# Amazon-Web-Services

## Exam Questions SAP-C02

AWS Certified Solutions Architect - Professional

**NEW QUESTION 1**
- (Exam Topic 1)
A company is developing a new serverless API by using Amazon API Gateway and AWS Lambda. The company integrated the Lambda functions with API Gateway to use several shared libraries and custom classes.
A solutions architect needs to simplify the deployment of the solution and optimize for code reuse. Which solution will meet these requirements?

A. Deploy the shared libraries and custom classes into a Docker imag
B. Store the image in an S3 bucket.Create a Lambda layer that uses the Docker image as the sourc
C. Deploy the API's Lambda functions as Zip package
D. Configure the packages to use the Lambda layer.
E. Deploy the shared libraries and custom classes to a Docker imag
F. Upload the image to Amazon Elastic Container Registry (Amazon ECR). Create a Lambda layer that uses the Docker image as the sourc
G. Deploy the API's Lambda functions as Zip package
H. Configure the packages to use the Lambda layer.
I. Deploy the shared libraries and custom classes to a Docker container in Amazon Elastic Container Service (Amazon ECS) by using the AWS Fargate launch typ
J. Deploy the API's Lambda functions as Zip package
K. Configure the packages to use the deployed container as a Lambda layer.
L. Deploy the shared libraries, custom classes, and code for the API's Lambda functions to a Docker imag
M. Upload the image to Amazon Elastic Container Registry (Amazon ECR). Configure the API's Lambda functions to use the Docker image as the deployment package.

**Answer:** B

**Explanation:**
Deploying the shared libraries and custom classes to a Docker image and uploading the image to Amazon Elastic Container Registry (Amazon ECR) and creating a Lambda layer that uses the Docker image as the source. Then, deploying the API's Lambda functions as Zip packages and configuring the packages to use the Lambda layer would meet the requirements for simplifying the deployment and optimizing for code reuse.
A Lambda layer is a distribution mechanism for libraries, custom runtimes, and other function dependencies. It allows you to manage your in-development function code separately from your dependencies, this way you can easily update your dependencies without having to update your entire function code.
By deploying the shared libraries and custom classes to a Docker image and uploading the image to Amazon Elastic Container Registry (ECR), it makes it easy to manage and version the dependencies. This way, the company can use the same version of the dependencies across different Lambda functions.
By creating a Lambda layer that uses the Docker image as the source, the company can configure the API's Lambda functions to use the layer, reducing the need to include the dependencies in each function package, and making it easy to update the dependencies across all functions at once.
Reference:
AWS Lambda Layers documentation: https://docs.aws.amazon.com/lambda/latest/dg/configuration-layers.html
AWS Elastic Container Registry (ECR) documentation: https://aws.amazon.com/ecr/ Building Lambda Layers with Docker documentation:
https://aws.amazon.com/blogs/compute/building-lambda-layers-with-docker/

**NEW QUESTION 2**
- (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 3**
- (Exam Topic 1)
A company runs a content management application on a single Windows Amazon EC2 instance in a development environment. The application reads and writes static content to a 2 TB Amazon Elastic Block Store (Amazon EBS) volume that is attached to the instance as the root device. The company plans to deploy this application in production as a highly available and fault-tolerant solution that runs on at least three EC2 instances across multiple Availability Zones.
A solutions architect must design a solution that joins all the instances that run the application to an Active Directory domain. The solution also must implement Windows ACLs to control access to file contents. The application always must maintain exactly the same content on all running instances at any given point in time.
Which solution will meet these requirements with the LEAST management overhead?

A. Create an Amazon Elastic File System (Amazon EFS) file shar

B. Create an Auto Scaling group that extends across three Availability Zones and maintains a minimum size of three instance
C. Implement a user data script to install the application, join the instance to the AD domain, and mount the EFS file share.
D. Create a new AMI from the current EC2 instance that is runnin
E. Create an Amazon FSx for Lustre file syste
F. Create an Auto Scaling group that extends across three Availability Zones and maintains a minimum size of three instance
G. Implement a user data script to join the instance to the AD domain and mount the FSx for Lustre file system.
H. Create an Amazon FSx for Windows File Server file syste
I. Create an Auto Scaling group that extends across three Availability Zones and maintains a minimum size of three instance
J. Implement a user data script to install the application and mount the FSx for Windows File Server file syste
K. Perform a seamless domain join to join the instance to the AD domain.
L. Create a new AMI from the current EC2 instance that is runnin
M. Create an Amazon Elastic File System (Amazon EFS) file syste
N. Create an Auto Scaling group that extends across three Availability Zones and maintains a minimum size of three instance
O. Perform a seamless domain join to join the instance to the AD domain.

**Answer:** C

**Explanation:**
https://docs.aws.amazon.com/fsx/latest/WindowsGuide/what-is.html https://docs.aws.amazon.com/directoryservice/latest/admin-guide/ms_ad_join_instance.html


**NEW QUESTION 4**
- (Exam Topic 1)
A company runs an IoT platform on AWS IoT 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 5**
- (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 6**
- (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 upload
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 7
- (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 8
- (Exam Topic 1)
A company has a serverless application comprised of Amazon CloudFront, Amazon API Gateway, and AWS Lambda functions. The current deployment process of the application code is to create a new version number of the Lambda function and run an AWS CLI script to update. If the new function version has errors, another CLI script reverts by deploying the previous working version of the function. The company would like to decrease the time to deploy new versions of the application logic provided by the Lambda functions, and also reduce the time to detect and revert when errors are identified.

How can this be accomplished?

A. Create and deploy nested AWS CloudFormation stacks with the parent stack consisting of the AWS CloudFront distribution and API Gateway, and the child stack containing the Lambda functio
B. For changes to Lambda, create an AWS CloudFormation change set and deploy; if errors are triggered, revert the AWS CloudFormation change set to the previous version.
C. Use AWS SAM and built-in AWS CodeDeploy to deploy the new Lambda version, gradually shift traffic to the new version, and use pre-traffic and post-traffic test functions to verify cod
D. Rollback if Amazon CloudWatch alarms are triggered.
E. Refactor the AWS CLI scripts into a single script that deploys the new Lambda versio
F. When deployment is completed, the script tests execut
G. If errors are detected, revert to the previous Lambda version.
H. Create and deploy an AWS CloudFormation stack that consists of a new API Gateway endpoint that references the new Lambda versio
I. Change the CloudFront origin to the new API Gateway endpoint, monitor errors and if detected, change the AWS CloudFront origin to the previous API Gateway endpoint.

**Answer:** B

**Explanation:**
https://aws.amazon.com/about-aws/whats-new/2017/11/aws-lambda-supports-traffic-shifting-and-phased-deploy

**NEW QUESTION 9**
- (Exam Topic 1)
A company has an organization in AWS Organizations that has a large number of AWS accounts. One of the AWS accounts is designated as a transit account and has a transit gateway that is shared with all of the other AWS accounts AWS Site-to-Site VPN connections are configured between ail of the company's global offices and the transit account The company has AWS Config enabled on all of its accounts.
The company's networking team needs to centrally manage a list of internal IP address ranges that belong to the global offices Developers Will reference this list to gain access to applications securely.
Which solution meets these requirements with the LEAST amount of operational overhead?

A. Create a JSON file that is hosted in Amazon S3 and that lists all of the internal IP address ranges Configure an Amazon Simple Notification Service (Amazon SNS) topic in each of the accounts that can be involved when the JSON file is update
B. Subscribe an AWS Lambda function to the SNS topic to update all relevant security group rules with Vie updated IP address ranges.
C. Create a new AWS Config managed rule that contains all of the internal IP address ranges Use the rule to check the security groups in each of the accounts to ensure compliance with the list of IP address range
D. Configure the rule to automatically remediate any noncompliant security group that is detected.
E. In the transit account, create a VPC prefix list with all of the internal IP address range
F. Use AWS Resource Access Manager to share the prefix list with all of the other account
G. Use the shared prefix list to configure security group rules is the other accounts.
H. In the transit account create a security group with all of the internal IP address range
I. Configure the security groups in me other accounts to reference the transit account's securitygroup by using a nested security group reference of *<transit-account-id>./sg-1a2b3c4d".

**Answer:** C

**Explanation:**
Customer-managed prefix lists — Sets of IP address ranges that you define and manage. You can share your prefix list with other AWS accounts, enabling those accounts to reference the prefix list in their own resources. https://docs.aws.amazon.com/vpc/latest/userguide/managed-prefix-lists.html
a VPC prefix list is created in the transit account with all of the internal IP address ranges, and then shared to all of the other accounts using AWS Resource Access Manager. This allows for central management of the IP address ranges, and eliminates the need for manual updates to security group rules in each account. This solution also allows for compliance checks to be run using AWS Config and for any non-compliant security groups to be automatically remediated.

