Amazon-Web-Services

Exam Questions SAP-C02

AWS Certified Solutions Architect - Professional

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 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 3
- (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 4
- (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 5
- (Exam Topic 1)
A company is running a traditional web application on Amazon EC2 instances. The company needs to refactor the application as microservices that run on containers. Separate versions of the application exist in two distinct environments: production and testing. Load for the application is variable, but the minimum load and the maximum load are known. A solutions architect needs to design the updated application with a serverless architecture that minimizes operational complexity.
Which solution will meet these requirements MOST cost-effectively?

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

Answer: B

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

NEW QUESTION 6
- (Exam Topic 1)
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 7
- (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 8
- (Exam Topic 1)
A company uses Amazon S3 to store files and images in a variety of storage classes. The company's S3 costs have increased substantially during the past year.
A solutions architect needs to review data trends for the past 12 months and identity the appropriate storage class for the objects.
Which solution will meet these requirements?

A. Download AWS Cost and Usage Reports for the last 12 months of S3 usag
B. Review AWS Trusted Advisor recommendations for cost savings.
C. Use S3 storage class analysi
D. Import data trends into an Amazon QuickSight dashboard to analyze storage trends.
E. Use Amazon S3 Storage Len
F. Upgrade the default dashboard to include advanced metrics for storage trends.
G. Use Access Analyzer for S3. Download the Access Analyzer for S3 report for the last 12 month
H. Import the csvfile to an Amazon QuickSight dashboard.

Answer: B

Explanation:
https://docs.aws.amazon.com/AmazonS3/latest/userguide/storage_lens.html

NEW QUESTION 9
- (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 10
- (Exam Topic 1)
A company has developed a web application. The company is hosting the application on a group of Amazon EC2 instances behind an Application Load Balancer. The company wants to improve the security posture of the application and plans to use AWS WAF web ACLs. The solution must not adversely affect legitimate traffic to the application.
How should a solutions architect configure the web ACLs to meet these requirements?

A. Set the action of the web ACL rules to Coun
B. Enable AWS WAF logging Analyze the requests for false positives Modify the rules to avoid any false positive Over time change the action of the web ACL rules from Count to Block.
C. Use only rate-based rules in the web ACL
D. and set the throttle limit as high as possible Temporarily block all requests that exceed the limi
E. Define nested rules to narrow the scope of the rate tracking.
F. Set the action o' the web ACL rules to Bloc
G. Use only AWS managed rule groups in the web ACLs Evaluate the rule groups by using Amazon CloudWatch metrics with AWS WAF sampled requests or AWS WAF logs.
H. Use only custom rule groups in the web ACL
I. and set the action to Allow Enable AWS WAF logging Analyze the requests tor false positives Modify the rules to avoid any false positive Over time, change the action of the web ACL rules from Allow to Block.

Answer: A

Explanation:
https://aws.amazon.com/premiumsupport/knowledge-center/waf-analyze-count-action-rules/

NEW QUESTION 10
- (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 12
- (Exam Topic 1)
A company's solutions architect is reviewing a new internally developed application in a sandbox AWS account The application uses an AWS Auto Scaling group of Amazon EC2 instances that have an IAM instance profile attached Part of the application logic creates and accesses secrets from AWS Secrets Manager The company has an AWS Lambda function that calls the application API to test the functionality The company also has created an AWS CloudTrail trail in the account The application's developer has attached the SecretsManagerReadWnte AWS managed IAM policy to an IAM role The IAM role is associated with the instance profile that is attached to the EC2 instances The solutions architect has invoked the Lambda function for testing
The solutions architect must replace the SecretsManagerReadWnte policy with a new policy that provides least privilege access to the Secrets Manager actions that the application requires
What is the MOST operationally efficient solution that meets these requirements?

A. Generate a policy based on CloudTrail events for the IAM role Use the generated policy output to create a new IAM policy Use the newly generated IAM policy to replace the SecretsManagerReadWnte policy that is attached to the IAM role
B. Create an analyzer in AWS Identity and Access Management Access Analyzer Use the IAM role's Access Advisor findings to create a new IAM policy Use the newly created IAM policy to replace the SecretsManagerReadWnte policy that is attached to the IAM role
C. Use the aws cloudtrail lookup-events AWS CLI command to filter and export CloudTrail events that are related to Secrets Manager Use a new IAM policy that contains the actions from CloudTrail to replace the SecretsManagerReadWnte policy that is attached to the IAM role
D. Use the IAM policy simulator to generate an IAM policy for the IAM role Use the newly generated IAM policy to replace the SecretsManagerReadWnte policy that is attached to the IAM role

Answer: B

Explanation:
The IAM policy simulator will generate a policy that contains only the necessary permissions for the application to access Secrets Manager, providing the least privilege necessary to get the job done. This is the most efficient solution as it will not require additional steps such as analyzing CloudTrail events or manually creating and testing an IAM policy.
You can use the IAM policy simulator to generate an IAM policy for an IAM role by specifying the role and the API actions and resources that the application or service requires. The simulator will then generate an IAM policy that grants the least privilege access to those actions and resources.
Once you have generated an IAM policy using the simulator, you can replace the existing SecretsManagerReadWnte policy that is attached to the IAM role with the newly generated policy. This will ensure that the application or service has the least privilege access to the Secrets Manager actions that it requires.
You can access the IAM policy simulator through the IAM console, AWS CLI, and AWS SDKs. Here is the link for more information:
https://docs.aws.amazon.com/IAM/latest/UserGuide/access_policies_simulator.html

NEW QUESTION 16
- (Exam Topic 1)
A company is creating a sequel for a popular online game. A large number of users from all over the world will play the game within the first week after launch. Currently, the game consists of the following components deployed in a single AWS Region:
• Amazon S3 bucket that stores game assets
• Amazon DynamoDB table that stores player scores
A solutions architect needs to design a multi-Region solution that will reduce latency improve reliability, and require the least effort to implement
What should the solutions architect do to meet these requirements?

