Amazon

Exam Questions AWS-Certified-Solutions-Architect-Professional

Amazon AWS Certified Solutions Architect Professional

NEW QUESTION 1
- (Exam Topic 2)
A company has set up its entire infrastructure on AWS. The company uses Amazon EC2 instances to host its ecommerce website and uses Amazon S3 to store static data. Three engineers at the company handle the cloud administration and development through one AWS account. Occasionally, an engineer alters an EC2 security group configuration of another engineer and causes noncompliance issues in the environment.
A solutions architect must set up a system that tracks changes that the engineers make. The system must send alerts when the engineers make noncompliant changes to the security settings for the EC2 instances.
What is the FASTEST way for the solutions architect to meet these requirements?

A. Set up AWS Organizations for the compan
B. Apply SCPs to govern and track noncompliant security group changes that are made to the AWS account.
C. Enable AWS CloudTrail to capture the changes to EC2 security group
D. Enable Amazon CtoudWatch rules to provide alerts when noncompliant security settings are detected.
E. Enable SCPs on the AWS account to provide alerts when noncompliant security group changes are made to the environment.
F. Enable AWS Config on the EC2 security groups to track any noncompliant changes Send the changes as alerts through an Amazon Simple Notification Service (Amazon SNS) topic.

Answer: D

Explanation:
https://aws.amazon.com/es/blogs/industries/how-to-monitor-alert-and-remediate-non-compliant-hipaa-findings

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

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

Answer: B

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

NEW QUESTION 3
- (Exam Topic 2)
A company has deployed its database on an Amazon RDS for MySQL DB instance in the us-east-1 Region. The company needs to make its data available to customers in Europe. The customers in Europe must have access to the same data as customers in the United States (US) and will not tolerate high application latency or stale data. The customers in Europe and the customers in the US need to write to the database. Both groups of customers need to see updates from the other group in real time.
Which solution will meet these requirements?

A. Create an Amazon Aurora MySQL replica of the RDS for MySQL DB instanc
B. Pause application writes to the RDS DB instanc
C. Promote the Aurora Replica to a standalone DB cluste
D. Reconfigure the application to use the Aurora database and resume write
E. Add eu-west-1 as a secondary Region to the 06 cluste
F. Enable write forwarding on the DB cluste
G. Deploy the application in eu-west-1. Configure the application to use the Aurora MySQL endpoint in eu- west-1.
H. Add a cross-Region replica in eu-west-1 for the RDS for MySQL DB instanc
I. Configure the replica to replicate write queries back to the primary DB instanc
J. Deploy the application in eu-west-1. Configure the application to use the RDS for MySQL endpoint in eu-west-1.
K. Copy the most recent snapshot from the RDS for MySQL DB instance to eu-west-1. Create a new RDS for MySQL DB instance in eu-west-1 from the snapsho
L. Configure MySQL logical replication fromus-east-1 to eu-west-1. Enable write forwarding on the DB cluste
M. Deploy the application in eu-west-1. Configure the application to use the RDS for MySQL endpoint in eu-west-1.
N. Convert the RDS for MySQL DB instance to an Amazon Aurora MySQL DB cluste
O. Add eu-west-1 as a secondary Region to the DB cluste

P. Enable write forwarding on the DB cluste
Q. Deploy the application in eu-west-1. Configure the application to use the Aurora MySQL endpoint in eu-west-1.

Answer: D

Explanation:
The company should use AWS Amplify to create a static website for uploads of media files. The company should use Amplify Hosting to serve the website through Amazon CloudFront. The company should use Amazon S3 to store the uploaded media files. The company should use Amazon Cognito to authenticate users. This solution will meet the requirements with the least operational overhead because AWS Amplify is a complete solution that lets frontend web and mobile developers easily build, ship, and host full-stack applications on AWS, with the flexibility to leverage the breadth of AWS services as use cases evolve. No cloud expertise needed1. By using AWS Amplify, the company can refactor the application to a serverless architecture that reduces operational complexity and costs. AWS Amplify offers the following features and benefits:

- Amplify Studio: A visual interface that enables you to build and deploy a full-stack app quickly, including frontend UI and backend.
- Amplify CLI: A local toolchain that enables you to configure and manage an app backend with just a few commands.
- Amplify Libraries: Open-source client libraries that enable you to build cloud-powered mobile and web apps.
- Amplify UI Components: Open-source design system with cloud-connected components for building feature-rich apps fast.
- Amplify Hosting: Fully managed CI/CD and hosting for fast, secure, and reliable static and server-side rendered apps.

By using AWS Amplify to create a static website for uploads of media files, the company can leverage Amplify Studio to visually build a pixel-perfect UI and connect it to a cloud backend in clicks. By using Amplify Hosting to serve the website through Amazon CloudFront, the company can easily deploy its web app or website to the fast, secure, and reliable AWS content delivery network (CDN), with hundreds of points of presence globally. By using Amazon S3 to store the uploaded media files, the company can benefit from a highly scalable, durable, and cost-effective object storage service that can handle any amount of data2. By using Amazon Cognito to authenticate users, the company can add user sign-up, sign-in, and access control to its web app with a fully managed service that scales to support millions of users3.
The other options are not correct because:

- Using AWS Application Migration Service to migrate the application server to Amazon EC2 instances would not refactor the application or accelerate development. AWS Application Migration Service (AWS MGN) is a service that enables you to migrate physical servers, virtual machines (VMs), or cloud servers from any source infrastructure to AWS without requiring agents or specialized tools. However, this would not address the challenges of overutilization and data uploads failures. It would also not reduce operational overhead or costs compared to a serverless architecture.
- Creating a static website for uploads of media files and using AWS AppSync to create an API would not be as simple or fast as using AWS Amplify. AWS AppSync is a service that enables you to create flexible APIs for securely accessing, manipulating, and combining data from one or more data sources. However, this would require more configuration and management than using Amplify Studio and Amplify Hosting. It would also not provide authentication features like Amazon Cognito.
- Setting up AWS IAM Identity Center (AWS Single Sign-On) to give users the ability to sign in to the application would not be as suitable as using Amazon Cognito. AWS Single Sign-On (AWS SSO) is a
service that enables you to centrally manage SSO access and user permissions across multiple AWS accounts and business applications. However, this service is designed for enterprise customers who need to manage access for employees or partners across multiple resources. It is not intended for authenticating end users of web or mobile apps.

References:

- https://aws.amazon.com/amplify/
- https://aws.amazon.com/s3/
- https://aws.amazon.com/cognito/
- https://aws.amazon.com/mgn/
- https://aws.amazon.com/appsync/
- https://aws.amazon.com/single-sign-on/

NEW QUESTION 4
- (Exam Topic 2)
A company is running a two-tier web-based application in an on-premises data center. The application layer consists of a single server running a stateful application. The application connects to a PostgreSQL database running on a separate server. The application's user base is expected to grow significantly, so the company is migrating the application and database to AWS. The solution will use Amazon Aurora PostgreSQL, Amazon EC2 Auto Scaling, and Elastic Load Balancing.
Which solution will provide a consistent user experience that will allow the application and database tiers to scale?

A. Enable Aurora Auto Scaling for Aurora Replica
B. Use a Network Load Balancer with the least outstanding requests routing algorithm and sticky sessions enabled.
C. Enable Aurora Auto Scaling for Aurora writer
D. Use an Application Load Balancer with the round robin routing algorithm and sticky sessions enabled.
E. Enable Aurora Auto Scaling for Aurora Replica
F. Use an Application Load Balancer with the round robin routing and sticky sessions enabled.
G. Enable Aurora Scaling for Aurora writer
H. Use a Network Load Balancer with the least outstanding requests routing algorithm and sticky sessions enabled.

Answer: C

Explanation:
Aurora Auto Scaling enables your Aurora DB cluster to handle sudden increases in connectivity or workload. When the connectivity or workload decreases, Aurora Auto Scaling removes unnecessary Aurora Replicas so that you don't pay for unused provisioned DB instances

NEW QUESTION 5
- (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 6
- (Exam Topic 2)
A company needs to optimize the cost of an AWS environment that contains multiple accounts in an organization in AWS Organizations. The company conducted cost optimization activities 3 years ago and purchased Amazon EC2 Standard Reserved Instances that recently expired.
The company needs EC2 instances for 3 more years. Additionally, the company has deployed a new serverless workload.
Which strategy will provide the company with the MOST cost savings?