**NEW QUESTION 10**
- (Exam Topic 1)
A life sciences company is using a combination of open source tools to manage data analysis workflows and Docker containers running on servers in its on-premises data center to process genomics data Sequencing data is generated and stored on a local storage area network (SAN), and then the data is processed. The research and development teams are running into capacity issues and have decided to re-architect their genomics analysis platform on AWS to scale based on workload demands and reduce the turnaround time from weeks to days
The company has a high-speed AWS Direct Connect connection Sequencers will generate around 200 GB of data for each genome, and individual jobs can take several hours to process the data with ideal compute capacity. The end result will be stored in Amazon S3. The company is expecting 10-15 job requests each day
Which solution meets these requirements?

A. Use regularly scheduled AWS Snowball Edge devices to transfer the sequencing data into AWS When AWS receives the Snowball Edge device and the data is loaded into Amazon S3 use S3 events to trigger an AWS Lambda function to process the data
B. Use AWS Data Pipeline to transfer the sequencing data to Amazon S3 Use S3 events to trigger an Amazon EC2 Auto Scaling group to launch custom-AMI EC2 instances running the Docker containers to process the data
C. Use AWS DataSync to transfer the sequencing data to Amazon S3 Use S3 events to trigger an AWS Lambda function that starts an AWS Step Functions workflow Store the Docker images in Amazon Elastic Container Registry (Amazon ECR) and trigger AWS Batch to run the container and process the sequencing data
D. Use an AWS Storage Gateway file gateway to transfer the sequencing data to Amazon S3 Use S3 events to trigger an AWS Batch job that runs on Amazon EC2 instances running the Docker containers to process the data

**Answer:** C

**Explanation:**
AWS DataSync can be used to transfer the sequencing data to Amazon S3, which is a more efficient and faster method than using Snowball Edge devices. Once the data is in S3, S3 events can trigger an AWS Lambda function that starts an AWS Step Functions workflow. The Docker images can be stored in Amazon Elastic Container Registry (Amazon ECR) and AWS Batch can be used to run the container and process the sequencing data.

**NEW QUESTION 10**

- (Exam Topic 1)
A company is planning to migrate its business-critical applications from an on-premises data center to AWS. The company has an on-premises installation of a Microsoft SQL Server Always On cluster. The company wants to migrate to an AWS managed database service. A solutions architect must design a heterogeneous database migration on AWS.
Which solution will meet these requirements?

A. Migrate the SQL Server databases to Amazon RDS for MySQL by using backup and restore utilities.
B. Use an AWS Snowball Edge Storage Optimized device to transfer data to Amazon S3. Set up Amazon RDS for MySQ
C. Use S3 integration with SQL Server features, such as BULK INSERT.
D. Use the AWS Schema Conversion Tool to translate the database schema to Amazon RDS for MeSQ
E. Then use AWS Database Migration Service (AWS DMS) to migrate the data from on-premises databases to Amazon RDS.
F. Use AWS DataSync to migrate data over the network between on-premises storage and Amazon S3. Set up Amazon RDS for MySQ
G. Use S3 integration with SQL Server features, such as BULK INSERT.

**Answer:** C

**Explanation:**
https://aws.amazon.com/dms/schema-conversion-tool/
AWS Schema Conversion Tool (SCT) can automatically convert the database schema from Microsoft SQL Server to Amazon RDS for MySQL. This allows for a smooth transition of the database schema without any manual intervention. AWS DMS can then be used to migrate the data from the on-premises databases to the newly created Amazon RDS for MySQL instance. This service can perform a one-time migration of the data or can set up ongoing replication of data changes to keep the on-premises and AWS databases in sync.

## NEW QUESTION 11
- (Exam Topic 1)
An enterprise company wants to allow its developers to purchase third-party software through AWS Marketplace. The company uses an AWS Organizations account structure with full features enabled, and has a shared services account in each organizational unit (OU) that will be used by procurement managers. The procurement team's policy indicates that developers should be able to obtain third-party software from an approved list only and use Private Marketplace in AWS Marketplace to achieve this requirement . The procurement team wants administration of Private Marketplace to be restricted to a role named procurement-manager-role, which could be assumed by procurement managers Other IAM users groups, roles, and account administrators in the company should be denied Private Marketplace administrative access
What is the MOST efficient way to design an architecture to meet these requirements?

A. Create an IAM role named procurement-manager-role in all AWS accounts in the organization Add the PowerUserAccess managed policy to the role Apply an inline policy to all IAM users and roles in every AWS account to deny permissions on the AWSPrivateMarketplaceAdminFullAccess managed policy.
B. Create an IAM role named procurement-manager-role in all AWS accounts in the organization Add the AdministratorAccess managed policy to the role Define a permissions boundary with the AWSPrivateMarketplaceAdminFullAccess managed policy and attach it to all the developer roles.
C. Create an IAM role named procurement-manager-role in all the shared services accounts in the organization Add the AWSPrivateMarketplaceAdminFullAccess managed policy to the role Create an organization root-level SCP to deny permissions to administer Private Marketplace to everyone except the role named procurement-manager-role Create another organization root-level SCP to deny permissions to create an IAM role named procurement-manager-role to everyone in the organization.
D. Create an IAM role named procurement-manager-role in all AWS accounts that will be used by developer
E. Add the AWSPrivateMarketplaceAdminFullAccess managed policy to the rol
F. Create an SCP in Organizations to deny permissions to administer Private Marketplace to everyone except the role named procurement-manager-rol
G. Apply the SCP to all the shared services accounts in the organization.

**Answer:** C

**Explanation:**
SCP to deny permissions to administer Private Marketplace to everyone except the role named procurement-manager-role.
https://aws.amazon.com/blogs/awsmarketplace/controlling-access-to-a-well-architected-private-marketplace-usi
This approach allows the procurement managers to assume the procurement-manager-role in shared services accounts, which have the AWSPrivateMarketplaceAdminFullAccess managed policy attached to it and can then manage the Private Marketplace. The organization root-level SCP denies the permission to administer Private Marketplace to everyone except the role named procurement-manager-role and another SCP denies the permission to create an IAM role named procurement-manager-role to everyone in the organization, ensuring that only the procurement team can assume the role and manage the Private Marketplace. This approach provides a centralized way to manage and restrict access to Private Marketplace while maintaining a high level of security.

## NEW QUESTION 14
- (Exam Topic 1)
A company has created an OU in AWS Organizations for each of its engineering teams Each OU owns multiple AWS accounts. The organization has hundreds of AWS accounts A solutions architect must design a solution so that each OU can view a breakdown of usage costs across its AWS accounts. Which solution meets these requirements?

A. Create an AWS Cost and Usage Report (CUR) for each OU by using AWS Resource Access Manager Allow each team to visualize the CUR through an Amazon QuickSight dashboard.
B. Create an AWS Cost and Usage Report (CUR) from the AWS Organizations management account- Allow each team to visualize the CUR through an Amazon QuickSight dashboard
C. Create an AWS Cost and Usage Report (CUR) in each AWS Organizations member account Allow each team to visualize the CUR through an Amazon QuickSight dashboard.
D. Create an AWS Cost and Usage Report (CUR) by using AWS Systems Manager Allow each team to visualize the CUR through Systems Manager OpsCenter dashboards

**Answer:** B

**Explanation:**
https://docs.aws.amazon.com/cur/latest/userguide/billing-cur-limits.html

## NEW QUESTION 17
- (Exam Topic 1)
A security engineer determined that an existing application retrieves credentials to an Amazon RDS for MySQL database from an encrypted file in Amazon S3. For

the next version of the application, the security engineer wants to implement the following application design changes to improve security:

➢ The database must use strong, randomly generated passwords stored in a secure AWS managed service.
➢ The application resources must be deployed through AWS CloudFormation.
➢ The application must rotate credentials for the database every 90 days.

A solutions architect will generate a CloudFormation template to deploy the application.
Which resources specified in the CloudFormation template will meet the security engineer's requirements with the LEAST amount of operational overhead?

A. Generate the database password as a secret resource using AWS Secrets Manage
B. Create an AWS Lambda function resource to rotate the database passwor
C. Specify a Secrets Manager RotationSchedule resource to rotate the database password every 90 days.
D. Generate the database password as a SecureString parameter type using AWS Systems Manager Parameter Stor
E. Create an AWS Lambda function resource to rotate the database passwor
F. Specify a Parameter Store RotationSchedule resource to rotate the database password every 90 days.
G. Generate the database password as a secret resource using AWS Secrets Manage
H. Create an AWS Lambda function resource to rotate the database passwor
I. Create an Amazon EventBridge scheduled rule resource to trigger the Lambda function password rotation every 90 days.
J. Generate the database password as a SecureString parameter type using AWS Systems Manager Parameter Stor
K. Specify an AWS AppSync DataSource resource to automatically rotate the database password every 90 days.

**Answer:** B

**Explanation:**
https://aws.amazon.com/blogs/security/how-to-securely-provide-database-credentials-to-lambda-functions-by-us
https://docs.aws.amazon.com/secretsmanager/latest/userguide/rotating-secrets.html
https://docs.aws.amazon.com/secretsmanager/latest/userguide/integrating_cloudformation.html

**NEW QUESTION 20**
- (Exam Topic 1)
A company uses Amazon S3 to store files and images in a variety of storage classes. The company's S3 costs have increased substantially during the past year.
A solutions architect needs to review data trends for the past 12 months and identity the appropriate storage class for the objects.
Which solution will meet these requirements?

