Exam Questions SAP-C02

AWS Certified Solutions Architect - Professional

https://www.2passeasy.com/dumps/SAP-C02/

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

A. Configure AWS Budgets in each account and configure budget alerts that are grouped by application, environment, and owne
B. Add each business unit to an Amazon SNS topic for each aler
C. Use Cost Explorer in each account to create monthly reports for each business unit.
D. Configure AWS Budgets in the organization's master account and configure budget alerts that are grouped by application, environment, and owne
E. Add each business unit to an Amazon SNS topic for each aler
F. Use Cost Explorer in the organization's master account to create monthly reports for each business unit.
G. Configure AWS Budgets in each account and configure budget alerts lhat are grouped by application, environment, and owne
H. Add each business unit to an Amazon SNS topic for each aler
I. Use the AWS Billing and Cost Management dashboard in each account to create monthly reports for each businessunit.
J. Enable AWS Cost and Usage Reports in the organization's master account and configure reports grouped by application, environment, and owne
K. Create an AWS Lambda function that processes AWS Cost and Usage Reports, sends budget alerts, and sends monthly reports to each business unit's email list.

Answer: B

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

NEW QUESTION 2
- (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 3
- (Exam Topic 1)
A company is refactoring its on-premises order-processing platform in the AWS Cloud. The platform includes a web front end that is hosted on a fleet of VMs RabbitMQ to connect the front end to the backend, and a Kubernetes cluster to run a containerized backend system to process the orders. The company does not want to make any major changes to the application
Which solution will meet these requirements with the LEAST operational overhead?

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

Answer: A

Explanation:

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

NEW QUESTION 4
- (Exam Topic 1)
An international delivery company hosts a delivery management system on AWS. Drivers use the system to upload confirmation of delivery. Confirmation includes the recipient's signature or a photo of the package with the recipient. The driver's handheld device uploads signatures and photos through FTP to a single Amazon EC2 instance. Each handheld device saves a file in a directory based on the signed-in user, and the file name matches the delivery number. The EC2 instance then adds metadata to the file after querying a central database to pull delivery information. The file is then placed in Amazon S3 for archiving.
As the company expands, drivers report that the system is rejecting connections. The FTP server is having problems because of dropped connections and memory issues. In response to these problems, a system engineer schedules a cron task to reboot the EC2 instance every 30 minutes. The billing team reports that files are not always in the archive and that the central system is not always updated.
A solutions architect needs to design a solution that maximizes scalability to ensure that the archive always receives the files and that systems are always updated. The handheld devices cannot be modified, so the company cannot deploy a new application.
Which solution will meet these requirements?

A. Create an AMI of the existing EC2 instanc
B. Create an Auto Scaling group of EC2 instances behind an Application Load Balance
C. Configure the Auto Scaling group to have a minimum of three instances.
D. Use AWS Transfer Family to create an FTP server that places the files in Amazon Elastic File System (Amazon EFS). Mount the EFS volume to the existing EC2 instanc
E. Point the EC2 instance to the new path for file processing.
F. Use AWS Transfer Family to create an FTP server that places the files in Amazon S3. Use an S3 event notification through Amazon Simple Notification Service (Amazon SNS) to invoke an AWS Lambda functio
G. Configure the Lambda function to add the metadata and update the delivery system.
H. Update the handheld devices to place the files directly in Amazon S3. Use an S3 event notification through Amazon Simple Queue Service (Amazon SQS) to invoke an AWS Lambda functio
I. Configure the Lambda function to add the metadata and update the delivery system.

Answer: C

Explanation:
Using AWS Transfer Family to create an FTP server that places the files in Amazon S3 and using S3 event notifications through Amazon Simple Notification Service (Amazon SNS) to invoke an AWS Lambda function will ensure that the archive always receives the files and that the central system is always updated. This solution maximizes scalability and eliminates the need for manual intervention, such as rebooting the EC2 instance.

NEW QUESTION 5
- (Exam Topic 1)
A company is hosting a three-tier web application in an on-premises environment. Due to a recent surge in traffic that resulted in downtime and a significant financial impact, company management has ordered that the application be moved to AWS. The application is written in .NET and has a dependency on a MySQL database A solutions architect must design a scalable and highly available solution to meet the demand of 200000 daily users.
Which steps should the solutions architect take to design an appropriate solution?

A. Use AWS Elastic Beanstalk to create a new application with a web server environment and an Amazon RDS MySQL Multi-AZ DB instance The environment should launch a Network Load Balancer (NLB) in front of an Amazon EC2 Auto Scaling group in multiple Availability Zones Use an Amazon Route 53 alias record to route traffic from the company's domain to the NLB.
B. Use AWS CloudFormation to launch a stack containing an Application Load Balancer (ALB) in front of an Amazon EC2 Auto Scaling group spanning three Availability Zone
C. The stack should launch a Multi-AZ deployment of an Amazon Aurora MySQL DB cluster with a Retain deletion polic
D. Use an Amazon Route 53 alias record to route traffic from the company's domain to the ALB
E. Use AWS Elastic Beanstalk to create an automatically scaling web server environment that spans two separate Regions with an Application Load Balancer (ALB) in each Regio
F. Create a Multi-AZ deployment of an Amazon Aurora MySQL DB cluster with a cross-Region read replica Use Amazon Route 53 with a geoproximity routing policy to route traffic between the two Regions.
G. Use AWS CloudFormation to launch a stack containing an Application Load Balancer (ALB) in front of an Amazon ECS cluster of Spot Instances spanning three Availability Zones The stack should launch an Amazon RDS MySQL DB instance with a Snapshot deletion policy Use an Amazon Route 53 alias record to route traffic from the company's domain to the ALB

Answer: C

Explanation:
Using AWS CloudFormation to launch a stack with an Application Load Balancer (ALB) in front of an Amazon EC2 Auto Scaling group spanning three Availability Zones, a Multi-AZ deployment of an Amazon Aurora MySQL DB cluster with a Retain deletion policy, and an Amazon Route 53 alias record to route traffic from the company's domain to the ALB will ensure that

NEW QUESTION 6
- (Exam Topic 1)
A software company hosts an application on AWS with resources in multiple AWS accounts and Regions. The application runs on a group of Amazon EC2 instances in an application VPC located in the us-east-1 Region with an IPv4 CIDR block of 10.10.0.0/16. In a different AWS account, a shared services VPC is located in the us-east-2 Region with an IPv4 CIDR block of 10.10.10.0/24. When a cloud engineer uses AWS CloudFormation to attempt to peer the application VPC with the shared services VPC, an error message indicates a peering failure. Which factors could cause this error? (Choose two.)

