Amazon

Exam Questions AWS-Certified-Solutions-Architect-Professional

Amazon AWS Certified Solutions Architect Professional

About Exambible

Your Partner of IT Exam

Found in 1998

Exambible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, Exambible has its unique advantages that other companies could not achieve.

Our Advances

* 99.9% Uptime

All examinations will be up to date.

* 24/7 Quality Support

We will provide service round the clock.

* 100% Pass Rate

Our guarantee that you will pass the exam.

* Unique Gurantee

If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

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 uses a load balancer to distribute traffic to Amazon EC2 instances in a single Availability Zone. The company is concerned about security and wants a solutions architect to re-architect the solution to meet the following requirements:
• Inbound requests must be filtered for common vulnerability attacks.
• Rejected requests must be sent to a third-party auditing application.
• All resources should be highly available. Which solution meets these requirements?

A. Configure a Multi-AZ Auto Scaling group using the application's AM
B. Create an Application Load Balancer (ALB) and select the previously created Auto Scaling group as the targe
C. Use Amazon Inspector to monitor traffic to the ALB and EC2 instance
D. Create a web ACL in WA
E. Create an AWS WAF using the web ACL and AL
F. Use an AWS Lambda function to frequently push the Amazon Inspector report to the third-party auditing application.
G. Configure an Application Load Balancer (ALB) and add the EC2 instances as targets Create a web ACL in WA
H. Create an AWS WAF using the web ACL and ALB name and enable logging with Amazon CloudWatch Log
I. Use an AWS Lambda function to frequently push the logs to the third-party auditing application.
J. Configure an Application Load Balancer (ALB) along with a target group adding the EC2 instances as target
K. Create an Amazon Kinesis Data Firehose with the destination of the third-party auditing applicatio
L. Create a web ACL in WA
M. Create an AWS WAF using the web ACL and ALB then enable logging by selecting the Kinesis Data Firehose as the destinatio
N. Subscribe to AWS Managed Rules in AWS Marketplace, choosing the WAF as the subscriber.
O. Configure a Multi-AZ Auto Scaling group using the application's AM
P. Create an Application Load Balancer (ALB) and select the previously created Auto Scaling group as the targe
Q. Create an Amazon Kinesis Data Firehose with a destination of the third-party auditing applicatio
R. Create a web ACL inWA
S. Create an AWS WAF using the WebACL and ALB then enable logging by selecting the Kinesis Data Firehose as the destinatio
T. Subscribe to AWS Managed Rules in AWS Marketplace, choosing the WAF as the subscriber.

Answer: D

Explanation:
https://docs.aws.amazon.com/waf/latest/developerguide/marketplace-managed-rule-groups.html

NEW QUESTION 6
- (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 7
- (Exam Topic 2)
A company needs to optimize the cost of an AWS environment that contains multiple accounts in an organization in AWS Organizations. The company conducted cost optimization activities 3 years ago and purchased Amazon EC2 Standard Reserved Instances that recently expired.
The company needs EC2 instances for 3 more years. Additionally, the company has deployed a new serverless workload.
Which strategy will provide the company with the MOST cost savings?

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

Answer: A

Explanation:
The company should purchase the same Reserved Instances for an additional 3-year term with All Upfront payment. The company should purchase a 3-year Compute Savings Plan with All Upfront payment in the management account to cover any additional compute costs. This solution will provide the company with the most cost savings because Reserved Instances and Savings Plans are both pricing models that offer significant discounts compared to On-Demand pricing. Reserved Instances are commitments to use a specific instance type and size in a single Region for a one- or three-year term. You can choose between three payment options:
No Upfront, Partial Upfront, or All Upfront. The more you pay upfront, the greater the discoun1t. Savings Plans are flexible pricing models that offer low prices on EC2 instances, Fargate, and Lambda usage, in exchange for a commitment to a consistent amount of usage (measured in $/hour) for a one- or three-year term. You can choose between two types of Savings Plans: Compute Savings Plans and EC2 Instance Savings Plans. Compute Savings Plans apply to any EC2 instance regardless of Region, instance family, operating system, or tenancy, including those that are part of EMR, ECS, or EKS clusters, or launched by Fargate or Lambda. EC2 Instance Savings Plans apply to a specific instance family within a Region and provide the most savings2. By purchasing the same Reserved Instances for an additional 3-year term with All Upfront payment, the company can lock in the lowest possible price for its EC2 instances that run continuously for 3 years. By purchasing a 3-year Compute Savings Plan with All Upfront payment in the management account, the company can benefit from additional discounts on any other compute usage across its member accounts.
The other options are not correct because:
> Purchasing a 1-year Compute Savings Plan with No Upfront payment in each member account would not provide as much cost savings as purchasing a 3-year Compute Savings Plan with All Upfront payment in the management account. A 1-year term offers lower discounts than a 3-year term, and a No Upfront payment

option offers lower discounts than an All Upfront payment option. Also, purchasing a Savings Plan in each member account would not allow the company to share the benefits of unused Savings Plan discounts across its organization.