A. Purchase the same Reserved Instances for an additional 3-year term with All Upfront paymen
B. Purchase a 3-year Compute Savings Plan with All Upfront payment in the management account to cover any additional compute costs.
C. Purchase a l-year Compute Savings Plan with No Upfront payment in each member accoun
D. Use the Savings Plans recommendations in the AWS Cost Management console to choose the Compute Savings Plan.
E. Purchase a 3-year EC2 Instance Savings Plan with No Upfront payment in the management account to cover EC2 costs in each AWS Regio
F. Purchase a 3 year Compute Savings Plan with No Upfront payment in the management account to cover any additional compute costs.
G. Purchase a 3-year EC2 Instance Savings Plan with All Upfront payment in each member accoun
H. Use the Savings Plans recommendations in the AWS Cost Management console to choose the EC2 Instance Savings Plan.

Answer: A

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

- Purchasing a 1-year Compute Savings Plan with No Upfront payment in each member account would not provide as much cost savings as purchasing a 3-year Compute Savings Plan with All Upfront payment in the management account. A 1-year term offers lower discounts than a 3-year term, and a No Upfront payment option offers lower discounts than an All Upfront payment option. Also, purchasing a Savings Plan in each member account would not allow the company to share the benefits of unused Savings Plan discounts across its organization.
- Purchasing a 3-year EC2 Instance Savings Plan with No Upfront payment in the management account to cover EC2 costs in each AWS Region would not provide as much cost savings as purchasing Reserved Instances for an additional 3-year term with All Upfront payment. An EC2 Instance Savings Plan offers lower discounts than Reserved Instances for the same instance family and Region. Also, a No Upfront payment option offers lower discounts than an All Upfront payment option.
- Purchasing a 3-year EC2 Instance Savings Plan with All Upfront payment in each member account would not provide as much flexibility or cost savings as purchasing a 3-year Compute Savings Plan with All Upfront payment in the management account. An EC2 Instance Savings Plan applies only to a specific instance family within a Region and does not cover Fargate or Lambda usage. Also, purchasing a Savings Plan in each member account would not allow the company to share the benefits of unused Savings Plan discounts across its organization.

References:

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

NEW QUESTION 7
- (Exam Topic 2)
A solutions architect is designing an AWS account structure for a company that consists of multiple teams. All the teams will work in the same AWS Region. The company needs a VPC that is connected to the on-premises network. The company expects less than 50 Mbps of total traffic to and from the on-premises network.
Which combination of steps will meet these requirements MOST cost-effectively? (Select TWO.)

A. Create an AWS Cloud Formation template that provisions a VPC and the required subnet
B. Deploy the template to each AWS account.
C. Create an AWS Cloud Formation template that provisions a VPC and the required subnet
D. Deploy the template to a shared services account Share the subnets by using AWS Resource Access Manager.
E. Use AWS Transit Gateway along with an AWS Site-to-Site VPN for connectivity to the on-premises networ
F. Share the transit gateway by using AWS Resource Access Manager.
G. Use AWS Site-to-Site VPN for connectivity to the on-premises network.
H. Use AWS Direct Connect for connectivity to the on-premises network.

Answer: BD

NEW QUESTION 8

- (Exam Topic 2)
A solutions architect is redesigning a three-tier application that a company hosts on premises. The application provides personalized recommendations based on user profiles. The company already has an AWS account and has configured a VPC to host the application.
The frontend is a Java-based application that runs in on-premises VMs. The company hosts a personalization model on a physical application server and uses TensorFlow to implement the model. The personalization model uses artificial intelligence and machine learning (AI/ML). The company stores user information in a Microsoft SQL Server database. The web application calls the personalization model, which reads the user profiles from the database and provides recommendations.
The company wants to migrate the redesigned application to AWS.
Which solution will meet this requirement with the LEAST operational overhead?

A. Use AWS Server Migration Service (AWS SMS) to migrate the on-premises physical application server and the web application VMs to AW
B. Use AWS Database Migration Service (AWS DMS) to migrate the SQL Server database to Amazon RDS for SQL Server.
C. Export the personalization mode
D. Store the model artifacts in Amazon S3. Deploy the model to Amazon SageMaker and create an endpoin
E. Host the Java application in AWS Elastic Beanstal
F. Use AWS Database Migration Service {AWS DMS) to migrate the SQL Server database to Amazon RDS for SQL Server.
G. Use AWS Application Migration Service to migrate the on-premises personalization model and VMs to Amazon EC2 instances in Auto Scaling group
H. Use AWS Database Migration Service (AWS DMS) to migrate the SQL Server database to an EC2 instance.
I. Containerize the personalization model and the Java applicatio
J. Use Amazon Elastic Kubernetes Service (Amazon EKS) managed node groups to deploy the model and the application to Amazon EKS Host the node groups in a VP
K. Use AWS Database Migration Service (AWS DMS) to migrate the SQL Server database to Amazon RDS for SQL Server.

Answer: B

Explanation:
Amazon SageMaker is a fully managed machine learning service that allows users to build, train, and deploy machine learning models quickly and easily1. Users can export their existing TensorFlow models and store the model artifacts in Amazon S3, a highly scalable and durable object storage service2. Users can then deploy the model to Amazon SageMaker and create an endpoint that can be invoked by the web application to provide recommendations3. This way, the solution can leverage the AI/ML capabilities of Amazon SageMaker without having to rewrite the personalization model.
AWS Elastic Beanstalk is a service that allows users to deploy and manage web applications without worrying about the infrastructure that runs those applications. Users can host their Java application in AWS Elastic Beanstalk and configure it to communicate with the Amazon SageMaker endpoint. This way, the solution can reduce the operational overhead of managing servers, load balancers, scaling, and application health monitoring.
AWS Database Migration Service (AWS DMS) is a service that helps users migrate databases to AWS quickly and securely. Users can use AWS DMS to migrate their SQL Server database to Amazon RDS for SQL Server, a fully managed relational database service that offers high availability, scalability, security, and compatibility. This way, the solution can reduce the operational overhead of managing database servers, backups, patches, and upgrades.
Option A is incorrect because using AWS Server Migration Service (AWS SMS) to migrate the on-premises physical application server and the web application VMs to AWS is not cost-effective or scalable. AWS SMS is a service that helps users migrate on-premises workloads to AWS. However, for this use case, migrating the physical application server and the web application VMs to AWS will not take advantage of the AI/ML capabilities of Amazon SageMaker or the managed services of AWS Elastic Beanstalk and Amazon RDS.
Option C is incorrect because using AWS Application Migration Service to migrate the on-premises personalization model and VMs to Amazon EC2 instances in Auto Scaling groups is not cost-effective or scalable. AWS Application Migration Service is a service that helps users migrate applications from on-premises or other clouds to AWS without making any changes to their applications. However, for this use case, migrating the personalization model and VMs to EC2 instances will not take advantage of the AI/ML capabilities of Amazon SageMaker or the managed services of AWS Elastic Beanstalk and Amazon RDS.
Option D is incorrect because containerizing the personalization model and the Java application and using Amazon Elastic Kubernetes Service (Amazon EKS) managed node groups to deploy them to Amazon EKS is not necessary or cost-effective. Amazon EKS is a service that allows users to run Kubernetes on AWS without needing to install, operate, and maintain their own Kubernetes control plane or nodes. However, for this use case, containerizing and deploying the personalization model and the Java application will not take advantage of the AI/ML capabilities of Amazon SageMaker or the managed services of AWS Elastic Beanstalk. Moreover, using S3 Glacier Deep Archive as a storage class for images will incur a high retrieval fee and latency for accessing them.

NEW QUESTION 9
- (Exam Topic 2)
A company wants to refactor its retail ordering web application that currently has a load-balanced Amazon EC2 instance fleet for web hosting, database API services, and business logic. The company needs to create a decoupled, scalable architecture with a mechanism for retaining failed orders while also minimizing operational costs.
Which solution will meet these requirements?