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

Answer: AE

Explanation:

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

NEW QUESTION 7
- (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 8
- (Exam Topic 1)
A company wants to migrate its workloads from on premises to AWS. The workloads run on Linux and Windows. The company has a large on-premises intra structure that consists of physical machines and VMs that host numerous applications.
The company must capture details about the system configuration. system performance. running processure and network coi.net lions of its o. -premises ,on boards. The company also must divide the on-premises applications into groups for AWS migrations. The company needs recommendations for Amazon EC2 instance types so that the company can run its workloads on AWS in the most cost-effective manner.
Which combination of steps should a solutions architect take to meet these requirements? (Select THREE.)

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

Answer: ADE

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

NEW QUESTION 9
- (Exam Topic 1)
A company has a serverless application comprised of Amazon CloudFront, Amazon API Gateway, and AWS Lambda functions. The current deployment process of the application code is to create a new version number of the Lambda function and run an AWS CLI script to update. If the new function version has errors, another CLI script reverts by deploying the previous working version of the function. The company would like to decrease the time to deploy new versions of the application logic provided by the Lambda functions, and also reduce the time to detect and revert when errors are identified.
How can this be accomplished?

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

Answer: B

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

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

A solutions architect needs to implement a solution to minimize the cost of the table. Which solution will meet these requirements?

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

Answer: D

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

NEW QUESTION 10
- (Exam Topic 1)
A company is running applications on AWS in a multi-account environment. The company's sales team and marketing team use separate AWS accounts in AWS Organizations.
The sales team stores petabytes of data in an Amazon S3 bucket. The marketing team uses Amazon QuickSight for data visualizations. The marketing team needs access to data that the sales team stores in the S3 bucket. The company has encrypted the S3 bucket with an AWS Key Management Service (AWS KMS) key. The marketing team has already created the IAM service role for QuickSight to provide QuickSight access in the marketing AWS account. The company needs a solution that will provide secure access to the data in the S3 bucket across AWS accounts.
Which solution will meet these requirements with the LEAST operational overhead?

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

Answer: D

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

➢ AWS IAM Roles
➢ AWS KMS - Key Grants
➢ AWS RAM

NEW QUESTION 12
- (Exam Topic 1)
A company has 50 AWS accounts that are members of an organization in AWS Organizations Each account contains multiple VPCs The company wants to use AWS Transit Gateway to establish connectivity between the VPCs in each member account Each time a new member account is created, the company wants to automate the process of creating a new VPC and a transit gateway attachment.
Which combination of steps will meet these requirements? (Select TWO)

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

Answer: AC

Explanation:

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

NEW QUESTION 15
- (Exam Topic 1)
A company is migrating some of its applications to AWS. The company wants to migrate and modernize the applications quickly after it finalizes networking and security strategies. The company has set up an AWS Direct Connection connection in a central network account.
The company expects to have hundreds of AWS accounts and VPCs in the near future. The corporate network must be able to access the resources on AWS seamlessly and also must be able to communicate with all the VPCs. The company also wants to route its cloud resources to the internet through its on-premises data center.
Which combination of steps will meet these requirements? (Choose three.)

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

Answer: BDF

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

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

▷ Option C is incorrect because provisioning an internet gateway, attaching the internet gateway to subnets, and allowing internet traffic through the gateway does not meet the requirement of routing cloud resources to the internet through its on-premises data center. An internet gateway is a horizontally scaled, redundant, and highly available VPC component that allows communication between your VPC and the internet. An internet gateway serves two purposes: to provide a target in your VPC route tables for internet-routable traffic, and to perform network address translation (NAT) for instances that have been assigned public IPv4 addresses. By using an internet gateway, you are routing cloud resources directly to the internet, not through your on-premises data center.

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

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

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

NEW QUESTION 20
- (Exam Topic 1)
A publishing company's design team updates the icons and other static assets that an ecommerce web application uses. The company serves the icons and assets from an Amazon S3 bucket that is hosted in the company's production account. The company also uses a development account that members of the design team can access.
After the design team tests the static assets in the development account, the design team needs to load the assets into the S3 bucket in the production account. A

solutions architect must provide the design team with access to the production account without exposing other parts of the web application to the risk of unwanted changes.
Which combination of steps will meet these requirements? (Select THREE.)

A. In the production account, create a new IAM policy that allows read and write access to the S3 bucket.
B. In the development account, create a new IAM policy that allows read and write access to the S3 bucket.
C. In the production account, create a rol
D. Attach the new policy to the rol
E. Define the development account as a trusted entity.
F. In the development account, create a rol
G. Attach the new policy to the rol
H. Define the production account as a trusted entity.
I. In the development account, create a group that contains all the IAM users of the design tea
J. Attach a different IAM policy to the group to allow the sts:AssumeRole action on the role in the production account.
K. In the development account, create a group that contains all tfje IAM users of the design tea
L. Attach a different IAM policy to the group to allow the sts;AssumeRole action on the role in the development account.

Answer: ACE

Explanation:

≫ A. In the production account, create a new IAM policy that allows read and write access to the S3 bucket. The policy grants the necessary permissions to access the assets in the production S3 bucket.

≫ C. In the production account, create a role. Attach the new policy to the role. Define the development account as a trusted entity. By creating a role and attaching the policy, and then defining the development account as a trusted entity, the development account can assume the role and access the production S3 bucket with the read and write permissions.