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

References:

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

NEW QUESTION 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 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 10
- (Exam Topic 2)
A solutions architect needs to define a reference architecture for a solution for three-tier applications with web. application, and NoSQL data layers. The reference architecture must meet the following requirements:
• High availability within an AWS Region
• Able to fail over in 1 minute to another AWS Region for disaster recovery
• Provide the most efficient solution while minimizing the impact on the user experience Which combination of steps will meet these requirements? (Select THREE.)

A. Use an Amazon Route 53 weighted routing policy set to 100/0 across the two selected Region
B. Set Time to Live (TTL) to 1 hour.
C. Use an Amazon Route 53 failover routing policy for failover from the primary Region to the disaster recovery Regio
D. Set Time to Live (TTL) to 30 seconds.
E. Use a global table within Amazon DynamoDB so data can be accessed in the two selected Regions.
F. Back up data from an Amazon DynamoDB table in the primary Region every 60 minutes and then write the data to Amazon S3. Use S3 Cross-Region replication to copy the data from the primary Region to the disaster recovery Regio
G. Have a script import the data into DynamoDB in a disaster recovery scenario.
H. Implement a hot standby model using Auto Scaling groups for the web and application layers across multiple Availability Zones in the Region
I. Use zonal Reserved Instances for the minimum number of servers and On-Demand Instances for any additional resources.
J. Use Auto Scaling groups for the web and application layers across multiple Availability Zones in the Region
K. Use Spot Instances for the required resources.

Answer: BCE

Explanation:
The requirements can be achieved by using an Amazon DynamoDB database with a global table. DynamoDB is a NoSQL database so it fits the requirements. A global table also allows both reads and writes to occur in both Regions. For the web and application tiers Auto Scaling groups should be configured. Due to the 1-minute RTO these must be configured in an active/passive state. The best pricing model to lower price but ensure resources are available when needed is to use a combination of zonal reserved instances and on-demand instances. To failover between the Regions, a Route 53 failover routing policy can be configured with a TTL configured on the record of 30 seconds. This will mean clients must resolve against Route 53 every 30 seconds to get the latest record. In a failover scenario the clients would be redirected to the secondary site if the primary site is unhealthy.

NEW QUESTION 15
- (Exam Topic 2)
A financial services company loaded millions of historical stock trades into an Amazon DynamoDB table. The table uses on-demand capacity mode. Once each day at midnight, a few million new records are loaded into the table. Application read activity against the table happens in bursts throughout the day. and a limited set of keys are repeatedly looked up. The company needs to reduce costs associated with DynamoDB.
Which strategy should a solutions architect recommend to meet this requirement?

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

Answer: D

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

NEW QUESTION 17
- (Exam Topic 2)
A company has developed a hybrid solution between its data center and AWS. The company uses Amazon VPC and Amazon EC2 instances that send application togs to Amazon CloudWatch. The EC2 instances read data from multiple relational databases that are hosted on premises.
The company wants to monitor which EC2 instances are connected to the databases in near-real time. The company already has a monitoring solution that uses Splunk on premises. A solutions architect needs to determine how to send networking traffic to Splunk.
How should the solutions architect meet these requirements?

A. Enable VPC flows logs, and send them to CloudWatc
B. Create an AWS Lambda function to periodically export the CloudWatch logs to an Amazon S3 bucket by using the pre-defined export functio
C. Generate ACCESS_KEY and SECRET_KEY AWS credential
D. Configure Splunk to pull the logs from the S3 bucket by using those credentials.
E. Create an Amazon Kinesis Data Firehose delivery stream with Splunk as the destinatio
F. Configure a pre-processing AWS Lambda function with a Kinesis Data Firehose stream processor that extracts individual log events from records sent by CloudWatch Logs subscription filter
G. Enable VPC flows logs, and send them to CloudWatc
H. Create a CloudWatch Logs subscription that sends log events to the Kinesis Data Firehose delivery stream.
I. Ask the company to log every request that is made to the databases along with the EC2 instance IP addres
J. Export the CloudWatch logs to an Amazon S3 bucke
K. Use Amazon Athena to query the logs grouped by database nam
L. Export Athena results to another S3 bucke
M. Invoke an AWS Lambdafunction to automatically send any new file that is put in the S3 bucket to Splunk.
N. Send the CloudWatch logs to an Amazon Kinesis data stream with Amazon Kinesis Data Analytics for SOL Application
O. Configure a 1 -minute sliding window to collect the event
P. Create a SQL query that uses the anomaly detection template to monitor any networking traffic anomalies in near-real tim
Q. Send the result to an Amazon Kinesis Data Firehose delivery stream with Splunk as the destination.

