Amazon-Web-Services

Exam Questions SAP-C02

AWS Certified Solutions Architect - Professional

NEW QUESTION 1
- (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 2
- (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 3
- (Exam Topic 1)
A company is developing and hosting several projects in the AWS Cloud. The projects are developed across multiple AWS accounts under the same organization in AWS Organizations. The company requires the cost lor cloud infrastructure to be allocated to the owning project. The team responsible for all of the AWS accounts has discovered that several Amazon EC2 instances are lacking the Project tag used for cost allocation.
Which actions should a solutions architect take to resolve the problem and prevent it from happening in the future? (Select THREE.)

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

Answer: ABE

Explanation:
https://docs.aws.amazon.com/config/latest/developerguide/config-rule-multi-account-deployment.html
https://docs.aws.amazon.com/config/latest/developerguide/aggregate-data.html
https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps_examples_tagging.htm

NEW QUESTION 4
- (Exam Topic 1)
A company is building a solution in the AWS Cloud. Thousands or devices will connect to the solution and send data. Each device needs to be able to send and receive data in real time over the MQTT protocol. Each device must authenticate by using a unique X.509 certificate.
Which solution will meet these requirements with the LEAST operational overhead?

A. Set up AWS loT Cor
B. For each device, create a corresponding Amazon MQ queue and provision a certificat
C. Connect each device to Amazon MQ.
D. Create a Network Load Balancer (NLB) and configure it with an AWS Lambda authorize
E. Run an MQTT broker on Amazon EC2 instances in an Auto Scaling grou
F. Set the Auto Scaling group as the target for the NL
G. Connect each device to the NLB.
H. Set up AWS loT Cor
I. For each device, create a corresponding AWS loT thing and provision a certificat
J. Connect each device to AWS loT Core.
K. Set up an Amazon API Gateway HTTP API and a Network Load Balancer (NLB). Create integration between API Gateway and the NL
L. Configure a mutual TLS certificate authorizer on the HTTP AP
M. Run an MQTT broker on an Amazon EC2 instance that the NLB target
N. Connect each device to the NLB.

Answer: D

Explanation:
This solution requires minimal operational overhead, as it only requires setting up AWS IoT Core and creating a thing for each device. (Reference: AWS Certified Solutions Architect - Professional Official Amazon Text Book, Page 537)
AWS IoT Core is a fully managed service that enables secure, bi-directional communication between internet-connected devices and the AWS Cloud. It supports the MQTT protocol and includes built-in device
authentication and access control. By using AWS IoT Core, the company can easily provision and manage the X.509 certificates for each device, and connect the devices to the service with minimal operational overhead.

NEW QUESTION 5
- (Exam Topic 1)
A company is running a traditional web application on Amazon EC2 instances. The company needs to refactor the application as microservices that run on containers. Separate versions of the application exist in two distinct environments: production and testing. Load for the application is variable, but the minimum load and the maximum load are known. A solutions architect needs to design the updated application with a serverless architecture that minimizes operational complexity.
Which solution will meet these requirements MOST cost-effectively?

A. Upload the container images to AWS Lambda as function
B. Configure a concurrency limit for the associated Lambda functions to handle the expected peak loa
C. Configure two separate Lambda integrations within Amazon API Gateway: one for production and one for testing.
D. Upload the container images to Amazon Elastic Container Registry (Amazon ECR). Configure two auto scaled Amazon Elastic Container Service (Amazon ECS) clusters with the Fargate launch type to handle the expected loa
E. Deploy tasks from the ECR image
F. Configure two separate Application Load Balancers to direct traffic to the ECS clusters.
G. Upload the container images to Amazon Elastic Container Registry (Amazon ECR). Configure two auto scaled Amazon Elastic Kubernetes Service (Amazon EKS) clusters with the Fargate launch type to handle the expected loa
H. Deploy tasks from the ECR image
I. Configure two separate Application Load Balancers to direct traffic to the EKS clusters.
J. Upload the container images to AWS Elastic Beanstal
K. In Elastic Beanstalk, create separate environments and deployments for production and testin
L. Configure two separate Application Load Balancers to direct traffic to the Elastic Beanstalk deployments.

Answer: B

Explanation:
minimizes operational + microservices that run on containers = AWS Elastic Beanstalk

NEW QUESTION 6
- (Exam Topic 1)
An AWS partner company is building a service in AWS Organizations using Its organization named org. This service requires the partner company to have access to AWS resources in a customer account, which is in a separate organization named org2 The company must establish least privilege security access using an API or command line tool to the customer account
What is the MOST secure way to allow org1 to access resources h org2?

A. The customer should provide the partner company with their AWS account access keys to log in and perform the required tasks
B. The customer should create an IAM user and assign the required permissions to the IAM user The customer should then provide the credentials to the partner company to log In and perform the required tasks.
C. The customer should create an IAM role and assign the required permissions to the IAM rol
D. The partner company should then use the IAM rote's Amazon Resource Name (ARN) when requesting access to perform the required tasks
E. The customer should create an IAM rote and assign the required permissions to the IAM rot
F. The partner company should then use the IAM rote's Amazon Resource Name (ARN). Including the external ID in the IAM role's trust pokey, when requesting access to perform the required tasks

Answer: C

Explanation:
https://docs.aws.amazon.com/IAM/latest/UserGuide/confused-deputy.html
This is the most secure way to allow org1 to access resources in org2 because it allows for least privilege security access. The customer should create an IAM role

and assign the required permissions to the IAM role. The partner company should then use the IAM role's Amazon Resource Name (ARN) and include the external ID in the IAM role's trust policy when requesting access to perform the required tasks. This ensures that the partner company can only access the resources that it needs and only from the specific customer account.

NEW QUESTION 7
- (Exam Topic 1)
A company is using an on-premises Active Directory service for user authentication. The company wants to use the same authentication service to sign in to the company's AWS accounts, which are using AWS Organizations. AWS Site-to-Site VPN connectivity already exists between the on-premises environment and all the company's AWS accounts.
The company's security policy requires conditional access to the accounts based on user groups and roles. User identities must be managed in a single location. Which solution will meet these requirements?

A. Configure AWS Single Sign-On (AWS SSO) to connect to Active Directory by using SAML 2.0.Enable automatic provisioning by using the System for Cross-domain Identity Management (SCIM) v2.0 protoco
B. Grant access to the AWS accounts by using attribute-based access controls (ABACs).
C. Configure AWS Single Sign-On (AWS SSO) by using AWS SSO as an identity sourc
D. Enable automatic provisioning by using the System for Cross-domain Identity Management (SCIM) v2.0 protoco
E. Grant access to the AWS accounts by using AWS SSO permission sets.
F. In one of the company's AWS accounts, configure AWS Identity and Access Management (IAM) to use a SAML 2.0 identity provide
G. Provision IAM users that are mapped to the federated user
H. Grant access that corresponds to appropriate groups in Active Director
I. Grant access to the required AWS accounts by using cross-account IAM users.
J. In one of the company's AWS accounts, configure AWS Identity and Access Management (IAM) to use an OpenID Connect (OIDC) identity provide
K. Provision IAM roles that grant access to the AWS account for the federated users that correspond to appropriate groups in Active Director
L. Grant access to the required AWS accounts by using cross-account IAM roles.