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

**Answer:** B

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

**NEW QUESTION 25**
- (Exam Topic 1)
A company developed a pilot application by using AWS Elastic Beanstalk and Java. To save costs during development, the company's development team deployed the application into a single-instance environment. Recent tests indicate that the application consumes more CPU than expected. CPU utilization is regularly greater than 85%, which causes some performance bottlenecks.
A solutions architect must mitigate the performance issues before the company launches the application to production.
Which solution will meet these requirements with the LEAST operational overhead?

A. Create a new Elastic Beanstalk applicatio
B. Select a load-balanced environment typ
C. Select all Availability Zone
D. Add a scale-out rule that will run if the maximum CPU utilization is over 85% for 5 minutes.
E. Create a second Elastic Beanstalk environmen
F. Apply the traffic-splitting deployment polic
G. Specify a percentage of incoming traffic to direct to the new environment in the average CPU utilization is over 85% for 5 minutes.
H. Modify the existing environment's capacity configuration to use a load-balanced environment type.Select all Availability Zone
I. Add a scale-out rule that will run if the average CPU utilization is over 85% for 5 minutes.
J. Select the Rebuild environment action with the load balancing option Select an Availability Zones Add a scale-out rule that will run if the sum CPU utilization is over 85% for 5 minutes.

**Answer:** C

**Explanation:**
This solution will meet the requirements with the least operational overhead because it allows the company to modify the existing environment's capacity configuration, so it becomes a load-balanced environment type. By selecting all availability zones, the company can ensure that the application is running in multiple availability zones, which can help to improve the availability and scalability of the application. The company can also add a scale-out rule that will run if the average CPU utilization is over 85% for 5 minutes, which can help to mitigate the performance issues. This solution does not require creating new Elastic Beanstalk environments or rebuilding the existing one, which reduces the operational overhead.
You can refer to the AWS Elastic Beanstalk documentation for more information on how to use this service: https://aws.amazon.com/elasticbeanstalk/ You can refer to the AWS documentation for more information on how to use autoscaling: https://aws.amazon.com/autoscaling/

**NEW QUESTION 30**
- (Exam Topic 1)

A company is subject to regulatory audits of its financial information. External auditors who use a single AWS account need access to the company's AWS account. A solutions architect must provide the auditors with secure, read-only access to the company's AWS account. The solution must comply with AWS security best practices.
Which solution will meet these requirements?

A. In the company's AWS account, create resource policies for all resources in the account to grant access to the auditors' AWS accoun
B. Assign a unique external ID to the resource policy.
C. In the company's AWS account create an IAM role that trusts the auditors' AWS account Create an IAM policy that has the required permission
D. Attach the policy to the rol
E. Assign a unique external ID to the role's trust policy.
F. In the company's AWS account, create an IAM use
G. Attach the required IAM policies to the IAM user.Create API access keys for the IAM use
H. Share the access keys with the auditors.
I. In the company's AWS account, create an IAM group that has the required permissions Create an IAM user in the company s account for each audito
J. Add the IAM users to the IAM group.

**Answer:** B

**Explanation:**
This solution will allow the external auditors to have read-only access to the company's AWS account while being compliant with AWS security best practices. By creating an IAM role, which is a secure and flexible way of granting access to AWS resources, and trusting the auditors' AWS account, the company can ensure that the auditors only have the permissions that are required for their role and nothing more. Assigning a unique external ID to the role's trust policy, it will ensure that only the auditors' AWS account can assume the role.
Reference:
AWS IAM Roles documentation: https://aws.amazon.com/iam/features/roles/ AWS IAM Best practices: https://aws.amazon.com/iam/security-best-practices/


**NEW QUESTION 33**
- (Exam Topic 1)
A company is using an on-premises Active Directory service for user authentication. The company wants to use the same authentication service to sign in to the company's AWS accounts, which are using AWS Organizations. AWS Site-to-Site VPN connectivity already exists between the on-premises environment and all the company's AWS accounts.
The company's security policy requires conditional access to the accounts based on user groups and roles. User identities must be managed in a single location.
Which solution will meet these requirements?

A. Configure AWS Single Sign-On (AWS SSO) to connect to Active Directory by using SAML 2.0.Enable automatic provisioning by using the System for Cross-domain Identity Management (SCIM) v2.0 protoco
B. Grant access to the AWS accounts by using attribute-based access controls (ABACs).
C. Configure AWS Single Sign-On (AWS SSO) by using AWS SSO as an identity sourc
D. Enable automatic provisioning by using the System for Cross-domain Identity Management (SCIM) v2.0 protoco
E. Grant access to the AWS accounts by using AWS SSO permission sets.
F. In one of the company's AWS accounts, configure AWS Identity and Access Management (IAM) to use a SAML 2.0 identity provide
G. Provision IAM users that are mapped to the federated user
H. Grant access that corresponds to appropriate groups in Active Director
I. Grant access to the required AWS accounts by using cross-account IAM users.
J. In one of the company's AWS accounts, configure AWS Identity and Access Management (IAM) to use an OpenID Connect (OIDC) identity provide
K. Provision IAM roles that grant access to the AWS account for the federated users that correspond to appropriate groups in Active Director
L. Grant access to the required AWS accounts by using cross-account IAM roles.

**Answer:** D

**Explanation:**
https://aws.amazon.com/blogs/aws/new-attributes-based-access-control-with-aws-single-sign-on/


**NEW QUESTION 34**
- (Exam Topic 1)
A company has a latency-sensitive trading platform that uses Amazon DynamoDB as a storage backend. The company configured the DynamoDB table to use on-demand capacity mode. A solutions architect needs to design a solution to improve the performance of the trading platform. The new solution must ensure high availability for the trading platform.
Which solution will meet these requirements with the LEAST latency?

A. Create a two-node DynamoDB Accelerator (DAX) cluster Configure an application to read and write data by using DAX.
B. Create a three-node DynamoDB Accelerator (DAX) cluste
C. Configure an application to read data by using DAX and to write data directly to the DynamoDB table.
D. Create a three-node DynamoDB Accelerator (DAX) cluste
E. Configure an application to read data directly from the DynamoDB table and to write data by using DAX.
F. Create a single-node DynamoD8 Accelerator (DAX) cluste
G. Configure an application to read data by using DAX and to write data directly to the DynamoD8 table.

**Answer:** B

**Explanation:**
A DAX cluster can be deployed with one or two nodes for development or test workloads. One- and two-node clusters are not fault-tolerant, and we don't recommend using fewer than three nodes for production use. If a one- or two-node cluster encounters software or hardware errors, the cluster can become unavailable or lose cached data.A DAX cluster can be deployed with one or two nodes for development or test workloads. One and two-node clusters are not fault-tolerant, and we don't recommend using fewer than three nodes for production use. If a one- or two-node cluster encounters software or hardware errors, the cluster can become unavailable or lose cached data.
https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/DAX.concepts.cluster.html


**NEW QUESTION 39**
- (Exam Topic 1)

A company recently acquired several other companies. Each company has a separate AWS account with a different billing and reporting method. The acquiring company has consolidated all the accounts into one organization in AWS Organizations. However, the acquiring company has found it difficult to generate a cost report that contains meaningful groups for all the teams.

The acquiring company's finance team needs a solution to report on costs for all the companies through a self-managed application.
Which solution will meet these requirements?

A. Create an AWS Cost and Usage Report for the organizatio
B. Define tags and cost categories in the repor
C. Create a table in Amazon Athen
D. Create an Amazon QuickSight dataset based on the Athena tabl
E. Share the dataset with the finance team.
F. Create an AWS Cost and Usage Report for the organizatio
G. Define tags and cost categories in the repor
H. Create a specialized template in AWS Cost Explorer that the finance department will use to build reports.
I. Create an Amazon QuickSight dataset that receives spending information from the AWS Price List Query AP
J. Share the dataset with the finance team.
K. Use the AWS Price List Query API to collect account spending informatio
L. Create a specialized template in AWS Cost Explorer that the finance department will use to build reports.

**Answer:** A

**Explanation:**
Creating an AWS Cost and Usage Report for the organization and defining tags and cost categories in the report will allow for detailed cost reporting for the different companies that have been consolidated into one organization. By creating a table in Amazon Athena and an Amazon QuickSight dataset based on the Athena table, the finance team will be able to easily query and generate reports on the costs for all the companies. The dataset can then be shared with the finance team for them to use for their reporting needs.

**NEW QUESTION 43**
- (Exam Topic 1)
A company hosts a Git repository in an on-premises data center. The company uses webhooks to invoke functionality that runs in the AWS Cloud. The company hosts the webhook logic on a set of Amazon EC2 instances in an Auto Scaling group that the company set as a target for an Application Load Balancer (ALB). The Git server calls the ALB for the configured webhooks. The company wants to move the solution to a serverless architecture.
Which solution will meet these requirements with the LEAST operational overhead?

