# Amazon

## Exam Questions AWS-Certified-Solutions-Architect-Professional

Amazon AWS Certified Solutions Architect Professional

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

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

**Answer:** D

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

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

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

**Answer:** B

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

**NEW QUESTION 3**
- (Exam Topic 2)
A company needs to optimize the cost of backups for Amazon Elastic File System (Amazon EFS). A solutions architect has already configured a backup plan in AWS Backup for the EFS backups. The backup plan contains a rule with a lifecycle configuration to transition EFS backups to cold storage after 7 days and to keep the backups for an additional 90 days.
After I month, the company reviews its EFS storage costs and notices an increase in the EFS backup costs. The EFS backup cold storage produces almost double the cost of the EFS warm backup storage.
What should the solutions architect do to optimize the cost?

A. Modify the backup rule's lifecycle configuration to move the EFS backups to cold storage after 1 day.Set the backup retention period to 30 days.
B. Modify the backup rule's lifecycle configuration to move the EFS backups to cold storage after 8 days.Set the backup retention period to 30 days.
C. Modify the backup rule's lifecycle configuration to move the EFS backups to cold storage after 1 day.Set the backup retention period to 90 days.
D. Modify the backup rule's lifecycle configuration to move the EFS backups to cold storage after 8 days.Set the backup retention period to 98 days.

**Answer:** A

**Explanation:**
The cost of EFS backup cold storage is $0.01 per GB-month, whereas the cost of EFS backup warm storage is $0.05 per GB-month1. Therefore, moving the backups to cold storage as soon as possible will reduce the storage cost. However, cold storage backups must be retained for a minimum of 90 day2s, otherwise they incur a pro-rated charge equal to the storage charge for the remaining days1. Therefore, setting the backup retention period to 30 days will incur a penalty of 60 days of cold storage cost for each backup deleted. This penalty will still be lower than keeping the backups in warm storage for 7 days and then in cold storage for 83 days, which is the current configuration. Therefore, option A is the most cost-effective solution.

**NEW QUESTION 4**
- (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 5**
- (Exam Topic 2)
A company has loT sensors that monitor traffic patterns throughout a large city. The company wants to read and collect data from the sensors and perform aggregations on the data.
A solutions architect designs a solution in which the loT devices are streaming to Amazon Kinesis Data Streams. Several applications are reading from the stream. However, several consumers are experiencing throttling and are periodically and are periodically encountering a RealProvisioned Throughput Exceeded error.
Which actions should the solution architect take to resolve this issue? (Select THREE.)

A. Reshard the stream to increase the number of shards s in the stream.
B. Use the Kinesis Producer Library KPL). Adjust the polling frequency.
C. Use consumers with the enhanced fan-out feature.
D. Reshard the stream to reduce the number of shards in the stream.
E. Use an error retry and exponential backoff mechanism in the consumer logic.
F. Configure the stream to use dynamic partitioning.

**Answer:** ACE

**Explanation:**
https://repost.aws/knowledge-center/kinesis-readprovisionedthroughputexceeded Follow Data Streams best practices
To mitigate ReadProvisionedThroughputExceeded exceptions, apply these best practices:
• Reshard your stream to increase the number of shards in the stream.
• Use consumers with enhanced fan-out. For more information about enhanced fan-out, see Developing custom consumers with dedicated throughput (enhanced fan-out).
• Use an error retry and exponential backoff mechanism in the consumer logic if ReadProvisionedThroughputExceeded exceptions are encountered. For consumer applications that use an AWS SDK, the requests are retried by default.

**NEW QUESTION 6**
- (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 7**
- (Exam Topic 2)
A company wants to use AWS for disaster recovery for an on-premises application. The company has hundreds of Windows-based servers that run the application. All the servers mount a common share.
The company has an RTO of 15 minutes and an RPO of 5 minutes. The solution must support native failover and fallback capabilities.
Which solution will meet these requirements MOST cost-effectively?

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

**Answer:** D

**NEW QUESTION 8**
- (Exam Topic 2)
A company has an application that runs on Amazon EC2 instances in an Amazon EC2 Auto Scaling group. The company uses AWS CodePipeline to deploy the application. The instances that run in the Auto Scaling group are constantly changing because of scaling events.
When the company deploys new application code versions, the company installs the AWS CodeDeploy agent on any new target EC2 instances and associates the instances with the CodeDeploy deployment group. The application is set to go live within the next 24 hours.
What should a solutions architect recommend to automate the application deployment process with the LEAST amount of operational overhead?

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

**Answer:** D

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

**NEW QUESTION 9**
- (Exam Topic 2)
A company is storing sensitive data in an Amazon S3 bucket. The company must log all activities for objects in the S3 bucket and must keep the logs for 5 years.
The company's security team also must receive an email notification every time there is an attempt to delete data in the S3 bucket.
Which combination of steps will meet these requirements MOST cost-effectively? (Select THREE.)

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

F. Configure a new S3 bucket to store the logs with an S3 Lifecycle policy.

**Answer:** ADF

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

**NEW QUESTION 10**
- (Exam Topic 2)
A company is building a call center by using Amazon Connect. The company's operations team is defining a disaster recovery (DR) strategy across AWS Regions. The contact center has dozens of contact flows, hundreds of users, and dozens of claimed phone numbers.
Which solution will provide DR with the LOWEST RTO?