Answer: D

Explanation:
https://aws.amazon.com/blogs/aws/new-attributes-based-access-control-with-aws-single-sign-on/

NEW QUESTION 8
- (Exam Topic 1)
A company has a latency-sensitive trading platform that uses Amazon DynamoDB as a storage backend. The company configured the DynamoDB table to use on-demand capacity mode. A solutions architect needs to design a solution to improve the performance of the trading platform. The new solution must ensure high availability for the trading platform.
Which solution will meet these requirements with the LEAST latency?

A. Create a two-node DynamoDB Accelerator (DAX) cluster Configure an application to read and write data by using DAX.
B. Create a three-node DynamoDB Accelerator (DAX) cluste
C. Configure an application to read data by using DAX and to write data directly to the DynamoDB table.
D. Create a three-node DynamoDB Accelerator (DAX) cluste
E. Configure an application to read data directly from the DynamoDB table and to write data by using DAX.
F. Create a single-node DynamoD8 Accelerator (DAX) cluste
G. Configure an application to read data by using DAX and to write data directly to the DynamoD8 table.

Answer: B

Explanation:
A DAX cluster can be deployed with one or two nodes for development or test workloads. One- and two-node clusters are not fault-tolerant, and we don't recommend using fewer than three nodes for production use. If a one- or two-node cluster encounters software or hardware errors, the cluster can become unavailable or lose cached data.A DAX cluster can be deployed with one or two nodes for development or test workloads. One and two-node clusters are not fault-tolerant, and we don't recommend using fewer than three nodes for production use. If a one- or two-node cluster encounters software or hardware errors, the cluster can become unavailable or lose cached data.
https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/DAX.concepts.cluster.html

NEW QUESTION 9
- (Exam Topic 1)
A company runs a Java application that has complex dependencies on VMs that are in the company's data center. The application is stable. but the company wants to modernize the technology stack. The company wants to migrate the application to AWS and minimize the administrative overhead to maintain the servers.
Which solution will meet these requirements with the LEAST code changes?

A. Migrate the application to Amazon Elastic Container Service (Amazon ECS) on AWS Fargate by using AWS App2Containe
B. Store container images in Amazon Elastic Container Registry (Amazon ECR). Grant the ECS task execution role permission 10 access the ECR image repositor
C. Configure Amazon ECS to use an Application Load Balancer (ALB). Use the ALB to interact with the application.
D. Migrate the application code to a container that runs in AWS Lambd
E. Build an Amazon API Gateway REST API with Lambda integratio
F. Use API Gateway to interact with the application.
G. Migrate the application to Amazon Elastic Kubernetes Service (Amazon EKS) on EKS managed node groups by using AWS App2Containe
H. Store container images in Amazon Elastic Container Registry (Amazon ECR). Give the EKS nodes permission to access the ECR image repositor
I. Use Amazon API Gateway to interact with the application.
J. Migrate the application code to a container that runs in AWS Lambd
K. Configure Lambda to use an Application Load Balancer (ALB). Use the ALB to interact with the application.

Answer: A

Explanation:
According to the AWS documentation1, AWS App2Container (A2C) is a command line tool for migrating and modernizing Java and .NET web applications into container format. AWS A2C analyzes and builds an inventory of applications running in bare metal, virtual machines, Amazon Elastic Compute Cloud (EC2)

instances, or in the cloud. You can use AWS A2C to generate container images for your applications and deploy them on Amazon ECS or Amazon EKS. Option A meets the requirements of the scenario because it allows you to migrate your existing Java application to AWS and minimize the administrative overhead to maintain the servers. You can use AWS A2C to analyze your application dependencies, extract application artifacts, and generate a Dockerfile. You can then store your container images in Amazon ECR, which is a fully managed container registry service. You can use AWS Fargate as the launch type for your Amazon ECS cluster, which is a serverless compute engine that eliminates the need to provision and manage servers for your containers. You can grant the ECS task execution role permission to access the ECR image repository, which allows your tasks to pull images from ECR. You can configure Amazon ECS to use an ALB, which is a load balancecer that distributes traffic across multiple targets in multiple Availability Zones using HTTP or HTTPS protocols. You can use the ALB to interact with your application.

NEW QUESTION 10
- (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 10
- (Exam Topic 1)
A company has applications in an AWS account that is named Source. The account is in an organization in AWS Organizations. One of the applications uses AWS Lambda functions and store's inventory data in an Amazon Aurora database. The application deploys the Lambda functions by using a deployment package. The company has configured automated backups for Aurora.
The company wants to migrate the Lambda functions and the Aurora database to a new AWS account that is named Target. The application processes critical data, so the company must minimize downtime.
Which solution will meet these requirements?

A. Download the Lambda function deployment package from the Source accoun
B. Use the deployment package and create new Lambda functions in the Target accoun
C. Share the automated Aurora DB cluster snapshot with the Target account.
D. Download the Lambda function deployment package from the Source accoun
E. Use the deployment package and create new Lambda functions in the Target account Share the Aurora DB cluster with the Target account by using AWS Resource Access Manager {AWS RAM). Grant the Target account permission to clone the Aurora DB cluster.
F. Use AWS Resource Access Manager (AWS RAM) to share the Lambda functions and the Aurora DB cluster with the Target accoun
G. Grant the Target account permission to clone the Aurora DB cluster.
H. Use AWS Resource Access Manager (AWS RAM) to share the Lambda functions with the Target accoun
I. Share the automated Aurora DB cluster snapshot with the Target account.

Answer: C

Explanation:
This solution uses a combination of AWS Resource Access Manager (RAM) and automated backups to migrate the Lambda functions and the Aurora database to the Target account while minimizing downtime. In this solution, the Lambda function deployment package is downloaded from the Source account and used to create new Lambda functions in the Target account. The Aurora DB cluster is shared with the Target account using AWS RAM and the Target account is granted permission to clone the Aurora DB cluster, allowing for a new copy of the Aurora database to be created in the Target account. This approach allows for the data to be migrated to the Target account while minimizing downtime, as the Target account can use the cloned Aurora database while the original Aurora database continues to be used in the Source account.

NEW QUESTION 11
- (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 15
- (Exam Topic 1)
A company with global offices has a single 1 Gbps AWS Direct Connect connection to a single AWS Region. The company's on-premises network uses the connection to communicate with the company's resources in the AWS Cloud. The connection has a single private virtual interface that connects to a single VPC.
A solutions architect must implement a solution that adds a redundant Direct Connect connection in the same Region. The solution also must provide connectivity to other Regions through the same pair of Direct Connect connections as the company expands into other Regions.
Which solution meets these requirements?

A. Provision a Direct Connect gatewa
B. Delete the existing private virtual interface from the existing connectio
C. Create the second Direct Connect connectio
D. Create a new private virtual interlace on each connection, and connect both private victual interfaces to the Direct Connect gatewa
E. Connect the Direct Connect gateway to the single VPC.
F. Keep the existing private virtual interfac
G. Create the second Direct Connect connectio
H. Create a new private virtual interface on the new connection, and connect the new private virtual interface to the single VPC.
I. Keep the existing private virtual interfac
J. Create the second Direct Connect connectio
K. Create a new public virtual interface on the new connection, and connect the new public virtual interface to the single VPC.
L. Provision a transit gatewa
M. Delete the existing private virtual interface from the existing connection.Create the second Direct Connect connectio
N. Create a new private virtual interface on each connection, and connect both private virtual interfaces to the transit gatewa
O. Associate the transit gateway with the single VPC.

