Amazon-Web-Services

Exam Questions SAP-C02

AWS Certified Solutions Architect - Professional

About Exambible

Your Partner of IT Exam

Found in 1998

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

Our Advances

* 99.9% Uptime

    All examinations will be up to date.

* 24/7 Quality Support

    We will provide service round the clock.

* 100% Pass Rate

    Our guarantee that you will pass the exam.

* Unique Gurantee

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

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

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

Answer: CE

Explanation:
➤ Option C is correct because leveraging Amazon Kinesis Data Streams and AWS Lambda to ingest and process the raw data resolves the issues permanently and enable growth as new sensors are provisioned. Amazon Kinesis Data Streams is a serverless streaming data service that simplifies the capture, processing, and storage of data streams at any scale. Kinesis Data Streams can handle any amount of streaming data and process data from hundreds of thousands of sources with very low latency. AWS Lambda is a serverless compute service that lets you run code without provisioning or managing servers. Lambda can be triggered by Kinesis Data Streams events and process the data records in real time. Lambda can also scale automatically based on the incoming data volume. By using Kinesis Data Streams and Lambda, the company can reduce the load on the API servers and improve the performance and scalability of the data ingestion and processing layer3

➤ Option E is correct because re-architecting the database tier to use Amazon DynamoDB instead of an RDS MySQL DB instance resolves the issues permanently and enable growth as new sensors are provisioned. Amazon DynamoDB is a fully managed key-value and document database that delivers single-digit millisecond performance at any scale. DynamoDB supports auto scaling, which automatically adjusts read and write capacity based on actual traffic patterns. DynamoDB also supports on-demand capacity mode, which instantly accommodates up to double the previous peak traffic on a table. By using DynamoDB instead of RDS MySQL DB instance, the company can eliminate high write latency and improve scalability and performance of the database tier.
References: 1: https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-volume-types.html 2: https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/CHAP_AuroraOverview.html 3: https://docs.aws.amazon.com/streams/latest/dev/introduction.html : https://docs.aws.amazon.com/lambda/latest/dg/welcome.html : https://docs.aws.amazon.com/xray/latest/devguide/aws-xray.html : https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Introduction.html :

NEW QUESTION 3
- (Exam Topic 1)
A company is running an application in the AWS Cloud. The company's security team must approve the creation of all new IAM users. When a new IAM user is created, all access for the user must be removed automatically. The security team must then receive a notification to approve the user. The company has a multi-Region AWS CloudTrail trail In the AWS account.
Which combination of steps will meet these requirements? (Select THREE.)

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

Answer: ADE

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

NEW QUESTION 4
- (Exam Topic 1)
A software as a service (SaaS) based company provides a case management solution to customers A3 part of the solution. The company uses a standalone Simple Mail Transfer Protocol (SMTP) server to send email messages from an application. The application also stores an email template for acknowledgement email messages that populate customer data before the application sends the email message to the customer.
The company plans to migrate this messaging functionality to the AWS Cloud and needs to minimize operational overhead.
Which solution will meet these requirements MOST cost-effectively?

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

Answer: D

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

NEW QUESTION 5
- (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 6
- (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 7
- (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 8
- (Exam Topic 1)
A company wants to migrate its workloads from on premises to AWS. The workloads run on Linux and Windows. The company has a large on-premises intra structure that consists of physical machines and VMs that host numerous applications.
The company must capture details about the system configuration. system performance. running processes and network coi.net lions of its o. -premises ,on boards. The company also must divide the on-premises applications into groups for AWS migrations. The company needs recommendations for Amazon EC2 instance types so that the company can run its workloads on AWS in the most cost-effective manner.
Which combination of steps should a solutions architect take to meet these requirements? (Select THREE.)

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

Answer: ADE

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

NEW QUESTION 9
- (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 10
- (Exam Topic 1)
A company manages multiple AWS accounts by using AWS Organizations. Under the root OU. the company has two OUs: Research and DataOps.
Because of regulatory requirements, all resources that the company deploys in the organization must reside in the ap-northeast-1 Region. Additionally. EC2 instances that the company deploys in the DataOps OU must use a predefined list of instance types
A solutions architect must implement a solution that applies these restrictions. The solution must maximize operational efficiency and must minimize ongoing maintenance
Which combination of steps will meet these requirements? (Select TWO)

A. Create an IAM role in one account under the DataOps OU Use the ec2 Instance Type condition key in an inline policy on the role to restrict access to specific instance types.
B. Create an IAM user in all accounts under the root OU Use the aws RequestedRegion condition key in an inline policy on each user to restrict access to all AWS Regions except ap-northeast-1.
C. Create an SCP Use the aws:RequestedRegion condition key to restrict access to all AWS Regions except ap-northeast-1 Apply the SCP to the root OU.
D. Create an SCP Use the ec2Reo»on condition key to restrict access to all AWS Regions except ap-northeast-1. Apply the SCP to the root O
E. the DataOps O
F. and the Research OU.
G. Create an SCP Use the ec2:InstanceType condition key to restrict access to specific instance types Apply the SCP to the DataOps OU.

