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)
A company wants to migrate its data analytics environment from on premises to AWS The environment consists of two simple Node js applications One of the applications collects sensor data and loads it into a MySQL database The other application aggregates the data into reports When the aggregation jobs run. some of the load jobs fail to run correctly
The company must resolve the data loading issue The company also needs the migration to occur without interruptions or changes for the company's customers What should a solutions architect do to meet these requirements?

A. Set up an Amazon Aurora MySQL database as a replication target for the on-premises database Create an Aurora Replica for the Aurora MySQL database, and move the aggregation jobs to run against the Aurora Replica Set up collection endpomts as AWS Lambda functions behind a Network Load Balancer (NLB). and use Amazon RDS Proxy to wnte to the Aurora MySQL database When the databases are synced disable the replication job and restart the Aurora Replica as the primary instanc
B. Point the collector DNS record to the NLB.
C. Set up an Amazon Aurora MySQL database Use AWS Database Migration Service (AWS DMS) to perform continuous data replication from the on-premises database to Aurora Move the aggregation jobs to run against the Aurora MySQL database Set up collection endpomts behind an Application Load Balancer (ALB) as Amazon EC2 instances in an Auto Scaling group When the databases are synced, point the collector DNS record to the ALB Disable the AWS DMS sync task after the cutover from on premises to AWS
D. Set up an Amazon Aurora MySQL database Use AWS Database Migration Service (AWS DMS) to perform continuous data replication from the on-premises database to Aurora Create an Aurora Replica for the Aurora MySQL database and move the aggregation jobs to run against the Aurora Replica Set up collection endpoints as AWS Lambda functions behind an Application Load Balancer (ALB) and use Amazon RDS Proxy to write to the Aurora MySQL database When the databases are synced, point the collector DNS record to the ALB Disable the AWS DMS sync task after the cutover from on premises to AWS
E. Set up an Amazon Aurora MySQL database Create an Aurora Replica for the Aurora MySQL database and move the aggregation jobs to run against the Aurora Replica Set up collection endpoints as an Amazon Kinesis data stream Use Amazon Kinesis Data Firehose to replicate the data to the Aurora MySQL database When the databases are synced disable the replication job and restart the Aurora Replica as the primary instance Point the collector DNS record to the Kinesis data stream.

Answer: C

Explanation:
Set up an Amazon Aurora MySQL database. Use AWS Database Migration Service (AWS DMS) to perform continuous data replication from the on-premises database to Aurora. Create an Aurora Replica for the Aurora MySQL database, and move the aggregation jobs to run against the Aurora Replica. Set up collection endpoints as AWS Lambda functions behind an Application Load Balancer (ALB), and use Amazon RDS Proxy to write to the Aurora MySQL database. When the databases are synced, point the collector DNS record to the ALB. Disable the AWS DMS sync task after the cutover from on premises to AWS.
Amazon RDS Proxy allows applications to pool and share connections established with the database, improving database efficiency and application scalability.
With RDS Proxy, failover times for Aurora and RDS databases are reduced by up to 66%

NEW QUESTION 3
- (Exam Topic 1)
A company wants to change its internal cloud billing strategy for each of its business units. Currently, the cloud governance team shares reports for overall cloud spending with the head of each business unit. The company uses AWS Organizations lo manage the separate AWS accounts for each business unit. The existing tagging standard in Organizations includes the application, environment, and owner. The cloud governance team wants a centralized solution so each business

unit receives monthly reports on its cloud spending. The solution should also send notifications for any cloud spending that exceeds a set threshold.
Which solution is the MOST cost-effective way to meet these requirements?

A. Configure AWS Budgets in each account and configure budget alerts that are grouped by application, environment, and owne
B. Add each business unit to an Amazon SNS topic for each aler
C. Use Cost Explorer in each account to create monthly reports for each business unit.
D. Configure AWS Budgets in the organization's master account and configure budget alerts that are grouped by application, environment, and owne
E. Add each business unit to an Amazon SNS topic for each aler
F. Use Cost Explorer in the organization's master account to create monthly reports for each business unit.
G. Configure AWS Budgets in each account and configure budget alerts lhat are grouped by application, environment, and owne
H. Add each business unit to an Amazon SNS topic for each aler
I. Use the AWS Billing and Cost Management dashboard in each account to create monthly reports for each businessunit.
J. Enable AWS Cost and Usage Reports in the organization's master account and configure reports grouped by application, environment, and owne
K. Create an AWS Lambda function that processes AWS Cost and Usage Reports, sends budget alerts, and sends monthly reports to each business unit's email list.

Answer: B

Explanation:
Configure AWS Budgets in the organization€™s master account and configure budget alerts that are grouped by application, environment, and owner. Add each business unit to an Amazon SNS topic for each alert. Use Cost Explorer in the organization€™s master account to create monthly reports for each business unit. https://aws.amazon.com/about-aws/whats-new/2019/07/introducing-aws-budgets-reports/#:~:text=AWS%20Bud

NEW QUESTION 4
- (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 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 is running an application in the AWS Cloud. The company's security team must approve the creation of all new IAM users. When a new IAM user is created, all access for the user must be removed automatically. The security team must then receive a notification to approve the user. The company has a multi-Region AWS CloudTrail trail In the AWS account.

Which combination of steps will meet these requirements? (Select THREE.)

A. Create an Amazon EventBridge (Amazon CloudWatch Events) rul
B. Define a pattern with the detail-type value set to AWS API Call via CloudTrail and an eventName of CreateUser.
C. Configure CloudTrail to send a notification for the CreateUser event to an Amazon Simple Notification Service (Amazon SNS) topic.
D. Invoke a container that runs in Amazon Elastic Container Service (Amazon ECS) with AWS Fargate technology to remove access
E. Invoke an AWS Step Functions state machine to remove access.
F. Use Amazon Simple Notification Service (Amazon SNS) to notify the security team.
G. Use Amazon Pinpoint to notify the security team.

Answer: ADE

Explanation:
https://docs.aws.amazon.com/prescriptive-guidance/latest/patterns/send-a-notification-when-an-iam-user-is-crea

NEW QUESTION 7
- (Exam Topic 1)
A software as a service (SaaS) based company provides a case management solution to customers A3 part of the solution. The company uses a standalone Simple Mail Transfer Protocol (SMTP) server to send email messages from an application. The application also stores an email template for acknowledgement email messages that populate customer data before the application sends the email message to the customer.
The company plans to migrate this messaging functionality to the AWS Cloud and needs to minimize operational overhead.
Which solution will meet these requirements MOST cost-effectively?

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

Answer: D

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

NEW QUESTION 8
- (Exam Topic 1)
A company has 10 accounts that are part of an organization in AWS Organizations AWS Config is configured in each account All accounts belong to either the Prod OU or the NonProd OU
The company has set up an Amazon EventBridge rule in each AWS account to notify an Amazon Simple Notification Service (Amazon SNS) topic when an Amazon EC2 security group inbound rule is created with 0.0.0.0/0 as the source The company's security team is subscribed to the SNS topic
For all accounts in the NonProd OU the security team needs to remove the ability to create a security group inbound rule that includes 0.0.0.0/0 as the source
Which solution will meet this requirement with the LEAST operational overhead?

A. Modify the EventBridge rule to invoke an AWS Lambda function to remove the security group inbound rule and to publish to the SNS topic Deploy the updated rule to the NonProd OU
B. Add the vpc-sg-open-only-to-authorized-ports AWS Config managed rule to the NonProd OU
C. Configure an SCP to allow the ec2 AulhonzeSecurityGroupIngress action when the value of the aws Sourcelp condition key is not 0.0.0.0/0 Apply the SCP to the NonProd OU
D. Configure an SCP to deny the ec2 AuthorizeSecurityGroupIngress action when the value of the aws Sourcelp condition key is 0.0.0.0/0 Apply the SCP to the NonProd OU

Answer: D

Explanation:
This solution will meet the requirement with the least operational overhead because it directly denies the creation of the security group inbound rule with 0.0.0.0/0 as the source, which is the exact requirement. Additionally, it does not require any additional steps or resources such as invoking a Lambda function or adding a Config rule.
An SCP (Service Control Policy) is a policy that you can use to set fine-grained permissions for your AWS
accounts within your organization. You can use SCPs to set permissions for the root user of an account and to delegate permissions to IAM users and roles in the accounts. You can use SCPs to set permissions that allow or deny access to specific services, actions, and resources.
To implement this solution, you would need to create an SCP that denies the ec2:AuthorizeSecurityGroupIngress action when the value of the aws:Sourcelp condition key is 0.0.0.0/0. This SCP would then be applied to the NonProd OU. This would ensure that any security group inbound rule that includes 0.0.0.0/0 as the source will be denied, thus meeting the requirement.
Reference: https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scp.html
https://docs.aws.amazon.com/IAM/latest/UserGuide/access_policies_condition-keys.html

NEW QUESTION 9
- (Exam Topic 1)
A software company hosts an application on AWS with resources in multiple AWS accounts and Regions. The application runs on a group of Amazon EC2 instances in an application VPC located in the us-east-1 Region with an IPv4 CIDR block of 10.10.0.0/16. In a different AWS account, a shared services VPC is located in the us-east-2 Region with an IPv4 CIDR block of 10.10.10.0/24. When a cloud engineer uses AWS CloudFormation to attempt to peer the application VPC with the shared services VPC, an error message indicates a peering failure. Which factors could cause this error? (Choose two.)