Answer: A

Explanation:
A Direct Connect gateway is a globally available resource. You can create the Direct Connect gateway in any Region and access it from all other Regions. The following describe scenarios where you can use a Direct Connect gateway.
https://docs.aws.amazon.com/directconnect/latest/UserGuide/direct-connect-gateways-intro.html

NEW QUESTION 19
- (Exam Topic 1)
A company built an application based on AWS Lambda deployed in an AWS CloudFormation stack. The last production release of the web application introduced an issue that resulted in an outage lasting several minutes. A solutions architect must adjust the deployment process to support a canary release.
Which solution will meet these requirements?

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

Answer: A

Explanation:
https://aws.amazon.com/blogs/compute/implementing-canary-deployments-of-aws-lambda-functions-with-alias-
https://docs.aws.amazon.com/lambda/latest/dg/configuration-aliases.html

NEW QUESTION 21
- (Exam Topic 1)
A retail company has an on-premises data center in Europe. The company also has a multi-Region AWS presence that includes the eu-west-1 and us-east-1 Regions. The company wants to be able to route network traffic from its on-premises infrastructure into VPCs in either of those Regions. The company also needs to support traffic that is routed directly between VPCs in those Regions. No single points of failure can exist on the network.
The company already has created two 1 Gbps AWS Direct Connect connections from its on-premises data center. Each connection goes into a separate Direct Connect location in Europe for high availability. These two locations are named DX-A and DX-B, respectively. Each Region has a single AWS Transit Gateway that is configured to route all inter-VPC traffic within that Region.

Which solution will meet these requirements?

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

Answer: D

Explanation:
in this solution, two transit VIFs are created - one from the DX-A connection and one from the DX-B connection - into the same Direct Connect gateway for high availability. Both the eu-west-1 and us-east-1 transit gateways are then associated with this Direct Connect gateway. The transit gateways are then peered with each other to support cross-Region routing. This solution meets the requirements of the company by creating a highly available connection between the on-premises data center and the VPCs in both the eu-west-1 and us-east-1 regions, and by enabling direct traffic routing between VPCs in those regions.

NEW QUESTION 24
- (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 25
- (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 26
- (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 27
- (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 30
- (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 31
- (Exam Topic 1)
A company has a legacy monolithic application that is critical to the company's business. The company hosts the application on an Amazon EC2 instance that runs Amazon Linux 2. The company's application team receives a directive from the legal department to back up the data from the instance's encrypted Amazon Elastic Block Store (Amazon EBS) volume to an Amazon S3 bucket. The application team does not have the administrative SSH key pair for the instance. The application must continue to serve the users.
Which solution will meet these requirements?

A. Attach a role to the instance with permission to write to Amazon S3. Use the AWS Systems Manager Session Manager option to gain access to the instance and run commands to copy data into Amazon S3.
B. Create an image of the instance with the reboot option turned o
C. Launch a new EC2 instance from the imag
D. Attach a role to the new instance with permission to write to Amazon S3. Run a command to copy data into Amazon S3.
E. Take a snapshot of the EBS volume by using Amazon Data Lifecycle Manager (Amazon DLM). Copy the data to Amazon S3.
F. Create an image of the instanc
G. Launch a new EC2 instance from the imag
H. Attach a role to the new instance with permission to write to Amazon S3. Run a command to copy data into Amazon S3.

Answer: C

Explanation:
Taking a snapshot of the EBS volume using Amazon Data Lifecycle Manager (DLM) will meet the requirements because it allows you to create a backup of the volume without the need to access the instance or its SSH key pair. Additionally, DLM allows you to schedule the backups to occur at specific intervals and also enables you to copy the snapshots to an S3 bucket. This approach will not impact the running application as the backup is performed on the EBS volume level.

NEW QUESTION 34
- (Exam Topic 1)
A company used Amazon EC2 instances to deploy a web fleet to host a blog site The EC2 instances are behind an Application Load Balancer (ALB) and are configured in an Auto ScaSng group The web application stores all blog content on an Amazon EFS volume.
The company recently added a feature 'or Moggers to add video to their posts, attracting 10 times the previous user traffic At peak times of day. users report buffering and timeout issues while attempting to reach the site or watch videos
Which is the MOST cost-efficient and scalable deployment that win resolve the issues for users?

A. Reconfigure Amazon EFS to enable maximum I/O.
B. Update the Nog site to use instance store volumes tor storag
C. Copy the site contents to the volumes atlaunch and to Amazon S3 al shutdown.
D. Configure an Amazon CloudFront distributio
E. Point the distribution to an S3 bucket, and migrate the videos from EFS to Amazon S3.
F. Set up an Amazon CloudFront distribution for all site contents, and point the distribution at the ALB.

Answer: C

Explanation:
https://aws.amazon.com/premiumsupport/knowledge-center/cloudfront-https-connection-fails/ Using an Amazon S3 bucket
Using a MediaStore container or a MediaPackage channel Using an Application Load Balancer
Using a Lambda function URL
Using Amazon EC2 (or another custom origin)
Using CloudFront origin groups https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/restrict-access-to-load-balancer.html

NEW QUESTION 35
- (Exam Topic 1)
A company runs a serverless application in a single AWS Region. The application accesses external URLs and extracts metadata from those sites. The company uses an Amazon Simple Notification Service (Amazon SNS) topic to publish URLs to an Amazon Simple Queue Service (Amazon SQS) queue An AWS Lambda function uses the queue as an event source and processes the URLs from the queue Results are saved to an Amazon S3 bucket
The company wants to process each URL other Regions to compare possible differences in site localization URLs must be published from the existing Region. Results must be written to the existing S3 bucket in the current Region.
Which combination of changes will produce multi-Region deployment that meets these requirements? (Select TWO.)

A. Deploy the SOS queue with the Lambda function to other Regions.
B. Subscribe the SNS topic in each Region to the SQS queue.
C. Subscribe the SQS queue in each Region to the SNS topics in each Region.
D. Configure the SQS queue to publish URLs to SNS topics in each Region.
E. Deploy the SNS topic and the Lambda function to other Regions.

Answer: AC

Explanation:
https://docs.aws.amazon.com/sns/latest/dg/sns-cross-region-delivery.html

NEW QUESTION 36
- (Exam Topic 1)
A global media company is planning a multi-Region deployment of an application. Amazon DynamoDB global tables will back the deployment to keep the user experience consistent across the two continents where users are concentrated. Each deployment will have a public Application Load Balancer (ALB). The company manages public DNS internally. The company wants to make the application available through an apex domain.
Which solution will meet these requirements with the LEAST effort?