A. Create an Amazon CloudFront distribution to serve assets from the S3 bucket Configure S3Cross-Region Replication Create a new DynamoDB able in a new Region Use the new table as a replica target tor DynamoDB global tables.
B. Create an Amazon CloudFront distribution to serve assets from the S3 bucke
C. Configure S3Same-Region Replicatio
D. Create a new DynamoDB able m a new Regio
E. Configure asynchronous replication between the DynamoDB tables by using AWS Database Migration Service (AWS DMS) with change data capture (CDC)

F. Create another S3 bucket in a new Region and configure S3 Cross-Region Replication between the buckets Create an Amazon CloudFront distribution and configure origin failover with two origins accessing the S3 buckets in each Regio
G. Configure DynamoDB global tables by enabling Amazon DynamoDB Streams, and add a replica table in a new Region.
H. Create another S3 bucket in the same Region, and configure S3 Same-Region Replication between the buckets- Create an Amazon CloudFront distribution and configure origin failover with two origin accessing the S3 buckets Create a new DynamoDB table m a new Region Use the new table as a replica target for DynamoDB global tables.

Answer: C

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

NEW QUESTION 19
- (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 21
- (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 22
- (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 24
- (Exam Topic 1)
An AWS customer has a web application that runs on premises. The web application fetches data from a third-party API that is behind a firewall. The third party accepts only one public CIDR block in each client's allow list.
The customer wants to migrate their web application to the AWS Cloud. The application will be hosted on a set of Amazon EC2 instances behind an Application Load Balancer (ALB) in a VPC. The ALB is located in public subnets. The EC2 instances are located in private subnets. NAT gateways provide internet access to the private subnets.
How should a solutions architect ensure that the web application can continue to call the third-parly API after the migration?

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

Answer: B

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

NEW QUESTION 27
- (Exam Topic 1)
A retail company has an on-premises data center in Europe. The company also has a multi-Region AWS presence that includes the eu-west-1 and us-east-1 Regions. The company wants to be able to route network traffic from its on-premises infrastructure into VPCs in either of those Regions. The company also needs to support traffic that is routed directly between VPCs in those Regions. No single points of failure can exist on the network.
The company already has created two 1 Gbps AWS Direct Connect connections from its on-premises data center. Each connection goes into a separate Direct Connect location in Europe for high availability. These two locations are named DX-A and DX-B, respectively. Each Region has a single AWS Transit Gateway that is configured to route all inter-VPC traffic within that Region.
Which solution will meet these requirements?

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

Answer: D

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

NEW QUESTION 28
- (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 29
- (Exam Topic 1)
A company is planning to host a web application on AWS and works to load balance the traffic across a group of Amazon EC2 instances. One of the security requirements is to enable end-to-end encryption in transit between the client and the web server.
Which solution will meet this requirement?

A. Place the EC2 instances behind an Application Load Balancer (ALB) Provision an SSL certificate using AWS Certificate Manager (ACM), and associate the SSL certificate with the AL
B. Export the SSL certificate and install it on each EC2 instanc
C. Configure the ALB to listen on port 443 and to forward traffic to port 443 on the instances.
D. Associate the EC2 instances with a target grou
E. Provision an SSL certificate using AWS Certificate Manager (ACM). Create an Amazon CloudFront distribution and configure It to use the SSL certificat
F. Set CloudFront to use the target group as the origin server
G. Place the EC2 instances behind an Application Load Balancer (ALB). Provision an SSL certificate using AWS Certificate Manager (ACM), and associate the SSL certificate with the AL
H. Provision athird-party SSL certificate and install it on each EC2 instanc
I. Configure the ALB to listen on port 443 and to forward traffic to port 443 on the instances.
J. Place the EC2 instances behind a Network Load Balancer (NLB). Provision a third-party SSL certificate and install it on the NLB and on each EC2 instanc
K. Configure the NLB to listen on port 443 and to forward traffic to port 443 on the instances.

Answer: A

Explanation:
Option A is correct because placing the EC2 instances behind an Application Load Balancer (ALB) and associating an SSL certificate from AWS Certificate Manager (ACM) with the ALB enables encryption in transit between the client and the ALB. Exporting the SSL certificate and installing it on each EC2 instance enables encryption in transit between the ALB and the web server. Configuring the ALB to listen on port 443 and to forward traffic to port 443 on the instances ensures that HTTPS is used for both connections. This solution achieves end-to-end encryption in transit for the web applicatio1n2
References: 1: https://docs.aws.amazon.com/elasticloadbalancing/latest/application/introduction.html 2: https://docs.aws.amazon.com/acm/latest/userguide/acm-overview.html 3: https://docs.aws.amazon.com/elasticloadbalancing/latest/application/load-balancer-target-groups.html : https://aws.amazon.com/certificate-manager/faqs/ : https://docs.aws.amazon.com/elasticloadbalancing/latest/network/introduction.html

NEW QUESTION 34
- (Exam Topic 1)
A company has an asynchronous HTTP application that is hosted as an AWS Lambda function. A public Amazon API Gateway endpoint invokes the Lambda function. The Lambda function and the API Gateway endpoint reside in the us-east-1 Region. A solutions architect needs to redesign the application to support failover to another AWS Region.
Which solution will meet these requirements?

