Amazon-Web-Services

Exam Questions SAA-C03

AWS Certified Solutions Architect - Associate (SAA-C03)

NEW QUESTION 1
- (Topic 1)
A company hosts a containerized web application on a fleet of on-premises servers that process incoming requests. The number of requests is growing quickly. The on-premises servers cannot handle the increased number of requests. The company wants to move the application to AWS with minimum code changes and minimum development effort.
Which solution will meet these requirements with the LEAST operational overhead?

A. Use AWS Fargate on Amazon Elastic Container Service (Amazon ECS) to run the containerized web application with Service Auto Scalin
B. Use an Application Load Balancer to distribute the incoming requests.
C. Use two Amazon EC2 instances to host the containerized web applicatio
D. Use an Application Load Balancer to distribute the incoming requests
E. Use AWS Lambda with a new code that uses one of the supported language
F. Create multiple Lambda functions to support the loa
G. Use Amazon API Gateway as an entry point to the Lambda functions.
H. Use a high performance computing (HPC) solution such as AWS ParallelClusterto establish an HPC cluster that can process the incoming requests at the appropriate scale.

Answer: A

Explanation:
AWS Fargate is a serverless compute engine that lets users run containers without having to manage servers or clusters of Amazon EC2 instances1. Users can use AWS Fargate on Amazon Elastic Container Service (Amazon ECS) to run the containerized web application with Service Auto Scaling. Amazon ECS is a fully managed container orchestration service that supports both Docker and Kubernetes2. Service Auto Scaling is a feature that allows users to adjust the desired number of tasks in an ECS service based on CloudWatch metrics, such as CPU utilization or request count3. Users can use AWS Fargate on
Amazon ECS to migrate the application to AWS with minimum code changes and minimum development effort, as they only need to package their application in containers and specify the CPU and memory requirements.
Users can also use an Application Load Balancer to distribute the incoming requests. An Application Load Balancer is a load balancer that operates at the application layer and routes traffic to targets based on the content of the request. Users can register their ECS tasks as targets for an Application Load Balancer and configure listener rules to route requests to different target groups based on path or host headers. Users can use an Application Load Balancer to improve the availability and performance of their web
application.

NEW QUESTION 2
- (Topic 1)
A company has an application that generates a large number of files, each approximately 5 MB in size. The files are stored in Amazon S3. Company policy requires the files to be stored for 4 years before they can be deleted Immediate accessibility is always required as the files contain critical business data that is not easy to reproduce. The files are frequently accessed in the first 30 days of the object creation but are rarely accessed after the first 30 days
Which storage solution is MOST cost-effective?

A. Create an S3 bucket lifecycle policy to move Mm from S3 Standard to S3 Glacier 30 days from object creation Delete the Tiles 4 years after object creation
B. Create an S3 bucket lifecycle policy to move tiles from S3 Standard to S3 One Zone- infrequent Access (S3 One Zone-IA] 30 days from object creatio
C. Delete the fees 4 years after object creation
D. Create an S3 bucket lifecycle policy to move files from S3 Standard-infrequent Access (S3 Standard -IA) 30 from object creatio
E. Delete the ties 4 years after object creation
F. Create an S3 bucket Lifecycle policy to move files from S3 Standard to S3 Standard- Infrequent Access (S3 Standard-IA) 30 days from object creation Move the files to S3 Glacier 4 years after object carton.

Answer: B

Explanation:
https://aws.amazon.com/s3/storage-
classes/?trk=66264cd8-3b73-416c-9693-ea7cf4fe846a&sc_channel=ps&s_kwcid=AL!4422!3!536452716950!p!!g!!aws%20s3%20pri
cing&ef_id=Cj0KCQjwnbmaBhD- ARIsAGTPcfVHUZN5_BMrzl5zBcaC8KnqpnNZvjbZzqPkH6k7q4JcYO5KFLx0YYgaAm6nE
ALw_wcB:G:s&s_kwcid=AL!4422!3!536452716950!p!!g!!aws%20s3%20pricing

NEW QUESTION 3
- (Topic 1)
A company has an application that provides marketing services to stores. The services are based on previous purchases by store customers. The stores upload transaction data to the company through SFTP, and the data is processed and analyzed to generate new marketing offers. Some of the files can exceed 200 GB in size.
Recently, the company discovered that some of the stores have uploaded files that contain personally identifiable information (PII) that should not have been included. The company wants administrators to be alerted if PII is shared again. The company also wants to automate remediation.
What should a solutions architect do to meet these requirements with the LEAST development effort?

A. Use an Amazon S3 bucket as a secure transfer poin
B. Use Amazon Inspector to scan me objects in the bucke
C. If objects contain Pl
D. trigger an S3 Lifecycle policy to remove the objects that contain Pll.
E. Use an Amazon S3 bucket as a secure transfer poin
F. Use Amazon Macie to scan the objects in the bucke
G. If objects contain Pl
H. Use Amazon Simple Notification Service (Amazon SNS) to trigger a notification to the administrators to remove the objects mat contain Pll.
I. Implement custom scanning algorithms in an AWS Lambda functio
J. Trigger the function when objects are loaded into the bucke
K. It objects contain Rl
L. use Amazon Simple Notification Service (Amazon SNS) to trigger a notification to the administrators to remove the objects that contain Pll.
M. Implement custom scanning algorithms in an AWS Lambda functio
N. Trigger the function when objects are loaded into the bucke
O. If objects contain Pl

P. use Amazon Simple Email Service (Amazon STS) to trigger a notification to the administrators and trigger on S3 Lifecycle policy to remove the objects mot contain PII.