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

Answer: C

Explanation:
AWS Global Accelerator is a service that directs traffic to optimal endpoints (in this case, the Application Load Balancer) based on the health of the endpoints and network routing. It allows you to create an accelerator that directs traffic to multiple endpoint groups, one for each Region where the application is deployed. The accelerator uses the AWS global network to optimize the traffic routing to the healthy endpoint.
By using Global Accelerator, the company can use a single static IP address for the apex domain, and traffic will be directed to the optimal endpoint based on the user's location, without the need for additional load balancers or routing policies.
Reference:
AWS Global Accelerator documentation: https://aws.amazon.com/global-accelerator/ Routing User Traffic to the Optimal AWS Region using Global Accelerator documentation:
https://aws.amazon.com/blogs/networking-and-content-delivery/routing-user-traffic-to-the-optimal-aws-region-u

NEW QUESTION 37
- (Exam Topic 1)
A company has its cloud infrastructure on AWS A solutions architect needs to define the infrastructure as code. The infrastructure is currently deployed in one AWS Region. The company's business expansion plan includes deployments in multiple Regions across multiple AWS accounts
What should the solutions architect do to meet these requirements?

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

Answer: C

Explanation:
https://aws.amazon.com/blogs/aws/new-use-aws-cloudformation-stacksets-for-multiple-accounts-in-an-aws-orga AWS Organizations allows the management of multiple AWS accounts as a single entity and AWS
CloudFormation StackSets allows creating, updating, and deleting stacks across multiple accounts and regions in an organization. This solution allows creating a single CloudFormation template that can be deployed across multiple accounts and regions, and also allows for the management of access and permissions for the different accounts through the use of IAM roles and policies in the management account.

NEW QUESTION 41
- (Exam Topic 1)
A delivery company needs to migrate its third-party route planning application to AWS. The third party supplies a supported Docker image from a public registry. The image can run in as many containers as required to generate the route map.
The company has divided the delivery area into sections with supply hubs so that delivery drivers travel the shortest distance possible from the hubs to the customers. To reduce the time necessary to generate route maps, each section uses its own set of Docker containers with a custom configuration that processes orders only in the section's area.
The company needs the ability to allocate resources cost-effectively based on the number of running containers.
Which solution will meet these requirements with the LEAST operational overhead?

A. Create an Amazon Elastic Kubernetes Service (Amazon EKS) cluster on Amazon EC2. Use the Amazon EKS CLI to launch the planning application in pods by using the -tags option to assign a custom tag to the pod.
B. Create an Amazon Elastic Kubernetes Service (Amazon EKS) cluster on AWS Fargat
C. Use the Amazon EKS CLI to launch the planning applicatio
D. Use the AWS CLI tag-resource API call to assign a custom tag to the pod.
E. Create an Amazon Elastic Container Service (Amazon ECS) cluster on Amazon EC2. Use the AWS CLI with run-tasks set to true to launch the planning application by using the -tags option to assign a custom tag to the task.
F. Create an Amazon Elastic Container Service (Amazon ECS) cluster on AWS Fargat
G. Use the AWS CLI run-task command and set enableECSManagedTags to true to launch the planning applicatio
H. Use the --tags option to assign a custom tag to the task.

Answer: D

Explanation:
Amazon Elastic Container Service (ECS) on AWS Fargate is a fully managed service that allows you to run containers without having to manage the underlying

infrastructure. When you launch tasks on Fargate,
resources are automatically allocated based on the number of tasks running, which reduces the operational overhead.
Using ECS on Fargate allows you to assign custom tags to tasks using the --tags option in the run-task command, as described in the documentation:
https://docs.aws.amazon.com/cli/latest/reference/ecs/run-task.html You can also set enableECSManagedTags to true, which allows the service to automatically add the cluster name and service name as tags.
https://docs.aws.amazon.com/AmazonECS/latest/developerguide/task-placement-constraints.html#tag-based-sch

NEW QUESTION 43
- (Exam Topic 1)
A health insurance company stores personally identifiable information (PII) in an Amazon S3 bucket. The company uses server-side encryption with S3 managed encryption keys (SSE-S3) to encrypt the objects. According to a new requirement, all current and future objects in the S3 bucket must be encrypted by keys that the company's security team manages. The S3 bucket does not have versioning enabled. Which solution will meet these requirements?

A. In the S3 bucket properties, change the default encryption to SSE-S3 with a customer managed ke
B. Use the AWS CLI to re-upload all objects in the S3 bucke
C. Set an S3 bucket policy to deny unencrypted PutObject requests.
D. In the S3 bucket properties, change the default encryption to server-side encryption with AWS KMS managed encryption keys (SSE-KMS). Set an S3 bucket policy to deny unencrypted PutObject request
E. Use the AWS CLI to re-upload all objects in the S3 bucket.
F. In the S3 bucket properties, change the default encryption to server-side encryption with AWS KMS managed encryption keys (SSE-KMS). Set an S3 bucket policy to automatically encrypt objects on GetObject and PutObject requests.
G. In the S3 bucket properties, change the default encryption to AES-256 with a customer managed key.Attach a policy to deny unencrypted PutObject requests to any entities that access the S3 bucke
H. Use the AWS CLI to re-upload all objects in the S3 bucket.

Answer: D

Explanation:
https://docs.aws.amazon.com/AmazonS3/latest/userguide/ServerSideEncryptionCustomerKeys.html Clearly says we need following header for SSE-C x-amz-server-side-encryption-customer-algorithm Use this header to specify the encryption algorithm. The header value must be AES256.

NEW QUESTION 46
- (Exam Topic 1)
A company wants to use AWS to create a business continuity solution in case the company's main on-premises application fails. The application runs on physical servers that also run other applications. The on-premises application that the company is planning to migrate uses a MySQL database as a data store. All the company's on-premises applications use operating systems that are compatible with Amazon EC2.
Which solution will achieve the company's goal with the LEAST operational overhead?

A. Install the AWS Replication Agent on the source servers, including the MySQL server
B. Set up replication for all server
C. Launch test instances for regular drill
D. Cut over to the test instances to fail over the workload in the case of a failure event.
E. Install the AWS Replication Agent on the source servers, including the MySQL server
F. Initialize AWS Elastic Disaster Recovery in the target AWS Regio
G. Define the launch setting
H. Frequently perform failover and fallback from the most recent point in time.
I. Create AWS Database Migration Service (AWS DMS) replication servers and a target Amazon Aurora MySQL DB cluster to host the databas
J. Create a DMS replication task to copy the existing data to the target DB cluste
K. Create a local AWS Schema Conversion Tool (AWS SCT) change data capture (CDC) task to keep the data synchronize
L. Install the rest of the software on EC2 instances by starting with a compatible base AMI.
M. Deploy an AWS Storage Gateway Volume Gateway on premise
N. Mount volumes on all on-premises server
O. Install the application and the MySQL database on the new volume
P. Take regular snapshot
Q. Install all the software on EC2 Instances by starting with a compatible base AM
R. Launch a Volume Gateway on an EC2 instanc
S. Restore the volumes from the latest snapsho
T. Mount the new volumes on the EC2 instances in the case of a failure event.