A. Create an API Gateway endpoint in the us-west-2 Region to direct traffic to the Lambda function in us-east-1. Configure Amazon Route 53 to use a failover routing policy to route traffic for the two API Gateway endpoints.
B. Create an Amazon Simple Queue Service (Amazon SQS) queu
C. Configure API Gateway to direct traffic to the SQS queue instead of to the Lambda functio
D. Configure the Lambda function to pull messages from the queue for processing.
E. Deploy the Lambda function to the us-west-2 Regio
F. Create an API Gateway endpoint in us-west-2 to direct traffic to the Lambda function in us-west-2. Configure AWS Global Accelerator and an Application Load Balancer to manage traffic across the two API Gateway endpoints.
G. Deploy the Lambda function and an API Gateway endpoint to the us-west-2 Regio
H. Configure Amazon Route 53 to use a failover routing policy to route traffic for the two API Gateway endpoints.

Answer: B

Explanation:
This solution allows for deploying the Lambda function and API Gateway endpoint to another region, providing a failover option in case of any issues in the primary region. Using Route 53's failover routing policy allows for automatic routing of traffic to the healthy endpoint, ensuring that the application is available even in case of issues in one region. This solution provides a cost-effective and simple way to implement failover while minimizing operational overhead.

NEW QUESTION 35
- (Exam Topic 1)
A video streaming company recently launched a mobile app for video sharing. The app uploads various files to an Amazon S3 bucket in the us-east-1 Region. The files range in size from 1 GB to 10 GB.
Users who access the app from Australia have experienced uploads that take long periods of time Sometimes the files fail to completely upload for these users . A solutions architect must improve the app' performance for these uploads
Which solutions will meet these requirements? (Select TWO.)

A. Enable S3 Transfer Acceleration on the S3 bucket Configure the app to use the Transfer Acceleration endpoint for uploads
B. Configure an S3 bucket in each Region to receive the upload
C. Use S3 Cross-Region Replication to copy the files to the distribution S3 bucket.
D. Set up Amazon Route 53 with latency-based routing to route the uploads to the nearest S3 bucket Region.
E. Configure the app to break the video files into chunks Use a multipart upload to transfer files to Amazon S3.
F. Modify the app to add random prefixes to the files before uploading

Answer: AD

Explanation:
https://aws.amazon.com/premiumsupport/knowledge-center/s3-upload-large-files/
Enabling S3 Transfer Acceleration on the S3 bucket and configuring the app to use the Transfer Acceleration endpoint for uploads will improve the app's performance for these uploads by leveraging Amazon CloudFront's globally distributed edge locations to accelerate the uploads. Breaking the video files into chunks and using a multipart upload to transfer files to Amazon S3 will also improve the app's performance by allowing parts of the file to be uploaded in parallel, reducing the overall upload time.

NEW QUESTION 38
- (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 41
- (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 44
- (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 46
- (Exam Topic 1)
A company has registered 10 new domain names. The company uses the domains for online marketing. The company needs a solution that will redirect online visitors to a specific URL for each domain. All domains and target URLs are defined in a JSON document. All DNS records are managed by Amazon Route 53.
A solutions architect must implement a redirect service that accepts HTTP and HTTPS requests.
Which combination of steps should the solutions architect take to meet these requirements with the LEAST amount of operational effort? (Choose three.)

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

Answer: CEF

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

NEW QUESTION 51
- (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 56
- (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 57
- (Exam Topic 1)
A company is storing data on premises on a Windows file server. The company produces 5 GB of new data daily.
The company migrated part of its Windows-based workload to AWS and needs the data to be available on a file system in the cloud. The company already has established an AWS Direct Connect connection between the on-premises network and AWS.
Which data migration strategy should the company use?

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

Answer: B

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

NEW QUESTION 61
- (Exam Topic 1)
A delivery company needs to migrate its third-party route planning application to AWS. The third party supplies a supported Docker image from a public registry. The image can run in as many containers as required to generate the route map.
The company has divided the delivery area into sections with supply hubs so that delivery drivers travel the shortest distance possible from the hubs to the customers. To reduce the time necessary to generate route maps, each section uses its own set of Docker containers with a custom configuration that processes orders only in the section's area.
The company needs the ability to allocate resources cost-effectively based on the number of running containers.
Which solution will meet these requirements with the LEAST operational overhead?

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

Answer: D

Explanation:
Amazon Elastic Container Service (ECS) on AWS Fargate is a fully managed service that allows you to run containers without having to manage the underlying infrastructure. When you launch tasks on Fargate,
resources are automatically allocated based on the number of tasks running, which reduces the operational overhead.
Using ECS on Fargate allows you to assign custom tags to tasks using the --tags option in the run-task command, as described in the documentation:
https://docs.aws.amazon.com/cli/latest/reference/ecs/run-task.html You can also set enableECSManagedTags to true, which allows the service to automatically add the cluster name and service name as tags.
https://docs.aws.amazon.com/AmazonECS/latest/developerguide/task-placement-constraints.html#tag-based-sch

NEW QUESTION 64

- (Exam Topic 1)
A company is storing data in several Amazon DynamoDB tables. A solutions architect must use a serverless architecture to make the data accessible publicly through a simple API over HTTPS. The solution must scale automatically in response to demand.
Which solutions meet these requirements? (Choose two.)

