amazon ElastiCache to reduce overhead on DynamoDB.
I. Deploy an AWS WAF web ACL that includes an appropriate rule grou
J. Associate the web ACL with the Amazon CloudFront distribution.

Answer: AE

Explanation:
The company should create an Amazon CloudFront distribution with the ALB as the origin. The company should add a custom header and random value on the CloudFront domain. The company should configure the ALB to conditionally forward traffic if the header and value match. The company should also deploy an AWS WAF web ACL that includes an appropriate rule group. The company should associate the web ACL with the Amazon CloudFront distribution. This solution will meet the requirements most cost-effectively because Amazon CloudFront is a fast content delivery network (CDN) service that securely delivers data, videos, applications, and APIs to customers globally with low latency, high transfer speeds, all within a
developer-friendly environment1. By creating an Amazon CloudFront distribution with the ALB as the origin, the company can improve the performance and availability of its application by caching static content at edge locations closer to end users. By adding a custom header and random value on the CloudFront domain, the company can prevent direct access to the ALB and ensure that only requests from CloudFront are forwarded to the ECS tasks. By configuring the ALB to conditionally forward traffic if the header and value match, the company can implement origin access identity (OAI) for its ALB origin. OAI is a feature that enables you to restrict access to your content by requiring users to access your content through CloudFront URLs2. By deploying an AWS WAF web ACL that includes an appropriate rule group, the company can prevent attacks and ensure business continuity with minimal service interruptions during an ongoing attack. AWS WAF is a web application firewall that lets you monitor and control web requests that are forwarded to your web applications. You can use AWS WAF to define customizable web security rules that control which traffic can access your web applications and which traffic should be blocked3. By associating the web ACL with the Amazon CloudFront distribution, the company can apply the web security rules to all requests that are forwarded by CloudFront.
The other options are not correct because:

- Deploying the application in two AWS Regions and configuring Amazon Route 53 to route to both Regions with equal weight would not prevent attacks or ensure business continuity. Amazon Route 53 is a highly available and scalable cloud Domain Name System (DNS) web service that routes end users to Internet applications by translating names like www.example.com into numeric IP addresses4. However, routing traffic to multiple Regions would not protect against attacks or provide failover in case of an outage. It would also increase operational complexity and costs compared to using CloudFront and
AWS WAF.
- Configuring auto scaling for Amazon ECS tasks and creating a DynamoDB Accelerator (DAX) cluster would not prevent attacks or ensure business continuity. Auto scaling is a feature that enables you to automatically adjust your ECS tasks based on demand or a schedule. DynamoDB Accelerator (DAX) is a fully managed, highly available, in-memory cache for DynamoDB that delivers up to a 10x performance improvement. However, these features would not protect against attacks or provide failover in case of an outage. They would also increase operational complexity and costs compared to using CloudFront and AWS WAF.
- Configuring Amazon ElastiCache to reduce overhead on DynamoDB would not prevent attacks or ensure business continuity. Amazon ElastiCache is a fully managed in-memory data store service that makes it easy to deploy, operate, and scale popular open-source compatible in-memory data stores. However, this service would not protect against attacks or provide failover in case of an outage. It would also increase operational complexity and costs compared to using

CloudFront and AWS WAF.
References:

- https://aws.amazon.com/cloudfront/
- https://aws.amazon.com/waf/
- https://aws.amazon.com/route53/
- https://aws.amazon.com/dynamodb/dax/
- https://aws.amazon.com/elasticache/

NEW QUESTION 33
- (Exam Topic 2)
A solutions architect needs to review the design of an Amazon EMR cluster that is using the EMR File System (EMRFS). The cluster performs tasks that are critical to business needs. The cluster is running Amazon EC2 On-Demand Instances at all times tor all task, primary, and core nodes. The EMR tasks run each morning, starting at 1 ;00 AM. and take 6 hours to finish running. The amount of time to complete the processing is not a priority because the data is not referenced until late in the day.
The solutions architect must review the architecture and suggest a solution to minimize the compute costs. Which solution should the solutions architect recommend to meet these requirements?

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

Answer: A

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

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

NEW QUESTION 38
- (Exam Topic 2)
A company wants to optimize AWS data-transfer costs and compute costs across developer accounts within the company's organization in AWS Organizations Developers can configure VPCs and launch Amazon EC2 instances in a single AWS Region The EC2 instances retrieve approximately 1 TB of data each day from Amazon S3
The developer activity leads to excessive monthly data-transfer charges and NAT gateway processing charges between EC2 instances and S3 buckets, along with high compute costs The company wants to proactively enforce approved architectural patterns for any EC2 instance and VPC infrastructure that developers deploy within the AWS accounts The company does not want this enforcement to negatively affect the speed at which the developers can perform their tasks
Which solution will meet these requirements MOST cost-effectively?