A. Use Amazon S3 for web hosting with Amazon API Gateway for database API service
B. Use Amazon Simple Queue Service (Amazon SQS) for order queuin
C. Use Amazon Elastic Container Service (Amazon ECS) for business logic with Amazon SQS long polling for retaining failed orders.
D. Use AWS Elastic Beanstalk for web hosting with Amazon API Gateway for database API service
E. Use Amazon MQ for order queuin
F. Use AWS Step Functionsfor business logic with Amazon S3 Glacier Deep Archive for retaining failed orders.
G. Use Amazon S3 for web hosting with AWS AppSync for database API service
H. Use Amazon Simple Queue Service (Amazon SQS) for order queuin
I. Use AWS Lambda for business logic with an Amazon SQS dead-letter queue for retaining failed orders.
J. Use Amazon Lightsail for web hosting with AWS AppSync for database API service
K. Use Amazon Simple Email Service (Amazon SES) for order queuin
L. UseAmazon Elastic Kubernetes Service (Amazon EKS) for business logic with Amazon OpenSearch Service for retaining failed orders.

Answer: C

Explanation:
•Use Amazon S3 for web hosting with AWS AppSync for database API services. Use Amazon Simple Queue Service (Amazon SQS) for order queuing. Use AWS Lambda for business logic with an Amazon SQS dead-letter queue for retaining failed orders.
This solution will allow you to:
•Host a static website on Amazon S3 without provisioning or managing servers1.
•Use AWS AppSync to create a scalable GraphQL API that connects to your database and other data sources1.
•Use Amazon SQS to decouple and scale your order processing microservices1.
•Use AWS Lambda to run code for your business logic without provisioning or managing servers1.
•Use an Amazon SQS dead-letter queue to retain messages that can’t be processed by your Lambda function1.

NEW QUESTION 10
- (Exam Topic 2)
A company wants to use AWS for disaster recovery for an on-premises application. The company has hundreds of Windows-based servers that run the application. All the servers mount a common share.
The company has an RTO of 15 minutes and an RPO of 5 minutes. The solution must support native failover and fallback capabilities.
Which solution will meet these requirements MOST cost-effectively?

A. Create an AWS Storage Gateway File Gatewa
B. Schedule daily Windows server backup
C. Save the data lo Amazon S3. During a disaster, recover the on-premises servers from the backu
D. During failbac
E. run the on-premises servers on Amazon EC2 instances.
F. Create a set of AWS CloudFormation templates to create infrastructur
G. Replicate all data to Amazon Elastic File System (Amazon EFS) by using AWS DataSyn
H. During a disaster, use AWS CodePipeline to deploy the templates to restore the on-premises server
I. Fail back the data by using DataSync.
J. Create an AWS Cloud Development Kit (AWS CDK) pipeline to stand up a multi-site active-active environment on AW
K. Replicate data into Amazon S3 by using the s3 sync comman
L. During a disaster, swap DNS endpoints to point to AW
M. Fail back the data by using the s3 sync command.
N. Use AWS Elastic Disaster Recovery to replicate the on-premises server
O. Replicate data to an Amazon FSx for Windows File Server file system by using AWS DataSyn
P. Mount the file system to AWS server
Q. During a disaster, fail over the on-premises servers to AW
R. Fail back to new or existing servers by using Elastic Disaster Recovery.

Answer: D

NEW QUESTION 10
- (Exam Topic 2)
A solutions architect needs to improve an application that is hosted in the AWS Cloud. The application uses an Amazon Aurora MySQL DB instance that is experiencing overloaded connections. Most of the application's operations insert records into the database. The application currently stores credentials in a text-based configuration file.
The solutions architect needs to implement a solution so that the application can handle the current connection load. The solution must keep the credentials secure and must provide the ability to rotate the credentials automatically on a regular basis.
Which solution will meet these requirements?

A. Deploy an Amazon RDS Proxy layer in front of the DB instanc
B. Store the connection credentials as a secret in AWS Secrets Manager.
C. Deploy an Amazon RDS Proxy layer in front of the DB instanc
D. Store the connection credentials in AWS Systems Manager Parameter Store.
E. Create an Aurora Replic
F. Store the connection credentials as a secret in AWS Secrets Manager.
G. Create an Aurora Replic
H. Store the connection credentials in AWS Systems Manager Parameter Store.

Answer: A

Explanation:
https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/rds-proxy.html

NEW QUESTION 15
- (Exam Topic 2)
A company has an application that runs on Amazon EC2 instances in an Amazon EC2 Auto Scaling group. The company uses AWS CodePipeline to deploy the application. The instances that run in the Auto Scaling group are constantly changing because of scaling events.
When the company deploys new application code versions, the company installs the AWS CodeDeploy agent on any new target EC2 instances and associates the instances with the CodeDeploy deployment group. The application is set to go live within the next 24 hours.
What should a solutions architect recommend to automate the application deployment process with the LEAST amount of operational overhead?

A. Configure Amazon EventBridge to invoke an AWS Lambda function when a new EC2 instance is launched into the Auto Scaling grou
B. Code the Lambda function to associate the EC2 instances with the CodeDeploy deployment group.
C. Write a script to suspend Amazon EC2 Auto Scaling operations before the deployment of new code When the deployment is complete, create a new AMI and configure the Auto Scaling group's launch template to use the new AMI for new launche
D. Resume Amazon EC2 Auto Scaling operations.
E. Create a new AWS CodeBuild project that creates a new AMI that contains the new code Configure CodeBuild to update the Auto Scaling group's launch template to the new AM
F. Run an Amazon EC2 Auto Scaling instance refresh operation.
G. Create a new AMI that has the CodeDeploy agent installe
H. Configure the Auto Scaling group's launch template to use the new AM
I. Associate the CodeDeploy deployment group with the Auto Scaling group instead of the EC2 instances.

Answer: D

Explanation:
https://docs.aws.amazon.com/codedeploy/latest/userguide/integrations-aws-auto-scaling.html

NEW QUESTION 16
- (Exam Topic 2)
A company recently started hosting new application workloads in the AWS Cloud. The company is using Amazon EC2 instances, Amazon Elastic File System (Amazon EFS) file systems, and Amazon RDS DB instances.

To meet regulatory and business requirements, the company must make the following changes for data backups:
* Backups must be retained based on custom daily, weekly, and monthly requirements.
* Backups must be replicated to at least one other AWS Region immediately after capture.
* The backup solution must provide a single source of backup status across the AWS environment.
* The backup solution must send immediate notifications upon failure of any resource backup.
Which combination of steps will meet this requirement with the LEAST amount of operational overhead? (Select THREE.)

A. Create an AWS Backup plan with a backup rule for each of the retention requirements.
B. Configure an AWS backup plan to copy backups to another Region.
C. Create an AWS Lambda function to replicate backups to another Region and send notification if a failure occurs.
D. Add an Amazon Simple Notification Service (Amazon SNS) topic to the backup plan to send a notification for finished jobs that have any status except BACKUP- JOB- COMPLETED.
E. Create an Amazon Data Lifecycle Manager (Amazon DLM) snapshot lifecycle policy for each of the retention requirements.
F. Set up RDS snapshots on each database.

Answer: ABD

Explanation:
Cross region with AWS Backup:
https://docs.aws.amazon.com/aws-backup/latest/devguide/cross-region-backup.html

NEW QUESTION 21
- (Exam Topic 2)
A company is designing a new website that hosts static content. The website will give users the ability to upload and download large files. According to company requirements, all data must be encrypted in transit and at rest. A solutions architect is building the solution by using Amazon S3 and Amazon CloudFront.
Which combination of steps will meet the encryption requirements? (Select THREE.)

A. Turn on S3 server-side encryption for the S3 bucket that the web application uses.
B. Add a policy attribute of "aws:SecureTransport": "true" for read and write operations in the S3 ACLs.
C. Create a bucket policy that denies any unencrypted operations in the S3 bucket that the web application uses.
D. Configure encryption at rest on CloudFront by using server-side encryption with AWS KMS keys (SSE-KMS).
E. Configure redirection of HTTP requests to HTTPS requests in CloudFront.
F. Use the RequireSSL option in the creation of presigned URLs for the S3 bucket that the web application uses.