A. Create an Amazon API Gateway REST AP
B. Configure this API with direct integrations to DynamoDB by using API Gateway's AWS integration type.
C. Create an Amazon API Gateway HTTP AP
D. Configure this API with direct integrations to Dynamo DB by using API Gateway's AWS integration type.
E. Create an Amazon API Gateway HTTP AP
F. Configure this API with integrations to AWS Lambda functions that return data from the DynamoDB tables.
G. Create an accelerator in AWS Global Accelerato
H. Configure this accelerator with AWS Lambda@Edge function integrations that return data from the DynamoDB tables.
I. Create a Network Load Balance
J. Configure listener rules to forward requests to the appropriate AWS Lambda functions

Answer: AC

Explanation:
https://docs.aws.amazon.com/apigateway/latest/developerguide/api-gateway-overview-developer-experience.htm

NEW QUESTION 69
- (Exam Topic 1)
A video processing company has an application that downloads images from an Amazon S3 bucket, processes the images, stores a transformed image in a second S3 bucket, and updates metadata about the image in an Amazon DynamoDB table. The application is written in Node.js and runs by using an AWS Lambda function. The Lambda function is invoked when a new image is uploaded to Amazon S3.
The application ran without incident for a while. However, the size of the images has grown significantly. The Lambda function is now failing frequently with timeout errors. The function timeout is set to its maximum value. A solutions architect needs to refactor the application's architecture to prevent invocation failures. The company does not want to manage the underlying infrastructure.
Which combination of steps should the solutions architect take to meet these requirements? (Choose two.)

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

Answer: AB

Explanation:
A. Modify the application deployment by building a Docker image that contains the application code. Publish the image to Amazon Elastic Container Registry (Amazon ECR). - This step is necessary to package the application code in a container and make it available for running on ECS. B. Create a new Amazon Elastic Container Service (Amazon ECS) task definition with a compatibility type of AWS Fargate. Configure the task definition to use the new image in Amazon Elastic Container Registry (Amazon ECR). Adjust the Lambda function to invoke an ECS task by using the ECS task definition when a new file arrives in Amazon S3.

NEW QUESTION 70
- (Exam Topic 1)
A company has an organization that has many AWS accounts in AWS Organizations. A solutions architect must improve how the company manages common security group rules for the AWS accounts in the organization.
The company has a common set of IP CIDR ranges in an allow list in each AWS account to allow access to
and from the company's on-premises network.
Developers within each account are responsible for adding new IP CIDR ranges to their security groups. The security team has its own AWS account. Currently, the security team notifies the owners of the other AWS accounts when changes are made to the allow list.
The solutions architect must design a solution that distributes the common set of CIDR ranges across all accounts.
Which solution meets these requirements with the LEAST amount of operational overhead?

A. Set up an Amazon Simple Notification Service (Amazon SNS) topic in the security team's AWS accoun
B. Deploy an AWS Lambda function in each AWS accoun
C. Configure the Lambda function to run every time an SNS topic receives a messag
D. Configure the Lambda function to take an IP address as input and add it to a list of security groups in the accoun
E. Instruct the security team to distribute changes by publishing messages to its SNS topic.
F. Create new customer-managed prefix lists in each AWS account within the organizatio
G. Populate the prefix lists in each account with all internal CIDR range
H. Notify the owner of each AWS account to allow the new customer-managed prefix list IDs in their accounts in their security group
I. Instruct the security team to share updates with each AWS account owner.
J. Create a new customer-managed prefix list in the security team's AWS accoun
K. Populate the customer-managed prefix list with all internal CIDR range
L. Share the customer-managed prefix listwith the organization by using AWS Resource Access Manage
M. Notify the owner of each AWS account to allow the new customer-managed prefix list ID in their security groups.
N. Create an IAM role in each account in the organizatio
O. Grant permissions to update security groups.Deploy an AWS Lambda function in the security team's AWS accoun
P. Configure the Lambda function to take a list of internal IP addresses as input, assume a role in each organization account, and add the list of IP addresses to the security groups in each account.

Answer: C

Explanation:
Create a new customer-managed prefix list in the security team's AWS account. Populate the customer-managed prefix list with all internal CIDR ranges. Share the customer-managed prefix list with the organization by using AWS Resource Access Manager. Notify the owner of each AWS account to allow the new customer-managed prefix list ID in their security groups. This solution meets the requirements with the least amount of operational overhead as it requires the security team to create and maintain a single customer-managed prefix list, and share it with the organization using AWS Resource Access Manager. The owners of each AWS account are then responsible for allowing the prefix list in their security groups, which eliminates the need for the security team to manually notify each account owner when changes are made. This solution also eliminates the need for a separate AWS Lambda function in each account, reducing the overall complexity of the solution.

NEW QUESTION 74
- (Exam Topic 1)
A start up company hosts a fleet of Amazon EC2 instances in private subnets using the latest Amazon Linux 2 AMI. The company's engineers rely heavily on SSH access to the instances for troubleshooting.
The company's existing architecture includes the following:
• A VPC with private and public subnets, and a NAT gateway
• Site-to-Site VPN for connectivity with the on-premises environment
• EC2 security groups with direct SSH access from the on-premises environment
The company needs to increase security controls around SSH access and provide auditing of commands executed by the engineers.
Which strategy should a solutions architect use?