Answer: B

Explanation:
https://docs.aws.amazon.com/firehose/latest/dev/creating-the-stream-to-splunk.html

NEW QUESTION 21
- (Exam Topic 2)
A company is migrating a document processing workload to AWS. The company has updated many applications to natively use the Amazon S3 API to store, retrieve, and modify documents that a processing server generates at a rate of approximately 5 documents every second. After the document processing is finished, customers can download the documents directly from Amazon S3.
During the migration, the company discovered that it could not immediately update the processing server that generates many documents to support the S3 API. The server runs on Linux and requires fast local access to the files that the server generates and modifies. When the server finishes processing, the files must be available to the public for download within 30 minutes.
Which solution will meet these requirements with the LEAST amount of effort?

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

Answer: C

Explanation:
The company should configure Amazon FSx for Lustre with an import and export policy. The company should link the new file system to an S3 bucket. The company should install the Lustre client and mount the document store to an Amazon EC2 instance by using NFS. This solution will meet the requirements with the least amount of effort because Amazon FSx for Lustre is a fully managed service that provides a
high-performance file system optimized for fast processing of workloads such as machine learning, high
performance computing, video processing, financial modeling, and electronic design automation1. Amazon FSx for Lustre can be linked to an S3 bucket and can import data from and export data to the bucket2. The import and export policy can be configured to automatically import new or changed objects from S3 and export new or changed files to S33. This will ensure that the files are available to the public for download within 30 minutes. Amazon FSx for Lustre supports NFS version 3.0 protocol for Linux clients.
The other options are not correct because:
- Migrating the application to an AWS Lambda function would require a lot of effort and may not be feasible for the existing server that generates many documents. Lambda functions have limitations on execution time, memory, disk space, and network bandwidth.
- Setting up an Amazon S3 File Gateway would not work because S3 File Gateway does not support write-back caching, which means that files written to the file

share are uploaded to S3 immediately and are not available locally until they are downloaded again. This would not provide fast local access to the files that the server generates and modifies.

Configuring AWS DataSync to connect to an Amazon EC2 instance would not meet the requirement of making the files available to the public for download within 30 minutes. DataSync is a service that transfers data between on-premises storage systems and AWS storage services over the internet or AWS Direct Connect. DataSync tasks can be scheduled to run at specific times or intervals, but they are not triggered by file changes.
References:

- https://aws.amazon.com/fsx/lustre/
- https://docs.aws.amazon.com/fsx/latest/LustreGuide/create-fs-linked-data-repo.html
- https://docs.aws.amazon.com/fsx/latest/LustreGuide/import-export-data-repositories.html
- https://docs.aws.amazon.com/fsx/latest/LustreGuide/mounting-on-premises.html
- https://docs.aws.amazon.com/lambda/latest/dg/gettingstarted-limits.html
- https://docs.aws.amazon.com/storagegateway/latest/userguide/StorageGatewayConcepts.html
- https://docs.aws.amazon.com/datasync/latest/userguide/what-is-datasync.html

NEW QUESTION 25
- (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 26
- (Exam Topic 2)
A company uses AWS Organizations with a single OU named Production to manage multiple accounts All accounts are members of the Production OU Administrators use deny list SCPs in the root of the organization to manage access to restricted services.
The company recently acquired a new business unit and invited the new unit's existing AWS account to the organization Once onboarded the administrators of the new business unit discovered that they are not able to update existing AWS Config rules to meet the company's policies.
Which option will allow administrators to make changes and continue to enforce the current policies without introducing additional long-term maintenance?

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

Answer: D

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

NEW QUESTION 27
- (Exam Topic 2)
A company is running an application in the AWS Cloud. The application collects and stores a large amount of unstructured data in an Amazon S3 bucket. The S3 bucket contains several terabytes of data and uses the S3 Standard storage class. The data increases in size by several gigabytes every day.
The company needs to query and analyze the data. The company does not access data that is more than 1 year old. However, the company must retain all the data indefinitely for compliance reasons.
Which solution will meet these requirements MOST cost-effectively?

