Amazon

Exam Questions AWS-Certified-Solutions-Architect-Professional

Amazon AWS Certified Solutions Architect Professional

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

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

Answer: A

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

NEW QUESTION 2
- (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 3
- (Exam Topic 2)
A software-as-a-service (SaaS) provider exposes APIs through an Application Load Balancer (ALB). The ALB connects to an Amazon Elastic Kubernetes Service (Amazon EKS) cluster that is deployed in the
us-east-I Region. The exposed APIs contain usage of a few non-standard REST methods: LINK, UNLINK, LOCK, and UNLOCK.
Users outside the United States are reporting long and inconsistent response times for these APIs. A solutions architect needs to resolve this problem with a solution that minimizes operational overhead.
Which solution meets these requirements?

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

Answer: C

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

NEW QUESTION 4
- (Exam Topic 2)
A solutions architect must create a business case for migration of a company's on-premises data center to the AWS Cloud. The solutions architect will use a configuration management database (CMDB) export of all the company's servers to create the case.
Which solution will meet these requirements MOST cost-effectively?

A. Use AWS Well-Architected Tool to import the CMDB data to perform an analysis and generate recommendations.
B. Use Migration Evaluator to perform an analysi
C. Use the data import template to upload the data from the CMDB export.
D. Implement resource matching rule
E. Use the CMDB export and the AWS Price List Bulk API to query CMDB data against AWS services in bulk.

F. Use AWS Application Discovery Service to import the CMDB data to perform an analysis.

Answer: B

Explanation:
https://aws.amazon.com/blogs/architecture/accelerating-your-migration-to-aws/ Build a business case with AWS Migration Evaluator The foundation for a successful migration starts with a defined business objective (for example, growth or new offerings). In order to enable the business drivers, the established business case must then be aligned to a technical capability (increased security and elasticity). AWS Migration Evaluator (formerly known as TSO Logic) can help you meet these objectives. To get started, you can choose to upload exports from third-party tools such as Configuration Management Database (CMDB) or install a collector agent to monitor. You will receive an assessment after data collection, which includes a projected cost estimate and savings of running your on-premises workloads in the AWS Cloud. This estimate will provide a summary of the projected costs to re-host on AWS based on usage patterns. It will show the breakdown of costs by infrastructure and software licenses. With this information, you can make the business case and plan next steps.

NEW QUESTION 5
- (Exam Topic 2)
A company is running an application on Amazon EC2 instances in the AWS Cloud. The application is using a MongoDB database with a replica set as its data tier. The MongoDB database is installed on systems in the company's on-premises data center and is accessible through an AWS Direct Connect connection to the data center environment.
A solutions architect must migrate the on-premises MongoDB database to Amazon DocumentDB (with MongoDB compatibility).
Which strategy should the solutions architect choose to perform this migration?

A. Create a fleet of EC2 instance
B. Install MongoDB Community Edition on the EC2 instances, and create a databas
C. Configure continuous synchronous replication with the database that is running in theon-premises data center.
D. Create an AWS Database Migration Service (AWS DMS) replication instanc
E. Create a source endpoint for the on-premises MongoDB database by using change data capture (CDC). Create a target endpoint for the Amazon DocumentDB databas
F. Create and run a DMS migration task.
G. Create a data migration pipeline by using AWS Data Pipelin
H. Define data nodes for the on-premises MongoDB database and the Amazon DocumentDB databas
I. Create a scheduled task to run the data pipeline.
J. Create a source endpoint for the on-premises MongoDB database by using AWS Glue crawlers.Configure continuous asynchronous replication between the MongoDB database and the Amazon DocumentDB database.

Answer: B

Explanation:
https://aws.amazon.com/getting-started/hands-on/move-to-managed/migrate-mongodb-to-documentdb/

NEW QUESTION 6
- (Exam Topic 2)
A company runs an application on AWS. The company curates data from several different sources. The company uses proprietary algorithms to perform data transformations and aggregations. After the company performs E TL processes, the company stores the results in Amazon Redshift tables. The company sells this data to other companies. The company downloads the data as files from the Amazon Redshift tables and transmits the files to several data customers by using FTP. The number of data customers has grown significantly. Management of the data customers has become difficult.
The company will use AWS Data Exchange to create a data product that the company can use to share data with customers. The company wants to confirm the identities of the customers before the company shares data.
The customers also need access to the most recent data when the company publishes the data. Which solution will meet these requirements with the LEAST operational overhead?

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

Answer: C

Explanation:
The company should download the data from the Amazon Redshift tables to an Amazon S3 bucket periodically and use AWS Data Exchange for S3 to share data with customers. The company should configure subscription verification and require the data customers to subscribe to the data product. This solution will meet the requirements with the least operational overhead because AWS Data Exchange for S3 is a feature that enables data subscribers to access third-party data files directly from data providers' Amazon S3 buckets. Subscribers can easily use these files for their data analysis with AWS services without needing to create or manage data copies. Data providers can easily set up AWS Data Exchange for S3 on top of their existing S3 buckets to share direct access to an entire S3 bucket or specific prefixes and S3 objects. AWS Data Exchange automatically manages subscriptions, entitlements, billing, and payment1.
The other options are not correct because:
- Using AWS Data Exchange for APIs to share data with customers would not work because AWS Data Exchange for APIs is a feature that enables data subscribers to access third-party APIs directly from data providers' AWS accounts. Subscribers can easily use these APIs for their data analysis with AWS services without needing to manage API keys or tokens. Data providers can easily set up AWS Data Exchange for APIs on top of their existing API Gateway resources to share direct access to an entire API or specific routes and stages2. However, this feature is not suitable for sharing data from Amazon Redshift

tables, which are not exposed as APIs.

