Amazon-Web-Services

Exam Questions SAP-C02

AWS Certified Solutions Architect - Professional

NEW QUESTION 1
- (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 2
- (Exam Topic 1)
An application is using an Amazon RDS for MySQL Multi-AZ DB instance in the us-east-1 Region. After a failover test, the application lost the connections to the database and could not re-establish the connections. After a restart of the application, the application re-established the connections.
A solutions architect must implement a solution so that the application can re-establish connections to the database without requiring a restart.
Which solution will meet these requirements?

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

Answer: B

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

NEW QUESTION 3
- (Exam Topic 1)
A large mobile gaming company has successfully migrated all of its on-premises infrastructure to the AWS Cloud. A solutions architect is reviewing the environment to ensure that it was built according to the design and that it is running in alignment with the Well-Architected Framework.
While reviewing previous monthly costs in Cost Explorer, the solutions architect notices that the creation and subsequent termination of several large instance types account for a high proportion of the costs. The solutions architect finds out that the company's developers are launching new Amazon EC2 instances as part of their testing and that the developers are not using the appropriate instance types.
The solutions architect must implement a control mechanism to limit the instance types that only the developers can launch.
Which solution will meet these requirements?

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

Answer: C

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

NEW QUESTION 4
- (Exam Topic 1)
A solutions architect needs to advise a company on how to migrate its on-premises data processing application to the AWS Cloud. Currently, users upload input files through a web portal. The web server then stores the uploaded files on NAS and messages the processing server over a message queue. Each media file can take up to 1 hour to process. The company has determined that the number of media files awaiting processing is significantly higher during business hours, with the number of files rapidly declining after business hours.
What is the MOST cost-effective migration recommendation?

A. Create a queue using Amazon SQ
B. Configure the existing web server to publish to the new queue.When there are messages in the queue, invoke an AWS Lambda function to pull requests from the queue and process the file
C. Store the processed files in an Amazon S3 bucket.
D. Create a queue using Amazon
E. Configure the existing web server to publish to the new queu
F. When there are messages in the queue, create a new Amazon EC2 instance to pull requests from the queue and process the file
G. Store the processed files in Amazon EF
H. Shut down the EC2 instance after the task is complete.
I. Create a queue using Amazon M
J. Configure the existing web server to publish to the new queue.When there are messages in the queue, invoke an AWS Lambda function to pull requests from the queue and process the file
K. Store the processed files in Amazon EFS.
L. Create a queue using Amazon SO
M. Configure the existing web server to publish to the new queu
N. Use Amazon EC2 instances in an EC2 Auto Scaling group to pull requests from the queue and process the file
O. Scale the EC2 instances based on the SOS queue lengt
P. Store the processed files in an Amazon S3 bucket.

Answer: D

Explanation:
https://aws.amazon.com/blogs/compute/operating-lambda-performance-optimization-part-1/

NEW QUESTION 5
- (Exam Topic 1)
A company runs an loT platform on AWS loT sensors in various locations send data to the company's Node js API servers on Amazon EC2 instances running behind an Application Load Balancer The data is stored in an Amazon RDS MySQL DB instance that uses a 4 TB General Purpose SSD volume
The number of sensors the company has deployed in the field has increased over time and is expected to grow significantly The API servers are consistently overloaded and RDS metrics show high write latency
Which of the following steps together will resolve the issues permanently and enable growth as new sensors are provisioned, while keeping this platform cost-efficient? {Select TWO.)

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

Answer: CE

Explanation:
➢ Option C is correct because leveraging Amazon Kinesis Data Streams and AWS Lambda to ingest and process the raw data resolves the issues permanently and enable growth as new sensors are provisioned. Amazon Kinesis Data Streams is a serverless streaming data service that simplifies the capture, processing, and storage of data streams at any scale. Kinesis Data Streams can handle any amount of streaming data and process data from hundreds of thousands of sources with very low latency. AWS Lambda is a serverless compute service that lets you run code without provisioning or managing servers. Lambda can be triggered by Kinesis Data Streams events and process the data records in real time. Lambda can also scale automatically based on the incoming data volume. By using Kinesis Data Streams and Lambda, the company can reduce the load on the API servers and improve the performance and scalability of the data ingestion and processing layer3
➢ Option E is correct because re-architecting the database tier to use Amazon DynamoDB instead of an RDS MySQL DB instance resolves the issues permanently and enable growth as new sensors are provisioned. Amazon DynamoDB is a fully managed key-value and document database that delivers single-digit millisecond performance at any scale. DynamoDB supports auto scaling, which automatically adjusts read and write capacity based on actual traffic patterns. DynamoDB also supports on-demand capacity mode, which instantly accommodates up to double the previous peak traffic on a table. By using DynamoDB instead of RDS MySQL DB instance, the company can eliminate high write latency and improve scalability and performance of the database tier.
References: 1: https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-volume-types.html 2:
https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/CHAP_AuroraOverview.html 3:
https://docs.aws.amazon.com/streams/latest/dev/introduction.html : https://docs.aws.amazon.com/lambda/latest/dg/welcome.html :
https://docs.aws.amazon.com/xray/latest/devguide/aws-xray.html : https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Introduction.html :

NEW QUESTION 6
- (Exam Topic 1)
An international delivery company hosts a delivery management system on AWS. Drivers use the system to upload confirmation of delivery. Confirmation includes the recipient's signature or a photo of the package with the recipient. The driver's handheld device uploads signatures and photos through FTP to a single Amazon EC2 instance. Each handheld device saves a file in a directory based on the signed-in user, and the file name matches the delivery number. The EC2 instance then adds metadata to the file after querying a central database to pull delivery information. The file is then placed in Amazon S3 for archiving.
As the company expands, drivers report that the system is rejecting connections. The FTP server is having problems because of dropped connections and memory issues. In response to these problems, a system engineer schedules a cron task to reboot the EC2 instance every 30 minutes. The billing team reports

that files are not always in the archive and that the central system is not always updated.
A solutions architect needs to design a solution that maximizes scalability to ensure that the archive always receives the files and that systems are always updated. The handheld devices cannot be modified, so the company cannot deploy a new application.
Which solution will meet these requirements?

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

Answer: C

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

NEW QUESTION 7
- (Exam Topic 1)
A company has migrated an application from on premises to AWS. The application frontend is a static website that runs on two Amazon EC2 instances behind an Application Load Balancer (ALB). The application backend is a Python application that runs on three EC2 instances behind another ALB. The EC2 instances are large, general purpose On-Demand Instances that were sized to meet the on-premises specifications for peak usage of the application.
The application averages hundreds of thousands of requests each month. However, the application is used mainly during lunchtime and receives minimal traffic during the rest of the day.
A solutions architect needs to optimize the infrastructure cost of the application without negatively affecting the application availability.
Which combination of steps will meet these requirements? (Choose two.)

A. Change all the EC2 instances to compute optimized instances that have the same number of cores as the existing EC2 instances.
B. Move the application frontend to a static website that is hosted on Amazon S3.
C. Deploy the application frontend by using AWS Elastic Beanstal
D. Use the same instance type for the nodes.
E. Change all the backend EC2 instances to Spot Instances.
F. Deploy the backend Python application to general purpose burstable EC2 instances that have the same number of cores as the existing EC2 instances.