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

**Answer:** D

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

**NEW QUESTION 10**
- (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 13**
- (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 14**
- (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 19**
- (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 24**
- (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 26**
- (Exam Topic 2)
A solutions architect is planning to migrate critical Microsoft SOL Server databases to AWS. Because the databases are legacy systems, the solutions architect will move the databases to a modern data architecture. The solutions architect must migrate the databases with near-zero downtime.
Which solution will meet these requirements?

A. Use AWS Application Migration Service and the AWS Schema Conversion Tool (AWS SCT). Perform an In-place upgrade before the migratio
B. Export the migrated data to Amazon Aurora Serverless after cutove
C. Repoint the applications to Amazon Aurora.
D. Use AWS Database Migration Service (AWS DMS) to Rehost the databas
E. Set Amazon S3 as a target.Set up change data capture (CDC) replicatio
F. When the source and destination are fully synchronized, load the data from Amazon S3 into an Amazon RDS for Microsoft SQL Server DB Instance.
G. Use native database high availability tools Connect the source system to an Amazon RDS for Microsoft SQL Server DB instance Configure replication accordingl
H. When data replication is finished, transition the workload to an Amazon RDS for Microsoft SQL Server DB instance.
I. Use AWS Application Migration Servic
J. Rehost the database server on Amazon EC2. When data replication is finished, detach the database and move the database to an Amazon RDS for Microsoft SQL Server DB instanc
K. Reattach the database and then cut over all networking.

**Answer:** B

**Explanation:**
AWS DMS can migrate data from a source database to a target database in AWS, using change data capture (CDC) to replicate ongoing changes and keep the databases in sync. Setting Amazon S3 as a target allows storing the migrated data in a durable and cost-effective storage service. When the source and destination are fully synchronized, the data can be loaded from Amazon S3 into an Amazon RDS for Microsoft SQL Server DB instance, which is a managed database service that simplifies database administration tasks. References:

- https://docs.aws.amazon.com/dms/latest/userguide/CHAP_Source.SQLServer.html
- https://docs.aws.amazon.com/dms/latest/userguide/CHAP_Target.S3.html
- https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_SQLServer.html

**NEW QUESTION 31**
- (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 33**
- (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 35**
- (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 39**
- (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 42**
- (Exam Topic 2)
A company is running a critical stateful web application on two Linux Amazon EC2 instances behind an Application Load Balancer (ALB) with an Amazon RDS for MySQL database The company hosts the DNS records for the application in Amazon Route 53 A solutions architect must recommend a solution to improve the resiliency of the application
The solution must meet the following objectives:
• Application tier RPO of 2 minutes. RTO of 30 minutes
• Database tier RPO of 5 minutes RTO of 30 minutes
The company does not want to make significant changes to the existing application architecture The company must ensure optimal latency after a failover
Which solution will meet these requirements?

A. Configure the EC2 instances to use AWS Elastic Disaster Recovery Create a cross-Region read replica for the RDS DB instance Create an ALB in a second AWS Region Create an AWS Global Accelerator endpoint and associate the endpoint with the ALBs Update DNS records to point to the Global Accelerator endpoint
B. Configure the EC2 instances to use Amazon Data Lifecycle Manager (Amazon DLM) to take snapshots of the EBS volumes Configure RDS automated backups Configure backup replication to a second AWS Region Create an ALB in the second Region Create an AWS Global Accelerator endpoint, and associate the endpoint with the ALBs Update DNS records to point to the Global Accelerator endpoint
C. Create a backup plan in AWS Backup for the EC2 instances and RDS DB instance Configure backup replication to a second AWS Region Create an ALB in the second Region Configure an Amazon CloudFront distribution in front of the ALB Update DNS records to point to CloudFront
D. Configure the EC2 instances to use Amazon Data Lifecycle Manager (Amazon DLM) to take snapshots of the EBS volumes Create a cross-Region read replica for the RDS DB instance Create an ALB in a second AWS Region Create an AWS Global Accelerator endpoint and associate the endpoint with the ALBs

**Answer:** B

**Explanation:**
This option meets the RPO and RTO requirements for both the application and database tiers and uses tools like Amazon DLM and RDS automated backups to create and manage the backups. Additionally, it uses Global Accelerator to ensure low latency after failover by directing traffic to the closest healthy endpoint.

**NEW QUESTION 47**
- (Exam Topic 2)
A company hosts a blog post application on AWS using Amazon API Gateway, Amazon DynamoDB, and AWS Lambda. The application currently does not use API keys to authorize requests. The API model is as follows: GET/posts/[postid] to get post details GET/users[userid] to get user details GET/comments/[commentid] to get comments details
The company has noticed users are actively discussing topics in the comments section, and the company wants to increase user engagement by marking the comments appears in real time.
Which design should be used to reduce comment latency and improve user experience?

A. Use edge-optimized API with Amazon CloudFront to cache API responses.
B. Modify the blog application code to request GET comment[commented] every 10 seconds.
C. Use AWS AppSync and leverage WebSockets to deliver comments.
D. Change the concurrency limit of the Lambda functions to lower the API response time.

**Answer:** C

**Explanation:**
https://docs.aws.amazon.com/appsync/latest/devguide/graphql-overview.html
AWS AppSync is a fully managed GraphQL service that allows applications to securely access, manipulate, and receive data as well as real-time updates from multiple data sources1. AWS AppSync supports GraphQL subscriptions to perform real-time operations and can push data to clients that choose to listen to specific events from the backend1. AWS AppSync uses WebSockets to establish and maintain a secure connection between the clients and the API endpoint2. Therefore, using AWS AppSync and leveraging WebSockets is a suitable design to reduce comment latency and improve user experience.