A. Use S3 Select to query the dat
B. Create an S3 Lifecycle policy to transition data that is more than 1 year old to S3 Glacier Deep Archive.
C. Use Amazon Redshift Spectrum to query the dat
D. Create an S3 Lifecycle policy to transition data that is more than 1 year old to S3 Glacier Deep Archive.
E. Use an AWS Glue Data Catalog and Amazon Athena to query the dat
F. Create an S3 Lifecycle policy to transition data that is more than 1 year old to S3 Glacier Deep Archive.
G. Use Amazon Redshift Spectrum to query the dat

H. Create an S3 Lifecycle policy to transition data that is more than 1 year old to S3 Intelligent-Tiering.

Answer: C

Explanation:
Generally, unstructured data should be converted structured data before querying them. AWS Glue can do that.
https://docs.aws.amazon.com/glue/latest/dg/schema-relationalize.html https://docs.aws.amazon.com/athena/latest/ug/glue-athena.html

NEW QUESTION 29
- (Exam Topic 2)
A company is running an application that uses an Amazon ElastiCache for Redis cluster as a caching layer A recent security audit revealed that the company has configured encryption at rest for ElastiCache However the company did not configure ElastiCache to use encryption in transit Additionally, users can access the cache without authentication
A solutions architect must make changes to require user authentication and to ensure that the company is using end-to-end encryption
Which solution will meet these requirements?

A. Create an AUTH token Store the token in AWS System Manager Parameter Store, as an encrypted parameter Create a new cluster with AUTH and configure encryption in transit Update the application toretrieve the AUTH token from Parameter Store when necessary and to use the AUTH token for authentication
B. Create an AUTH token Store the token in AWS Secrets Manager Configure the existing cluster to use the AUTH token and configure encryption in transit Update the application to retrieve the AUTH token from Secrets Manager when necessary and to use the AUTH token for authentication.
C. Create an SSL certificate Store the certificate in AWS Secrets Manager Create a new cluster and configure encryption in transit Update the application to retrieve the SSL certificate from Secrets Manager when necessary and to use the certificate for authentication.
D. Create an SSL certificate Store the certificate in AWS Systems Manager Parameter Store, as an encrypted advanced parameter Update the existing cluster to configure encryption in transit Update the application to retrieve the SSL certificate from Parameter Store when necessary and to use the certificate for authentication

Answer: B

Explanation:
Creating an AUTH token and storing it in AWS Secrets Manager and configuring the existing cluster to use the AUTH token and configure encryption in transit, and updating the application to retrieve the AUTH token from Secrets Manager when necessary and to use the AUTH token for authentication, would meet the requirements for user authentication and end-to-end encryption.
AWS Secrets Manager is a service that enables you to easily rotate, manage, and retrieve database credentials, API keys, and other secrets throughout their lifecycle. Secrets Manager also enables you to encrypt the data and ensure that only authorized users and applications can access it.
By configuring the existing cluster to use the AUTH token and encryption in transit, all data will be encrypted as it is sent over the network, providing additional security for the data stored in ElastiCache.
Additionally, by updating the application to retrieve the AUTH token from Secrets Manager when necessary and to use the AUTH token for authentication, it ensures that only authorized users and applications can access the cache.
Reference:
AWS Secrets Manager documentation: https://aws.amazon.com/secrets-manager/ Encryption in transit for ElastiCache:
https://docs.aws.amazon.com/AmazonElastiCache/latest/red-ug/encryption.html
Authentication and Authorization for ElastiCache: https://docs.aws.amazon.com/AmazonElastiCache/latest/red-ug/accessing-elasticache.html

NEW QUESTION 31
- (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 33
- (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 34
- (Exam Topic 2)
A company is implementing a serverless architecture by using AWS Lambda functions that need to access a Microsoft SQL Server DB instance on Amazon RDS. The company has separate environments for development and production, including a clone of the database system.
The company's developers are allowed to access the credentials for the development database. However, the credentials for the production database must be encrypted with a key that only members of the IT security team's IAM user group can access. This key must be rotated on a regular basis.
What should a solutions architect do in the production environment to meet these requirements?

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

Answer: D

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

NEW QUESTION 35
- (Exam Topic 2)
A solutions architect must create a business case for migration of a company's on-premises data center to the AWS Cloud. The solutions architect will use a configuration management database (CMDB) export of all the company's servers to create the case.
Which solution will meet these requirements MOST cost-effectively?