≫ E. In the development account, create a group that contains all the IAM users of the design team. Attach a different IAM policy to the group to allow the sts:AssumeRole action on the role in the production account. The IAM policy attached to the group allows the design team members to assume the role created in the production account, thereby giving them access to the production S3 bucket.
Step 1: Create a role in the Production Account; create the role in the Production account and specify the Development account as a trusted entity. You also limit the role permissions to only read and write access to the productionapp bucket. Anyone granted permission to use the role can read and write to the productionapp bucket. Step 2: Grant access to the role Sign in as an administrator in the Development account and allow the AssumeRole action on the UpdateApp role in the Production account. So, recap, production account you create the policy for S3, and you set development account as a trusted entity. Then on the development account you allow the sts:assumeRole action on the role in production account. https://docs.aws.amazon.com/IAM/latest/UserGuide/tutorial_cross-account-with-roles.html

---

NEW QUESTION 22
- (Exam Topic 1)
A company is running an application on several Amazon EC2 instances in an Auto Scaling group behind an Application Load Balancer. The load on the application varies throughout the day, and EC2 instances are scaled in and out on a regular basis. Log files from the EC2 instances are copied to a central Amazon S3 bucket every 15 minutes. The security team discovers that log files are missing from some of the terminated EC2 instances.
Which set of actions will ensure that log files are copied to the central S3 bucket from the terminated EC2 instances?

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

Answer: B

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

---

NEW QUESTION 25
- (Exam Topic 1)
A finance company hosts a data lake in Amazon S3. The company receives financial data records over SFTP each night from several third parties. The company runs its own SFTP server on an Amazon EC2 instance in a public subnet of a VPC. After the files ate uploaded, they are moved to the data lake by a cron job that runs on the same instance. The SFTP server is reachable on DNS sftp.examWe.com through the use of Amazon Route 53.
What should a solutions architect do to improve the reliability and scalability of the SFTP solution?

A. Move the EC2 instance into an Auto Scaling grou
B. Place the EC2 instance behind an Application Load Balancer (ALB). Update the DNS record sftp.example.com in Route 53 to point to the ALB.
C. Migrate the SFTP server to AWS Transfer for SFT

D. Update the DNS record sftp.example.com in Route 53 to point to the server endpoint hostname.
E. Migrate the SFTP server to a file gateway in AWS Storage Gatewa
F. Update the DNS record sflp.example.com in Route 53 to point to the file gateway endpoint.
G. Place the EC2 instance behind a Network Load Balancer (NLB). Update the DNS record sftp.example.com in Route 53 to point to the NLB.

Answer: B

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

NEW QUESTION 29
- (Exam Topic 1)
A company wants to migrate to AWS. The company wants to use a multi-account structure with centrally managed access to all accounts and applications. The company also wants to keep the traffic on a private network. Multi-factor authentication (MFA) is required at login, and specific roles are assigned to user groups. The company must create separate accounts for development. staging, production, and shared network. The production account and the shared network account must have connectivity to all accounts. The development account and the staging account must have access only to each other.
Which combination of steps should a solutions architect take 10 meet these requirements? (Choose three.)

A. Deploy a landing zone environment by using AWS Control Towe
B. Enroll accounts and invite existing accounts into the resulting organization in AWS Organizations.
C. Enable AWS Security Hub in all accounts to manage cross-account acces
D. Collect findings through AWS CloudTrail to force MFA login.
E. Create transit gateways and transit gateway VPC attachments in each accoun
F. Configure appropriate route tables.
G. Set up and enable AWS IAM Identity Center (AWS Single Sign-On). Create appropriate permission sets with required MFA for existing accounts.
H. Enable AWS Control Tower in all Recounts to manage routing between account
I. Collect findings through AWS CloudTrail to force MFA login.
J. Create IAM users and group
K. Configure MFA for all user
L. Set up Amazon Cognito user pools andidentity pools to manage access to accounts and between accounts.

Answer: ACD

Explanation:
The correct answer would be options A, C and D, because they address the requirements outlined in the question. A. Deploying a landing zone environment using AWS Control Tower and enrolling accounts in an organization in AWS Organizations allows for a centralized management of access to all accounts and applications. C. Creating transit gateways and transit gateway VPC attachments in each account and configuring appropriate route tables allows for private network traffic, and ensures that the production account and shared network account have connectivity to all accounts, while the development and staging accounts have access only to each other. D. Setting up and enabling AWS IAM Identity Center (AWS Single Sign-On) and creating appropriate permission sets with required MFA for existing accounts allows for multi-factor authentication at login and specific roles to be assigned to user groups.

NEW QUESTION 32
- (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 33
- (Exam Topic 1)
A company has an organization in AWS Organizations that has a large number of AWS accounts. One of the AWS accounts is designated as a transit account and has a transit gateway that is shared with all of the other AWS accounts AWS Site-to-Site VPN connections are configured between ail of the company's global offices and the transit account The company has AWS Config enabled on all of its accounts.
The company's networking team needs to centrally manage a list of internal IP address ranges that belong to the global offices Developers Will reference this list to gain access to applications securely.
Which solution meets these requirements with the LEAST amount of operational overhead?

A. Create a JSON file that is hosted in Amazon S3 and that lists all of the internal IP address ranges Configure an Amazon Simple Notification Service (Amazon

SNS) topic in each of the accounts that can be involved when the JSON file is update

B. Subscribe an AWS Lambda function to the SNS topic to update all relevant security group rules with Vie updated IP address ranges.

C. Create a new AWS Config managed rule that contains all of the internal IP address ranges Use the rule to check the security groups in each of the accounts to ensure compliance with the list of IP address range

D. Configure the rule to automatically remediate any noncompliant security group that is detected.

E. In the transit account, create a VPC prefix list with all of the internal IP address range

F. Use AWS Resource Access Manager to share the prefix list with all of the other account

G. Use the shared prefix list to configure security group rules is the other accounts.

H. In the transit account create a security group with all of the internal IP address range

I. Configure the security groups in me other accounts to reference the transit account's securitygroup by using a nested security group reference of *<transit-account-id>./sg-1a2b3c4d".

Answer: C

Explanation:
Customer-managed prefix lists — Sets of IP address ranges that you define and manage. You can share your prefix list with other AWS accounts, enabling those accounts to reference the prefix list in their own resources. https://docs.aws.amazon.com/vpc/latest/userguide/managed-prefix-lists.html
a VPC prefix list is created in the transit account with all of the internal IP address ranges, and then shared to all of the other accounts using AWS Resource Access Manager. This allows for central management of the IP address ranges, and eliminates the need for manual updates to security group rules in each account. This solution also allows for compliance checks to be run using AWS Config and for any non-compliant security groups to be automatically remediated.

