Amazon-Web-Services

Exam Questions SAP-C02

AWS Certified Solutions Architect - Professional

NEW QUESTION 1
- (Exam Topic 1)
A company is developing a new serverless API by using Amazon API Gateway and AWS Lambda. The company integrated the Lambda functions with API Gateway to use several shared libraries and custom classes.
A solutions architect needs to simplify the deployment of the solution and optimize for code reuse. Which solution will meet these requirements?

A. Deploy the shared libraries and custom classes into a Docker imag
B. Store the image in an S3 bucket.Create a Lambda layer that uses the Docker image as the sourc
C. Deploy the API's Lambda functions as Zip package
D. Configure the packages to use the Lambda layer.
E. Deploy the shared libraries and custom classes to a Docker imag
F. Upload the image to Amazon Elastic Container Registry (Amazon ECR). Create a Lambda layer that uses the Docker image as the sourc
G. Deploy the API's Lambda functions as Zip package
H. Configure the packages to use the Lambda layer.
I. Deploy the shared libraries and custom classes to a Docker container in Amazon Elastic Container Service (Amazon ECS) by using the AWS Fargate launch typ
J. Deploy the API's Lambda functions as Zip package
K. Configure the packages to use the deployed container as a Lambda layer.
L. Deploy the shared libraries, custom classes, and code for the API's Lambda functions to a Docker imag
M. Upload the image to Amazon Elastic Container Registry (Amazon ECR). Configure the API's Lambda functions to use the Docker image as the deployment package.

Answer: B

Explanation:
Deploying the shared libraries and custom classes to a Docker image and uploading the image to Amazon Elastic Container Registry (Amazon ECR) and creating a Lambda layer that uses the Docker image as the source. Then, deploying the API's Lambda functions as Zip packages and configuring the packages to use the Lambda layer would meet the requirements for simplifying the deployment and optimizing for code reuse.
A Lambda layer is a distribution mechanism for libraries, custom runtimes, and other function dependencies. It allows you to manage your in-development function code separately from your dependencies, this way you can easily update your dependencies without having to update your entire function code.
By deploying the shared libraries and custom classes to a Docker image and uploading the image to Amazon Elastic Container Registry (ECR), it makes it easy to manage and version the dependencies. This way, the company can use the same version of the dependencies across different Lambda functions.
By creating a Lambda layer that uses the Docker image as the source, the company can configure the API's Lambda functions to use the layer, reducing the need to include the dependencies in each function package, and making it easy to update the dependencies across all functions at once.
Reference:
AWS Lambda Layers documentation: https://docs.aws.amazon.com/lambda/latest/dg/configuration-layers.html
AWS Elastic Container Registry (ECR) documentation: https://aws.amazon.com/ecr/ Building Lambda Layers with Docker documentation:
https://aws.amazon.com/blogs/compute/building-lambda-layers-with-docker/

NEW QUESTION 2
- (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 3
- (Exam Topic 1)
An application is using an Amazon RDS for MySQL Multi-AZ DB instance in the us-east-1 Region. After a failover test, the application lost the connections to the database and could not re-establish the connections. After a restart of the application, the application re-established the connections.
A solutions architect must implement a solution so that the application can re-establish connections to the database without requiring a restart.
Which solution will meet these requirements?

A. Create an Amazon Aurora MySQL Serverless v1 DB instanc
B. Migrate the RDS DB instance to the Aurora Serverless v1 DB instanc
C. Update the connection settings in the application to point to the Aurora reader endpoint.
D. Create an RDS prox
E. Configure the existing RDS endpoint as a targe
F. Update the connection settings in the application to point to the RDS proxy endpoint.

G. Create a two-node Amazon Aurora MySQL DB cluste
H. Migrate the RDS DB instance to the Aurora DB cluste
I. Create an RDS prox
J. Configure the existing RDS endpoint as a targe
K. Update the connection settings in the application to point to the RDS proxy endpoint.
L. Create an Amazon S3 bucke
M. Export the database to Amazon S3 by using AWS Database Migration Service (AWS DMS). Configure Amazon Athena to use the S3 bucket as a data stor
N. Install the latest Open Database Connectivity (ODBC) driver for the applicatio
O. Update the connection settings in the application to point to the Athena endpoint

Answer: B

Explanation:
Amazon RDS Proxy is a fully managed database proxy service for Amazon Relational Database Service (RDS) that makes applications more scalable, resilient, and secure. It allows applications to pool and share connections to an RDS database, which can help reduce database connection overhead, improve scalability, and provide automatic failover and high availability.

NEW QUESTION 4
- (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 5
- (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 6
- (Exam Topic 1)
A company runs an IoT platform on AWS IoT sensors in various locations send data to the company's Node js API servers on Amazon EC2 instances running behind an Application Load Balancer The data is stored in an Amazon RDS MySQL DB instance that uses a 4 TB General Purpose SSD volume
The number of sensors the company has deployed in the field has increased over time and is expected to grow significantly The API servers are consistently overloaded and RDS metrics show high write latency
Which of the following steps together will resolve the issues permanently and enable growth as new sensors are provisioned, while keeping this platform cost-efficient? {Select TWO.)

A. Resize the MySQL General Purpose SSD storage to 6 TB to improve the volume's IOPS
B. Re-architect the database tier to use Amazon Aurora instead of an RDS MySQL DB instance and add read replicas
C. Leverage Amazon Kinesis Data Streams and AWS Lambda to ingest and process the raw data
D. Use AWS X-Ray to analyze and debug application issues and add more API servers to match the load
E. Re-architect the database tier to use Amazon DynamoDB instead of an RDS MySQL DB instance

Answer: CE

Explanation:
⟩⟩ Option C is correct because leveraging Amazon Kinesis Data Streams and AWS Lambda to ingest and process the raw data resolves the issues permanently and enable growth as new sensors are provisioned. Amazon Kinesis Data Streams is a serverless streaming data service that simplifies the capture, processing, and storage of data streams at any scale. Kinesis Data Streams can handle any amount of streaming data and process data from hundreds of thousands of sources with very low latency. AWS Lambda is a serverless compute service that lets you run code without provisioning or managing servers. Lambda can be triggered by Kinesis Data Streams events and process the data records in real time. Lambda can also scale automatically based on the incoming data volume. By using Kinesis Data Streams and Lambda, the company can reduce the load on the API servers and improve the performance and scalability of the data ingestion and processing layer3

⟩⟩ Option E is correct because re-architecting the database tier to use Amazon DynamoDB instead of an RDS MySQL DB instance resolves the issues permanently and enable growth as new sensors are provisioned. Amazon DynamoDB is a fully managed key-value and document database that delivers single-digit millisecond performance at any scale. DynamoDB supports auto scaling, which automatically adjusts read and write capacity based on actual traffic patterns. DynamoDB also supports on-demand capacity mode, which instantly accommodates up to double the previous peak traffic on a table. By using DynamoDB instead of RDS MySQL DB instance, the company can eliminate high write latency and improve scalability and performance of the database tier.
References: 1: https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-volume-types.html 2:
https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/CHAP_AuroraOverview.html 3:
https://docs.aws.amazon.com/streams/latest/dev/introduction.html : https://docs.aws.amazon.com/lambda/latest/dg/welcome.html :
https://docs.aws.amazon.com/xray/latest/devguide/aws-xray.html : https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Introduction.html :

NEW QUESTION 7
- (Exam Topic 1)
A company is refactoring its on-premises order-processing platform in the AWS Cloud. The platform includes a web front end that is hosted on a fleet of VMs RabbitMQ to connect the front end to the backend, and a Kubernetes cluster to run a containerized backend system to process the orders. The company does not want to make any major changes to the application
Which solution will meet these requirements with the LEAST operational overhead?

A. Create an AMI of the web server VM Create an Amazon EC2 Auto Scaling group that uses the AMI and an Application Load Balancer Set up Amazon MQ to replace the on-premises messaging queue Configure Amazon Elastic Kubernetes Service (Amazon EKS) to host the order-processing backend
B. Create a custom AWS Lambda runtime to mimic the web server environment Create an Amazon API Gateway API to replace the front-end web servers Set up Amazon MQ to replace the on-premises messaging queue Configure Amazon Elastic Kubernetes Service (Amazon EKS) to host theorder-processing backend
C. Create an AMI of the web server VM Create an Amazon EC2 Auto Scaling group that uses the AMI and an Application Load Balancer Set up Amazon MQ to replace the on-premises messaging queue Install Kubernetes on a fleet of different EC2 instances to host the order-processing backend
D. Create an AMI of the web server VM Create an Amazon EC2 Auto Scaling group that uses the AMI and an Application Load Balancer Set up an Amazon Simple Queue Service (Amazon SQS) queue to replace the on-premises messaging queue Configure Amazon Elastic Kubernetes Service (Amazon EKS) to host the order-processing backend

Answer: A

Explanation:
https://aws.amazon.com/about-aws/whats-new/2020/11/announcing-amazon-mq-rabbitmq/

NEW QUESTION 8
- (Exam Topic 1)
A company uses 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 9
- (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 10
- (Exam Topic 1)
A company wants to migrate its workloads from on premises to AWS. The workloads run on Linux and Windows. The company has a large on-premises intra structure that consists of physical machines and VMs that host numerous applications.
The company must capture details about the system configuration. system performance. running processes and network coi.net lions of its o. -premises ,on boards. The company also must divide the on-premises applications into groups for AWS migrations. The company needs recommendations for Amazon EC2 instance types so that the company can run its workloads on AWS in the most cost-effective manner.
Which combination of steps should a solutions architect take to meet these requirements? (Select THREE.)

A. Assess the existing applications by installing AWS Application Discovery Agent on the physical machines and VMs.
B. Assess the existing applications by installing AWS Systems Manager Agent on the physical machines and VMs
C. Group servers into applications for migration by using AWS Systems Manager Application Manager.
D. Group servers into applications for migration by using AWS Migration Hub.
E. Generate recommended instance types and associated costs by using AWS Migration Hub.
F. Import data about server sizes into AWS Trusted Adviso
G. Follow the recommendations for cost optimization.