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

Answer: B

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

NEW QUESTION 37
- (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 41
- (Exam Topic 2)
A company is migrating a legacy application from an on-premises data center to AWS. The application uses MongoDB as a key-value database According to the company's technical guidelines, all Amazon EC2 instances must be hosted in a private subnet without an internet connection. In addition, all connectivity between applications and databases must be encrypted. The database must be able to scale based on demand.
Which solution will meet these requirements?

A. Create new Amazon DocumentDB (with MongoDB compatibility) tables for the application with Provisioned IOPS volume
B. Use the instance endpoint to connect to Amazon DocumentDB.
C. Create new Amazon DynamoDB tables for the application with on-demand capacit
D. Use a gateway VPC endpoint for DynamoDB to connect to the DynamoDB tables
E. Create new Amazon DynamoDB tables for the application with on-demand capacit
F. Use an interface VPC endpoint for DynamoDB to connect to the DynamoDB tables.
G. Create new Amazon DocumentDB (with MongoDB compatibility) tables for the application with Provisioned IOPS volumes Use the cluster endpoint to connect to Amazon DocumentDB

Answer: A

Explanation:
A is the correct answer because it uses Amazon DocumentDB (with MongoDB compatibility) as a key-value database that can scale based on demand and supports encryption in transit and at rest. Amazon DocumentDB is a fully managed document database service that is designed to be compatible with the MongoDB API. It is a NoSQL database that is optimized for storing, indexing, and querying JSON data. Amazon DocumentDB supports encryption in transit using TLS and encryption at rest using AWS Key Management Service (AWS KMS). Amazon DocumentDB also supports provisioned IOPS volumes that can scale up to 64 TiB of storage and 256,000 IOPS per cluster. To connect to Amazon DocumentDB, you can use the instance endpoint, which connects to a specific instance in the cluster, or the cluster endpoint, which connects to the primary instance or one of the replicas in the cluster. Using the cluster endpoint is recommended for high availability and load balancing purposes. References:

- https://docs.aws.amazon.com/documentdb/latest/developerguide/what-is.html
- https://docs.aws.amazon.com/documentdb/latest/developerguide/security.encryption.html
- https://docs.aws.amazon.com/documentdb/latest/developerguide/limits.html
- https://docs.aws.amazon.com/documentdb/latest/developerguide/connecting.html

NEW QUESTION 45
- (Exam Topic 2)
A solutions architect needs to review the design of an Amazon EMR cluster that is using the EMR File System (EMRFS). The cluster performs tasks that are critical to business needs. The cluster is running Amazon EC2 On-Demand Instances at all times tor all task, primary, and core nodes. The EMR tasks run each morning, starting at 1 ;00 AM. and take 6 hours to finish running. The amount of time to complete the processing is not a priority because the data is not referenced until late in the day.

The solutions architect must review the architecture and suggest a solution to minimize the compute costs. Which solution should the solutions architect recommend to meet these requirements?

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

Answer: A

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

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

NEW QUESTION 47
- (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 49
- (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 50
- (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 55
- (Exam Topic 2)
A company uses an AWS CodeCommit repository The company must store a backup copy of the data that is in the repository in a second AWS Region
Which solution will meet these requirements?

A. Configure AWS Elastic Disaster Recovery to replicate the CodeCommit repository data to the second Region
B. Use AWS Backup to back up the CodeCommit repository on an hourly schedule Create a cross-Region copy in the second Region
C. Create an Amazon EventBridge rule to invoke AWS CodeBuild when the company pushes code to the repository Use CodeBuild to clone the repository Create a zip file of the content Copy the file to an S3 bucket in the second Region
D. Create an AWS Step Functions workflow on an hourly schedule to take a snapshot of the CodeCommit repository Configure the workflow to copy the snapshot to an S3 bucket in the second Region

Answer: B

Explanation:
AWS Backup is a fully managed service that makes it easy to centralize and automate the creation, retention, and restoration of backups across AWS services. It provides a way to schedule automatic backups for CodeCommit repositories on an hourly basis. Additionally, it also supports cross-Region replication, which allows you to copy the backups to a second Region for disaster recovery.
By using AWS Backup, the company can set up an automatic and regular backup schedule for the CodeCommit repository, ensuring that the data is regularly backed up and stored in a second Region. This can provide a way to recover quickly from any disaster event that might occur.
Reference:
AWS Backup documentation: https://aws.amazon.com/backup/ AWS Backup for AWS CodeCommit documentation: https://aws.amazon.com/about-aws/whats-new/2020/07/aws-backup-now-supports-aws-codecommit-repositorie

NEW QUESTION 57
- (Exam Topic 2)
A company has an on-premises Microsoft SOL Server database that writes a nightly 200 GB export to a local drive. The company wants to move the backups to more robust cloud storage on Amazon S3. The company has set up a 10 Gbps AWS Direct Connect connection between the on-premises data center and AWS.
Which solution meets these requirements MOST cost-effectively?