Answer: CE

Explanation:
https://docs.aws.amazon.com/IAM/latest/UserGuide/reference_policies_examples_aws_deny-requested-region.h
https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps_examples_ec2.html

NEW QUESTION 10
- (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 15
- (Exam Topic 1)
A publishing company's design team updates the icons and other static assets that an ecommerce web application uses. The company serves the icons and assets from an Amazon S3 bucket that is hosted in the company's production account. The company also uses a development account that members of the design team can access.
After the design team tests the static assets in the development account, the design team needs to load the assets into the S3 bucket in the production account. A solutions architect must provide the design team with access to the production account without exposing other parts of the web application to the risk of unwanted changes.
Which combination of steps will meet these requirements? (Select THREE.)

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

Answer: ACE

Explanation:

A. In the production account, create a new IAM policy that allows read and write access to the S3 bucket. The policy grants the necessary permissions to access the assets in the production S3 bucket.

C. In the production account, create a role. Attach the new policy to the role. Define the development account as a trusted entity. By creating a role and attaching the policy, and then defining the development account as a trusted entity, the development account can assume the role and access the production S3 bucket with the read and write permissions.

E. In the development account, create a group that contains all the IAM users of the design team. Attach a different IAM policy to the group to allow the sts:AssumeRole action on the role in the production account. The IAM policy attached to the group allows the design team members to assume the role created in the production account, thereby giving them access to the production S3 bucket.

Step 1: Create a role in the Production Account; create the role in the Production account and specify the Development account as a trusted entity. You also limit the role permissions to only read and write access to the productionapp bucket. Anyone granted permission to use the role can read and write to the productionapp bucket. Step 2: Grant access to the role Sign in as an administrator in the Development account and allow the AssumeRole action on the UpdateApp role in the Production account. So, recap, production account you create the policy for S3, and you set development account as a trusted entity. Then on the development account you allow the sts:assumeRole action on the role in production account. https://docs.aws.amazon.com/IAM/latest/UserGuide/tutorial_cross-account-with-roles.html

NEW QUESTION 17
- (Exam Topic 1)
A software company has deployed an application that consumes a REST API by using Amazon API Gateway. AWS Lambda functions, and an Amazon DynamoDB table. The application is showing an increase in the number of errors during PUT requests. Most of the PUT calls come from a small number of clients that are authenticated with specific API keys.
A solutions architect has identified that a large number of the PUT requests originate from one client. The API is noncritical, and clients can tolerate retries of unsuccessful calls. However, the errors are displayed to customers and are causing damage to the API's reputation.
What should the solutions architect recommend to improve the customer experience?

A. Implement retry logic with exponential backoff and irregular variation in the client applicatio
B. Ensure that the errors are caught and handled with descriptive error messages.
C. Implement API throttling through a usage plan at the API Gateway leve
D. Ensure that the client application handles code 429 replies without error.
E. Turn on API caching to enhance responsiveness for the production stag
F. Run 10-minute load tests.Verify that the cache capacity is appropriate for the workload.
G. Implement reserved concurrency at the Lambda function level to provide the resources that are needed during sudden increases in traffic.

Answer: B

Explanation:
https://aws.amazon.com/premiumsupport/knowledge-center/aws-batch-requests-error/ https://aws.amazon.com/premiumsupport/knowledge-center/api-gateway-429-limit/

NEW QUESTION 22
- (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 27
- (Exam Topic 1)
A company has an organization in AWS Organizations that has a large number of AWS accounts. One of the AWS accounts is designated as a transit account and has a transit gateway that is shared with all of the other AWS accounts AWS Site-to-Site VPN connections are configured between ail of the company's global offices and the transit account The company has AWS Config enabled on all of its accounts.

The company's networking team needs to centrally manage a list of internal IP address ranges that belong to the global offices Developers Will reference this list to gain access to applications securely.
Which solution meets these requirements with the LEAST amount of operational overhead?

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

Answer: C

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

NEW QUESTION 31
- (Exam Topic 1)
A solutions architect has developed a web application that uses an Amazon API Gateway Regional endpoint and an AWS Lambda function. The consumers of the web application are all close to the AWS Region where the application will be deployed. The Lambda function only queries an Amazon Aurora MySQL database. The solutions architect has configured the database to have three read replicas.
During testing, the application does not meet performance requirements. Under high load, the application opens a large number of database connections. The solutions architect must improve the application's performance.
Which actions should the solutions architect take to meet these requirements? (Choose two.)