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

Answer: AE

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

NEW QUESTION 10
- (Exam Topic 1)
A company wants to migrate its workloads from on premises to AWS. The workloads run on Linux and Windows. The company has a large on-premises intra structure that consists of physical machines and VMs that host numerous applications.
The company must capture details about the system configuration. system performance. running processure and network coi.net lions of its o. -premises ,on boards. The company also must divide the on-premises applications into groups for AWS migrations. The company needs recommendations for Amazon EC2 instance types so that the company can run its workloads on AWS in the most cost-effective manner.
Which combination of steps should a solutions architect take to meet these requirements? (Select THREE.)

A. Assess the existing applications by installing AWS Application Discovery Agent on the physical machines and VMs.
B. Assess the existing applications by installing AWS Systems Manager Agent on the physical machines and VMs
C. Group servers into applications for migration by using AWS Systems Manager Application Manager.
D. Group servers into applications for migration by using AWS Migration Hub.
E. Generate recommended instance types and associated costs by using AWS Migration Hub.
F. Import data about server sizes into AWS Trusted Adviso
G. Follow the recommendations for cost optimization.

Answer: ADE

Explanation:
https://docs.aws.amazon.com/application-discovery/latest/userguide/discovery-agent.html
https://docs.aws.amazon.com/migrationhub/latest/ug/ec2-recommendations.html

NEW QUESTION 10
- (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 15
- (Exam Topic 1)
A company has 50 AWS accounts that are members of an organization in AWS Organizations Each account contains multiple VPCs The company wants to use AWS Transit Gateway to establish connectivity between the VPCs in each member account Each time a new member account is created, the company wants to automate the process of creating a new VPC and a transit gateway attachment.
Which combination of steps will meet these requirements? (Select TWO)

A. From the management account, share the transit gateway with member accounts by using AWS Resource Access Manager
B. Prom the management account, share the transit gateway with member accounts by using an AWS Organizations SCP
C. Launch an AWS CloudFormation stack set from the management account that automatical^/ creates a new VPC and a VPC transit gateway attachment in a member accoun
D. Associate the attachment with the transit gateway in the management account by using the transit gateway ID.

E. Launch an AWS CloudFormation stack set from the management account that automatical^ creates a new VPC and a peering transit gateway attachment in a member accoun
F. Share the attachment with the transit gateway in the management account by using a transit gateway service-linked role.
G. From the management account, share the transit gateway with member accounts by using AWS Service Catalog

Answer: AC

Explanation:
https://aws.amazon.com/blogs/mt/self-service-vpcs-in-aws-control-tower-using-aws-service-catalog/ https://docs.aws.amazon.com/vpc/latest/tgw/tgw-transit-gateways.html
https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-resource-ec2-transitgatewayattachme

NEW QUESTION 18
- (Exam Topic 1)
A company is building a solution in the AWS Cloud. Thousands or devices will connect to the solution and send data. Each device needs to be able to send and receive data in real time over the MQTT protocol. Each device must authenticate by using a unique X.509 certificate.
Which solution will meet these requirements with the LEAST operational overhead?

A. Set up AWS loT Cor
B. For each device, create a corresponding Amazon MQ queue and provision a certificat
C. Connect each device to Amazon MQ.
D. Create a Network Load Balancer (NLB) and configure it with an AWS Lambda authorize
E. Run an MQTT broker on Amazon EC2 instances in an Auto Scaling grou
F. Set the Auto Scaling group as the target for the NL
G. Connect each device to the NLB.
H. Set up AWS loT Cor
I. For each device, create a corresponding AWS loT thing and provision a certificat
J. Connect each device to AWS loT Core.
K. Set up an Amazon API Gateway HTTP API and a Network Load Balancer (NLB). Create integration between API Gateway and the NL
L. Configure a mutual TLS certificate authorizer on the HTTP AP
M. Run an MQTT broker on an Amazon EC2 instance that the NLB target
N. Connect each device to the NLB.

Answer: D

Explanation:
This solution requires minimal operational overhead, as it only requires setting up AWS IoT Core and creating a thing for each device. (Reference: AWS Certified Solutions Architect - Professional Official Amazon Text Book, Page 537)
AWS IoT Core is a fully managed service that enables secure, bi-directional communication between internet-connected devices and the AWS Cloud. It supports the MQTT protocol and includes built-in device
authentication and access control. By using AWS IoT Core, the company can easily provision and manage the X.509 certificates for each device, and connect the devices to the service with minimal operational overhead.

NEW QUESTION 22
- (Exam Topic 1)
A publishing company's design team updates the icons and other static assets that an ecommerce web application uses. The company serves the icons and assets from an Amazon S3 bucket that is hosted in the company's production account. The company also uses a development account that members of the design team can access.
After the design team tests the static assets in the development account, the design team needs to load the assets into the S3 bucket in the production account. A solutions architect must provide the design team with access to the production account without exposing other parts of the web application to the risk of unwanted changes.
Which combination of steps will meet these requirements? (Select THREE.)

A. In the production account, create a new IAM policy that allows read and write access to the S3 bucket.
B. In the development account, create a new IAM policy that allows read and write access to the S3 bucket.
C. In the production account, create a rol
D. Attach the new policy to the rol
E. Define the development account as a trusted entity.
F. In the development account, create a rol
G. Attach the new policy to the rol
H. Define the production account as a trusted entity.
I. In the development account, create a group that contains all the IAM users of the design tea
J. Attach a different IAM policy to the group to allow the sts:AssumeRole action on the role in the production account.
K. In the development account, create a group that contains all tfje IAM users of the design tea
L. Attach a different IAM policy to the group to allow the sts;AssumeRole action on the role in the development account.

Answer: ACE

Explanation:
➤ A. In the production account, create a new IAM policy that allows read and write access to the S3 bucket. The policy grants the necessary permissions to access the assets in the production S3 bucket.
➤ C. In the production account, create a role. Attach the new policy to the role. Define the development account as a trusted entity. By creating a role and attaching the policy, and then defining the development account as a trusted entity, the development account can assume the role and access the production S3 bucket with the read and write permissions.
➤ E. In the development account, create a group that contains all the IAM users of the design team. Attach a different IAM policy to the group to allow the sts:AssumeRole action on the role in the production account. The IAM policy attached to the group allows the design team members to assume the role created in the production account, thereby giving them access to the production S3 bucket.
Step 1: Create a role in the Production Account; create the role in the Production account and specify the Development account as a trusted entity. You also limit the role permissions to only read and write access to the productionapp bucket. Anyone granted permission to use the role can read and write to the productionapp bucket. Step 2: Grant access to the role Sign in as an administrator in the Development account and allow the AssumeRole action on the UpdateApp role in the

Production account. So, recap, production account you create the policy for S3, and you set development account as a trusted entity. Then on the development account you allow the sts:assumeRole action on the role in production account. https://docs.aws.amazon.com/IAM/latest/UserGuide/tutorial_cross-account-with-roles.html

NEW QUESTION 23
- (Exam Topic 1)
A company wants to migrate to AWS. The company wants to use a multi-account structure with centrally managed access to all accounts and applications. The company also wants to keep the traffic on a private network. Multi-factor authentication (MFA) is required at login, and specific roles are assigned to user groups. The company must create separate accounts for development. staging, production, and shared network. The production account and the shared network account must have connectivity to all accounts. The development account and the staging account must have access only to each other.
Which combination of steps should a solutions architect take 10 meet these requirements? (Choose three.)