**NEW QUESTION 48**
- (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 53**
- (Exam Topic 2)
A company has several AWS accounts. A development team is building an automation framework for cloud governance and remediation processes. The automation framework uses AWS Lambda functions in a centralized account. A solutions architect must implement a least privilege permissions policy that allows the Lambda functions to run in each of the company's AWS accounts.
Which combination of steps will meet these requirements? (Choose two.)

A. In the centralized account, create an IAM role that has the Lambda service as a trusted entit
B. Add an inline policy to assume the roles of the other AWS accounts.
C. In the other AWS accounts, create an IAM role that has minimal permission
D. Add the centralized account's Lambda IAM role as a trusted entity.
E. In the centralized account, create an IAM role that has roles of the other accounts as trusted entities.Provide minimal permissions.
F. In the other AWS accounts, create an IAM role that has permissions to assume the role of the centralized accoun
G. Add the Lambda service as a trusted entity.
H. In the other AWS accounts, create an IAM role that has minimal permission
I. Add the Lambda service as a trusted entity.

**Answer:** AB

**Explanation:**
https://medium.com/@it.melnichenko/invoke-a-lambda-across-multiple-aws-accounts-8c094b2e70be

**NEW QUESTION 56**
- (Exam Topic 2)
A company is running a web application in a VPC. The web application runs on a group of Amazon EC2 instances behind an Application Load Balancer (ALB). The ALB is using AWS WAF.
An external customer needs to connect to the web application. The company must provide IP addresses to all external customers.
Which solution will meet these requirements with the LEAST operational overhead?

A. Replace the ALB with a Network Load Balancer (NLB). Assign an Elastic IP address to the NLB.
B. Allocate an Elastic IP addres
C. Assign the Elastic IP address to the ALProvide the Elastic IP address to the customer.
D. Create an AWS Global Accelerator standard accelerato
E. Specify the ALB as the accelerator's endpoint.Provide the accelerator's IP addresses to the customer.
F. Configure an Amazon CloudFront distributio
G. Set the ALB as the origi
H. Ping the distribution's DNS name to determine the distribution's public IP addres
I. Provide the IP address to the customer.

**Answer:** C

**Explanation:**
https://docs.aws.amazon.com/global-accelerator/latest/dg/about-accelerators.alb-accelerator.html Option A is wrong. AWS WAF does not support associating with NLB.
https://docs.aws.amazon.com/waf/latest/developerguide/waf-chapter.html Option B is wrong. An ALB does not support an Elastic IP address.
https://aws.amazon.com/elasticloadbalancing/features/

**NEW QUESTION 60**
- (Exam Topic 2)
A company runs an application on AWS. The company curates data from several different sources. The company uses proprietary algorithms to perform data transformations and aggregations. After the company performs E TL processes, the company stores the results in Amazon Redshift tables. The company sells this data to other companies. The company downloads the data as files from the Amazon Redshift tables and transmits the files to several data customers by using FTP. The number of data customers has grown significantly. Management of the data customers has become difficult.
The company will use AWS Data Exchange to create a data product that the company can use to share data with customers. The company wants to confirm the identities of the customers before the company shares data.
The customers also need access to the most recent data when the company publishes the data. Which solution will meet these requirements with the LEAST operational overhead?

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

**Answer:** C

**Explanation:**
The company should download the data from the Amazon Redshift tables to an Amazon S3 bucket periodically and use AWS Data Exchange for S3 to share data with customers. The company should configure subscription verification and require the data customers to subscribe to the data product. This solution will meet the requirements with the least operational overhead because AWS Data Exchange for S3 is a feature that enables data subscribers to access third-party data files directly from data providers' Amazon S3 buckets. Subscribers can easily use these files for their data analysis with AWS services without needing to create or manage data copies. Data providers can easily set up AWS Data Exchange for S3 on top of their existing S3 buckets to share direct access to an entire S3 bucket or specific prefixes and S3 objects. AWS Data Exchange automatically manages subscriptions, entitlements, billing, and payment1.
The other options are not correct because:
- Using AWS Data Exchange for APIs to share data with customers would not work because AWS Data Exchange for APIs is a feature that enables data subscribers to access third-party APIs directly from data providers' AWS accounts. Subscribers can easily use these APIs for their data analysis with AWS services without needing to manage API keys or tokens. Data providers can easily set up AWS Data Exchange for APIs on top of their existing API Gateway resources to share direct access to an entire API or specific routes and stages2. However, this feature is not suitable for sharing data from Amazon Redshift tables, which are not exposed as APIs.
- Creating an Amazon API Gateway Data API service integration with Amazon Redshift would not work because the Data API is a feature that enables you to query your Amazon Redshift cluster using HTTP requests, without needing a persistent connection or a SQL client3. It is useful for building applications that interact with Amazon Redshift, but not for sharing data files with customers.
- Creating an AWS Data Exchange datashare by connecting AWS Data Exchange to the Redshift cluster would not work because AWS Data Exchange does not support datashares for Amazon Redshift clusters. A datashare is a feature that enables you to share live and secure access to your Amazon Redshift data across your accounts or with third parties without copying or moving the underlying data4. It is useful for sharing query results and views with other users, but not for sharing data files with customers.
- Publishing the Amazon Redshift data to an Open Data on AWS Data Exchange would not work because Open Data on AWS Data Exchange is a feature that enables you to find and use free and public datasets from AWS customers and partners. It is useful for accessing open and free data, but not for confirming the identities of the customers or charging them for the data.
References:
- https://aws.amazon.com/data-exchange/why-aws-data-exchange/s3/
- https://aws.amazon.com/data-exchange/why-aws-data-exchange/api/
- https://docs.aws.amazon.com/redshift/latest/mgmt/data-api.html
- https://docs.aws.amazon.com/redshift/latest/dg/datashare-overview.html
- https://aws.amazon.com/data-exchange/open-data/