- Creating an Amazon API Gateway Data API service integration with Amazon Redshift would not work because the Data API is a feature that enables you to query your Amazon Redshift cluster using HTTP requests, without needing a persistent connection or a SQL client3. It is useful for building applications that interact with Amazon Redshift, but not for sharing data files with customers.
- Creating an AWS Data Exchange datashare by connecting AWS Data Exchange to the Redshift cluster would not work because AWS Data Exchange does not support datashares for Amazon Redshift clusters. A datashare is a feature that enables you to share live and secure access to your Amazon Redshift data across your accounts or with third parties without copying or moving the underlying data4. It is useful for sharing query results and views with other users, but not for sharing data files with customers.
- Publishing the Amazon Redshift data to an Open Data on AWS Data Exchange would not work because Open Data on AWS Data Exchange is a feature that enables you to find and use free and public datasets from AWS customers and partners. It is useful for accessing open and free data, but not for confirming the identities of the customers or charging them for the data.

References:

- https://aws.amazon.com/data-exchange/why-aws-data-exchange/s3/
- https://aws.amazon.com/data-exchange/why-aws-data-exchange/api/
- https://docs.aws.amazon.com/redshift/latest/mgmt/data-api.html
- https://docs.aws.amazon.com/redshift/latest/dg/datashare-overview.html
- https://aws.amazon.com/data-exchange/open-data/

NEW QUESTION 7

- (Exam Topic 1)

A company is developing a new service that will be accessed using TCP on a static port A solutions architect must ensure that the service is highly available, has redundancy across Availability Zones, and is accessible using the DNS name myservice.com, which is publicly accessible The service must use fixed address assignments so other companies can add the addresses to their allow lists.

Assuming that resources are deployed in multiple Availability Zones in a single Region, which solution will meet these requirements?

A. Create Amazon EC2 instances with an Elastic IP address for each instance Create a Network Load Balancer (NLB) and expose the static TCP port Register EC2instances with the NLB Create a new name server record set named my service com, and assign the Elastic IP addresses of the EC2 instances to the record set Provide the Elastic IP addresses of the EC2 instances to the other companies to add to their allow lists

B. Create an Amazon ECS cluster and a service definition for the application Create and assign public IP addresses for the ECS cluster Create a Network Load Balancer (NLB) and expose the TCP port Create atarget group and assign the ECS cluster name to the NLB Create a new A record set named my service com and assign the public IP addresses of the ECS cluster to the record set Provide the public IP addresses of the ECS cluster to the other companies to add to their allow lists

C. Create Amazon EC2 instances for the service Create one Elastic IP address for each Availability Zone Create a Network Load Balancer (NLB) and expose the assigned TCP port Assign the Elastic IP addresses to the NLB for each Availability Zone Create a target group and register the EC2 instances with the NLB Create a new A (alias) record set named my service com, and assign the NLB DNS name to the record set.

D. Create an Amazon ECS cluster and a service definition for the application Create and assign public IP address for each host in the cluster Create an Application Load Balancer (ALB) and expose the static TCP port Create a target group and assign the ECS service definition name to the ALB Create a new CNAME record set and associate the public IP addresses to the record set Provide the Elastic IP addresses of the Amazon EC2 instances to the other companies to add to their allow lists

Answer: C

Explanation:

https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-to-elb-load-balancer.html

Create a Network Load Balancer (NLB) and expose the assigned TCP port. Assign the Elastic IP addresses to the NLB for each Availability Zone. Create a target group and register the EC2 instances with the NLB. Create a new A (alias) record set named my.service.com, and assign the NLB DNS name to the record set. As it uses the NLB as the resource in the A-record, traffic will be routed through the NLB, and it will automatically route the traffic to the healthy instances based on the health checks and also it provides the fixed address assignments as the other companies can add the NLB's Elastic IP addresses to their allow lists.

NEW QUESTION 8

- (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 9
- (Exam Topic 1)
A company manages multiple AWS accounts by using AWS Organizations. Under the root OU. the company has two OUs: Research and DataOps.
Because of regulatory requirements, all resources that the company deploys in the organization must reside in the ap-northeast-1 Region. Additionally. EC2 instances that the company deploys in the DataOps OU must use a predefined list of instance types
A solutions architect must implement a solution that applies these restrictions. The solution must maximize operational efficiency and must minimize ongoing maintenance
Which combination of steps will meet these requirements? (Select TWO)

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

Answer: CE

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

NEW QUESTION 10
- (Exam Topic 1)
A company is in the process of implementing AWS Organizations to constrain its developers to use only Amazon EC2. Amazon S3 and Amazon DynamoDB. The developers account resides In a dedicated organizational unit (OU). The solutions architect has implemented the following SCP on the developers account:

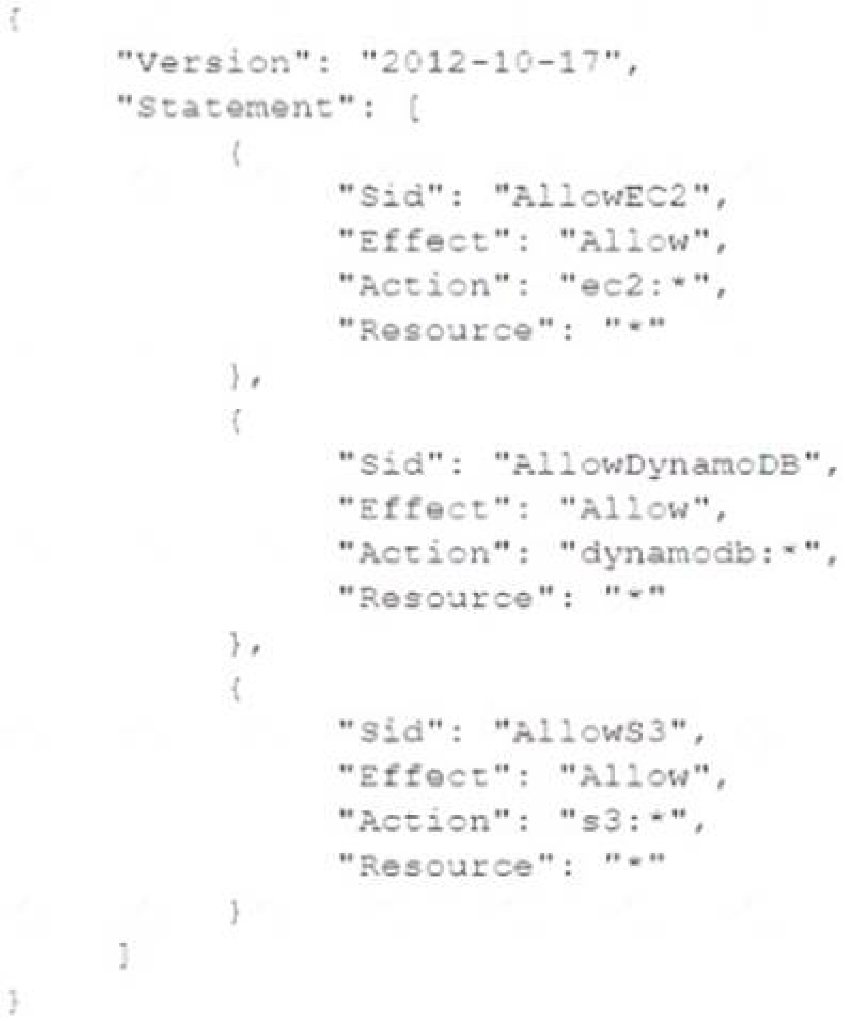

```
{
      "Version": "2012-10-17",
      "Statement": [
            {
                  "Sid": "AllowEC2",
                  "Effect": "Allow",
                  "Action": "ec2:*",
                  "Resource": "*"
            },
            {
                  "Sid": "AllowDynamoDB",
            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

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

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

Answer: B

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

NEW QUESTION 11
- (Exam Topic 1)
A financial services company in North America plans to release a new online web application to its customers on AWS . The company will launch the application in the us-east-1 Region on Amazon EC2 instances. The application must be highly available and must dynamically scale to meet user traffic. The company also wants to implement a disaster recovery environment for the application in the us-west-1 Region by using active-passive failover.
Which solution will meet these requirements?

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

Answer: C

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

NEW QUESTION 12
- (Exam Topic 1)
A company has its cloud infrastructure on AWS A solutions architect needs to define the infrastructure as code. The infrastructure is currently deployed in one AWS Region. The company's business expansion plan includes deployments in multiple Regions across multiple AWS accounts
What should the solutions architect do to meet these requirements?

A. Use AWS CloudFormation templates Add IAM policies to control the various accounts Deploy the templates across the multiple Regions
B. Use AWS Organizations Deploy AWS CloudFormation templates from the management account Use AWS Control Tower to manage deployments across accounts

C. Use AWS Organizations and AWS CloudFormation StackSets Deploy a CloudFormation template from an account that has the necessary IAM permissions
D. Use nested stacks with AWS CloudFormation templates Change the Region by using nested stacks

Answer: C

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

NEW QUESTION 13
- (Exam Topic 1)
An adventure company has launched a new feature on its mobile app. Users can use the feature to upload their hiking and ratting photos and videos anytime. The photos and videos are stored in Amazon S3 Standard storage in an S3 bucket and are served through Amazon CloudFront.
The company needs to optimize the cost of the storage. A solutions architect discovers that most of the uploaded photos and videos are accessed infrequently after 30 days. However, some of the uploaded photos and videos are accessed frequently after 30 days. The solutions architect needs to implement a solution that maintains millisecond retrieval availability of the photos and videos at the lowest possible cost.
Which solution will meet these requirements?

A. Configure S3 Intelligent-Tiering on the S3 bucket.
B. Configure an S3 Lifecycle policy to transition image objects and video objects from S3 Standard to S3 Glacier Deep Archive after 30 days.
C. Replace Amazon S3 with an Amazon Elastic File System (Amazon EFS) file system that is mounted on Amazon EC2 instances.
D. Add a Cache-Control: max-age header to the S3 image objects and S3 video object
E. Set the header to 30 days.

Answer: A

Explanation:
Amazon S3 Intelligent-Tiering is a storage class that automatically moves objects between two access tiers based on changing access patterns. Objects that are accessed frequently are stored in the frequent access tier and objects that are accessed infrequently are stored in the infrequent access tier. This allows for cost optimization without requiring manual intervention. This makes it an ideal solution for the scenario described, as it can automatically move objects that are infrequently accessed after 30 days to a lower-cost storage tier while still maintaining millisecond retrieval availability.

NEW QUESTION 14
- (Exam Topic 1)
A health insurance company stores personally identifiable information (PII) in an Amazon S3 bucket. The company uses server-side encryption with S3 managed encryption keys (SSE-S3) to encrypt the objects. According to a new requirement, all current and future objects in the S3 bucket must be encrypted by keys that the company's security team manages. The S3 bucket does not have versioning enabled. Which solution will meet these requirements?

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

Answer: D

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

NEW QUESTION 15
- (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 16
- (Exam Topic 3)
A company is running a workload that consists of thousands of Amazon EC2 instances. The workload is running in a VPC that contains several public subnets and private subnets. The public subnets have a route for 0.0.0.0/0 to an existing internet gateway. The private subnets have a route for 0.0.0.0/0 to an existing NAT gateway.
A solutions architect needs to migrate the entire fleet of EC2 instances to use IPv6. The EC2 instances that are in private subnets must not be accessible from the public internet.
What should the solutions architect do to meet these requirements?

A. Update the existing VPC, and associate a custom IPv6 CIDR block with the VPC and all subnets.Update all the VPC route tables, and add a route for ::/0 to the internet gateway.
B. Update the existing VPC, and associate an Amazon-provided IPv6 CIDR block with the VPC and all subnet
C. Update the VPC route tables for all private subnets, and add a route for ::/0 to the NAT gateway.
D. Update the existing VPC, and associate an Amazon-provided IPv6 CIDR block with the VPC and all subnet
E. Create an egress-only internet gatewa
F. Update the VPC route tables for all private subnets, and add a route for ::/0 to the egress-only internet gateway.
G. Update the existing VPC, and associate a custom IPv6 CIDR block with the VPC and all subnet