Answer: B

Explanation:
To meet the requirements of detecting and alerting the administrators when PII is shared and automating remediation with the least development effort, the best approach would be to use Amazon S3 bucket as a secure transfer point and scan the objects in the bucket with Amazon Macie. Amazon Macie is a fully managed data security and data privacy service that uses machine learning and pattern matching to discover and protect sensitive data stored in Amazon S3. It can be used to classify sensitive data, monitor access to sensitive data, and automate remediation actions.
In this scenario, after uploading the files to the Amazon S3 bucket, the objects can be scanned for PII by Amazon Macie, and if it detects any PII, it can trigger an Amazon Simple Notification Service (SNS) notification to alert the administrators to remove the objects containing PII. This approach requires the least development effort, as Amazon Macie already has pre-built data classification rules that can detect PII in various formats. Hence, option B is the correct answer.
References:
? Amazon Macie User Guide: https://docs.aws.amazon.com/macie/latest/userguide/what-is-macie.html
? AWS Well-Architected Framework - Security Pillar: https://docs.aws.amazon.com/wellarchitected/latest/security-pillar/welcome.html

NEW QUESTION 4
- (Topic 1)
A company stores call transcript files on a monthly basis. Users access the files randomly within 1 year of the call, but users access the files infrequently after 1 year. The company wants to optimize its solution by giving users the ability to query and retrieve files that are less than 1-year-old as quickly as possible. A delay in retrieving older files is acceptable.
Which solution will meet these requirements MOST cost-effectively?

A. Store individual files with tags in Amazon S3 Glacier Instant Retrieva
B. Query the tags to retrieve the files from S3 Glacier Instant Retrieval.
C. Store individual files in Amazon S3 Intelligent-Tierin
D. Use S3 Lifecycle policies to move the files to S3 Glacier Flexible Retrieval after 1 yea
E. Query and retrieve the files that are in Amazon S3 by using Amazon Athen
F. Query and retrieve the files that are in S3 Glacier by using S3 Glacier Select.
G. Store individual files with tags in Amazon S3 Standard storag
H. Store search metadata for each archive in Amazon S3 Standard storag
I. Use S3 Lifecycle policies to move the files to S3 Glacier Instant Retrieval after 1 yea
J. Query and retrieve the files by searching for metadata from Amazon S3.
K. Store individual files in Amazon S3 Standard storag
L. Use S3 Lifecycle policies to move the files to S3 Glacier Deep Archive after 1 yea
M. Store search metadata in Amazon RD
N. Query the files from Amazon RD
O. Retrieve the files from S3 Glacier Deep Archive.

Answer: B

Explanation:
"For archive data that needs immediate access, such as medical images, news media assets, or genomics data, choose the S3 Glacier Instant Retrieval storage class, an archive storage class that delivers the lowest cost storage with milliseconds retrieval. For archive data that does not require immediate access but needs the flexibility to retrieve large sets of data at no cost, such as backup or disaster recovery use cases, choose S3 Glacier Flexible Retrieval (formerly S3 Glacier), with retrieval in minutes or free bulk retrievals in 5- 12 hours." https://aws.amazon.com/about-aws/whats-new/2021/11/amazon-s3-glacier-instant-retrieval-storage-class/

NEW QUESTION 5
- (Topic 1)
A company has a production workload that runs on 1,000 Amazon EC2 Linux instances. The workload is powered by third-party software. The company needs to patch the third- party software on all EC2 instances as quickly as possible to remediate a critical security vulnerability.
What should a solutions architect do to meet these requirements?

A. Create an AWS Lambda function to apply the patch to all EC2 instances.
B. Configure AWS Systems Manager Patch Manager to apply the patch to all EC2 instances.
C. Schedule an AWS Systems Manager maintenance window to apply the patch to all EC2 instances.
D. Use AWS Systems Manager Run Command to run a custom command that applies the patch to all EC2 instances.

Answer: B

Explanation:
https://docs.aws.amazon.com/systems-manager/latest/userguide/about-windows-app-patching.html

NEW QUESTION 6
- (Topic 1)
A company hosts its multi-tier applications on AWS. For compliance, governance, auditing, and security, the company must track configuration changes on its AWS resources and record a history of API calls made to these resources.
What should a solutions architect do to meet these requirements?

A. Use AWS CloudTrail to track configuration changes and AWS Config to record API calls
B. Use AWS Config to track configuration changes and AWS CloudTrail to record API calls
C. Use AWS Config to track configuration changes and Amazon CloudWatch to record API calls
D. Use AWS CloudTrail to track configuration changes and Amazon CloudWatch to record API calls

Answer: B

Explanation:
AWS Config is a fully managed service that allows the company to assess, audit, and evaluate the configurations of its AWS resources. It provides a detailed

inventory of the resources in use and tracks changes to resource configurations. AWS Config can detect configuration changes and alert the company when changes occur. It also provides a historical view of changes, which is essential for compliance and governance purposes. AWS CloudTrail is a fully managed service that provides a detailed history of API calls made to the company's AWS resources. It records all API activity in the AWS account, including who made the API call, when the call was made, and what resources were affected by the call. This information is critical for security and auditing purposes, as it allows the company to investigate any suspicious activity that might occur on its AWS resources.