Answer: BD

Explanation:
Moving the application frontend to a static website that is hosted on Amazon S3 will save cost as S3 is cheaper than running EC2 instances.
Using Spot instances for the backend EC2 instances will also save cost, as they are significantly cheaper than On-Demand instances. This will be suitable for the application, as it has minimal traffic during the rest of the day, and the availability of spot instances will not negatively affect the application's availability.
Reference:
Amazon S3 pricing: https://aws.amazon.com/s3/pricing/
Amazon EC2 Spot Instances documentation: https://aws.amazon.com/ec2/spot/ AWS Elastic Beanstalk documentation: https://aws.amazon.com/elasticbeanstalk/
Amazon Elastic Compute Cloud (EC2) pricing: https://aws.amazon.com/ec2/pricing/

NEW QUESTION 8
- (Exam Topic 1)
A company is hosting a three-tier web application in an on-premises environment. Due to a recent surge in traffic that resulted in downtime and a significant financial impact, company management has ordered that the application be moved to AWS. The application is written in .NET and has a dependency on a MySQL database A solutions architect must design a scalable and highly available solution to meet the demand of 200000 daily users.
Which steps should the solutions architect take to design an appropriate solution?

A. Use AWS Elastic Beanstalk to create a new application with a web server environment and an Amazon RDS MySQL Multi-AZ DB instance The environment should launch a Network Load Balancer (NLB) in front of an Amazon EC2 Auto Scaling group in multiple Availability Zones Use an Amazon Route 53 alias record to route traffic from the company's domain to the NLB.
B. Use AWS CloudFormation to launch a stack containing an Application Load Balancer (ALB) in front of an Amazon EC2 Auto Scaling group spanning three Availability Zone
C. The stack should launch a Multi-AZ deployment of an Amazon Aurora MySQL DB cluster with a Retain deletion polic
D. Use an Amazon Route 53 alias record to route traffic from the company's domain to the ALB
E. Use AWS Elastic Beanstalk to create an automatically scaling web server environment that spans two separate Regions with an Application Load Balancer (ALB) in each Regio
F. Create a Multi-AZ deployment of an Amazon Aurora MySQL DB cluster with a cross-Region read replica Use Amazon Route 53 with a geoproximity routing policy to route traffic between the two Regions.
G. Use AWS CloudFormation to launch a stack containing an Application Load Balancer (ALB) in front of an Amazon ECS cluster of Spot Instances spanning three Availability Zones The stack should launch an Amazon RDS MySQL DB instance with a Snapshot deletion policy Use an Amazon Route 53 alias record to route traffic from the company's domain to the ALB

Answer: C

Explanation:
Using AWS CloudFormation to launch a stack with an Application Load Balancer (ALB) in front of an Amazon EC2 Auto Scaling group spanning three Availability Zones, a Multi-AZ deployment of an Amazon Aurora MySQL DB cluster with a Retain deletion policy, and an Amazon Route 53 alias record to route traffic from the company’s domain to the ALB will ensure that

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 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 10
- (Exam Topic 1)
A company recently deployed an application on AWS. The application uses Amazon DynamoDB. The company measured the application load and configured the RCUs and WCUs on the DynamoDB table to
match the expected peak load. The peak load occurs once a week for a 4-hour period and is double the average load. The application load is close to the average load tor the rest of the week. The access pattern includes many more writes to the table than reads of the table.
A solutions architect needs to implement a solution to minimize the cost of the table. Which solution will meet these requirements?

A. Use AWS Application Auto Scaling to increase capacity during the peak perio
B. Purchase reserved RCUs and WCUs to match the average load.
C. Configure on-demand capacity mode for the table.
D. Configure DynamoDB Accelerator (DAX) in front of the tabl
E. Reduce the provisioned read capacity to match the new peak load on the table.
F. Configure DynamoDB Accelerator (DAX) in front of the tabl
G. Configure on-demand capacity mode for the table.

Answer: D

Explanation:
This solution meets the requirements by using Application Auto Scaling to automatically increase capacity during the peak period, which will handle the double the average load. And by purchasing reserved RCUs and WCUs to match the average load, it will minimize the cost of the table for the rest of the week when the load is close to the average.

NEW QUESTION 13
- (Exam Topic 1)
A company is running applications on AWS in a multi-account environment. The company's sales team and marketing team use separate AWS accounts in AWS Organizations.
The sales team stores petabytes of data in an Amazon S3 bucket. The marketing team uses Amazon QuickSight for data visualizations. The marketing team needs access to data that the sates team stores in the S3 bucket. The company has encrypted the S3 bucket with an AWS Key Management Service (AWS KMS) key. The marketing team has already created the IAM service role for QuickSight to provide QuickSight access in the marketing AWS account. The company needs a solution that will provide secure access to the data in the S3 bucket across AWS accounts.
Which solution will meet these requirements with the LEAST operational overhead?

A. Create a new S3 bucket in the marketing accoun

B. Create an S3 replication rule in the sales account to copy the objects to the new S3 bucket in the marketing accoun
C. Update the QuickSight permissions in the marketing account to grant access to the new S3 bucket.
D. Create an SCP to grant access to the S3 bucket to the marketing accoun
E. Use AWS Resource Access Manager (AWS RAM) to share the KMS key from the sates account with the marketing accoun
F. Update the QuickSight permissions in the marketing account to grant access to the S3 bucket.
G. Update the S3 bucket policy in the marketing account to grant access to the QuickSight rol
H. Create a KMS grant for the encryption key that is used in the S3 bucke
I. Grant decrypt access to the QuickSight rol
J. Update the QuickSight permissions in the marketing account to grant access to the S3 bucket.
K. Create an IAM role in the sales account and grant access to the S3 bucke
L. From the marketing account, assume the IAM role in the sales account to access the S3 bucke
M. Update the QuickSight rote, to create a trust relationship with the new IAM role in the sales account.

Answer: D

Explanation:
Create an IAM role in the sales account and grant access to the S3 bucket. From the marketing account, assume the IAM role in the sales account to access the S3 bucket. Update the QuickSight role, to create a trust relationship with the new IAM role in the sales account.
This approach is the most secure way to grant cross-account access to the data in the S3 bucket while minimizing operational overhead. By creating an IAM role in the sales account, the marketing team can assume the role in their own account, and have access to the S3 bucket. And updating the QuickSight role, to create a trust relationship with the new IAM role in the sales account will grant the marketing team to access the data in the S3 bucket and use it for data visualization using QuickSight.
AWS Resource Access Manager (AWS RAM) also allows sharing of resources between accounts, but it would require additional management and configuration to set up the sharing, which would increase operational overhead.
Using S3 replication would also replicate the data to the marketing account, but it would not provide the marketing team access to the original data, and also it would increase operational overhead with managing the replication process.
IAM roles and policies, KMS grants and trust relationships are a powerful combination for managing
cross-account access in a secure and efficient manner. References:

- AWS IAM Roles
- AWS KMS - Key Grants
- AWS RAM

NEW QUESTION 15
- (Exam Topic 1)
A company is running an application in the AWS Cloud. The application runs on containers in an Amazon Elastic Container Service (Amazon ECS) cluster. The ECS tasks use the Fargate launch type. The application's data is relational and is stored in Amazon Aurora MySQL. To meet regulatory requirements, the application must be able to recover to a separate AWS Region in the event of an application failure. In case of a failure, no data can be lost. Which solution will meet these requirements with the LEAST amount of operational overhead?