Answer: ADE

Explanation:
https://docs.aws.amazon.com/application-discovery/latest/userguide/discovery-agent.html
https://docs.aws.amazon.com/migrationhub/latest/ug/ec2-recommendations.html

NEW QUESTION 10
- (Exam Topic 1)
A company recently deployed an application on AWS. The application uses Amazon DynamoDB. The company measured the application load and configured the RCUs and WCUs on the DynamoDB table to
match the expected peak load. The peak load occurs once a week for a 4-hour period and is double the average load. The application load is close to the average load tor the rest of the week. The access pattern includes many more writes to the table than reads of the table.
A solutions architect needs to implement a solution to minimize the cost of the table. Which solution will meet these requirements?

A. Use AWS Application Auto Scaling to increase capacity during the peak perio
B. Purchase reserved RCUs and WCUs to match the average load.
C. Configure on-demand capacity mode for the table.
D. Configure DynamoDB Accelerator (DAX) in front of the tabl
E. Reduce the provisioned read capacity to match the new peak load on the table.
F. Configure DynamoDB Accelerator (DAX) in front of the tabl
G. Configure on-demand capacity mode for the table.

Answer: D

Explanation:
This solution meets the requirements by using Application Auto Scaling to automatically increase capacity during the peak period, which will handle the double the average load. And by purchasing reserved RCUs and WCUs to match the average load, it will minimize the cost of the table for the rest of the week when the load is close to the average.

NEW QUESTION 12
- (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 17
- (Exam Topic 1)
A company is migrating some of its applications to AWS. The company wants to migrate and modernize the applications quickly after it finalizes networking and security strategies. The company has set up an AWS Direct Connection connection in a central network account.
The company expects to have hundreds of AWS accounts and VPCs in the near future. The corporate network must be able to access the resources on AWS seamlessly and also must be able to communicate with all the VPCs. The company also wants to route its cloud resources to the internet through its on-premises data center.
Which combination of steps will meet these requirements? (Choose three.)

A. Create a Direct Connect gateway in the central accoun
B. In each of the accounts, create an association proposal by using the Direct Connect gateway and the account ID for every virtual private gateway.
C. Create a Direct Connect gateway and a transit gateway in the central network accoun
D. Attach the transit gateway to the Direct Connect gateway by using a transit VIF.
E. Provision an internet gatewa
F. Attach the internet gateway to subnet
G. Allow internet traffic through the gateway.
H. Share the transit gateway with other account
I. Attach VPCs to the transit gateway.
J. Provision VPC peering as necessary.
K. Provision only private subnet
L. Open the necessary route on the transit gateway and customer gatewayto allow outbound internet traffic from AWS to flow through NAT services that run in the data center.

Answer: BDF

Explanation:
➤ Option A is incorrect because creating a Direct Connect gateway in the central account and creating an association proposal by using the Direct Connect gateway and the account ID for every virtual private gateway does not enable active-passive failover between the regions. A Direct Connect gateway is a globally available resource that enables you to connect your AWS Direct Connect connection over a private virtual interface (VIF) to one or more VPCs in any AWS Region. A virtual private gateway is the VPN concentrator on the Amazon side of a VPN connection. You can associate a Direct Connect gateway with either a transit gateway or a virtual private gateway. However, a Direct Connect gateway does not provide any load balancing or failover capabilities by itself1

➤ Option B is correct because creating a Direct Connect gateway and a transit gateway in the central network account and attaching the transit gateway to the Direct Connect gateway by using a transit VIF meets the requirement of enabling the corporate network to access the resources on AWS seamlessly and also to communicate with all the VPCs. A transit VIF is a type of private VIF that you can use to connect your AWS Direct Connect connection to a transit gateway or a Direct Connect gateway. A transit gateway is a network transit hub that you can use to interconnect your VPCs and on-premises networks. By using a transit VIF, you can route traffic between your on-premises network and multiple
VPCs across different AWS accounts and Regions through a single connection23

➤ Option C is incorrect because provisioning an internet gateway, attaching the internet gateway to subnets, and allowing internet traffic through the gateway does not meet the requirement of routing cloud resources to the internet through its on-premises data center. An internet gateway is a horizontally scaled, redundant, and highly available VPC component that allows communication between your VPC and the internet. An internet gateway serves two purposes: to provide a target in your VPC route tables for internet-routable traffic, and to perform network address translation (NAT) for instances that have been assigned

public IPv4 addresses. By using an internet gateway, you are routing cloud resources directly to the internet, not through your on-premises data center.

⟩ Option D is correct because sharing the transit gateway with other accounts and attaching VPCs to the transit gateway meets the requirement of enabling the corporate network to access the resources on AWS seamlessly and also to communicate with all the VPCs. You can share your transit gateway with other AWS accounts within the same organization by using AWS Resource Access Manager (AWS RAM). This allows you to centrally manage connectivity from multiple accounts without having to create individual peering connections between VPCs or duplicate network appliances in each account. You can attach VPCs from different accounts and Regions to your shared transit gateway and enable routing between them.

⟩ Option E is incorrect because provisioning VPC peering as necessary does not meet the requirement of enabling the corporate network to access the resources on AWS seamlessly and also to communicate with all the VPCs. VPC peering is a networking connection between two VPCs that enables you to route traffic between them using private IPv4 addresses or IPv6 addresses. You can create a VPC peering connection between your own VPCs, or with a VPC in another AWS account within a single Region. However, VPC peering does not allow you to route traffic from your on-premises network to your VPCs or between multiple Regions. You would need to create multiple VPN connections or Direct Connect connections for each VPC peering connection, which increases operational complexity and costs.

⟩ Option F is correct because provisioning only private subnets, opening the necessary route on the transit gateway and customer gateway to allow outbound internet traffic from AWS to flow through NAT services that run in the data center meets the requirement of routing cloud resources to the internet through its on-premises data center. A private subnet is a subnet that's associated with a route table that has no route to an internet gateway. Instances in a private subnet can communicate with other instances in the same VPC but cannot access resources on the internet directly. To enable outbound internet access from instances in private subnets, you can use NAT devices such as NAT gateways or NAT instances that are deployed in public subnets. A public subnet is a subnet that's associated with a route table that has a route to an internet gateway. Alternatively, you can use your on-premises data center as a NAT device by configuring routes on your transit gateway and customer gateway that direct outbound internet traffic from your private subnets through your VPN connection or Direct Connect connection. This way, you can route cloud resources to the internet through your on-premises data center instead of using an internet gateway.
References: 1:
https://docs.aws.amazon.com/directconnect/latest/UserGuide/direct-connect-gateways-intro.html 2:
https://docs.aws.amazon.com/directconnect/latest/UserGuide/direct-connect-transit-virtual-interfaces.html 3: https://docs.aws.amazon.com/vpc/latest/tgw/what-is-transit-gateway.html : https://docs.aws.amazon.com/vpc/latest/userguide/VPC_Internet_Gateway.html : https://docs.aws.amazon.com/vpc/latest/tgw/tgw-sharing.html : https://docs.aws.amazon.com/vpc/latest/peering/what-is-vpc-peering.html : https://docs.aws.amazon.com/vpc/latest/userguide/VPC_Scenario2.html : https://docs.aws.amazon.com/vpc/latest/userguide/VPC_Scenario3.html : https://docs.aws.amazon.com/vpc/latest/userguide/VPC_NAT_Instance.html : https://docs.aws.amazon.com/vpc/latest/userguide/VPC_NAT_Gateway.html

NEW QUESTION 20
- (Exam Topic 1)
A company manages multiple AWS accounts by using AWS Organizations. Under the root OU. the company has two OUs: Research and DataOps.
Because of regulatory requirements, all resources that the company deploys in the organization must reside in the ap-northeast-1 Region. Additionally. EC2 instances that the company deploys in the DataOps OU must use a predefined list of instance types
A solutions architect must implement a solution that applies these restrictions. The solution must maximize operational efficiency and must minimize ongoing maintenance
Which combination of steps will meet these requirements? (Select TWO)

A. Create an IAM role in one account under the DataOps OU Use the ec2 Instance Type condition key in an inline policy on the role to restrict access to specific instance types.
B. Create an IAM user in all accounts under the root OU Use the aws RequestedRegion condition key in an inline policy on each user to restrict access to all AWS Regions except ap-northeast-1.
C. Create an SCP Use the aws:RequestedRegion condition key to restrict access to all AWS Regions except ap-northeast-1 Apply the SCP to the root OU.
D. Create an SCP Use the ec2Reo»on condition key to restrict access to all AWS Regions except ap-northeast-1. Apply the SCP to the root O
E. the DataOps O
F. and the Research OU.
G. Create an SCP Use the ec2:lnstanceType condition key to restrict access to specific instance types Apply the SCP to the DataOps OU.

Answer: CE

Explanation:
https://docs.aws.amazon.com/IAM/latest/UserGuide/reference_policies_examples_aws_deny-requested-region.h
https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps_examples_ec2.html

NEW QUESTION 22
- (Exam Topic 1)
A company is running an application in the AWS Cloud. The application runs on containers in an Amazon Elastic Container Service (Amazon ECS) cluster. The ECS tasks use the Fargate launch type. The application's data is relational and is stored in Amazon Aurora MySQL. To meet regulatory requirements, the application must be able to recover to a separate AWS Region in the event of an application failure. In case of a failure, no data can be lost. Which solution will meet these requirements with the LEAST amount of operational overhead?

A. Provision an Aurora Replica in a different Region.
B. Set up AWS DataSync for continuous replication of the data to a different Region.
C. Set up AWS Database Migration Service (AWS DMS) to perform a continuous replication of the data to a different Region.
D. Use Amazon Data Lifecycle Manager {Amazon DLM) to schedule a snapshot every 5 minutes.

Answer: A

Explanation:
Provision an Aurora Replica in a different Region will meet the requirement of the application being able to recover to a separate AWS Region in the event of an application failure, and no data can be lost, with the least amount of operational overhead.