H. Create a new NAT gateway, and enable IPv6 suppor
I. Update the VPC route tables for all private subnets, and add a route for ::/0 to the IPv6-enabled NAT gateway.

Answer: C

NEW QUESTION 17
- (Exam Topic 3)
A company deploys a new web application. As pari of the setup, the company configures AWS WAF to log to Amazon S3 through Amazon Kinesis Data Firehose. The company develops an Amazon Athena query that runs once daily to return AWS WAF log data from the previous 24 hours. The volume of daily logs is constant. However, over time, the same query is taking more time to run.
A solutions architect needs to design a solution to prevent the query time from continuing to increase. The solution must minimize operational overhead.
Which solution will meet these requirements?

A. Create an AWS Lambda function that consolidates each day's AWS WAF logs into one log file.
B. Reduce the amount of data scanned by configuring AWS WAF to send logs to a different S3 bucket each day.
C. Update the Kinesis Data Firehose configuration to partition the data in Amazon S3 by date and time.Create external tables for Amazon Redshif
D. Configure Amazon Redshift Spectrum to query the data source.
E. Modify the Kinesis Data Firehose configuration and Athena table definition to partition the data by date and tim
F. Change the Athena query to view the relevant partitions.

Answer: D

Explanation:
The best solution is to modify the Kinesis Data Firehose configuration and Athena table definition to partition the data by date and time. This will reduce the amount of data scanned by Athena and improve the query performance. Changing the Athena query to view the relevant partitions will also help to filter out unnecessary data. This solution requires minimal operational overhead as it does not involve creating additional resources or changing the log format. References: [AWS WAF Developer Guide], [Amazon Kinesis Data Firehose Use Guide], [Amazon Athena User Guide]

NEW QUESTION 18
- (Exam Topic 3)
A company has an organization in AWS Organizations that includes a separate AWS account for each of the company's departments. Application teams from different departments develop and deploy solutions independently.
The company wants to reduce compute costs and manage costs appropriately across departments. The company also wants to improve visibility into billing for individual departments. The company does not want to lose operational flexibility when the company selects compute resources.
Which solution will meet these requirements?

A. Use AWS Budgets for each departmen
B. Use Tag Editor to apply tags to appropriate resource
C. Purchase EC2 Instance Savings Plans.
D. Configure AWS Organizations to use consolidated billin
E. Implement a tagging strategy that identifies department
F. Use SCPs to apply tags to appropriate resource
G. Purchase EC2 Instance Savings Plans.
H. Configure AWS Organizations to use consolidated billin
I. Implement a tagging strategy that identifies department
J. Use Tag Editor to apply tags to appropriate resource
K. Purchase Compute Savings Plans.
L. Use AWS Budgets for each departmen
M. Use SCPs to apply tags to appropriate resource
N. Purchase Compute Savings Plans.

Answer: C

NEW QUESTION 19

- (Exam Topic 3)
A company has AWS accounts that are in an organization in AWS Organizations. The company wants to track Amazon EC2 usage as a metric. The company's architecture team must receive a daily alert if the EC2 usage is more than 10% higher than the average EC2 usage from the last 30 days.
Which solution will meet these requirements?

A. Configure AWS Budgets in the organization's management accoun
B. Specify a usage type of EC2 running hour
C. Specify a daily perio
D. Set the budget amount to be 10% more than the reported average usage for the last 30 days from AWS Cost Explore
E. Configure an alert to notify the architecture team if the usage threshold is met.
F. Configure AWS Cost Anomaly Detection in the organization's management accoun
G. Configure a monitor type of AWS Servic
H. Apply a filter of Amazon EC2. Configure an alert subscription to notify the architecture team if the usage is 10% more than the average usage for the last 30 days.
I. Enable AWS Trusted Advisor in the organization's management accoun
J. Configure a cost optimization advisory alert to notify the architecture team if the EC2 usage is 10% more than the reported average usage for the last 30 days.
K. Configure Amazon Detective in the organization's management accoun
L. Configure an EC2 usage anomaly alert to notify the architecture team if Detective identifies a usage anomaly of more than 10%.

Answer: B

Explanation:
AWS Cost Anomaly Detection is a feature of the AWS Cost Management suite that leverages machine learning to enable continuous monitoring of your AWS costs and usage, allowing you to identify unexpected
and abnormal spending1. You can create cost monitors that evaluate specific AWS services, member accounts cost allocation tags, or cost categories based on your AWS account structure2. You can also configure alert subscriptions that notify you when a cost monitor detects an anomaly that meets your threshold2. In this case, you can create a cost monitor with a monitor type of AWS Service and apply a filter of Amazon EC2 to track the EC2 usage as a metric. You can then configure an alert subscription to notify the architecture team if the usage is 10% more than the average usage for the last 30 days, which is the anomaly detection period used by AWS Cost Anomaly Detection3.

NEW QUESTION 20
- (Exam Topic 3)
A solutions architect has implemented a SAML 2 0 federated identity solution with their company's
on-premises identity provider (IdP) to authenticate users' access to the AWS environment. When the solutions architect tests authentication through the federated identity web portal, access to the AWS environment is granted However when test users attempt to authenticate through the federated identity web portal, they are not able to access the AWS environment