A. Deploy a landing zone environment by using AWS Control Towe
B. Enroll accounts and invite existing accounts into the resulting organization in AWS Organizations.
C. Enable AWS Security Hub in all accounts to manage cross-account acces
D. Collect findings through AWS CloudTrail to force MFA login.
E. Create transit gateways and transit gateway VPC attachments in each accoun
F. Configure appropriate route tables.
G. Set up and enable AWS IAM Identity Center (AWS Single Sign-On). Create appropriate permission sets with required MFA for existing accounts.
H. Enable AWS Control Tower in all Recounts to manage routing between account
I. Collect findings through AWS CloudTrail to force MFA login.
J. Create IAM users and group
K. Configure MFA for all user
L. Set up Amazon Cognito user pools andidentity pools to manage access to accounts and between accounts.

Answer: ACD

Explanation:
The correct answer would be options A, C and D, because they address the requirements outlined in the question. A. Deploying a landing zone environment using AWS Control Tower and enrolling accounts in an organization in AWS Organizations allows for a centralized management of access to all accounts and applications. C. Creating transit gateways and transit gateway VPC attachments in each account and configuring appropriate route tables allows for private network traffic, and ensures that the production account and shared network account have connectivity to all accounts, while the development and staging accounts have access only to each other. D. Setting up and enabling AWS IAM Identity Center (AWS Single Sign-On) and creating appropriate permission sets with required MFA for existing accounts allows for multi-factor authentication at login and specific roles to be assigned to user groups.

NEW QUESTION 25
- (Exam Topic 1)
A company is running a critical application that uses an Amazon RDS for MySQL database to store data. The RDS DB instance is deployed in Multi-AZ mode.
A recent RDS database failover test caused a 40-second outage to the application A solutions architect needs to design a solution to reduce the outage time to less than 20 seconds.
Which combination of steps should the solutions architect take to meet these requirements? (Select THREE.)

A. Use Amazon ElastiCache for Memcached in front of the database
B. Use Amazon ElastiCache for Redis in front of the database.
C. Use RDS Proxy in front of the database
D. Migrate the database to Amazon Aurora MySQL
E. Create an Amazon Aurora Replica
F. Create an RDS for MySQL read replica

Answer: CDE

Explanation:
Migrate the database to Amazon Aurora MySQL. - Create an Amazon Aurora Replica. - Use RDS Proxy in front of the database. - These options are correct because they address the requirement of reducing the failover time to less than 20 seconds. Migrating to Amazon Aurora MySQL and creating an Aurora replica can reduce the failover time to less than 20 seconds. Aurora has a built-in, fault-tolerant storage system that can automatically detect and repair failures. Additionally, Aurora has a feature called "Aurora Global Database" which allows you to create read-only replicas across multiple AWS regions which can further help to reduce the failover time. Creating an Aurora replica can also help to reduce the failover time as it can take over as the primary DB instance in case of a failure. Using RDS proxy can also help to reduce the failover time as it can route the queries to the healthy DB instance, it also helps to balance the load across multiple DB instances.

NEW QUESTION 27
- (Exam Topic 1)
A solutions architect has developed a web application that uses an Amazon API Gateway Regional endpoint and an AWS Lambda function. The consumers of the web application are all close to the AWS Region where the application will be deployed. The Lambda function only queries an Amazon Aurora MySQL database. The solutions architect has configured the database to have three read replicas.
During testing, the application does not meet performance requirements. Under high load, the application opens a large number of database connections. The solutions architect must improve the application's performance.
Which actions should the solutions architect take to meet these requirements? (Choose two.)

A. Use the cluster endpoint of the Aurora database.
B. Use RDS Proxy to set up a connection pool to the reader endpoint of the Aurora database.
C. Use the Lambda Provisioned Concurrency feature.
D. Move the code for opening the database connection in the Lambda function outside of the event handler.
E. Change the API Gateway endpoint to an edge-optimized endpoint.

Answer: BD

Explanation:
Connect to RDS outside of Lambda handler method to improve performance https://awstut.com/en/2022/04/30/connect-to-rds-outside-of-lambda-handler-method-to-improve-performance-en

Using RDS Proxy, you can handle unpredictable surges in database traffic. Otherwise, these surges might cause issues due to oversubscribing connections or creating new connections at a fast rate. RDS Proxy establishes a database connection pool and reuses connections in this pool. This approach avoids the memory and CPU overhead of opening a new database connection each time. To protect the database against oversubscription, you can control the number of database connections that are created. https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/rds-proxy.html

NEW QUESTION 30
- (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 exceptthe 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 33
- (Exam Topic 1)
An AWS partner company is building a service in AWS Organizations using Its organization named org. This service requires the partner company to have access to AWS resources in a customer account, which is in a separate organization named org2 The company must establish least privilege security access using an API or command line tool to the customer account
What is the MOST secure way to allow org1 to access resources h org2?

A. The customer should provide the partner company with their AWS account access keys to log in and perform the required tasks
B. The customer should create an IAM user and assign the required permissions to the IAM user The customer should then provide the credentials to the partner company to log In and perform the required tasks.
C. The customer should create an IAM role and assign the required permissions to the IAM rol
D. The partner company should then use the IAM rote's Amazon Resource Name (ARN) when requesting access to perform the required tasks
E. The customer should create an IAM rote and assign the required permissions to the IAM rot
F. The partner company should then use the IAM rote's Amazon Resource Name (ARN). Including the external ID in the IAM role's trust pokey, when requesting access to perform the required tasks

Answer: C

Explanation:
https://docs.aws.amazon.com/IAM/latest/UserGuide/confused-deputy.html
This is the most secure way to allow org1 to access resources in org2 because it allows for least privilege security access. The customer should create an IAM role and assign the required permissions to the IAM role. The partner company should then use the IAM role's Amazon Resource Name (ARN) and include the external ID in the IAM role's trust policy when requesting access to perform the required tasks. This ensures that the partner company can only access the resources that it needs and only from the specific customer account.

NEW QUESTION 36
- (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 41
- (Exam Topic 1)
A company that uses AWS Organizations allows developers to experiment on AWS. As part of the landing zone that the company has deployed, developers use their company email address to request an account. The company wants to ensure that developers are not launching costly services or running services unnecessarily. The company must give developers a fixed monthly budget to limit their AWS costs.
Which combination of steps will meet these requirements? (Choose three.)

A. Create an SCP to set a fixed monthly account usage limi
B. Apply the SCP to the developer accounts.
C. Use AWS Budgets to create a fixed monthly budget for each developer's account as part of the account creation process.
D. Create an SCP to deny access to costly services and component
E. Apply the SCP to the developer accounts.
F. Create an IAM policy to deny access to costly services and component
G. Apply the IAM policy to the developer accounts.
H. Create an AWS Budgets alert action to terminate services when the budgeted amount is reached.Configure the action to terminate all services.
I. Create an AWS Budgets alert action to send an Amazon Simple Notification Service (Amazon SNS) notification when the budgeted amount is reache
J. Invoke an AWS Lambda function to terminate all services.

Answer: BCF

Explanation:
▷ Option A is incorrect because creating an SCP to set a fixed monthly account usage limit is not possible.
SCPs are policies that specify the services and actions that users and roles can use in the member accounts of an AWS Organization. SCPs cannot enforce budget limits or prevent users from launching
costly services or running services unnecessarily1
▷ Option B is correct because using AWS Budgets to create a fixed monthly budget for each developer's account as part of the account creation process meets the requirement of giving developers a fixed monthly budget to limit their AWS costs. AWS Budgets allows you to plan your service usage, service costs, and instance reservations. You can create budgets that alert you when your costs or usage exceed (or are forecasted to exceed) your budgeted amount2
▷ Option C is correct because creating an SCP to deny access to costly services and components meets the requirement of ensuring that developers are not launching costly services or running services
unnecessarily. SCPs can restrict access to certain AWS services or actions based on conditions such as region, resource tags, or request time. For example, an SCP can deny access to Amazon Redshift clusters or Amazon EC2 instances with certain instance types1
▷ Option D is incorrect because creating an IAM policy to deny access to costly services and components is not sufficient to meet the requirement of ensuring that developers are not launching costly services or running services unnecessarily. IAM policies can only control access to resources within a single AWS account. If developers have multiple accounts or can create new accounts, they can bypass the IAM policy restrictions. SCPs can apply across multiple accounts within an AWS Organization and prevent users from creating new accounts that do not comply with the SCP rules3
▷ Option E is incorrect because creating an AWS Budgets alert action to terminate services when the budgeted amount is reached is not possible. AWS Budgets alert actions can only perform one of the following actions: apply an IAM policy, apply an SCP, or send a notification through Amazon SNS. AWS Budgets alert actions cannot terminate services directly.
▷ Option F is correct because creating an AWS Budgets alert action to send an Amazon SNS notification when the budgeted amount is reached and invoking an AWS Lambda function to terminate all services meets the requirement of giving developers a fixed monthly budget to limit their AWS costs. AWS Budgets alert actions can send notifications through Amazon SNS when a budget threshold is breached. Amazon SNS can trigger an AWS Lambda function that can perform custom logic such as terminating all services in the developer's account. This way, developers cannot exceed their budget limit and incur additional costs.
References: 1: https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps.html 2
: https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/budgets-create.html 3: https://docs.aws.amazon.com/IAM/latest/UserGuide/introduction.html :
https://docs.aws.amazon.com/cost-management/latest/userguide/budgets-actions.html : https://docs.aws.amazon.com/sns/latest/dg/sns-lambda.html :
https://docs.aws.amazon.com/lambda/latest/dg/welcome.html