Answer: B

Explanation:
https://docs.aws.amazon.com/drs/latest/userguide/what-is-drs.html https://docs.aws.amazon.com/drs/latest/userguide/recovery-workflow-gs.html

NEW QUESTION 48
- (Exam Topic 1)
A company recently completed the migration from an on-premises data center to the AWS Cloud by using a replatforming strategy. One of the migrated servers is running a legacy Simple Mail Transfer Protocol (SMTP) service that a critical application relies upon. The application sends outbound email messages to the company's customers. The legacy SMTP server does not support TLS encryption and uses TCP port 25. The application can use SMTP only.
The company decides to use Amazon Simple Email Service (Amazon SES) and to decommission the legacy SMTP server. The company has created and validated the SES domain. The company has lifted the SES limits.
What should the company do to modify the application to send email messages from Amazon SES?

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

H. Create an IAM role that has ses:SendEmail and ses:SendRawEmail permission
I. Use the IAM role as a service role for Amazon SES.
J. Configure the application to use AWS SDKs to send email message
K. Create an IAM user for Amazon SE
L. Generate API access key
M. Use the access keys to authenticate with Amazon SES.

Answer: B

Explanation:
To set up a STARTTLS connection, the SMTP client connects to the Amazon SES SMTP endpoint on port 25, 587, or 2587, issues an EHLO command, and waits for the server to announce that it supports the STARTTLS SMTP extension. The client then issues the STARTTLS command, initiating TLS negotiation. When negotiation is complete, the client issues an EHLO command over the new encrypted connection, and the SMTP session proceeds normally To set up a TLS Wrapper connection, the SMTP client connects to the Amazon SES SMTP endpoint on port 465 or 2465. The server presents its certificate, the client issues an EHLO command, and the SMTP session proceeds normally.
https://docs.aws.amazon.com/ses/latest/dg/smtp-connect.html

NEW QUESTION 53
- (Exam Topic 1)
A solutions architect is investigating an issue in which a company cannot establish new sessions in Amazon Workspaces. An initial analysis indicates that the issue involves user profiles. The Amazon Workspaces environment is configured to use Amazon FSx for Windows File Server as the profile share storage. The FSx for Windows File Server file system is configured with 10 TB of storage.
The solutions architect discovers that the file system has reached its maximum capacity. The solutions architect must ensure that users can regain access. The solution also must prevent the problem from occurring again.
Which solution will meet these requirements?

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

Answer: B

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

NEW QUESTION 58
- (Exam Topic 2)
A company needs to audit the security posture of a newly acquired AWS account. The company's data security team requires a notification only when an Amazon S3 bucket becomes publicly exposed. The company has already established an Amazon Simple Notification Service (Amazon SNS) topic that has the data security team's email address subscribed.
Which solution will meet these requirements?

A. Create an S3 event notification on all S3 buckets for the isPublic even
B. Select the SNS topic as the target for the event notifications.
C. Create an analyzer in AWS Identity and Access Management Access Analyze
D. Create an Amazon EventBridge rule for the event type "Access Analyzer Finding" with a filter for "isPublic: true." Select the SNS topic as the EventBridge rule target.
E. Create an Amazon EventBridge rule for the event type "Bucket-Level API Call via CloudTrail" with a filter for "PutBucketPolicy." Select the SNS topic as the EventBridge rule target.
F. Activate AWS Config and add the cloudtrail-s3-dataevents-enabled rul
G. Create an Amazon EventBridge rule for the event type "Config Rules Re-evaluation Status" with a filter for "NON_COMPLIANT." Select the SNS topic as the EventBridge rule target.

Answer: B

Explanation:
Access Analyzer is to assess the access policy. https://docs.aws.amazon.com/ja_jp/AmazonS3/latest/userguide/access-control-block-public-access.html

NEW QUESTION 59
- (Exam Topic 2)
A company has five development teams that have each created five AWS accounts to develop and host applications. To track spending, the development teams log in to each account every month, record the current cost from the AWS Billing and Cost Management console, and provide the information to the company's finance team.
The company has strict compliance requirements and needs to ensure that resources are created only in AWS Regions in the United States. However, some resources have been created in other Regions.
A solutions architect needs to implement a solution that gives the finance team the ability to track and consolidate expenditures for all the accounts. The solution also must ensure that the company can create resources only in Regions in the United States.
Which combination of steps will meet these requirements in the MOST operationally efficient way? (Select THREE.)

A. Create a new account to serve as a management accoun
B. Create an Amazon S3 bucket for the finance learn Use AWS Cost and Usage Reports to create monthly reports and to store the data in the finance team's S3

bucket.
C. Create a new account to serve as a management accoun
D. Deploy an organization in AWS Organizations with all features enable
E. Invite all the existing accounts to the organizatio
F. Ensure that each account accepts the invitation.
G. Create an OU that includes all the development team
H. Create an SCP that allows the creation of resources only in Regions that are in the United State
I. Apply the SCP to the OU.
J. Create an OU that includes all the development team
K. Create an SCP that denies (he creation of resources in Regions that are outside the United State
L. Apply the SCP to the OU.
M. Create an 1AM role in the management account Attach a policy that includes permissions to view the Billing and Cost Management consol
N. Allow the finance learn users to assume the rol
O. Use AWS Cost Explorer and the Billing and Cost Management console to analyze cost.
P. Create an 1AM role in each AWS accoun
Q. Attach a policy that includes permissions to view the Billing and Cost Management consol
R. Allow the finance team users to assume the role.

Answer: BCE

Explanation:
AWS Organizations is a service that enables you to consolidate multiple AWS accounts into an organization that you create and centrally manage. By creating a management account and inviting all the existing accounts to join the organization, the solutions architect can track and consolidate expenditures for all the accounts using AWS Cost Management tools such as AWS Cost Explorer and AWS Budgets. An organizational unit (OU) is a group of accounts within an organization that can be used to apply policies and simplify management. A service control policy (SCP) is a type of policy that you can use to manage permissions in your organization. By creating an OU that includes all the development teams and applying an SCP that allows the creation of resources only in Regions that are in the United States, the solutions architect can ensure that the company meets its compliance requirements and avoids unwanted charges from other Regions. An IAM role is an identity with permission policies that determine what the identity can and cannot do in AWS. By creating an IAM role in the management account and allowing the finance team users to assume it, the solutions architect can give them access to view the Billing and Cost Management console without sharing credentials or creating additional users. References:

≫ https://docs.aws.amazon.com/organizations/latest/userguide/orgs_introduction.html

≫ https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scp.html

≫ https://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles.html

≫ https://docs.aws.amazon.com/aws-cost-management/latest/userguide/what-is-costmanagement.html

NEW QUESTION 61
- (Exam Topic 2)
A solutions architect is designing a solution to process events. The solution must have the ability to scale in and out based on the number of events that the solution receives. If a processing error occurs, the event must move into a separate queue for review.
Which solution will meet these requirements?