NEW QUESTION 7
- (Topic 1)
A company is developing an application that provides order shipping statistics for retrieval by a REST API. The company wants to extract the shipping statistics, organize the data into an easy-to-read HTML format, and send the report to several email addresses at the same time every morning.
Which combination of steps should a solutions architect take to meet these requirements? (Choose two.)

A. Configure the application to send the data to Amazon Kinesis Data Firehose.
B. Use Amazon Simple Email Service (Amazon SES) to format the data and to send the report by email.
C. Create an Amazon EventBridge (Amazon CloudWatch Events) scheduled event that invokes an AWS Glue job to query the application's API for the data.
D. Create an Amazon EventBridge (Amazon CloudWatch Events) scheduled event that invokes an AWS Lambda function to query the application's API for the data.
E. Store the application data in Amazon S3. Create an Amazon Simple Notification Service (Amazon SNS) topic as an S3 event destination to send the report by

Answer: BD

Explanation:
https://docs.aws.amazon.com/ses/latest/dg/send-email-formatted.html
* D. Create an Amazon EventBridge (Amazon CloudWatch Events) scheduled event that invokes an AWS Lambda function to query the application's API for the data. This step can be done using AWS Lambda to extract the shipping statistics and organize the data into an HTML format.
* B. Use Amazon Simple Email Service (Amazon SES) to format the data and send the report by email. This step can be done by using Amazon SES to send the report to multiple email addresses at the same time every morning.
Therefore, options D and B are the correct choices for this question. Option A is incorrect because Kinesis Data Firehose is not necessary for this use case. Option C is incorrect because AWS Glue is not required to query the application's API. Option E is incorrect because S3 event notifications cannot be used to send the report by email.

NEW QUESTION 8
- (Topic 1)
A company uses 50 TB of data for reporting. The company wants to move this data from on premises to AWS A custom application in the company's data center runs a weekly data transformation job. The company plans to pause the application until the data transfer is complete and needs to begin the transfer process as soon as possible.
The data center does not have any available network bandwidth for additional workloads A solutions architect must transfer the data and must configure the transformation job to continue to run in the AWS Cloud
Which solution will meet these requirements with the LEAST operational overhead?

A. Use AWS DataSync to move the data Create a custom transformation job by using AWS Glue
B. Order an AWS Snowcone device to move the data Deploy the transformation application to the device
C. Order an AWS Snowball Edge Storage Optimized devic
D. Copy the data to the devic
E. Create a custom transformation job by using AWS Glue
F. Order an AWS
G. Snowball Edge Storage Optimized device that includes Amazon EC2 compute Copy the data to the device Create a new EC2 instance on AWS to run the transformation application

Answer: D

Explanation:
AWS Snowball Edge is a type of Snowball device with on-board storage and compute power for select AWS capabilities. Snowball Edge can do local processing and edge- computing workloads in addition to transferring data between your local environment and the AWS Cloud1. Users can order an AWS Snowball Edge Storage Optimized device that includes Amazon EC2 compute to move 50 TB of data from on premises to AWS. The Storage Optimized device has 80 TB of usable storage and 40 vCPUs of compute power2. Users can copy the data to the device using the AWS OpsHub graphical user interface or the Snowball client command line tool3. Users can also create and run Amazon EC2 instances on the device using Amazon Machine Images (AMIs) that are compatible with the sbe1 instance type. Users can use the Snowball Edge device to transfer the data and run the transformation job locally without using any network bandwidth.
Users can also create a new EC2 instance on AWS to run the transformation application after the data transfer is complete. Amazon EC2 is a web service that provides secure, resizable compute capacity in the cloud. Users can launch an EC2 instance in the same AWS Region where they send their Snowball Edge device and choose an AMI that matches their application requirements. Users can use the EC2 instance to continue running the transformation job in the AWS Cloud.

NEW QUESTION 9
- (Topic 1)
A company provides a Voice over Internet Protocol (VoIP) service that uses UDP connections. The service consists of Amazon EC2 instances that run in an Auto Scaling group. The company has deployments across multiple AWS Regions.
The company needs to route users to the Region with the lowest latency. The company also needs automated failover between Regions.
Which solution will meet these requirements?

A. Deploy a Network Load Balancer (NLB) and an associated target grou
B. Associate the target group with the Auto Scaling grou
C. Use the NLB as an AWS Global Accelerator endpoint in each Region.
D. Deploy an Application Load Balancer (ALB) and an associated target grou
E. Associate the target group with the Auto Scaling grou
F. Use the ALB as an AWS Global Accelerator endpoint in each Region.
G. Deploy a Network Load Balancer (NLB) and an associated target grou
H. Associate the target group with the Auto Scaling grou
I. Create an Amazon Route 53 latency record that points to aliases for each NL
J. Create an Amazon CloudFront distribution that uses the latency record as an origin.

K. Deploy an Application Load Balancer (ALB) and an associated target grou
L. Associate the target group with the Auto Scaling grou
M. Create an Amazon Route 53 weighted record that points to aliases for each AL
N. Deploy an Amazon CloudFront distribution that uses the weighted record as an origin.

Answer: D

Explanation:
https://aws.amazon.com/global-accelerator/faqs/
HTTP /HTTPS - ALB ; TCP and UDP - NLB; Lowest latency routing and more throughput. Also supports failover, uses Anycast Ip addressing - Global Accelerator Caching at Egde Locations – Cloutfront
WS Global Accelerator automatically checks the health of your applications and routes user traffic only to healthy application endpoints. If the health status changes or you make configuration updates, AWS Global Accelerator reacts instantaneously to route your users to the next available endpoint..