NEW QUESTION 37
- (Exam Topic 1)
A solutions architect has developed a web application that uses an Amazon API Gateway Regional endpoint and an AWS Lambda function. The consumers of the web application are all close to the AWS Region where the application will be deployed. The Lambda function only queries an Amazon Aurora MySQL database. The solutions architect has configured the database to have three read replicas.
During testing, the application does not meet performance requirements. Under high load, the application opens a large number of database connections. The solutions architect must improve the application's performance.
Which actions should the solutions architect take to meet these requirements? (Choose two.)

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

Answer: BD

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

NEW QUESTION 38
- (Exam Topic 1)
A company is planning to migrate its business-critical applications from an on-premises data center to AWS. The company has an on-premises installation of a Microsoft SQL Server Always On cluster. The company wants to migrate to an AWS managed database service. A solutions architect must design a heterogeneous database migration on AWS.
Which solution will meet these requirements?

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

Answer: C

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

NEW QUESTION 40
- (Exam Topic 1)
An enterprise company wants to allow its developers to purchase third-party software through AWS Marketplace. The company uses an AWS Organizations account structure with full features enabled, and has a shared services account in each organizational unit (OU) that will be used by procurement managers. The procurement team's policy indicates that developers should be able to obtain third-party software from an approved list only and use Private Marketplace in AWS Marketplace to achieve this requirement . The procurement team wants administration of Private Marketplace to be restricted to a role named procurement-manager-role, which could be assumed by procurement managers Other IAM users groups, roles, and account administrators in the company should be denied Private Marketplace administrative access
What is the MOST efficient way to design an architecture to meet these requirements?

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

Answer: C

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

NEW QUESTION 44
- (Exam Topic 1)
A company has developed APIs that use Amazon API Gateway with Regional endpoints. The APIs call AWS Lambda functions that use API Gateway authentication mechanisms. After a design review, a solutions architect identifies a set of APIs that do not require public access.
The solutions architect must design a solution to make the set of APIs accessible only from a VPC. All APIs need to be called with an authenticated user.
Which solution will meet these requirements with the LEAST amount of effort?

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

Answer: C

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

NEW QUESTION 49
- (Exam Topic 1)
A company has created an OU in AWS Organizations for each of its engineering teams Each OU owns multiple AWS accounts. The organization has hundreds of AWS accounts A solutions architect must design a solution so that each OU can view a breakdown of usage costs across its AWS accounts. Which solution meets these requirements?

A. Create an AWS Cost and Usage Report (CUR) for each OU by using AWS Resource Access Manager Allow each team to visualize the CUR through an Amazon QuickSight dashboard.
B. Create an AWS Cost and Usage Report (CUR) from the AWS Organizations management account- Allow each team to visualize the CUR through an Amazon QuickSight dashboard
C. Create an AWS Cost and Usage Report (CUR) in each AWS Organizations member account Allow each team to visualize the CUR through an Amazon QuickSight dashboard.
D. Create an AWS Cost and Usage Report (CUR) by using AWS Systems Manager Allow each team to visualize the CUR through Systems Manager OpsCenter dashboards

Answer: B

Explanation:
https://docs.aws.amazon.com/cur/latest/userguide/billing-cur-limits.html

NEW QUESTION 51
- (Exam Topic 1)
A security engineer determined that an existing application retrieves credentials to an Amazon RDS for MySQL database from an encrypted file in Amazon S3. For the next version of the application, the security engineer wants to implement the following application design changes to improve security:
The database must use strong, randomly generated passwords stored in a secure AWS managed service.

⟫ The application resources must be deployed through AWS CloudFormation.

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

A solutions architect will generate a CloudFormation template to deploy the application.

Which resources specified in the CloudFormation template will meet the security engineer's requirements with the LEAST amount of operational overhead?

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

Answer: B

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

NEW QUESTION 54
- (Exam Topic 1)
A company gives users the ability to upload images from a custom application. The upload process invokes an AWS Lambda function that processes and stores the image in an Amazon S3 bucket. The application invokes the Lambda function by using a specific function version ARN.

The Lambda function accepts image processing parameters by using environment variables. The company often adjusts the environment variables of the Lambda function to achieve optimal image processing output. The company tests different parameters and publishes a new function version with the updated environment variables after validating results. This update process also requires frequent changes to the custom application to invoke the new function version ARN. These changes cause interruptions for users.

A solutions architect needs to simplify this process to minimize disruption to users. Which solution will meet these requirements with the LEAST operational overhead?

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

Answer: D

Explanation:
A Lambda function alias allows you to point to a specific version of a function and also can be updated to point to a new version of the function without modifying the client application. This way, the company can test different versions of the function with different environment variables and, once the optimal parameters are found, update the alias to point to the new version, without the need to update the client application.

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

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

NEW QUESTION 57
- (Exam Topic 1)
A retail company is hosting an ecommerce website on AWS across multiple AWS Regions. The company wants the website to be operational at all times for online purchases. The website stores data in an Amazon RDS for MySQL DB instance.

Which solution will provide the HIGHEST availability for the database?

A. Configure automated backups on Amazon RD
B. In the case of disruption, promote an automated backup to be a standalone DB instanc
C. Direct database traffic to the promoted DB instanc
D. Create a replacement read replica that has the promoted DB instance as its source.
E. Configure global tables and read replicas on Amazon RD
F. Activate the cross-Region scop
G. In the case of disruption, use AWS Lambda to copy the read replicas from one Region to another Region.
H. Configure global tables and automated backups on Amazon RD
I. In the case of disruption, use AWS Lambda to copy the read replicas from one Region to another Region.
J. Configure read replicas on Amazon RD
K. In the case of disruption, promote a cross-Region and read replica to be a standalone DB instanc
L. Direct database traffic to the promoted DB instanc
M. Create areplacement read replica that has the promoted DB instance as its source.

Answer: D

Explanation:
This solution will provide the highest availability for the database, as the read replicas will allow the database to be available in multiple Regions, thus reducing the chances of disruption. Additionally, the promotion of the cross-Region read replica to become a standalone DB instance will ensure that the database is still available even if one of the Regions experiences disruptions.