Which items should the solutions architect check to ensure identity federation is properly configured? (Select THREE)

A. The 1AM user's permissions policy has allowed the use of SAML federation for that user
B. The 1AM roles created for the federated users' or federated groups' trust policy have set the SAML provider as the principal
C. Test users are not in the AWSFederatedUsers group in the company's IdP
D. The web portal calls the AWS STS AssumeRoleWithSAML API with the ARN of the SAML provider, the ARN of the 1AM role, and the SAML assertion from IdP
E. The on-premises IdP's DNS hostname is reachable from the AWS environment VPCs
F. The company's IdP defines SAML assertions that properly map users or groups in the company to 1AM roles with appropriate permissions

Answer: BDF

NEW QUESTION 21
- (Exam Topic 3)
A company that provisions job boards for a seasonal workforce is seeing an increase in traffic and usage. The backend services run on a pair of Amazon EC2 instances behind an Application Load Balancer with Amazon DynamoDB as the datastore. Application read and write traffic is slow during peak seasons.
Which option provides a scalable application architecture to handle peak seasons with the LEAST development effort?

A. Migrate the backend services to AWS Lambd
B. Increase the read and write capacity of DynamoDB.
C. Migrate the backend services to AWS Lambd
D. Configure DynamoDB to use global tables.
E. Use Auto Scaling groups for the backend service
F. Use DynamoDB auto scaling.
G. Use Auto Scaling groups for the backend service
H. Use Amazon Simple Queue Service (Amazon SQS) and an AWS Lambda function to write to DynamoDB.

Answer: C

Explanation:
➢ Option C is correct because using Auto Scaling groups for the backend services allows the company to scale up or down the number of EC2 instances based on the demand and traffic. This way, the backend services can handle more requests during peak seasons without compromising performance or availability. Using DynamoDB auto scaling allows the company to adjust the provisioned read and write capacity of the table or index automatically based on the actual traffic patterns. This way, the table or index can handle sudden increases or decreases in workload without throttling or overprovisioning1.
➢ Option A is incorrect because migrating the backend services to AWS Lambda may require significant development effort to rewrite the code and test the functionality. Moreover, increasing the read and write capacity of DynamoDB manually may not be efficient or cost-effective, as it does not account for the variability of the workload. The company may end up paying for unused capacity or experiencing throttling if the workload exceeds the provisioned capacity1.
➢ Option B is incorrect because migrating the backend services to AWS Lambda may require significant development effort to rewrite the code and test the functionality. Moreover, configuring DynamoDB to use global tables may not be necessary or beneficial for the company, as global tables are mainly used for replicating data across multiple AWS Regions for fast local access and disaster recovery. Global tables do not automatically scale the provisioned capacity of each replica table; they still require manual or auto scaling settings2.
➢ Option D is incorrect because using Amazon Simple Queue Service (Amazon SQS) and an AWS Lambda function to write to DynamoDB may introduce additional complexity and latency to the application architecture. Amazon SQS is a message queue service that decouples and coordinates the components of a distributed system. AWS Lambda is a serverless compute service that runs code in response to events. Using these services may require significant development

effort to integrate them with the backend services and DynamoDB. Moreover, they may not improve the read performance of DynamoDB, which may also be affected by high traffic3.
References:

- Auto Scaling groups
- DynamoDB auto scaling
- AWS Lambda
- DynamoDB global tables
- AWS Lambda vs EC2: Comparison of AWS Compute Resources - Simform
- Managing throughput capacity automatically with DynamoDB auto scaling - Amazon DynamoDB
- AWS Aurora Global Database vs. DynamoDB Global Tables
- Amazon Simple Queue Service (SQS)

NEW QUESTION 22
- (Exam Topic 3)
A live-events company is designing a scaling solution for its ticket application on AWS. The application has high peaks of utilization during sale events. Each sale event is a one-time event that is scheduled. The application runs on Amazon EC2 instances that are in an Auto Scaling group.
The application uses PostgreSQL for the database layer.
The company needs a scaling solution to maximize availability during the sale events. Which solution will meet these requirements?

A. Use a predictive scaling policy for the EC2 instance
B. Host the database on an Amazon Aurora PostgreSQL Serverless v2 Multi-AZ DB instance with automatically scaling read replica
C. Create an AWS Step Functions state machine to run parallel AWS Lambda functions to pre-warm the database before a sale even
D. Create an Amazon EventBridge rule to invoke the state machine.
E. Use a scheduled scaling policy for the EC2 instance
F. Host the database on an Amazon RDS for PostgreSQL Multi-AZ DB instance with automatically scaling read replica
G. Create an Amazon EventBridge rule that invokes an AWS Lambda function to create a larger read replica before a sale even
H. Fail over to the larger read replic
I. Create another EventBridge rule that invokes another Lambda function to scale down the read replica after the sale event.
J. Use a predictive scaling policy for the EC2 instance
K. Host the database on an Amazon RDS for PostgreSQL Multi-AZ DB instance with automatically scaling read replica
L. Create an AWS Step Functions state machine to run parallel AWS Lambda functions to pre-warm the database before a sale even
M. Create an Amazon EventBridge rule to invoke the state machine.