NEW QUESTION 10
- (Topic 1)
A company hosts an application on AWS Lambda functions mat are invoked by an Amazon API Gateway API The Lambda functions save customer data to an Amazon Aurora MySQL database Whenever the company upgrades the database, the Lambda functions fail to establish database connections until the upgrade is complete The result is that customer data Is not recorded for some of the event
A solutions architect needs to design a solution that stores customer data that is created during database upgrades
Which solution will meet these requirements?

A. Provision an Amazon RDS proxy to sit between the Lambda functions and the database Configure the Lambda functions to connect to the RDS proxy
B. Increase the run time of me Lambda functions to the maximum Create a retry mechanism in the code that stores the customer data in the database
C. Persist the customer data to Lambda local storag
D. Configure new Lambda functions to scan the local storage to save the customer data to the database.
E. Store the customer data m an Amazon Simple Queue Service (Amazon SOS) FIFO queue Create a new Lambda function that polls the queue and stores the customer data in the database

Answer: D

Explanation:
https://www.learnaws.org/2020/12/13/aws-rds-proxy-deep-dive/
RDS proxy can improve application availability in such a situation by waiting for the new database instance to be functional and maintaining any requests received from the application during this time. The end result is that the application is more resilient to issues with the underlying database.
This will enable solution to hold data till the time DB comes back to normal. RDS proxy is to optimally utilize the connection between Lambda and DB. Lambda can open multiple connection concurrently which can be taxing on DB compute resources, hence RDS proxy was introduced to manage and leverage these connections efficiently.

NEW QUESTION 10
- (Topic 1)
A company maintains a searchable repository of items on its website. The data is stored in an Amazon RDS for MySQL database table that contains more than 10 million rows The database has 2 TB of General Purpose SSD storage There are millions of updates against this data every day through the company's website
The company has noticed that some insert operations are taking 10 seconds or longer The company has determined that the database storage performance is the problem
Which solution addresses this performance issue?

A. Change the storage type to Provisioned IOPS SSD
B. Change the DB instance to a memory optimized instance class
C. Change the DB instance to a burstable performance instance class
D. Enable Multi-AZ RDS read replicas with MySQL native asynchronous replication.

Answer: A

Explanation:
https://aws.amazon.com/ebs/features/
"Provisioned IOPS volumes are backed by solid-state drives (SSDs) and are the highest performance EBS volumes designed for your critical, I/O intensive database applications.
These volumes are ideal for both IOPS-intensive and throughput-intensive workloads that require extremely low latency."
https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/CHAP_Storage.html

NEW QUESTION 12
- (Topic 1)
A company recently migrated a message processing system to AWS. The system receives messages into an ActiveMQ queue running on an Amazon EC2 instance. Messages are processed by a consumer application running on Amazon EC2. The consumer application processes the messages and writes results to a MySQL database funning on Amazon EC2. The company wants this application to be highly available with tow operational complexity
Which architecture otters the HGHEST availability?

A. Add a second ActiveMQ server to another Availably Zone Add an additional consumer EC2 instance in another Availability Zon
B. Replicate the MySQL database to another Availability Zone.
C. Use Amazon MO with active/standby brokers configured across two Availability Zones Add an additional consumer EC2 instance in another Availability Zon
D. Replicate the MySQL database to another Availability Zone.
E. Use Amazon MO with active/standby blotters configured across two Availability Zone
F. Add an additional consumer EC2 instance in another Availability Zon
G. Use Amazon ROS tor MySQL with Multi-AZ enabled.
H. Use Amazon MQ with active/standby brokers configured across two Availability Zones Add an Auto Scaling group for the consumer EC2 instances across two Availability Zone
I. Use Amazon RDS for MySQL with Multi-AZ enabled.

Answer: D

Explanation:
Amazon S3 is a highly scalable and durable object storage service that can store and retrieve any amount of data from anywhere on the web1. Users can configure the application to upload images directly from each user's browser to Amazon S3 through the use of a presigned URL. A presigned URL is a URL that gives access to an object in an S3 bucket for a limited time and with a specific action, such as uploading an object2. Users can generate a presigned URL programmatically using the AWS SDKs or AWS CLI. By using a presigned URL, users can reduce coupling within the application and improve website performance, as they do not need to send the images to the web server first. AWS Lambda is a serverless compute service that runs code in response to events and automatically manages the underlying compute resources3. Users can configure S3 Event Notifications to invoke an AWS Lambda function when an image is uploaded. S3 Event Notifications is a feature that allows users to receive notifications when certain events happen in an S3 bucket, such as object creation or deletion. Users can configure S3 Event Notifications to invoke a Lambda function that resizes the image and stores it back in the same or a different S3 bucket. This way, users can offload the image resizing task from the web server to Lambda.

NEW QUESTION 14
- (Topic 1)
A company has a large Microsoft SharePoint deployment running on-premises that requires Microsoft Windows shared file storage. The company wants to migrate this workload to the AWS Cloud and is considering various storage options. The storage solution must be highly available and integrated with Active Directory for access control.
Which solution will satisfy these requirements?