A. Create a new S3 bucke
B. Deploy an AWS Storage Gateway file gateway within the VPC that Is connected to the Direct Connect connectio
C. Create a new SMB file shar
D. Write nightly database exports to the new SMB file share.
E. Create an Amazon FSx for Windows File Server Single-AZ file system within the VPC that is connected to the Direct Connect connectio
F. Create a new SMB file shar

G. Write nightly database exports to an SMB file share on the Amazon FSx file syste
H. Enable nightly backups.
I. Create an Amazon FSx for Windows File Server Multi-AZ file system within the VPC that is connected to the Direct Connect connectio
J. Create a new SMB file shar
K. Write nightly database exports to an SMB file share on the Amazon FSx file syste
L. Enable nightly backups.
M. Create a new S3 bucke
N. Deploy an AWS Storage Gateway volume gateway within the VPC that Is connected to the Direct Connect connectio
O. Create a new SMB file shar
P. Write nightly database exports to the new SMB file share on the volume gateway, and automate copies of this data to an S3 bucket.

Answer: A

Explanation:
https://docs.aws.amazon.com/filegateway/latest/files3/CreatingAnSMBFileShare.html

NEW QUESTION 60
- (Exam Topic 2)
A company runs an application in an on-premises data center. The application gives users the ability to upload media files. The files persist in a file server. The web application has many users. The application server is overutilized, which causes data uploads to fail occasionally. The company frequently adds new storage to the file server. The company wants to resolve these challenges by migrating the application to AWS.
Users from across the United States and Canada access the application. Only authenticated users should have the ability to access the application to upload files. The company will consider a solution that refactors the application, and the company needs to accelerate application development.
Which solution will meet these requirements with the LEAST operational overhead?
A. Use AWS Application Migration Service to migrate the application server to Amazon EC2 instances.Create an Auto Scaling group for the EC2 instance
B. Use an Application Load Balancer to distribute the request
C. Modify the application to use Amazon S3 to persist the file
D. Use Amazon Cognito to authenticate users.
E. Use AWS Application Migration Service to migrate the application server to Amazon EC2 instances.Create an Auto Scaling group for the EC2 instance
F. Use an Application Load Balancer to distribute the request
G. Set up AWS IAM Identity Center (AWS Single Sign-On) to give users the ability to sign in to the applicatio
H. Modify the application to use Amazon S3 to persist the files.
I. Create a static website for uploads of media file
J. Store the static assets in Amazon S3. Use AWS AppSync to create an AP
K. Use AWS Lambda resolvers to upload the media files to Amazon S3. Use Amazon Cognito to authenticate users.
L. Use AWS Amplify to create a static website for uploads of media file
M. Use Amplify Hosting to serve the website through Amazon CloudFron
N. Use Amazon S3 to store the uploaded media file
O. Use Amazon Cognito to authenticate users.

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 64
- (Exam Topic 2)
A company operates a proxy server on a fleet of Amazon EC2 instances. Partners in different countries use the proxy server to test the company's functionality. The EC2 instances are running in a VPC. and the instances have access to the internet.
The company's security policy requires that partners can access resources only from domains that the company owns.
Which solution will meet these requirements?

A. Create an Amazon Route 53 Resolver DNS Firewall domain list that contains the allowed domains.Configure a DNS Firewall rule group with a rule that has a high numeric value that blocks all request
B. Configure a rule that has a low numeric value that allows requests for domains in the allowed lis
C. Associate the rule group with the VPC.
D. Create an Amazon Route 53 Resolver DNS Firewall domain list that contains the allowed domains.Configure a Route 53 outbound endpoin
E. Associate the outbound endpoint with the VP
F. Associate the domain list with the outbound endpoint.
G. Create an Amazon Route 53 traffic flow policy to match the allowed domain
H. Configure the traffic flow policy to forward requests that match to the Route 53 Resolve
I. Associate the traffic flow policy with the VPC.
J. Create an Amazon Route 53 outbound endpoin
K. Associate the outbound endpoint with the VP
L. Configure a Route 53 traffic flow policy to forward requests for allowed domains to the outbound endpoin
M. Associate the traffic flow policy with the VPC.