Answer: ACE

Explanation:
Turning on S3 server-side encryption for the S3 bucket that the web application uses will enable encrypting the data at rest using Amazon S3 managed keys (SSE-S3)1. Creating a bucket policy that denies any unencrypted operations in the S3 bucket that the web application uses will enable enforcing encryption for all requests to the bucket2. Configuring redirection of HTTP requests to HTTPS requests in CloudFront will enable encrypting the data in transit using SSL/TLS3.

NEW QUESTION 25
- (Exam Topic 2)
A company is storing sensitive data in an Amazon S3 bucket. The company must log all activities for objects in the S3 bucket and must keep the logs for 5 years. The company's security team also must receive an email notification every time there is an attempt to delete data in the S3 bucket.
Which combination of steps will meet these requirements MOST cost-effectively? (Select THREE.)

A. Configure AWS CloudTrail to log S3 data events.
B. Configure S3 server access logging for the S3 bucket.
C. Configure Amazon S3 to send object deletion events to Amazon Simple Email Service (Amazon SES).
D. Configure Amazon S3 to send object deletion events to an Amazon EventBridge event bus that publishes to an Amazon Simple Notification Service (Amazon SNS) topic.
E. Configure Amazon S3 to send the logs to Amazon Timestream with data storage tiering.
F. Configure a new S3 bucket to store the logs with an S3 Lifecycle policy.

Answer: ADF

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

NEW QUESTION 30
- (Exam Topic 2)
A company runs an IoT application in the AWS Cloud. The company has millions of sensors that collect data from houses in the United States. The sensors use the MOTT protocol to connect and send data to a custom MQTT broker. The MQTT broker stores the data on a single Amazon EC2 instance. The sensors connect to the broker through the domain named iot.example.com. The company uses Amazon Route 53 as its DNS service. The company stores the data in Amazon DynamoDB.
On several occasions, the amount of data has overloaded the MOTT broker and has resulted in lost sensor data. The company must improve the reliability of the solution.
Which solution will meet these requirements?

A. Create an Application Load Balancer (ALB) and an Auto Scaling group for the MOTT broke
B. Use the Auto Scaling group as the target for the AL
C. Update the DNS record in Route 53 to an alias recor
D. Point the alias record to the AL
E. Use the MQTT broker to store the data.

F. Set up AWS loT Core to receive the sensor dat
G. Create and configure a custom domain to connect to AWS loT Cor
H. Update the DNS record in Route 53 to point to the AWS loT Core Data-ATS endpoin
I. Configure an AWS loT rule to store the data.
J. Create a Network Load Balancer (NLB). Set the MQTT broker as the targe
K. Create an AWS Global Accelerator accelerato
L. Set the NLB as the endpoint for the accelerato
M. Update the DNS record in Route 53 to a multivalue answer recor
N. Set the Global Accelerator IP addresses as value
O. Use the MQTT broker to store the data.
P. Set up AWS loT Greengrass to receive the sensor dat
Q. Update the DNS record in Route 53 to point to the AWS loT Greengrass endpoin
R. Configure an AWS loT rule to invoke an AWS Lambda function to store the data.

Answer: A

Explanation:
it describes a solution that uses an Application Load Balancer (ALB) and an Auto Scaling group for the MQTT broker. The ALB distributes incoming traffic across the instances in the Auto Scaling group and allows for automatic scaling based on incoming traffic. The use of an alias record in Route 53 allows for easy updates to the DNS record without changing the IP address. This solution improves the reliability of the MQTT broker by allowing it to automatically scale based on incoming traffic, reducing the likelihood of lost data due to broker overload.
Reference: https://aws.amazon.com/elasticloadbalancing/applicationloadbalancer/ https://aws.amazon.com/autoscaling/ https://aws.amazon.com/route53/

NEW QUESTION 33
- (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 38
- (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 41
- (Exam Topic 2)
A large company runs workloads in VPCs that are deployed across hundreds of AWS accounts. Each VPC consists to public subnets and private subnets that span across multiple Availability Zones. NAT gateways are deployed in the public subnets and allow outbound connectivity to the internet from the private subnets. A solutions architect is working on a hub-and-spoke design. All private subnets in the spoke VPCs must route traffic to the internet through an egress VPC. The solutions architect already has deployed a NAT gateway in an egress VPC in a central AWS account.
Which set of additional steps should the solutions architect take to meet these requirements?

A. Create peering connections between the egress VPC and the spoke VPC
B. Configure the required routing to allow access to the internet.
C. Create a transit gateway, and share it with the existing AWS account
D. Attach existing VPCs to the transit gateway Configure the required routing to allow access to the internet.
E. Create a transit gateway in every accoun
F. Attach the NAT gateway to the transit gateway
G. Configure the required routing to allow access to the internet.
H. Create an AWS PrivateLink connection between the egress VPC and the spoke VPC
I. Configure the required routing to allow access to the internet

Answer: B

Explanation:
https://d1.awsstatic.com/architecture-diagrams/ArchitectureDiagrams/NAT-gateway-centralized-egress-ra.pdf?d

NEW QUESTION 44
- (Exam Topic 2)
A company needs to migrate its customer transactions database from on premises to AWS. The database resides on an Oracle DB instance that runs on a Linux server. According to a new security requirement, the company must rotate the database password each year.
Which solution will meet these requirements with the LEAST operational overhead?

A. Convert the database to Amazon DynamoDB by using the AWS Schema Conversion Tool (AWS SCT).Store the password in AWS Systems Manager Parameter Stor
B. Create an Amazon CloudWatch alarm to invoke an AWS Lambda function for yearly password rotation.
C. Migrate the database to Amazon RDS for Oracl
D. Store the password in AWS Secrets Manage
E. Turn on automatic rotatio
F. Configure a yearly rotation schedule.
G. Migrate the database to an Amazon EC2 instanc
H. Use AWS Systems Manager Parameter Store to keep and rotate the connection string by using an AWS Lambda function on a yearly schedule
I. Migrate the database to Amazon Neptune by using the AWS Schema Conversion Tool {AWS SCT).Create an Amazon CloudWatch alarm to invoke an AWS Lambda function for yearly password rotation.

Answer: B

NEW QUESTION 45
- (Exam Topic 2)
A company is using AWS Organizations to manage multiple AWS accounts. For security purposes, the company requires the creation of an Amazon Simple Notification Service (Amazon SNS) topic that enables integration with a third-party alerting system in all the Organizations member accounts.
A solutions architect used an AWS CloudFormation template to create the SNS topic and stack sets to automate the deployment of Cloud Formation stacks. Trusted access has been enabled in Organizations.
What should the solutions architect do to deploy the CloudFormation StackSets in all AWS accounts?

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

Answer: C

Explanation:
https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/stacksets-orgs-manage-auto-deployment.h

NEW QUESTION 46
- (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 51
- (Exam Topic 2)
A company has a new application that needs to run on five Amazon EC2 instances in a single AWS Region. The application requires high-through put. low-latency network connections between all to the EC2 instances where the application will run. There is no requirement for the application to be fault tolerant.
Which solution will meet these requirements?

A. Launch five new EC2 instances into a cluster placement grou
B. Ensure that the EC2 instance type supports enhanced networking.
C. Launch five new EC2 instances into an Auto Scaling group in the same Availability Zon
D. Attach an extra elastic network interface to each EC2 instance.
E. Launch five new EC2 instances into a partition placement grou
F. Ensure that the EC2 instance type supports enhanced networking.
G. Launch five new EC2 instances into a spread placement group Attach an extra elastic network interface to each EC2 instance.

Answer: A

Explanation:
https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/placement-groups.html#placement- groups-cluster

NEW QUESTION 53
- (Exam Topic 2)
A company is updating an application that customers use to make online orders. The number of attacks on the application by bad actors has increased recently. The company will host the updated application on an Amazon Elastic Container Service (Amazon ECS) cluster. The company will use Amazon DynamoDB to store application data. A public Application Load Balancer (ALB) will provide end users with access to the application. The company must prevent prevent attacks and ensure business continuity with minimal service interruptions during an ongoing attack.
Which combination of steps will meet these requirements MOST cost-effectively? (Select TWO.)