**NEW QUESTION 64**
- (Exam Topic 2)
A company runs an intranet application on premises. The company wants to configure a cloud backup of the application. The company has selected AWS Elastic Disaster Recovery for this solution.
The company requires that replication traffic does not travel through the public internet. The application also must not be accessible from the internet. The company does not want this solution to consume all available network bandwidth because other applications require bandwidth.
Which combination of steps will meet these requirements? (Select THREE.)

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

F. During configuration of the launch settings for the target servers, select the option to ensure that the Recovery instance's private IP address matches the source server's private IP address.

**Answer:** BDE

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

**NEW QUESTION 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 71**
- (Exam Topic 2)
A company has many separate AWS accounts and uses no central billing or management. Each AWS account hosts services for different departments in the company. The company has a Microsoft Azure Active Directory that is deployed.
A solution architect needs to centralize billing and management of the company's AWS accounts. The company wants to start using identify federation instead of manual user management. The company also wants to use temporary credentials instead of long-lived access keys.
Which combination of steps will meet these requirements? (Select THREE)

A. Create a new AWS account to serve as a management accoun
B. Deploy an organization in AWS Organization
C. Invite each existing AWS account to join the organizatio
D. Ensure that each account accepts the invitation.
E. Configure each AWS Account's email address to be aws+<account id>@example.com so that account management email messages and invoices are sent to the same place.
F. Deploy AWS IAM Identity Center (AWS Single Sign-On) in the management accoun
G. Connect IAM Identity Center to the Azure Active Director
H. Configure IAM Identity Center for automatic synchronization of users and groups.
I. Deploy an AWS Managed Microsoft AD directory in the management accoun
J. Share the directory with all other accounts in the organization by using AWS Resource Access Manager (AWS RAM).
K. Create AWS IAM Identity Center (AWS Single Sign-On) permission set
L. Attach the permission sets to the appropriate IAM Identity Center groups and AWS accounts.
M. Configure AWS Identity and Access Management (IAM) in each AWS account to use AWS Managed Microsoft AD for authentication and authorization.

**Answer:** ACE

**NEW QUESTION 74**
- (Exam Topic 2)
A company has millions of objects in an Amazon S3 bucket. The objects are in the S3 Standard storage class. All the S3 objects are accessed frequently. The number of users and applications that access the objects is increasing rapidly. The objects are encrypted with server-side encryption with AWS KMS Keys (SSE-KMS).
A solutions architect reviews the company's monthly AWS invoice and notices that AWS KMS costs are increasing because of the high number of requests from Amazon S3. The solutions architect needs to optimize costs with minimal changes to the application.
Which solution will meet these requirements with the LEAST operational overhead?

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

**Answer:** B

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

**NEW QUESTION 78**
- (Exam Topic 2)
A company needs to establish a connection from its on-premises data center to AWS. The company needs to connect all of its VPCs that are located in different AWS Regions with transitive routing capabilities between VPC networks. The company also must reduce network outbound traffic costs, increase bandwidth throughput, and provide a consistent network experience for end users.
Which solution will meet these requirements?

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

**Answer:** B

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

**NEW QUESTION 80**
- (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 81**
- (Exam Topic 2)
A company uses AWS Organizations for a multi-account setup in the AWS Cloud. The company's finance team has a data processing application that uses AWS Lambda and Amazon DynamoDB. The company's marketing team wants to access the data that is stored in the DynamoDB table.
The DynamoDB table contains confidential data. The marketing team can have access to only specific attributes of data in the DynamoDB table. The fi-nance team and the marketing team have separate AWS accounts.
What should a solutions architect do to provide the marketing team with the appropriate access to the DynamoDB table?

A. Create an SCP to grant the marketing team's AWS account access to the specific attributes of the DynamoDB tabl
B. Attach the SCP to the OU of the finance team.
C. Create an IAM role in the finance team's account by using IAM policy conditions for specific DynamoDB attributes (fine-grained access con-trol). Establish trust with the marketing team's accoun
D. In the mar-keting team's account, create an IAM role that has permissions to as-sume the IAM role in the finance team's account.

E. Create a resource-based IAM policy that includes conditions for spe-cific DynamoDB attributes (fine-grained access control). Attach the policy to the DynamoDB tabl
F. In the marketing team'saccount, create an IAM role that has permissions to access the DynamoDB table in the finance team's account.
G. Create an IAM role in the finance team's account to access the Dyna-moDB tabl
H. Use an IAM permissions boundary to limit the access to the specific attribute
I. In the marketing team's account, create an IAM role that has permissions to assume the IAM role in the finance team's account.

**Answer:** C

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

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

References:

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

**NEW QUESTION 84**
- (Exam Topic 2)
A company has a website that runs on Amazon EC2 instances behind an Application Load Balancer (ALB). The instances are in an Auto Scaling group. The ALB is associated with an AWS WAF web ACL.
The website often encounters attacks in the application layer. The attacks produce sudden and significant increases in traffic on the application server. The access logs show that each attack originates from different IP addresses. A solutions architect needs to implement a solution to mitigate these attacks.
Which solution will meet these requirements with the LEAST operational overhead?