A. Configure Amazon EFS storage and set the Active Directory domain for authentication
B. Create an SMB Me share on an AWS Storage Gateway tile gateway in two Availability Zones
C. Create an Amazon S3 bucket and configure Microsoft Windows Server to mount it as a volume
D. Create an Amazon FSx for Windows File Server file system on AWS and set the Active Directory domain for authentication

Answer: D

NEW QUESTION 15
- (Topic 1)
An image-processing company has a web application that users use to upload images. The application uploads the images into an Amazon S3 bucket. The company has set up S3 event notifications to publish the object creation events to an Amazon Simple Queue Service (Amazon SQS) standard queue. The SQS queue serves as the event source for an AWS Lambda function that processes the images and sends the results to users through email.
Users report that they are receiving multiple email messages for every uploaded image. A solutions architect determines that SQS messages are invoking the Lambda function more than once, resulting in multiple email messages.
What should the solutions architect do to resolve this issue with the LEAST operational overhead?

A. Set up long polling in the SQS queue by increasing the ReceiveMessage wait time to 30 seconds.
B. Change the SQS standard queue to an SQS FIFO queu
C. Use the message deduplication ID to discard duplicate messages.
D. Increase the visibility timeout in the SQS queue to a value that is greater than the total of the function timeout and the batch window timeout.
E. Modify the Lambda function to delete each message from the SQS queue immediately after the message is read before processing.

Answer: C

NEW QUESTION 18
- (Topic 1)
A company has applications that run on Amazon EC2 instances in a VPC. One of the applications needs to call the Amazon S3 API to store and read objects. According to the company's security regulations, no traffic from the applications is allowed to travel across the internet.
Which solution will meet these requirements?

A. Configure an S3 interface endpoint.
B. Configure an S3 gateway endpoint.
C. Create an S3 bucket in a private subnet.
D. Create an S3 bucket in the same Region as the EC2 instance.

Answer: B

Explanation:
https://docs.aws.amazon.com/AmazonS3/latest/userguide/privatelink-interface-endpoints.html#types-of-vpc-endpoints-for-s3
https://docs.aws.amazon.com/vpc/latest/userguide/vpc-endpoints-s3.html

NEW QUESTION 21
- (Topic 1)
A company is running an SMB file server in its data center. The file server stores large files that are accessed frequently for the first few days after the files are created. After 7 days the files are rarely accessed.
The total data size is increasing and is close to the company's total storage capacity. A solutions architect must increase the company's available storage space without losing low-latency access to the most recently accessed files. The solutions architect must also provide file lifecycle management to avoid future storage issues.
Which solution will meet these requirements?

A. Use AWS DataSync to copy data that is older than 7 days from the SMB file server to AWS.
B. Create an Amazon S3 File Gateway to extend the company's storage spac
C. Create an S3 Lifecycle policy to transition the data to S3 Glacier Deep Archive after 7 days.
D. Create an Amazon FSx for Windows File Server file system to extend the company's storage space.
E. Install a utility on each user's computer to access Amazon S3. Create an S3 Lifecycle policy to transition the data to S3 Glacier Flexible Retrieval after 7 days.

Answer: B

Explanation:
Amazon S3 File Gateway is a hybrid cloud storage service that enables on- premises applications to seamlessly use Amazon S3 cloud storage. It provides a file interface to Amazon S3 and supports SMB and NFS protocols. It also supports S3 Lifecycle policies that can automatically transition data from S3 Standard to S3

Glacier Deep Archive after a specified period of time. This solution will meet the requirements of increasing the company's available storage space without losing low-latency access to the most recently accessed files and providing file lifecycle management to avoid future storage issues.
Reference:
https://docs.aws.amazon.com/storagegateway/latest/userguide/WhatIsStorageGateway.ht ml

NEW QUESTION 25
- (Topic 1)
A company is storing sensitive user information in an Amazon S3 bucket The company wants to provide secure access to this bucket from the application tier running on Ama2on EC2 instances inside a VPC.
Which combination of steps should a solutions architect take to accomplish this? (Select TWO.)

A. Configure a VPC gateway endpoint for Amazon S3 within the VPC
B. Create a bucket policy to make the objects to the S3 bucket public
C. Create a bucket policy that limits access to only the application tier running in the VPC
D. Create an IAM user with an S3 access policy and copy the IAM credentials to the EC2 instance
E. Create a NAT instance and have the EC2 instances use the NAT instance to access the S3 bucket

Answer: AC

Explanation:
https://aws.amazon.com/premiumsupport/knowledge-center/s3-private-connection-no-authentication/

NEW QUESTION 30
- (Topic 1)
A company collects temperature, humidity, and atmospheric pressure data in cities across multiple continents. The average volume of data collected per site each day is 500 GB. Each site has a high-speed internet connection. The company's weather forecasting applications are based in a single Region and analyze the data daily.
What is the FASTEST way to aggregate data from all of these global sites?

A. Enable Amazon S3 Transfer Acceleration on the destination bucke
B. Use multipart uploads to directly upload site data to the destination bucket.
C. Upload site data to an Amazon S3 bucket in the closest AWS Regio
D. Use S3 cross- Region replication to copy objects to the destination bucket.
E. Schedule AWS Snowball jobs daily to transfer data to the closest AWS Regio
F. Use S3 cross-Region replication to copy objects to the destination bucket.
G. Upload the data to an Amazon EC2 instance in the closest Regio
H. Store the data in an Amazon Elastic Block Store (Amazon EBS) volum
I. Once a day take an EBS snapshot and copy it to the centralized Regio
J. Restore the EBS volume in the centralized Region and run an analysis on the data daily.