NEW QUESTION 23
- (Exam Topic 1)
A software company has deployed an application that consumes a REST API by using Amazon API Gateway. AWS Lambda functions. and an Amazon DynamoDB table. The application is showing an increase in the number of errors during PUT requests. Most of the PUT calls come from a small number of clients that are authenticated with specific API keys.
A solutions architect has identified that a large number of the PUT requests originate from one client. The API is noncritical, and clients can tolerate retries of unsuccessful calls. However, the errors are displayed to customers and are causing damage to the API's reputation.
What should the solutions architect recommend to improve the customer experience?

A. Implement retry logic with exponential backoff and irregular variation in the client applicatio
B. Ensure that the errors are caught and handled with descriptive error messages.
C. Implement API throttling through a usage plan at the API Gateway leve
D. Ensure that the client application handles code 429 replies without error.
E. Turn on API caching to enhance responsiveness for the production stag
F. Run 10-minute load tests.Verify that the cache capacity is appropriate for the workload.
G. Implement reserved concurrency at the Lambda function level to provide the resources that are needed during sudden increases in traffic.

Answer: B

Explanation:
https://aws.amazon.com/premiumsupport/knowledge-center/aws-batch-requests-error/ https://aws.amazon.com/premiumsupport/knowledge-center/api-gateway-429-limit/

NEW QUESTION 26
- (Exam Topic 1)
A company is running an application on several Amazon EC2 instances in an Auto Scaling group behind an Application Load Balancer. The load on the application varies throughout the day, and EC2 instances are scaled in and out on a regular basis. Log files from the EC2 instances are copied to a central Amazon S3 bucket every 15 minutes. The security team discovers that log files are missing from some of the terminated EC2 instances.
Which set of actions will ensure that log files are copied to the central S3 bucket from the terminated EC2 instances?

A. Create a script to copy log files to Amazon S3, and store the script in a file on the EC2 instanc
B. Create an Auto Scaling lifecycle hook and an Amazon EventBridge (Amazon CloudWatch Events) rule to detect lifecycle events from the Auto Scaling grou
C. Invoke an AWS Lambda function on the autoscaling:EC2_INSTANCE_TERMINATING transition to send ABANDON to the Auto Scaling group to prevent termination, run the script to copy the log files, and terminate the instance using the AWS SDK.
D. Create an AWS Systems Manager document with a script to copy log files to Amazon S3. Create an Auto Scaling lifecycle hook and an Amazon EventBridge (Amazon CloudWatch Events) rule to detect lifecycle events from the Auto Scaling grou
E. Invoke an AWS Lambda function on the autoscaling:EC2_INSTANCE_TERMINATING transition to call the AWS Systems Manager API SendCommand operation to run the document to copy the log files and send CONTINUE to the Auto Scaling group to terminate the instance.
F. Change the log delivery rate to every 5 minute
G. Create a script to copy log files to Amazon S3, and add the script to EC2 instance user dat
H. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to detect EC2 instance terminatio
I. Invoke an AWS Lambda function from the EventBridge (CloudWatch Events) rule that uses the AWS CLI to run the user-data script to copy the log files and terminate the instance.
J. Create an AWS Systems Manager document with a script to copy log files to Amazon S3. Create an Auto Scaling lifecycle hook that publishes a message to an Amazon Simple Notification Service (Amazon SNS) topi
K. From the SNS notification, call the AWS Systems Manager API SendCommand operation to run the document to copy the log files and send ABANDON to the Auto Scaling group to terminate the instance.

Answer: B

Explanation:
https://docs.aws.amazon.com/autoscaling/ec2/userguide/adding-lifecycle-hooks.html
- Refer to Default Result section - If the instance is terminating, both abandon and continue allow the instance to terminate. However, abandon stops any remaining actions, such as other lifecycle hooks, and continue allows any other lifecycle hooks to complete.
https://aws.amazon.com/blogs/infrastructure-and-automation/run-code-before-terminating-an-ec2-auto-scaling-i https://github.com/aws-samples/aws-lambda-lifecycle-hooks-function
https://github.com/aws-samples/aws-lambda-lifecycle-hooks-function/blob/master/cloudformation/template.yam

NEW QUESTION 31
- (Exam Topic 1)
A finance company hosts a data lake in Amazon S3. The company receives financial data records over SFTP each night from several third parties. The company runs its own SFTP server on an Amazon EC2 instance in a public subnet of a VPC. After the files ate uploaded, they are moved to the data lake by a cron job that runs on the same instance. The SFTP server is reachable on DNS sftp.examWe.com through the use of Amazon Route 53.
What should a solutions architect do to improve the reliability and scalability of the SFTP solution?

A. Move the EC2 instance into an Auto Scaling grou
B. Place the EC2 instance behind an Application Load Balancer (ALB). Update the DNS record sftp.example.com in Route 53 to point to the ALB.
C. Migrate the SFTP server to AWS Transfer for SFT
D. Update the DNS record sftp.example.com in Route 53 to point to the server endpoint hostname.
E. Migrate the SFTP server to a file gateway in AWS Storage Gatewa
F. Update the DNS record sflp.example.com in Route 53 to point to the file gateway endpoint.
G. Place the EC2 instance behind a Network Load Balancer (NLB). Update the DNS record sftp.example.com in Route 53 to point to the NLB.

Answer: B

Explanation:
https://aws.amazon.com/aws-transfer-family/faqs/ https://docs.aws.amazon.com/transfer/latest/userguide/what-is-aws-transfer-family.html
https://aws.amazon.com/about-aws/whats-new/2018/11/aws-transfer-for-sftp-fully-managed-sftp-for-s3/?nc1=h_

NEW QUESTION 32
- (Exam Topic 1)
A company has an application that runs on Amazon EC2 instances. A solutions architect is designing VPC infrastructure in an AWS Region where the application needs to access an Amazon Aurora DB cluster. The EC2 instances are all associated with the same security group. The DB cluster is associated with its own security group.
The solutions architect needs to add rules to the security groups to provide the application with least privilege access to the DB cluster.
Which combination of steps will meet these requirements? (Select TWO.)

A. Add an inbound rule to the EC2 instances' security grou
B. Specify the DB cluster's security group as the source over the default Aurora port.

C. Add an outbound rule to the EC2 instances' security grou
D. Specify the DB cluster's security group as the destination over the default Aurora port.
E. Add an inbound rule to the DB cluster's security grou
F. Specify the EC2 instances' security group as the source over the default Aurora port.
G. Add an outbound rule to the DB cluster's security grou
H. Specify the EC2 instances' security group as the destination over the default Aurora port.
I. Add an outbound rule to the DB cluster's security grou
J. Specify the EC2 instances' security group as the destination over the ephemeral ports.

Answer: AB

Explanation:
* B. Add an outbound rule to the EC2 instances' security group. Specify the DB cluster's security group as the destination over the default Aurora port. This allows the instances to make outbound connections to the DB cluster on the default Aurora port. C. Add an inbound rule to the DB cluster's security group. Specify the EC2 instances' security group as the source over the default Aurora port. This allows connections to the DB cluster from the EC2 instances on the default Aurora port.

NEW QUESTION 35
- (Exam Topic 1)
A solutions architect has developed a web application that uses an Amazon API Gateway Regional endpoint and an AWS Lambda function. The consumers of the web application are all close to the AWS Region where the application will be deployed. The Lambda function only queries an Amazon Aurora MySQL database. The solutions architect has configured the database to have three read replicas.
During testing, the application does not meet performance requirements. Under high load, the application opens a large number of database connections. The solutions architect must improve the application's performance.
Which actions should the solutions architect take to meet these requirements? (Choose two.)

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

Answer: BD

Explanation:
Connect to RDS outside of Lambda handler method to improve performance https://awstut.com/en/2022/04/30/connect-to-rds-outside-of-lambda-handler-method-to-improve-performance-en
Using RDS Proxy, you can handle unpredictable surges in database traffic. Otherwise, these surges might cause issues due to oversubscribing connections or creating new connections at a fast rate. RDS Proxy establishes a database connection pool and reuses connections in this pool. This approach avoids the memory and CPU overhead of opening a new database connection each time. To protect the database against oversubscription, you can control the number of database connections that are created. https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/rds-proxy.html

NEW QUESTION 39
- (Exam Topic 1)
A life sciences company is using a combination of open source tools to manage data analysis workflows and Docker containers running on servers in its on-premises data center to process genomics data Sequencing data is generated and stored on a local storage area network (SAN), and then the data is processed. The research and development teams are running into capacity issues and have decided to re-architect their genomics analysis platform on AWS to scale based on workload demands and reduce the turnaround time from weeks to days
The company has a high-speed AWS Direct Connect connection Sequencers will generate around 200 GB of data for each genome, and individual jobs can take several hours to process the data with ideal compute capacity. The end result will be stored in Amazon S3. The company is expecting 10-15 job requests each day
Which solution meets these requirements?

A. Use regularly scheduled AWS Snowball Edge devices to transfer the sequencing data into AWS When AWS receives the Snowball Edge device and the data is loaded into Amazon S3 use S3 events to trigger an AWS Lambda function to process the data
B. Use AWS Data Pipeline to transfer the sequencing data to Amazon S3 Use S3 events to trigger an Amazon EC2 Auto Scaling group to launch custom-AMI EC2 instances running the Docker containers to process the data
C. Use AWS DataSync to transfer the sequencing data to Amazon S3 Use S3 events to trigger an AWS Lambda function that starts an AWS Step Functions workflow Store the Docker images in Amazon Elastic Container Registry (Amazon ECR) and trigger AWS Batch to run the container and process the sequencing data
D. Use an AWS Storage Gateway file gateway to transfer the sequencing data to Amazon S3 Use S3 events to trigger an AWS Batch job that runs on Amazon EC2 instances running the Docker containers to process the data

Answer: C

Explanation:
AWS DataSync can be used to transfer the sequencing data to Amazon S3, which is a more efficient and faster method than using Snowball Edge devices. Once the data is in S3, S3 events can trigger an AWS Lambda function that starts an AWS Step Functions workflow. The Docker images can be stored in Amazon Elastic Container Registry (Amazon ECR) and AWS Batch can be used to run the container and process the sequencing data.