A. Install and configure EC2 Instance Connect on the fleet of EC2 instance
B. Remove all security group rules attached to EC2 instances that allow inbound TCP on port 22. Advise the engineers to remotely access the instances by using the EC2 Instance Connect CLI.
C. Update the EC2 security groups to only allow inbound TCP on port 22 to the IP addresses of the engineer's device
D. Install the Amazon CloudWatch agent on all EC2 instances and send operating system audit logs to CloudWatch Logs.
E. Update the EC2 security groups to only allow inbound TCP on port 22 to the IP addresses of the engineer's device
F. Enable AWS Config for EC2 security group resource change
G. Enable AWS Firewall Manager and apply a security group policy that automatically remediates changes to rules.
H. Create an IAM role with the AmazonSSMManagedInstanceCore managed policy attache
I. Attach the IAM role to all the EC2 instance
J. Remove all security group rules attached to the EC2 instances that allow inbound TCP on port 22. Have the engineers install the AWS Systems Manager Session Manager plugin for their devices and remotely access the instances by using the start-session API call from Systems Manager.

Answer: D

Explanation:
Allows client machines to be able to connect to Session Manager using the AWS CLI instead of going through the AWS EC2 or AWS Server Manager console. https://docs.aws.amazon.com/systems-manager/latest/userguide/session-manager-working-with-install-plugin.ht https://docs.aws.amazon.com/systems-manager/latest/userguide/session-manager-working-with-install-plugin.ht

NEW QUESTION 78
- (Exam Topic 1)
A solutions architect is investigating an issue in which a company cannot establish new sessions in Amazon Workspaces. An initial analysis indicates that the issue involves user profiles. The Amazon Workspaces environment is configured to use Amazon FSx for Windows File Server as the profile share storage. The FSx for Windows File Server file system is configured with 10 TB of storage.
The solutions architect discovers that the file system has reached its maximum capacity. The solutions architect must ensure that users can regain access. The solution also must prevent the problem from occurring again.
Which solution will meet these requirements?

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

Answer: B

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

NEW QUESTION 83
- (Exam Topic 1)
A company wants to use a third-party software-as-a-service (SaaS) application. The third-party SaaS application is consumed through several API calls. The third-party SaaS application also runs on AWS inside a VPC.
The company will consume the third-party SaaS application from inside a VPC. The company has internal security policies that mandate the use of private connectivity that does not traverse the internet. No resources that run in the company VPC are allowed to be accessed from outside the company's VPC. All permissions must conform to the principles of least privilege.
Which solution meets these requirements?

A. Create an AWS PrivateLink interface VPC endpoin
B. Connect this endpoint to the endpoint service that the third-party SaaS application provide
C. Create a security group to limit the access to the endpoin
D. Associate the security group with the endpoint.

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

Answer: A

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

NEW QUESTION 84
- (Exam Topic 2)
A company processes environment data. The has a set up sensors to provide a continuous stream of data from different areas in a city. The data is available in JSON format.
The company wants to use an AWS solution to send the data to a database that does not require fixed schemas for storage. The data must be send in real time.
Which solution will meet these requirements?

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

Answer: B

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

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

NEW QUESTION 85
- (Exam Topic 2)
A company manufactures smart vehicles. The company uses a custom application to collect vehicle data. The vehicles use the MQTT protocol to connect to the application.
The company processes the data in 5-minute intervals. The company then copies vehicle telematics data to on-premises storage. Custom applications analyze this data to detect anomalies.
The number of vehicles that send data grows constantly. Newer vehicles generate high volumes of data. The on-premises storage solution is not able to scale for peak traffic, which results in data loss. The company must modernize the solution and migrate the solution to AWS to resolve the scaling challenges.
Which solution will meet these requirements with the LEAST operational overhead?

A. Use AWS IOT Greengrass to send the vehicle data to Amazon Managed Streaming for Apache Kafka (Amazon MSK). Create an Apache Kafka application to store the data in Amazon S3. Use a pretrained model in Amazon SageMaker to detect anomalies.
B. Use AWS IOT Core to receive the vehicle dat
C. Configure rules to route data to an Amazon Kinesis Data Firehose delivery stream that stores the data in Amazon S3. Create an Amazon Kinesis Data Analytics application that reads from the delivery stream to detect anomalies.
D. Use AWS IOT FleetWise to collect the vehicle dat
E. Send the data to an Amazon Kinesis data stream.Use an Amazon Kinesis Data Firehose delivery stream to store the data in Amazon S3. Use the built-in machine learning transforms in AWS Glue to detect anomalies.
F. Use Amazon MQ for RabbitMQ to collect the vehicle dat
G. Send the data to an Amazon Kinesis Data Firehose delivery stream to store the data in Amazon S3. Use Amazon Lookout for Metrics to detect anomalies.

Answer: B

Explanation:
Using AWS IoT Core to receive the vehicle data will enable connecting the smart vehicles to the cloud using the MQTT protocol1. AWS IoT Core is a platform that enables you to connect devices to AWS Services and other devices, secure data and interactions, process and act upon device data, and enable applications to interact with devices even when they are offline2. Configuring rules to route data to an Amazon Kinesis Data Firehose delivery stream that stores the data in Amazon S3 will enable processing and storing the vehicle data in a scalable and reliable way3. Amazon Kinesis Data Firehose is a fully managed service that delivers real-time streaming data to destinations such as Amazon S3. Creating an Amazon Kinesis Data Analytics application that reads from the delivery stream to detect anomalies will enable analyzing the vehicle data using SQL queries or Apache Flink applications. Amazon Kinesis Data Analytics is a fully managed service that enables you to process and analyze streaming data using SQL or Java.