Answer: A

Explanation:
The company should create an Amazon Route 53 Resolver DNS Firewall domain list that contains the allowed domains. The company should configure a DNS Firewall rule group with a rule that has a high numeric value that blocks all requests. The company should configure a rule that has a low numeric value that allows requests for domains in the allowed list. The company should associate the rule group with the VPC. This solution will meet the requirements because Amazon Route 53 Resolver DNS Firewall is a feature that enables you to filter and regulate outbound DNS traffic for your VPC. You can create reusable collections of filtering rules in DNS Firewall rule groups and associate them with your VPCs. You can specify lists of domain names to allow or block, and you can customize the responses for the DNS queries that you block1. By creating a domain list with the allowed domains and a rule group with rules to allow or block requests based on the domain list, the company can enforce its security policy and control access to sites.
The other options are not correct because:
- Configuring a Route 53 outbound endpoint and associating it with the VPC would not help with filtering outbound DNS traffic. A Route 53 outbound endpoint is a resource that enables you to forward DNS
queries from your VPC to your network over AWS Direct Connect or VPN connections2. It does not provide any filtering capabilities.
- Creating a Route 53 traffic flow policy to match the allowed domains would not help with filtering outbound DNS traffic. A Route 53 traffic flow policy is a resource that enables you to route traffic based on multiple criteria, such as endpoint health, geographic location, and latency3. It does not provide any filtering capabilities.
- Creating a Gateway Load Balancer (GWLB) would not help with filtering outbound DNS traffic. A GWLB is a service that enables you to deploy, scale, and manage third-party virtual appliances such as firewalls, intrusion detection and prevention systems, and deep packet inspection systems in the cloud4. It does not provide any filtering capabilities.
References:
- https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/resolver-dns-firewall.html
- https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/resolver-outbound-endpoints.html
- https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/traffic-flow.html
- https://docs.aws.amazon.com/elasticloadbalancing/latest/gateway/introduction.html

NEW QUESTION 65
- (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 69
- (Exam Topic 2)
A company is creating a centralized logging service running on Amazon EC2 that will receive and analyze logs from hundreds of AWS accounts. AWS PrivateLink is being used to provide connectivity between the client services and the logging service.
In each AWS account with a client, an interface endpoint has been created for the logging service and is available. The logging service running on EC2 instances with a Network Load Balancer (NLB) are deployed in different subnets. The clients are unable to submit logs using the VPC endpoint.
Which combination of steps should a solutions architect take to resolve this issue? (Select TWO.)

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

Answer: AC

NEW QUESTION 73
- (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 76
- (Exam Topic 2)
A company operates an on-premises software-as-a-service (SaaS) solution that ingests several files daily. The company provides multiple public SFTP endpoints to its customers to facilitate the file transfers. The customers add the SFTP endpoint IP addresses to their firewall allow list for outbound traffic. Changes to the SFTP endmost IP addresses are not permitted.
The company wants to migrate the SaaS solution to AWS and decrease the operational overhead of the file transfer service.
Which solution meets these requirements?

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

Answer: A

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

NEW QUESTION 80
- (Exam Topic 2)
A company has an application that runs as a ReplicaSet of multiple pods in an Amazon Elastic Kubernetes Service (Amazon EKS) cluster. The EKS cluster has nodes in multiple Availability Zones. The application generates many small files that must be accessible across all running instances of the application. The company needs to back up the files and retain the backups for 1 year.
Which solution will meet these requirements while providing the FASTEST storage performance?

A. Create an Amazon Elastic File System (Amazon EFS) file system and a mount target for each subnet that contains nodes in the EKS cluste
B. Configure the ReplicaSet to mount the file syste
C. Direct the application to store files in the file syste
D. Configure AWS Backup to back up and retain copies of the data for 1 year.
E. Create an Amazon Elastic Block Store (Amazon EBS) volum
F. Enable the EBS Multi-Attach feature.Configure the ReplicaSet to mount the EBS volum
G. Direct the application to store files in the EBS volum
H. Configure AWS Backup to back up and retain copies of the data for 1 year.
I. Create an Amazon S3 bucke
J. Configure the ReplicaSet to mount the S3 bucke
K. Direct the application to store files in the S3 bucke
L. Configure S3 Versioning to retain copies of the dat
M. Configure an S3 Lifecycle policy to delete objects after 1 year.
N. Configure the ReplicaSet to use the storage available on each of the running application pods to store the files locall
O. Use a third-party tool to back up the EKS cluster for 1 year.