A. Create an Amazon CloudWatch alarm that monitors server acces
B. Set a threshold based on access by IP addres
C. Configure an alarm action that adds the IP address to the web ACL's deny list.
D. Deploy AWS Shield Advanced in addition to AWS WA
E. Add the ALB as a protected resource.
F. Create an Amazon CloudWatch alarm that monitors user IP addresse
G. Set a threshold based on access by IP addres
H. Configure the alarm to invoke an AWS Lambda function to add a deny rule in the application server's subnet route table for any IP addresses that activate the alarm.
I. Inspect access logs to find a pattern of IP addresses that launched the attack
J. Use an Amazon Route 53 geolocation routing policy to deny traffic from the countries that host those IP addresses.

**Answer:** C

**Explanation:**
"The AWS WAF API supports security automation such as blacklisting IP addresses that exceed request limits, which can be useful for mitigating HTTP flood attacks." >
https://aws.amazon.com/blogs/security/how-to-protect-dynamic-web-applications-against-ddos-attacks-by-using

**NEW QUESTION 86**
- (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 90**
- (Exam Topic 1)
A company is running an event ticketing platform on AWS and wants to optimize the platform's
cost-effectiveness. The platform is deployed on Amazon Elastic Kubernetes Service (Amazon EKS) with Amazon EC2 and is backed by an Amazon RDS for MySQL DB instance. The company is developing new application features to run on Amazon EKS with AWS Fargate.
The platform experiences infrequent high peaks in demand. The surges in demand depend on event dates. Which solution will provide the MOST cost-effective setup for the platform?

A. Purchase Standard Reserved Instances for the EC2 instances that the EKS cluster uses in its baseline loa
B. Scale the cluster with Spot Instances to handle peak
C. Purchase 1-year All Upfront Reserved Instances for the database to meet predicted peak load for the year.
D. Purchase Compute Savings Plans for the predicted medium load of the EKS cluste
E. Scale the cluster with On-Demand Capacity Reservations based on event dates for peak
F. Purchase 1-year No Upfront Reserved Instances for the database to meet the predicted base loa
G. Temporarily scale out database read replicas during peaks.
H. Purchase EC2 Instance Savings Plans for the predicted base load of the EKS cluste
I. Scale the cluster with Spot Instances to handle peak
J. Purchase 1-year All Upfront Reserved Instances for the database to meet the predicted base loa
K. Temporarily scale up the DB instance manually during peaks.
L. Purchase Compute Savings Plans for the predicted base load of the EKS cluste
M. Scale the cluster with Spot Instances to handle peak
N. Purchase 1-year All Upfront Reserved Instances for the database to meet the predicted base loa
O. Temporarily scale up the DB instance manually during peaks.

**Answer:** B

**Explanation:**
They all mention using spot instances and EKS based on EC2. A spot instance is not appropriate for a production server and the company is developing new application designed for AWS Fargate, which means we must plan the future cost improvement including AWS Fargate.
https://aws.amazon.com/savingsplans/compute-pricing/

**NEW QUESTION 91**
- (Exam Topic 1)
A company is developing a new service that will be accessed using TCP on a static port A solutions architect must ensure that the service is highly available, has redundancy across Availability Zones, and is accessible using the DNS name myservice.com, which is publicly accessible The service must use fixed address assignments so other companies can add the addresses to their allow lists.
Assuming that resources are deployed in multiple Availability Zones in a single Region, which solution will meet these requirements?

A. Create Amazon EC2 instances with an Elastic IP address for each instance Create a Network Load Balancer (NLB) and expose the static TCP port Register EC2instances with the NLB Create a new name server record set named my service com, and assign the Elastic IP addresses of the EC2 instances to the record set Provide the Elastic IP addresses of the EC2 instances to the other companies to add to their allow lists
B. Create an Amazon ECS cluster and a service definition for the application Create and assign public IP addresses for the ECS cluster Create a Network Load Balancer (NLB) and expose the TCP port Create atarget group and assign the ECS cluster name to the NLB Create a new A record set named my service com and assign the public IP addresses of the ECS cluster to the record set Provide the public IP addresses of the ECS cluster to the other companies to add to their allow lists
C. Create Amazon EC2 instances for the service Create one Elastic IP address for each Availability Zone Create a Network Load Balancer (NLB) and expose the assigned TCP port Assign the Elastic IP addresses to the NLB for each Availability Zone Create a target group and register the EC2 instances with the NLB Create a new A (alias) record set named my service com, and assign the NLB DNS name to the record set.
D. Create an Amazon ECS cluster and a service definition for the application Create and assign public IP address for each host in the cluster Create an Application Load Balancer (ALB) and expose the static TCP port Create a target group and assign the ECS service definition name to the ALB Create a new CNAME record set and associate the public IP addresses to the record set Provide the Elastic IP addresses of the Amazon EC2 instances to the other companies to add to their allow lists

**Answer:** C

**Explanation:**
https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-to-elb-load-balancer.html
Create a Network Load Balancer (NLB) and expose the assigned TCP port. Assign the Elastic IP addresses to the NLB for each Availability Zone. Create a target group and register the EC2 instances with the NLB. Create a new A (alias) record set named my.service.com, and assign the NLB DNS name to the record set. As it uses the NLB as the resource in the A-record, traffic will be routed through the NLB, and it will automatically route the traffic to the healthy instances based on the health checks and also it provides the fixed address assignments as the other companies can add the NLB's Elastic IP addresses to their allow lists.