NEW QUESTION 43
- (Exam Topic 1)
A company is planning to migrate its business-critical applications from an on-premises data center to AWS. The company has an on-premises installation of a Microsoft SQL Server Always On cluster. The company wants to migrate to an AWS managed database service. A solutions architect must design a heterogeneous database migration on AWS.
Which solution will meet these requirements?

A. Migrate the SQL Server databases to Amazon RDS for MySQL by using backup and restore utilities.
B. Use an AWS Snowball Edge Storage Optimized device to transfer data to Amazon S3. Set up Amazon RDS for MySQ
C. Use S3 integration with SQL Server features, such as BULK INSERT.
D. Use the AWS Schema Conversion Tool to translate the database schema to Amazon RDS for MeSQ

E. Then use AWS Database Migration Service (AWS DMS) to migrate the data from on-premises databases to Amazon RDS.
F. Use AWS DataSync to migrate data over the network between on-premises storage and Amazon S3. Set up Amazon RDS for MySQ
G. Use S3 integration with SQL Server features, such as BULK INSERT.

Answer: C

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

NEW QUESTION 47
- (Exam Topic 1)
An enterprise company wants to allow its developers to purchase third-party software through AWS Marketplace. The company uses an AWS Organizations account structure with full features enabled, and has a shared services account in each organizational unit (OU) that will be used by procurement managers. The procurement team's policy indicates that developers should be able to obtain third-party software from an approved list only and use Private Marketplace in AWS Marketplace to achieve this requirement . The procurement team wants administration of Private Marketplace to be restricted to a role named procurement-manager-role, which could be assumed by procurement managers Other IAM users groups, roles, and account administrators in the company should be denied Private Marketplace administrative access
What is the MOST efficient way to design an architecture to meet these requirements?

A. Create an IAM role named procurement-manager-role in all AWS accounts in the organization Add the PowerUserAccess managed policy to the role Apply an inline policy to all IAM users and roles in every AWS account to deny permissions on the AWSPrivateMarketplaceAdminFullAccess managed policy.
B. Create an IAM role named procurement-manager-role in all AWS accounts in the organization Add the AdministratorAccess managed policy to the role Define a permissions boundary with the AWSPrivateMarketplaceAdminFullAccess managed policy and attach it to all the developer roles.
C. Create an IAM role named procurement-manager-role in all the shared services accounts in the organization Add the AWSPrivateMarketplaceAdminFullAccess managed policy to the role Create an organization root-level SCP to deny permissions to administer Private Marketplace to everyone exceptthe role named procurement-manager-role Create another organization root-level SCP to deny permissions to create an IAM role named procurement-manager-role to everyone in the organization.
D. Create an IAM role named procurement-manager-role in all AWS accounts that will be used by developer
E. Add the AWSPrivateMarketplaceAdminFullAccess managed policy to the rol
F. Create an SCP in Organizations to deny permissions to administer Private Marketplace to everyone except the role named procurement-manager-rol
G. Apply the SCP to all the shared services accounts in the organization.

Answer: C

Explanation:
SCP to deny permissions to administer Private Marketplace to everyone except the role named procurement-manager-role.
https://aws.amazon.com/blogs/awsmarketplace/controlling-access-to-a-well-architected-private-marketplace-usi
This approach allows the procurement managers to assume the procurement-manager-role in shared services accounts, which have the AWSPrivateMarketplaceAdminFullAccess managed policy attached to it and can then manage the Private Marketplace. The organization root-level SCP denies the permission to administer Private Marketplace to everyone except the role named procurement-manager-role and another SCP denies the permission to create an IAM role named procurement-manager-role to everyone in the organization, ensuring that only the procurement team can assume the role and manage the Private Marketplace. This approach provides a centralized way to manage and restrict access to Private Marketplace while maintaining a high level of security.

NEW QUESTION 52
- (Exam Topic 1)
A company has developed APIs that use Amazon API Gateway with Regional endpoints. The APIs call AWS Lambda functions that use API Gateway authentication mechanisms. After a design review, a solutions architect identifies a set of APIs that do not require public access.
The solutions architect must design a solution to make the set of APIs accessible only from a VPC. All APIs need to be called with an authenticated user.
Which solution will meet these requirements with the LEAST amount of effort?

A. Create an internal Application Load Balancer (ALB). Create a target grou
B. Select the Lambda function to cal
C. Use the ALB DNS name to call the API from the VPC.
D. Remove the DNS entry that is associated with the API in API Gatewa
E. Create a hosted zone in Amazon Route 53. Create a CNAME record in the hosted zon
F. Update the API in API Gateway with the CNAME recor
G. Use the CNAME record to call the API from the VPC.
H. Update the API endpoint from Regional to private in API Gatewa
I. Create an interface VPC endpoint in the VP
J. Create a resource policy, and attach it to the AP
K. Use the VPC endpoint to call the API from the VPC.
L. Deploy the Lambda functions inside the VP
M. Provision an EC2 instance, and install an Apache server.From the Apache server, call the Lambda function
N. Use the internal CNAME record of the EC2 instance to call the API from the VPC.

Answer: C

Explanation:
This solution requires the least amount of effort as it only requires to update the API endpoint to private in API Gateway and create an interface VPC endpoint. Then create a resource policy and attach it to the API. This will make the API only accessible from the VPC and still keep the authentication mechanism intact.
Reference:
➤ https://aws.amazon.com/api-gateway/features/

NEW QUESTION 57
- (Exam Topic 1)

A security engineer determined that an existing application retrieves credentials to an Amazon RDS for MySQL database from an encrypted file in Amazon S3. For the next version of the application, the security engineer wants to implement the following application design changes to improve security:

The database must use strong, randomly generated passwords stored in a secure AWS managed service.

The application resources must be deployed through AWS CloudFormation.

The application must rotate credentials for the database every 90 days.

A solutions architect will generate a CloudFormation template to deploy the application.
Which resources specified in the CloudFormation template will meet the security engineer's requirements with the LEAST amount of operational overhead?

A. Generate the database password as a secret resource using AWS Secrets Manage
B. Create an AWS Lambda function resource to rotate the database passwor
C. Specify a Secrets Manager RotationSchedule resource to rotate the database password every 90 days.
D. Generate the database password as a SecureString parameter type using AWS Systems Manager Parameter Stor
E. Create an AWS Lambda function resource to rotate the database passwor
F. Specify a Parameter Store RotationSchedule resource to rotate the database password every 90 days.
G. Generate the database password as a secret resource using AWS Secrets Manage
H. Create an AWS Lambda function resource to rotate the database passwor
I. Create an Amazon EventBridge scheduled rule resource to trigger the Lambda function password rotation every 90 days.
J. Generate the database password as a SecureString parameter type using AWS Systems Manager Parameter Stor
K. Specify an AWS AppSync DataSource resource to automatically rotate the database password every 90 days.

Answer: B

Explanation:
https://aws.amazon.com/blogs/security/how-to-securely-provide-database-credentials-to-lambda-functions-by-us
https://docs.aws.amazon.com/secretsmanager/latest/userguide/rotating-secrets.html
https://docs.aws.amazon.com/secretsmanager/latest/userguide/integrating_cloudformation.html

NEW QUESTION 58
- (Exam Topic 1)
A company wants to migrate an application to Amazon EC2 from VMware Infrastructure that runs in an
on-premises data center. A solutions architect must preserve the software and configuration settings during the migration.
What should the solutions architect do to meet these requirements?

A. Configure the AWS DataSync agent to start replicating the data store to Amazon FSx for Windows FileServer Use the SMB share to host the VMware data stor
B. Use VM Import/Export to move the VMs to Amazon EC2.
C. Use the VMware vSphere client to export the application as an image in Open Virealization Format (OVF) format Create an Amazon S3 bucket to store the image in the destination AWS Regio
D. Create and apply an IAM role for VM Import Use the AWS CLI to run the EC2 import command.
E. . Configure AWS Storage Gateway for files service to export a Common Internet File System (CIFSJ shar
F. Create a backup copy to the shared folde
G. Sign in to the AWS Management Console and create an AMI from the backup copy Launch an EC2 instance that is based on the AMI.
H. Create a managed-instance activation for a hybrid environment in AWS Systems Manage
I. Download and install Systems Manager Agent on the on-premises VM Register the VM with Systems Manager to be a managed instance Use AWS Backup to create a snapshot of the VM and create an AM
J. Launch an EC2 instance that is based on the AMI

Answer: D

Explanation:
https://docs.aws.amazon.com/vm-import/latest/userguide/vmimport-image-import.html
- Export an OVF Template
- Create / use an Amazon S3 bucket for storing the exported images. The bucket must be in the Region where you want to import your VMs.
- Create an IAM role named vmimport.
- You'll use AWS CLI to run the import commands. https://aws.amazon.com/premiumsupport/knowledge-center/import-instances/

NEW QUESTION 60
- (Exam Topic 1)
A company gives users the ability to upload images from a custom application. The upload process invokes an AWS Lambda function that processes and stores the image in an Amazon S3 bucket. The application invokes the Lambda function by using a specific function version ARN.
The Lambda function accepts image processing parameters by using environment variables. The company often adjusts the environment variables of the Lambda function to achieve optimal image processing output. The company tests different parameters and publishes a new function version with the updated environment variables after validating results. This update process also requires frequent changes to the custom application to invoke the new function version ARN. These changes cause interruptions for users.
A solutions architect needs to simplify this process to minimize disruption to users. Which solution will meet these requirements with the LEAST operational overhead?

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

Answer: D

Explanation:
A Lambda function alias allows you to point to a specific version of a function and also can be updated to point to a new version of the function without modifying

the client application. This way, the company can test different versions of the function with different environment variables and, once the optimal parameters are found, update the alias to point to the new version, without the need to update the client application.