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

Answer: BD

Explanation:
Connect to RDS outside of Lambda handler method to improve performance https://awstut.com/en/2022/04/30/connect-to-rds-outside-of-lambda-handler-method-to-improve-performance-en
Using RDS Proxy, you can handle unpredictable surges in database traffic. Otherwise, these surges might cause issues due to oversubscribing connections or creating new connections at a fast rate. RDS Proxy establishes a database connection pool and reuses connections in this pool. This approach avoids the memory and CPU overhead of opening a new database connection each time. To protect the database against oversubscription, you can control the number of database connections that are created. https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/rds-proxy.html

NEW QUESTION 36
- (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 41
- (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 42
- (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 46
- (Exam Topic 1)
A company wants to migrate an application to Amazon EC2 from VMware Infrastructure that runs in an
on-premises data center. A solutions architect must preserve the software and configuration settings during the migration.
What should the solutions architect do to meet these requirements?

A. Configure the AWS DataSync agent to start replicating the data store to Amazon FSx for Windows FileServer Use the SMB share to host the VMware data stor
B. Use VM Import/Export to move the VMs to Amazon EC2.
C. Use the VMware vSphere client to export the application as an image in Open Virealization Format (OVF) format Create an Amazon S3 bucket to store the image in the destination AWS Regio
D. Create and apply an IAM role for VM Import Use the AWS CLI to run the EC2 import command.
E. . Configure AWS Storage Gateway for files service to export a Common Internet File System (CIFSJ shar
F. Create a backup copy to the shared folde
G. Sign in to the AWS Management Console and create an AMI from the backup copy Launch an EC2 instance that is based on the AMI.
H. Create a managed-instance activation for a hybrid environment in AWS Systems Manage
I. Download and install Systems Manager Agent on the on-premises VM Register the VM with Systems Manager to be a managed instance Use AWS Backup to create a snapshot of the VM and create an AM
J. Launch an EC2 instance that is based on the AMI

Answer: D

Explanation:
https://docs.aws.amazon.com/vm-import/latest/userguide/vmimport-image-import.html
- Export an OVF Template
- Create / use an Amazon S3 bucket for storing the exported images. The bucket must be in the Region where you want to import your VMs.
- Create an IAM role named vmimport.
- You'll use AWS CLI to run the import commands. https://aws.amazon.com/premiumsupport/knowledge-center/import-instances/

NEW QUESTION 48
- (Exam Topic 1)
A retail company is hosting an ecommerce website on AWS across multiple AWS Regions. The company wants the website to be operational at all times for online purchases. The website stores data in an Amazon RDS for MySQL DB instance.
Which solution will provide the HIGHEST availability for the database?

A. Configure automated backups on Amazon RD
B. In the case of disruption, promote an automated backup to be a standalone DB instanc
C. Direct database traffic to the promoted DB instanc
D. Create a replacement read replica that has the promoted DB instance as its source.
E. Configure global tables and read replicas on Amazon RD
F. Activate the cross-Region scop
G. In the case of disruption, use AWS Lambda to copy the read replicas from one Region to another Region.
H. Configure global tables and automated backups on Amazon RD
I. In the case of disruption, use AWS Lambda to copy the read replicas from one Region to another Region.
J. Configure read replicas on Amazon RD
K. In the case of disruption, promote a cross-Region and read replica to be a standalone DB instanc
L. Direct database traffic to the promoted DB instanc
M. Create areplacement read replica that has the promoted DB instance as its source.

Answer: D

Explanation:
This solution will provide the highest availability for the database, as the read replicas will allow the database to be available in multiple Regions, thus reducing the chances of disruption. Additionally, the promotion of the cross-Region read replica to become a standalone DB instance will ensure that the database is still available even if one of the Regions experiences disruptions.

NEW QUESTION 51
- (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 52
- (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 54
- (Exam Topic 1)
A company developed a pilot application by using AWS Elastic Beanstalk and Java. To save costs during development, the company's development team deployed the application into a single-instance environment. Recent tests indicate that the application consumes more CPU than expected. CPU utilization is regularly greater than 85%, which causes some performance bottlenecks.
A solutions architect must mitigate the performance issues before the company launches the application to production.
Which solution will meet these requirements with the LEAST operational overhead?

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

Answer: C

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

NEW QUESTION 58
- (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 61
- (Exam Topic 1)
A company has developed a web application. The company is hosting the application on a group of Amazon EC2 instances behind an Application Load Balancer. The company wants to improve the security posture of the application and plans to use AWS WAF web ACLs. The solution must not adversely affect legitimate traffic to the application.
How should a solutions architect configure the web ACLs to meet these requirements?

A. Set the action of the web ACL rules to Coun
B. Enable AWS WAF logging Analyze the requests for false positives Modify the rules to avoid any false positive Over time change the action of the web ACL rules from Count to Block.
C. Use only rate-based rules in the web ACL
D. and set the throttle limit as high as possible Temporarily block all requests that exceed the limi
E. Define nested rules to narrow the scope of the rate tracking.
F. Set the action o' the web ACL rules to Bloc
G. Use only AWS managed rule groups in the web ACLs Evaluate the rule groups by using Amazon CloudWatch metrics with AWS WAF sampled requests or AWS WAF logs.
H. Use only custom rule groups in the web ACL
I. and set the action to Allow Enable AWS WAF logging Analyze the requests tor false positives Modify the rules to avoid any false positive Over time, change the action of the web ACL rules from Allow to Block.