**NEW QUESTION 92**
- (Exam Topic 1)
An application is using an Amazon RDS for MySQL Multi-AZ DB instance in the us-east-1 Region. After a failover test, the application lost the connections to the database and could not re-establish the connections. After a restart of the application, the application re-established the connections.
A solutions architect must implement a solution so that the application can re-establish connections to the database without requiring a restart.
Which solution will meet these requirements?

A. Create an Amazon Aurora MySQL Serverless v1 DB instanc
B. Migrate the RDS DB instance to the Aurora Serverless v1 DB instanc
C. Update the connection settings in the application to point to the Aurora reader endpoint.
D. Create an RDS prox
E. Configure the existing RDS endpoint as a targe
F. Update the connection settings in the application to point to the RDS proxy endpoint.
G. Create a two-node Amazon Aurora MySQL DB cluste
H. Migrate the RDS DB instance to the Aurora DB cluste
I. Create an RDS prox
J. Configure the existing RDS endpoint as a targe
K. Update the connection settings in the application to point to the RDS proxy endpoint.
L. Create an Amazon S3 bucke
M. Export the database to Amazon S3 by using AWS Database Migration Service (AWS DMS). Configure Amazon Athena to use the S3 bucket as a data stor
N. Install the latest Open Database Connectivity (ODBC) driver for the applicatio
O. Update the connection settings in the application to point to the Athena endpoint

**Answer:** B

**Explanation:**
Amazon RDS Proxy is a fully managed database proxy service for Amazon Relational Database Service (RDS) that makes applications more scalable, resilient, and secure. It allows applications to pool and share connections to an RDS database, which can help reduce database connection overhead, improve scalability, and provide automatic failover and high availability.

**NEW QUESTION 94**
- (Exam Topic 1)
A large mobile gaming company has successfully migrated all of its on-premises infrastructure to the AWS Cloud. A solutions architect is reviewing the environment to ensure that it was built according to the design and that it is running in alignment with the Well-Architected Framework.
While reviewing previous monthly costs in Cost Explorer, the solutions architect notices that the creation and subsequent termination of several large instance types account for a high proportion of the costs. The solutions architect finds out that the company's developers are launching new Amazon EC2 instances as part of their testing and that the developers are not using the appropriate instance types.
The solutions architect must implement a control mechanism to limit the instance types that only the developers can launch.
Which solution will meet these requirements?

A. Create a desired-instance-type managed rule in AWS Confi
B. Configure the rule with the instance types that are allowe
C. Attach the rule to an event to run each time a new EC2 instance is launched.
D. In the EC2 console, create a launch template that specifies the instance types that are allowe
E. Assign the launch template to the developers' IAM accounts.
F. Create a new IAM polic
G. Specify the instance types that are allowe
H. Attach the policy to an IAM group that contains the IAM accounts for the developers
I. Use EC2 Image Builder to create an image pipeline for the developers and assist them in the creation of a golden image.

**Answer:** C

**Explanation:**
This is doable with IAM policy creation to restrict users to specific instance types. Found the below article. https://blog.vizuri.com/limiting-allowed-aws-instance-type-with-iam-policy

**NEW QUESTION 97**
- (Exam Topic 1)
A company runs an loT platform on AWS loT sensors in various locations send data to the company's Node js API servers on Amazon EC2 instances running behind an Application Load Balancer The data is stored in an Amazon RDS MySQL DB instance that uses a 4 TB General Purpose SSD volume
The number of sensors the company has deployed in the field has increased over time and is expected to grow significantly The API servers are consistently overloaded and RDS metrics show high write latency
Which of the following steps together will resolve the issues permanently and enable growth as new sensors are provisioned, while keeping this platform cost-efficient? {Select TWO.)

A. Resize the MySQL General Purpose SSD storage to 6 TB to improve the volume's IOPS
B. Re-architect the database tier to use Amazon Aurora instead of an RDS MySQL DB instance and add read replicas
C. Leverage Amazon Kinesis Data Streams and AWS Lambda to ingest and process the raw data
D. Use AWS X-Ray to analyze and debug application issues and add more API servers to match the load
E. Re-architect the database tier to use Amazon DynamoDB instead of an RDS MySQL DB instance

**Answer:** CE

**Explanation:**
Option C is correct because leveraging Amazon Kinesis Data Streams and AWS Lambda to ingest and process the raw data resolves the issues permanently and enable growth as new sensors are provisioned. Amazon Kinesis Data Streams is a serverless streaming data service that simplifies the capture, processing, and storage of data streams at any scale. Kinesis Data Streams can handle any amount of streaming data and process data from hundreds of thousands of sources with very low latency. AWS Lambda is a serverless compute service that lets you run code without provisioning or managing servers. Lambda can be triggered by Kinesis Data Streams events and process the data records in real time. Lambda can also scale automatically based on the incoming data volume. By using Kinesis Data Streams and Lambda, the company can reduce the load on the API servers and improve the performance and scalability of the data ingestion and processing layer3
Option E is correct because re-architecting the database tier to use Amazon DynamoDB instead of an RDS MySQL DB instance resolves the issues permanently and enable growth as new sensors are provisioned. Amazon DynamoDB is a fully managed key-value and document database that delivers single-digit millisecond performance at any scale. DynamoDB supports auto scaling, which automatically adjusts read and write capacity based on actual traffic patterns. DynamoDB also supports on-demand capacity mode, which instantly accommodates up to double the previous peak traffic on a table. By using DynamoDB instead of RDS MySQL DB instance, the company can eliminate high write latency and improve scalability and performance of the database tier.
References: 1: https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-volume-types.html 2:

https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/CHAP_AuroraOverview.html 3:
https://docs.aws.amazon.com/streams/latest/dev/introduction.html : https://docs.aws.amazon.com/lambda/latest/dg/welcome.html :
https://docs.aws.amazon.com/xray/latest/devguide/aws-xray.html : https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Introduction.html :