By using this approach, the company can simplify the process of updating the environment variables, minimize disruption to users, and reduce the operational overhead.

Reference:

AWS Lambda documentation: https://aws.amazon.com/lambda/

AWS Lambda Aliases documentation: https://docs.aws.amazon.com/lambda/latest/dg/aliases-intro.html AWS Lambda versioning and aliases documentation: https://aws.amazon.com/blogs/compute/versioning-aliases-in-aws-lambda/

NEW QUESTION 63
- (Exam Topic 1)
A company is subject to regulatory audits of its financial information. External auditors who use a single AWS account need access to the company's AWS account. A solutions architect must provide the auditors with secure, read-only access to the company's AWS account. The solution must comply with AWS security best practices.

Which solution will meet these requirements?

A. In the company's AWS account, create resource policies for all resources in the account to grant access to the auditors' AWS accoun
B. Assign a unique external ID to the resource policy.
C. In the company's AWS account create an IAM role that trusts the auditors' AWS account Create an IAM policy that has the required permission
D. Attach the policy to the rol
E. Assign a unique external ID to the role's trust policy.
F. In the company's AWS account, create an IAM use
G. Attach the required IAM policies to the IAM user.Create API access keys for the IAM use
H. Share the access keys with the auditors.
I. In the company's AWS account, create an IAM group that has the required permissions Create an IAM user in the company s account for each audito
J. Add the IAM users to the IAM group.

Answer: B

Explanation:
This solution will allow the external auditors to have read-only access to the company's AWS account while being compliant with AWS security best practices. By creating an IAM role, which is a secure and flexible way of granting access to AWS resources, and trusting the auditors' AWS account, the company can ensure that the auditors only have the permissions that are required for their role and nothing more. Assigning a unique external ID to the role's trust policy, it will ensure that only the auditors' AWS account can assume the role.

Reference:

AWS IAM Roles documentation: https://aws.amazon.com/iam/features/roles/ AWS IAM Best practices: https://aws.amazon.com/iam/security-best-practices/

NEW QUESTION 68
- (Exam Topic 1)
A company is planning to store a large number of archived documents and make the documents available to employees through the corporate intranet. Employees will access the system by connecting through a client VPN service that is attached to a VPC. The data must not be accessible to the public.

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

Which solution will meet these requirements at the LOWEST cost?

A. Create an Amazon S3 bucke
B. Configure the S3 bucket to use the S3 One Zone-Infrequent Access (S3 One Zone-IA) storage class as defaul
C. Configure the S3 bucket for website hostin
D. Create an S3 interface endpoin
E. Configure the S3 bucket to allow access only through that endpoint.
F. Launch an Amazon EC2 instance that runs a web serve
G. Attach an Amazon Elastic File System (Amazon EFS) file system to store the archived data in the EFS One Zone-Infrequent Access (EFS One Zone-IA) storage class Configure the instance security groups to allow access only from private networks.
H. Launch an Amazon EC2 instance that runs a web server Attach an Amazon Elastic Block Store (Amazon EBS) volume to store the archived dat
I. Use the Cold HDD (sc1) volume typ
J. Configure the instance security groups to allow access only from private networks.
K. Create an Amazon S3 bucke
L. Configure the S3 bucket to use the S3 Glacier Deep Archive storage class as defaul
M. Configure the S3 bucket for website hostin
N. Create an S3 interface endpoin
O. Configure the S3 bucket to allow access only through that endpoint.

Answer: D

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

NEW QUESTION 70
- (Exam Topic 1)
A company plans to refactor a monolithic application into a modern application designed deployed or AWS. The CLCD pipeline needs to be upgraded to support the modem design for the application with the following requirements
• It should allow changes to be released several times every hour.
* It should be able to roll back the changes as quickly as possible Which design will meet these requirements?

A. Deploy a CI-CD pipeline that incorporates AMIs to contain the application and their configurationsDeploy the application by replacing Amazon EC2 instances

B. Specify AWS Elastic Beanstak to sage in a secondary environment as the deployment target for the CI/CD pipeline of the applicatio
C. To deploy swap the staging and production environment URLs.
D. Use AWS Systems Manager to re-provision the infrastructure for each deployment Update the Amazon EC2 user data to pull the latest code art-fact from Amazon S3 and use Amazon Route 53 weighted routing to point to the new environment
E. Roll out At application updates as pan of an Auto Scaling event using prebuilt AMI
F. Use new versions of the AMIs to add instances, and phase out all instances that use the previous AMI version with the configured termination policy during a deployment event.

Answer: B

Explanation:
It is the fastest when it comes to rollback and deploying changes every hour

NEW QUESTION 74
- (Exam Topic 1)
A company runs its application in the eu-west-1 Region and has one account for each of its environments development, testing, and production All the environments are running 24 hours a day 7 days a week by using stateful Amazon EC2 instances and Amazon RDS for MySQL databases The databases are between 500 GB and 800 GB in size
The development team and testing team work on business days during business hours, but the production environment operates 24 hours a day. 7 days a week. The company wants to reduce costs AH resources are tagged with an environment tag with either development, testing, or production as the key. What should a solutions architect do to reduce costs with the LEAST operational effort?

A. Create an Amazon EventBridge (Amazon CloudWatch Events) rule that runs once every day Configure the rule to invoke one AWS Lambda function that starts or stops instances based on the tag day and time.
B. Create an Amazon EventBridge (Amazon CloudWatch Events) rule that runs every business day in the evenin
C. Configure the rule to invoke an AWS Lambda function that stops instances based on thetag-Create a second EventBridge (CloudWatch Events) rule that runs every business day in the morning Configure the second rule to invoke another Lambda function that starts instances based on the tag
D. Create an Amazon EventBridge (Amazon CloudWatch Events) rule that runs every business day in the evening Configure the rule to invoke an AWS Lambda function that terminates instances based on the tag Create a second EventBridge (CloudWatch Events) rule that runs every business day in the morning Configure the second rule to invoke another Lambda function that restores the instances from their last backup based on the tag.
E. Create an Amazon EventBridge rule that runs every hou
F. Configure the rule to invoke one AWS Lambda function that terminates or restores instances from their last backup based on the ta
G. day, and time.

Answer: B

Explanation:
Creating an Amazon EventBridge rule that runs every business day in the evening to stop instances and another rule that runs every business day in the morning to start instances based on the tag will reduce costs with the least operational effort. This approach allows for instances to be stopped during non-business hours when they are not in use, reducing the costs associated with running them. It also allows for instances to be started again in the morning when the development and testing teams need to use them.

NEW QUESTION 78
- (Exam Topic 1)
A company runs a Python script on an Amazon EC2 instance to process data. The script runs every
10 minutes. The script ingests files from an Amazon S3 bucket and processes the files. On average, the script takes approximately 5 minutes to process each file The script will not reprocess a file that the script has already processed.
The company reviewed Amazon CloudWatch metrics and noticed that the EC2 instance is idle for approximately 40% of the time because of the file processing speed. The company wants to make the workload highly available and scalable. The company also wants to reduce long-term management overhead.
Which solution will meet these requirements MOST cost-effectively?

A. Migrate the data processing script to an AWS Lambda functio
B. Use an S3 event notification to invoke the Lambda function to process the objects when the company uploads the objects.
C. Create an Amazon Simple Queue Service (Amazon SQS) queu
D. Configure Amazon S3 to send event notifications to the SQS queu
E. Create an EC2 Auto Scaling group with a minimum size of one instanc
F. Update the data processing script to poll the SQS queu
G. Process the S3 objects that the SQS message identifies.
H. Migrate the data processing script to a container imag
I. Run the data processing container on an EC2 instanc
J. Configure the container to poll the S3 bucket for new objects and to process the resulting objects.
K. Migrate the data processing script to a container image that runs on Amazon Elastic Container Service (Amazon ECS) on AWS Fargat
L. Create an AWS Lambda function that calls the Fargate RunTaskAPI operation when the container processes the fil
M. Use an S3 event notification to invoke the Lambda function.

Answer: D

Explanation:
migrating the data processing script to an AWS Lambda function and using an S3 event notification to invoke the Lambda function to process the objects when the company uploads the objects. This solution meets the company's requirements of high availability and scalability, as well as reducing long-term management overhead, and is likely to be the most cost-effective option.

NEW QUESTION 82
- (Exam Topic 1)
A company is in the process of implementing AWS Organizations to constrain its developers to use only Amazon EC2. Amazon S3 and Amazon DynamoDB. The developers account resides In a dedicated organizational unit (OU). The solutions architect has implemented the following SCP on the developers account:

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

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

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

Answer: B

Explanation:
https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_inheritance_auth.html

NEW QUESTION 85
- (Exam Topic 1)
A company recently acquired several other companies. Each company has a separate AWS account with a different billing and reporting method. The acquiring company has consolidated all the accounts into one organization in AWS Organizations. However, the acquiring company has found it difficult to generate a cost report that contains meaningful groups for all the teams.
The acquiring company's finance team needs a solution to report on costs for all the companies through a self-managed application.
Which solution will meet these requirements?

A. Create an AWS Cost and Usage Report for the organizatio
B. Define tags and cost categories in the repor
C. Create a table in Amazon Athen
D. Create an Amazon QuickSight dataset based on the Athena tabl
E. Share the dataset with the finance team.
F. Create an AWS Cost and Usage Report for the organizatio
G. Define tags and cost categories in the repor
H. Create a specialized template in AWS Cost Explorer that the finance department will use to build reports.
I. Create an Amazon QuickSight dataset that receives spending information from the AWS Price List Query AP
J. Share the dataset with the finance team.
K. Use the AWS Price List Query API to collect account spending informatio
L. Create a specialized template in AWS Cost Explorer that the finance department will use to build reports.

Answer: A

Explanation:
Creating an AWS Cost and Usage Report for the organization and defining tags and cost categories in the report will allow for detailed cost reporting for the different companies that have been consolidated into one organization. By creating a table in Amazon Athena and an Amazon QuickSight dataset based on the Athena table, the finance team will be able to easily query and generate reports on the costs for all the companies. The dataset can then be shared with the finance team for them to use for their reporting needs.