NEW QUESTION 44
- (Exam Topic 1)
A company's solutions architect is reviewing a new internally developed application in a sandbox AWS account The application uses an AWS Auto Scaling group of Amazon EC2 instances that have an IAM instance profile attached Part of the application logic creates and accesses secrets from AWS Secrets Manager The company has an AWS Lambda function that calls the application API to test the functionality The company also has created an AWS CloudTrail trail in the account The application's developer has attached the SecretsManagerReadWnte AWS managed IAM policy to an IAM role The IAM role is associated with the instance profile that is attached to the EC2 instances The solutions architect has invoked the Lambda function for testing
The solutions architect must replace the SecretsManagerReadWnte policy with a new policy that provides least privilege access to the Secrets Manager actions that the application requires
What is the MOST operationally efficient solution that meets these requirements?

A. Generate a policy based on CloudTrail events for the IAM role Use the generated policy output to create a new IAM policy Use the newly generated IAM policy to replace the SecretsManagerReadWnte policy that is attached to the IAM role
B. Create an analyzer in AWS Identity and Access Management Access Analyzer Use the IAM role's Access Advisor findings to create a new IAM policy Use the newly created IAM policy to replace the SecretsManagerReadWnte policy that is attached to the IAM role
C. Use the aws cloudtrail lookup-events AWS CLI command to filter and export CloudTrail events that are related to Secrets Manager Use a new IAM policy that contains the actions from CloudTrail to replace the SecretsManagerReadWnte policy that is attached to the IAM role
D. Use the IAM policy simulator to generate an IAM policy for the IAM role Use the newly generated IAM policy to replace the SecretsManagerReadWnte policy that is attached to the IAM role

Answer: B

Explanation:
The IAM policy simulator will generate a policy that contains only the necessary permissions for the application to access Secrets Manager, providing the least privilege necessary to get the job done. This is the most efficient solution as it will not require additional steps such as analyzing CloudTrail events or manually creating and testing an IAM policy.
You can use the IAM policy simulator to generate an IAM policy for an IAM role by specifying the role and the API actions and resources that the application or service requires. The simulator will then generate an IAM policy that grants the least privilege access to those actions and resources.
Once you have generated an IAM policy using the simulator, you can replace the existing SecretsManagerReadWnte policy that is attached to the IAM role with the newly generated policy. This will ensure that the application or service has the least privilege access to the Secrets Manager actions that it requires.
You can access the IAM policy simulator through the IAM console, AWS CLI, and AWS SDKs. Here is the link for more information:
https://docs.aws.amazon.com/IAM/latest/UserGuide/access_policies_simulator.html

NEW QUESTION 49
- (Exam Topic 1)
A finance company is running its business-critical application on current-generation Linux EC2 instances The application includes a self-managed MySQL database performing heavy I/O operations. The application is working fine to handle a moderate amount of traffic during the month. However, it slows down during the final three days of each month due to month-end reporting, even though the company is using Elastic Load Balancers and Auto Scaling within its infrastructure to meet the increased demand.
Which of the following actions would allow the database to handle the month-end load with the LEAST impact on performance?

A. Pre-warming Elastic Load Balancers, using a bigger instance type, changing all Amazon EBS volumes to GP2 volumes.
B. Performing a one-time migration of the database cluster to Amazon RD
C. and creating several additional read replicas to handle the load during end of month
D. Using Amazon CioudWatch with AWS Lambda to change the typ
E. size, or IOPS of Amazon EBS volumes in the cluster based on a specific CloudWatch metric
F. Replacing all existing Amazon EBS volumes with new PIOPS volumes that have the maximum available storage size and I/O per second by taking snapshots before the end of the month and reverting back afterwards.

Answer: B

Explanation:
In this scenario, the Amazon EC2 instances are in an Auto Scaling group already which means that the database read operations is the possible bottleneck especially during the month-end wherein the reports are generated. This can be solved by creating RDS read replicas.

NEW QUESTION 51
- (Exam Topic 1)
A company runs a Python script on an Amazon EC2 instance to process data. The script runs every
10 minutes. The script ingests files from an Amazon S3 bucket and processes the files. On average, the script takes approximately 5 minutes to process each file The script will not reprocess a file that the script has already processed.
The company reviewed Amazon CloudWatch metrics and noticed that the EC2 instance is idle for approximately 40% of the time because of the file processing speed. The company wants to make the workload highly available and scalable. The company also wants to reduce long-term management overhead.
Which solution will meet these requirements MOST cost-effectively?

A. Migrate the data processing script to an AWS Lambda functio
B. Use an S3 event notification to invoke the Lambda function to process the objects when the company uploads the objects.
C. Create an Amazon Simple Queue Service (Amazon SQS) queu
D. Configure Amazon S3 to send event notifications to the SQS queu
E. Create an EC2 Auto Scaling group with a minimum size of one instanc
F. Update the data processing script to poll the SQS queu
G. Process the S3 objects that the SQS message identifies.
H. Migrate the data processing script to a container imag
I. Run the data processing container on an EC2 instanc
J. Configure the container to poll the S3 bucket for new objects and to process the resulting objects.
K. Migrate the data processing script to a container image that runs on Amazon Elastic Container Service (Amazon ECS) on AWS Fargat
L. Create an AWS Lambda function that calls the Fargate RunTaskAPI operation when the container processes the fil
M. Use an S3 event notification to invoke the Lambda function.

Answer: D

Explanation:
migrating the data processing script to an AWS Lambda function and using an S3 event notification to invoke the Lambda function to process the objects when the company uploads the objects. This solution meets the company's requirements of high availability and scalability, as well as reducing long-term management overhead, and is likely to be the most cost-effective option.

NEW QUESTION 54
- (Exam Topic 1)
A company is in the process of implementing AWS Organizations to constrain its developers to use only Amazon EC2. Amazon S3 and Amazon DynamoDB. The developers account resides In a dedicated organizational unit (OU). The solutions architect has implemented the following SCP on the developers account:

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

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

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

Answer: B

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

NEW QUESTION 55
- (Exam Topic 1)
A company is using AWS Organizations lo manage multiple AWS accounts For security purposes, the company requires the creation of an Amazon Simple Notification Service (Amazon SNS) topic that enables integration with a third-party alerting system in all the Organizations member accounts
A solutions architect used an AWS CloudFormation template to create the SNS topic and stack sets to automate the deployment of CloudFormation stacks Trusted access has been enabled in Organizations
What should the solutions architect do to deploy the CloudFormation StackSets in all AWS accounts?

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

Answer: C

Explanation:
https://aws.amazon.com/blogs/aws/use-cloudformation-stacksets-to-provision-resources-across-multiple-aws-ac

NEW QUESTION 56
- (Exam Topic 1)
A company is hosting a critical application on a single Amazon EC2 instance. The application uses an Amazon ElastiCache for Redis single-node cluster for an in-memory data store. The application uses an Amazon RDS for MariaDB DB instance for a relational database. For the application to function, each piece of the infrastructure must be healthy and must be in an active state.
A solutions architect needs to improve the application's architecture so that the infrastructure can automatically recover from failure with the least possible downtime.
Which combination of steps will meet these requirements? (Select THREE.)