A. Provision an Aurora Replica in a different Region.
B. Set up AWS DataSync for continuous replication of the data to a different Region.
C. Set up AWS Database Migration Service (AWS DMS) to perform a continuous replication of the data to a different Region.
D. Use Amazon Data Lifecycle Manager {Amazon DLM) to schedule a snapshot every 5 minutes.

Answer: A

Explanation:
Provision an Aurora Replica in a different Region will meet the requirement of the application being able to recover to a separate AWS Region in the event of an application failure, and no data can be lost, with the least amount of operational overhead.

NEW QUESTION 19
- (Exam Topic 1)
A company is running an application on several Amazon EC2 instances in an Auto Scaling group behind an Application Load Balancer. The load on the application varies throughout the day, and EC2 instances are scaled in and out on a regular basis. Log files from the EC2 instances are copied to a central Amazon S3 bucket every 15 minutes. The security team discovers that log files are missing from some of the terminated EC2 instances.
Which set of actions will ensure that log files are copied to the central S3 bucket from the terminated EC2 instances?

A. Create a script to copy log files to Amazon S3, and store the script in a file on the EC2 instanc
B. Create an Auto Scaling lifecycle hook and an Amazon EventBridge (Amazon CloudWatch Events) rule to detect lifecycle events from the Auto Scaling grou
C. Invoke an AWS Lambda function on the autoscaling:EC2_INSTANCE_TERMINATING transition to send ABANDON to the Auto Scaling group to prevent termination, run the script to copy the log files, and terminate the instance using the AWS SDK.
D. Create an AWS Systems Manager document with a script to copy log files to Amazon S3. Create an Auto Scaling lifecycle hook and an Amazon EventBridge (Amazon CloudWatch Events) rule to detect lifecycle events from the Auto Scaling grou
E. Invoke an AWS Lambda function on the autoscaling:EC2_INSTANCE_TERMINATING transition to call the AWS Systems Manager API SendCommand operation to run the document to copy the log files and send CONTINUE to the Auto Scaling group to terminate the instance.
F. Change the log delivery rate to every 5 minute
G. Create a script to copy log files to Amazon S3, and add the script to EC2 instance user dat
H. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to detect EC2 instance terminatio
I. Invoke an AWS Lambda function from the EventBridge (CloudWatch Events) rule that uses the AWS CLI to run the user-data script to copy the log files and terminate the instance.
J. Create an AWS Systems Manager document with a script to copy log files to Amazon S3. Create an Auto Scaling lifecycle hook that publishes a message to an Amazon Simple Notification Service (Amazon SNS) topi
K. From the SNS notification, call the AWS Systems Manager API SendCommand operation to run the document to copy the log files and send ABANDON to the Auto Scaling group to terminate the instance.

Answer: B

Explanation:
https://docs.aws.amazon.com/autoscaling/ec2/userguide/adding-lifecycle-hooks.html
- Refer to Default Result section - If the instance is terminating, both abandon and continue allow the instance to terminate. However, abandon stops any

remaining actions, such as other lifecycle hooks, and continue allows any other lifecycle hooks to complete.
https://aws.amazon.com/blogs/infrastructure-and-automation/run-code-before-terminating-an-ec2-auto-scaling-i https://github.com/aws-samples/aws-lambda-lifecycle-hooks-function
https://github.com/aws-samples/aws-lambda-lifecycle-hooks-function/blob/master/cloudformation/template.yam

NEW QUESTION 24
- (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 26
- (Exam Topic 1)
A life sciences company is using a combination of open source tools to manage data analysis workflows and Docker containers running on servers in its on-premises data center to process genomics data Sequencing data is generated and stored on a local storage area network (SAN), and then the data is processed. The research and development teams are running into capacity issues and have decided to re-architect their genomics analysis platform on AWS to scale based on workload demands and reduce the turnaround time from weeks to days
The company has a high-speed AWS Direct Connect connection Sequencers will generate around 200 GB of data for each genome, and individual jobs can take several hours to process the data with ideal compute capacity. The end result will be stored in Amazon S3. The company is expecting 10-15 job requests each day
Which solution meets these requirements?

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

Answer: C

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

NEW QUESTION 29
- (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 31
- (Exam Topic 1)
A company has developed APIs that use Amazon API Gateway with Regional endpoints. The APIs call AWS Lambda functions that use API Gateway authentication mechanisms. After a design review, a solutions architect identifies a set of APIs that do not require public access.
The solutions architect must design a solution to make the set of APIs accessible only from a VPC. All APIs need to be called with an authenticated user.
Which solution will meet these requirements with the LEAST amount of effort?

A. Create an internal Application Load Balancer (ALB). Create a target grou
B. Select the Lambda function to cal
C. Use the ALB DNS name to call the API from the VPC.
D. Remove the DNS entry that is associated with the API in API Gatewa
E. Create a hosted zone in Amazon Route 53. Create a CNAME record in the hosted zon
F. Update the API in API Gateway with the CNAME recor
G. Use the CNAME record to call the API from the VPC.
H. Update the API endpoint from Regional to private in API Gatewa
I. Create an interface VPC endpoint in the VP
J. Create a resource policy, and attach it to the AP
K. Use the VPC endpoint to call the API from the VPC.
L. Deploy the Lambda functions inside the VP
M. Provision an EC2 instance, and install an Apache server.From the Apache server, call the Lambda function
N. Use the internal CNAME record of the EC2 instance to call the API from the VPC.

Answer: C

Explanation:
This solution requires the least amount of effort as it only requires to update the API endpoint to private in API Gateway and create an interface VPC endpoint. Then create a resource policy and attach it to the API. This will make the API only accessible from the VPC and still keep the authentication mechanism intact.
Reference:
- https://aws.amazon.com/api-gateway/features/

NEW QUESTION 34
- (Exam Topic 1)
A company has created an OU in AWS Organizations for each of its engineering teams Each OU owns multiple AWS accounts. The organization has hundreds of AWS accounts A solutions architect must design a solution so that each OU can view a breakdown of usage costs across its AWS accounts. Which solution meets these requirements?

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

Answer: B

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

NEW QUESTION 38
- (Exam Topic 1)
A security engineer determined that an existing application retrieves credentials to an Amazon RDS for MySQL database from an encrypted file in Amazon S3. For the next version of the application, the security engineer wants to implement the following application design changes to improve security:
- The database must use strong, randomly generated passwords stored in a secure AWS managed service.
- The application resources must be deployed through AWS CloudFormation.
- The application must rotate credentials for the database every 90 days.

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

A. Generate the database password as a secret resource using AWS Secrets Manage

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

Answer: B

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

NEW QUESTION 39
- (Exam Topic 1)
A company gives users the ability to upload images from a custom application. The upload process invokes an AWS Lambda function that processes and stores the image in an Amazon S3 bucket. The application invokes the Lambda function by using a specific function version ARN.
The Lambda function accepts image processing parameters by using environment variables. The company often adjusts the environment variables of the Lambda function to achieve optimal image processing output. The company tests different parameters and publishes a new function version with the updated environment variables after validating results. This update process also requires frequent changes to the custom application to invoke the new function version ARN. These changes cause interruptions for users.
A solutions architect needs to simplify this process to minimize disruption to users. Which solution will meet these requirements with the LEAST operational overhead?