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

Answer: A

Explanation:
Amazon Simple Notification Service (Amazon SNS) is a fully managed pub/sub messaging service that enables users to send messages to multiple subscribers1. Users can send event details to an Amazon SNS topic and configure an AWS Lambda function as a subscriber to the SNS topic to process the events. Lambda is a serverless compute service that runs code in response to events and automatically manages the underlying compute resources2. Users can add an on-failure destination to the function and set an Amazon Simple Queue
Service (Amazon SQS) queue as the target. Amazon SQS is a fully managed message queuing service that enables users to decouple and scale microservices, distributed systems, and serverless applications3. This way, if a processing error occurs, the event will move into the separate queue for review.
Option B is incorrect because publishing events to an Amazon SQS queue and creating an Amazon EC2 Auto Scaling group will not have the ability to scale in and out based on the number of events that the solution receives. Amazon EC2 is a web service that provides secure, resizable compute capacity in the cloud. Auto Scaling is a feature that helps users maintain application availability and allows them to scale their EC2 capacity up or down automatically according to conditions they define. However, for this use case, using SQS and EC2 will not take advantage of the serverless capabilities of Lambda and SNS.
Option C is incorrect because writing events to an Amazon DynamoDB table and configuring a DynamoDB stream for the table will not have the ability to move events into a separate queue for review if a processing error occurs. Amazon DynamoDB is a fully managed key-value and document database that delivers single-digit millisecond performance at any scale. DynamoDB Streams is a feature that captures data
modification events in DynamoDB tables. Users can configure the stream to invoke a Lambda function, but they cannot configure an on-failure destination for the function.
Option D is incorrect because publishing events to an Amazon EventBridge event bus and setting an Application Load Balancer (ALB) as the event bus target will not have the ability to move events into a separate queue for review if a processing error occurs. Amazon EventBridge is a serverless event bus service that makes it easy to connect applications with data from a variety of sources. An ALB is a load balancer that distributes incoming application traffic across multiple

targets, such as EC2 instances, containers, IP addresses, Lambda functions, and virtual appliances. Users can configure EventBridge to retry events, but they cannot configure an on-failure destination for the ALB.

NEW QUESTION 62
- (Exam Topic 2)
A company runs its sales reporting application in an AWS Region in the United States. The application uses an Amazon API Gateway Regional API and AWS Lambda functions to generate on-demand reports from data in an Amazon RDS for MySQL database. The frontend of the application is hosted on Amazon S3 and is accessed by users through an Amazon CloudFront distribution. The company is using Amazon Route 53 as the DNS service for the domain. Route 53 is configured with a simple routing policy to route traffic to the API Gateway API.
In the next 6 months, the company plans to expand operations to Europe. More than 90% of the database traffic is read-only traffic. The company has already deployed an API Gateway API and Lambda functions in the new Region.
A solutions architect must design a solution that minimizes latency for users who download reports. Which solution will meet these requirements?

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

Answer: C

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

≫ Using AWS Database Migration Service (AWS DMS) to replicate the primary database in the original Region to the database in the new Region would not be as cost-effective or simple as using a
cross-Region read replica. AWS DMS is a service that enables you to migrate relational databases, data
warehouses, NoSQL databases, and other types of data stores. You can use AWS DMS to perform one-time migrations or continuous data replication with high availability and consolidate databases into
a petabyte-scale data warehouse3. However, AWS DMS requires more configuration and management than creating a cross-Region read replica, which is fully managed by Amazon RDS. AWS DMS also incurs additional charges for replication instances and tasks.

≫ Creating an Amazon API Gateway Data API service integration with Amazon Redshift would not help with disaster recovery or minimizing latency. The Data API is a feature that enables you to query your Amazon Redshift cluster using HTTP requests, without needing a persistent connection or a SQL client. It is useful for building applications that interact with Amazon Redshift, but not for replicating or recovering data from an RDS database.

≫ Creating an AWS Data Exchange datashare by connecting AWS Data Exchange to the Redshift cluster would not help with disaster recovery or minimizing latency. AWS Data Exchange is a service that makes it easy for AWS customers to exchange data in the cloud. You can use AWS Data Exchange to subscribe to a diverse selection of third-party data products or offer your own data products to other AWS customers. A datashare is a feature that enables you to share live and secure access to your Amazon Redshift data across your accounts or with third parties without copying or moving the underlying data. It is useful for sharing query results and views with other users, but not for replicating or recovering data from an RDS database.
References:
≫ https://aws.amazon.com/dms/
≫ https://docs.aws.amazon.com/redshift/latest/mgmt/data-api.html
≫ https://aws.amazon.com/data-exchange/
≫ https://docs.aws.amazon.com/redshift/latest/dg/datashare-overview.html

NEW QUESTION 66
- (Exam Topic 2)
A company wants to containerize a multi-tier web application and move the application from an on-premises data center to AWS. The application includes web. application, and database tiers. The company needs to make the application fault tolerant and scalable. Some frequently accessed data must always be available across application servers. Frontend web servers need session persistence and must scale to meet increases in traffic.
Which solution will meet these requirements with the LEAST ongoing operational overhead?

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

Answer: D

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

≫ https://docs.aws.amazon.com/eks/latest/userguide/managed-node-groups.html

≫ https://docs.aws.amazon.com/eks/latest/userguide/deployments.html

≫ https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Introduction.html

≫ https://docs.aws.amazon.com/efs/latest/ug/mounting-fs.html

NEW QUESTION 71
- (Exam Topic 2)
A company operates an on-premises software-as-a-service (SaaS) solution that ingests several files daily. The company provides multiple public SFTP endpoints to its customers to facilitate the file transfers. The customers add the SFTP endpoint IP addresses to their firewall allow list for outbound traffic. Changes to the SFTP endmost IP addresses are not permitted.
The company wants to migrate the SaaS solution to AWS and decrease the operational overhead of the file transfer service.
Which solution meets these requirements?

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

Answer: A

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

NEW QUESTION 75
- (Exam Topic 2)
A company is creating a centralized logging service running on Amazon EC2 that will receive and analyze logs from hundreds of AWS accounts. AWS PrivateLink is being used to provide connectivity between the client services and the logging service.
In each AWS account with a client, an interface endpoint has been created for the logging service and is available. The logging service running on EC2 instances with a Network Load Balancer (NLB) are deployed in different subnets. The clients are unable to submit logs using the VPC endpoint.
Which combination of steps should a solutions architect take to resolve this issue? (Select TWO.)

A. Check that the NACL is attached to the logging service subnet to allow communications to and from the NLB subnet
B. Check that the NACL is attached to the NLB subnet to allow communications to and from the logging service subnets running on EC2 instances.
C. Check that the NACL is attached to the logging service subnets to allow communications to and from the interface endpoint subnet
D. Check that the NACL is attached to the interface endpoint subnet to allow communications to and from the logging service subnets running on EC2 instances.
E. Check the security group for the logging service running on the EC2 instances to ensure it allows Ingress from the NLB subnets.
F. Check the security group for the loggia service running on EC2 instances to ensure it allows ingress from the clients.
G. Check the security group for the NLB to ensure it allows ingress from the interlace endpoint subnets.

Answer: AC

NEW QUESTION 76
- (Exam Topic 2)
A company uses a Grafana data visualization solution that runs on a single Amazon EC2 instance to monitor the health of the company's AWS workloads. The company has invested time and effort to create dashboards that the company wants to preserve. The dashboards need to be highly available and cannot be down for longer than 10 minutes. The company needs to minimize ongoing maintenance.
Which solution will meet these requirements with the LEAST operational overhead?

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

K. Configure AWS Backup to back up the EC2 instance that runs Grafana once each hou
L. Restore the EC2 instance from the most recent snapshot in an alternate Availability Zone when required.