Answer: A

Explanation:
https://aws.amazon.com/premiumsupport/knowledge-center/waf-analyze-count-action-rules/

NEW QUESTION 64
- (Exam Topic 1)
A company plans to refactor a monolithic application into a modern application designed deployed or AWS. The CLCD pipeline needs to be upgraded to support the modem design for the application with the following requirements
• It should allow changes to be released several times every hour.
* It should be able to roll back the changes as quickly as possible Which design will meet these requirements?

A. Deploy a CI-CD pipeline that incorporates AMIs to contain the application and their configurationsDeploy the application by replacing Amazon EC2 instances
B. Specify AWS Elastic Beanstak to sage in a secondary environment as the deployment target for the CI/CD pipeline of the applicatio
C. To deploy swap the staging and production environment URLs.
D. Use AWS Systems Manager to re-provision the infrastructure for each deployment Update the Amazon EC2 user data to pull the latest code art-fact from Amazon S3 and use Amazon Route 53 weighted routing to point to the new environment
E. Roll out At application updates as pan of an Auto Scaling event using prebuilt AMI
F. Use new versions of the AMIs to add instances, and phase out all instances that use the previous AMI version with the configured termination policy during a deployment event.

Answer: B

Explanation:
It is the fastest when it comes to rollback and deploying changes every hour

NEW QUESTION 65
- (Exam Topic 1)
A company has a latency-sensitive trading platform that uses Amazon DynamoDB as a storage backend. The company configured the DynamoDB table to use on-demand capacity mode. A solutions architect needs to design a solution to improve the performance of the trading platform. The new solution must ensure high availability for the trading platform.
Which solution will meet these requirements with the LEAST latency?

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

Answer: B

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

NEW QUESTION 68
- (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 73
- (Exam Topic 1)
A finance company is running its business-critical application on current-generation Linux EC2 instances The application includes a self-managed MySQL database performing heavy I/O operations. The application is working fine to handle a moderate amount of traffic during the month. However, it slows down during the final three days of each month due to month-end reporting, even though the company is using Elastic Load Balancers and Auto Scaling within its infrastructure to meet the increased demand.
Which of the following actions would allow the database to handle the month-end load with the LEAST impact on performance?

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

Answer: B

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

NEW QUESTION 75
- (Exam Topic 1)
A financial services company receives a regular data feed from its credit card servicing partner Approximately 5.1 records are sent every 15 minutes in plaintext, delivered over HTTPS directly into an Amazon S3 bucket with server-side encryption. This feed contains sensitive credit card primary account number (PAN) data The company needs to automatically mask the PAN before sending the data to another S3 bucket for additional internal processing. The company also needs to remove and merge specific fields, and then transform the record into JSON format Additionally, extra feeds are likely to be added in the future, so any design needs to be easily expandable.
Which solutions will meet these requirements?

A. Trigger an AWS Lambda function on file delivery that extracts each record and writes it to an Amazon SQS queu
B. Trigger another Lambda function when new messages arrive in the SQS queue to process the records, writing the results to a temporary location in Amazon S3. Trigger a final Lambda function once the SQS queue is empty to transform the records into JSON format and send the results to another S3 bucket for internal processing.
C. Trigger an AWS Lambda function on file delivery that extracts each record and writes it to an Amazon SQS queu
D. Configure an AWS Fargate container application to automatically scale to a single instance when the SQS queue contains message
E. Have the application process each record, and transform the record into JSON forma
F. When the queue is empty, send the results to another S3 bucket for internal processing and scale down the AWS Fargate instance.
G. Create an AWS Glue crawler and custom classifier based on the data feed formats and build a table definition to matc
H. Trigger an AWS Lambda function on file delivery to start an AWS Glue ETL job to transform the entire record according to the processing and transformation requirement
I. Define the output format as JSO
J. Once complete, have the ETL job send the results to another S3 bucket for internal processing.
K. Create an AWS Glue crawler and custom classifier based upon the data feed formats and build a table definition to matc
L. Perform an Amazon Athena query on file delivery to start an Amazon EMR ETLjob to transform the entire record according to the processing and transformation requirement
M. Define the output format as JSO
N. Once complete, send the results to another S3 bucket for internal processing and scale down the EMR cluster.

Answer: C

Explanation:
You can use a Glue crawler to populate the AWS Glue Data Catalog with tables. The Lambda function can be triggered using S3 event notifications when object create events occur. The Lambda function will then trigger the Glue ETL job to transform the records masking the sensitive data and modifying the output format to JSON. This solution meets all requirements.

NEW QUESTION 80
- (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 81
- (Exam Topic 1)
A company recently acquired several other companies. Each company has a separate AWS account with a different billing and reporting method. The acquiring company has consolidated all the accounts into one organization in AWS Organizations. However, the acquiring company has found it difficult to generate a cost report that contains meaningful groups for all the teams.
The acquiring company's finance team needs a solution to report on costs for all the companies through a self-managed application.
Which solution will meet these requirements?