Answer: A

Explanation:
You might want to use Transfer Acceleration on a bucket for various reasons, including the following:
You have customers that upload to a centralized bucket from all over the world. You transfer gigabytes to terabytes of data on a regular basis across continents.
You are unable to utilize all of your available bandwidth over the Internet when uploading to Amazon S3.
https://docs.aws.amazon.com/AmazonS3/latest/dev/transfer-acceleration.html
https://aws.amazon.com/s3/transfer-acceleration/#:~:text=S3%20Transfer%20Acceleration%20(S3TA)%20reduces,to%20S3% 20for%20remote%20applications:
"Amazon S3 Transfer Acceleration can speed up content transfers to and from Amazon S3 by as much as 50-500% for long-distance transfer of larger objects. Customers who have either web or mobile applications with widespread users or applications hosted far away from their S3 bucket can experience long and variable upload and download speeds over the Internet"
https://docs.aws.amazon.com/AmazonS3/latest/userguide/mpuoverview.html
"Improved throughput - You can upload parts in parallel to improve throughput."

NEW QUESTION 35
- (Topic 1)
A company is building an application in the AWS Cloud. The application will store data in Amazon S3 buckets in two AWS Regions. The company must use an AWS Key Management Service (AWS KMS) customer managed key to encrypt all data that is stored in the S3 buckets. The data in both S3 buckets must be encrypted and decrypted with the same KMS key. The data and the key must be stored in each of the two Regions.
Which solution will meet these requirements with the LEAST operational overhead?

A. Create an S3 bucket in each Region Configure the S3 buckets to use server-side encryption with Amazon S3 managed encryption keys (SSE-S3) Configure replication between the S3 buckets.
B. Create a customer managed multi-Region KMS ke
C. Create an S3 bucket in each Regio
D. Configure replication between the S3 bucket
E. Configure the application to use the KMS key with client-side encryption.
F. Create a customer managed KMS key and an S3 bucket in each Region Configure the S3 buckets to use server-side encryption with Amazon S3 managed encryption keys (SSE- S3) Configure replication between the S3 buckets.
G. Create a customer managed KMS key and an S3 bucket m each Region Configure the S3 buckets to use server-side encryption with AWS KMS keys (SSE-KMS) Configure replication between the S3 buckets.

Answer: B

Explanation:
From https://docs.aws.amazon.com/kms/latest/developerguide/custom-key-store- overview.html
For most users, the default AWS KMS key store, which is protected by FIPS 140-2 validated cryptographic modules, fulfills their security requirements. There is no need to add an extra layer of maintenance responsibility or a dependency on an additional service. However, you might consider creating a custom key store if your organization has any of the following requirements: Key material cannot be stored in a shared environment. Key material must be subject to a secondary, independent audit path. The HSMs that generate and store key material must be certified at FIPS 140-2 Level 3.

https://docs.aws.amazon.com/kms/latest/developerguide/custom-key-store-overview.html
https://docs.aws.amazon.com/kms/latest/developerguide/multi-region-keys-overview.html

NEW QUESTION 37

- (Topic 1)

An Amazon EC2 administrator created the following policy associated with an IAM group containing several users