Answer: C

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

NEW QUESTION 81
- (Exam Topic 2)
A company has multiple business units that each have separate accounts on AWS. Each business unit manages its own network with several VPCs that have CIDR ranges that overlap. The company's marketing team has created a new internal application and wants to make the application accessible to all the other business units. The solution must use private IP addresses only.
Which solution will meet these requirements with the LEAST operational overhead?

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

Answer: C

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

NEW QUESTION 86
- (Exam Topic 2)
A company is creating a REST API to share information with six of its partners based in the United States. The company has created an Amazon API Gateway Regional endpoint. Each of the six partners will access the API once per day to post daily sales figures.
After initial deployment, the company observes 1.000 requests per second originating from 500 different IP addresses around the world. The company believes this traffic is originating from a botnet and wants to secure its API while minimizing cost.
Which approach should the company take to secure its API?

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

Answer: D

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

NEW QUESTION 87

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

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

Answer: C

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

NEW QUESTION 89
- (Exam Topic 2)
A company is running a critical stateful web application on two Linux Amazon EC2 instances behind an Application Load Balancer (ALB) with an Amazon RDS for MySQL database The company hosts the DNS records for the application in Amazon Route 53 A solutions architect must recommend a solution to improve the resiliency of the application
The solution must meet the following objectives:
• Application tier RPO of 2 minutes. RTO of 30 minutes
• Database tier RPO of 5 minutes RTO of 30 minutes
The company does not want to make significant changes to the existing application architecture The company must ensure optimal latency after a failover
Which solution will meet these requirements?

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

Answer: B

Explanation:
This option meets the RPO and RTO requirements for both the application and database tiers and uses tools like Amazon DLM and RDS automated backups to create and manage the backups. Additionally, it uses Global Accelerator to ensure low latency after failover by directing traffic to the closest healthy endpoint.

NEW QUESTION 93
- (Exam Topic 2)
A software-as-a-service (SaaS) provider exposes APIs through an Application Load Balancer (ALB). The ALB connects to an Amazon Elastic Kubernetes Service (Amazon EKS) cluster that is deployed in the
us-east-I Region. The exposed APIs contain usage of a few non-standard REST methods: LINK, UNLINK, LOCK, and UNLOCK.
Users outside the United States are reporting long and inconsistent response times for these APIs. A solutions architect needs to resolve this problem with a solution that minimizes operational overhead.
Which solution meets these requirements?

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

Answer: C

Explanation:

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

NEW QUESTION 96
- (Exam Topic 2)
A company has a few AWS accounts for development and wants to move its production application to AWS. The company needs to enforce Amazon Elastic Block Store (Amazon EBS) encryption at rest current production accounts and future production accounts only. The company needs a solution that includes built-in blueprints and guardrails.
Which combination of steps will meet these requirements? (Choose three.)

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

Answer: CDF

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

NEW QUESTION 101
- (Exam Topic 2)
A company is implementing a serverless architecture by using AWS Lambda functions that need to access a Microsoft SQL Server DB instance on Amazon RDS. The company has separate environments for development and production, including a clone of the database system.
The company's developers are allowed to access the credentials for the development database. However, the credentials for the production database must be encrypted with a key that only members of the IT security team's IAM user group can access. This key must be rotated on a regular basis.
What should a solutions architect do in the production environment to meet these requirements?

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

Answer: D

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

NEW QUESTION 105
- (Exam Topic 2)
A company uses AWS Organizations with a single OU named Production to manage multiple accounts All accounts are members of the Production OU Administrators use deny list SCPs in the root of the organization to manage access to restricted services.
The company recently acquired a new business unit and invited the new unit's existing AWS account to the organization Once onboarded the administrators of the new business unit discovered that they are not able to update existing AWS Config rules to meet the company's policies.
Which option will allow administrators to make changes and continue to enforce the current policies without introducing additional long-term maintenance?

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

Answer: D

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

NEW QUESTION 109
- (Exam Topic 2)
A company is migrating a document processing workload to AWS. The company has updated many applications to natively use the Amazon S3 API to store, retrieve, and modify documents that a processing server generates at a rate of approximately 5 documents every second. After the document processing is finished, customers can download the documents directly from Amazon S3.
During the migration, the company discovered that it could not immediately update the processing server that generates many documents to support the S3 API. The server runs on Linux and requires fast local access to the files that the server generates and modifies. When the server finishes processing, the files must be available to the public for download within 30 minutes.
Which solution will meet these requirements with the LEAST amount of effort?

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

Answer: C

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

≫ Migrating the application to an AWS Lambda function would require a lot of effort and may not be feasible for the existing server that generates many documents. Lambda functions have limitations on execution time, memory, disk space, and network bandwidth.

≫ Setting up an Amazon S3 File Gateway would not work because S3 File Gateway does not support write-back caching, which means that files written to the file share are uploaded to S3 immediately and are not available locally until they are downloaded again. This would not provide fast local access to the
files that the server generates and modifies.

≫ Configuring AWS DataSync to connect to an Amazon EC2 instance would not meet the requirement of making the files available to the public for download within 30 minutes. DataSync is a service that transfers data between on-premises storage systems and AWS storage services over the internet or AWS Direct Connect. DataSync tasks can be scheduled to run at specific times or intervals, but they are not triggered by file changes.
References:
≫ https://aws.amazon.com/fsx/lustre/
≫ https://docs.aws.amazon.com/fsx/latest/LustreGuide/create-fs-linked-data-repo.html
≫ https://docs.aws.amazon.com/fsx/latest/LustreGuide/import-export-data-repositories.html
≫ https://docs.aws.amazon.com/fsx/latest/LustreGuide/mounting-on-premises.html
≫ https://docs.aws.amazon.com/lambda/latest/dg/gettingstarted-limits.html
≫ https://docs.aws.amazon.com/storagegateway/latest/userguide/StorageGatewayConcepts.html
≫ https://docs.aws.amazon.com/datasync/latest/userguide/what-is-datasync.html

NEW QUESTION 114
- (Exam Topic 2)
A company is building a call center by using Amazon Connect. The company's operations team is defining a disaster recovery (DR) strategy across AWS Regions. The contact center has dozens of contact flows, hundreds of users, and dozens of claimed phone numbers.
Which solution will provide DR with the LOWEST RTO?

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

N. Provision a new Amazon Connect instance with all existing users and contact flows in a second Region.Create an Amazon Route 53 health check for the URL of the Amazon Connect instanc

O. Create an Amazon CloudWatch alarm for failed health check

P. Create an AWS Lambda function to deploy an AWS CloudFormation template that provisions claimed phone number

Q. Configure the alarm to invoke the Lambda function.

Answer: D

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

NEW QUESTION 115
- (Exam Topic 2)
A company is storing sensitive data in an Amazon S3 bucket. The company must log all activities for objects in the S3 bucket and must keep the logs for 5 years. The company's security team also must receive an email notification every time there is an attempt to delete data in the S3 bucket.
Which combination of steps will meet these requirements MOST cost-effectively? (Select THREE.)

A. Configure AWS CloudTrail to log S3 data events.

B. Configure S3 server access logging for the S3 bucket.

C. Configure Amazon S3 to send object deletion events to Amazon Simple Email Service (Amazon SES).

D. Configure Amazon S3 to send object deletion events to an Amazon EventBridge event bus that publishes to an Amazon Simple Notification Service (Amazon SNS) topic.