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

Answer: A

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

NEW QUESTION 86
- (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 88
- (Exam Topic 1)
A company uses an on-premises data analytics platform. The system is highly available in a fully redundant configuration across 12 servers in the company's data center.
The system runs scheduled jobs, both hourly and daily, in addition to one-time requests from users. Scheduled jobs can take between 20 minutes and 2 hours to finish running and have tight SLAs. The scheduled jobs account for 65% of the system usage. User jobs typically finish running in less than 5 minutes and have no SLA. The user jobs account for 35% of system usage. During system failures, scheduled jobs must continue to meet SLAs. However, user jobs can be delayed.
A solutions architect needs to move the system to Amazon EC2 instances and adopt a consumption-based model to reduce costs with no long-term commitments. The solution must maintain high availability and must not affect the SLAs.
Which solution will meet these requirements MOST cost-effectively?

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

Answer: D

Explanation:
By splitting the 12 instances across three Availability Zones, the system can maintain high availability and availability of resources in case of a failure. Option D also uses a combination of On-Demand Instances with Capacity Reservations and Spot Instances, which allows for scheduled jobs to be run on the On-Demand instances with guaranteed capacity, while also taking advantage of the cost savings from Spot Instances for the user jobs which have lower SLA requirements.

NEW QUESTION 93
- (Exam Topic 1)
A company is using multiple AWS accounts The DNS records are stored in a private hosted zone for Amazon Route 53 in Account A The company's applications and databases are running in Account B.
A solutions architect win deploy a two-net application In a new VPC To simplify the configuration, the db.example com CNAME record set tor the Amazon RDS endpoint was created in a private hosted zone for Amazon Route 53.
During deployment, the application failed to start. Troubleshooting revealed that db.example com is not resolvable on the Amazon EC2 instance The solutions architect confirmed that the record set was created correctly in Route 53.
Which combination of steps should the solutions architect take to resolve this issue? (Select TWO)

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

Answer: CE

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

NEW QUESTION 97
- (Exam Topic 1)
A company has applications in an AWS account that is named Source. The account is in an organization in AWS Organizations. One of the applications uses AWS Lambda functions and store's inventory data in an Amazon Aurora database. The application deploys the Lambda functions by using a deployment package. The company has configured automated backups for Aurora.
The company wants to migrate the Lambda functions and the Aurora database to a new AWS account that is named Target. The application processes critical data, so the company must minimize downtime.
Which solution will meet these requirements?

A. Download the Lambda function deployment package from the Source accoun
B. Use the deployment package and create new Lambda functions in the Target accoun
C. Share the automated Aurora DB cluster snapshot with the Target account.
D. Download the Lambda function deployment package from the Source accoun
E. Use the deployment package and create new Lambda functions in the Target account Share the Aurora DB cluster with the Target account by using AWS Resource Access Manager {AWS RAM). Grant the Target account permission to clone the Aurora DB cluster.
F. Use AWS Resource Access Manager (AWS RAM) to share the Lambda functions and the Aurora DB cluster with the Target accoun
G. Grant the Target account permission to clone the Aurora DB cluster.
H. Use AWS Resource Access Manager (AWS RAM) to share the Lambda functions with the Target accoun
I. Share the automated Aurora DB cluster snapshot with the Target account.

Answer: C

Explanation:
This solution uses a combination of AWS Resource Access Manager (RAM) and automated backups to migrate the Lambda functions and the Aurora database to the Target account while minimizing downtime. In this solution, the Lambda function deployment package is downloaded from the Source account and used to create new Lambda functions in the Target account. The Aurora DB cluster is shared with the Target account using AWS RAM and the Target account is granted permission to clone the Aurora DB cluster, allowing for a new copy of the Aurora database to be created in the Target account. This approach allows for the data to be migrated to the Target account while minimizing downtime, as the Target account can use the cloned Aurora database while the original Aurora database continues to be used in the Source account.

NEW QUESTION 101
- (Exam Topic 1)
A company is building a software-as-a-service (SaaS) solution on AWS. The company has deployed an Amazon API Gateway REST API with AWS Lambda integration in multiple AWS Regions and in the same production account.
The company offers tiered pricing that gives customers the ability to pay for the capacity to make a certain number of API calls per second. The premium tier offers up to 3,000 calls per second, and customers are identified by a unique API key. Several premium tier customers in various Regions report that they receive error responses of 429 Too Many Requests from multiple API methods during peak usage hours. Logs indicate that the Lambda function is never invoked.
What could be the cause of the error messages for these customers?

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

Answer: C

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

NEW QUESTION 106
- (Exam Topic 1)
A company has a multi-tier web application that runs on a fleet of Amazon EC2 instances behind an Application Load Balancer (ALB). The instances are in an Auto Scaling group. The ALB and the Auto Scaling group are replicated in a backup AWS Region. The minimum value and the maximum value for the Auto Scaling group are set to zero. An Amazon RDS Multi-AZ DB instance stores the application's data. The DB instance has a read replica in the backup Region. The application presents an endpoint to end users by using an Amazon Route 53 record.
The company needs to reduce its RTO to less than 15 minutes by giving the application the ability to automatically fail over to the backup Region. The company does not have a large enough budget for an active-active strategy.
What should a solutions architect recommend to meet these requirements?