NEW QUESTION 87
- (Exam Topic 1)
A company is creating a sequel for a popular online game. A large number of users from all over the world will play the game within the first week after launch.
Currently, the game consists of the following components deployed in a single AWS Region:
• Amazon S3 bucket that stores game assets
• Amazon DynamoDB table that stores player scores
A solutions architect needs to design a multi-Region solution that will reduce latency improve reliability, and require the least effort to implement
What should the solutions architect do to meet these requirements?

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

Answer: C

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

NEW QUESTION 92
- (Exam Topic 1)
A company is using AWS Organizations lo manage multiple AWS accounts For security purposes, the company requires the creation of an Amazon Simple Notification Service (Amazon SNS) topic that enables integration with a third-party alerting system in all the Organizations member accounts
A solutions architect used an AWS CloudFormation template to create the SNS topic and stack sets to automate the deployment of CloudFormation stacks Trusted access has been enabled in Organizations
What should the solutions architect do to deploy the CloudFormation StackSets in all AWS accounts?

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

Answer: C

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

NEW QUESTION 94
- (Exam Topic 1)
A company is building a software-as-a-service (SaaS) solution on AWS. The company has deployed an Amazon API Gateway REST API with AWS Lambda integration in multiple AWS Regions and in the same production account.
The company offers tiered pricing that gives customers the ability to pay for the capacity to make a certain number of API calls per second. The premium tier offers up to 3,000 calls per second, and customers are identified by a unique API key. Several premium tier customers in various Regions report that they receive error responses of 429 Too Many Requests from multiple API methods during peak usage hours. Logs indicate that the Lambda function is never invoked.
What could be the cause of the error messages for these customers?

A. The Lambda function reached its concurrency limit.
B. The Lambda function its Region limit for concurrency.
C. The company reached its API Gateway account limit for calls per second.
D. The company reached its API Gateway default per-method limit for calls per second.

Answer: C

Explanation:
https://docs.aws.amazon.com/apigateway/latest/developerguide/api-gateway-request-throttling.html#apig-reques

NEW QUESTION 95
- (Exam Topic 1)
A company is running a web application in the AWS Cloud. The application consists of dynamic content that is created on a set of Amazon EC2 instances. The EC2 instances run in an Auto Scaling group that is configured as a target group for an Application Load Balancer (ALB).
The company is using an Amazon CloudFront distribution to distribute the application globally. The CloudFront distribution uses the ALB as an origin. The company uses Amazon Route 53 for DNS and has created an A record of www.example.com for the CloudFront distribution.
A solutions architect must configure the application so that itis highly available and fault tolerant. Which solution meets these requirements?

A. Provision a full, secondary application deployment in a different AWS Regio
B. Update the Route 53 A record to be a failover recor
C. Add both of the CloudFront distributions as value
D. Create Route 53 health checks.
E. Provision an ALB, an Auto Scaling group, and EC2 instances in a different AWS Regio
F. Update the CloudFront distribution, and create a second origin for the new AL

G. Create an origin group for the two origin
H. Configure one origin as primary and one origin as secondary.
I. Provision an Auto Scaling group and EC2 instances in a different AWS Regio
J. Create a second target for the new Auto Scaling group in the AL
K. Set up the failover routing algorithm on the ALB.
L. Provision a full, secondary application deployment in a different AWS Regio
M. Create a second CloudFront distribution, and add the new application setup as an origi
N. Create an AWS Global Accelerator accelerato
O. Add both of the CloudFront distributions as endpoints.

Answer: B

Explanation:
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/DownloadDistS3AndCustomOrigins.h
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/high_availability_origin_failover.html
You can set up CloudFront with origin failover for scenarios that require high availability. To get started, you create an origin group with two origins: a primary and
a secondary. If the primary origin is unavailable, or returns specific HTTP response status codes that indicate a failure, CloudFront automatically switches to the
secondary origin.

NEW QUESTION 100
- (Exam Topic 1)
A company that has multiple AWS accounts is using AWS Organizations. The company's AWS accounts host VPCs, Amazon EC2 instances, and containers.
The company's compliance team has deployed a security tool in each VPC where the company has deployments. The security tools run on EC2 instances and
send information to the AWS account that is dedicated for the compliance team. The company has tagged all the compliance-related resources with a key of
"costCenter" and a value or "compliance".
The company wants to identify the cost of the security tools that are running on the EC2 instances so that the company can charge the compliance team's AWS
account. The cost calculation must be as accurate as possible.
What should a solutions architect do to meet these requirements?

A. In the management account of the organization, activate the costCenter user-defined ta
B. Configure monthly AWS Cost and Usage Reports to save to an Amazon S3 bucket in the management accoun
C. Use the tag breakdown in the report to obtain the total cost for the costCenter tagged resources.
D. In the member accounts of the organization, activate the costCenter user-defined ta
E. Configure monthly AWS Cost and Usage Reports to save to an Amazon S3 bucket in the management accoun
F. Schedule a monthly AWS Lambda function to retrieve the reports and calculate the total cost for the costCenter tagged resources.
G. In the member accounts of the organization activate the costCenter user-defined ta
H. From the management account, schedule a monthly AWS Cost and Usage Repor
I. Use the tag breakdown in the report to calculate the total cost for the costCenter tagged resources.
J. Create a custom report in the organization view in AWS Trusted Adviso
K. Configure the report to generate a monthly billing summary for the costCenter tagged resources in the compliance team's AWS account.

Answer: A

Explanation:
https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/custom-tags.html
https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/configurecostallocreport.html

NEW QUESTION 105
- (Exam Topic 1)
An AWS customer has a web application that runs on premises. The web application fetches data from a third-party API that is behind a firewall. The third party
accepts only one public CIDR block in each client's allow list.
The customer wants to migrate their web application to the AWS Cloud. The application will be hosted on a set of Amazon EC2 instances behind an Application
Load Balancer (ALB) in a VPC. The ALB is located in public subnets. The EC2 instances are located in private subnets. NAT gateways provide internet access to
the private subnets.
How should a solutions architect ensure that the web application can continue to call the third-parly API after the migration?

A. Associate a block of customer-owned public IP addresses to the VP
B. Enable public IP addressing for public subnets in the VPC.
C. Register a block of customer-owned public IP addresses in the AWS accoun
D. Create Elastic IP addresses from the address block and assign them lo the NAT gateways in the VPC.
E. Create Elastic IP addresses from the block of customer-owned IP addresse
F. Assign the static Elastic IP addresses to the ALB.
G. Register a block of customer-owned public IP addresses in the AWS accoun
H. Set up AWS Global Accelerator to use Elastic IP addresses from the address bloc
I. Set the ALB as the accelerator endpoint.

Answer: B

Explanation:
When EC2 instances reach third-party API through internet, their privates IP addresses will be masked by NAT Gateway public IP address.
https://aws.amazon.com/blogs/networking-and-content-delivery/introducing-bring-your-own-ip-byoip-for-amaz

NEW QUESTION 106
- (Exam Topic 1)
A company has a web application that allows users to upload short videos. The videos are stored on Amazon EBS volumes and analyzed by custom recognition
software for categorization.
The website contains stat c content that has variable traffic with peaks in certain months. The architecture consists of Amazon EC2 instances running in an Auto
Scaling group for the web application and EC2
instances running in an Auto Scaling group to process an Amazon SQS queue The company wants to
re-architect the application to reduce operational overhead using AWS managed services where possible and remove dependencies on third-party software.

Which solution meets these requirements?

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

Answer: C

Explanation:

Option C is correct because hosting the web application in Amazon S3, storing the uploaded videos in Amazon S3, and using S3 event notifications to publish events to the SQS queue reduces the operational overhead of managing EC2 instances and EBS volumes. Amazon S3 can serve static content such as HTML, CSS, JavaScript, and media files directly from S3 buckets. Amazon S3 can also trigger AWS Lambda functions through S3 event notifications when new objects are created or existing objects are updated or deleted. AWS Lambda can process the SQS queue with an AWS Lambda function that calls the Amazon Rekognition API to categorize the videos. This solution eliminates the need for custom recognition software and third-party dependencies[345]
References: 1: https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-spot-instances.html 2: https://aws.amazon.com/efs/pricing/ 3: https://docs.aws.amazon.com/AmazonS3/latest/userguide/WebsiteHosting.html 4: https://docs.aws.amazon.com/AmazonS3/latest/userguide/NotificationHowTo.html 5: https://docs.aws.amazon.com/rekognition/latest/dg/what-is.html 6: https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/Welcome.html

NEW QUESTION 108
- (Exam Topic 1)
A company is processing videos in the AWS Cloud by using Amazon EC2 instances in an Auto Scaling group. It takes 30 minutes to process a video. Several EC2 instances scale in and out depending on the number of videos in an Amazon Simple Queue Service (Amazon SQS) queue.
The company has configured the SQS queue with a redrive policy that specifies a target dead-letter queue and a maxReceiveCount of 1. The company has set the visibility timeout for the SQS queue to 1 hour. The company has set up an Amazon CloudWatch alarm to notify the development team when there are messages in the dead-letter queue.
Several times during the day, the development team receives notification that messages are in the dead-letter queue and that videos have not been processed properly. An investigation finds no errors in the application logs.
How can the company solve this problem?

A. Turn on termination protection for the EC2 instances.
B. Update the visibility timeout for the SOS queue to 3 hours.
C. Configure scale-in protection for the instances during processing.
D. Update the redrive policy and set maxReceiveCount to 0.

Answer: B

Explanation:
The best solution for this problem is to update the visibility timeout for the SQS queue to 3 hours. This is because when the visibility timeout is set to 1 hour, it means that if the EC2 instance doesn't process the message within an hour, it will be moved to the dead-letter queue. By increasing the visibility timeout to 3 hours, this should give the EC2 instance enough time to process the message before it gets moved to the dead-letter queue. Additionally, configuring scale-in protection for the EC2 instances during processing will help to ensure that the instances are not terminated while the messages are being processed.