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

Answer: C

Explanation:
This solution allows developers to quickly launch resources using pre-approved configurations and instance types, while also ensuring that the resources launched comply with the company's architectural patterns. This can help reduce data transfer and compute costs associated with the resources. Using AWS Service Catalog also allows the company to control access to the approved configurations and resources through the use of IAM roles, while also allowing developers to quickly provision resources without negatively affecting their ability to perform their tasks.
Reference:
AWS Service Catalog: https://aws.amazon.com/service-catalog/ AWS Service Catalog Constraints: https://docs.aws.amazon.com/servicecatalog/latest/adminguide/constraints.html
AWS Service Catalog Launch Constraints: https://docs.aws.amazon.com/servicecatalog/latest/adminguide/launch-constraints.html

NEW QUESTION 41
- (Exam Topic 2)
A solutions architect is reviewing a company's process for taking snapshots of Amazon RDS DB instances. The company takes automatic snapshots every day and retains the snapshots for 7 days.
The solutions architect needs to recommend a solution that takes snapshots every 6 hours and retains the snapshots for 30 days. The company uses AWS Organizations to manage all of its AWS accounts. The company needs a consolidated view of the health of the RDS snapshots.
Which solution will meet these requirements with the LEAST operational overhead?

A. Turn on the cross-account management feature in AWS Backu
B. Create a backup plan that specifies the frequency and retention requirement
C. Add a tag to the DB instance
D. Apply the backup plan by using tag
E. Use AWS Backup to monitor the status of the backups.
F. Turn on the cross-account management feature in Amazon RD
G. Create a snapshot global policy that specifies the frequency and retention requirement
H. Use the RDS console in the management account to monitor the status of the backups.
I. Turn on the cross-account management feature in AWS CloudFormatio
J. From the management account, deploy a CloudFormation stack set that contains a backup plan from AWS Backup thatspecifies the frequency and retention requirement
K. Create an AWS Lambda function in the management account to monitor the status of the backup
L. Create an Amazon EventBridge rule in each account to run the Lambda function on a schedule.
M. Configure AWS Backup in each accoun
N. Create an Amazon Data Lifecycle Manager lifecycle policy that specifies the frequency and retention requirement
O. Specify the DB instances as the target resourc
P. Use the Amazon Data Lifecycle Manager console in each member account to monitor the status of the backups.

Answer: A

Explanation:
Turning on the cross-account management feature in AWS Backup will enable managing and monitoring backups across multiple AWS accounts that belong to the same organization in AWS Organizations1. Creating a backup plan that specifies the frequency and retention requirements will enable taking snapshots every 6 hours and retaining them for 30 days2. Adding a tag to the DB instances will enable applying the backup plan by using tags2. Using AWS Backup to monitor the status of the backups will enable having a consolidated view of the health of the RDS snapshots1.

NEW QUESTION 44
- (Exam Topic 2)
A company has multiple business units that each have separate accounts on AWS. Each business unit manages its own network with several VPCs that have CIDR ranges that overlap. The company’s marketing team has created a new internal application and wants to make the application accessible to all the other business units. The solution must use private IP addresses only.
Which solution will meet these requirements with the LEAST operational overhead?

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

Answer: C

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

NEW QUESTION 46
- (Exam Topic 2)
A company uses an AWS CodeCommit repository The company must store a backup copy of the data that is in the repository in a second AWS Region
Which solution will meet these requirements?

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

Answer: B

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

allows you to copy the backups to a second Region for disaster recovery.
By using AWS Backup, the company can set up an automatic and regular backup schedule for the CodeCommit repository, ensuring that the data is regularly backed up and stored in a second Region. This can provide a way to recover quickly from any disaster event that might occur.
Reference:
AWS Backup documentation: https://aws.amazon.com/backup/ AWS Backup for AWS CodeCommit documentation: https://aws.amazon.com/about-aws/whats-new/2020/07/aws-backup-now-supports-aws-codecommit-repositorie

NEW QUESTION 50
- (Exam Topic 2)
A company has several AWS accounts. A development team is building an automation framework for cloud governance and remediation processes. The automation framework uses AWS Lambda functions in a centralized account. A solutions architect must implement a least privilege permissions policy that allows the Lambda functions to run in each of the company's AWS accounts.
Which combination of steps will meet these requirements? (Choose two.)

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

Answer: AB

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

NEW QUESTION 54
- (Exam Topic 2)
A company has an on-premises Microsoft SOL Server database that writes a nightly 200 GB export to a local drive. The company wants to move the backups to more robust cloud storage on Amazon S3. The company has set up a 10 Gbps AWS Direct Connect connection between the on-premises data center and AWS.
Which solution meets these requirements MOST cost-effectively?

A. Create a new S3 bucke
B. Deploy an AWS Storage Gateway file gateway within the VPC that Is connected to the Direct Connect connectio
C. Create a new SMB file shar
D. Write nightly database exports to the new SMB file share.
E. Create an Amazon FSx for Windows File Server Single-AZ file system within the VPC that is connected to the Direct Connect connectio
F. Create a new SMB file shar
G. Write nightly database exports to an SMB file share on the Amazon FSx file syste
H. Enable nightly backups.
I. Create an Amazon FSx for Windows File Server Multi-AZ file system within the VPC that is connected to the Direct Connect connectio
J. Create a new SMB file shar
K. Write nightly database exports to an SMB file share on the Amazon FSx file syste
L. Enable nightly backups.
M. Create a new S3 bucke
N. Deploy an AWS Storage Gateway volume gateway within the VPC that Is connected to the Direct Connect connectio
O. Create a new SMB file shar
P. Write nightly database exports to the new SMB file share on the volume gateway, and automate copies of this data to an S3 bucket.