**NEW QUESTION 100**
- (Exam Topic 1)
A company is refactoring its on-premises order-processing platform in the AWS Cloud. The platform includes a web front end that is hosted on a fleet of VMs RabbitMQ to connect the front end to the backend, and a Kubernetes cluster to run a containerized backend system to process the orders. The company does not want to make any major changes to the application
Which solution will meet these requirements with the LEAST operational overhead?

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

**Answer:** A

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

**NEW QUESTION 104**
- (Exam Topic 1)
A company has migrated its forms-processing application to AWS. When users interact with the application, they upload scanned forms as files through a web application. A database stores user metadata and references to files that are stored in Amazon S3. The web application runs on Amazon EC2 instances and an Amazon RDS for PostgreSQL database.
When forms are uploaded, the application sends notifications to a team through Amazon Simple Notification Service (Amazon SNS). A team member then logs in and processes each form. The team member performs data validation on the form and extracts relevant data before entering the information into another system that uses an API.
A solutions architect needs to automate the manual processing of the forms. The solution must provide accurate form extraction, minimize time to market, and minimize long-term operational overhead.
Which solution will meet these requirements?

A. Develop custom libraries to perform optical character recognition (OCR) on the form
B. Deploy the libraries to an Amazon Elastic Kubernetes Service (Amazon EKS) cluster as an application tie
C. Use this tier to process the forms when forms are uploade
D. Store the output in Amazon S3. Parse this output by extracting the data into an Amazon DynamoDB tabl
E. Submit the data to the target system's AP
F. Host the new application tier on EC2 instances.
G. Extend the system with an application tier that uses AWS Step Functions and AWS Lambd
H. Configure this tier to use artificial intelligence and machine learning (AI/ML) models that are trained and hosted on an EC2 instance to perform optical character recognition (OCR) on the forms when forms are uploade
I. Store the output in Amazon S3. Parse this output by extracting the data that is required within the application tie
J. Submit the data to the target system's API.
K. Host a new application tier on EC2 instance
L. Use this tier to call endpoints that host artificial intelligence and machine learning (AI/ML) models that are trained and hosted in Amazon SageMaker to perform optical character recognition (OCR) on the form
M. Store the output in Amazon ElastiCach
N. Parse this output by extracting the data that is required within the application tie
O. Submit the data to the target system's API.
P. Extend the system with an application tier that uses AWS Step Functions and AWS Lambd
Q. Configure this tier to use Amazon Textract and Amazon Comprehend to perform optical character recognition (OCR) on the forms when forms are uploade
R. Store the output in Amazon S3. Parse this output by extracting the data that is required within the application tie
S. Submit the data to the target system's API.

**Answer:** D

**Explanation:**
Extend the system with an application tier that uses AWS Step Functions and AWS Lambda. Configure this tier to use Amazon Textract and Amazon Comprehend to perform optical character recognition (OCR) on the forms when forms are uploaded. Store the output in Amazon S3. Parse this output by extracting the data that is required within the application tier. Submit the data to the target system's API. This solution meets the requirements of accurate form extraction, minimal time to market, and minimal long-term operational overhead. Amazon Textract and Amazon Comprehend are fully managed and serverless services that can perform OCR and extract relevant data from the forms, which eliminates the need to develop custom libraries or train and host models. Using AWS Step Functions and Lambda allows for easy automation of the process and the ability to scale as needed.

**NEW QUESTION 109**
- (Exam Topic 1)
A software as a service (SaaS) based company provides a case management solution to customers A3 part of the solution. The company uses a standalone Simple Mail Transfer Protocol (SMTP) server to send email messages from an application. The application also stores an email template for acknowledgement email messages that populate customer data before the application sends the email message to the customer.
The company plans to migrate this messaging functionality to the AWS Cloud and needs to minimize operational overhead.
Which solution will meet these requirements MOST cost-effectively?

A. Set up an SMTP server on Amazon EC2 instances by using an AMI from the AWS Marketplac
B. Store the email template in an Amazon S3 bucke

C. Create an AWS Lambda function to retrieve the template from the S3 bucket and to merge the customer data from the application with the templat
D. Use an SDK in the Lambda function to send the email message.
E. Set up Amazon Simple Email Service (Amazon SES) to send email message
F. Store the email template in an Amazon S3 bucke
G. Create an AWS Lambda function to retrieve the template from the S3 bucket and to merge the customer data from the application with the templat
H. Use an SDK in the Lambda function to send the email message.
I. Set up an SMTP server on Amazon EC2 instances by using an AMI from the AWS Marketplac
J. Store the email template in Amazon Simple Email Service (Amazon SES) with parameters for the customer dat
K. Create an AWS Lambda function to call the SES template and to pass customer data to replace the parameter
L. Use the AWS Marketplace SMTP server to send the email message.
M. Set up Amazon Simple Email Service (Amazon SES) to send email message
N. Store the email template on Amazon SES with parameters for the customer dat
O. Create an AWS Lambda function to call the SendTemplatedEmail API operation and to pass customer data to replace the parameters and the email destination.