NEW QUESTION 89
- (Exam Topic 2)
A company is migrating a document processing workload to AWS. The company has updated many applications to natively use the Amazon S3 API to store, retrieve, and modify documents that a processing server generates at a rate of approximately 5 documents every second. After the document processing is finished, customers can download the documents directly from Amazon S3.
During the migration, the company discovered that it could not immediately update the processing server that generates many documents to support the S3 API. The server runs on Linux and requires fast local access to the files that the server generates and modifies. When the server finishes processing, the files must be available to the public for download within 30 minutes.
Which solution will meet these requirements with the LEAST amount of effort?

A. Migrate the application to an AWS Lambda functio
B. Use the AWS SDK for Java to generate, modify, and access the files that the company stores directly in Amazon S3.

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

Answer: C

Explanation:
Amazon FSx for Lustre is a fully managed service that provides cost-effective, high-performance, scalable storage for compute workloads. Powered by Lustre, the world's most popular high-performance file system, FSx for Lustre offers shared storage with sub-ms latencies, up to terabytes per second of throughput, and millions of IOPS. FSx for Lustre file systems can also be linked to Amazon Simple Storage Service (S3) buckets, allowing you to access and process data concurrently from both a high-performance file system and from the S3 API.

NEW QUESTION 93
- (Exam Topic 2)
A company operates an on-premises software-as-a-service (SaaS) solution that ingests several files daily. The company provides multiple public SFTP endpoints to its customers to facilitate the file transfers. The customers add the SFTP endpoint IP addresses to their firewall allow list for outbound traffic. Changes to the SFTP endmost IP addresses are not permitted.
The company wants to migrate the SaaS solution to AWS and decrease the operational overhead of the file transfer service.
Which solution meets these requirements?

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

Answer: A

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

NEW QUESTION 98
- (Exam Topic 2)
A company needs to establish a connection from its on-premises data center to AWS. The company needs to connect all of its VPCs that are located in different AWS Regions with transitive routing capabilities between VPC networks. The company also must reduce network outbound traffic costs, increase bandwidth throughput, and provide a consistent network experience for end users.
Which solution will meet these requirements?

A. Create an AWS Site-to-Site VPN connection between the on-premises data center and a new central VP
B. Create VPC peering connections that initiate from the central VPC to all other VPCs.
C. Create an AWS Direct Connect connection between the on-premises data center and AW
D. Provision a transit VIF, and connect it to a Direct Connect gatewa
E. Connect the Direct Connect gateway to all the other VPCs by using a transit gateway in each Region.
F. Create an AWS Site-to-Site VPN connection between the on-premises data center and a new central VP
G. Use a transit gateway with dynamic routin
H. Connect the transit gateway to all other VPCs.
I. Create an AWS Direct Connect connection between the on-premises data center and AWS Establish an AWS Site-to-Site VPN connection between all VPCs in each Regio
J. Create VPC peering connections that initiate from the central VPC to all other VPCs.

Answer: B

Explanation:
Transit GW + Direct Connect GW + Transit VIF + enabled SiteLink if two different DX locations https://aws.amazon.com/blogs/networking-and-content-delivery/introducing-aws-direct-connect-sitelink/

NEW QUESTION 103
- (Exam Topic 2)
A company has millions of objects in an Amazon S3 bucket. The objects are in the S3 Standard storage class. All the S3 objects are accessed frequently. The number of users and applications that access the objects is increasing rapidly. The objects are encrypted with server-side encryption with AWS KMS Keys (SSE-KMS).
A solutions architect reviews the company's monthly AWS invoice and notices that AWS KMS costs are increasing because of the high number of requests from Amazon S3. The solutions architect needs to optimize costs with minimal changes to the application.
Which solution will meet these requirements with the LEAST operational overhead?

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

Answer: B

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

NEW QUESTION 107
- (Exam Topic 2)
A company runs an intranet application on premises. The company wants to configure a cloud backup of the application. The company has selected AWS Elastic Disaster Recovery for this solution.
The company requires that replication traffic does not travel through the public internet. The application also must not be accessible from the internet. The company does not want this solution to consume all available network bandwidth because other applications require bandwidth.
Which combination of steps will meet these requirements? (Select THREE.)

A. Create a VPC that has at least two private subnets, two NAT gateways, and a virtual private gateway.
B. Create a VPC that has at least two public subnets, a virtual private gateway, and an internet gateway.
C. Create an AWS Site-to-Site VPN connection between the on-premises network and the target AWS network.
D. Create an AWS Direct Connect connection and a Direct Connect gateway between the on-premises network and the target AWS network.
E. During configuration of the replication servers, select the option to use private IP addresses for data replication.
F. During configuration of the launch settings for the target servers, select the option to ensure that the Recovery instance's private IP address matches the source server's private IP address.