Answer: A

Explanation:
https://docs.aws.amazon.com/filegateway/latest/files3/CreatingAnSMBFileShare.html

NEW QUESTION 58
- (Exam Topic 2)
A company's public API runs as tasks on Amazon Elastic Container Service (Amazon ECS). The tasks run on AWS Fargate behind an Application Load Balancer (ALB) and are configured with Service Auto Scaling for the tasks based on CPU utilization. This service has been running well for several months.
Recently, API performance slowed down and made the application unusable. The company discovered that a significant number of SQL injection attacks had occurred against the API and that the API service had scaled to its maximum amount.
A solutions architect needs to implement a solution that prevents SQL injection attacks from reaching the ECS API service. The solution must allow legitimate traffic through and must maximize operational efficiency. Which solution meets these requirements?

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

Answer: C

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

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

References:

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

NEW QUESTION 61
- (Exam Topic 2)
A company operates a proxy server on a fleet of Amazon EC2 instances. Partners in different countries use the proxy server to test the company's functionality. The EC2 instances are running in a VPC. and the instances have access to the internet.
The company's security policy requires that partners can access resources only from domains that the company owns.
Which solution will meet these requirements?

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

Answer: A

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

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

References:

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

NEW QUESTION 64
- (Exam Topic 2)
A company runs an application on AWS. The company curates data from several different sources. The company uses proprietary algorithms to perform data transformations and aggregations. After the company performs E TL processes, the company stores the results in Amazon Redshift tables. The company sells this data to other companies. The company downloads the data as files from the Amazon Redshift tables and transmits the files to several data customers by using FTP. The number of data customers has grown significantly. Management of the data customers has become difficult.
The company will use AWS Data Exchange to create a data product that the company can use to share data with customers. The company wants to confirm the identities of the customers before the company shares data.
The customers also need access to the most recent data when the company publishes the data. Which solution will meet these requirements with the LEAST operational overhead?

A. Use AWS Data Exchange for APIs to share data with customer
B. Configure subscription verification In the AWS account of the company that produces the data, create an Amazon API Gateway Data API service integration with Amazon Redshif
C. Require the data customers to subscribe to the data product In the AWS account of the company that produces the data, create an AWS Data Exchange datashare by connecting AWS Data Exchange to the Redshift
D. cluste
E. Configure subscription verificatio
F. Require the data customers to subscribe to the data product.
G. Download the data from the Amazon Redshift tables to an Amazon S3 bucket periodicall
H. Use AWS Data Exchange for S3 to share data with customers.
I. Configure subscription verificatio
J. Require the data customers to subscribe to the data product Publish the Amazon Redshift data to an Open Data on AWS Data Exchang
K. Require the customers to subscribe to the data product in AWS Data Exchang
L. In the AWS account of the company that produces the data, attach IAM resource-based policies to the Amazon Redshift tables to allow access only to verified AWS accounts.

Answer: C

Explanation:
The company should download the data from the Amazon Redshift tables to an Amazon S3 bucket periodically and use AWS Data Exchange for S3 to share data with customers. The company should configure subscription verification and require the data customers to subscribe to the data product. This solution will meet the requirements with the least operational overhead because AWS Data Exchange for S3 is a feature that enables data subscribers to access third-party data files directly from data providers' Amazon S3 buckets. Subscribers can easily use these files for their data analysis with AWS services without needing to create or manage data copies. Data providers can easily set up AWS Data Exchange for S3 on top of their existing S3 buckets to share direct access to an entire S3 bucket or specific prefixes and S3 objects. AWS Data Exchange automatically manages subscriptions, entitlements, billing, and payment1.
The other options are not correct because:
- Using AWS Data Exchange for APIs to share data with customers would not work because AWS Data Exchange for APIs is a feature that enables data subscribers to access third-party APIs directly from data providers' AWS accounts. Subscribers can easily use these APIs for their data analysis with AWS services without needing to manage API keys or tokens. Data providers can easily set up AWS Data Exchange for APIs on top of their existing API Gateway resources to share direct access to an entire API or specific routes and stages2. However, this feature is not suitable for sharing data from Amazon Redshift tables, which are not exposed as APIs.
- Creating an Amazon API Gateway Data API service integration with Amazon Redshift would not work because the Data API is a feature that enables you to query your Amazon Redshift cluster using HTTP requests, without needing a persistent connection or a SQL client3. It is useful for building applications that interact with Amazon Redshift, but not for sharing data files with customers.
- Creating an AWS Data Exchange datashare by connecting AWS Data Exchange to the Redshift cluster would not work because AWS Data Exchange does not support datashares for Amazon Redshift clusters. A datashare is a feature that enables you to share live and secure access to your Amazon Redshift data across your accounts or with third parties without copying or moving the underlying data4. It is useful for sharing query results and views with other users, but not for sharing data files with customers.
- Publishing the Amazon Redshift data to an Open Data on AWS Data Exchange would not work because Open Data on AWS Data Exchange is a feature that enables you to find and use free and public datasets from AWS customers and partners. It is useful for accessing open and free data, but not for confirming the identities of the customers or charging them for the data.
References:
- https://aws.amazon.com/data-exchange/why-aws-data-exchange/s3/
- https://aws.amazon.com/data-exchange/why-aws-data-exchange/api/
- https://docs.aws.amazon.com/redshift/latest/mgmt/data-api.html
- https://docs.aws.amazon.com/redshift/latest/dg/datashare-overview.html
- https://aws.amazon.com/data-exchange/open-data/

NEW QUESTION 68
- (Exam Topic 2)
A company is migrating its development and production workloads to a new organization in AWS Organizations. The company has created a separate member account for development and a separate member account for production. Consolidated billing is linked to the management account. In the management account, a solutions architect needs to create an 1AM user that can stop or terminate resources in both member accounts.
Which solution will meet this requirement?