A. Use an Elastic Load Balancer to distribute traffic across multiple EC2 instance
B. Ensure that the EC2 instances are part of an Auto Scaling group that has a minimum capacity of two instances.
C. Use an Elastic Load Balancer to distribute traffic across multiple EC2 instances Ensure that the EC2 instances are configured in unlimited mode.
D. Modify the DB instance to create a read replica in the same Availability Zon
E. Promote the read replica to be the primary DB instance in failure scenarios.
F. Modify the DB instance to create a Multi-AZ deployment that extends across two Availability Zones.
G. Create a replication group for the ElastiCache for Redis cluste
H. Configure the cluster to use an Auto Scaling group that has a minimum capacity of two instances.
I. Create a replication group for the ElastiCache for Redis cluste
J. Enable Multi-AZ on the cluster.

Answer: ADF

Explanation:

⟩ Option A is correct because using an Elastic Load Balancer and an Auto Scaling group with a minimum capacity of two instances can improve the availability and scalability of the EC2 instances that host the application. The load balancer can distribute traffic across multiple instances and the Auto Scaling group can replace any unhealthy instances automatically1

⟩ Option D is correct because modifying the DB instance to create a Multi-AZ deployment that extends across two Availability Zones can improve the availability and durability of the RDS for MariaDB
database. Multi-AZ deployments provide enhanced data protection and minimize downtime by automatically failing over to a standby replica in another Availability Zone in case of a planned or unplanned outage4

⟩ Option F is correct because creating a replication group for the ElastiCache for Redis cluster and enabling Multi-AZ on the cluster can improve the availability and fault tolerance of the in-memory data store. A replication group consists of a primary node and up to five read-only replica nodes that are synchronized with the primary node using asynchronous replication. Multi-AZ allows automatic failove to one of the replicas if the primary node fails or becomes unreachable6
References: 1:
https://docs.aws.amazon.com/elasticloadbalancing/latest/userguide/how-elastic-load-balancing-works.html 2:
https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/burstable-performance-instances-unlimited-mode.htm 3:
https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_ReadRepl.html 4:
https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Concepts.MultiAZ.html 5:
https://docs.aws.amazon.com/AmazonElastiCache/latest/red-ug/AutoScaling.html 6: https://docs.aws.amazon.com/AmazonElastiCache/latest/red-ug/Replication.Redis.Groups.html

NEW QUESTION 59
- (Exam Topic 1)
A company is hosting a monolithic REST-based API for a mobile app on five Amazon EC2 instances in public subnets of a VPC. Mobile clients connect to the API by using a domain name that is hosted on Amazon Route 53. The company has created a Route 53 multivalue answer routing policy with the IP addresses of all the EC2 instances. Recently, the app has been overwhelmed by large and sudden increases to traffic. The app has not been able to keep up with the traffic.
A solutions architect needs to implement a solution so that the app can handle the new and varying load. Which solution will meet these requirements with the LEAST operational overhead?

A. Separate the API into individual AWS Lambda function
B. Configure an Amazon API Gateway REST API with Lambda integration for the backen
C. Update the Route 53 record to point to the API Gateway API.
D. Containerize the API logi
E. Create an Amazon Elastic Kubernetes Service (Amazon EKS) cluste
F. Run the containers in the cluster by using Amazon EC2. Create a Kubernetes ingres
G. Update the Route 53 record to point to the Kubernetes ingress.
H. Create an Auto Scaling grou
I. Place all the EC2 instances in the Auto Scaling grou
J. Configure the Auto Scaling group to perform scaling actions that are based on CPU utilizatio
K. Create an AWS Lambda function that reacts to Auto Scaling group changes and updates the Route 53 record.
L. Create an Application Load Balancer (ALB) in front of the AP
M. Move the EC2 instances to private subnets in the VP
N. Add the EC2 instances as targets for the AL
O. Update the Route 53 record to point to the ALB.

Answer: D

Explanation:
By breaking down the monolithic API into individual Lambda functions and using API Gateway to handle the incoming requests, the solution can automatically scale to handle the new and varying load without the need for manual scaling actions. Additionally, this option will automatically handle the traffic without the need of having EC2 instances running all the time and only pay for the number of requests and the duration of the execution of the Lambda function.
By updating the Route 53 record to point to the API Gateway, the solution can handle the traffic and also it will direct the traffic to the correct endpoint.

NEW QUESTION 62
- (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 65
- (Exam Topic 1)
A company is using multiple AWS accounts The DNS records are stored in a private hosted zone for Amazon Route 53 in Account A The company's applications and databases are running in Account B.
A solutions architect win deploy a two-net application In a new VPC To simplify the configuration, the db.example com CNAME record set tor the Amazon RDS endpoint was created in a private hosted zone for Amazon Route 53.
During deployment, the application failed to start. Troubleshooting revealed that db.example com is not resolvable on the Amazon EC2 instance The solutions architect confirmed that the record set was created correctly in Route 53.
Which combination of steps should the solutions architect take to resolve this issue? (Select TWO)

A. Deploy the database on a separate EC2 instance in the new VPC Create a record set for the instance's private IP in the private hosted zone
B. Use SSH to connect to the application tier EC2 instance Add an RDS endpoint IP address to the/eto/resolv.conf file
C. Create an authorization lo associate the private hosted zone in Account A with the new VPC In Account B
D. Create a private hosted zone for the example.com domain m Account B Configure Route 53 replication between AWS accounts
E. Associate a new VPC in Account B with a hosted zone in Account
F. Delete the association authorization In Account A.

Answer: CE

Explanation:
https://aws.amazon.com/premiumsupport/knowledge-center/private-hosted-zone-different-account/

NEW QUESTION 68
- (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 71
- (Exam Topic 1)
A company is using Amazon OpenSearch Service to analyze data. The company loads data into an OpenSearch Service cluster with 10 data nodes from an Amazon S3 bucket that uses S3 Standard storage. The data resides in the cluster for 1 month for read-only analysis. After 1 month, the company deletes the index that contains the data from the cluster. For compliance purposes, the company must retain a copy of all input data.
The company is concerned about ongoing costs and asks a solutions architect to recommend a new solution. Which solution will meet these requirements MOST cost-effectively?

A. Replace all the data nodes with UltraWarm nodes to handle the expected capacit
B. Transition the input data from S3 Standard to S3 Glacier Deep Archive when the company loads the data into the cluster.
C. Reduce the number of data nodes in the cluster to 2 Add UltraWarm nodes to handle the expected capacit

D. Configure the indexes to transition to UltraWarm when OpenSearch Service ingests the dat
E. Transition the input data to S3 Glacier Deep Archive after 1 month by using an S3 Lifecycle policy.
F. Reduce the number of data nodes in the cluster to 2. Add UltraWarm nodes to handle the expected capacit
G. Configure the indexes to transition to UltraWarm when OpenSearch Service ingests the dat
H. Add cold storage nodes to the cluster Transition the indexes from UltraWarm to cold storag
I. Delete the input data from the S3 bucket after 1 month by using an S3 Lifecycle policy.
J. Reduce the number of data nodes in the cluster to 2. Add instance-backed data nodes to handle the expected capacit
K. Transition the input data from S3 Standard to S3 Glacier Deep Archive when the company loads the data into the cluster.

Answer: B

Explanation:
By reducing the number of data nodes in the cluster to 2 and adding UltraWarm nodes to handle the expected capacity, the company can reduce the cost of running the cluster. Additionally, configuring the indexes to transition to UltraWarm when OpenSearch Service ingests the data will ensure that the data is stored in the most cost-effective manner. Finally, transitioning the input data to S3 Glacier Deep Archive after 1 month by using an S3 Lifecycle policy will ensure that the data is retained for compliance purposes, while also reducing the ongoing costs.

NEW QUESTION 74
- (Exam Topic 1)
A retail company has an on-premises data center in Europe. The company also has a multi-Region AWS presence that includes the eu-west-1 and us-east-1 Regions. The company wants to be able to route network traffic from its on-premises infrastructure into VPCs in either of those Regions. The company also needs to support traffic that is routed directly between VPCs in those Regions. No single points of failure can exist on the network.
The company already has created two 1 Gbps AWS Direct Connect connections from its on-premises data center. Each connection goes into a separate Direct Connect location in Europe for high availability. These two locations are named DX-A and DX-B, respectively. Each Region has a single AWS Transit Gateway that is configured to route all inter-VPC traffic within that Region.
Which solution will meet these requirements?