Answer: BDE

Explanation:
AWS Elastic Disaster Recovery (AWS DRS) is a service that minimizes downtime and data loss with fast, reliable recovery of on-premises and cloud-based applications using affordable storage, minimal compute, and point-in-time recovery1. Users can set up AWS DRS on their source servers to initiate secure data replication to a staging area subnet in their AWS account, in the AWS Region they select. Users can then launch recovery instances on AWS within minutes, using the most up-to-date server state or a previous point in time.
To configure a cloud backup of the application with AWS DRS, users need to create a VPC that has at least
two public subnets, a virtual private gateway, and an internet gateway. A VPC is a logically isolated section of the AWS Cloud where users can launch AWS resources in a virtual network that they define2. A public subnet is a subnet that has a route to an internet gateway3. A virtual private gateway is the VPN concentrator on the Amazon side of the Site-to-Site VPN connection4. An internet gateway is a horizontally scaled, redundant, and highly available VPC component that allows communication between instances in the VPC and the internet. Users need to create at least two public subnets for redundancy and high availability. Users need to create a virtual private gateway and attach it to the VPC to enable VPN connectivity between the on-premises network and the target AWS network. Users need to create an internet gateway and attach it to the VPC to enable internet access for the replication servers.
To ensure that replication traffic does not travel through the public internet, users need to create an AWS Direct Connect connection and a Direct Connect gateway between the on-premises network and the target AWS network. AWS Direct Connect is a service that establishes a dedicated network connection from an on-premises network to one or more VPCs. A Direct Connect gateway is a globally available resource that allows users to connect multiple VPCs across different Regions to their on-premises networks using one or more Direct Connect connections. Users need to create an AWS Direct Connect connection between their on-premises network and an AWS Region. Users need to create a Direct Connect gateway and associate it with their VPC and their Direct Connect connection.
To ensure that the application is not accessible from the internet, users need to select the option to use private IP addresses for data replication during configuration of the replication servers. This option configures the replication servers with private IP addresses only, without assigning any public IP addresses or Elastic IP addresses. This way, the replication servers can only communicate with other resources within the VPC or through VPN connections.
Option A is incorrect because creating a VPC that has at least two private subnets, two NAT gateways, and a virtual private gateway is not necessary or cost-effective. A private subnet is a subnet that does not have a route to an internet gateway3. A NAT gateway is a highly available, managed Network Address Translation (NAT) service that enables instances in a private subnet to connect to the internet or other AWS services, but prevents the internet from initiating connections with those instances. Users do not need to create private subnets or NAT gateways for this use case, as they can use public subnets with private IP addresses for data replication.
Option C is incorrect because creating an AWS Site-to-Site VPN connection between the on-premises network and the target AWS network will not ensure that replication traffic does not travel through the public
internet. A Site-to-Site VPN connection consists of two VPN tunnels between an on-premises customer
gateway device and a virtual private gateway in your VPC4. The VPN tunnels are encrypted using IPSec protocols, but they still use public IP addresses for communication. Users need to use AWS Direct Connect instead of Site-to-Site VPN for this use case.
Option F is incorrect because selecting the option to ensure that the Recovery instance's private IP address matches the source server's private IP address during configuration of the launch settings for the target servers will not ensure that the application is not accessible from the internet. This option configures the Recovery instance with an identical private IP address as its source server when launched in drills or recovery mode. However, this option does not prevent assigning public IP addresses or Elastic IP addresses to the Recovery instance. Users need to select the option to use private IP addresses for data replication instead.

NEW QUESTION 112
- (Exam Topic 2)
A company has a critical application in which the data tier is deployed in a single AWS Region. The data tier uses an Amazon DynamoDB table and an Amazon Aurora MySQL DB cluster. The current Aurora MySQL engine version supports a global database. The application tier is already deployed in two Regions.
Company policy states that critical applications must have application tier components and data tier components deployed across two Regions. The RTO and RPO must be no more than a few minutes each. A solutions architect must recommend a solution to make the data tier compliant with company policy.
Which combination of steps will meet these requirements? (Choose two.)

A. Add another Region to the Aurora MySQL DB cluster
B. Add another Region to each table in the Aurora MySQL DB cluster
C. Set up scheduled cross-Region backups for the DynamoDB table and the Aurora MySQL DB cluster
D. Convert the existing DynamoDB table to a global table by adding another Region to its configuration
E. Use Amazon Route 53 Application Recovery Controller to automate database backup and recovery to the secondary Region

Answer: AD

Explanation:
The company should use Amazon Aurora global database and Amazon DynamoDB global table to deploy the data tier components across two Regions. Amazon Aurora global database is a feature that allows a single
Aurora database to span multiple AWS Regions, enabling low-latency global reads and fast recovery from Region-wide outages1. Amazon DynamoDB global table is a feature that allows a single DynamoDB table to span multiple AWS Regions, enabling low-latency global reads and writes and fast recovery from Region-wide outages2.
References:

- https://aws.amazon.com/rds/aurora/global-database/
- https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/globaltables_HowItWorks.html
- https://aws.amazon.com/route53/application-recovery-controller/

NEW QUESTION 117
- (Exam Topic 2)
A company's public API runs as tasks on Amazon Elastic Container Service (Amazon ECS). The tasks run on AWS Fargate behind an Application Load Balancer (ALB) and are configured with Service Auto Scaling for the tasks based on CPU utilization. This service has been running well for several months.
Recently, API performance slowed down and made the application unusable. The company discovered that a significant number of SQL injection attacks had occurred against the API and that the API service had scaled to its maximum amount.
A solutions architect needs to implement a solution that prevents SQL injection attacks from reaching the ECS API service. The solution must allow legitimate traffic through and must maximize operational efficiency. Which solution meets these requirements?

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

Answer: C

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

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

References:

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

NEW QUESTION 121
- (Exam Topic 2)
A company is creating a REST API to share information with six of its partners based in the United States. The company has created an Amazon API Gateway Regional endpoint. Each of the six partners will access the API once per day to post daily sales figures.

After initial deployment, the company observes 1.000 requests per second originating from 500 different IP addresses around the world. The company believes this traffic is originating from a botnet and wants to secure its API while minimizing cost.
Which approach should the company take to secure its API?