A. Directly modify the environment variables of the published Lambda function versio
B. Use theSLATEST version to test image processing parameters.
C. Create an Amazon DynamoDB table to store the image processing parameter
D. Modify the Lambda function to retrieve the image processing parameters from the DynamoDB table.
E. Directly code the image processing parameters within the Lambda function and remove the environment variable
F. Publish a new function version when the company updates the parameters.
G. Create a Lambda function alia
H. Modify the client application to use the function alias AR
I. Reconfigure the Lambda alias to point to new versions of the function when the company finishes testing.

Answer: D

Explanation:
A Lambda function alias allows you to point to a specific version of a function and also can be updated to point to a new version of the function without modifying the client application. This way, the company can test different versions of the function with different environment variables and, once the optimal parameters are found, update the alias to point to the new version, without the need to update the client application.
By using this approach, the company can simplify the process of updating the environment variables, minimize disruption to users, and reduce the operational overhead.
Reference:
AWS Lambda documentation: https://aws.amazon.com/lambda/
AWS Lambda Aliases documentation: https://docs.aws.amazon.com/lambda/latest/dg/aliases-intro.html AWS Lambda versioning and aliases documentation: https://aws.amazon.com/blogs/compute/versioning-aliases-in-aws-lambda/

NEW QUESTION 44
- (Exam Topic 1)
A company uses Amazon S3 to store files and images in a variety of storage classes. The company's S3 costs have increased substantially during the past year.
A solutions architect needs to review data trends for the past 12 months and identity the appropriate storage class for the objects.
Which solution will meet these requirements?

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

Answer: B

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

NEW QUESTION 48
- (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 52
- (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 55
- (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 58

- (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 63
- (Exam Topic 1)
A company runs a Java application that has complex dependencies on VMs that are in the company's data center. The application is stable. but the company wants to modernize the technology stack. The company wants to migrate the application to AWS and minimize the administrative overhead to maintain the servers.
Which solution will meet these requirements with the LEAST code changes?

A. Migrate the application to Amazon Elastic Container Service (Amazon ECS) on AWS Fargate by using AWS App2Containe
B. Store container images in Amazon Elastic Container Registry (Amazon ECR). Grant the ECS task execution role permission 10 access the ECR image repositor
C. Configure Amazon ECS to use an Application Load Balancer (ALB). Use the ALB to interact with the application.
D. Migrate the application code to a container that runs in AWS Lambd
E. Build an Amazon API Gateway REST API with Lambda integratio
F. Use API Gateway to interact with the application.
G. Migrate the application to Amazon Elastic Kubernetes Service (Amazon EKS) on EKS managed node groups by using AWS App2Containe
H. Store container images in Amazon Elastic Container Registry (Amazon ECR). Give the EKS nodes permission to access the ECR image repositor
I. Use Amazon API Gateway to interact with the application.
J. Migrate the application code to a container that runs in AWS Lambd
K. Configure Lambda to use an Application Load Balancer (ALB). Use the ALB to interact with the application.

Answer: A

Explanation:
According to the AWS documentation1, AWS App2Container (A2C) is a command line tool for migrating and modernizing Java and .NET web applications into container format. AWS A2C analyzes and builds an inventory of applications running in bare metal, virtual machines, Amazon Elastic Compute Cloud (EC2) instances, or in the cloud. You can use AWS A2C to generate container images for your applications and deploy them on Amazon ECS or Amazon EKS.
Option A meets the requirements of the scenario because it allows you to migrate your existing Java application to AWS and minimize the administrative overhead to maintain the servers. You can use AWS A2C to analyze your application dependencies, extract application artifacts, and generate a Dockerfile. You can then store your container images in Amazon ECR, which is a fully managed container registry service. You can use AWS Fargate as the launch type for your Amazon ECS cluster, which is a serverless compute engine that eliminates the need to provision and manage servers for your containers. You can grant the ECS task execution role permission to access the ECR image repository, which allows your tasks to pull images from ECR. You can configure Amazon ECS to use an ALB, which is a load balancer that distributes traffic across multiple targets in multiple Availability Zones using HTTP or HTTPS protocols. You can use the ALB to interact with your application.

NEW QUESTION 64
- (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 68
- (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:

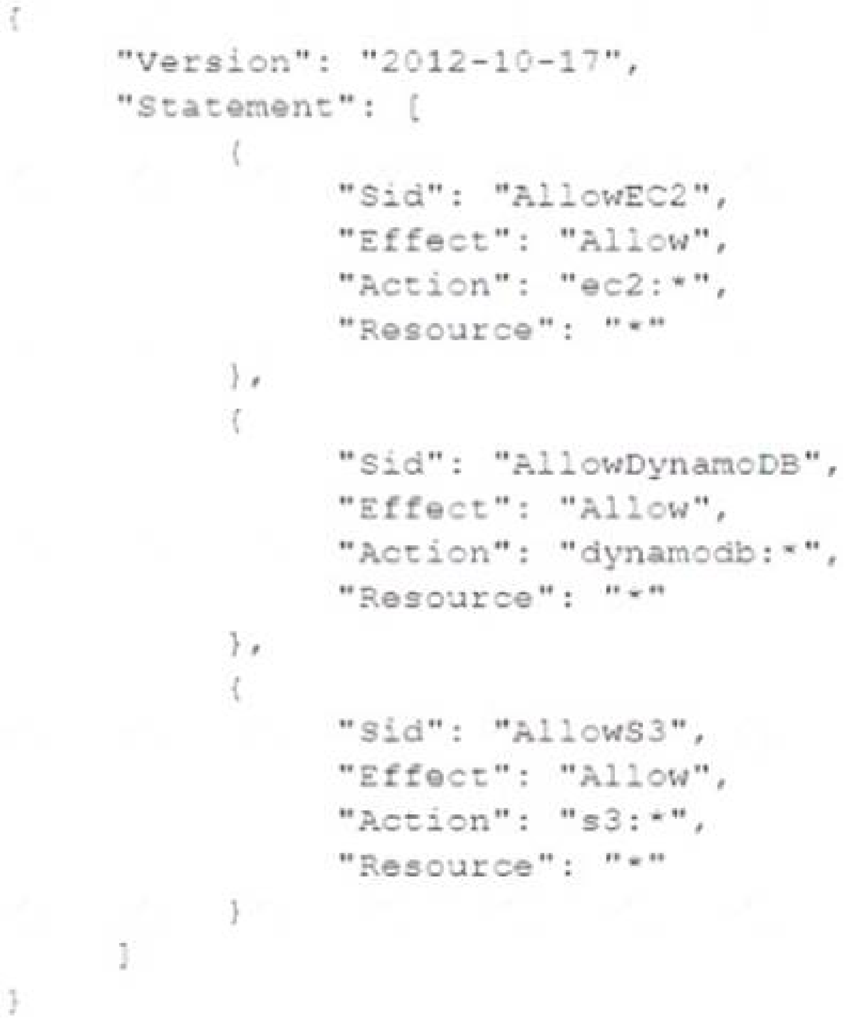

```
{
      "Version": "2012-10-17",
      "Statement": [
            {
                  "Sid": "AllowEC2",
                  "Effect": "Allow",
                  "Action": "ec2:*",
                  "Resource": "*"
            },
            {
                  "Sid": "AllowDynamoDB",
                  "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 72
- (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 76
- (Exam Topic 1)
A company hosts a Git repository in an on-premises data center. The company uses webhooks to invoke functionality that runs in the AWS Cloud. The company hosts the webhook logic on a set of Amazon EC2 instances in an Auto Scaling group that the company set as a target for an Application Load Balancer (ALB). The Git server calls the ALB for the configured webhooks. The company wants to move the solution to a serverless architecture.
Which solution will meet these requirements with the LEAST operational overhead?

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

Answer: B

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

NEW QUESTION 81
- (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 82

- (Exam Topic 1)
A company with global offices has a single 1 Gbps AWS Direct Connect connection to a single AWS Region. The company's on-premises network uses the connection to communicate with the company's resources in the AWS Cloud. The connection has a single private virtual interface that connects to a single VPC.
A solutions architect must implement a solution that adds a redundant Direct Connect connection in the same Region. The solution also must provide connectivity to other Regions through the same pair of Direct Connect connections as the company expands into other Regions.
Which solution meets these requirements?

A. Provision a Direct Connect gatewa
B. Delete the existing private virtual interface from the existing connectio
C. Create the second Direct Connect connectio
D. Create a new private virtual interlace on each connection, and connect both private victual interfaces to the Direct Connect gatewa
E. Connect the Direct Connect gateway to the single VPC.
F. Keep the existing private virtual interfac
G. Create the second Direct Connect connectio
H. Create a new private virtual interface on the new connection, and connect the new private virtual interface to the single VPC.
I. Keep the existing private virtual interfac
J. Create the second Direct Connect connectio
K. Create a new public virtual interface on the new connection, and connect the new public virtual interface to the single VPC.
L. Provision a transit gatewa
M. Delete the existing private virtual interface from the existing connection.Create the second Direct Connect connectio
N. Create a new private virtual interface on each connection, and connect both private virtual interfaces to the transit gatewa
O. Associate the transit gateway with the single VPC.

Answer: A

Explanation:
A Direct Connect gateway is a globally available resource. You can create the Direct Connect gateway in any Region and access it from all other Regions. The following describe scenarios where you can use a Direct Connect gateway.
https://docs.aws.amazon.com/directconnect/latest/UserGuide/direct-connect-gateways-intro.html

NEW QUESTION 83
- (Exam Topic 1)
A company has its cloud infrastructure on AWS A solutions architect needs to define the infrastructure as code. The infrastructure is currently deployed in one AWS Region. The company's business expansion plan includes deployments in multiple Regions across multiple AWS accounts
What should the solutions architect do to meet these requirements?

A. Use AWS CloudFormation templates Add IAM policies to control the various accounts Deploy the templates across the multiple Regions
B. Use AWS Organizations Deploy AWS CloudFormation templates from the management account Use AWS Control Tower to manage deployments across accounts
C. Use AWS Organizations and AWS CloudFormation StackSets Deploy a CloudFormation template from an account that has the necessary IAM permissions
D. Use nested stacks with AWS CloudFormation templates Change the Region by using nested stacks

Answer: C

Explanation:
https://aws.amazon.com/blogs/aws/new-use-aws-cloudformation-stacksets-for-multiple-accounts-in-an-aws-orga AWS Organizations allows the management of multiple AWS accounts as a single entity and AWS
CloudFormation StackSets allows creating, updating, and deleting stacks across multiple accounts and regions in an organization. This solution allows creating a single CloudFormation template that can be deployed across multiple accounts and regions, and also allows for the management of access and permissions for the different accounts through the use of IAM roles and policies in the management account.

NEW QUESTION 85
- (Exam Topic 1)
A solutions architect needs to copy data from an Amazon S3 bucket m an AWS account to a new S3 bucket in a new AWS account. The solutions architect must implement a solution that uses the AWS CLI.
Which combination of steps will successfully copy the data? (Choose three.)

A. Create a bucket policy to allow the source bucket to list its contents and to put objects and set object ACLs in the destination bucke
B. Attach the bucket policy to the destination bucket.
C. Create a bucket policy to allow a user In the destination account to list the source bucket's contents and read the source bucket's object
D. Attach the bucket policy to the source bucket.
E. Create an IAM policy in the source accoun
F. Configure the policy to allow a user In the source account to list contents and get objects In the source bucket, and to list contents, put objects, and set object ACLs in the destination bucke
G. Attach the policy to the user _
H. Create an IAM policy in the destination accoun
I. Configure the policy to allow a user In the destination account to list contents and get objects In the source bucket, and to list contents, put objects, and set objectACLs in the destination bucke
J. Attach the policy to the user.
K. Run the aws s3 sync command as a user in the source accoun
L. Specify' the source and destination buckets to copy the data.
M. Run the aws s3 sync command as a user in the destination accoun
N. Specify' the source and destination buckets to copy the data.

Answer: BDF

Explanation:
Step B is necessary so that the user in the destination account has the necessary permissions to access the source bucket and list its contents, read its objects. Step D is needed so that the user in the destination account has the necessary permissions to access the destination bucket and list contents, put objects, and set object ACLs Step F is necessary because the aws s3 sync command needs to be run using the IAM user credentials from the destination account, so that the objects will have the appropriate permissions for the user in the destination account once they are copied.

NEW QUESTION 88
- (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 90
- (Exam Topic 1)
A company is processing videos in the AWS Cloud by using Amazon EC2 instances in an Auto Scaling group. It takes 30 minutes to process a video. Several EC2 instances scale in and out depending on the number of videos in an Amazon Simple Queue Service (Amazon SQS) queue.
The company has configured the SQS queue with a redrive policy that specifies a target dead-letter queue and a maxReceiveCount of 1. The company has set the visibility timeout for the SQS queue to 1 hour. The company has set up an Amazon CloudWatch alarm to notify the development team when there are messages in the dead-letter queue.
Several times during the day, the development team receives notification that messages are in the dead-letter queue and that videos have not been processed properly. An investigation finds no errors in the application logs.
How can the company solve this problem?

A. Turn on termination protection for the EC2 instances.
B. Update the visibility timeout for the SOS queue to 3 hours.
C. Configure scale-in protection for the instances during processing.
D. Update the redrive policy and set maxReceiveCount to 0.

Answer: B

Explanation:
The best solution for this problem is to update the visibility timeout for the SQS queue to 3 hours. This is because when the visibility timeout is set to 1 hour, it means that if the EC2 instance doesn't process the message within an hour, it will be moved to the dead-letter queue. By increasing the visibility timeout to 3 hours, this should give the EC2 instance enough time to process the message before it gets moved to the dead-letter queue. Additionally, configuring scale-in protection for the EC2 instances during processing will help to ensure that the instances are not terminated while the messages are being processed.

NEW QUESTION 94
- (Exam Topic 1)
A solutions architect needs to implement a client-side encryption mechanism for objects that will be stored in a new Amazon S3 bucket. The solutions architect created a CMK that is stored in AWS Key Management Service (AWS KMS) for this purpose.
The solutions architect created the following IAM policy and attached it to an IAM role:

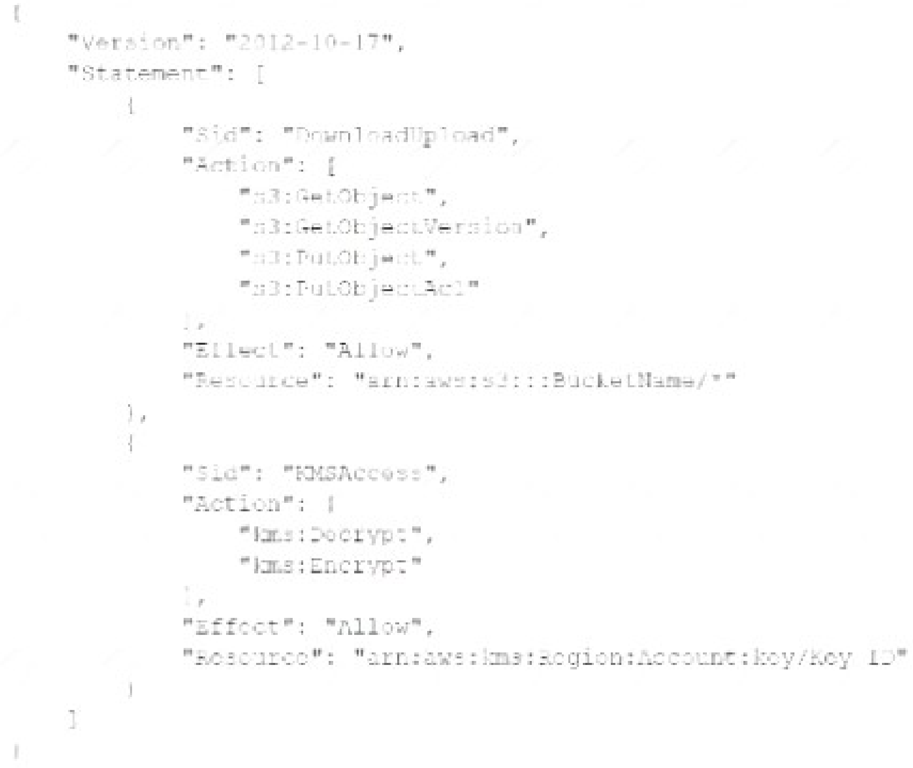

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "DownloadUpload",
            "Action": [
                "s3:GetObject",
                "s3:GetObjectVersion",
                "s3:PutObject",
                "s3:PutObjectAcl"
            ],
            "Effect": "Allow",
            "Resource": "arn:aws:s3:::BucketName/*"
        },
        {
            "Sid": "KMSAccess",
            "Action": [
                "kms:Decrypt",
                "kms:Encrypt"
            ],
            "Effect": "Allow",
            "Resource": "arn:aws:kms:Region:Account:key/Key ID"
        }
    ]
}
```