A. Create a private VIF from the DX-A connection into a Direct Connect gatewa
B. Create a private VIF from the DX-B connection into the same Direct Connect gateway for high availabilit
C. Associate both the eu-west-1 and us-east-1 transit gateways with the Direct Connect gatewa
D. Peer the transit gatewayswith each other to support cross-Region routing.
E. Create a transit VIF from the DX-A connection into a Direct Connect gatewa
F. Associate the eu-west-1 transit gateway with this Direct Connect gatewa
G. Create a transit VIF from the DX-B connection into a separate Direct Connect gatewa
H. Associate the us-east-1 transit gateway with this separate Direct Connect gatewa
I. Peer the Direct Connect gateways with each other to support high availability and cross-Region routing.
J. Create a transit VIF from the DX-A connection into a Direct Connect gatewa
K. Create a transit VIF from the DX-B connection into the same Direct Connect gateway for high availabilit
L. Associate both the eu-west-1 and us-east-1 transit gateways with this Direct Connect gatewa
M. Configure the Direct Connect gateway to route traffic between the transit gateways.
N. Create a transit VIF from the DX-A connection into a Direct Connect gatewa
O. Create a transit VIF from the DX-B connection into the same Direct Connect gateway for high availabilit
P. Associate both the eu-west-1 and us-east-1 transit gateways with this Direct Connect gatewa
Q. Peer the transit gateways with each other to support cross-Region routing.

Answer: D

Explanation:
in this solution, two transit VIFs are created - one from the DX-A connection and one from the DX-B connection - into the same Direct Connect gateway for high availability. Both the eu-west-1 and us-east-1 transit gateways are then associated with this Direct Connect gateway. The transit gateways are then peered with each other to support cross-Region routing. This solution meets the requirements of the company by creating a highly available connection between the on-premises data center and the VPCs in both the eu-west-1 and us-east-1 regions, and by enabling direct traffic routing between VPCs in those regions.

NEW QUESTION 78
- (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 82
- (Exam Topic 1)
A company is planning to host a web application on AWS and works to load balance the traffic across a group of Amazon EC2 instances. One of the security requirements is to enable end-to-end encryption in transit between the client and the web server.
Which solution will meet this requirement?

A. Place the EC2 instances behind an Application Load Balancer (ALB) Provision an SSL certificate using AWS Certificate Manager (ACM), and associate the SSL certificate with the AL
B. Export the SSL certificate and install it on each EC2 instanc
C. Configure the ALB to listen on port 443 and to forward traffic to port 443 on the instances.
D. Associate the EC2 instances with a target grou
E. Provision an SSL certificate using AWS Certificate Manager (ACM). Create an Amazon CloudFront distribution and configure It to use the SSL certificat
F. Set CloudFront to use the target group as the origin server
G. Place the EC2 instances behind an Application Load Balancer (ALB). Provision an SSL certificate using AWS Certificate Manager (ACM), and associate the SSL certificate with the AL
H. Provision athird-party SSL certificate and install it on each EC2 instanc
I. Configure the ALB to listen on port 443 and to forward traffic to port 443 on the instances.
J. Place the EC2 instances behind a Network Load Balancer (NLB). Provision a third-party SSL certificate and install it on the NLB and on each EC2 instanc
K. Configure the NLB to listen on port 443 and to forward traffic to port 443 on the instances.

Answer: A

Explanation:
 Option A is correct because placing the EC2 instances behind an Application Load Balancer (ALB) and associating an SSL certificate from AWS Certificate Manager (ACM) with the ALB enables encryption in transit between the client and the ALB. Exporting the SSL certificate and installing it on each EC2 instance enables encryption in transit between the ALB and the web server. Configuring the ALB to listen on port 443 and to forward traffic to port 443 on the instances ensures that HTTPS is used for both connections. This solution achieves end-to-end encryption in transit for the web applicatio1n2
References: 1: https://docs.aws.amazon.com/elasticloadbalancing/latest/application/introduction.html 2:
https://docs.aws.amazon.com/acm/latest/userguide/acm-overview.html 3: https://docs.aws.amazon.com/elasticloadbalancing/latest/application/load-balancer-target-groups.html : https://aws.amazon.com/certificate-manager/faqs/ : https://docs.aws.amazon.com/elasticloadbalancing/latest/network/introduction.html

NEW QUESTION 87
- (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 91
- (Exam Topic 1)
A solutions architect is designing the data storage and retrieval architecture for a new application that a company will be launching soon. The application is designed to ingest millions of small records per minute from devices all around the world. Each record is less than 4 KB in size and needs to be stored in a durable location where it can be retrieved with low latency. The data is ephemeral and the company is required to store the data for 120 days only, after which the data can be deleted.
The solutions architect calculates that, during the course of a year, the storage requirements would be about 10-15 TB.
Which storage strategy is the MOST cost-effective and meets the design requirements?

A. Design the application to store each incoming record as a single .csv file in an Amazon S3 bucket to allow for indexed retrieva
B. Configure a lifecycle policy to delete data older than 120 days.
C. Design the application to store each incoming record in an Amazon DynamoDB table properly configured for the scal
D. Configure the DynamoOB Time to Live (TTL) feature to delete records older than 120 days.
E. Design the application to store each incoming record in a single table in an Amazon RDS MySQL databas
F. Run a nightly cron job that executes a query to delete any records older than 120 days.
G. Design the application to batch incoming records before writing them to an Amazon S3 bucke
H. Updatethe metadata for the object to contain the list of records in the batch and use the Amazon S3 metadata search feature to retrieve the dat
I. Configure a lifecycle policy to delete the data after 120 days.

Answer: B

Explanation:
DynamoDB with TTL, cheaper for sustained throughput of small items + suited for fast retrievals. S3 cheaper for storage only, much higher costs with writes. RDS not designed for this use case.

NEW QUESTION 96
- (Exam Topic 1)
A company is planning to store a large number of archived documents and make the documents available to employees through the corporate intranet. Employees will access the system by connecting through a client VPN service that is attached to a VPC. The data must not be accessible to the public.
The documents that the company is storing are copies of data that is held on physical media elsewhere. The number of requests will be low. Availability and speed of retrieval are not concerns of the company.
Which solution will meet these requirements at the LOWEST cost?

A. Create an Amazon S3 bucke
B. Configure the S3 bucket to use the S3 One Zone-Infrequent Access (S3 One Zone-IA) storage class as defaul
C. Configure the S3 bucket for website hostin
D. Create an S3 interface endpoin
E. Configure the S3 bucket to allow access only through that endpoint.
F. Launch an Amazon EC2 instance that runs a web serve
G. Attach an Amazon Elastic File System (Amazon EFS) file system to store the archived data in the EFS One Zone-Infrequent Access (EFS One Zone-IA) storage class Configure the instance security groups to allow access only from private networks.
H. Launch an Amazon EC2 instance that runs a web server Attach an Amazon Elastic Block Store (Amazon EBS) volume to store the archived dat
I. Use the Cold HDD (sc1) volume typ
J. Configure the instance security groups to allow access only from private networks.
K. Create an Amazon S3 bucke
L. Configure the S3 bucket to use the S3 Glacier Deep Archive storage class as defaul
M. Configure the S3 bucket for website hostin
N. Create an S3 interface endpoin
O. Configure the S3 bucket to allow access only through that endpoint.

Answer: D

Explanation:
The S3 Glacier Deep Archive storage class is the lowest-cost storage class offered by Amazon S3, and it is designed for archival data that is accessed infrequently and for which retrieval time of several hours is acceptable. S3 interface endpoint for the VPC ensures that access to the bucket is only from resources within the VPC and this will meet the requirement of not being accessible to the public. And also, S3 bucket can be configured for website hosting, and this will allow employees to access the documents through the corporate intranet. Using an EC2 instance and a file system or block store would be more expensive and unnecessary because the number of requests to the data will be low and availability and speed of retrieval are not concerns. Additionally, using Amazon S3 bucket will provide durability, scalability and availability of data.

NEW QUESTION 101
- (Exam Topic 1)
A company has registered 10 new domain names. The company uses the domains for online marketing. The company needs a solution that will redirect online visitors to a specific URL for each domain. All domains and target URLs are defined in a JSON document. All DNS records are managed by Amazon Route 53.
A solutions architect must implement a redirect service that accepts HTTP and HTTPS requests.
Which combination of steps should the solutions architect take to meet these requirements with the LEAST amount of operational effort? (Choose three.)