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

Answer: D

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

NEW QUESTION 126
- (Exam Topic 2)
A global manufacturing company plans to migrate the majority of its applications to AWS. However, the company is concerned about applications that need to remain within a specific country or in the company's central on-premises data center because of data regulatory requirements or requirements for latency of single-digit milliseconds. The company also is concerned about the applications that it hosts in some of its factory sites, where limited network infrastructure exists.
The company wants a consistent developer experience so that its developers can build applications once and deploy on premises, in the cloud, or in a hybrid architecture.
The developers must be able to use the same tools, APIs, and services that are familiar to them. Which solution will provide a consistent hybrid experience to meet these requirements?

A. Migrate all applications to the closest AWS Region that is complian
B. Set up an AWS Direct Connect connection between the central on-premises data center and AW
C. Deploy a Direct Connect gateway.
D. Use AWS Snowball Edge Storage Optimized devices for the applications that have data regulatory requirements or requirements for latency of single-digit millisecond
E. Retain the devices on premise
F. Deploy AWS Wavelength to host the workloads in the factory sites.
G. Install AWS Outposts for the applications that have data regulatory requirements or requirements for latency of single-digit millisecond
H. Use AWS Snowball Edge Compute Optimized devices to host the workloads in the factory sites.
I. Migrate the applications that have data regulatory requirements or requirements for latency ofsingle-digit milliseconds to an AWS Local Zon
J. Deploy AWS Wavelength to host the workloads in the factory sites.

Answer: C

Explanation:
Installing AWS Outposts for the applications that have data regulatory requirements or requirements for latency of single-digit milliseconds will provide a fully managed service that extends AWS infrastructure, services, APIs, and tools to customer premises1. AWS Outposts allows customers to run some AWS services locally and connect to a broad range of services available in the local AWS Region1. Using AWS Snowball Edge Compute Optimized devices to host the workloads in the factory sites will provide local compute and storage resources for locations with limited network infrastructure2. AWS Snowball Edge devices can run Amazon EC2 instances and AWS Lambda functions locally and sync data with AWS when network connectivity is available2.

NEW QUESTION 127
- (Exam Topic 2)
A solutions architect needs to review the design of an Amazon EMR cluster that is using the EMR File System (EMRFS). The cluster performs tasks that are critical to business needs. The cluster is running Amazon EC2 On-Demand Instances at all times tor all task, primary, and core nodes. The EMR tasks run each morning, starting at 1 ;00 AM. and take 6 hours to finish running. The amount of time to complete the processing is not a priority because the data is not referenced until late in the day.
The solutions architect must review the architecture and suggest a solution to minimize the compute costs. Which solution should the solutions architect recommend to meet these requirements?

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

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

Answer: A

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

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

NEW QUESTION 131
- (Exam Topic 2)
A solutions architect at a large company needs to set up network security tor outbound traffic to the internet from all AWS accounts within an organization in AWS Organizations. The organization has more than 100 AWS accounts, and the accounts route to each other by using a centralized AWS Transit Gateway. Each account has both an internet gateway and a NAT gateway tor outbound traffic to the internet The company deploys resources only into a single AWS Region. The company needs the ability to add centrally managed rule-based filtering on all outbound traffic to the internet for all AWS accounts in the organization. The peak load of outbound traffic will not exceed 25 Gbps in each Availability Zone.
Which solution meets these requirements?

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

Answer: B

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

NEW QUESTION 135
- (Exam Topic 2)
A company's interactive web application uses an Amazon CloudFront distribution to serve images from an Amazon S3 bucket. Occasionally, third-party tools ingest corrupted images into the S3 bucket. This image corruption causes a poor user experience in the application later. The company has successfully implemented and tested Python logic to detect corrupt images.
A solutions architect must recommend a solution to integrate the detection logic with minimal latency between the ingestion and serving.
Which solution will meet these requirements?

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

Answer: B

Explanation:
This solution will allow the detection logic to be run as soon as the image is uploaded to the S3 bucket, before it is served to users via the CloudFront distribution. This way, the detection logic can quickly identify any corrupted images and prevent them from being served to users, minimizing latency between ingestion and serving.
Reference: AWS Lambda@Edge documentation:
https://docs.aws.amazon.com/lambda/latest/dg/lambda-edge.html You can use Lambda@Edge to run your code in response to CloudFront events, such as a viewer request, an origin request, a response, or an error.

NEW QUESTION 140
- (Exam Topic 2)
A financial services company loaded millions of historical stock trades into an Amazon DynamoDB table. The table uses on-demand capacity mode. Once each day at midnight, a few million new records are loaded into the table. Application read activity against the table happens in bursts throughout the day. and a limited set of keys are repeatedly looked up. The company needs to reduce costs associated with DynamoDB.
Which strategy should a solutions architect recommend to meet this requirement?

A. Deploy an Amazon ElastiCache cluster in front of the DynamoDB table.
B. Deploy DynamoDB Accelerator (DAX). Configure DynamoDB auto scalin
C. Purchase Savings Plans in Cost Explorer
D. Use provisioned capacity mod
E. Purchase Savings Plans in Cost Explorer.
F. Deploy DynamoDB Accelerator (DAX). Use provisioned capacity mod
G. Configure DynamoDB auto scaling.

Answer: D

Explanation:
https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/HowItWorks.ReadWriteCapacityMode.h

NEW QUESTION 143
- (Exam Topic 2)
A company is building a call center by using Amazon Connect. The company's operations team is defining a disaster recovery (DR) strategy across AWS Regions. The contact center has dozens of contact flows, hundreds of users, and dozens of claimed phone numbers.
Which solution will provide DR with the LOWEST RTO?

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

Answer: D

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

NEW QUESTION 148
.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questons and Answers in PDF Format

SAP-C02 Practice Exam Features:

* SAP-C02 Questions and Answers Updated Frequently

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

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

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