During tests, me solutions architect was able to successfully get existing test objects m the S3 bucket However, attempts to upload a new object resulted in an error message. The error message stated that me action was forbidden.
Which action must me solutions architect add to the IAM policy to meet all the requirements?

A. Kms:GenerateDataKey
B. KmsGetKeyPolpcy
C. kmsGetPubKKey
D. kms:SKjn

Answer: A

Explanation:
https://aws.amazon.com/premiumsupport/knowledge-center/s3-access-denied-error-kms/
"An error occurred (AccessDenied) when calling the PutObject operation: Access Denied" This error message indicates that your IAM user or role needs permission for the kms:GenerateDataKey action.

NEW QUESTION 96
- (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 98
- (Exam Topic 1)
A company is building an electronic document management system in which users upload their documents. The application stack is entirely serverless and runs on AWS in the eu-central-1 Region. The system includes a web application that uses an Amazon CloudFront distribution for delivery with Amazon S3 as the origin. The web application communicates with Amazon API Gateway Regional endpoints. The API Gateway APIs call AWS Lambda functions that store metadata in an Amazon Aurora Serverless database and put the documents into an S3 bucket.
The company is growing steadily and has completed a proof of concept with its largest customer. The company must improve latency outside of Europe.
Which combination of actions will meet these requirements? (Select TWO.)

A. Enable S3 Transfer Acceleration on the S3 bucke
B. Ensure that the web application uses the Transfer Acceleration signed URLs.
C. Create an accelerator in AWS Global Accelerato
D. Attach the accelerator to the CloudFront distribution.
E. Change the API Gateway Regional endpoints to edge-optimized endpoints.
F. Provision the entire stack in two other locations that are spread across the worl
G. Use global databases on the Aurora Serverless cluster.
H. Add an Amazon RDS proxy between the Lambda functions and the Aurora Serverless database.

Answer: AC

Explanation:
https://aws.amazon.com/global-accelerator/faqs/

NEW QUESTION 103
- (Exam Topic 1)
A company is running several workloads in a single AWS account. A new company policy states that engineers can provision only approved resources and that engineers must use AWS CloudFormation to provision these resources. A solutions architect needs to create a solution to enforce the new restriction on the IAM role that the engineers use for access.
What should the solutions architect do to create the solution?

A. Upload AWS CloudFormation templates that contain approved resources to an Amazon S3 bucket.Update the IAM policy for the engineers' IAM role to only allow access to Amazon S3 and AWS CloudFormatio

B. Use AWS CloudFormation templates to provision resources.
C. Update the IAM policy for the engineers' IAM role with permissions to only allow provisioning of approved resources and AWS CloudFormatio
D. Use AWS CloudFormation templates to create stacks with approved resources.
E. Update the IAM policy for the engineers' IAM role with permissions to only allow AWS CloudFormation action
F. Create a new IAM policy with permission to provision approved resources, and assign the policy to a new IAM service rol
G. Assign the IAM service role to AWS CloudFormation during stack creation.
H. Provision resources in AWS CloudFormation stack
I. Update the IAM policy for the engineers' IAM role to only allow access to their own AWS CloudFormation stack.

Answer: B

Explanation:
https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/security-best-practices.html#use-iam-to-c
https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/using-iam-servicerole.html

NEW QUESTION 105
- (Exam Topic 1)
A digital marketing company has multiple AWS accounts that belong to various teams. The creative team uses an Amazon S3 bucket in its AWS account to securely store images and media files that are used as content for the company's marketing campaigns. The creative team wants to share the S3 bucket with the strategy team so that the strategy team can view the objects.
A solutions architect has created an IAM role that is named strategy_reviewer in the Strategy account. The solutions architect also has set up a custom AWS Key Management Service (AWS KMS) key in the Creative account and has associated the key with the S3 bucket. However, when users from the Strategy account assume the IAM role and try to access objects in the S3 bucket, they receive an Account.
The solutions architect must ensure that users in the Strategy account can access the S3 bucket. The solution must provide these users with only the minimum permissions that they need.
Which combination of steps should the solutions architect take to meet these requirements? (Select THREE.)

A. Create a bucket policy that includes read permissions for the S3 bucke
B. Set the principal of the bucket policy to the account ID of the Strategy account
C. Update the strategy_reviewer IAM role to grant full permissions for the S3 bucket and to grant decrypt permissions for the custom KMS key.
D. Update the custom KMS key policy in the Creative account to grant decrypt permissions to the strategy_reviewer IAM role.
E. Create a bucket policy that includes read permissions for the S3 bucke
F. Set the principal of the bucket policy to an anonymous user.
G. Update the custom KMS key policy in the Creative account to grant encrypt permissions to the strategy_reviewer IAM role.
H. Update the strategy_reviewer IAM role to grant read permissions for the S3 bucket and to grant decrypt permissions for the custom KMS key

Answer: ACF

Explanation:
https://aws.amazon.com/premiumsupport/knowledge-center/cross-account-access-denied-error-s3/

NEW QUESTION 109
- (Exam Topic 1)
A company needs to implement a patching process for its servers. The on-premises servers and Amazon EC2 instances use a variety of tools to perform patching. Management requires a single report showing the patch status of all the servers and instances.
Which set of actions should a solutions architect take to meet these requirements?