NEW QUESTION 60
- (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 63
- (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 65
- (Exam Topic 1)
A company plans to refactor a monolithic application into a modern application designed deployed or AWS. The CLCD pipeline needs to be upgraded to support the modem design for the application with the following requirements
• It should allow changes to be released several times every hour.
* It should be able to roll back the changes as quickly as possible Which design will meet these requirements?

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

Answer: B

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

NEW QUESTION 70

- (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 71
- (Exam Topic 1)
A company that uses AWS Organizations allows developers to experiment on AWS. As part of the landing zone that the company has deployed, developers use their company email address to request an account. The company wants to ensure that developers are not launching costly services or running services unnecessarily. The company must give developers a fixed monthly budget to limit their AWS costs.
Which combination of steps will meet these requirements? (Choose three.)

A. Create an SCP to set a fixed monthly account usage limi
B. Apply the SCP to the developer accounts.
C. Use AWS Budgets to create a fixed monthly budget for each developer's account as part of the account creation process.
D. Create an SCP to deny access to costly services and component
E. Apply the SCP to the developer accounts.
F. Create an IAM policy to deny access to costly services and component
G. Apply the IAM policy to the developer accounts.
H. Create an AWS Budgets alert action to terminate services when the budgeted amount is reached.Configure the action to terminate all services.
I. Create an AWS Budgets alert action to send an Amazon Simple Notification Service (Amazon SNS) notification when the budgeted amount is reache
J. Invoke an AWS Lambda function to terminate all services.

Answer: BCF

Explanation:
▷ Option A is incorrect because creating an SCP to set a fixed monthly account usage limit is not possible.
SCPs are policies that specify the services and actions that users and roles can use in the member accounts of an AWS Organization. SCPs cannot enforce budget limits or prevent users from launching
costly services or running services unnecessarily1

▷ Option B is correct because using AWS Budgets to create a fixed monthly budget for each developer's account as part of the account creation process meets the requirement of giving developers a fixed monthly budget to limit their AWS costs. AWS Budgets allows you to plan your service usage, service costs, and instance reservations. You can create budgets that alert you when your costs or usage exceed (or are forecasted to exceed) your budgeted amount2

▷ Option C is correct because creating an SCP to deny access to costly services and components meets the requirement of ensuring that developers are not launching costly services or running services
unnecessarily. SCPs can restrict access to certain AWS services or actions based on conditions such as region, resource tags, or request time. For example, an SCP can deny access to Amazon Redshift clusters or Amazon EC2 instances with certain instance types1

▷ Option D is incorrect because creating an IAM policy to deny access to costly services and components is not sufficient to meet the requirement of ensuring that developers are not launching costly services or running services unnecessarily. IAM policies can only control access to resources within a single AWS account. If developers have multiple accounts or can create new accounts, they can bypass the IAM policy restrictions. SCPs can apply across multiple accounts within an AWS Organization and prevent users from creating new accounts that do not comply with the SCP rules3

▷ Option E is incorrect because creating an AWS Budgets alert action to terminate services when the budgeted amount is reached is not possible. AWS Budgets alert actions can only perform one of the following actions: apply an IAM policy, apply an SCP, or send a notification through Amazon SNS. AWS Budgets alert actions cannot terminate services directly.

▷ Option F is correct because creating an AWS Budgets alert action to send an Amazon SNS notification when the budgeted amount is reached and invoking an AWS Lambda function to terminate all services meets the requirement of giving developers a fixed monthly budget to limit their AWS costs. AWS Budgets alert actions can send notifications through Amazon SNS when a budget threshold is breached. Amazon SNS can trigger an AWS Lambda function that can perform custom logic such as terminating all services in the developer's account. This way, developers cannot exceed their budget limit and incur additional costs.
References: 1: https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps.html 2
: https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/budgets-create.html 3: https://docs.aws.amazon.com/IAM/latest/UserGuide/introduction.html :
https://docs.aws.amazon.com/cost-management/latest/userguide/budgets-actions.html : https://docs.aws.amazon.com/sns/latest/dg/sns-lambda.html :
https://docs.aws.amazon.com/lambda/latest/dg/welcome.html

NEW QUESTION 72
- (Exam Topic 1)
A finance company is running its business-critical application on current-generation Linux EC2 instances The application includes a self-managed MySQL database performing heavy I/O operations. The application is working fine to handle a moderate amount of traffic during the month. However, it slows down during the final three days of each month due to month-end reporting, even though the company is using Elastic Load Balancers and Auto Scaling within its infrastructure to meet the increased demand.
Which of the following actions would allow the database to handle the month-end load with the LEAST impact on performance?

A. Pre-warming Elastic Load Balancers, using a bigger instance type, changing all Amazon EBS volumes to GP2 volumes.
B. Performing a one-time migration of the database cluster to Amazon RD
C. and creating several additional read replicas to handle the load during end of month
D. Using Amazon CioudWatch with AWS Lambda to change the typ
E. size, or IOPS of Amazon EBS volumes in the cluster based on a specific CloudWatch metric
F. Replacing all existing Amazon EBS volumes with new PIOPS volumes that have the maximum available storage size and I/O per second by taking snapshots before the end of the month and reverting back afterwards.

Answer: B

Explanation:
In this scenario, the Amazon EC2 instances are in an Auto Scaling group already which means that the database read operations is the possible bottleneck especially during the month-end wherein the reports are generated. This can be solved by creating RDS read replicas.

NEW QUESTION 74
- (Exam Topic 1)
A company runs a Python script on an Amazon EC2 instance to process data. The script runs every
10 minutes. The script ingests files from an Amazon S3 bucket and processes the files. On average, the script takes approximately 5 minutes to process each file The script will not reprocess a file that the script has already processed.
The company reviewed Amazon CloudWatch metrics and noticed that the EC2 instance is idle for approximately 40% of the time because of the file processing speed. The company wants to make the workload highly available and scalable. The company also wants to reduce long-term management overhead.
Which solution will meet these requirements MOST cost-effectively?