A. Create an IAM user and a cross-account role in the management accoun
B. Configure the cross-account role with least privilege access to the member accounts.
C. Create an IAM user in each member accoun
D. In the management account, create a cross-account role that has least privilege acces
E. Grant the IAM users access to the cross-account role by using a trust policy.
F. Create an IAM user in the management accoun
G. In the member accounts, create an IAM group that has least privilege acces
H. Add the IAM user from the management account to each IAM group in the member accounts.
I. Create an IAM user in the management accoun
J. In the member accounts, create cross-account roles that have least privilege acces
K. Grant the IAM user access to the roles by using a trust policy.

Answer: D

Explanation:
Cross account role should be created in destination(member) account. The role has trust entity to master account.

NEW QUESTION 69
- (Exam Topic 2)
A company is designing an AWS Organizations structure. The company wants to standardize a process to apply tags across the entire organization. The company will require tags with specific values when a user creates a new resource. Each of the company's OUs will have unique tag values.
Which solution will meet these requirements?

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

Answer: A

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

NEW QUESTION 72
- (Exam Topic 2)
A company is deploying a new web-based application and needs a storage solution for the Linux application servers. The company wants to create a single location for updates to application data for all instances. The active dataset will be up to 100 GB in size. A solutions architect has determined that peak operations will occur for 3 hours daily and will require a total of 225 MiBps of read throughput.
The solutions architect must design a Multi-AZ solution that makes a copy of the data available in another AWS Region for disaster recovery (DR). The DR copy has an RPO of less than 1 hour.
Which solution will meet these requirements?

A. Deploy a new Amazon Elastic File System (Amazon EFS) Multi-AZ file syste
B. Configure the file system for 75 MiBps of provisioned throughpu
C. Implementreplication to a file system in the DR Region.
D. Deploy a new Amazon FSx for Lustre file syste
E. Configure Bursting Throughput mode for the file syste
F. Use AWS Backup to back up the file system to the DR Region.
G. Deploy a General Purpose SSD (gp3) Amazon Elastic Block Store (Amazon EBS) volume with 225 MiBps of throughpu
H. Enable Multi-Attach for the EBSvolum
I. Use AWS Elastic Disaster Recovery to replicate the EBS volume to the DR Region.
J. Deploy an Amazon FSx for OpenZFS file system in both the production Region and the DR Region.Create an AWS DataSync scheduled task to replicate thedata from the production file system to the DR file system every 10 minutes.

Answer: A

Explanation:
The company should deploy a new Amazon Elastic File System (Amazon EFS) Multi-AZ file system. The company should configure the file system for 75 MiBps of provisioned throughput. The company should implement replication to a file system in the DR Region. This solution will meet the requirements because Amazon EFS is a serverless, fully elastic file storage service that lets you share file data without provisioning or managing storage capacity and performance. Amazon EFS is built to scale on demand to petabytes without disrupting applications, growing and shrinking automatically as you add and remove files1. By deploying a new Amazon EFS Multi-AZ file system, the company can create a single location for updates to application data for all instances. A Multi-AZ file system replicates data across multiple Availability Zones (AZs) within a Region, providing high availability and durability2. By configuring the file system for 75 MiBps of provisioned throughput, the company can ensure that it meets the peak operations requirement of 225 MiBps of read throughput. Provisioned throughput is a feature that enables you to specify a level of throughput that the file system can drive independent of the file system's size or burst credit balance3. By implementing replication to a file system in the DR Region, the company can make a copy of the data available in another AWS Region for disaster recovery. Replication is a feature that enables you to replicate data from one EFS file system to another EFS file system across AWS Regions. The replication process has an RPO of less than 1 hour.
The other options are not correct because:

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

replication of data. AWS DataSync is a service that enables you to transfer data between on-premises storage and AWS services, or between AWS services. However, it does not support continuous data replication or sub-minute RPOs.
References:

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

NEW QUESTION 75
- (Exam Topic 2)
A company has millions of objects in an Amazon S3 bucket. The objects are in the S3 Standard storage class. All the S3 objects are accessed frequently. The number of users and applications that access the objects is increasing rapidly. The objects are encrypted with server-side encryption with AWS KMS Keys (SSE-KMS).
A solutions architect reviews the company's monthly AWS invoice and notices that AWS KMS costs are increasing because of the high number of requests from Amazon S3. The solutions architect needs to optimize costs with minimal changes to the application.
Which solution will meet these requirements with the LEAST operational overhead?