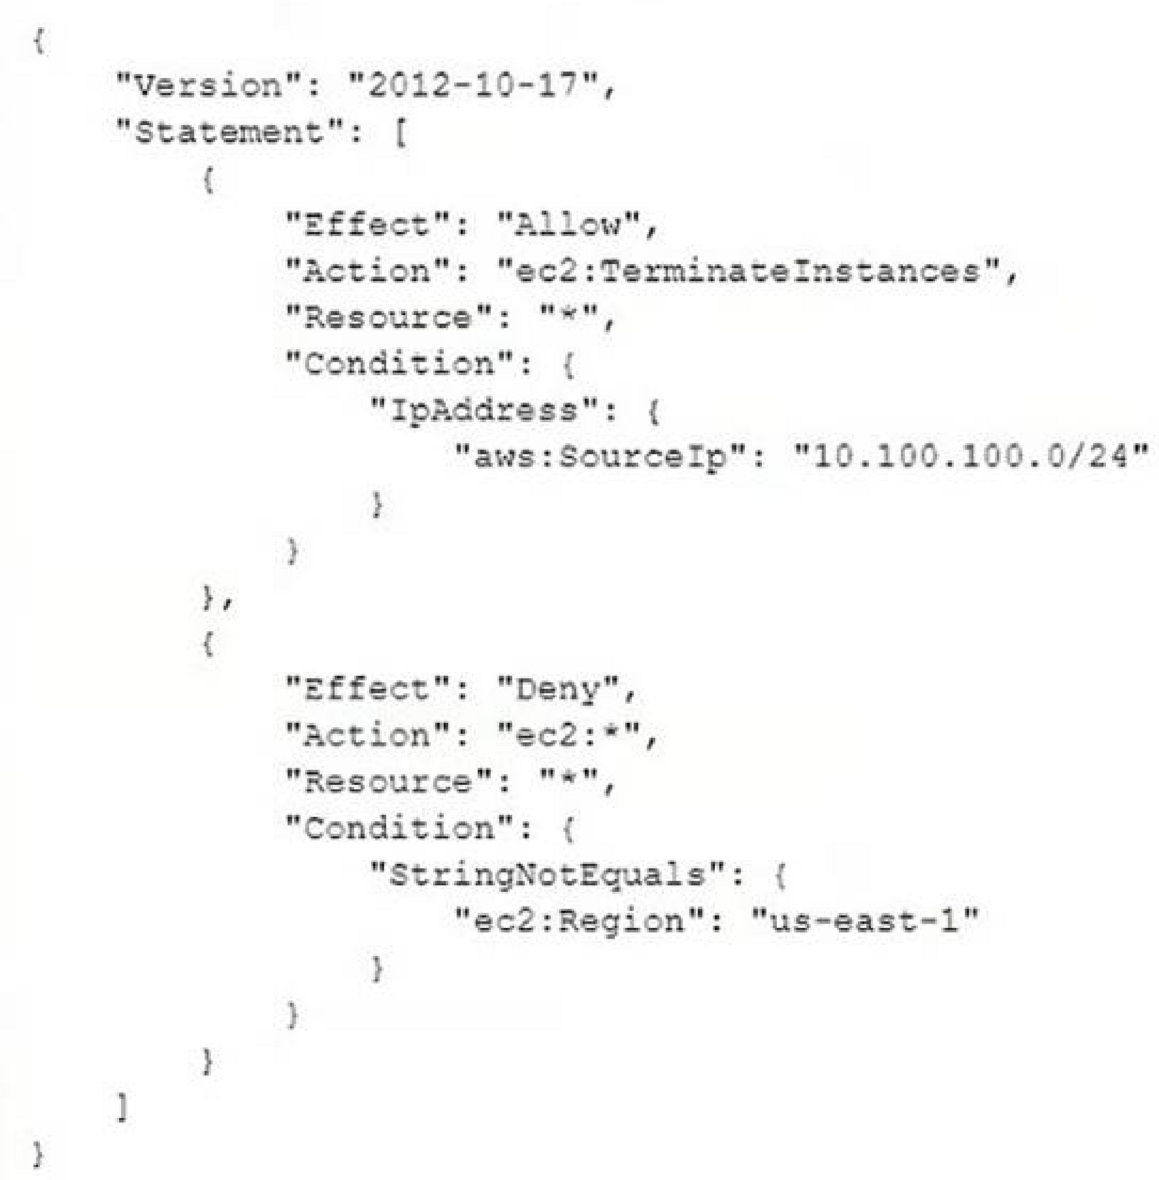

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": "ec2:TerminateInstances",
            "Resource": "*",
            "Condition": {
                "IpAddress": {
                    "aws:SourceIp": "10.100.100.0/24"
                }
            }
        },
        {
            "Effect": "Deny",
            "Action": "ec2:*",
            "Resource": "*",
            "Condition": {
                "StringNotEquals": {
                    "ec2:Region": "us-east-1"
                }
            }
        }
    ]
}
```

What is the effect of this policy?

A. Users can terminate an EC2 instance in any AWS Region except us-east-1.
B. Users can terminate an EC2 instance with the IP address 10 100 100 1 in the us-east-1 Region
C. Users can terminate an EC2 instance in the us-east-1 Region when the user's source IP is 10.100.100.254.
D. Users cannot terminate an EC2 instance in the us-east-1 Region when the user's source IP is 10.100 100 254

Answer: C

Explanation:

as the policy prevents anyone from doing any EC2 action on any region except us-east-1 and allows only users with source ip 10.100.100.0/24 to terminate instances. So user with source ip 10.100.100.254 can terminate instances in us-east-1 region.

NEW QUESTION 40

- (Topic 1)

A solutions architect is designing a new hybrid architecture to extend a company s on- premises infrastructure to AWS The company requires a highly available connection with consistent low latency to an AWS Region. The company needs to minimize costs and is willing to accept slower traffic if the primary connection fails.
What should the solutions architect do to meet these requirements?

A. Provision an AWS Direct Connect connection to a Region Provision a VPN connection as a backup if the primary Direct Connect connection fails.
B. Provision a VPN tunnel connection to a Region for private connectivit
C. Provision a second VPN tunnel for private connectivity and as a backup if the primary VPN connection fails.
D. Provision an AWS Direct Connect connection to a Region Provision a second Direct Connect connection to the same Region as a backup if the primary Direct Connect connection fails.
E. Provision an AWS Direct Connect connection to a Region Use the Direct Connect failover attribute from the AWS CLI to automatically create a backup connection if the primary Direct Connect connection fails.

Answer: A

Explanation:

"In some cases, this connection alone is not enough. It is always better to guarantee a fallback connection as the backup of DX. There are several options, but implementing it with an AWS Site-To-Site VPN is a real cost-effective solution that can be exploited to reduce costs or, in the meantime, wait for the setup of a second DX." https://www.proud2becloud.com/hybrid-cloud-networking-backup-aws-direct-connect-network-connection-with-aws-site-to-site-vpn/

NEW QUESTION 43

- (Topic 1)

A company has an Amazon S3 bucket that contains critical data. The company must protect the data from accidental deletion.
Which combination of steps should a solutions architect take to meet these requirements?

(Choose two.)

A. Enable versioning on the S3 bucket.
B. Enable MFA Delete on the S3 bucket.
C. Create a bucket policy on the S3 bucket.
D. Enable default encryption on the S3 bucket.
E. Create a lifecycle policy for the objects in the S3 bucket.

Answer: AB

Explanation:
To protect data in an S3 bucket from accidental deletion, versioning should be enabled, which enables you to preserve, retrieve, and restore every version of every object in an S3 bucket. Additionally, enabling MFA (multi-factor authentication) Delete on the S3 bucket adds an extra layer of protection by requiring an authentication token in addition to the user's access keys to delete objects in the bucket.
Reference:
AWS S3 Versioning documentation: https://docs.aws.amazon.com/AmazonS3/latest/dev/Versioning.html
AWS S3 MFA Delete documentation: https://docs.aws.amazon.com/AmazonS3/latest/dev/UsingMFADelete.html

NEW QUESTION 44
- (Topic 1)
A company has more than 5 TB of file data on Windows file servers that run on premises Users and applications interact with the data each day
The company is moving its Windows workloads to AWS. As the company continues this process, the company requires access to AWS and on-premises file storage with minimum latency The company needs a solution that minimizes operational overhead and requires no significant changes to the existing file access patterns. The company uses an AWS Site-to-Site VPN connection for connectivity to AWS
What should a solutions architect do to meet these requirements?

A. Deploy and configure Amazon FSx for Windows File Server on AW
B. Move the on- premises file data to FSx for Windows File Serve
C. Reconfigure the workloads to use FSx for Windows File Server on AWS.