A. Create a dynamic webpage that runs on an Amazon EC2 instanc
B. Configure the webpage to use the JSON document in combination with the event message to look up and respond with a redirect URL.
C. Create an Application Load Balancer that includes HTTP and HTTPS listeners.
D. Create an AWS Lambda function that uses the JSON document in combination with the event message to look up and respond with a redirect URL.
E. Use an Amazon API Gateway API with a custom domain to publish an AWS Lambda function.
F. Create an Amazon CloudFront distributio
G. Deploy a Lambda@Edge function.
H. Create an SSL certificate by using AWS Certificate Manager (ACM). Include the domains as Subject Alternative Names.

Answer: CEF

Explanation:
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/lambda-edge-how-it-works-tutorial.ht

NEW QUESTION 103
- (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 105
- (Exam Topic 1)
A company has a legacy monolithic application that is critical to the company's business. The company hosts the application on an Amazon EC2 instance that runs Amazon Linux 2. The company's application team receives a directive from the legal department to back up the data from the instance's encrypted Amazon Elastic Block Store (Amazon EBS) volume to an Amazon S3 bucket. The application team does not have the administrative SSH key pair for the instance. The application must continue to serve the users.
Which solution will meet these requirements?

A. Attach a role to the instance with permission to write to Amazon S3. Use the AWS Systems Manager Session Manager option to gain access to the instance and run commands to copy data into Amazon S3.
B. Create an image of the instance with the reboot option turned o
C. Launch a new EC2 instance from the imag
D. Attach a role to the new instance with permission to write to Amazon S3. Run a command to copy data into Amazon S3.
E. Take a snapshot of the EBS volume by using Amazon Data Lifecycle Manager (Amazon DLM). Copy the data to Amazon S3.
F. Create an image of the instanc
G. Launch a new EC2 instance from the imag
H. Attach a role to the new instance with permission to write to Amazon S3. Run a command to copy data into Amazon S3.

Answer: C

Explanation:
Taking a snapshot of the EBS volume using Amazon Data Lifecycle Manager (DLM) will meet the requirements because it allows you to create a backup of the volume without the need to access the instance or its SSH key pair. Additionally, DLM allows you to schedule the backups to occur at specific intervals and also enables you to copy the snapshots to an S3 bucket. This approach will not impact the running application as the backup is performed on the EBS volume level.

NEW QUESTION 108
- (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 112
- (Exam Topic 1)
A company used Amazon EC2 instances to deploy a web fleet to host a blog site The EC2 instances are behind an Application Load Balancer (ALB) and are configured in an Auto ScaSng group The web application stores all blog content on an Amazon EFS volume.
The company recently added a feature 'or Moggers to add video to their posts, attracting 10 times the previous user traffic At peak times of day. users report buffering and timeout issues while attempting to reach the site or watch videos
Which is the MOST cost-efficient and scalable deployment that win resolve the issues for users?

A. Reconfigure Amazon EFS to enable maximum I/O.
B. Update the Nog site to use instance store volumes tor storag
C. Copy the site contents to the volumes atlaunch and to Amazon S3 al shutdown.
D. Configure an Amazon CloudFront distributio
E. Point the distribution to an S3 bucket, and migrate the videos from EFS to Amazon S3.
F. Set up an Amazon CloudFront distribution for all site contents, and point the distribution at the ALB.

Answer: C

Explanation:
https://aws.amazon.com/premiumsupport/knowledge-center/cloudfront-https-connection-fails/ Using an Amazon S3 bucket
Using a MediaStore container or a MediaPackage channel Using an Application Load Balancer
Using a Lambda function URL

Using Amazon EC2 (or another custom origin)
Using CloudFront origin groups https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/restrict-access-to-load-balancer.html

NEW QUESTION 116
- (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 118
- (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 122
- (Exam Topic 1)
A video processing company has an application that downloads images from an Amazon S3 bucket, processes the images, stores a transformed image in a second S3 bucket, and updates metadata about the image in an Amazon DynamoDB table. The application is written in Node.js and runs by using an AWS Lambda function. The Lambda function is invoked when a new image is uploaded to Amazon S3.
The application ran without incident for a while. However, the size of the images has grown significantly. The Lambda function is now failing frequently with timeout errors. The function timeout is set to its maximum value. A solutions architect needs to refactor the application's architecture to prevent invocation failures. The company does not want to manage the underlying infrastructure.
Which combination of steps should the solutions architect take to meet these requirements? (Choose two.)

A. Modify the application deployment by building a Docker image that contains the application code.Publish the image to Amazon Elastic Container Registry (Amazon ECR).
B. Create a new Amazon Elastic Container Service (Amazon ECS) task definition with a compatibility type of AWS Fargat
C. Configure the task definition to use the new image in Amazon Elastic Container Registry (Amazon ECR). Adjust the Lambda function to invoke an ECS task by using the ECS task definition when a new file arrives in Amazon S3.
D. Create an AWS Step Functions state machine with a Parallel state to invoke the Lambda function.Increase the provisioned concurrency of the Lambda function.
E. Create a new Amazon Elastic Container Service (Amazon ECS) task definition with a compatibility type of Amazon EC2. Configure the task definition to use the new image in Amazon Elastic Container Registry (Amazon ECR). Adjust the Lambda function to invoke an ECS task by using the ECS task definition when a new file arrives in Amazon S3.
F. Modify the application to store images on Amazon Elastic File System (Amazon EFS) and to store metadata on an Amazon RDS DB instanc

G. Adjust the Lambda function to mount the EFS file share.

Answer: AB

Explanation:
A. Modify the application deployment by building a Docker image that contains the application code. Publish the image to Amazon Elastic Container Registry (Amazon ECR). - This step is necessary to package the application code in a container and make it available for running on ECS. B. Create a new Amazon Elastic Container Service (Amazon ECS) task definition with a compatibility type of AWS Fargate. Configure the task definition to use the new image in Amazon Elastic Container Registry (Amazon ECR). Adjust the Lambda function to invoke an ECS task by using the ECS task definition when a new file arrives in Amazon S3.

NEW QUESTION 125
- (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 128
- (Exam Topic 1)
A company is developing a new service that will be accessed using TCP on a static port A solutions architect must ensure that the service is highly available, has redundancy across Availability Zones, and is accessible using the DNS name myservice.com, which is publicly accessible The service must use fixed address assignments so other companies can add the addresses to their allow lists.
Assuming that resources are deployed in multiple Availability Zones in a single Region, which solution will meet these requirements?

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

Answer: C

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

NEW QUESTION 131
- (Exam Topic 1)
A company is hosting an image-processing service on AWS in a VPC. The VPC extends across two Availability Zones. Each Availability Zone contains one public subnet and one private subnet.
The service runs on Amazon EC2 instances in the private subnets. An Application Load Balancer in the public subnets is in front of the service. The service needs to communicate with the internet and does so through two NAT gateways. The service uses Amazon S3 for image storage. The EC2 instances retrieve approximately 1 ¢' of data from an S3 bucket each day.
The company has promoted the service as highly secure. A solutions architect must reduce cloud expenditures as much as possible without compromising the service's security posture or increasing the time spent on ongoing operations.
Which solution will meet these requirements?

A. Replace the NAT gateways with NAT instance
B. In the VPC route table, create a route from the private subnets to the NAT instances.
C. Move the EC2 instances to the public subnet
D. Remove the NAT gateways.
E. Set up an S3 gateway VPC endpoint in the VP
F. Attach an endpoint policy to the endpoint to allow the required actions on the S3 bucket.
G. Attach an Amazon Elastic File System (Amazon EFS) volume to the EC2 instance
H. Host the image on the EFS volume.

Answer: C

Explanation:
Create Amazon S3 gateway endpoint in the VPC and add a VPC endpoint policy. This VPC endpoint policy will have a statement that allows S3 access only via access points owned by the organization.

NEW QUESTION 132
- (Exam Topic 1)
A start up company hosts a fleet of Amazon EC2 instances in private subnets using the latest Amazon Linux 2 AMI. The company's engineers rely heavily on SSH access to the instances for troubleshooting.
The company's existing architecture includes the following:
• A VPC with private and public subnets, and a NAT gateway
• Site-to-Site VPN for connectivity with the on-premises environment
• EC2 security groups with direct SSH access from the on-premises environment
The company needs to increase security controls around SSH access and provide auditing of commands executed by the engineers.
Which strategy should a solutions architect use?