A. For each webhook, create and configure an AWS Lambda function UR
B. Update the Git servers to call the individual Lambda function URLs.
C. Create an Amazon API Gateway HTTP AP
D. Implement each webhook logic in a separate AWS Lambda functio
E. Update the Git servers to call the API Gateway endpoint.
F. Deploy the webhook logic to AWS App Runne
G. Create an ALB, and set App Runner as the target.Update the Git servers to call the ALB endpoint.
H. Containerize the webhook logi
I. Create an Amazon Elastic Container Service (Amazon ECS) cluster, and run the webhook logic in AWS Fargat
J. Create an Amazon API Gateway REST API, and set Fargate as the targe
K. Update the Git servers to call the API Gateway endpoint.

**Answer:** B

**Explanation:**
https://aws.amazon.com/solutions/implementations/git-to-s3-using-webhooks/ https://medium.com/mindorks/building-webhook-is-easy-using-aws-lambda-and-api-gateway-56f5e5c3a596

**NEW QUESTION 47**
- (Exam Topic 1)
A company uses an on-premises data analytics platform. The system is highly available in a fully redundant configuration across 12 servers in the company's data center.
The system runs scheduled jobs, both hourly and daily, in addition to one-time requests from users. Scheduled jobs can take between 20 minutes and 2 hours to finish running and have tight SLAs. The scheduled jobs account for 65% of the system usage. User jobs typically finish running in less than 5 minutes and have no SLA. The user jobs account for 35% of system usage. During system failures, scheduled jobs must continue to meet SLAs. However, user jobs can be delayed.
A solutions architect needs to move the system to Amazon EC2 instances and adopt a consumption-based model to reduce costs with no long-term commitments. The solution must maintain high availability and must not affect the SLAs.
Which solution will meet these requirements MOST cost-effectively?

A. Split the 12 instances across two Availability Zones in the chosen AWS Regio
B. Run two instances in each Availability Zone as On-Demand Instances with Capacity Reservation
C. Run four instances in each Availability Zone as Spot Instances.
D. Split the 12 instances across three Availability Zones in the chosen AWS Regio
E. In one of the Availability Zones, run all four instances as On-Demand Instances with Capacity Reservation
F. Run the remaining instances as Spot Instances.
G. Split the 12 instances across three Availability Zones in the chosen AWS Regio
H. Run two instances in each Availability Zone as On-Demand Instances with a Savings Pla
I. Run two instances in each Availability Zone as Spot Instances.
J. Split the 12 instances across three Availability Zones in the chosen AWS Regio
K. Run three instances in each Availability Zone as On-Demand Instances with Capacity Reservation
L. Run one instance in each Availability Zone as a Spot Instance.

**Answer:** D

**Explanation:**
By splitting the 12 instances across three Availability Zones, the system can maintain high availability and availability of resources in case of a failure. Option D also uses a combination of On-Demand Instances with Capacity Reservations and Spot Instances, which allows for scheduled jobs to be run on the On-Demand

instances with guaranteed capacity, while also taking advantage of the cost savings from Spot Instances for the user jobs which have lower SLA requirements.

**NEW QUESTION 48**
- (Exam Topic 1)
A financial services company in North America plans to release a new online web application to its customers on AWS . The company will launch the application in the us-east-1 Region on Amazon EC2 instances. The application must be highly available and must dynamically scale to meet user traffic. The company also wants to implement a disaster recovery environment for the application in the us-west-1 Region by using active-passive failover.
Which solution will meet these requirements?

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

**Answer:** C

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

**NEW QUESTION 52**
- (Exam Topic 1)
A company is running a data-intensive application on AWS. The application runs on a cluster of hundreds of Amazon EC2 instances. A shared file system also runs on several EC2 instances that store 200 TB of data. The application reads and modifies the data on the shared file system and generates a report. The job runs once monthly, reads a subset of the files from the shared file system, and takes about 72 hours to complete. The compute instances scale in an Auto Scaling group, but the instances that host the shared file system run continuously. The compute and storage instances are all in the same AWS Region.
A solutions architect needs to reduce costs by replacing the shared file system instances. The file system must provide high performance access to the needed data for the duration of the 72-hour run.
Which solution will provide the LARGEST overall cost reduction while meeting these requirements?

A. Migrate the data from the existing shared file system to an Amazon S3 bucket that uses the S3 Intelligent-Tiering storage clas
B. Before the job runs each month, use Amazon FSx for Lustre to create a new file system with the data from Amazon S3 by using lazy loadin
C. Use the new file system as the shared storage for the duration of the jo
D. Delete the file system when the job is complete.
E. Migrate the data from the existing shared file system to a large Amazon Elastic Block Store (Amazon EBS) volume with Multi-Attach enable
F. Attach the EBS volume to each of the instances by using a user data script in the Auto Scaling group launch templat
G. Use the EBS volume as the shared storage for the duration of the jo
H. Detach the EBS volume when the job is complete.
I. Migrate the data from the existing shared file system to an Amazon S3 bucket that uses the S3 Standard storage clas
J. Before the job runs each month, use Amazon FSx for Lustre to create a new file system with the data from Amazon S3 by using batch loadin
K. Use the new file system as the shared storage for the duration of the jo
L. Delete the file system when the job is complete.
M. Migrate the data from the existing shared file system to an Amazon S3 bucke
N. Before the job runs each month, use AWS Storage Gateway to create a file gateway with the data from Amazon S3. Use the file gateway as the shared storage for the jo
O. Delete the file gateway when the job is complete.

**Answer:** A

**Explanation:**
https://aws.amazon.com/blogs/storage/new-enhancements-for-moving-data-between-amazon-fsx-for-lustre-and

**NEW QUESTION 54**
- (Exam Topic 1)
A company is building a software-as-a-service (SaaS) solution on AWS. The company has deployed an Amazon API Gateway REST API with AWS Lambda integration in multiple AWS Regions and in the same production account.
The company offers tiered pricing that gives customers the ability to pay for the capacity to make a certain number of API calls per second. The premium tier offers up to 3,000 calls per second, and customers are identified by a unique API key. Several premium tier customers in various Regions report that they receive error responses of 429 Too Many Requests from multiple API methods during peak usage hours. Logs indicate that the Lambda function is never invoked.
What could be the cause of the error messages for these customers?

A. The Lambda function reached its concurrency limit.
B. The Lambda function its Region limit for concurrency.
C. The company reached its API Gateway account limit for calls per second.
D. The company reached its API Gateway default per-method limit for calls per second.

**Answer:** C

**Explanation:**

https://docs.aws.amazon.com/apigateway/latest/developerguide/api-gateway-request-throttling.html#apig-reques

**NEW QUESTION 57**
- (Exam Topic 1)
A company has purchased appliances from different vendors. The appliances all have IoT sensors. The sensors send status information in the vendors' proprietary formats to a legacy application that parses the information into JSON. The parsing is simple, but each vendor has a unique format. Once daily, the application parses all the JSON records and stores the records in a relational database for analysis.
The company needs to design a new data analysis solution that can deliver faster and optimize costs. Which solution will meet these requirements?

A. Connect the IoT sensors to AWS IoT Cor
B. Set a rule to invoke an AWS Lambda function to parse the information and save a .csv file to Amazon S3. Use AWS Glue to catalog the file
C. Use Amazon Athena and Amazon OuickSight for analysis.
D. Migrate the application server to AWS Fargate, which will receive the information from IoT sensors and parse the information into a relational forma
E. Save the parsed information to Amazon Redshift for analysis.
F. Create an AWS Transfer for SFTP serve
G. Update the IoT sensor code to send the information as a .csv file through SFTP to the serve
H. Use AWS Glue to catalog the file
I. Use Amazon Athena for analysis.
J. Use AWS Snowball Edge to collect data from the IoT sensors directly to perform local analysis.Periodically collect the data into Amazon Redshift to perform global analysis.

**Answer:** A

**Explanation:**
⟩ Connect the IoT sensors to AWS IoT Core. Set a rule to invoke an AWS Lambda function to parse the information and save a .csv file to Amazon S3. Use AWS Glue to catalog the files. Use Amazon Athena and Amazon QuickSight for analysis. This solution meets the requirement of faster analysis and cost optimization by using AWS IoT Core to collect data from the IoT sensors in real-time and then using AWS Glue and Amazon Athena for efficient data analysis. This solution involves connecting the IoT sensors to the AWS IoT Core, setting a rule to invoke an AWS Lambda function to parse the information, and saving a .csv file to Amazon S3. AWS Glue can be used to catalog the files and Amazon Athena and Amazon QuickSight can be used for analysis. This solution will enable faster and more cost-effective data analysis.
This solution is in line with the official Amazon Textbook and Resources for the AWS Certified Solutions Architect - Professional certification. In particular, the book states that: "AWS IoT Core can be used to ingest and process the data, AWS Lambda can be used to process and transform the data, and Amazon S3 can be used to store the data. AWS Glue can be used to catalog and access the data, Amazon Athena can be used to query the data, and Amazon QuickSight can be used to visualize the data." (Source: https://d1.awsstatic.com/training-and-certification/docs-sa-pro/AWS_Certified_Solutions_Architect_Professiona

**NEW QUESTION 59**
- (Exam Topic 1)
An AWS customer has a web application that runs on premises. The web application fetches data from a third-party API that is behind a firewall. The third party accepts only one public CIDR block in each client's allow list.
The customer wants to migrate their web application to the AWS Cloud. The application will be hosted on a set of Amazon EC2 instances behind an Application Load Balancer (ALB) in a VPC. The ALB is located in public subnets. The EC2 instances are located in private subnets. NAT gateways provide internet access to the private subnets.
How should a solutions architect ensure that the web application can continue to call the third-parly API after the migration?