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

Answer: B

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

NEW QUESTION 78
- (Exam Topic 2)
A company has built a high performance computing (HPC) cluster in AWS tor a tightly coupled workload that generates a large number of shared files stored in Amazon EFS. The cluster was performing well when the number of Amazon EC2 instances in the cluster was 100. However, when the company increased the cluster size to 1,000 EC2 instances, overall performance was well below expectations.
Which collection of design choices should a solutions architect make to achieve the maximum performance from the HPC cluster? (Select THREE.)

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

Answer: ACF

Explanation:
* A. High performance computing (HPC) workload cluster should be in a single AZ.
* C. Elastic Fabric Adapter (EFA) is a network device that you can attach to your Amazon EC2 instances to accelerate High Performance Computing (HPC)
* F. Amazon FSx for Lustre - Use it for workloads where speed matters, such as machine learning, high performance computing (HPC), video processing, and financial modeling.
Cluster – packs instances close together inside an Availability Zone. This strategy enables workloads to achieve the low-latency network performance necessary for tightly-coupled node-to-node communication that is typical of HPC applications.
https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/placement-groups.html

NEW QUESTION 82
- (Exam Topic 2)
A company runs its sales reporting application in an AWS Region in the United States. The application uses an Amazon API Gateway Regional API and AWS Lambda functions to generate on-demand reports from data in an Amazon RDS for MySQL database. The frontend of the application is hosted on Amazon S3 and is accessed by users through an Amazon CloudFront distribution. The company is using Amazon Route 53 as the DNS service for the domain. Route 53 is configured with a simple routing policy to route traffic to the API Gateway API.
In the next 6 months, the company plans to expand operations to Europe. More than 90% of the database traffic is read-only traffic. The company has already deployed an API Gateway API and Lambda functions in the new Region.
A solutions architect must design a solution that minimizes latency for users who download reports. Which solution will meet these requirements?

A. Use an AWS Database Migration Service (AWS DMS) task with full load to replicate the primary database in the original Region to the database in the new

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

Answer: C

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

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

References:

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

NEW QUESTION 87
- (Exam Topic 2)
A company processes environment data. The has a set up sensors to provide a continuous stream of data from different areas in a city. The data is available in JSON format.
The company wants to use an AWS solution to send the data to a database that does not require fixed schemas for storage. The data must be send in real time.
Which solution will meet these requirements?

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

Answer: B

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

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

NEW QUESTION 92
- (Exam Topic 2)
A company's solutions architect is analyzing costs of a multi-application environment. The environment is deployed across multiple Availability Zones in a single AWS Region. After a recent acquisition, the company manages two organizations in AWS Organizations. The company has created multiple service provider applications as AWS PrivateLink-powered VPC endpoint services in one organization. The company has created multiple service consumer applications in the other organization.
Data transfer charges are much higher than the company expected, and the solutions architect needs to reduce the costs. The solutions architect must recommend guidelines for developers to follow when they deploy services. These guidelines must minimize data transfer charges for the whole environment.
Which guidelines meet these requirements? (Select TWO.)

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

E. Create a Savings Plan that provides adequate coverage for the organization's planned inter-Availability Zone data transfer usage.

Answer: CD

Explanation:
Cross-zone load balancing enables traffic to be distributed evenly across all registered instances in all enabled Availability Zones. However, this also increases data transfer charges between Availability Zones. By turning off cross-zone load balancing, the service provider applications can reduce inter-Availability Zone data transfer costs. Similarly, by using the Availability Zone-specific endpoint service, the service consumer applications can ensure that they connect to the nearest service provider application in the same Availability Zone, avoiding cross-Availability Zone data transfer charges. References:
https://docs.aws.amazon.com/vpc/latest/userguide/vpce-interface.html#vpce-interface-dns

NEW QUESTION 95
- (Exam Topic 1)
A company wants to use a third-party software-as-a-service (SaaS) application. The third-party SaaS application is consumed through several API calls. The third-party SaaS application also runs on AWS inside a VPC.
The company will consume the third-party SaaS application from inside a VPC. The company has internal security policies that mandate the use of private connectivity that does not traverse the internet. No resources that run in the company VPC are allowed to be accessed from outside the company's VPC. All permissions must conform to the principles of least privilege.
Which solution meets these requirements?

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

Answer: A

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

NEW QUESTION 99
- (Exam Topic 1)
A solutions architect is investigating an issue in which a company cannot establish new sessions in Amazon Workspaces. An initial analysis indicates that the issue involves user profiles. The Amazon Workspaces environment is configured to use Amazon FSx for Windows File Server as the profile share storage. The FSx for Windows File Server file system is configured with 10 TB of storage.
The solutions architect discovers that the file system has reached its maximum capacity. The solutions architect must ensure that users can regain access. The solution also must prevent the problem from occurring again.
Which solution will meet these requirements?

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

Answer: B

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

NEW QUESTION 100
- (Exam Topic 1)
A company is hosting an image-processing service on AWS in a VPC. The VPC extends across two Availability Zones. Each Availability Zone contains one public subnet and one private subnet.
The service runs on Amazon EC2 instances in the private subnets. An Application Load Balancer in the public subnets is in front of the service. The service needs to communicate with the internet and does so through two NAT gateways. The service uses Amazon S3 for image storage. The EC2 instances retrieve approximately 1 ¢' of data from an S3 bucket each day.
The company has promoted the service as highly secure. A solutions architect must reduce cloud expenditures as much as possible without compromising the service's security posture or increasing the time spent on ongoing operations.
Which solution will meet these requirements?