NEW QUESTION 111
- (Exam Topic 1)
A solutions architect needs to implement a client-side encryption mechanism for objects that will be stored in a new Amazon S3 bucket. The solutions architect created a CMK that is stored in AWS Key Management Service (AWS KMS) for this purpose.
The solutions architect created the following IAM policy and attached it to an IAM role:

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "DownloadUpload",
            "Action": [
                "s3:GetObject",
                "s3:GetObjectVersion",
                "s3:PutObject",
                "s3:PutObjectAcl"
            ],
            "Effect": "Allow",
            "Resource": "arn:aws:s3:::BucketName/*"
        },
        {
            "Sid": "KMSAccess",
            "Action": [
                "kms:Decrypt",
                "kms:Encrypt"
            ],
            "Effect": "Allow",
            "Resource": "arn:aws:kms:region:Account:key/Key ID"
        }
    ]
}
```

During tests, me solutions architect was able to successfully get existing test objects m the S3 bucket However, attempts to upload a new object resulted in an error message. The error message stated that me action was forbidden.
Which action must me solutions architect add to the IAM policy to meet all the requirements?

A. Kms:GenerateDataKey
B. KmsGetKeyPolpcy
C. kmsGetPubKKey
D. kms:SKjn

Answer: A

Explanation:
https://aws.amazon.com/premiumsupport/knowledge-center/s3-access-denied-error-kms/
"An error occurred (AccessDenied) when calling the PutObject operation: Access Denied" This error message indicates that your IAM user or role needs permission for the kms:GenerateDataKey action.

NEW QUESTION 113
- (Exam Topic 1)
A company with several AWS accounts is using AWS Organizations and service control policies (SCPs). An Administrator created the following SCP and has attached it to an organizational unit (OU) that contains AWS account 1111-1111-1111:

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "AllowsAllActions",
            "Effect": "Allow",
            "Action": "*",
            "Resource": "*"
        },
        {
            "Sid": "DenyCloudTrail",
            "Effect": "Deny",
            "Action": "cloudtrail:*",
            "Resource": "*"
        }
    ]
}
```

Developers working in account 1111-1111-1111 complain that they cannot create Amazon S3 buckets. How should the Administrator address this problem?

A. Add s3:CreateBucket with €Allow€ effect to the SCP.
B. Remove the account from the OU, and attach the SCP directly to account 1111-1111-1111.
C. Instruct the Developers to add Amazon S3 permissions to their IAM entities.
D. Remove the SCP from account 1111-1111-1111.

Answer: C

Explanation:
However A's explanation is incorrect - https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps.html
"SCPs are similar to AWS Identity and Access Management (IAM) permission policies and use almost the
same syntax. However, an SCP never grants permissions."
SCPs alone are not sufficient to granting permissions to the accounts in your organization. No permissions are granted by an SCP. An SCP defines a guardrail, or sets limits, on the actions that the account's administrator can delegate to the IAM users and roles in the affected accounts. The administrator must still attach identity-based or resource-based policies to IAM users or roles, or to the resources in your accounts to actually grant permissions. The effective permissions are the logical intersection between what is allowed by the SCP and what is allowed by the IAM and resource-based policies.

NEW QUESTION 117
- (Exam Topic 1)
A company's solutions architect is reviewing a web application that runs on AWS. The application references static assets in an Amazon S3 bucket in the us-east-1 Region. The company needs resiliency across multiple AWS Regions. The company already has created an S3 bucket in a second Region.
Which solution will meet these requirements with the LEAST operational overhead?

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

Answer: C

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

NEW QUESTION 119

- (Exam Topic 1)

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

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

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

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

Answer: AD

Explanation:

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

NEW QUESTION 120

- (Exam Topic 1)

A large company is running a popular web application. The application runs on several Amazon EC2 Linux Instances in an Auto Scaling group in a private subnet. An Application Load Balancer is targeting the Instances In the Auto Scaling group in the private subnet. AWS Systems Manager Session Manager Is configured, and AWS Systems Manager Agent is running on all the EC2 instances.

The company recently released a new version of the application Some EC2 instances are now being marked as unhealthy and are being terminated As a result, the application is running at reduced capacity A solutions architect tries to determine the root cause by analyzing Amazon CloudWatch logs that are collected from the application, but the logs are inconclusive

How should the solutions architect gain access to an EC2 instance to troubleshoot the issue1?

A. Suspend the Auto Scaling group's HealthCheck scaling proces
B. Use Session Manager to log in to an instance that is marked as unhealthy
C. Enable EC2 instance termination protection Use Session Manager to log In to an instance that is marked as unhealthy.
D. Set the termination policy to Oldestinstance on the Auto Scaling grou
E. Use Session Manager to log in to an instance that is marked as unhealthy
F. Suspend the Auto Scaling group's Terminate proces
G. Use Session Manager to log in to an instance thatis marked as unhealthy

Answer: D

Explanation:

https://docs.aws.amazon.com/autoscaling/ec2/userguide/as-suspend-resume-processes.html

NEW QUESTION 124

- (Exam Topic 1)

A company has introduced a new policy that allows employees to work remotely from their homes if they connect by using a VPN The company Is hosting Internal applications with VPCs in multiple AWS accounts Currently the applications are accessible from the company's on-premises office network through an AWS Site-to-Site VPN connection The VPC in the company's main AWS account has peering connections established with VPCs in other AWS accounts.

A solutions architect must design a scalable AWS Client VPN solution for employees to use while they work from home

What is the MOST cost-effective solution that meets these requirements?

A. Create a Client VPN endpoint in each AWS account Configure required routing that allows access to internal applications
B. Create a Client VPN endpoint in the mam AWS account Configure required routing that allows access to internal applications
C. Create a Client VPN endpoint in the main AWS account Provision a transit gateway that is connected to each AWS account Configure required routing that allows access to internal applications
D. Create a Client VPN endpoint in the mam AWS account Establish connectivity between the Client VPN endpoint and the AWS Site-to-Site VPN

Answer: C

Explanation:

https://docs.aws.amazon.com/vpn/latest/clientvpn-admin/scenario-peered.html

NEW QUESTION 126

- (Exam Topic 1)

A company is building an electronic document management system in which users upload their documents. The application stack is entirely serverless and runs on AWS in the eu-central-1 Region. The system includes a web application that uses an Amazon CloudFront distribution for delivery with Amazon S3 as the origin. The web application communicates with Amazon API Gateway Regional endpoints. The API Gateway APIs call AWS Lambda functions that store metadata in an Amazon Aurora Serverless database and put the documents into an S3 bucket.

The company is growing steadily and has completed a proof of concept with its largest customer. The company must improve latency outside of Europe.

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

A. Enable S3 Transfer Acceleration on the S3 bucke
B. Ensure that the web application uses the Transfer Acceleration signed URLs.
C. Create an accelerator in AWS Global Accelerato
D. Attach the accelerator to the CloudFront distribution.
E. Change the API Gateway Regional endpoints to edge-optimized endpoints.

F. Provision the entire stack in two other locations that are spread across the worl
G. Use global databases on the Aurora Serverless cluster.
H. Add an Amazon RDS proxy between the Lambda functions and the Aurora Serverless database.

Answer: AC

Explanation:
https://aws.amazon.com/global-accelerator/faqs/

NEW QUESTION 129
- (Exam Topic 1)
A solutions architect is designing the data storage and retrieval architecture for a new application that a company will be launching soon. The application is designed to ingest millions of small records per minute from devices all around the world. Each record is less than 4 KB in size and needs to be stored in a durable location where it can be retrieved with low latency. The data is ephemeral and the company is required to store the data for 120 days only, after which the data can be deleted.
The solutions architect calculates that, during the course of a year, the storage requirements would be about 10-15 TB.
Which storage strategy is the MOST cost-effective and meets the design requirements?

A. Design the application to store each incoming record as a single .csv file in an Amazon S3 bucket to allow for indexed retrieva
B. Configure a lifecycle policy to delete data older than 120 days.
C. Design the application to store each incoming record in an Amazon DynamoDB table properly configured for the scal
D. Configure the DynamoOB Time to Live (TTL) feature to delete records older than 120 days.
E. Design the application to store each incoming record in a single table in an Amazon RDS MySQL databas
F. Run a nightly cron job that executes a query to delete any records older than 120 days.
G. Design the application to batch incoming records before writing them to an Amazon S3 bucke
H. Updatethe metadata for the object to contain the list of records in the batch and use the Amazon S3 metadata search feature to retrieve the dat
I. Configure a lifecycle policy to delete the data after 120 days.

Answer: B

Explanation:
DynamoDB with TTL, cheaper for sustained throughput of small items + suited for fast retrievals. S3 cheaper for storage only, much higher costs with writes. RDS not designed for this use case.

NEW QUESTION 131
- (Exam Topic 1)
A company is running several workloads in a single AWS account. A new company policy states that engineers can provision only approved resources and that engineers must use AWS CloudFormation to provision these resources. A solutions architect needs to create a solution to enforce the new restriction on the IAM role that the engineers use for access.
What should the solutions architect do to create the solution?

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

Answer: B

Explanation:
https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/security-best-practices.html#use-iam-to-c
https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/using-iam-servicerole.html

NEW QUESTION 132
- (Exam Topic 1)
A company is planning to store a large number of archived documents and make the documents available to employees through the corporate intranet. Employees will access the system by connecting through a client VPN service that is attached to a VPC. The data must not be accessible to the public.
The documents that the company is storing are copies of data that is held on physical media elsewhere. The number of requests will be low. Availability and speed of retrieval are not concerns of the company.
Which solution will meet these requirements at the LOWEST cost?