A. Create an Amazon CloudFront distribution with the ALB as the origi
B. Add a custom header and random value on the CloudFront domai
C. Configure the ALB to conditionally forward traffic if the header and value match.
D. Deploy the application in two AWS Region
E. Configure Amazon Route 53 to route to both Regions with equal weight.
F. Configure auto scaling for Amazon ECS task
G. Create a DynamoDB Accelerator (DAX) cluster.
H. Configure Amazon ElastiCache to reduce overhead on DynamoDB.
I. Deploy an AWS WAF web ACL that includes an appropriate rule grou
J. Associate the web ACL with the Amazon CloudFront distribution.

Answer: AE

Explanation:
The company should create an Amazon CloudFront distribution with the ALB as the origin. The company should add a custom header and random value on the CloudFront domain. The company should configure the ALB to conditionally forward traffic if the header and value match. The company should also deploy an AWS WAF web ACL that includes an appropriate rule group. The company should associate the web ACL with the Amazon CloudFront distribution. This solution will meet the requirements most cost-effectively because Amazon CloudFront is a fast content delivery network (CDN) service that securely delivers data, videos, applications, and APIs to customers globally with low latency, high transfer speeds, all within a
developer-friendly environment1. By creating an Amazon CloudFront distribution with the ALB as the origin, the company can improve the performance and availability of its application by caching static content at edge locations closer to end users. By adding a custom header and random value on the CloudFront domain, the company can prevent direct access to the ALB and ensure that only requests from CloudFront are forwarded to the ECS tasks. By configuring the ALB to conditionally forward traffic if the header and value match, the company can implement origin access identity (OAI) for its ALB origin. OAI is a feature that enables you to restrict access to your content by requiring users to access your content through CloudFront URLs2. By deploying an AWS WAF web ACL that includes an appropriate rule group, the company can prevent attacks and ensure business continuity with minimal service interruptions during an ongoing attack. 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 blocked3. By associating the web ACL with the Amazon CloudFront distribution, the company can apply the web security rules to all requests that are forwarded by CloudFront.
The other options are not correct because:
Deploying the application in two AWS Regions and configuring Amazon Route 53 to route to both Regions with equal weight would not prevent attacks or ensure business continuity. Amazon Route 53 is a highly available and scalable cloud Domain Name System (DNS) web service that routes end users to Internet applications by translating names like www.example.com into numeric IP addresses4. However, routing traffic to multiple Regions would not protect against attacks or provide failover in case of an outage. It would also increase operational complexity and costs compared to using CloudFront and
AWS WAF.

- Configuring auto scaling for Amazon ECS tasks and creating a DynamoDB Accelerator (DAX) cluster would not prevent attacks or ensure business continuity. Auto scaling is a feature that enables you to automatically adjust your ECS tasks based on demand or a schedule. DynamoDB Accelerator (DAX) is a fully managed, highly available, in-memory cache for DynamoDB that delivers up to a 10x performance improvement. However, these features would not protect against attacks or provide failover in case of an outage. They would also increase operational complexity and costs compared to using CloudFront and AWS WAF.
- Configuring Amazon ElastiCache to reduce overhead on DynamoDB would not prevent attacks or ensure business continuity. Amazon ElastiCache is a fully managed in-memory data store service that makes it easy to deploy, operate, and scale popular open-source compatible in-memory data stores. However, this service would not protect against attacks or provide failover in case of an outage. It would also increase operational complexity and costs compared to using CloudFront and AWS WAF.

References:

- https://aws.amazon.com/cloudfront/
- https://aws.amazon.com/waf/
- https://aws.amazon.com/route53/
- https://aws.amazon.com/dynamodb/dax/
- https://aws.amazon.com/elasticache/

NEW QUESTION 56
- (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 60
- (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 61
- (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 64
- (Exam Topic 2)
A company wants to optimize AWS data-transfer costs and compute costs across developer accounts within the company's organization in AWS Organizations Developers can configure VPCs and launch Amazon EC2 instances in a single AWS Region The EC2 instances retrieve approximately 1 TB of data each day from Amazon S3
The developer activity leads to excessive monthly data-transfer charges and NAT gateway processing charges between EC2 instances and S3 buckets, along with high compute costs The company wants to proactively enforce approved architectural patterns for any EC2 instance and VPC infrastructure that developers deploy within the AWS accounts The company does not want this enforcement to negatively affect the speed at which the developers can perform their tasks
Which solution will meet these requirements MOST cost-effectively?

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

Answer: C

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

NEW QUESTION 66
- (Exam Topic 2)
A company has multiple business units that each have separate accounts on AWS. Each business unit manages its own network with several VPCs that have CIDR ranges that overlap. The company’s marketing team has created a new internal application and wants to make the application accessible to all the other business units. The solution must use private IP addresses only.
Which solution will meet these requirements with the LEAST operational overhead?

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

Answer: C

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

NEW QUESTION 70
- (Exam Topic 2)
A company is running a containerized application in the AWS Cloud. The application is running by using Amazon Elastic Container Service (Amazon ECS) on a set of Amazon EC2 instances. The EC2 instances run in an Auto Scaling group.
The company uses Amazon Elastic Container Registry (Amazon ECR) to store its container images. When a new image version is uploaded, the new image version receives a unique tag.
The company needs a solution that inspects new image versions for common vulnerabilities and exposures. The solution must automatically delete new image tags that have Critical or High severity findings. The solution also must notify the development team when such a deletion occurs.
Which solution meets these requirements?

A. Configure scan on push on the repository Use Amazon EventBridge to invoke an AWS Step Functions state machine when a scan is complete for images that have Critical or High severity finding
B. Use the Step Functions state machine to delete the image tag for those images and to notify the development team through Amazon Simple Notification Service (Amazon SNS).
C. Configure scan on push on the repository Configure scan results to be pushed to an Amazon Simple Queue Service (Amazon SQS) queu
D. Invoke an AWS Lambda function when a new message is added to the SQS queu
E. Use the Lambda function to delete the image tag for images that have Critical or High seventy finding
F. Notify the development team by using Amazon Simple Email Service (Amazon SES).
G. Schedule an AWS Lambda function to start a manual image scan every hou
H. Configure Amazon EventBridge to invoke another Lambda function when a scan is complet
I. Use the second Lambda function to delete the image tag for images that have Critical or High severity finding
J. Notify the development team by using Amazon Simple Notification Service (Amazon SNS).
K. Configure periodic image scan on the repositor
L. Configure scan results to be added lo an Amazon Simple Queue Service (Amazon SQS) queu
M. Invoke an AWS Step Functions state machine when a new message is added to the SQS queu
N. Use the Step Functions state machine to delete the image tag for images that have Critical or High severity finding
O. Notify the development team by using Amazon Simple Email Service (Amazon SES).

Answer: A

Explanation:
https://docs.aws.amazon.com/AmazonECR/latest/userguide/ecr-eventbridge.html "Activating an AWS Step Functions state machine"
https://docs.aws.amazon.com/step-functions/latest/dg/tutorial-creating-lambda-state-machine.html

NEW QUESTION 72
- (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 75
- (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 76
- (Exam Topic 2)
A company uses a Grafana data visualization solution that runs on a single Amazon EC2 instance to monitor the health of the company's AWS workloads. The company has invested time and effort to create dashboards that the company wants to preserve. The dashboards need to be highly available and cannot be down for longer than 10 minutes. The company needs to minimize ongoing maintenance.
Which solution will meet these requirements with the LEAST operational overhead?

A. Migrate to Amazon CloudWatch dashboard
B. Recreate the dashboards to match the existing Grafana dashboard
C. Use automatic dashboards where possible.
D. Create an Amazon Managed Grafana workspac
E. Configure a new Amazon CloudWatch data source.Export dashboards from the existing Grafana instanc
F. Import the dashboards into the new workspace.
G. Create an AMI that has Grafana pre-installe
H. Store the existing dashboards in Amazon Elastic File System (Amazon EFS). Create an Auto Scaling group that uses the new AM
I. Set the Auto Scaling group's minimum, desired, and maximum number of instances to on
J. Create an Application Load Balancer that serves at least two Availability Zones.
K. Configure AWS Backup to back up the EC2 instance that runs Grafana once each hou
L. Restore the EC2 instance from the most recent snapshot in an alternate Availability Zone when required.