A. Reconfigure the application's Route 53 record with a latency-based routing policy that load balances traffic between the two ALB
B. Create an AWS Lambda function in the backup Region to promote the read replica and modify the Auto Scaling group value
C. Create an Amazon CloudWatch alarm that is based on the HTTPCode_Target_5XX_Count metric for the ALB in the primary Regio
D. Configure the CloudWatch alarm to invoke the Lambda function.
E. Create an AWS Lambda function in the backup Region to promote the read replica and modify the Auto Scaling group value
F. Configure Route 53 with a health check that monitors the web application and sends an Amazon Simple Notification Service (Amazon SNS) notification to the Lambda function when the health check status is unhealth
G. Update the application's Route 53 record with a failover policy that routes traffic to the ALB in the backup Region when a health check failure occurs.
H. Configure the Auto Scaling group in the backup Region to have the same values as the Auto Scaling group in the primary Regio
I. Reconfigure the application's Route 53 record with a latency-based routing policy that load balances traffic between the two ALB
J. Remove the read replic
K. Replace the read replica with a standalone RDS DB instanc
L. Configure Cross-Region Replication between the RDS DB instances by using snapshots and Amazon S3.
M. Configure an endpoint in AWS Global Accelerator with the two ALBs as equal weighted target
N. Create an AWS Lambda function in the backup Region to promote the read replica and modify the Auto Scaling group value
O. Create an Amazon CloudWatch alarm that is based on the HTTPCode_Target_5XX_Count metric for the ALB in the primary Regio

P. Configure the CloudWatch alarm to invoke the Lambda function.

Answer: B

Explanation:
an AWS Lambda function in the backup region to promote the read replica and modify the Auto Scaling group values, and then configuring Route 53 with a health check that monitors the web application and sends an Amazon SNS notification to the Lambda function when the health check status is unhealthy. Finally, the application's Route 53 record should be updated with a failover policy that routes traffic to the ALB in the backup region when a health check failure occurs. This approach provides automatic failover to the backup region when a health check failure occurs, reducing the RTO to less than 15 minutes. Additionally, this approach is cost-effective as it does not require an active-active strategy.

NEW QUESTION 109
- (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 113
- (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 118
- (Exam Topic 1)
An AWS customer has a web application that runs on premises. The web application fetches data from a third-party API that is behind a firewall. The third party accepts only one public CIDR block in each client's allow list.
The customer wants to migrate their web application to the AWS Cloud. The application will be hosted on a set of Amazon EC2 instances behind an Application Load Balancer (ALB) in a VPC. The ALB is located in public subnets. The EC2 instances are located in private subnets. NAT gateways provide internet access to the private subnets.
How should a solutions architect ensure that the web application can continue to call the third-parly API after the migration?

A. Associate a block of customer-owned public IP addresses to the VP

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

Answer: B

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

NEW QUESTION 122
- (Exam Topic 1)
A retail company has an on-premises data center in Europe. The company also has a multi-Region AWS presence that includes the eu-west-1 and us-east-1 Regions. The company wants to be able to route network traffic from its on-premises infrastructure into VPCs in either of those Regions. The company also needs to support traffic that is routed directly between VPCs in those Regions. No single points of failure can exist on the network.
The company already has created two 1 Gbps AWS Direct Connect connections from its on-premises data center. Each connection goes into a separate Direct Connect location in Europe for high availability. These two locations are named DX-A and DX-B, respectively. Each Region has a single AWS Transit Gateway that is configured to route all inter-VPC traffic within that Region.
Which solution will meet these requirements?

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

Answer: D

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

NEW QUESTION 124
- (Exam Topic 1)
A company uses a service to collect metadata from applications that the company hosts on premises. Consumer devices such as TVs and internet radios access the applications. Many older devices do not support certain HTTP headers and exhibit errors when these headers are present in responses. The company has configured an on-premises load balancer to remove the unsupported headers from responses sent to older devices, which the company identified by the User-Agent headers.
The company wants to migrate the service to AWS, adopt serverless technologies, and retain the ability to support the older devices. The company has already migrated the applications into a set of AWS Lambda functions.
Which solution will meet these requirements?