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

Answer: D

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

NEW QUESTION 79
- (Exam Topic 1)
A company is in the process of implementing AWS Organizations to constrain its developers to use only Amazon EC2. Amazon S3 and Amazon DynamoDB. The developers account resides In a dedicated organizational unit (OU). The solutions architect has implemented the following SCP on the developers account:

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

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

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

Answer: B

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

NEW QUESTION 80
- (Exam Topic 1)
A financial services company in North America plans to release a new online web application to its customers on AWS . The company will launch the application in the us-east-1 Region on Amazon EC2 instances. The application must be highly available and must dynamically scale to meet user traffic. The company also wants to implement a disaster recovery environment for the application in the us-west-1 Region by using active-passive failover.
Which solution will meet these requirements?

A. Create a VPC in us-east-1 and a VPC in us-west-1 Configure VPC peering In the us-east-1 VP
B. create an Application Load Balancer (ALB) that extends across multiple Availability Zones in both VPCs Create an Auto Scaling group that deploys the EC2 instances across the multiple Availability Zones in both VPCs Place the Auto Scaling group behind the ALB.
C. Create a VPC in us-east-1 and a VPC in us-west-1. In the us-east-1 VP
D. create an Application Load Balancer (ALB) that extends across multiple Availability Zones in that VP
E. Create an Auto Scaling group that deploys the EC2 instances across the multiple Availability Zones in the us-east-1 VPC Place the Auto Scaling group behind the ALB Set up the same configuration in the us-west-1 VP
F. Create an Amazon Route 53 hosted zone Create separate records for each ALB Enable health checks to ensure high availability between Regions.
G. Create a VPC in us-east-1 and a VPC in us-west-1 In the us-east-1 VP
H. create an Application Load Balancer (ALB) that extends across multiple Availability Zones in that VPC Create an Auto Scaling group that deploys the EC2 instances across the multiple Availability Zones in the us-east-1 VPC Place the Auto Scaling group behind the ALB Set up the same configuration in the us-west-1 VPC Create an Amazon Route 53 hosted zon
I. Create separate records for each ALB Enable health checks and configure a failover routing policy for each record.
J. Create a VPC in us-east-1 and a VPC in us-west-1 Configure VPC peering In the us-east-1 VP
K. create an Application Load Balancer (ALB) that extends across multiple Availability Zones in Create an Auto Scaling group that deploys the EC2 instances across the multiple Availability Zones in both VPCs Place the Auto Scaling group behind the ALB Create an Amazon Route 53 host.. Create a record for the ALB.

Answer: C

Explanation:
it's the one that handles failover while B (the one shown as the answer today) it almost the same but does not handle failover.

NEW QUESTION 83
- (Exam Topic 1)
A company is using multiple AWS accounts The DNS records are stored in a private hosted zone for Amazon Route 53 in Account A The company's applications and databases are running in Account B.
A solutions architect win deploy a two-net application In a new VPC To simplify the configuration, the db.example com CNAME record set tor the Amazon RDS endpoint was created in a private hosted zone for Amazon Route 53.
During deployment, the application failed to start. Troubleshooting revealed that db.example com is not resolvable on the Amazon EC2 instance The solutions architect confirmed that the record set was created correctly in Route 53.
Which combination of steps should the solutions architect take to resolve this issue? (Select TWO)

A. Deploy the database on a separate EC2 instance in the new VPC Create a record set for the instance's private IP in the private hosted zone
B. Use SSH to connect to the application tier EC2 instance Add an RDS endpoint IP address to the/eto/resolv.conf file
C. Create an authorization lo associate the private hosted zone in Account A with the new VPC In Account B
D. Create a private hosted zone for the example.com domain m Account B Configure Route 53 replication between AWS accounts
E. Associate a new VPC in Account B with a hosted zone in Account
F. Delete the association authorization In Account A.

Answer: CE

Explanation:
https://aws.amazon.com/premiumsupport/knowledge-center/private-hosted-zone-different-account/

NEW QUESTION 87
- (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 91
- (Exam Topic 1)
A company is building a software-as-a-service (SaaS) solution on AWS. The company has deployed an Amazon API Gateway REST API with AWS Lambda integration in multiple AWS Regions and in the same production account.
The company offers tiered pricing that gives customers the ability to pay for the capacity to make a certain number of API calls per second. The premium tier offers up to 3,000 calls per second, and customers are identified by a unique API key. Several premium tier customers in various Regions report that they receive error responses of 429 Too Many Requests from multiple API methods during peak usage hours. Logs indicate that the Lambda function is never invoked.
What could be the cause of the error messages for these customers?

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

Answer: C

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

NEW QUESTION 95
- (Exam Topic 1)
A company is running a web application in the AWS Cloud. The application consists of dynamic content that is created on a set of Amazon EC2 instances. The EC2 instances run in an Auto Scaling group that is configured as a target group for an Application Load Balancer (ALB).
The company is using an Amazon CloudFront distribution to distribute the application globally. The CloudFront distribution uses the ALB as an origin. The company uses Amazon Route 53 for DNS and has created an A record of www.example.com for the CloudFront distribution.
A solutions architect must configure the application so that it is 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 96
- (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 99
- (Exam Topic 1)
A company has purchased appliances from different vendors. The appliances all have IoT sensors. The sensors send status information in the vendors' proprietary formats to a legacy application that parses the information into JSON. The parsing is simple, but each vendor has a unique format. Once daily, the application parses all the JSON records and stores the records in a relational database for analysis.
The company needs to design a new data analysis solution that can deliver faster and optimize costs. Which solution will meet these requirements?

A. Connect the IoT sensors to AWS IoT Cor
B. Set a rule to invoke an AWS Lambda function to parse the information and save a .csv file to Amazon S3. Use AWS Glue to catalog the file
C. Use Amazon Athena and Amazon QuickSight for analysis.
D. Migrate the application server to AWS Fargate, which will receive the information from IoT sensors and parse the information into a relational forma
E. Save the parsed information to Amazon Redshift for analysis.
F. Create an AWS Transfer for SFTP serve
G. Update the IoT sensor code to send the information as a .csv file through SFTP to the serve
H. Use AWS Glue to catalog the file
I. Use Amazon Athena for analysis.
J. Use AWS Snowball Edge to collect data from the IoT sensors directly to perform local analysis.Periodically collect the data into Amazon Redshift to perform global analysis.