Answer: A

Explanation:
In the past, EBS can be attached only to one ec2 instance but not anymore but there are limitations like - it works only on io1/io2 instance types and many others as described here. https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-volumes-multi.html EFS has shareable storage
In terms of performance, Amazon EFS is optimized for workloads that require high levels of aggregate throughput and IOPS, whereas EBS is optimized for low-latency, random access I/O operations. Amazon EFS is designed to scale throughput and capacity automatically as your storage needs grow, while EBS volumes can be resized on demand.

NEW QUESTION 81
- (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 85
- (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 86
- (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 91
- (Exam Topic 2)
A company wants to containerize a multi-tier web application and move the application from an on-premises data center to AWS. The application includes web. application, and database tiers. The company needs to make the application fault tolerant and scalable. Some frequently accessed data must always be available across application servers. Frontend web servers need session persistence and must scale to meet increases in traffic.
Which solution will meet these requirements with the LEAST ongoing operational overhead?

A. Run the application on Amazon Elastic Container Service (Amazon ECS) on AWS Fargat
B. Use Amazon Elastic File System (Amazon EFS) for data that is frequently accessed between the web and application tier
C. Store the frontend web server session data in Amazon Simple Queue Service (Amazon SOS).
D. Run the application on Amazon Elastic Container Service (Amazon ECS) on Amazon EC2. Use Amazon ElastiCache for Redis to cache frontend web server session dat
E. Use Amazon Elastic Block Store (Amazon EBS) with Multi-Attach on EC2 instances that are distributed across multiple Availability Zones.
F. Run the application on Amazon Elastic Kubernetes Service (Amazon EKS). Configure Amazon EKS to use managed node group

G. Use ReplicaSets to run the web servers and application
H. Create an Amazon Elastic File System (Amazon EFS) Me syste
I. Mount the EFS file system across all EKS pods to store frontend web server session data.
J. Deploy the application on Amazon Elastic Kubernetes Service (Amazon EKS) Configure Amazon EKS to use managed node group
K. Run the web servers and application as Kubernetes deployments in the EKS cluste
L. Store the frontend web server session data in an Amazon DynamoDB tabl
M. Create an Amazon Elastic File System (Amazon EFS) volume that all applications will mount at the time of deployment.

Answer: D

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

- https://docs.aws.amazon.com/eks/latest/userguide/managed-node-groups.html
- https://docs.aws.amazon.com/eks/latest/userguide/deployments.html
- https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Introduction.html
- https://docs.aws.amazon.com/efs/latest/ug/mounting-fs.html

NEW QUESTION 94
- (Exam Topic 2)
A company runs a processing engine in the AWS Cloud The engine processes environmental data from logistics centers to calculate a sustainability index The company has millions of devices in logistics centers that are spread across Europe The devices send information to the processing engine through a RESTful API The API experiences unpredictable bursts of traffic The company must implement a solution to process all data that the devices send to the processing engine Data loss is unacceptable
Which solution will meet these requirements?

A. Create an Application Load Balancer (ALB) for the RESTful API Create an Amazon Simple Queue Service (Amazon SQS) queue Create a listener and a target group for the ALB Add the SQS queue as the target Use a container that runs in Amazon Elastic Container Service (Amazon ECS) with the Fargate launch type to process messages in the queue
B. Create an Amazon API Gateway HTTP API that implements the RESTful API Create an Amazon Simple Queue Service (Amazon SQS) queue Create an API Gateway service integration with the SQS queue Create an AWS Lambda function to process messages in the SQS queue
C. Create an Amazon API Gateway REST API that implements the RESTful API Create a fleet of Amazon EC2 instances in an Auto Scaling group Create an API Gateway Auto Scaling group proxy integration Use the EC2 instances to process incoming data
D. Create an Amazon CloudFront distribution for the RESTful API Create a data stream in Amazon Kinesis Data Streams Set the data stream as the origin for the distribution Create an AWS Lambda function to consume and process data in the data stream

Answer: A

Explanation:
it will use the ALB to handle the unpredictable bursts of traffic and route it to the SQS queue. The SQS queue will act as a buffer to store incoming data temporarily and the container running in Amazon ECS with the Fargate launch type will process messages in the queue. This approach will ensure that all data is processed and prevent data loss.

NEW QUESTION 95
.....

Relate Links

100% Pass Your AWS-Certified-Solutions-Architect-Professional Exam with Exambible Prep Materials

https://www.exambible.com/AWS-Certified-Solutions-Architect-Professional-exam/

Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - https://www.exambible.com/