A. Replace the NAT gateways with NAT instance
B. In the VPC route table, create a route from the private subnets to the NAT instances.
C. Move the EC2 instances to the public subnet
D. Remove the NAT gateways.

E. Set up an S3 gateway VPC endpoint in the VP
F. Attach an endpoint policy to the endpoint to allow the required actions on the S3 bucket.
G. Attach an Amazon Elastic File System (Amazon EFS) volume to the EC2 instance
H. Host the image on the EFS volume.

Answer: C

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

NEW QUESTION 101
- (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 businessunit.
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 104
- (Exam Topic 1)
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: C

Explanation:
https://docs.aws.amazon.com/fsx/latest/WindowsGuide/what-is.html https://docs.aws.amazon.com/directoryservice/latest/admin-guide/ms_ad_join_instance.html

NEW QUESTION 109
- (Exam Topic 1)
A solutions architect needs to advise a company on how to migrate its on-premises data processing application to the AWS Cloud. Currently, users upload input files through a web portal. The web server then stores the uploaded files on NAS and messages the processing server over a message queue. Each media file can take up to 1 hour to process. The company has determined that the number of media files awaiting processing is significantly higher during business hours, with the number of files rapidly declining after business hours.
What is the MOST cost-effective migration recommendation?

A. Create a queue using Amazon SQ
B. Configure the existing web server to publish to the new queue.When there are messages in the queue, invoke an AWS Lambda function to pull requests from

the queue and process the file
C. Store the processed files in an Amazon S3 bucket.
D. Create a queue using Amazon
E. Configure the existing web server to publish to the new queu
F. When there are messages in the queue, create a new Amazon EC2 instance to pull requests from the queue and process the file
G. Store the processed files in Amazon EF
H. Shut down the EC2 instance after the task is complete.
I. Create a queue using Amazon M
J. Configure the existing web server to publish to the new queue.When there are messages in the queue, invoke an AWS Lambda function to pull requests from the queue and process the file
K. Store the processed files in Amazon EFS.
L. Create a queue using Amazon SO
M. Configure the existing web server to publish to the new queu
N. Use Amazon EC2 instances in an EC2 Auto Scaling group to pull requests from the queue and process the file
O. Scale the EC2 instances based on the SOS queue lengt
P. Store the processed files in an Amazon S3 bucket.

Answer: D

Explanation:
https://aws.amazon.com/blogs/compute/operating-lambda-performance-optimization-part-1/

NEW QUESTION 113
- (Exam Topic 1)
A company is running an application in the AWS Cloud. The company's security team must approve the creation of all new IAM users. When a new IAM user is created, all access for the user must be removed automatically. The security team must then receive a notification to approve the user. The company has a multi-Region AWS CloudTrail trail In the AWS account.
Which combination of steps will meet these requirements? (Select THREE.)

A. Create an Amazon EventBridge (Amazon CloudWatch Events) rul
B. Define a pattern with the detail-type value set to AWS API Call via CloudTrail and an eventName of CreateUser.
C. Configure CloudTrail to send a notification for the CreateUser event to an Amazon Simple Notification Service (Amazon SNS) topic.
D. Invoke a container that runs in Amazon Elastic Container Service (Amazon ECS) with AWS Fargate technology to remove access
E. Invoke an AWS Step Functions state machine to remove access.
F. Use Amazon Simple Notification Service (Amazon SNS) to notify the security team.
G. Use Amazon Pinpoint to notify the security team.

Answer: ADE

Explanation:
https://docs.aws.amazon.com/prescriptive-guidance/latest/patterns/send-a-notification-when-an-iam-user-is-crea

NEW QUESTION 114
- (Exam Topic 1)
A company uses AWS Organizations for a multi-account setup in the AWS Cloud. The company uses AWS Control Tower for governance and uses AWS Transit Gateway for VPC connectivity across accounts.
In an AWS application account, the company's application team has deployed a web application that uses AWS Lambda and Amazon RDS. The company's database administrators have a separate DBA account and use the account to centrally manage all the databases across the organization. The database administrators use an Amazon EC2 instance that is deployed in the DBA account to access an RDS database that is deployed in the application account.
The application team has stored the database credentials as secrets in AWS Secrets Manager in the application account. The application team is manually sharing the secrets with the database administrators. The secrets are encrypted by the default AWS managed key for Secrets Manager in the application account. A solutions architect needs to implement a solution that gives the database administrators access to the database and eliminates the need to manually share the secrets.
Which solution will meet these requirements?

A. Use AWS Resource Access Manager (AWS RAM) to share the secrets from the application account with the DBA accoun
B. In the DBA account, create an IAM role that is named DBA-Admi
C. Grant the role the required permissions to access the shared secret
D. Attach the DBA-Admin role to the EC2 instance for access to the cross-account secrets.
E. In the application account, create an IAM role that is named DBA-Secre
F. Grant the role the required permissions to access the secret
G. In the DBA account, create an IAM role that is named DBA-Admi
H. Grant the DBA-Admin role the required permissions to assume the DBA-Secret role in the application accoun
I. Attach the DBA-Admin role to the EC2 instance for access to the cross-account secrets.
J. In the DBA account, create an IAM role that is named DBA-Admi
K. Grant the role the required permissions to access the secrets and the default AWS managed key in the application accoun
L. In the application account, attach resource-based policies to the key to allow access from the DBA accoun
M. Attach the DBA-Admin role to the EC2 instance for access to the cross-account secrets.
N. In the DBA account, create an IAM role that is named DBA-Admi
O. Grant the role the required permissions to access the secrets in the application accoun
P. Attach an SCP to the application account to allow access to the secrets from the DBA accoun
Q. Attach the DBA-Admin role to the EC2 instance for access to the cross-account secrets.