Answer: C

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

NEW QUESTION 79

- (Exam Topic 2)
A company is migrating its development and production workloads to a new organization in AWS Organizations. The company has created a separate member account for development and a separate member account for production. Consolidated billing is linked to the management account. In the management account, a solutions architect needs to create an 1AM user that can stop or terminate resources in both member accounts.
Which solution will meet this requirement?

A. Create an IAM user and a cross-account role in the management accoun
B. Configure the cross-account role with least privilege access to the member accounts.
C. Create an IAM user in each member accoun
D. In the management account, create a cross-account role that has least privilege acces
E. Grant the IAM users access to the cross-account role by using a trust policy.
F. Create an IAM user in the management accoun
G. In the member accounts, create an IAM group that has least privilege acces
H. Add the IAM user from the management account to each IAM group in the member accounts.
I. Create an IAM user in the management accoun
J. In the member accounts, create cross-account roles that have least privilege acces
K. Grant the IAM user access to the roles by using a trust policy.

Answer: D

Explanation:
Cross account role should be created in destination(member) account. The role has trust entity to master account.

NEW QUESTION 82
- (Exam Topic 2)
A company is designing an AWS Organizations structure. The company wants to standardize a process to apply tags across the entire organization. The company will require tags with specific values when a user creates a new resource. Each of the company's OUs will have unique tag values.
Which solution will meet these requirements?

A. Use an SCP to deny the creation of resources that do not have the required tag
B. Create a tag policy that Includes the tag values that the company has assigned to each O
C. Attach the tag policies to the OUs.
D. Use an SCP to deny the creation of resources that do not have the required tag
E. Create a tag policy that includes the tag values that the company has assigned to each O
F. Attach the tag policies to the organization's management account.
G. Use an SCP to allow the creation of resources only when the resources have the required tag
H. Create a tag policy that includes the tag values that the company has assigned to each O
I. Attach the tag policies to the OUs.
J. Use an SCP to deny the creation of resources that do not have the required tag
K. Define the list of tags.Attach the SCP to the OUs

Answer: A

Explanation:
https://aws.amazon.com/blogs/mt/implement-aws-resource-tagging-strategy-using-aws-tag-policies-and-service

NEW QUESTION 86
- (Exam Topic 2)
A company is deploying a new web-based application and needs a storage solution for the Linux application servers. The company wants to create a single location for updates to application data for all instances. The active dataset will be up to 100 GB in size. A solutions architect has determined that peak operations will occur for 3 hours daily and will require a total of 225 MiBps of read throughput.
The solutions architect must design a Multi-AZ solution that makes a copy of the data available in another AWS Region for disaster recovery (DR). The DR copy has an RPO of less than 1 hour.
Which solution will meet these requirements?

A. Deploy a new Amazon Elastic File System (Amazon EFS) Multi-AZ file syste
B. Configure the file system for 75 MiBps of provisioned throughpu
C. Implementreplication to a file system in the DR Region.
D. Deploy a new Amazon FSx for Lustre file syste
E. Configure Bursting Throughput mode for the file syste
F. Use AWS Backup to back up the file system to the DR Region.
G. Deploy a General Purpose SSD (gp3) Amazon Elastic Block Store (Amazon EBS) volume with 225 MiBps of throughpu
H. Enable Multi-Attach for the EBSvolum
I. Use AWS Elastic Disaster Recovery to replicate the EBS volume to the DR Region.
J. Deploy an Amazon FSx for OpenZFS file system in both the production Region and the DR Region.Create an AWS DataSync scheduled task to replicate thedata from the production file system to the DR file system every 10 minutes.

Answer: A

Explanation:
The company should deploy a new Amazon Elastic File System (Amazon EFS) Multi-AZ file system. The company should configure the file system for 75 MiBps of provisioned throughput. The company should implement replication to a file system in the DR Region. This solution will meet the requirements because Amazon EFS is a serverless, fully elastic file storage service that lets you share file data without provisioning or managing storage capacity and performance. Amazon EFS is built to scale on demand to petabytes without disrupting applications, growing and shrinking automatically as you add and remove files1. By deploying a new Amazon EFS Multi-AZ file system, the company can create a single location for updates to application data for all instances. A Multi-AZ file system replicates data across multiple Availability Zones (AZs) within a Region, providing high availability and durability2. By configuring the file system for 75 MiBps of provisioned throughput, the company can ensure that it meets the peak operations requirement of 225 MiBps of read throughput. Provisioned throughput is a feature that enables you to specify a level of throughput that the file system can drive independent of the file system’s size or burst credit balance3. By implementing replication to a file system in the DR Region, the company can make a copy of the data available in another AWS Region for disaster recovery. Replication is a feature that enables you to replicate data from one EFS file system to another EFS file system across AWS Regions. The replication process has an RPO of less than 1 hour.

The other options are not correct because:

- Deploying a new Amazon FSx for Lustre file system would not provide a single location for updates to application data for all instances. Amazon FSx for Lustre is a fully managed service that provides cost-effective, high-performance storage for compute workloads. However, it does not support concurrent write access from multiple instances. Using AWS Backup to back up the file system to the DR Region would not provide real-time replication of data. AWS Backup is a service that enables you to centralize and automate data protection across AWS services. However, it does not support continuous data replication or cross-Region disaster recovery.
- Deploying a General Purpose SSD (gp3) Amazon Elastic Block Store (Amazon EBS) volume with 225 MiBps of throughput would not provide a single location for updates to application data for all instances. Amazon EBS is a service that provides persistent block storage volumes for use with Amazon EC2 instances. However, it does not support concurrent access from multiple instances, unless Multi-Attach is enabled. Enabling Multi-Attach for the EBS volume would not provide Multi-AZ resilience or cross-Region replication. Multi-Attach is a feature that enables you to attach an EBS volume to multiple EC2 instances within the same Availability Zone. Using AWS Elastic Disaster Recovery to replicate the EBS volume to the DR Region would not provide real-time replication of data. AWS Elastic Disaster Recovery (AWS DRS) is a service that enables you to orchestrate and automate disaster recovery workflows across AWS Regions. However, it does not support continuous data replication or sub-hour RPOs.
- Deploying an Amazon FSx for OpenZFS file system in both the production Region and the DR Region would not be as simple or cost-effective as using Amazon EFS. Amazon FSx for OpenZFS is a fully managed service that provides high-performance storage with strong data consistency and advanced data management features for Linux workloads. However, it requires more configuration and management than Amazon EFS, which is serverless and fully elastic. Creating an AWS DataSync scheduled task to replicate the data from the production file system to the DR file system every 10 minutes would not provide real-time replication of data. AWS DataSync is a service that enables you to transfer data between on-premises storage and AWS services, or between AWS services. However, it does not support continuous data replication or sub-minute RPOs.

References:

- https://aws.amazon.com/efs/
- https://docs.aws.amazon.com/efs/latest/ug/how-it-works.html#how-it-works-azs
- https://docs.aws.amazon.com/efs/latest/ug/performance.html#provisioned-throughput
- https://docs.aws.amazon.com/efs/latest/ug/replication.html
- https://aws.amazon.com/fsx/lustre/
- https://aws.amazon.com/backup/
- https://aws.amazon.com/ebs/
- https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-volumes-multi.html

NEW QUESTION 90
- (Exam Topic 2)
A company is building a hybrid environment that includes servers in an on-premises data center and in the AWS Cloud. The company has deployed Amazon EC2 instances in three VPCs. Each VPC is in a different AWS Region. The company has established an AWS Direct Connect connection to the data center from the Region that is closest to the data center.
The company needs the servers in the on-premises data center to have access to the EC2 instances in all three VPCs. The servers in the on-premises data center also must have access to AWS public services.
Which combination of steps will meet these requirements with the LEAST cost? (Select TWO.)