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

**Answer:** B

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

**NEW QUESTION 64**
- (Exam Topic 1)
A company has a web application that allows users to upload short videos. The videos are stored on Amazon EBS volumes and analyzed by custom recognition software for categorization.
The website contains stat c content that has variable traffic with peaks in certain months. The architecture consists of Amazon EC2 instances running in an Auto Scaling group for the web application and EC2
instances running in an Auto Scaling group to process an Amazon SQS queue The company wants to
re-architect the application to reduce operational overhead using AWS managed services where possible and remove dependencies on third-party software.
Which solution meets these requirements?

A. Use Amazon ECS containers for the web application and Spot Instances for the Auto Scaling group that processes the SQS queu
B. Replace the custom software with Amazon Recognition to categorize the videos.
C. Store the uploaded videos n Amazon EFS and mount the file system to the EC2 instances for Te web applicatio
D. Process the SOS queue with an AWS Lambda function that calls the Amazon Rekognition API to categorize the videos.
E. Host the web application in Amazon S3. Store the uploaded videos in Amazon S3. Use S3 event notifications to publish events to the SQS queue Process the SQS queue with an AWS Lambda function that calls the Amazon Rekognition API to categorize the videos.
F. Use AWS Elastic Beanstalk to launch EC2 instances in an Auto Scaling group for the web application and launch a worker environment to process the SQS queue Replace the custom software with Amazon Rekognition to categorize the videos.

**Answer:** C

**Explanation:**

Option C is correct because hosting the web application in Amazon S3, storing the uploaded videos in Amazon S3, and using S3 event notifications to publish events to the SQS queue reduces the operational overhead of managing EC2 instances and EBS volumes. Amazon S3 can serve static content such as HTML, CSS, JavaScript, and media files directly from S3 buckets. Amazon S3 can also trigger AWS Lambda functions through S3 event notifications when new objects are created or existing objects are updated or deleted. AWS Lambda can process the SQS queue with an AWS Lambda function that calls the Amazon Rekognition API to categorize the videos. This solution eliminates the need for custom recognition software and third-party dependencies345

References: 1: https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-spot-instances.html 2: https://aws.amazon.com/efs/pricing/ 3: https://docs.aws.amazon.com/AmazonS3/latest/userguide/WebsiteHosting.html 4: https://docs.aws.amazon.com/AmazonS3/latest/userguide/NotificationHowTo.html 5: https://docs.aws.amazon.com/rekognition/latest/dg/what-is.html 6: https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/Welcome.html

## NEW QUESTION 65
- (Exam Topic 1)
A financial company is planning to migrate its web application from on premises to AWS. The company uses a third-party security tool to monitor the inbound traffic to the application. The company has used the security tool for the last 15 years, and the tool has no cloud solutions available from its vendor. The company's security team is concerned about how to integrate the security tool with AWS technology.

The company plans to deploy the application migration to AWS on Amazon EC2 instances. The EC2 instances will run in an Auto Scaling group in a dedicated VPC. The company needs to use the security tool to inspect all packets that come in and out of the VPC. This inspection must occur in real time and must not affect the application's performance. A solutions architect must design a target architecture on AWS that is highly available within an AWS Region.

Which combination of steps should the solutions architect take to meet these requirements? (Select TWO.)

A. Deploy the security tool on EC2 instances in a new Auto Scaling group in the existing VPC.
B. Deploy the web application behind a Network Load Balancer.
C. Deploy an Application Load Balancer in front of the security tool instances.
D. Provision a Gateway Load Balancer for each Availability Zone to redirect the traffic to the security tool.
E. Provision a transit gateway to facilitate communication between VPCs.

**Answer:** AD

**Explanation:**

Option A, Deploy the security tool on EC2 instances in a new Auto Scaling group in the existing VPC, allows the company to use its existing security tool while still running it within the AWS environment. This ensures that all packets coming in and out of the VPC are inspected by the security tool in real time. Option D, Provision a Gateway Load Balancer for each Availability Zone to redirect the traffic to the security tool, allows for high availability within an AWS Region. By provisioning a Gateway Load Balancer for each Availability Zone, the traffic is redirected to the security tool in the event of any failures or outages. This ensures that the security tool is always available to inspect the traffic, even in the event of a failure.

## NEW QUESTION 69
- (Exam Topic 1)
A large company is running a popular web application. The application runs on several Amazon EC2 Linux Instances in an Auto Scaling group in a private subnet. An Application Load Balancer is targeting the Instances In the Auto Scaling group in the private subnet. AWS Systems Manager Session Manager Is configured, and AWS Systems Manager Agent is running on all the EC2 instances.

The company recently released a new version of the application Some EC2 instances are now being marked as unhealthy and are being terminated As a result, the application is running at reduced capacity A solutions architect tries to determine the root cause by analyzing Amazon CloudWatch logs that are collected from the application, but the logs are inconclusive

How should the solutions architect gain access to an EC2 instance to troubleshoot the issue1?

A. Suspend the Auto Scaling group's HealthCheck scaling proces
B. Use Session Manager to log in to an instance that is marked as unhealthy
C. Enable EC2 instance termination protection Use Session Manager to log In to an instance that is marked as unhealthy.
D. Set the termination policy to Oldestinstance on the Auto Scaling grou
E. Use Session Manager to log in to an instance that is marked as unhealthy
F. Suspend the Auto Scaling group's Terminate proces
G. Use Session Manager to log in to an instance thatis marked as unhealthy

**Answer:** D

**Explanation:**

https://docs.aws.amazon.com/autoscaling/ec2/userguide/as-suspend-resume-processes.html

## NEW QUESTION 71
- (Exam Topic 1)
A video streaming company recently launched a mobile app for video sharing. The app uploads various files to an Amazon S3 bucket in the us-east-1 Region. The files range in size from 1 GB to 10 GB.
Users who access the app from Australia have experienced uploads that take long periods of time Sometimes the files fail to completely upload for these users . A solutions architect must improve the app' performance for these uploads
Which solutions will meet these requirements? (Select TWO.)

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

**Answer:** AD

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

**NEW QUESTION 74**
- (Exam Topic 1)
A company needs to implement a patching process for its servers. The on-premises servers and Amazon EC2 instances use a variety of tools to perform patching. Management requires a single report showing the patch status of all the servers and instances.
Which set of actions should a solutions architect take to meet these requirements?

A. Use AWS Systems Manager to manage patches on the on-premises servers and EC2 instance
B. Use Systems Manager to generate patch compliance reports.
C. Use AWS OpsWorks to manage patches on the on-premises servers and EC2 instance
D. Use Amazon QuickSight integration with OpsWorks to generate patch compliance reports.
E. Use an Amazon EventBridge (Amazon CloudWatch Events) rule to apply patches by scheduling an AWS Systems Manager patch remediation jo
F. Use Amazon Inspector to generate patch compliance reports.
G. Use AWS OpsWorks to manage patches on the on-premises servers and EC2 instance
H. Use AWS X-Ray to post the patch status to AWS Systems Manager OpsCenter to generate patch compliance reports.

**Answer:** A

**Explanation:**
https://docs.aws.amazon.com/systems-manager/latest/userguide/systems-manager-patch.html

**NEW QUESTION 78**
- (Exam Topic 1)
A company has hundreds of AWS accounts. The company recently implemented a centralized internal process for purchasing new Reserved Instances and modifying existing Reserved Instances. This process requires all business units that want to purchase or modify Reserved Instances to submit requests to a dedicated team for procurement. Previously, business units directly purchased or modified Reserved Instances in their own respective AWS accounts autonomously.
A solutions architect needs to enforce the new process in the most secure way possible.
Which combination of steps should the solutions architect take to meet these requirements? (Choose two.)

A. Ensure that all AWS accounts are part of an organization in AWS Organizations with all features enabled.
B. Use AWS Config to report on the attachment of an IAM policy that denies access to the ec2:PurchaseReservedInstancesOffering action and the ec2:ModifyReservedInstances action.
C. In each AWS account, create an IAM policy that denies the ec2:PurchaseReservedInstancesOffering action and the ec2:ModifyReservedInstances action.
D. Create an SCP that denies the ec2:PurchaseReservedInstancesOffering action and the ec2:ModifyReservedInstances actio
E. Attach the SCP to each OU of the organization.
F. Ensure that all AWS accounts are part of an organization in AWS Organizations that uses the consolidated billing feature.