Answer: B

Explanation:
Option B is correct because creating an IAM role in the application account that has permissions to access the secrets and creating an IAM role in the DBA account that has permissions to assume the role in the application account eliminates the need to manually share the secrets. This approach uses cross-account IAM roles to grant access to the secrets in the application account. The database administrators can assume the role in the application account from their EC2 instance in the DBA

account and retrieve the secrets without having to store them locally or share them manually2
References: 1: https://docs.aws.amazon.com/ram/latest/userguide/what-is.html 2:
https://docs.aws.amazon.com/IAM/latest/UserGuide/tutorial_cross-account-with-roles.html 3:
https://docs.aws.amazon.com/kms/latest/developerguide/concepts.html : https://docs.aws.amazon.com/secretsmanager/latest/userguide/tutorials_basic.html :
https://docs.aws.amazon.com/IAM/latest/UserGuide/introduction.html

NEW QUESTION 119
- (Exam Topic 1)
A company has migrated an application from on premises to AWS. The application frontend is a static website that runs on two Amazon EC2 instances behind an Application Load Balancer (ALB). The application backend is a Python application that runs on three EC2 instances behind another ALB. The EC2 instances are large, general purpose On-Demand Instances that were sized to meet the on-premises specifications for peak usage of the application.
The application averages hundreds of thousands of requests each month. However, the application is used mainly during lunchtime and receives minimal traffic during the rest of the day.
A solutions architect needs to optimize the infrastructure cost of the application without negatively affecting the application availability.
Which combination of steps will meet these requirements? (Choose two.)

A. Change all the EC2 instances to compute optimized instances that have the same number of cores as the existing EC2 instances.
B. Move the application frontend to a static website that is hosted on Amazon S3.
C. Deploy the application frontend by using AWS Elastic Beanstal
D. Use the same instance type for the nodes.
E. Change all the backend EC2 instances to Spot Instances.
F. Deploy the backend Python application to general purpose burstable EC2 instances that have the same number of cores as the existing EC2 instances.

Answer: BD

Explanation:
Moving the application frontend to a static website that is hosted on Amazon S3 will save cost as S3 is cheaper than running EC2 instances.
Using Spot instances for the backend EC2 instances will also save cost, as they are significantly cheaper than On-Demand instances. This will be suitable for the application, as it has minimal traffic during the rest of the day, and the availability of spot instances will not negatively affect the application's availability.
Reference:
Amazon S3 pricing: https://aws.amazon.com/s3/pricing/
Amazon EC2 Spot Instances documentation: https://aws.amazon.com/ec2/spot/ AWS Elastic Beanstalk documentation: https://aws.amazon.com/elasticbeanstalk/
Amazon Elastic Compute Cloud (EC2) pricing: https://aws.amazon.com/ec2/pricing/

NEW QUESTION 121
- (Exam Topic 1)
A software company hosts an application on AWS with resources in multiple AWS accounts and Regions. The application runs on a group of Amazon EC2 instances in an application VPC located in the us-east-1 Region with an IPv4 CIDR block of 10.10.0.0/16. In a different AWS account, a shared services VPC is located in the us-east-2 Region with an IPv4 CIDR block of 10.10.10.0/24. When a cloud engineer uses AWS CloudFormation to attempt to peer the application VPC with the shared services VPC, an error message indicates a peering failure. Which factors could cause this error? (Choose two.)

A. The IPv4 CIDR ranges of the two VPCs overlap
B. The VPCs are not in the same Region
C. One or both accounts do not have access to an Internet gateway
D. One of the VPCs was not shared through AWS Resource Access Manager
E. The IAM role in the peer accepter account does not have the correct permissions

Answer: AE

Explanation:
https://aws.amazon.com/about-aws/whats-new/2017/11/announcing-support-for-inter-region-vpc-peering/

NEW QUESTION 122
- (Exam Topic 1)
A company has a serverless application comprised of Amazon CloudFront, Amazon API Gateway, and AWS Lambda functions. The current deployment process of the application code is to create a new version number of the Lambda function and run an AWS CLI script to update. If the new function version has errors, another CLI script reverts by deploying the previous working version of the function. The company would like to decrease the time to deploy new versions of the application logic provided by the Lambda functions, and also reduce the time to detect and revert when errors are identified.
How can this be accomplished?