Answer: A

Explanation:
➤ Connect the IoT sensors to AWS IoT Core. Set a rule to invoke an AWS Lambda function to parse the information and save a .csv file to Amazon S3. Use AWS Glue to catalog the files. Use Amazon Athena and Amazon QuickSight for analysis. This solution meets the requirement of faster analysis and cost optimization by using AWS IoT Core to collect data from the IoT sensors in real-time and then using AWS Glue and Amazon Athena for efficient data analysis. This solution involves connecting the IoT sensors to the AWS IoT Core, setting a rule to invoke an AWS Lambda function to parse the information, and saving a .csv file to Amazon S3. AWS Glue can be used to catalog the files and Amazon Athena and Amazon QuickSight can be used for analysis. This solution will enable faster and more cost-effective data analysis.
This solution is in line with the official Amazon Textbook and Resources for the AWS Certified Solutions Architect - Professional certification. In particular, the book states that: "AWS IoT Core can be used to ingest and process the data, AWS Lambda can be used to process and transform the data, and Amazon S3 can be used to store the data. AWS Glue can be used to catalog and access the data, Amazon Athena can be used to query the data, and Amazon QuickSight can be used to visualize the data." (Source: https://d1.awsstatic.com/training-and-certification/docs-sa-pro/AWS_Certified_Solutions_Architect_Professiona

NEW QUESTION 101
- (Exam Topic 1)
A company is processing videos in the AWS Cloud by using Amazon EC2 instances in an Auto Scaling group. It takes 30 minutes to process a video. Several EC2 instances scale in and out depending on the number of videos in an Amazon Simple Queue Service (Amazon SQS) queue.
The company has configured the SQS queue with a redrive policy that specifies a target dead-letter queue and a maxReceiveCount of 1. The company has set the visibility timeout for the SQS queue to 1 hour. The company has set up an Amazon CloudWatch alarm to notify the development team when there are messages in the dead-letter queue.
Several times during the day, the development team receives notification that messages are in the dead-letter queue and that videos have not been processed properly. An investigation finds no errors in the application logs.
How can the company solve this problem?

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

Answer: B

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

NEW QUESTION 106
- (Exam Topic 1)
A retail company is operating its ecommerce application on AWS. The application runs on Amazon EC2 instances behind an Application Load Balancer (ALB). The company uses an Amazon RDS DB instance as the database backend. Amazon CloudFront is configured with one origin that points to the ALB. Static content is cached. Amazon Route 53 is used to host all public zones.
After an update of the application, the ALB occasionally returns a 502 status code (Bad Gateway) error. The root cause is malformed HTTP headers that are returned to the ALB. The webpage returns successfully when a solutions architect reloads the webpage immediately after the error occurs.
While the company is working on the problem, the solutions architect needs to provide a custom error page instead of the standard ALB error page to visitors.
Which combination of steps will meet this requirement with the LEAST amount of operational overhead? (Choose two.)

A. Create an Amazon S3 bucke
B. Configure the S3 bucket to host a static webpag
C. Upload the custom error pages to Amazon S3.
D. Create an Amazon CloudWatch alarm to invoke an AWS Lambda function if the ALB health check response Target.FailedHealthChecks is greater than 0. Configure the Lambda function to modify the forwarding rule at the ALB to point to a publicly accessible web server.
E. Modify the existing Amazon Route 53 records by adding health check
F. Configure a fallback target if the health check fail
G. Modify DNS records to point to a publicly accessible webpage.
H. Create an Amazon CloudWatch alarm to invoke an AWS Lambda function if the ALB health check response Elb.InternalError is greater than 0. Configure the

Lambda function to modify the forwarding rule at the ALB to point to a public accessible web server.
I. Add a custom error response by configuring a CloudFront custom error pag
J. Modify DNS records to point to a publicly accessible web page.

Answer: CE

Explanation:
"Save your custom error pages in a location that is accessible to CloudFront. We recommend that you store them in an Amazon S3 bucket, and that you don't store them in the same place as the rest of your website or application's content. If you store the custom error pages on the same origin as your website or application, and the origin starts to return 5xx errors, CloudFront can't get the custom error pages because the origin server is unavailable."
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/GeneratingCustomErrorResponses.htm

NEW QUESTION 107
- (Exam Topic 1)
A company is planning to migrate 1,000 on-premises servers to AWS. The servers run on several VMware clusters in the company's data center. As part of the migration plan, the company wants to gather server metrics such as CPU details, RAM usage, operating system information, and running processes. The company then wants to query and analyze the data.
Which solution will meet these requirements?

A. Deploy and configure the AWS Agentless Discovery Connector virtual appliance on the on-premises host
B. Configure Data Exploration in AWS Migration Hu
C. Use AWS Glue to perform an ETL job against the dat
D. Query the data by using Amazon S3 Select.
E. Export only the VM performance information from the on-premises host
F. Directly import the required data into AWS Migration Hu
G. Update any missing information in Migration Hu
H. Query the data by using Amazon QuickSight.
I. Create a script to automatically gather the server information from the on-premises host
J. Use the AWS CLI to run the put-resource-attributes command to store the detailed server data in AWS Migration Hu
K. Query the data directly in the Migration Hub console.
L. Deploy the AWS Application Discovery Agent to each on-premises serve
M. Configure Data Exploration in AWS Migration Hu
N. Use Amazon Athena to run predefined queries against the data in Amazon S3.

Answer: D

Explanation:
it covers all the requirements mentioned in the question, it will allow collecting the detailed metrics, including process information and it provides a way to query and analyze the data using Amazon Athena.

NEW QUESTION 111
- (Exam Topic 1)
A financial company is planning to migrate its web application from on premises to AWS. The company uses a third-party security tool to monitor the inbound traffic to the application. The company has used the security tool for the last 15 years, and the tool has no cloud solutions available from its vendor. The company's security team is concerned about how to integrate the security tool with AWS technology.
The company plans to deploy the application migration to AWS on Amazon EC2 instances. The EC2 instances will run in an Auto Scaling group in a dedicated VPC. The company needs to use the security tool to inspect all packets that come in and out of the VPC. This inspection must occur in real time and must not affect the application's performance. A solutions architect must design a target architecture on AWS that is highly available within an AWS Region.
Which combination of steps should the solutions architect take to meet these requirements? (Select TWO.)