**Answer:** AD

**Explanation:**
All features – The default feature set that is available to AWS Organizations. It includes all the functionality of consolidated billing, plus advanced features that give you more control over accounts in your organization. For example, when all features are enabled the management account of the organization has full control over what member accounts can do. The management account can apply SCPs to restrict the services and actions that users (including the root user) and roles in an account can access. https://docs.aws.amazon.com/organizations/latest/userguide/orgs_getting-started_concepts.html#feature-set

**NEW QUESTION 82**
- (Exam Topic 1)
A team collects and routes behavioral data for an entire company The company runs a Multi-AZ VPC
environment with public subnets, private subnets, and in internet gateway Each public subnet also contains a NAT gateway Most of the company's applications read from and write to Amazon Kinesis Data Streams. Most of the workloads am in private subnets.
A solutions architect must review the infrastructure The solutions architect needs to reduce costs and maintain the function of the applications The solutions architect uses Cost Explorer and notices that the cost in the EC2-Other category is consistently high A further review shows that NatGateway-Bytes charges are increasing the cost in the EC2-Other category.
What should the solutions architect do to meet these requirements?

A. Enable VPC Flow Log
B. Use Amazon Athena to analyze the logs for traffic that can be remove
C. Ensure that security groups are Mocking traffic that is responsible for high costs.
D. Add an interface VPC endpoint for Kinesis Data Streams to the VP
E. Ensure that applications have the correct IAM permissions to use the interface VPC endpoint.
F. Enable VPC Flow Logs and Amazon Detective Review Detective findings for traffic that is not related to Kinesis Data Streams Configure security groups to block that traffic
G. Add an interface VPC endpoint for Kinesis Data Streams to the VP
H. Ensure that the VPC endpoint policy allows traffic from the applications.

**Answer:** D

**Explanation:**
https://docs.aws.amazon.com/vpc/latest/privatelink/vpc-endpoints-access.html https://aws.amazon.com/premiumsupport/knowledge-center/vpc-reduce-nat-gateway-transfer-costs/
VPC endpoint policies enable you to control access by either attaching a policy to a VPC endpoint or by using additional fields in a policy that is attached to an IAM user, group, or role to restrict access to only occur via the specified VPC endpoint

**NEW QUESTION 85**
- (Exam Topic 1)
A company has many AWS accounts and uses AWS Organizations to manage all of them. A solutions architect must implement a solution that the company can use to share a common network across multiple accounts.
The company's infrastructure team has a dedicated infrastructure account that has a VPC. The infrastructure team must use this account to manage the network. Individual accounts cannot have the ability to manage their own networks. However, individual accounts must be able to create AWS resources within subnets.
Which combination of actions should the solutions architect perform to meet these requirements? (Select TWO.)

A. Create a transit gateway in the infrastructure account.
B. Enable resource sharing from the AWS Organizations management account.
C. Create VPCs in each AWS account within the organization in AWS Organization
D. Configure the VPCs to share the same CIDR range and subnets as the VPC in the infrastructure accoun
E. Peer the VPCs in each individual account with the VPC in the infrastructure account,
F. Create a resource share in AWS Resource Access Manager in the infrastructure accoun
G. Select the specific AWS Organizations OU that will use the shared networ
H. Select each subnet to associate with the resource share.
I. Create a resource share in AWS Resource Access Manager in the infrastructure accoun
J. Select the specific AWS Organizations OU that will use the shared networ
K. Select each prefix list to associate with the resource share.

**Answer:** AE

**Explanation:**
https://docs.aws.amazon.com/vpc/latest/userguide/sharing-managed-prefix-lists.html

**NEW QUESTION 89**
- (Exam Topic 1)
A solutions architect must analyze a company's Amazon EC2 Instances and Amazon Elastic Block Store (Amazon EBS) volumes to determine whether the company is using resources efficiently The company is running several large, high-memory EC2 instances lo host database dusters that are deployed in active/passive configurations The utilization of these EC2 instances varies by the applications that use the databases, and the company has not identified a pattern The solutions architect must analyze the environment and take action based on the findings. Which solution meets these requirements MOST cost-effectively?

A. Create a dashboard by using AWS Systems Manager OpsConter Configure visualizations tor Amazon CloudWatch metrics that are associated with the EC2 instances and their EBS volumes Review thedashboard periodically and identify usage patterns Right size the EC2 instances based on the peaks in the metrics
B. Turn on Amazon CloudWatch detailed monitoring for the EC2 instances and their EBS volumes Create and review a dashboard that is based on the metrics Identify usage patterns Right size the FC? instances based on the peaks In the metrics
C. Install the Amazon CloudWatch agent on each of the EC2 Instances Turn on AWS Compute Optimizer, and let it run for at least 12 hours Review the recommendations from Compute Optimizer, and right size the EC2 instances as directed
D. Sign up for the AWS Enterprise Support plan Turn on AWS Trusted Advisor Wait 12 hours Review the recommendations from Trusted Advisor, and rightsize the EC2 instances as directed

**Answer:** C

**Explanation:**
(https://aws.amazon.com/compute-optimizer/pricing/ , https://aws.amazon.com/systems-manager/pricing/ ). https://aws.amazon.com/compute-optimizer/

**NEW QUESTION 92**
- (Exam Topic 1)
A delivery company needs to migrate its third-party route planning application to AWS. The third party supplies a supported Docker image from a public registry. The image can run in as many containers as required to generate the route map.
The company has divided the delivery area into sections with supply hubs so that delivery drivers travel the shortest distance possible from the hubs to the customers. To reduce the time necessary to generate route maps, each section uses its own set of Docker containers with a custom configuration that processes orders only in the section's area.
The company needs the ability to allocate resources cost-effectively based on the number of running containers.
Which solution will meet these requirements with the LEAST operational overhead?

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

**Answer:** D

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

**NEW QUESTION 94**
- (Exam Topic 1)

A company wants to deploy an AWS WAF solution to manage AWS WAF rules across multiple AWS accounts. The accounts are managed under different OUs in AWS Organizations.
Administrators must be able to add or remove accounts or OUs from managed AWS WAF rule sets as needed Administrators also must have the ability to automatically update and remediate noncompliant AWS WAF rules in all accounts
Which solution meets these requirements with the LEAST amount of operational overhead?

A. Use AWS Firewall Manager to manage AWS WAF rules across accounts in the organizatio
B. Use an AWS Systems Manager Parameter Store parameter to store account numbers and OUs to manage Update the parameter as needed to add or remove accounts or OUs Use an Amazon EventBridge (Amazon CloudWatch Events) rule to identify any changes to the parameter and to invoke an AWS Lambda function to update the security policy in the Firewall Manager administrative account
C. Deploy an organization-wide AWS Config rule that requires all resources in the selected OUs to associate the AWS WAF rule
D. Deploy automated remediation actions by using AWS Lambda to fix noncompliant resources Deploy AWS WAF rules by using an AWS CloudFormation stack set to target the same OUs where the AWS Config rule is applied.
E. Create AWS WAF rules in the management account of the organization Use AWS Lambda environment variables to store account numbers and OUs to manage Update environment variables as needed to add or remove accounts or OUs Create cross-account IAM roles in member accounts Assume the roles by using AWS Security Token Service (AWS STS) in the Lambda function to create and update AWS WAF rules in the member accounts.
F. Use AWS Control Tower to manage AWS WAF rules across accounts in the organization Use AWS Key Management Service (AWS KMS) to store account numbers and OUs to manage Update AWS KMS as needed to add or remove accounts or OUs Create IAM users in member accounts Allow AWS Control Tower in the management account to use the access key and secret access key to create and update AWS WAF rules in the member accounts

**Answer:** A

**Explanation:**
https://aws.amazon.com/solutions/implementations/automations-for-aws-firewall-manager/
In this solution, AWS Firewall Manager is used to manage AWS WAF rules across accounts in the organization. An AWS Systems Manager Parameter Store parameter is used to store account numbers and OUs to manage. This parameter can be updated as needed to add or remove accounts or OUs. An Amazon EventBridge rule is used to identify any changes to the parameter and to invoke an AWS Lambda function to update the security policy in the Firewall Manager administrative account. This solution allows for easy management of AWS WAF rules across multiple accounts with minimal operational overhead

**NEW QUESTION 99**
- (Exam Topic 1)
A company runs a new application as a static website in Amazon S3. The company has deployed the application to a production AWS account and uses Amazon CloudFront to deliver the website. The website calls an Amazon API Gateway REST API. An AWS Lambda function backs each API method.
The company wants to create a CSV report every 2 weeks to show each API Lambda function's recommended configured memory, recommended cost, and the price difference between current configurations and the recommendations. The company will store the reports in an S3 bucket.
Which solution will meet these requirements with the LEAST development time?