N. Use a scheduled scaling policy for the EC2 instance
O. Host the database on an Amazon AuroraPostgreSQL Multi-AZ DB cluste
P. Create an Amazon EventBridge rule that invokes an AWS Lambda function to create a larger Aurora Replica before a sale even
Q. Fail over to the larger Aurora Replic
R. Create another EventBridge rule that invokes another Lambda function to scale down the Aurora Replica after the sale event.

Answer: D

Explanation:
The correct answer is D.
* D. This solution meets the requirements because it uses a scheduled scaling policy for the EC2 instances, which can adjust the capacity according to the known sale events. It also uses Amazon Aurora PostgreSQL Multi-AZ DB cluster, which provides high availability and durability for the database. It uses Amazon EventBridge rules and AWS Lambda functions to create a larger Aurora Replica before a sale event and fail over to it, which can improve the performance and handle the increased traffic. It also uses another EventBridge rule and Lambda function to scale down the Aurora Replica after the sale event, which can save costs123
* A. This solution is incorrect because it uses predictive scaling policy for the EC2 instances, which is not suitable for one-time events that are scheduled. Predictive scaling is based on historical data and machine learning, which may not accurately forecast the demand for sale events. It also uses Amazon Aurora PostgreSQL Serverless v2 Multi-AZ DB instance, which does not support read replicas. The use of AWS Step Functions state machine and Lambda functions to pre-warm the database is unnecessary and adds complexity45
* B. This solution is incorrect because it uses Amazon RDS for PostgreSQL Multi-AZ DB instance with automatically scaling read replicas, which may not provide enough performance improvement for the sale events. The use of EventBridge rules and Lambda functions to create a larger read replica and fail over to it is risky and may cause downtime or data loss. The use of another EventBridge rule and Lambda function to scale down the read replica is also risky and may cause inconsistency or data loss67
* C. This solution is incorrect because it uses predictive scaling policy for the EC2 instances, which is not suitable for one-time events that are scheduled. Predictive scaling is based on historical data and machine learning, which may not accurately forecast the demand for sale events. The use of AWS Step Functions state machine and Lambda functions to pre-warm the database is unnecessary and adds complexity45
References:
1: Scheduled scaling for Amazon EC2 Auto Scaling 2: Amazon Aurora PostgreSQL features 3: Amazon EventBridge rules 4: Predictive scaling for Amazon EC2 Auto Scaling 5: Amazon Aurora Serverless v2 6: Multi-AZ DB instance deployments - Amazon Relational Database Service 7: Working with PostgreSQL read replicas - Amazon Relational Database Service

NEW QUESTION 23
- (Exam Topic 3)
A company operates a fleet of servers on premises and operates a fleet of Amazon EC2 instances in its organization in AWS Organizations. The company's AWS accounts contain hundreds of VPCs. The company wants to connect its AWS accounts to its on-premises network. AWS Site-to-Site VPN connections are already established to a single AWS account. The company wants to control which VPCs can communicate with other VPCs.
Which combination of steps will achieve this level of control with the LEAST operational effort? (Choose three.)

A. Create a transit gateway in an AWS accoun
B. Share the transit gateway across accounts by using AWS Resource Access Manager (AWS RAM).
C. Configure attachments to all VPCs and VPNs.
D. Set up transit gateway route table
E. Associate the VPCs and VPNs with the route tables.
F. Configure VPC peering between the VPCs.
G. Configure attachments between the VPCs and VPNs.
H. Set up route tables on the VPCs and VPNs.

Answer: ABC

NEW QUESTION 24
- (Exam Topic 3)
A company is planning to migrate its on-premises VMware cluster of 120 VMS to AWS. The VMS have many different operating systems and many custom software packages installed. The company also has an on-premises NFS server that is 10 TB in size. The company has set up a 10 GbpsAWS Direct Connect connection to AWS for the migration
Which solution will complete the migration to AWS in the LEAST amount of time?

A. Export the on-premises VMS and copy them to an Amazon S3 bucke
B. Use VM Import/Export to create AMIS from the VM images that are stored in Amazon S3. Order an AWS Snowball Edge devic
C. Copy the NFS server data to the devic
D. Restore the NFS server data to an Amazon EC2 instance that has NFS configured.
E. Configure AWS Application Migration Service with a connection to the VMware cluste
F. Create a replication job for the VM
G. Create an Amazon Elastic File System (Amazon EFS) file syste
H. Configure AWS DataSync to copy the NFS server data to the EFS file system over the Direct Connect connection.
I. Recreate the VMS on AWS as Amazon EC2 instance
J. Install all the required software package
K. Create an Amazon FSx for Lustre file syste
L. Configure AWS DataSync to copy the NFS server data to the FSx for Lustre file system over the Direct Connect connection.
M. Order two AWS Snowball Edge device
N. Copy the VMS and the NFS server data to the device
O. Run VM Import/Export after the data from the devices is loaded to an Amazon S3 bucke
P. Create an Amazon Elastic File System (Amazon EFS) file syste
Q. Copy the NFS server data from Amazon S3 to the EFS file system.

Answer: B

Explanation:
This option will complete the migration to AWS in the least amount of time because it uses two AWS services that are designed to simplify and accelerate data transfers and migrations. AWS Application Migration Service (AWS MGN) is a highly automated lift-and-shift solution that helps you migrate applications from any source infrastructure that runs supported operating systems to AWS1. It replicates your source servers into your AWS account and automatically converts and launches them on AWS so you can quickly benefit from the cloud1. You can use AWS MGN to migrate your on-premises VMware VMs to AWS by configuring a connection to your VMware cluster and creating a replication job for the VMs2. This process will minimize the time-intensive, error-prone manual processes of exporting and importing VM images.