A. Install and configure EC2 Instance Connect on the fleet of EC2 instance
B. Remove all security group rules attached to EC2 instances that allow inbound TCP on port 22. Advise the engineers to remotely access the instances by using the EC2 Instance Connect CLI.
C. Update the EC2 security groups to only allow inbound TCP on port 22 to the IP addresses of the engineer's device
D. Install the Amazon CloudWatch agent on all EC2 instances and send operating system audit logs to CloudWatch Logs.
E. Update the EC2 security groups to only allow inbound TCP on port 22 to the IP addresses of the engineer's device
F. Enable AWS Config for EC2 security group resource change
G. Enable AWS Firewall Manager and apply a security group policy that automatically remediates changes to rules.
H. Create an IAM role with the AmazonSSMManagedInstanceCore managed policy attache
I. Attach the IAM role to all the EC2 instance
J. Remove all security group rules attached to the EC2 instances that allow inbound TCP on port 22. Have the engineers install the AWS Systems Manager Session Manager plugin for their devices and remotely access the instances by using the start-session API call from Systems Manager.

Answer: D

Explanation:
Allows client machines to be able to connect to Session Manager using the AWS CLI instead of going through the AWS EC2 or AWS Server Manager console. https://docs.aws.amazon.com/systems-manager/latest/userguide/session-manager-working-with-install-plugin.ht https://docs.aws.amazon.com/systems-manager/latest/userguide/session-manager-working-with-install-plugin.ht

NEW QUESTION 134
- (Exam Topic 1)
A company has deployed an application on AWS Elastic Beanstalk. The application uses Amazon Aurora for the database layer. An Amazon CloudFront distribution serves web requests and includes the Elastic Beanstalk domain name as the origin server. The distribution is configured with an alternate domain name that visitors use when they access the application.
Each week, the company takes the application out of service for routine maintenance. During the time that the application is unavailable, the company wants visitors to receive an informational message instead of a CloudFront error message.
A solutions architect creates an Amazon S3 bucket as the first step in the process.
Which combination of steps should the solutions architect take next to meet the requirements? (Choose three.)

A. Upload static informational content to the S3 bucket.
B. Create a new CloudFront distributio
C. Set the S3 bucket as the origin.
D. Set the S3 bucket as a second origin in the original CloudFront distributio
E. Configure the distribution and the S3 bucket to use an origin access identity (OAI).
F. During the weekly maintenance, edit the default cache behavior to use the S3 origi
G. Revert the change when the maintenance is complete.
H. During the weekly maintenance, create a cache behavior for the S3 origin on the new distributio
I. Set the path pattern to \ Set the precedence to 0. Delete the cache behavior when the maintenance is complete.
J. During the weekly maintenance, configure Elastic Beanstalk to serve traffic from the S3 bucket.

Answer: ACD

Explanation:
The company wants to serve static content from an S3 bucket during the maintenance period. To do this, the following steps are required:

➤ Upload static informational content to the S3 bucket. This will provide the source of the content that will be served to the visitors.

➤ Set the S3 bucket as a second origin in the original CloudFront distribution. Configure the distribution and the S3 bucket to use an origin access identity (OAI). This will allow CloudFront to access the S3 bucket securely and prevent public access to the bucket.

➤ During the weekly maintenance, edit the default cache behavior to use the S3 origin. Revert the change when the maintenance is complete. This will redirect all web requests to the S3 bucket instead of the Elastic Beanstalk domain name.
The other options are not correct because:

➤ Creating a new CloudFront distribution is not necessary and would require changing the alternate domain name configuration.

> Creating a cache behavior for the S3 origin on a new distribution would not work because the visitors would still access the original distribution using the alternate domain name.

> Configuring Elastic Beanstalk to serve traffic from the S3 bucket is not possible and would not achieve the desired result.
References:

> https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/distribution-web-values-specify.

NEW QUESTION 137
- (Exam Topic 2)
A company's solutions architect is analyzing costs of a multi-application environment. The environment is deployed across multiple Availability Zones in a single AWS Region. After a recent acquisition, the company manages two organizations in AWS Organizations. The company has created multiple service provider applications as AWS PrivateLink-powered VPC endpoint services in one organization. The company has created multiple service consumer applications in the other organization.
Data transfer charges are much higher than the company expected, and the solutions architect needs to reduce the costs. The solutions architect must recommend guidelines for developers to follow when they deploy services. These guidelines must minimize data transfer charges for the whole environment.
Which guidelines meet these requirements? (Select TWO.)

A. Use AWS Resource Access Manager to share the subnets that host the service provider applications with other accounts in the organization.
B. Place the service provider applications and the service consumer applications in AWS accounts in the same organization.
C. Turn off cross-zone load balancing for the Network Load Balancer in all service provider application deployments.
D. Ensure that service consumer compute resources use the Availability Zone-specific endpoint service by using the endpoint's local DNS name.
E. Create a Savings Plan that provides adequate coverage for the organization's planned inter-Availability Zone data transfer usage.

Answer: CD

Explanation:
Cross-zone load balancing enables traffic to be distributed evenly across all registered instances in all enabled Availability Zones. However, this also increases data transfer charges between Availability Zones. By turning off cross-zone load balancing, the service provider applications can reduce inter-Availability Zone data transfer costs. Similarly, by using the Availability Zone-specific endpoint service, the service consumer applications can ensure that they connect to the nearest service provider application in the same Availability Zone, avoiding cross-Availability Zone data transfer charges. References:

> https://docs.aws.amazon.com/vpc/latest/userguide/vpce-interface.html#vpce-interface-dns

NEW QUESTION 141
- (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-l 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 145
- (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 146
- (Exam Topic 2)
A company has five development teams that have each created five AWS accounts to develop and host applications. To track spending, the development teams log in to each account every month, record the current cost from the AWS Billing and Cost Management console, and provide the information to the company's finance team.
The company has strict compliance requirements and needs to ensure that resources are created only in AWS Regions in the United States. However, some resources have been created in other Regions.
A solutions architect needs to implement a solution that gives the finance team the ability to track and consolidate expenditures for all the accounts. The solution

also must ensure that the company can create resources only in Regions in the United States.
Which combination of steps will meet these requirements in the MOST operationally efficient way? (Select THREE.)

A. Create a new account to serve as a management accoun
B. Create an Amazon S3 bucket for the finance learn Use AWS Cost and Usage Reports to create monthly reports and to store the data in the finance team's S3 bucket.
C. Create a new account to serve as a management accoun
D. Deploy an organization in AWS Organizations with all features enable
E. Invite all the existing accounts to the organizatio
F. Ensure that each account accepts the invitation.
G. Create an OU that includes all the development team
H. Create an SCP that allows the creation of resources only in Regions that are in the United State
I. Apply the SCP to the OU.
J. Create an OU that includes all the development team
K. Create an SCP that denies (he creation of resources in Regions that are outside the United State
L. Apply the SCP to the OU.
M. Create an 1AM role in the management account Attach a policy that includes permissions to view the Billing and Cost Management consol
N. Allow the finance learn users to assume the rol
O. Use AWS Cost Explorer and the Billing and Cost Management console to analyze cost.
P. Create an 1AM role in each AWS accoun
Q. Attach a policy that includes permissions to view the Billing and Cost Management consol
R. Allow the finance team users to assume the role.

Answer: BCE

Explanation:
AWS Organizations is a service that enables you to consolidate multiple AWS accounts into an organization that you create and centrally manage. By creating a management account and inviting all the existing accounts to join the organization, the solutions architect can track and consolidate expenditures for all the accounts using AWS Cost Management tools such as AWS Cost Explorer and AWS Budgets. An organizational unit (OU) is a group of accounts within an organization that can be used to apply policies and simplify management. A service control policy (SCP) is a type of policy that you can use to manage permissions in your organization. By creating an OU that includes all the development teams and applying an SCP that allows the creation of resources only in Regions that are in the United States, the solutions architect can ensure that the company meets its compliance requirements and avoids unwanted charges from other Regions. An IAM role is an identity with permission policies that determine what the identity can and cannot do in AWS. By creating an IAM role in the management account and allowing the finance team users to assume it, the solutions architect can give them access to view the Billing and Cost Management console without sharing credentials or creating additional users. References:

> https://docs.aws.amazon.com/organizations/latest/userguide/orgs_introduction.html

> https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scp.html

> https://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles.html

> https://docs.aws.amazon.com/aws-cost-management/latest/userguide/what-is-costmanagement.html

NEW QUESTION 147
......

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