A. Create a Lambda function that extracts metrics data for each API Lambda function from Amazon CloudWatch Logs for the 2-week penod_ Collate the data into tabular forma
B. Store the data as a_csvfile in an S3 bucke
C. Create an Amazon Eventaridge rule to schedule the Lambda function to run every 2 weeks.
D. Opt in to AWS Compute Optimize
E. Create a Lambda function that calls the ExportLambdaFunctionRecommendatlons operatio
F. Export the _csv file to an S3 bucke
G. Create an Amazon Eventaridge rule to schedule the Lambda function to run every 2 weeks.
H. Opt in to AWS Compute Optimize
I. Set up enhanced infrastructure metric
J. Within the Compute Optimizer console, schedule a job to export the Lambda recommendations to a _csvfile_ Store the file in an S3 bucket every 2 weeks.
K. Purchase the AWS Business Support plan for the production accoun
L. Opt in to AWS Compute Optimizer for AWS Trusted Advisor check
M. In the Trusted Advisor console, schedule a job to export the cost optimization checks to a _csvfile_ Store the file in an S3 bucket every 2 weeks.

**Answer:** B

**Explanation:**
https://docs.aws.amazon.com/compute-optimizer/latest/APIReference/API_ExportLambdaFunctionRecommend

**NEW QUESTION 100**
- (Exam Topic 1)
A solutions architect is auditing the security setup of an AWS Lambda function for a company. The Lambda function retrieves the latest changes from an Amazon Aurora database. The Lambda function and the database run in the same VPC. Lambda environment variables are providing the database credentials to the Lambda function.
The Lambda function aggregates data and makes the data available in an Amazon S3 bucket that is configured for server-side encryption with AWS KMS managed encryption keys (SSE-KMS). The data must not travel across the internet. If any database credentials become compromised, the company needs a solution that minimizes the impact of the compromise.
What should the solutions architect recommend to meet these requirements?

A. Enable IAM database authentication on the Aurora DB cluste
B. Change the IAM role for the Lambda function to allow the function to access the database by using IAM database authenticatio
C. Deploy a gateway VPC endpoint for Amazon S3 in the VPC.
D. Enable IAM database authentication on the Aurora DB cluste
E. Change the IAM role for the Lambda function to allow the function to access the database by using IAM database authenticatio
F. Enforce HTTPS on the connection to Amazon S3 during data transfers.
G. Save the database credentials in AWS Systems Manager Parameter Stor
H. Set up password rotation on the credentials in Parameter Stor
I. Change the IAM role for the Lambda function to allow the function to access Parameter Stor
J. Modify the Lambda function to retrieve the credentials from Parameter Stor
K. Deploy a gateway VPC endpoint for Amazon S3 in the VPC.
L. Save the database credentials in AWS Secrets Manage
M. Set up password rotation on the credentials in Secrets Manage
N. Change the IAM role for the Lambda function to allow the function to access Secrets Manage

O. Modify the Lambda function to retrieve the credentials Om Secrets Manage
P. Enforce HTTPS on the connection to Amazon S3 during data transfers.

**Answer:** A

**Explanation:**
https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/UsingWithRDS.IAMDBAuth.html


**NEW QUESTION 103**
- (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 107**
- (Exam Topic 1)
A solutions architect is investigating an issue in which a company cannot establish new sessions in Amazon Workspaces. An initial analysis indicates that the
issue involves user profiles. The Amazon Workspaces environment is configured to use Amazon FSx for Windows File Server as the profile share storage. The
FSx for Windows File Server file system is configured with 10 TB of storage.
The solutions architect discovers that the file system has reached its maximum capacity. The solutions architect must ensure that users can regain access. The
solution also must prevent the problem from occurring again.
Which solution will meet these requirements?

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

**Answer:** B

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


**NEW QUESTION 108**
- (Exam Topic 2)
A company provides auction services for artwork and has users across North America and Europe. The company hosts its application in Amazon EC2 instances in
the us-east-1 Region. Artists upload photos of their work as large-size, high-resolution image files from their mobile phones to a centralized Amazon S3 bucket
created in the us-east-I Region. The users in Europe are reporting slow performance for their Image uploads.
How can a solutions architect improve the performance of the image upload process?

A. Redeploy the application to use S3 multipart uploads.
B. Create an Amazon CloudFront distribution and point to the application as a custom origin
C. Configure the buckets to use S3 Transfer Acceleration.
D. Create an Auto Scaling group for the EC2 instances and create a scaling policy.

**Answer:** C

**Explanation:**
Transfer acceleration. S3 Transfer Acceleration utilizes the Amazon CloudFront global network of edge
locations to accelerate the transfer of data to and from S3 buckets. By enabling S3 Transfer Acceleration on the centralized S3 bucket, the users in Europe will
experience faster uploads as their data will be routed through the closest CloudFront edge location.

**NEW QUESTION 111**
- (Exam Topic 2)
A company is running a compute workload by using Amazon EC2 Spot Instances that are in an Auto Scaling group. The launch template uses two placement
groups and a single instance type.
Recently, a monitoring system reported Auto Scaling instance launch failures that correlated with longer wait times for system users. The company needs to
improve the overall reliability of the workload.
Which solution will meet this requirement?

A. Replace the launch template with a launch configuration to use an Auto Scaling group that uses attribute-based instance type selection.
B. Create a new launch template version that uses attribute-based instance type selectio
C. Configure the Auto Scaling group to use the new launch template version.
D. Update the launch template Auto Scaling group to increase the number of placement groups.
E. Update the launch template to use a larger instance type.

**Answer:** B

**Explanation:**
https://docs.aws.amazon.com/autoscaling/ec2/userguide/create-asg-instance-type-requirements.html#use-attribut

**NEW QUESTION 114**
- (Exam Topic 2)
A company processes environment data. The has a set up sensors to provide a continuous stream of data from different areas in a city. The data is available in
JSON format.
The company wants to use an AWS solution to send the data to a database that does not require fixed schemas for storage. The data must be send in real time.
Which solution will meet these requirements?

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

**Answer:** B

**Explanation:**
Amazon Kinesis Data Streams is a service that enables real-time data ingestion and processing. Amazon DynamoDB is a NoSQL database that does not require
fixed schemas for storage. By using Kinesis Data Streams and DynamoDB, the company can send the JSON data to a database that can handle schemaless data
in real time. References:
> https://docs.aws.amazon.com/streams/latest/dev/introduction.html
> https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Introduction.html

**NEW QUESTION 117**
- (Exam Topic 2)
A solutions architect is designing a solution to process events. The solution must have the ability to scale in and out based on the number of events that the
solution receives. If a processing error occurs, the event must move into a separate queue for review.
Which solution will meet these requirements?

A. Send event details to an Amazon Simple Notification Service (Amazon SNS) topi
B. Configure an AWS Lambda function as a subscriber to the SNS topic to process the event
C. Add an on-failure destination to the functio
D. Set an Amazon Simple Queue Service (Amazon SQS) queue as the target.
E. Publish events to an Amazon Simple Queue Service (Amazon SQS) queu
F. Create an Amazon EC2 Auto Scaling grou
G. Configure the Auto Scaling group to scale in and out based on the ApproximateAgeOfOldestMessage metric of the queu
H. Configure the application to write failed messages to a dead-letter queue.
I. Write events to an Amazon DynamoDB tabl
J. Configure a DynamoDB stream for the tabl
K. Configure the stream to invoke an AWS Lambda functio
L. Configure the Lambda function to process the events.
M. Publish events to an Amazon EventBridge event bu
N. Create and run an application on an Amazon EC2 instance with an Auto Scaling group that isbehind an Application Load Balancer (ALB). Set the ALB as the
event bus targe
O. Configure the event bus to retry event
P. Write messages to a dead-letter queue if the application cannot process the messages.

**Answer:** A

**Explanation:**
Amazon Simple Notification Service (Amazon SNS) is a fully managed pub/sub messaging service that enables users to send messages to multiple subscribers1.
Users can send event details to an Amazon SNS topic and configure an AWS Lambda function as a subscriber to the SNS topic to process the events. Lambda is
a serverless compute service that runs code in response to events and automatically manages the underlying compute resources2. Users can add an on-failure
destination to the function and set an Amazon Simple Queue
Service (Amazon SQS) queue as the target. Amazon SQS is a fully managed message queuing service that enables users to decouple and scale microservices,
distributed systems, and serverless applications3. This way, if a processing error occurs, the event will move into the separate queue for review.
Option B is incorrect because publishing events to an Amazon SQS queue and creating an Amazon EC2 Auto Scaling group will not have the ability to scale in and

out based on the number of events that the solution receives. Amazon EC2 is a web service that provides secure, resizable compute capacity in the cloud. Auto Scaling is a feature that helps users maintain application availability and allows them to scale their EC2 capacity up or down automatically according to conditions they define. However, for this use case, using SQS and EC2 will not take advantage of the serverless capabilities of Lambda and SNS.

Option C is incorrect because writing events to an Amazon DynamoDB table and configuring a DynamoDB stream for the table will not have the ability to move events into a separate queue for review if a processing error occurs. Amazon DynamoDB is a fully managed key-value and document database that delivers single-digit millisecond performance at any scale. DynamoDB Streams is a feature that captures data modification events in DynamoDB tables. Users can configure the stream to invoke a Lambda function, but they cannot configure an on-failure destination for the function.