AWS DataSync is an online data movement and discovery service that simplifies and accelerates data migrations to AWS and helps you move data quickly and securely between on-premises storage, edge locations, other cloud providers, and AWS Storage3. It can transfer data between Network File System (NFS) shares, Server Message Block (SMB) shares, Hadoop Distributed File Systems (HDFS), self-managed object storage, AWS Snowcone, Amazon Simple Storage Service (Amazon S3) buckets, Amazon Elastic File System (Amazon EFS) file systems, Amazon FSx for Windows File Server file systems, Amazon FSx for Lustre file systems, Amazon FSx for OpenZFS file systems, and Amazon FSx for NetApp ONTAP file systems3. You can use AWS DataSync to copy your on-premises NFS server data to an Amazon EFS file system over the Direct Connect connection4. This process will leverage the high bandwidth and low latency of Direct Connect and the encryption and data integrity validation of DataSync.

NEW QUESTION 25
- (Exam Topic 3)
A company wants to use Amazon S3 to back up its on-premises file storage solution. The company's
on-premises file storage solution supports NFS, and the company wants its new solution to support NFS. The company wants to archive the backup files after 5 days. If the company needs archived files for disaster recovery, the company is willing to wait a few days for the retrieval of those files.
Which solution meets these requirements MOST cost-effectively?

A. Deploy an AWS Storage Gateway files gateway that is associated with an S3 bucke
B. Move the files from the on-premises file storage solution to the file gatewa
C. Create an S3 Lifecycle rule to move the file to S3 Standard-Infrequent Access (S3 Standard-IA) after 5 days.
D. Deploy an AWS Storage Gateway volume gateway that is associated with an S3 bucke
E. Move the filesfrom the on-premises file storage solution to the volume gatewa
F. Create an S3 Lifecycle rule to move the files to S3 Glacier Deep Archive after 5 days.
G. Deploy an AWS Storage Gateway tape gateway that is associated with an S3 bucke
H. Move the files from the on-premises file storage solution to the tape gatewa
I. Create an S3 Lifecycle rule to move the files to S3 Standard-Infrequent Access (S3 Standard-IA) after 5 days.
J. Deploy an AWS Storage Gateway file gateway that is associated with an S3 bucke
K. Move the files from the on-premises file storage solution to the tape gatewa
L. Create an S3 Lifecycle rule to move the files to S3 Standard-Infrequent Access (S3 Standard-IA) after 5 days.
M. Deploy an AWS Storage Gateway file gateway that is associated with an S3 bucke
N. Move the files from the on-premises file storage solution to the file gatewa
O. Create an S3 Lifecycle rule to move the files to S3 Glacier Deep Archive after 5 days.

Answer: E

NEW QUESTION 26
- (Exam Topic 3)
A company has developed a mobile game. The backend for the game runs on several virtual machines located in an on-premises data center. The business logic is exposed using a REST API with multiple functions. Player session data is stored in central file storage. Backend services use different API keys for throttling and to distinguish between live and test traffic.
The load on the game backend varies throughout the day. During peak hours, the server capacity is not sufficient. There are also latency issues when fetching player session data. Management has asked a solutions architect to present a cloud architecture that can handle the game's varying load and provide low-latency data access. The API model should not be changed.
Which solution meets these requirements?

A. Implement the REST API using a Network Load Balancer (NLB). Run the business logic on an Amazon EC2 instance behind the NL

B. Store player session data in Amazon Aurora Serverless.
C. Implement the REST API using an Application Load Balancer (ALB). Run the business logic in AWS Lambd
D. Store player session data in Amazon DynamoDB with on-demand capacity.
E. Implement the REST API using Amazon API Gatewa
F. Run the business logic in AWS Lambd
G. Store player session data in Amazon DynamoDB with on- demand capacity.
H. Implement the REST API using AWS AppSyn
I. Run the business logic in AWS Lambd
J. Store player session data in Amazon Aurora Serverless.

Answer: C

NEW QUESTION 27
- (Exam Topic 3)
A solutions architect must update an application environment within AWS Elastic Beanstalk using a blue/green deployment methodology The solutions architect creates an environment that is identical to the existing application environment and deploys the application to the new environment.
What should be done next to complete the update?

A. Redirect to the new environment using Amazon Route 53
B. Select the Swap Environment URLs option
C. Replace the Auto Scaling launch configuration
D. Update the DNS records to point to the green environment

Answer: B

Explanation:
https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features.CNAMESwap.html

NEW QUESTION 28
- (Exam Topic 3)
A company uses AWS Organizations to manage a multi-account structure. The company has hundreds of AWS accounts and expects the number of accounts to increase. The company is building a new application that uses Docker images. The company will push the Docker images to Amazon Elastic Container Registry (Amazon ECR). Only accounts that are within the company's organization should have access to the images.
The company has a CI/CD process that runs frequently. The company wants to retain all the tagged images. However, the company wants to retain only the five most recent untagged images.
Which solution will meet these requirements with the LEAST operational overhead?