**Answer:** D

**Explanation:**
In this solution, the company can use Amazon SES to send email messages, which will minimize operational overhead as SES is a fully managed service that handles sending and receiving email messages. The company can store the email template on Amazon SES with parameters for the customer data and use an AWS Lambda function to call the SendTemplatedEmail API operation, passing in the customer data to replace the parameters and the email destination. This solution eliminates the need to set up and manage an SMTP server on EC2 instances, which can be costly and time-consuming.

**NEW QUESTION 112**
- (Exam Topic 1)
A company uses AWS Organizations for a multi-account setup in the AWS Cloud. The company uses AWS Control Tower for governance and uses AWS Transit Gateway for VPC connectivity across accounts.
In an AWS application account, the company's application team has deployed a web application that uses AWS Lambda and Amazon RDS. The company's database administrators have a separate DBA account and use the account to centrally manage all the databases across the organization. The database administrators use an Amazon EC2 instance that is deployed in the DBA account to access an RDS database that is deployed in the application account.
The application team has stored the database credentials as secrets in AWS Secrets Manager in the application account. The application team is manually sharing the secrets with the database administrators. The secrets are encrypted by the default AWS managed key for Secrets Manager in the application account. A solutions architect needs to implement a solution that gives the database administrators access to the database and eliminates the need to manually share the secrets.
Which solution will meet these requirements?

A. Use AWS Resource Access Manager (AWS RAM) to share the secrets from the application account with the DBA accoun
B. In the DBA account, create an IAM role that is named DBA-Admi
C. Grant the role the required permissions to access the shared secret
D. Attach the DBA-Admin role to the EC2 instance for access to the cross-account secrets.
E. In the application account, create an IAM role that is named DBA-Secre
F. Grant the role the required permissions to access the secret
G. In the DBA account, create an IAM role that is named DBA-Admi
H. Grant the DBA-Admin role the required permissions to assume the DBA-Secret role in the application accoun
I. Attach the DBA-Admin role to the EC2 instance for access to the cross-account secrets.
J. In the DBA account, create an IAM role that is named DBA-Admi
K. Grant the role the required permissions to access the secrets and the default AWS managed key in the application accoun
L. In the application account, attach resource-based policies to the key to allow access from the DBA accoun
M. Attach the DBA-Admin role to the EC2 instance for access to the cross-account secrets.
N. In the DBA account, create an IAM role that is named DBA-Admi
O. Grant the role the required permissions to access the secrets in the application accoun
P. Attach an SCP to the application account to allow access to the secrets from the DBA accoun
Q. Attach the DBA-Admin role to the EC2 instance for access to the cross-account secrets.

**Answer:** B

**Explanation:**
Option B is correct because creating an IAM role in the application account that has permissions to access the secrets and creating an IAM role in the DBA account that has permissions to assume the role in the application account eliminates the need to manually share the secrets. This approach uses cross-account IAM roles to grant access to the secrets in the application account. The database administrators can assume the role in the application account from their EC2 instance in the DBA
account and retrieve the secrets without having to store them locally or share them manually2
References: 1: https://docs.aws.amazon.com/ram/latest/userguide/what-is.html 2: https://docs.aws.amazon.com/IAM/latest/UserGuide/tutorial_cross-account-with-roles.html 3: https://docs.aws.amazon.com/kms/latest/developerguide/concepts.html : https://docs.aws.amazon.com/secretsmanager/latest/userguide/tutorials_basic.html : https://docs.aws.amazon.com/IAM/latest/UserGuide/introduction.html

**NEW QUESTION 116**
- (Exam Topic 1)
An international delivery company hosts a delivery management system on AWS. Drivers use the system to upload confirmation of delivery. Confirmation includes the recipient's signature or a photo of the package with the recipient. The driver's handheld device uploads signatures and photos through FTP to a single Amazon EC2 instance. Each handheld device saves a file in a directory based on the signed-in user, and the file name matches the delivery number. The EC2 instance then adds metadata to the file after querying a central database to pull delivery information. The file is then placed in Amazon S3 for archiving.
As the company expands, drivers report that the system is rejecting connections. The FTP server is having problems because of dropped connections and memory issues. In response to these problems, a system engineer schedules a cron task to reboot the EC2 instance every 30 minutes. The billing team reports that files are not always in the archive and that the central system is not always updated.
A solutions architect needs to design a solution that maximizes scalability to ensure that the archive always receives the files and that systems are always updated. The handheld devices cannot be modified, so the company cannot deploy a new application.
Which solution will meet these requirements?

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

**Answer:** C

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

**NEW QUESTION 120**
- (Exam Topic 1)
A software company hosts an application on AWS with resources in multiple AWS accounts and Regions. The application runs on a group of Amazon EC2 instances in an application VPC located in the us-east-1 Region with an IPv4 CIDR block of 10.10.0.0/16. In a different AWS account, a shared services VPC is located in the us-east-2 Region with an IPv4 CIDR block of 10.10.10.0/24. When a cloud engineer uses AWS CloudFormation to attempt to peer the application VPC with the shared services VPC, an error message indicates a peering failure. Which factors could cause this error? (Choose two.)

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

**Answer:** AE

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

**NEW QUESTION 124**
.....

# Thank You for Trying Our Product

## We offer two products:

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

2nd - Questons and Answers in PDF Format

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

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

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

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

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