Option D is incorrect because publishing events to an Amazon EventBridge event bus and setting an Application Load Balancer (ALB) as the event bus target will not have the ability to move events into a separate queue for review if a processing error occurs. Amazon EventBridge is a serverless event bus service that makes it easy to connect applications with data from a variety of sources. An ALB is a load balancer that distributes incoming application traffic across multiple targets, such as EC2 instances, containers, IP addresses, Lambda functions, and virtual appliances. Users can configure EventBridge to retry events, but they cannot configure an on-failure destination for the ALB.

## NEW QUESTION 122
- (Exam Topic 2)
A company runs its sales reporting application in an AWS Region in the United States. The application uses an Amazon API Gateway Regional API and AWS Lambda functions to generate on-demand reports from data in an Amazon RDS for MySQL database. The frontend of the application is hosted on Amazon S3 and is accessed by users through an Amazon CloudFront distribution. The company is using Amazon Route 53 as the DNS service for the domain. Route 53 is configured with a simple routing policy to route traffic to the API Gateway API.

In the next 6 months, the company plans to expand operations to Europe. More than 90% of the database traffic is read-only traffic. The company has already deployed an API Gateway API and Lambda functions in the new Region.

A solutions architect must design a solution that minimizes latency for users who download reports. Which solution will meet these requirements?

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

**Answer:** C

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

> Using AWS Database Migration Service (AWS DMS) to replicate the primary database in the original Region to the database in the new Region would not be as cost-effective or simple as using a cross-Region read replica. AWS DMS is a service that enables you to migrate relational databases, data warehouses, NoSQL databases, and other types of data stores. You can use AWS DMS to perform one-time migrations or continuous data replication with high availability and consolidate databases into a petabyte-scale data warehouse3. However, AWS DMS requires more configuration and management than creating a cross-Region read replica, which is fully managed by Amazon RDS. AWS DMS also incurs additional charges for replication instances and tasks.

> Creating an Amazon API Gateway Data API service integration with Amazon Redshift would not help with disaster recovery or minimizing latency. The Data API is a feature that enables you to query your Amazon Redshift cluster using HTTP requests, without needing a persistent connection or a SQL client. It is useful for building applications that interact with Amazon Redshift, but not for replicating or recovering data from an RDS database.

> Creating an AWS Data Exchange datashare by connecting AWS Data Exchange to the Redshift cluster would not help with disaster recovery or minimizing latency. AWS Data Exchange is a service that makes it easy for AWS customers to exchange data in the cloud. You can use AWS Data Exchange to subscribe to a diverse selection of third-party data products or offer your own data products to other AWS customers. A datashare is a feature that enables you to share live and secure access to your Amazon Redshift data across your accounts or with third parties without copying or moving the underlying data. It is useful for sharing query results and views with other users, but not for replicating or recovering data from an RDS database.
References:
> https://aws.amazon.com/dms/
> https://docs.aws.amazon.com/redshift/latest/mgmt/data-api.html
> https://aws.amazon.com/data-exchange/
> https://docs.aws.amazon.com/redshift/latest/dg/datashare-overview.html

## NEW QUESTION 125
- (Exam Topic 2)
A company wants to containerize a multi-tier web application and move the application from an on-premises data center to AWS. The application includes web. application, and database tiers. The company needs to make the application fault tolerant and scalable. Some frequently accessed data must always be available across application servers. Frontend web servers need session persistence and must scale to meet increases in traffic.
Which solution will meet these requirements with the LEAST ongoing operational overhead?

A. Run the application on Amazon Elastic Container Service (Amazon ECS) on AWS Fargat
B. Use Amazon Elastic File System (Amazon EFS) for data that is frequently accessed between the web and application tier
C. Store the frontend web server session data in Amazon Simple Queue Service (Amazon SOS).

D. Run the application on Amazon Elastic Container Service (Amazon ECS) on Amazon EC2. Use Amazon ElastiCache for Redis to cache frontend web server session dat
E. Use Amazon Elastic Block Store (Amazon EBS) with Multi-Attach on EC2 instances that are distributed across multiple Availability Zones.
F. Run the application on Amazon Elastic Kubernetes Service (Amazon EKS). Configure Amazon EKS to use managed node group
G. Use ReplicaSets to run the web servers and application
H. Create an Amazon Elastic File System (Amazon EFS) Me syste
I. Mount the EFS file system across all EKS pods to store frontend web server session data.
J. Deploy the application on Amazon Elastic Kubernetes Service (Amazon EKS) Configure Amazon EKS to use managed node group
K. Run the web servers and application as Kubernetes deployments in the EKS cluste
L. Store the frontend web server session data in an Amazon DynamoDB tabl
M. Create an Amazon Elastic File System (Amazon EFS) volume that all applications will mount at the time of deployment.

**Answer:** D

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

≫ https://docs.aws.amazon.com/eks/latest/userguide/managed-node-groups.html

≫ https://docs.aws.amazon.com/eks/latest/userguide/deployments.html

≫ https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Introduction.html

≫ https://docs.aws.amazon.com/efs/latest/ug/mounting-fs.html


**NEW QUESTION 126**
- (Exam Topic 2)
A company is migrating a document processing workload to AWS. The company has updated many applications to natively use the Amazon S3 API to store, retrieve, and modify documents that a processing server generates at a rate of approximately 5 documents every second. After the document processing is finished, customers can download the documents directly from Amazon S3.
During the migration, the company discovered that it could not immediately update the processing server that generates many documents to support the S3 API. The server runs on Linux and requires fast local access to the files that the server generates and modifies. When the server finishes processing, the files must be available to the public for download within 30 minutes.
Which solution will meet these requirements with the LEAST amount of effort?

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

**Answer:** C

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


**NEW QUESTION 128**
- (Exam Topic 2)
A company operates an on-premises software-as-a-service (SaaS) solution that ingests several files daily. The company provides multiple public SFTP endpoints to its customers to facilitate the file transfers. The customers add the SFTP endpoint IP addresses to their firewall allow list for outbound traffic. Changes to the SFTP endmost IP addresses are not permitted.
The company wants to migrate the SaaS solution to AWS and decrease the operational overhead of the file transfer service.
Which solution meets these requirements?

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

**Answer:** A

**Explanation:**
Bring your own IP addresses (BYOIP) You can bring part or all of your publicly routable IPv4 or IPv6 address range from your on-premises network to your AWS account. You continue to own the address range, but AWS advertises it on the internet by default. After you bring the address range to AWS, it appears in your

AWS account as an address pool. https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-byoip.html AWS Transfer for SFTP enables you to easily move your file transfer workloads that use the Secure Shell File Transfer Protocol (SFTP) to AWS without needing to modify your applications or manage any SFTP servers. https://aws.amazon.com/about-aws/whats-new/2018/11/aws-transfer-for-sftp-fully-managed-sftp-for-s3/

**NEW QUESTION 132**
- (Exam Topic 2)
A company runs an application on a fleet of Amazon EC2 instances that are in private subnets behind an internet-facing Application Load Balancer (ALB). The ALB is the origin for an Amazon CloudFront distribution. An AWS WAF web ACL that contains various AWS managed rules is associated with the CloudFront distribution.
The company needs a solution that will prevent internet traffic from directly accessing the ALB. Which solution will meet these requirements with the LEAST operational overhead?

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

**Answer:** C

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

**NEW QUESTION 136**
- (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 138**
- (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 140**
- (Exam Topic 2)

A company is building a hybrid environment that includes servers in an on-premises data center and in the AWS Cloud. The company has deployed Amazon EC2 instances in three VPCs. Each VPC is in a different AWS Region. The company has established an AWS Direct Connect connection to the data center from the Region that is closest to the data center.

The company needs the servers in the on-premises data center to have access to the EC2 instances in all three VPCs. The servers in the on-premises data center also must have access to AWS public services.

Which combination of steps will meet these requirements with the LEAST cost? (Select TWO.)

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

**Answer:** AE

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


**NEW QUESTION 145**
- (Exam Topic 2)
A company is designing an AWS Organizations structure. The company wants to standardize a process to apply tags across the entire organization. The company will require tags with specific values when a user creates a new resource. Each of the company's OUs will have unique tag values.
Which solution will meet these requirements?

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

**Answer:** A

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


**NEW QUESTION 146**
- (Exam Topic 2)
A company uses a Grafana data visualization solution that runs on a single Amazon EC2 instance to monitor the health of the company's AWS workloads. The company has invested time and effort to create dashboards that the company wants to preserve. The dashboards need to be highly available and cannot be down for longer than 10 minutes. The company needs to minimize ongoing maintenance.
Which solution will meet these requirements with the LEAST operational overhead?

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

**Answer:** C

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


**NEW QUESTION 147**
- (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 150**
......

# Thank You for Trying Our Product

## We offer two products:

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

2nd - Questons and Answers in PDF Format

## SAP-C02 Practice Exam Features:

* SAP-C02 Questions and Answers Updated Frequently

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

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

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

## 100% Actual & Verified — Instant Download, Please Click
Order The SAP-C02 Practice Test Here