A. Create a Direct Connect gateway in the Region that is closest to the data cente
B. Attach the Direct Connect connection to the Direct Connect gatewa
C. Use the
D. Direct Connect gateway to connect the VPCs in the other two Regions.
E. Set up additional Direct Connect connections from the on-premises data center to the other two Regions.
F. Create a private VI
G. Establish an AWS Site-to-Site VPN connection over the private VIF to the VPCs in the other two Regions.
H. Create a public VI
I. Establish an AWS Site-to-Site VPN connection over the public VIF to the VPCs in the other two Regions.
J. Use VPC peering to establish a connection between the VPCs across the Region
K. Create a private VIF with the existing Direct Connect connection to connect to the peered VPCs.

Answer: AE

Explanation:
A Direct Connect gateway allows you to connect multiple VPCs across different Regions to a Direct Connect connection1. A public VIF allows you to access AWS public services such as EC21. A Site-to-Site VPN connection over the public VIF provides encryption and redundancy for the traffic between the on-premises data center and the VPCs2. This solution is cheaper than setting up additional Direct Connect connections or using a private VIF with VPC peering.

NEW QUESTION 95
- (Exam Topic 2)
A company runs an application on a fleet of Amazon EC2 instances that are in private subnets behind an internet-facing Application Load Balancer (ALB). The ALB is the origin for an Amazon CloudFront distribution. An AWS WAF web ACL that contains various AWS managed rules is associated with the CloudFront distribution.
The company needs a solution that will prevent internet traffic from directly accessing the ALB. Which solution will meet these requirements with the LEAST operational overhead?

A. Create a new web ACL that contains the same rules that the existing web ACL contain
B. Associate the new web ACL with the ALB.
C. Associate the existing web ACL with the ALB.
D. Add a security group rule to the ALB to allow traffic from the AWS managed prefix list for CloudFront only.
E. Add a security group rule to the ALB to allow only the various CloudFront IP address ranges.

Answer: C

Explanation:
https://aws.amazon.com/about-aws/whats-new/2022/02/amazon-cloudfront-managed-prefix-list/

NEW QUESTION 100
- (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 104
- (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 106
- (Exam Topic 2)
A company wants to run a custom network analysis software package to inspect traffic as traffic leaves and enters a VPC. The company has deployed the solution by using AWS Cloud Formation on three Amazon EC2 instances in an Auto Scaling group. All network routing has been established to direct traffic to the EC2 instances.
Whenever the analysis software stops working, the Auto Scaling group replaces an instance. The network routes are not updated when the instance replacement occurs.
Which combination of steps will resolve this issue? {Select THREE.)

A. Create alarms based on EC2 status check metrics that will cause the Auto Scaling group to replace the failed instance.
B. Update the Cloud Formation template to install the Amazon CloudWatch agent on the EC2 instances.Configure the CloudWatch agent to send process metrics for the application.
C. Update the Cloud Formation template to install AWS Systems Manager Agent on the EC2 instances.Configure Systems Manager Agent to send process metrics for the application.
D. Create an alarm for the custom metric in Amazon CloudWatch for the failure scenario
E. Configure the alarm to publish a message to an Amazon Simple Notification Service {Amazon SNS) topic.
F. Create an AWS Lambda function that responds to the Amazon Simple Notification Service (Amazon SNS) message to take the instance out of servic
G. Update the network routes to point to the replacement instance.
H. In the Cloud Formation template, write a condition that updates the network routes when a replacement instance is launched.

Answer: BDE

NEW QUESTION 109
- (Exam Topic 2)
A solutions architect wants to cost-optimize and appropriately size Amazon EC2 instances in a single AWS account. The solutions architect wants to ensure that the instances are optimized based on CPU, memory, and network metrics.
Which combination of steps should the solutions architect take to meet these requirements? (Choose two.)

A. Purchase AWS Business Support or AWS Enterprise Support for the account.
B. Turn on AWS Trusted Advisor and review any “Low Utilization Amazon EC2 Instances” recommendations.
C. Install the Amazon CloudWatch agent and configure memory metric collection on the EC2 instances.
D. Configure AWS Compute Optimizer in the AWS account to receive findings and optimization recommendations.
E. Create an EC2 Instance Savings Plan for the AWS Regions, instance families, and operating systems of interest.

Answer: BD

Explanation:
AWS Trusted Advisor is a service that provides real-time guidance to help users provision their resources following AWS best practices1. One of the Trusted Advisor checks is “Low Utilization Amazon EC2 Instances”, which identifies EC2 instances that appear to be underutilized based on CPU, network I/O, and disk I/O metrics1. This check can help users optimize the cost and size of their EC2 instances by recommending smaller or more appropriate instance types.
AWS Compute Optimizer is a service that analyzes the configuration and utilization metrics of AWS resources and generates optimization recommendations to reduce the cost and improve the performance of workloads2. Compute Optimizer supports four types of AWS resources: EC2 instances, EBS volumes, ECS services on AWS Fargate, and Lambda functions2. For EC2 instances, Compute Optimizer evaluates the vCPUs, memory, storage, and other specifications, as well as the CPU utilization, network in and out, disk read and write, and other utilization metrics of currently running instances3. It then recommends optimal instance types based on price-performance trade-offs.
Option A is incorrect because purchasing AWS Business Support or AWS Enterprise Support for the account will not directly help with cost-optimization and sizing of EC2 instances. However, these support plans do provide access to more Trusted Advisor checks than the basic support plan1.
Option C is incorrect because installing the Amazon CloudWatch agent and configuring memory metric collection on the EC2 instances will not provide any optimization recommendations by itself. However, memory metrics can be used by Compute Optimizer to enhance its recommendations if enabled3.
Option E is incorrect because creating an EC2 Instance Savings Plan for the AWS Regions, instance families, and operating systems of interest will not help with cost-optimization and sizing of EC2 instances. Savings Plans are a flexible pricing model that offer lower prices on Amazon EC2 usage in exchange for a commitment to a consistent amount of usage for a 1- or 3-year term4. Savings Plans do not affect the configuration or utilization of EC2 instances.

NEW QUESTION 113
- (Exam Topic 2)
A company has built a high performance computing (HPC) cluster in AWS tor a tightly coupled workload that generates a large number of shared files stored in Amazon EFS. The cluster was performing well when the number of Amazon EC2 instances in the cluster was 100. However, when the company increased the cluster size to 1,000 EC2 instances, overall performance was well below expectations.
Which collection of design choices should a solutions architect make to achieve the maximum performance from the HPC cluster? (Select THREE.)

A. Ensure the HPC cluster Is launched within a single Availability Zone.
B. Launch the EC2 instances and attach elastic network interfaces in multiples of four.
C. Select EC2 Instance types with an Elastic Fabric Adapter (EFA) enabled.
D. Ensure the cluster Is launched across multiple Availability Zones.
E. Replace Amazon EFS with multiple Amazon EBS volumes in a RAID array.
F. Replace Amazon EFS with Amazon FSx for Lustre.

Answer: ACF

Explanation:
* A. High performance computing (HPC) workload cluster should be in a single AZ.
* C. Elastic Fabric Adapter (EFA) is a network device that you can attach to your Amazon EC2 instances to accelerate High Performance Computing (HPC)
* F. Amazon FSx for Lustre - Use it for workloads where speed matters, such as machine learning, high performance computing (HPC), video processing, and financial modeling.
Cluster – packs instances close together inside an Availability Zone. This strategy enables workloads to achieve the low-latency network performance necessary for tightly-coupled node-to-node communication that is typical of HPC applications.
https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/placement-groups.html

NEW QUESTION 114
- (Exam Topic 2)
A telecommunications company is running an application on AWS. The company has set up an AWS Direct Connect connection between the company's on-premises data center and AWS. The company deployed the application on Amazon EC2 instances in multiple Availability Zones behind an internal Application Load Balancer (ALB). The company's clients connect from the on-premises network by using HTTPS. The TLS terminates in the ALB. The company has multiple target groups and uses path-based routing to forward requests based on the URL path.
The company is planning to deploy an on-premises firewall appliance with an allow list that is based on IP address. A solutions architect must develop a solution to allow traffic flow to AWS from the on-premises network so that the clients can continue to access the application.
Which solution will meet these requirements?