E. Configure Amazon S3 to send the logs to Amazon Timestream with data storage tiering.

F. Configure a new S3 bucket to store the logs with an S3 Lifecycle policy.

Answer: ADF

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

NEW QUESTION 119
- (Exam Topic 2)
A company has IoT sensors that monitor traffic patterns throughout a large city. The company wants to read and collect data from the sensors and perform aggregations on the data.
A solutions architect designs a solution in which the IoT devices are streaming to Amazon Kinesis Data Streams. Several applications are reading from the stream. However, several consumers are experiencing throttling and are periodically and are periodically encountering a RealProvisioned Throughput Exceeded error.
Which actions should the solution architect take to resolve this issue? (Select THREE.)

A. Reshard the stream to increase the number of shards s in the stream.

B. Use the Kinesis Producer Library KPL). Adjust the polling frequency.

C. Use consumers with the enhanced fan-out feature.

D. Reshard the stream to reduce the number of shards in the stream.

E. Use an error retry and exponential backoff mechanism in the consumer logic.

F. Configure the stream to use dynamic partitioning.

Answer: ACE

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

NEW QUESTION 124
- (Exam Topic 2)
A company needs to optimize the cost of an AWS environment that contains multiple accounts in an organization in AWS Organizations. The company conducted cost optimization activities 3 years ago and purchased Amazon EC2 Standard Reserved Instances that recently expired.
The company needs EC2 instances for 3 more years. Additionally, the company has deployed a new serverless workload.
Which strategy will provide the company with the MOST cost savings?

A. Purchase the same Reserved Instances for an additional 3-year term with All Upfront paymen

B. Purchase a 3-year Compute Savings Plan with All Upfront payment in the management account to cover any additional compute costs.

C. Purchase a I-year Compute Savings Plan with No Upfront payment in each member accoun

D. Use the Savings Plans recommendations in the AWS Cost Management console to choose the Compute Savings Plan.

E. Purchase a 3-year EC2 Instance Savings Plan with No Upfront payment in the management account to cover EC2 costs in each AWS Regio

F. Purchase a 3 year Compute Savings Plan with No Upfront payment in the management account to cover any additional compute costs.

G. Purchase a 3-year EC2 Instance Savings Plan with All Upfront payment in each member accoun

H. Use the Savings Plans recommendations in the AWS Cost Management console to choose the EC2 Instance Savings Plan.

Answer: A

Explanation:
The company should purchase the same Reserved Instances for an additional 3-year term with All Upfront payment. The company should purchase a 3-year Compute Savings Plan with All Upfront payment in the management account to cover any additional compute costs. This solution will provide the company with the most cost savings because Reserved Instances and Savings Plans are both pricing models that offer significant discounts compared to On-Demand pricing. Reserved Instances are commitments to use a specific instance type and size in a single Region for a one- or three-year term. You can choose between three payment options:
No Upfront, Partial Upfront, or All Upfront. The more you pay upfront, the greater the discoun1t. Savings Plans are flexible pricing models that offer low prices on EC2 instances, Fargate, and Lambda usage, in exchange for a commitment to a consistent amount of usage (measured in $/hour) for a one- or three-year term. You can choose between two types of Savings Plans: Compute Savings Plans and EC2 Instance Savings Plans. Compute Savings Plans apply to any EC2 instance regardless of Region, instance family, operating system, or tenancy, including those that are part of EMR, ECS, or EKS clusters, or launched by Fargate or Lambda. EC2 Instance Savings Plans apply to a specific instance family within a Region and provide the most savings2. By purchasing the same Reserved Instances for an additional 3-year term with All Upfront payment, the company can lock in the lowest possible price for its EC2 instances that run continuously for 3 years. By purchasing a 3-year Compute Savings Plan with All Upfront payment in the management account, the company can benefit from additional discounts on any other compute usage across its member accounts.
The other options are not correct because:

> Purchasing a 1-year Compute Savings Plan with No Upfront payment in each member account would not provide as much cost savings as purchasing a 3-year Compute Savings Plan with All Upfront payment in the management account. A 1-year term offers lower discounts than a 3-year term, and a No Upfront payment option offers lower discounts than an All Upfront payment option. Also, purchasing a Savings Plan in each member account would not allow the company to share the benefits of unused Savings Plan discounts across its organization.

> Purchasing a 3-year EC2 Instance Savings Plan with No Upfront payment in the management account to cover EC2 costs in each AWS Region would not provide as much cost savings as purchasing Reserved Instances for an additional 3-year term with All Upfront payment. An EC2 Instance Savings Plan offers lower discounts than Reserved Instances for the same instance family and Region. Also, a No Upfront payment option offers lower discounts than an All Upfront payment option.

> Purchasing a 3-year EC2 Instance Savings Plan with All Upfront payment in each member account would not provide as much flexibility or cost savings as purchasing a 3-year Compute Savings Plan with All Upfront payment in the management account. An EC2 Instance Savings Plan applies only to a specific instance family within a Region and does not cover Fargate or Lambda usage. Also, purchasing a Savings Plan in each member account would not allow the company to share the benefits of unused Savings Plan discounts across its organization.
References:

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

> https://aws.amazon.com/savingsplans/

NEW QUESTION 125
- (Exam Topic 2)
An external audit of a company's serverless application reveals IAM policies that grant too many permissions. These policies are attached to the company's AWS Lambda execution roles. Hundreds of the company's Lambda functions have broad access permissions, such as full access to Amazon S3 buckets and Amazon DynamoDB tables. The company wants each function to have only the minimum permissions that the function needs to complete its task.
A solutions architect must determine which permissions each Lambda function needs.
What should the solutions architect do to meet this requirement with the LEAST amount of effort?

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

Answer: B

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

NEW QUESTION 128
- (Exam Topic 2)
A company wants to migrate to AWS. The company is running thousands of VMs in a VMware ESXi environment. The company has no configuration management database and has little Knowledge about the utilization of the VMware portfolio.
A solutions architect must provide the company with an accurate inventory so that the company can plan for a cost-effective migration.
Which solution will meet these requirements with the LEAST operational overhead?

A. Use AWS Systems Manager Patch Manager to deploy Migration Evaluator to each V
B. Review the collected data in Amazon QuickSigh
C. Identify servers that have high utilizatio
D. Remove the servers that have high utilization from the migration lis
E. Import the data to AWS Migration Hub.

F. Export the VMware portfolio to a csv fil
G. Check the disk utilization for each serve
H. Remove servers that have high utilizatio
I. Export the data to AWS Application Migration Servic
J. Use AWS Server Migration Service (AWS SMS) to migrate the remaining servers.
K. Deploy the Migration Evaluator agentless collector to the ESXi hyperviso
L. Review the collected data in Migration Evaluato
M. Identify inactive server
N. Remove the inactive servers from the migration lis
O. Import the data to AWS Migration Hub.
P. Deploy the AWS Application Migration Service Agent to each V
Q. When the data is collected, useAmazon Redshift to import and analyze the dat
R. Use Amazon QuickSight for data visualization.

Answer: C

Explanation:
https://aws.amazon.com/migration-evaluator/features/

NEW QUESTION 131
......

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