A. Create an Amazon CloudFront distribution for the metadata servic
B. Create an Application Load Balancer (ALB). Configure the CloudFront distribution to forward requests to the AL
C. Configure the ALB to invoke the correct Lambda function for each type of reques
D. Create a CloudFront function to remove the problematic headers based on the value of the User-Agent header.
E. Create an Amazon API Gateway REST API for the metadata servic
F. Configure API Gateway to invoke the correct Lambda function for each type of reques
G. Modify the default gateway responses to remove the problematic headers based on the value of the User-Agent header.
H. Create an Amazon API Gateway HTTP API for the metadata servic
I. Configure API Gateway to invoke the correct Lambda function for each type of reques
J. Create a response mapping template to remove the problematic headers based on the value of the User-Agen
K. Associate the response data mapping with the HTTP API.
L. Create an Amazon CloudFront distribution for the metadata servic
M. Create an Application Load Balancer (ALB). Configure the CloudFront distribution to forward requests to the AL
N. Configure the ALB to invoke the correct Lambda function for each type of reques
O. Create a Lambda@Edge function that will remove the problematic headers in response to viewer requests based on the value of theUser-Agent header.

Answer: D

Explanation:
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/lambda-examples.html

NEW QUESTION 128
- (Exam Topic 1)
A company has a web application that allows users to upload short videos. The videos are stored on Amazon EBS volumes and analyzed by custom recognition software for categorization.
The website contains stat c content that has variable traffic with peaks in certain months. The architecture consists of Amazon EC2 instances running in an Auto Scaling group for the web application and EC2
instances running in an Auto Scaling group to process an Amazon SQS queue The company wants to
re-architect the application to reduce operational overhead using AWS managed services where possible and remove dependencies on third-party software.
Which solution meets these requirements?

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

Answer: C

Explanation:
Option C is correct because hosting the web application in Amazon S3, storing the uploaded videos in Amazon S3, and using S3 event notifications to publish events to the SQS queue reduces the operational overhead of managing EC2 instances and EBS volumes. Amazon S3 can serve static content such as HTML, CSS, JavaScript, and media files directly from S3 buckets. Amazon S3 can also trigger AWS Lambda functions through S3 event notifications when new objects are created or existing objects are updated or deleted. AWS Lambda can process the SQS queue with an AWS Lambda function that calls the Amazon Rekognition API to categorize the videos. This solution eliminates the need for custom recognition software and third-party dependencies345
References: 1: https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-spot-instances.html 2: https://aws.amazon.com/efs/pricing/ 3:
https://docs.aws.amazon.com/AmazonS3/latest/userguide/WebsiteHosting.html 4:
https://docs.aws.amazon.com/AmazonS3/latest/userguide/NotificationHowTo.html 5:
https://docs.aws.amazon.com/rekognition/latest/dg/what-is.html 6: https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/Welcome.html

NEW QUESTION 129
- (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 133
- (Exam Topic 1)
A retail company is operating its ecommerce application on AWS. The application runs on Amazon EC2 instances behind an Application Load Balancer (ALB). The company uses an Amazon RDS DB instance as the database backend. Amazon CloudFront is configured with one origin that points to the ALB. Static content is cached. Amazon Route 53 is used to host all public zones.
After an update of the application, the ALB occasionally returns a 502 status code (Bad Gateway) error. The root cause is malformed HTTP headers that are returned to the ALB. The webpage returns successfully when a solutions architect reloads the webpage immediately after the error occurs.
While the company is working on the problem, the solutions architect needs to provide a custom error page instead of the standard ALB error page to visitors.
Which combination of steps will meet this requirement with the LEAST amount of operational overhead? (Choose two.)

A. Create an Amazon S3 bucke
B. Configure the S3 bucket to host a static webpag
C. Upload the custom error pages to Amazon S3.
D. Create an Amazon CloudWatch alarm to invoke an AWS Lambda function if the ALB health check response Target.FailedHealthChecks is greater than 0.
Configure the Lambda function to modify the forwarding rule at the ALB to point to a publicly accessible web server.
E. Modify the existing Amazon Route 53 records by adding health check
F. Configure a fallback target if the health check fail
G. Modify DNS records to point to a publicly accessible webpage.
H. Create an Amazon CloudWatch alarm to invoke an AWS Lambda function if the ALB health check response Elb.InternalError is greater than 0. Configure the Lambda function to modify the forwarding rule at the ALB to point to a public accessible web server.
I. Add a custom error response by configuring a CloudFront custom error pag
J. Modify DNS records to point to a publicly accessible web page.

Answer: CE

Explanation:
"Save your custom error pages in a location that is accessible to CloudFront. We recommend that you store them in an Amazon S3 bucket, and that you don't store them in the same place as the rest of your website or application's content. If you store the custom error pages on the same origin as your website or application, and the origin starts to return 5xx errors, CloudFront can't get the custom error pages because the origin server is unavailable."
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/GeneratingCustomErrorResponses.htm

NEW QUESTION 134
- (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 135
- (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 137
- (Exam Topic 1)
A financial company is planning to migrate its web application from on premises to AWS. The company uses a third-party security tool to monitor the inbound traffic to the application. The company has used the security tool for the last 15 years, and the tool has no cloud solutions available from its vendor. The company's security team is concerned about how to integrate the security tool with AWS technology.
The company plans to deploy the application migration to AWS on Amazon EC2 instances. The EC2 instances will run in an Auto Scaling group in a dedicated VPC. The company needs to use the security tool to inspect all packets that come in and out of the VPC. This inspection must occur in real time and must not affect the application's performance. A solutions architect must design a target architecture on AWS that is highly available within an AWS Region.
Which combination of steps should the solutions architect take to meet these requirements? (Select TWO.)