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

Answer: A

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

NEW QUESTION 113
- (Exam Topic 1)
A weather service provides high-resolution weather maps from a web application hosted on AWS in the
eu-west-1 Region. The weather maps are updated frequently and stored in Amazon S3 along with static HTML content. The web application is fronted by Amazon CloudFront.
The company recently expanded to serve users in the us-east-1 Region, and these new users report that viewing their respective weather maps is slow from time to time.
Which combination of steps will resolve the us-east-1 performance issues? (Choose two.)

A. Configure the AWS Global Accelerator endpoint for the S3 bucket in eu-west-1. Configure endpoint groups for TCP ports 80 and 443 in us-east-1.
B. Create a new S3 bucket in us-east-1. Configure S3 cross-Region replication to synchronize from the S3 bucket in eu-west-1.
C. Use Lambda@Edge to modify requests from North America to use the S3 Transfer Acceleration endpoint in us-east-1.
D. Use Lambda@Edge to modify requests from North America to use the S3 bucket in us-east-1.
E. Configure the AWS Global Accelerator endpoint for us-east-1 as an origin on the CloudFront distributio
F. Use Lambda@Edge to modify requests from North America to use the new origin.

Answer: BD

Explanation:

https://aws.amazon.com/about-aws/whats-new/2016/04/transfer-files-into-amazon-s3-up-to-300-percent-faster/

NEW QUESTION 116
- (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 118
- (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 119
- (Exam Topic 1)
A company has an environment that has a single AWS account. A solutions architect is reviewing the environment to recommend what the company could improve specifically in terms of access to the AWS Management Console. The company's IT support workers currently access the console for administrative tasks, authenticating with named IAM users that have been mapped to their job role.
The IT support workers no longer want to maintain both their Active Directory and IAM user accounts. They want to be able to access the console by using their existing Active Directory credentials. The solutions architect is using AWS Single Sign-On (AWS SSO) to implement this functionality.
Which solution will meet these requirements MOST cost-effectively?

A. Create an organization in AWS Organization
B. Turn on the AWS SSO feature in Organizations Create and configure a directory in AWS Directory Service for Microsoft Active Directory (AWS Managed Microsoft AD) with a two-way trust to the company's on-premises Active Director
C. Configure AWS SSO and set the AWS Managed Microsoft AD directory as the identity sourc
D. Create permission sets and map them to the existing groups within the AWS Managed Microsoft AD directory.
E. Create an organization in AWS Organization
F. Turn on the AWS SSO feature in Organizations Create and configure an AD Connector to connect to the company's on-premises Active Director
G. Configure AWS SSO and select the AD Connector as the identity sourc
H. Create permission sets and map them to the existing groups within the company's Active Directory.
I. Create an organization in AWS Organization
J. Turn on all features for the organizatio
K. Create and configure a directory in AWS Directory Service for Microsoft Active Directory (AWS Managed Microsoft AD) with a two-way trust to the company's on-premises Active Director
L. Configure AWS SSO and select the AWS Managed Microsoft AD directory as the identity sourc
M. Create permission sets and map them to the existing groups within the AWS Managed Microsoft AD directory.
N. Create an organization in AWS Organization
O. Turn on all features for the organizatio
P. Create and configure an AD Connector to connect to the company's on-premises Active Director
Q. Configure AWS SSO and select the AD Connector as the identity sourc

R. Create permission sets and map them to the existing groups within the company's Active Directory.

Answer: D

Explanation:
https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_org_support-all-features.html
https://docs.aws.amazon.com/singlesignon/latest/userguide/get-started-prereqs-considerations.html

NEW QUESTION 121
- (Exam Topic 1)
A company is storing data on premises on a Windows file server. The company produces 5 GB of new data daily.
The company migrated part of its Windows-based workload to AWS and needs the data to be available on a file system in the cloud. The company already has established an AWS Direct Connect connection between the on-premises network and AWS.
Which data migration strategy should the company use?

A. Use the file gateway option in AWS Storage Gateway to replace the existing Windows file server, and point the existing file share to the new file gateway.
B. Use AWS DataSync to schedule a daily task to replicate data between the on-premises Windows file server and Amazon FSx.
C. Use AWS Data Pipeline to schedule a daily task to replicate data between the on-premises Windows file server and Amazon Elastic File System (Amazon EFS).
D. Use AWS DataSync to schedule a daily task lo replicate data between the on-premises Windows file server and Amazon Elastic File System (Amazon EFS),

Answer: B

Explanation:
https://aws.amazon.com/storagegateway/file/
https://docs.aws.amazon.com/fsx/latest/WindowsGuide/migrate-files-to-fsx-datasync.html https://docs.aws.amazon.com/systems-manager/latest/userguide/prereqs-operating-systems.html#prereqs-os-win

NEW QUESTION 125
- (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 128
- (Exam Topic 1)
A solutions architect must analyze a company's Amazon EC2 Instances and Amazon Elastic Block Store (Amazon EBS) volumes to determine whether the company is using resources efficiently The company is running several large, high-memory EC2 instances lo host database dusters that are deployed in active/passive configurations The utilization of these EC2 instances varies by the applications that use the databases, and the company has not identified a pattern The solutions architect must analyze the environment and take action based on the findings. Which solution meets these requirements MOST cost-effectively?

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

Answer: C

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

NEW QUESTION 132
- (Exam Topic 1)
A retail company has structured its AWS accounts to be part of an organization in AWS Organizations. The company has set up consolidated billing and has mapped its departments to the following OUs: Finance. Sales. Human Resources <HR). Marketing, and Operations. Each OU has multiple AWS accounts, one for each environment within a department. These environments are development, test, pre-production, and production.

The HR department is releasing a new system thai will launch in 3 months. In preparation, the HR department has purchased several Reserved Instances (RIs) in its production AWS account. The HR department will install the new application on this account. The HR department wants to make sure that other departments cannot share the Rl discounts.
Which solution will meet these requirements?

A. In the AWS Billing and Cost Management console for the HR department's production account, turn off R1 sharing.
B. Remove the HR department's production AWS account from the organizatio
C. Add the account to the consolidating billing configuration only.
D. In the AWS Billing and Cost Management console, use the organization's management account to turn off R1 sharing for the HR department's production AWS account.
E. Create an SCP in the organization to restrict access to the RI
F. Apply the SCP to the OUs of the other departments.

Answer: C

Explanation:
You can use the management account of the organization in AWS Billing and Cost Management console to turn off RI sharing for the HR department's production AWS account. This will prevent other departments from sharing the RI discounts and ensure that only the HR department can use the RIs purchased in their production account.

NEW QUESTION 134
- (Exam Topic 1)
A company runs a proprietary stateless ETL application on an Amazon EC2 Linux instance. The application is a Linux binary, and the source code cannot be modified. The application is single-threaded, uses 2 GB of RAM. and is highly CPU intensive The application is scheduled to run every 4 hours and runs for up to 20 minutes A solutions architect wants to revise the architecture for the solution.
Which strategy should the solutions architect use?