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

Answer: AD

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

NEW QUESTION 116
- (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 121
- (Exam Topic 1)
A digital marketing company has multiple AWS accounts that belong to various teams. The creative team uses an Amazon S3 bucket in its AWS account to securely store images and media files that are used as content for the company's marketing campaigns. The creative team wants to share the S3 bucket with the strategy team so that the strategy team can view the objects.
A solutions architect has created an IAM role that is named strategy_reviewer in the Strategy account. The solutions architect also has set up a custom AWS Key Management Service (AWS KMS) key in the Creative account and has associated the key with the S3 bucket. However, when users from the Strategy account assume the IAM role and try to access objects in the S3 bucket, they receive an Account.
The solutions architect must ensure that users in the Strategy account can access the S3 bucket. The solution must provide these users with only the minimum permissions that they need.
Which combination of steps should the solutions architect take to meet these requirements? (Select THREE.)

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

Answer: ACF

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

NEW QUESTION 122
- (Exam Topic 1)
A company is planning to store a large number of archived documents and make the documents available to employees through the corporate intranet. Employees will access the system by connecting through a client VPN service that is attached to a VPC. The data must not be accessible to the public.
The documents that the company is storing are copies of data that is held on physical media elsewhere. The number of requests will be low. Availability and speed of retrieval are not concerns of the company.
Which solution will meet these requirements at the LOWEST cost?

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

Answer: D

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

NEW QUESTION 123
- (Exam Topic 1)
The company needs to determine which costs on the monthly AWS bill are attributable to each application or team. The company also must be able to create reports to compare costs from the last 12 months and to help forecast costs for the next 12 months. A solutions architect must recommend an AWS Billing and Cost Management solution that provides these cost reports.
Which combination of actions will meet these requirements? (Select THREE.)

A. Activate the user-defined cost allocation tags that represent the application and the team.

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

Answer: ACF

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

NEW QUESTION 124
- (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 128
- (Exam Topic 1)
A weather service provides high-resolution weather maps from a web application hosted on AWS in the
eu-west-1 Region. The weather maps are updated frequently and stored in Amazon S3 along with static HTML content. The web application is fronted by Amazon CloudFront.
The company recently expanded to serve users in the us-east-1 Region, and these new users report that viewing their respective weather maps is slow from time to time.
Which combination of steps will resolve the us-east-1 performance issues? (Choose two.)

A. Configure the AWS Global Accelerator endpoint for the S3 bucket in eu-west-1. Configure endpoint groups for TCP ports 80 and 443 in us-east-1.
B. Create a new S3 bucket in us-east-1. Configure S3 cross-Region replication to synchronize from the S3 bucket in eu-west-1.
C. Use Lambda@Edge to modify requests from North America to use the S3 Transfer Acceleration endpoint in us-east-1.
D. Use Lambda@Edge to modify requests from North America to use the S3 bucket in us-east-1.
E. Configure the AWS Global Accelerator endpoint for us-east-1 as an origin on the CloudFront distributio
F. Use Lambda@Edge to modify requests from North America to use the new origin.

Answer: BD

Explanation:
https://aws.amazon.com/about-aws/whats-new/2016/04/transfer-files-into-amazon-s3-up-to-300-percent-faster/

NEW QUESTION 130
- (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 133
- (Exam Topic 1)
A team collects and routes behavioral data for an entire company The company runs a Multi-AZ VPC
environment with public subnets, private subnets, and in internet gateway Each public subnet also contains a NAT gateway Most of the company's applications read from and write to Amazon Kinesis Data Streams. Most of the workloads am in private subnets.
A solutions architect must review the infrastructure The solutions architect needs to reduce costs and maintain the function of the applications The solutions architect uses Cost Explorer and notices that the cost in the EC2-Other category is consistently high A further review shows that NatGateway-Bytes charges are increasing the cost in the EC2-Other category.
What should the solutions architect do to meet these requirements?

A. Enable VPC Flow Log
B. Use Amazon Athena to analyze the logs for traffic that can be remove
C. Ensure that security groups are Mocking traffic that is responsible for high costs.
D. Add an interface VPC endpoint for Kinesis Data Streams to the VP
E. Ensure that applications have the correct IAM permissions to use the interface VPC endpoint.
F. Enable VPC Flow Logs and Amazon Detective Review Detective findings for traffic that is not related to Kinesis Data Streams Configure security groups to block that traffic
G. Add an interface VPC endpoint for Kinesis Data Streams to the VP
H. Ensure that the VPC endpoint policy allows traffic from the applications.

Answer: D

Explanation:
https://docs.aws.amazon.com/vpc/latest/privatelink/vpc-endpoints-access.html https://aws.amazon.com/premiumsupport/knowledge-center/vpc-reduce-nat-gateway-transfer-costs/
VPC endpoint policies enable you to control access by either attaching a policy to a VPC endpoint or by using additional fields in a policy that is attached to an IAM user, group, or role to restrict access to only occur via the specified VPC endpoint

NEW QUESTION 135
- (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 137
- (Exam Topic 1)
A company has an environment that has a single AWS account. A solutions architect is reviewing the environment to recommend what the company could improve

specifically in terms of access to the AWS Management Console. The company's IT support workers currently access the console for administrative tasks, authenticating with named IAM users that have been mapped to their job role.

The IT support workers no longer want to maintain both their Active Directory and IAM user accounts. They want to be able to access the console by using their existing Active Directory credentials. The solutions architect is using AWS Single Sign-On (AWS SSO) to implement this functionality.

Which solution will meet these requirements MOST cost-effectively?

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

Answer: D

Explanation:
https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_org_support-all-features.html
https://docs.aws.amazon.com/singlesignon/latest/userguide/get-started-prereqs-considerations.html

NEW QUESTION 138
......

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual SAP-C02 Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the SAP-C02 Product From:

https://www.2passeasy.com/dumps/SAP-C02/

Money Back Guarantee

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