A. Create a private repository in Amazon EC
B. Create a permissions policy for the repository that allows only required ECR operation
C. Include a condition to allow the ECR operations if the value of the aws:PrincipalOrgID condition key is equal to the ID of the company's organizatio
D. Add a lifecycle rule to the ECR repository that deletes all untagged images over the count of five.
E. Create a public repository in Amazon EC
F. Create an IAM role in the ECR accoun
G. Set permissions so that any account can assume the role if the value of the aws:PrincipalOrgID condition key is equal to the ID of the company's organizatio
H. Add a lifecycle rule to the ECR repository that deletes all untagged images over the count of five.
I. Create a private repository in Amazon EC
J. Create a permissions policy for the repository that includes only required ECR operation
K. Include a condition to allow the ECR operations for all account IDs in the organizatio
L. Schedule a daily Amazon EventBridge rule to invoke an AWS Lambda function that deletes all untagged images over the count of five.
M. Create a public repository in Amazon EC
N. Configure Amazon ECR to use an interface VPC endpoint with an endpoint policy that includes the required permissions for images that the company needs to pul
O. Include a condition to allow the ECR operations for all account IDs in the company's organizatio
P. Schedule a daily Amazon EventBridge rule to invoke an AWS Lambda function that deletes all untagged images over the count of five.

Answer: A

Explanation:
This option allows the company to use a private repository in Amazon ECR to store and manage its Docker images securely and efficiently1. By creating a permissions policy for the repository that allows only required ECR operations, such as ecr:GetDownloadUrlForLayer, ecr:BatchGetImage, ecr:BatchCheckLayerAvailability, ecr:PutImage, and ecr:InitiateLayerUpload2, the company can restrict access to the repository and prevent unauthorized actions. By including a condition to allow the ECR operations if the value of the aws:PrincipalOrgID condition key is equal to the ID of the company's organization, the company can ensure that only accounts that are within its organization can access the images3. By adding a lifecycle rule to the ECR repository that deletes all untagged images over the count of five, the company can reduce storage costs and retain only the most recent untagged images4.
References:
- Amazon ECR private repositories
- Amazon ECR repository policies
- Restricting access to AWS Organizations members
- Amazon ECR lifecycle policies

NEW QUESTION 29
- (Exam Topic 2)
A company consists of two separate business units. Each business unit has its own AWS account within a single organization in AWS Organizations. The business units regularly share sensitive documents with each other. To facilitate sharing, the company created an Amazon S3 bucket in each account and configured two-way replication between the S3 buckets. The S3 buckets have millions of objects.
Recently, a security audit identified that neither S3 bucket has encryption at rest enabled. Company policy requires that all documents must be stored with encryption at rest. The company wants to implement server-side encryption with Amazon S3 managed encryption keys (SSE-S3).
What is the MOST operationally efficient solution that meets these requirements?

A. Turn on SSE-S3 on both S3 bucket
B. Use S3 Batch Operations to copy and encrypt the objects in the same location.
C. Create an AWS Key Management Service (AWS KMS) key in each accoun
D. Turn on server-side encryption with AWS KMS keys (SSE-KMS) on each S3 bucket by using the corresponding KMS key in that AWS accoun
E. Encrypt the existing objects by using an S3 copy command in the AWS CLI.
F. Turn on SSE-S3 on both S3 bucket
G. Encrypt the existing objects by using an S3 copy command in the AWS CLI.
H. Create an AWS Key Management Service (AWS KMS) key in each accoun
I. Turn on server-side encryption with AWS KMS keys (SSE-KMS) on each S3 bucket by using the corresponding KMS key in that AWS accoun
J. Use S3 Batch Operations to copy the objects into the same location.

Answer: A

Explanation:
"The S3 buckets have millions of objects" If there are million of objects then you should use Batch operations. https://aws.amazon.com/blogs/storage/encrypting-objects-with-amazon-s3-batch-operations/

NEW QUESTION 30
- (Exam Topic 2)
A company is providing weather data over a REST-based API to several customers. The API is hosted by Amazon API Gateway and is integrated with different AWS Lambda functions for each API operation. The company uses Amazon Route 53 for DNS and has created a resource record of weather.example.com. The company stores data for the API in Amazon DynamoDB tables. The company needs a solution that will give the API the ability to fail over to a different AWS Region.
Which solution will meet these requirements?

A. Deploy a new set of Lambda functions in a new Regio
B. Update the API Gateway API to use an edge-optimized API endpoint with Lambda functions from both Regions as target
C. Convert the DynamoDB tables to global tables.
D. Deploy a new API Gateway API and Lambda functions in another Regio
E. Change the Route 53 DNS record to a multivalue answe
F. Add both API Gateway APIs to the answe
G. Enable target health monitorin
H. Convert the DynamoDB tables to global tables.
I. Deploy a new API Gateway API and Lambda functions in another Regio
J. Change the Route 53 DNS record to a failover recor
K. Enable target health monitorin
L. Convert the DynamoDB tables to global tables.
M. Deploy a new API Gateway API in a new Regio
N. Change the Lambda functions to global functions.Change the Route 53 DNS record to a multivalue answe
O. Add both API Gateway APIs to the answe
P. Enable target health monitorin
Q. Convert the DynamoDB tables to global tables.

Answer: C

Explanation:
https://docs.aws.amazon.com/apigateway/latest/developerguide/dns-failover.html

NEW QUESTION 31
......