A. Use AWS Lambda to run the applicatio
B. Use Amazon CloudWatch Logs to invoke the Lambda function every 4 hours.
C. Use AWS Batch to run the applicatio
D. Use an AWS Step Functions state machine to invoke the AWS Batch job every 4 hours.
E. Use AWS Fargate to run the applicatio
F. Use Amazon EventBridge (Amazon CloudWatch Events) to invoke the Fargate task every 4 hours.
G. Use Amazon EC2 Spot Instances to run the applicatio
H. Use AWS CodeDeploy to deploy and run the application every 4 hours.

Answer: C

Explanation:
step function could run a scheduled task when triggered by eventbrige, but why would you add that layer of complexity just to run aws batch when you could directly invoke it through eventbridge. The link provided - https://aws.amazon.com/pt/blogs/compute/orchestrating-high-performance-computing-with-aws-step-functions- makes sense only for HPC, this is a single instance that needs to be run

NEW QUESTION 136
- (Exam Topic 1)
A company wants to use AWS to create a business continuity solution in case the company's main on-premises application fails. The application runs on physical servers that also run other applications. The on-premises application that the company is planning to migrate uses a MySQL database as a data store. All the company's on-premises applications use operating systems that are compatible with Amazon EC2.
Which solution will achieve the company's goal with the LEAST operational overhead?

A. Install the AWS Replication Agent on the source servers, including the MySQL server
B. Set up replication for all server
C. Launch test instances for regular drill
D. Cut over to the test instances to fail over the workload in the case of a failure event.
E. Install the AWS Replication Agent on the source servers, including the MySQL server
F. Initialize AWS Elastic Disaster Recovery in the target AWS Regio
G. Define the launch setting
H. Frequently perform failover and fallback from the most recent point in time.
I. Create AWS Database Migration Service (AWS DMS) replication servers and a target Amazon Aurora MySQL DB cluster to host the databas
J. Create a DMS replication task to copy the existing data to the target DB cluste
K. Create a local AWS Schema Conversion Tool (AWS SCT) change data capture (CDC) task to keep the data synchronize
L. Install the rest of the software on EC2 instances by starting with a compatible base AMI.
M. Deploy an AWS Storage Gateway Volume Gateway on premise
N. Mount volumes on all on-premises server
O. Install the application and the MySQL database on the new volume
P. Take regular snapshot
Q. Install all the software on EC2 Instances by starting with a compatible base AM
R. Launch a Volume Gateway on an EC2 instanc
S. Restore the volumes from the latest snapsho
T. Mount the new volumes on the EC2 instances in the case of a failure event.

Answer: B

Explanation:
https://docs.aws.amazon.com/drs/latest/userguide/what-is-drs.html https://docs.aws.amazon.com/drs/latest/userguide/recovery-workflow-gs.html

NEW QUESTION 141
- (Exam Topic 1)
A company is running an event ticketing platform on AWS and wants to optimize the platform's
cost-effectiveness. The platform is deployed on Amazon Elastic Kubernetes Service (Amazon EKS) with Amazon EC2 and is backed by an Amazon RDS for

MySQL DB instance. The company is developing new application features to run on Amazon EKS with AWS Fargate.
The platform experiences infrequent high peaks in demand. The surges in demand depend on event dates. Which solution will provide the MOST cost-effective setup for the platform?

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

Answer: B

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

NEW QUESTION 143
- (Exam Topic 1)
A company wants to use a third-party software-as-a-service (SaaS) application. The third-party SaaS application is consumed through several API calls. The third-party SaaS application also runs on AWS inside a VPC.
The company will consume the third-party SaaS application from inside a VPC. The company has internal security policies that mandate the use of private connectivity that does not traverse the internet. No resources that run in the company VPC are allowed to be accessed from outside the company's VPC. All permissions must conform to the principles of least privilege.
Which solution meets these requirements?

A. Create an AWS PrivateLink interface VPC endpoin
B. Connect this endpoint to the endpoint service that the third-party SaaS application provide
C. Create a security group to limit the access to the endpoin
D. Associate the security group with the endpoint.
E. Create an AWS Site-to-Site VPN connection between the third-party SaaS application and the company VP
F. Configure network ACLs to limit access across the VPN tunnels.
G. Create a VPC peering connection between the third-party SaaS application and the company VPUpdate route tables by adding the needed routes for the peering connection.
H. Create an AWS PrivateLink endpoint servic
I. Ask the third-party SaaS provider to create an interface VPC endpoint for this endpoint servic
J. Grant permissions for the endpoint service to the specific account of the third-party SaaS provider.

Answer: A

Explanation:
Reference architecture - https://docs.aws.amazon.com/vpc/latest/privatelink/privatelink-access-saas.html Note from documentation that Interface Endpoint is at client side

NEW QUESTION 145
- (Exam Topic 2)
A company runs a customer service center that accepts calls and automatically sends all customers a managed, interactive, two-way experience survey by text message.
The applications that support the customer service center run on machines that the company hosts in an on-premises data center. The hardware that the company uses is old, and the company is experiencing downtime with the system. The company wants to migrate the system to AWS to improve reliability.
Which solution will meet these requirements with the LEAST ongoing operational overhead?

A. Use Amazon Connect to replace the old call center hardwar
B. Use Amazon Pinpoint to send text message surveys to customers.
C. Use Amazon Connect to replace the old call center hardwar
D. Use Amazon Simple Notification Service (Amazon SNS) to send text message surveys to customers.
E. Migrate the call center software to Amazon EC2 instances that are in an Auto Scaling grou
F. Use the EC2 instances to send text message surveys to customers.
G. Use Amazon Pinpoint to replace the old call center hardware and to send text message surveys to customers.

Answer: A

Explanation:
Amazon Connect is a cloud-based contact center service that allows you to set up a virtual call center for your business. It provides an easy-to-use interface for managing customer interactions through voice and chat. Amazon Connect integrates with other AWS services, such as Amazon S3 and Amazon Kinesis, to help you collect, store, and analyze customer data for insights into customer behavior and trends. On the other hand, Amazon Pinpoint is a marketing automation and analytics service that allows you to engage with your customers across different channels, such as email, SMS, push notifications, and voice. It helps you create personalized campaigns based on user behavior and enables you to track user engagement and retention. While both services allow you to communicate with your customers, they serve different purposes. Amazon Connect is focused on customer support and service, while Amazon Pinpoint is focused on marketing and engagement.

NEW QUESTION 147
- (Exam Topic 2)
A company uses AWS Organizations for a multi-account setup in the AWS Cloud. The company's finance team has a data processing application that uses AWS Lambda and Amazon DynamoDB. The company's marketing team wants to access the data that is stored in the DynamoDB table.
The DynamoDB table contains confidential data. The marketing team can have access to only specific attributes of data in the DynamoDB table. The fi-nance team and the marketing team have separate AWS accounts.
What should a solutions architect do to provide the marketing team with the appropriate access to the DynamoDB table?

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

Answer: C

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

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

References:

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

NEW QUESTION 152
- (Exam Topic 2)
A solutions architect wants to cost-optimize and appropriately size Amazon EC2 instances in a single AWS account. The solutions architect wants to ensure that the instances are optimized based on CPU, memory, and network metrics.
Which combination of steps should the solutions architect take to meet these requirements? (Choose two.)

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

Answer: BD

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

commitment to a consistent amount of usage for a 1- or 3-year term4. Savings Plans do not affect the configuration or utilization of EC2 instances.

NEW QUESTION 153

......