A. Create an Amazon S3 bucke
B. Configure the S3 bucket to use the S3 One Zone-Infrequent Access (S3 One Zone-IA) storage class as defaul
C. Configure the S3 bucket for website hostin
D. Create an S3 interface endpoin
E. Configure the S3 bucket to allow access only through that endpoint.
F. Launch an Amazon EC2 instance that runs a web serve
G. Attach an Amazon Elastic File System (Amazon EFS) file system to store the archived data in the EFS One Zone-Infrequent Access (EFS One Zone-IA) storage class Configure the instance security groups to allow access only from private networks.
H. Launch an Amazon EC2 instance that runs a web server Attach an Amazon Elastic Block Store (Amazon EBS) volume to store the archived dat
I. Use the Cold HDD (sc1) volume typ
J. Configure the instance security groups to allow access only from private networks.
K. Create an Amazon S3 bucke
L. Configure the S3 bucket to use the S3 Glacier Deep Archive storage class as defaul
M. Configure the S3 bucket for website hostin

N. Create an S3 interface endpoin
O. Configure the S3 bucket to allow access only through that endpoint.

Answer: D

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

NEW QUESTION 136
- (Exam Topic 1)
A company has registered 10 new domain names. The company uses the domains for online marketing. The company needs a solution that will redirect online visitors to a specific URL for each domain. All domains and target URLs are defined in a JSON document. All DNS records are managed by Amazon Route 53.
A solutions architect must implement a redirect service that accepts HTTP and HTTPS requests.
Which combination of steps should the solutions architect take to meet these requirements with the LEAST amount of operational effort? (Choose three.)

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

Answer: CEF

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

NEW QUESTION 141
- (Exam Topic 1)
The company needs to determine which costs on the monthly AWS bill are attributable to each application or team. The company also must be able to create reports to compare costs from the last 12 months and to help forecast costs for the next 12 months. A solutions architect must recommend an AWS Billing and Cost Management solution that provides these cost reports.
Which combination of actions will meet these requirements? (Select THREE.)

A. Activate the user-defined cost allocation tags that represent the application and the team.
B. Activate the AWS generated cost allocation tags that represent the application and the team.
C. Create a cost category for each application in Billing and Cost Management.
D. Activate IAM access to Billing and Cost Management.
E. Create a cost budget.
F. Enable Cost Explorer.

Answer: ACF

Explanation:
https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/manage-cost-categories.html https://aws.amazon.com/premiumsupport/knowledge-center/cost-explorer-analyze-spending-and-usage/ https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/manage-cost-categories.html
https://docs.aws.amazon.com/cost-management/latest/userguide/ce-enable.html
The best combination of actions to meet the company's requirements is Options A, C, and F.
Option A involves activating the user-defined cost allocation tags that represent the application and the team. This will allow the company to assign costs to different applications or teams, and will allow them to be tracked in the monthly AWS bill.
Option C involves creating a cost category for each application in Billing and Cost Management. This will allow the company to easily identify and compare costs across different applications and teams.
Option F involves enabling Cost Explorer. This will allow the company to view the costs of their AWS resources over the last 12 months and to create forecasts for the next 12 months.
These recommendations are in line with the official Amazon Textbook and Resources for the AWS Certified Solutions Architect - Professional certification. In particular, the book states that "You can use cost allocation tags to group your costs by application, team, or other categories" (Source:
https://d1.awsstatic.com/training-and-certification/docs-sa-pro/AWS_Certified_Solutions_Architect_Professiona Additionally, the book states that "Cost Explorer enables you to view the costs of your AWS resources over the last 12 months and to create forecasts for the next 12 months" (Source:
https://d1.awsstatic.com/training-and-certification/docs-sa-pro/AWS_Certified_Solutions_Architect_Professiona

NEW QUESTION 146
- (Exam Topic 1)
A company is building a serverless application that runs on an AWS Lambda function that is attached to a VPC. The company needs to integrate the application with a new service from an external provider. The external provider supports only requests that come from public IPv4 addresses that are in an allow list.
The company must provide a single public IP address to the external provider before the application can start using the new service.
Which solution will give the application the ability to access the new service?

A. Deploy a NAT gatewa
B. Associate an Elastic IP address with the NAT gatewa
C. Configure the VPC to use the NAT gateway.
D. Deploy an egress-only internet gatewa
E. Associate an Elastic IP address with the egress-only internet gatewa
F. Configure the elastic network interface on the Lambda function to use the egress-only internet gateway.

G. Deploy an internet gatewa
H. Associate an Elastic IP address with the internet gatewa
I. Configure theLambda function to use the internet gateway.
J. Deploy an internet gatewa
K. Associate an Elastic IP address with the internet gatewa
L. Configure the default route in the public VPC route table to use the internet gateway.

Answer: A

Explanation:
This solution will give the Lambda function access to the internet by routing its outbound traffic through the NAT gateway, which has a public Elastic IP address.
This will allow the external provider to whitelist the single public IP address associated with the NAT gateway, and enable the application to access the new service
Deploying a NAT gateway and associating an Elastic IP address with it, and then configuring the VPC to use the NAT gateway, will give the application the ability to access the new service. This is because the NAT gateway will be the single public IP address that the external provider needs for the allow list. The NAT gateway will allow the application to access the service, while keeping the underlying Lambda functions private.
When configuring NAT gateways, you should ensure that the route table associated with the NAT gateway has a route to the internet gateway with a target of the internet gateway. Additionally, you should ensure that the security group associated with the NAT gateway allows outbound traffic from the Lambda functions.
References:

⟩ AWS Certified Solutions Architect Professional Official Amazon Text Book [1], page 456
https://docs.aws.amazon.com/vpc/latest/userguide/VPC_NAT_Gateway.html

NEW QUESTION 147
- (Exam Topic 1)
A company has hundreds of AWS accounts. The company recently implemented a centralized internal process for purchasing new Reserved Instances and modifying existing Reserved Instances. This process requires all business units that want to purchase or modify Reserved Instances to submit requests to a dedicated team for procurement. Previously, business units directly purchased or modified Reserved Instances in their own respective AWS accounts autonomously.
A solutions architect needs to enforce the new process in the most secure way possible.
Which combination of steps should the solutions architect take to meet these requirements? (Choose two.)

A. Ensure that all AWS accounts are part of an organization in AWS Organizations with all features enabled.
B. Use AWS Config to report on the attachment of an IAM policy that denies access to the ec2:PurchaseReservedInstancesOffering action and the ec2:ModifyReservedInstances action.
C. In each AWS account, create an IAM policy that denies the ec2:PurchaseReservedInstancesOffering action and the ec2:ModifyReservedInstances action.
D. Create an SCP that denies the ec2:PurchaseReservedInstancesOffering action and the ec2:ModifyReservedInstances actio
E. Attach the SCP to each OU of the organization.
F. Ensure that all AWS accounts are part of an organization in AWS Organizations that uses the consolidated billing feature.

Answer: AD

Explanation:
All features – The default feature set that is available to AWS Organizations. It includes all the functionality of consolidated billing, plus advanced features that give you more control over accounts in your organization. For example, when all features are enabled the management account of the organization has full control over what member accounts can do. The management account can apply SCPs to restrict the services and actions that users (including the root user) and roles in an account can access. https://docs.aws.amazon.com/organizations/latest/userguide/orgs_getting-started_concepts.html#feature-set

NEW QUESTION 151
- (Exam Topic 1)
A company has an environment that has a single AWS account. A solutions architect is reviewing the environment to recommend what the company could improve specifically in terms of access to the AWS Management Console. The company's IT support workers currently access the console for administrative tasks, authenticating with named IAM users that have been mapped to their job role.
The IT support workers no longer want to maintain both their Active Directory and IAM user accounts. They want to be able to access the console by using their existing Active Directory credentials. The solutions architect is using AWS Single Sign-On (AWS SSO) to implement this functionality.
Which solution will meet these requirements MOST cost-effectively?

A. Create an organization in AWS Organization
B. Turn on the AWS SSO feature in Organizations Create and configure a directory in AWS Directory Service for Microsoft Active Directory (AWS Managed Microsoft AD) with a two-way trust to the company's on-premises Active Director
C. Configure AWS SSO and set the AWS Managed Microsoft AD directory as the identity sourc
D. Create permission sets and map them to the existing groups within the AWS Managed Microsoft AD directory.
E. Create an organization in AWS Organization
F. Turn on the AWS SSO feature in Organizations Create and configure an AD Connector to connect to the company's on-premises Active Director
G. Configure AWS SSO and select the AD Connector as the identity sourc
H. Create permission sets and map them to the existing groups within the company's Active Directory.
I. Create an organization in AWS Organization
J. Turn on all features for the organizatio
K. Create and configure a directory in AWS Directory Service for Microsoft Active Directory (AWS Managed Microsoft AD) with a two-way trust to the company's on-premises Active Director
L. Configure AWS SSO and select the AWS Managed Microsoft AD directory as the identity sourc
M. Create permission sets and map them to the existing groups within the AWS Managed Microsoft AD directory.
N. Create an organization in AWS Organization
O. Turn on all features for the organizatio
P. Create and configure an AD Connector to connect to the company's on-premises Active Director
Q. Configure AWS SSO and select the AD Connector as the identity sourc
R. Create permission sets and map them to the existing groups within the company's Active Directory.

Answer: D

Explanation:
https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_org_support-all-features.html

https://docs.aws.amazon.com/singlesignon/latest/userguide/get-started-prereqs-considerations.html

NEW QUESTION 154
- (Exam Topic 1)
A company is storing data on premises on a Windows file server. The company produces 5 GB of new data daily.
The company migrated part of its Windows-based workload to AWS and needs the data to be available on a file system in the cloud. The company already has established an AWS Direct Connect connection between the on-premises network and AWS.
Which data migration strategy should the company use?

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

Answer: B

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

NEW QUESTION 159
......

Thank You for Trying Our Product

We offer two products:

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

2nd - Questons and Answers in PDF Format

SAP-C02 Practice Exam Features:

* SAP-C02 Questions and Answers Updated Frequently

* SAP-C02 Practice Questions Verified by Expert Senior Certified Staff

* SAP-C02 Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* SAP-C02 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The SAP-C02 Practice Test Here](https://www.certshared.com/exam/SAP-C02/)