A. Configure the existing ALB to use static IP addresse
B. Assign IP addresses in multiple Availability Zones to the AL
C. Add the ALB IP addresses to the firewall appliance.
D. Create a Network Load Balancer (NLB). Associate the NLB with one static IP addresses in multiple Availability Zone
E. Create an ALB-type target group for the NLB and add the existing ALAdd the NLB IP addresses to the firewall applianc
F. Update the clients to connect to the NLB.
G. Create a Network Load Balancer (NLB). Associate the LNB with one static IP addresses in multiple Availability Zone
H. Add the existing target groups to the NL
I. Update the clients to connect to the NL
J. Delete the ALB Add the NLB IP addresses to the firewall appliance.
K. Create a Gateway Load Balancer (GWLB). Assign static IP addresses to the GWLB in multiple Availability Zone
L. Create an ALB-type target group for the GWLB and add the existing AL
M. Add the GWLB IP addresses to the firewall applianc
N. Update the clients to connect to the GWLB.

Answer: B

Explanation:
The company should create a Network Load Balancer (NLB) and associate it with one static IP address in multiple Availability Zones. The company should also create an ALB-type target group for the NLB and add the existing ALB. The company should add the NLB IP addresses to the firewall appliance and update the clients to connect to the NLB. This solution will allow traffic flow to AWS from the on-premises network by using static IP addresses that can be added to the

firewall appliance's allow list. The NLB will forward requests to the ALB, which will use path-based routing to forward requests to the target groups.

NEW QUESTION 115
- (Exam Topic 2)
A company uses AWS Organizations for a multi-account setup in the AWS Cloud. The company's finance team has a data processing application that uses AWS Lambda and Amazon DynamoDB. The company's marketing team wants to access the data that is stored in the DynamoDB table.
The DynamoDB table contains confidential data. The marketing team can have access to only specific attributes of data in the DynamoDB table. The fi-nance team and the marketing team have separate AWS accounts.
What should a solutions architect do to provide the marketing team with the appropriate access to the DynamoDB table?

A. Create an SCP to grant the marketing team's AWS account access to the specific attributes of the DynamoDB tabl
B. Attach the SCP to the OU of the finance team.
C. Create an IAM role in the finance team's account by using IAM policy conditions for specific DynamoDB attributes (fine-grained access con-trol). Establish trust with the marketing team's accoun
D. In the mar-keting team's account, create an IAM role that has permissions to as-sume the IAM role in the finance team's account.
E. Create a resource-based IAM policy that includes conditions for spe-cific DynamoDB attributes (fine-grained access control). Attach the policy to the DynamoDB tabl
F. In the marketing team'saccount, create an IAM role that has permissions to access the DynamoDB table in the finance team's account.
G. Create an IAM role in the finance team's account to access the Dyna-moDB tabl
H. Use an IAM permissions boundary to limit the access to the specific attribute
I. In the marketing team's account, create an IAM role that has permissions to assume the IAM role in the finance team's account.

Answer: C

Explanation:
The company should create a resource-based IAM policy that includes conditions for specific DynamoDB attributes (fine-grained access control). The company should attach the policy to the DynamoDB table. In the marketing team's account, the company should create an IAM role that has permissions to access the DynamoDB table in the finance team's account. This solution will meet the requirements because a
resource-based IAM policy is a policy that you attach to an AWS resource (such as a DynamoDB table) to control who can access that resource and what actions they can perform on it. You can use IAM policy conditions to specify fine-grained access control for DynamoDB items and attributes. For example, you can allow or deny access to specific attributes of all items in a table by matching on attribute names1. By creating a resource-based policy that allows access to only specific attributes of the DynamoDB table and attaching it to the table, the company can restrict access to confidential data. By creating an IAM role in the marketing team's account that has permissions to access the DynamoDB table in the finance team's account, the company can enable cross-account access.
The other options are not correct because:

- Creating an SCP to grant the marketing team's AWS account access to the specific attributes of the DynamoDB table would not work because SCPs are policies that you can use with AWS Organizations to manage permissions in your organization's accounts. SCPs do not grant permissions; instead, they specify the maximum permissions that identities in an account can have2. SCPs cannot be used to specify fine-grained access control for DynamoDB items and attributes.
- Creating an IAM role in the finance team's account by using IAM policy conditions for specific DynamoDB attributes and establishing trust with the marketing team's account would not work because IAM roles are identities that you can create in your account that have specific permissions. You can use an IAM role to delegate access to users, applications, or services that don't normally have access to your AWS resources3. However, creating an IAM role in the finance team's account would not restrict access to specific attributes of the DynamoDB table; it would only allow cross-account access. The company would still need a resource-based policy attached to the table to enforce fine-grained access control.
- Creating an IAM role in the finance team's account to access the DynamoDB table and using an IAM permissions boundary to limit the access to the specific attributes would not work because IAM permissions boundaries are policies that you use to delegate permissions management to other
users. You can use permissions boundaries to limit the maximum permissions that an identity-based
policy can grant to an IAM entity (user or role)4. Permissions boundaries cannot be used to specify fine-grained access control for DynamoDB items and attributes.
References:

- https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/specifying-conditions.html
- https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps.html
- https://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles.html
- https://docs.aws.amazon.com/IAM/latest/UserGuide/access_policies_boundaries.html

NEW QUESTION 117
- (Exam Topic 2)
A company runs its sales reporting application in an AWS Region in the United States. The application uses an Amazon API Gateway Regional API and AWS Lambda functions to generate on-demand reports from data in an Amazon RDS for MySQL database. The frontend of the application is hosted on Amazon S3 and is accessed by users through an Amazon CloudFront distribution. The company is using Amazon Route 53 as the DNS service for the domain. Route 53 is configured with a simple routing policy to route traffic to the API Gateway API.
In the next 6 months, the company plans to expand operations to Europe. More than 90% of the database traffic is read-only traffic. The company has already deployed an API Gateway API and Lambda functions in the new Region.
A solutions architect must design a solution that minimizes latency for users who download reports. Which solution will meet these requirements?

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

Answer: C

Explanation:
The company should configure a cross-Region read replica for the RDS database in the new Region. The company should change the Route 53 record to latency-based routing to connect to the API Gateway API. This solution will meet the requirements because a cross-Region read replica is a feature that enables you to

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

- Using AWS Database Migration Service (AWS DMS) to replicate the primary database in the original Region to the database in the new Region would not be as cost-effective or simple as using a
cross-Region read replica. AWS DMS is a service that enables you to migrate relational databases, data
warehouses, NoSQL databases, and other types of data stores. You can use AWS DMS to perform one-time migrations or continuous data replication with high availability and consolidate databases into
a petabyte-scale data warehouse3. However, AWS DMS requires more configuration and management than creating a cross-Region read replica, which is fully managed by Amazon RDS. AWS DMS also incurs additional charges for replication instances and tasks.
- Creating an Amazon API Gateway Data API service integration with Amazon Redshift would not help with disaster recovery or minimizing latency. The Data API is a feature that enables you to query your Amazon Redshift cluster using HTTP requests, without needing a persistent connection or a SQL client. It is useful for building applications that interact with Amazon Redshift, but not for replicating or recovering data from an RDS database.
- Creating an AWS Data Exchange datashare by connecting AWS Data Exchange to the Redshift cluster would not help with disaster recovery or minimizing latency. AWS Data Exchange is a service that makes it easy for AWS customers to exchange data in the cloud. You can use AWS Data Exchange to subscribe to a diverse selection of third-party data products or offer your own data products to other AWS customers. A datashare is a feature that enables you to share live and secure access to your Amazon Redshift data across your accounts or with third parties without copying or moving the underlying data. It is useful for sharing query results and views with other users, but not for replicating or recovering data from an RDS database.

References:

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

NEW QUESTION 120
.....

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

AWS-Certified-Solutions-Architect-Professional Practice Exam Features:

* AWS-Certified-Solutions-Architect-Professional Questions and Answers Updated Frequently

* AWS-Certified-Solutions-Architect-Professional Practice Questions Verified by Expert Senior Certified Staff

* AWS-Certified-Solutions-Architect-Professional Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* AWS-Certified-Solutions-Architect-Professional Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

100% Actual & Verified — Instant Download, Please Click
Order The AWS-Certified-Solutions-Architect-Professional Practice Test Here