A. Create and deploy nested AWS CloudFormation stacks with the parent stack consisting of the AWS CloudFront distribution and API Gateway, and the child stack containing the Lambda functio
B. For changes to Lambda, create an AWS CloudFormation change set and deploy; if errors are triggered, revert the AWS CloudFormation change set to the previous version.
C. Use AWS SAM and built-in AWS CodeDeploy to deploy the new Lambda version, gradually shift traffic to the new version, and use pre-traffic and post-traffic test functions to verify cod
D. Rollback if Amazon CloudWatch alarms are triggered.
E. Refactor the AWS CLI scripts into a single script that deploys the new Lambda versio
F. When deployment is completed, the script tests execut
G. If errors are detected, revert to the previous Lambda version.
H. Create and deploy an AWS CloudFormation stack that consists of a new API Gateway endpoint that references the new Lambda versio
I. Change the CloudFront origin to the new API Gateway endpoint, monitor errors and if detected, change the AWS CloudFront origin to the previous API Gateway endpoint.

Answer: B

Explanation:
https://aws.amazon.com/about-aws/whats-new/2017/11/aws-lambda-supports-traffic-shifting-and-phased-deploy

NEW QUESTION 127
- (Exam Topic 1)
A company is running applications on AWS in a multi-account environment. The company's sales team and marketing team use separate AWS accounts in AWS Organizations.
The sales team stores petabytes of data in an Amazon S3 bucket. The marketing team uses Amazon QuickSight for data visualizations. The marketing team needs access to data that the sates team stores in the S3 bucket. The company has encrypted the S3 bucket with an AWS Key Management Service (AWS KMS) key. The marketing team has already created the IAM service role for QuickSight to provide QuickSight access in the marketing AWS account. The company needs a solution that will provide secure access to the data in the S3 bucket across AWS accounts.
Which solution will meet these requirements with the LEAST operational overhead?

A. Create a new S3 bucket in the marketing accoun
B. Create an S3 replication rule in the sales account to copy the objects to the new S3 bucket in the marketing accoun
C. Update the QuickSight permissions in the marketing account to grant access to the new S3 bucket.
D. Create an SCP to grant access to the S3 bucket to the marketing accoun
E. Use AWS Resource Access Manager (AWS RAM) to share the KMS key from the sates account with the marketing accoun
F. Update the QuickSight permissions in the marketing account to grant access to the S3 bucket.
G. Update the S3 bucket policy in the marketing account to grant access to the QuickSight rol
H. Create a KMS grant for the encryption key that is used in the S3 bucke
I. Grant decrypt access to the QuickSight rol
J. Update the QuickSight permissions in the marketing account to grant access to the S3 bucket.
K. Create an IAM role in the sales account and grant access to the S3 bucke
L. From the marketing account, assume the IAM role in the sales account to access the S3 bucke
M. Update the QuickSight rote, to create a trust relationship with the new IAM role in the sales account.

Answer: D

Explanation:
Create an IAM role in the sales account and grant access to the S3 bucket. From the marketing account, assume the IAM role in the sales account to access the S3 bucket. Update the QuickSight role, to create a trust relationship with the new IAM role in the sales account.
This approach is the most secure way to grant cross-account access to the data in the S3 bucket while minimizing operational overhead. By creating an IAM role in the sales account, the marketing team can assume the role in their own account, and have access to the S3 bucket. And updating the QuickSight role, to create a trust relationship with the new IAM role in the sales account will grant the marketing team to access the data in the S3 bucket and use it for data visualization using QuickSight.
AWS Resource Access Manager (AWS RAM) also allows sharing of resources between accounts, but it would require additional management and configuration to set up the sharing, which would increase operational overhead.
Using S3 replication would also replicate the data to the marketing account, but it would not provide the marketing team access to the original data, and also it would increase operational overhead with managing the replication process.
IAM roles and policies, KMS grants and trust relationships are a powerful combination for managing
cross-account access in a secure and efficient manner. References:
- AWS IAM Roles
- AWS KMS - Key Grants
- AWS RAM

NEW QUESTION 131
- (Exam Topic 1)
A company has 50 AWS accounts that are members of an organization in AWS Organizations Each account contains multiple VPCs The company wants to use AWS Transit Gateway to establish connectivity between the VPCs in each member account Each time a new member account is created, the company wants to automate the process of creating a new VPC and a transit gateway attachment.
Which combination of steps will meet these requirements? (Select TWO)

A. From the management account, share the transit gateway with member accounts by using AWS Resource Access Manager
B. Prom the management account, share the transit gateway with member accounts by using an AWS Organizations SCP
C. Launch an AWS CloudFormation stack set from the management account that automatical^/ creates a new VPC and a VPC transit gateway attachment in a member accoun
D. Associate the attachment with the transit gateway in the management account by using the transit gateway ID.
E. Launch an AWS CloudFormation stack set from the management account that automatical^ creates a new VPC and a peering transit gateway attachment in a member accoun
F. Share the attachment with the transit gateway in the management account by using a transit gateway service-linked role.
G. From the management account, share the transit gateway with member accounts by using AWS Service Catalog

Answer: AC

Explanation:
https://aws.amazon.com/blogs/mt/self-service-vpcs-in-aws-control-tower-using-aws-service-catalog/ https://docs.aws.amazon.com/vpc/latest/tgw/tgw-transit-gateways.html
https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-resource-ec2-transitgatewayattachme

NEW QUESTION 136
.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual AWS-Certified-Solutions-Architect-Professional Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the AWS-Certified-Solutions-Architect-Professional Product From:

https://www.2passeasy.com/dumps/AWS-Certified-Solutions-Architect-Professional/

Money Back Guarantee

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