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

Answer: AD

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

NEW QUESTION 142
- (Exam Topic 1)
A large company is running a popular web application. The application runs on several Amazon EC2 Linux Instances in an Auto Scaling group in a private subnet. An Application Load Balancer is targeting the Instances In the Auto Scaling group in the private subnet. AWS Systems Manager Session Manager Is configured, and AWS Systems Manager Agent is running on all the EC2 instances.
The company recently released a new version of the application Some EC2 instances are now being marked as unhealthy and are being terminated As a result, the application is running at reduced capacity A solutions architect tries to determine the root cause by analyzing Amazon CloudWatch logs that are collected from the application, but the logs are inconclusive
How should the solutions architect gain access to an EC2 instance to troubleshoot the issue1?

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

Answer: D

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

NEW QUESTION 143
- (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 147
- (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 151
- (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 154
- (Exam Topic 1)
A company has registered 10 new domain names. The company uses the domains for online marketing. The company needs a solution that will redirect online visitors to a specific URL for each domain. All domains and target URLs are defined in a JSON document. All DNS records are managed by Amazon Route 53.
A solutions architect must implement a redirect service that accepts HTTP and HTTPS requests.
Which combination of steps should the solutions architect take to meet these requirements with the LEAST amount of operational effort? (Choose three.)

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

Answer: CEF

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

NEW QUESTION 159
- (Exam Topic 1)
A company used Amazon EC2 instances to deploy a web fleet to host a blog site The EC2 instances are behind an Application Load Balancer (ALB) and are configured in an Auto ScaSng group The web application stores all blog content on an Amazon EFS volume.
The company recently added a feature 'or Moggers to add video to their posts, attracting 10 times the previous user traffic At peak times of day. users report buffering and timeout issues while attempting to reach the site or watch videos
Which is the MOST cost-efficient and scalable deployment that win resolve the issues for users?

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

Answer: C

Explanation:
https://aws.amazon.com/premiumsupport/knowledge-center/cloudfront-https-connection-fails/ Using an Amazon S3 bucket
Using a MediaStore container or a MediaPackage channel Using an Application Load Balancer
Using a Lambda function URL
Using Amazon EC2 (or another custom origin)

Using CloudFront origin groups https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/restrict-access-to-load-balancer.html

NEW QUESTION 162
- (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 163
- (Exam Topic 1)
A company runs a serverless application in a single AWS Region. The application accesses external URLs and extracts metadata from those sites. The company uses an Amazon Simple Notification Service (Amazon SNS) topic to publish URLs to an Amazon Simple Queue Service (Amazon SQS) queue An AWS Lambda function uses the queue as an event source and processes the URLs from the queue Results are saved to an Amazon S3 bucket
The company wants to process each URL other Regions to compare possible differences in site localization URLs must be published from the existing Region. Results must be written to the existing S3 bucket in the current Region.
Which combination of changes will produce multi-Region deployment that meets these requirements? (Select TWO.)

A. Deploy the SOS queue with the Lambda function to other Regions.
B. Subscribe the SNS topic in each Region to the SQS queue.
C. Subscribe the SQS queue in each Region to the SNS topics in each Region.
D. Configure the SQS queue to publish URLs to SNS topics in each Region.
E. Deploy the SNS topic and the Lambda function to other Regions.

Answer: AC

Explanation:
https://docs.aws.amazon.com/sns/latest/dg/sns-cross-region-delivery.html

NEW QUESTION 165
- (Exam Topic 1)
A company has an organization in AWS Organizations. The company is using AWS Control Tower to deploy a landing zone for the organization. The company wants to implement governance and policy enforcement. The company must implement a policy that will detect Amazon RDS DB instances that are not encrypted at rest in the company's production OU.
Which solution will meet this requirement?

A. Turn on mandatory guardrails in AWS Control Towe
B. Apply the mandatory guardrails to the production OU.
C. Enable the appropriate guardrail from the list of strongly recommended guardrails in AWS Control Towe
D. Apply the guardrail to the production OU.
E. Use AWS Config to create a new mandatory guardrai
F. Apply the rule to all accounts in the production OU.
G. Create a custom SCP in AWS Control Towe
H. Apply the SCP to the production OU.

Answer: B

Explanation:
AWS Control Tower provides a set of "strongly recommended guardrails" that can be enabled to implement governance and policy enforcement. One of these guardrails is "Encrypt Amazon RDS instances" which will
detect RDS DB instances that are not encrypted at rest. By enabling this guardrail and applying it to the production OU, the company will be able to enforce encryption for RDS instances in the production environment.

NEW QUESTION 168
- (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 169
......

Relate Links

100% Pass Your SAP-C02 Exam with Exambible Prep Materials

https://www.exambible.com/SAP-C02-exam/

Contact us

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

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