D. Deploy and configure an Amazon S3 File Gateway on premises Move the on-premises file data to the S3 File Gateway Reconfigure the on-premises workloads and the cloud workloads to use the S3 File Gateway
E. Deploy and configure an Amazon S3 File Gateway on premises Move the on-premises file data to Amazon S3 Reconfigure the workloads to use either Amazon S3 directly or the S3 File Gateway, depending on each workload's location
F. Deploy and configure Amazon FSx for Windows File Server on AWS Deploy and configure an Amazon FSx File Gateway on premises Move the on-premises file data to the FSx File Gateway Configure the cloud workloads to use FSx for Windows File Server on AWS Configure the on-premises workloads to use the FSx File Gateway

Answer: D

Explanation:
https://docs.aws.amazon.com/filegateway/latest/filefsxw/what-is-file-fsxw.html
To meet the requirements of the company to have access to both AWS and on-premises file storage with minimum latency, a hybrid cloud architecture can be used. One solution is to deploy and configure Amazon FSx for Windows File Server on AWS, which provides fully managed Windows file servers. The on-premises file data can be moved to the FSx File Gateway, which can act as a bridge between on-premises and AWS file storage. The cloud workloads can be configured to use FSx for Windows File Server on AWS, while the on-premises workloads can be configured to use the FSx File Gateway. This solution minimizes operational overhead and requires no significant changes to the existing file access patterns. The connectivity between on-premises and AWS can be established using an AWS Site-to-Site VPN connection.
Reference:
AWS FSx for Windows File Server: https://aws.amazon.com/fsx/windows/ AWS FSx File Gateway: https://aws.amazon.com/fsx/file-gateway/
AWS Site-to-Site VPN: https://aws.amazon.com/vpn/site-to-site-vpn/

NEW QUESTION 47
- (Topic 1)
A company that hosts its web application on AWS wants to ensure all Amazon EC2 instances. Amazon RDS DB instances. and Amazon Redshift clusters are configured with tags. The company wants to minimize the effort of configuring and operating this check.
What should a solutions architect do to accomplish this?

A. Use AWS Config rules to define and detect resources that are not properly tagged.
B. Use Cost Explorer to display resources that are not properly tagge
C. Tag those resources manually.
D. Write API calls to check all resources for proper tag allocatio
E. Periodically run the code on an EC2 instance.
F. Write API calls to check all resources for proper tag allocatio
G. Schedule an AWS Lambda function through Amazon CloudWatch to periodically run the code.

Answer: A

Explanation:
To ensure all Amazon EC2 instances, Amazon RDS DB instances, and Amazon Redshift clusters are configured with tags, a solutions architect should use AWS Config rules to define and detect resources that are not properly tagged. AWS Config rules are a set of customizable rules that AWS Config uses to evaluate AWS resource configurations for compliance with best practices and company policies. Using AWS Config rules can minimize the effort of configuring and operating this check because it automates the process of identifying non-compliant resources and notifying the responsible teams. Reference:
AWS Config Developer Guide: AWS Config Rules (https://docs.aws.amazon.com/config/latest/developerguide/evaluate-config_use-managed- rules.html)

NEW QUESTION 50
- (Topic 1)
A company is preparing to deploy a new serverless workload. A solutions architect must use the principle of least privilege to configure permissions that will be used to run an AWS Lambda function. An Amazon EventBridge (Amazon CloudWatch Events) rule will invoke the function.
Which solution meets these requirements?

A. Add an execution role to the function with lambda: InvokeFunction as the action and * asthe principal.
B. Add an execution role to the function with lambda: InvokeFunction as the action and Service:amazonaws.com as the principal.
C. Add a resource-based policy to the function with lambda:'* as the action and Service:events.amazonaws.com as the principal.
D. Add a resource-based policy to the function with lambda: InvokeFunction as the action and Service:events.amazonaws.com as the principal.

Answer: D

Explanation:
https://docs.aws.amazon.com/eventbridge/latest/userguide/resource-based- policies-eventbridge.html#lambda-permissions

NEW QUESTION 53
- (Topic 1)
A solutions architect is developing a multiple-subnet VPC architecture. The solution will consist of six subnets in two Availability Zones. The subnets are defined as public, private and dedicated for databases. Only the Amazon EC2 instances running in the private subnets should be able to access a database.
Which solution meets these requirements?

A. Create a now route table that excludes the route to the public subnets' CIDR block
B. Associate the route table to the database subnets.
C. Create a security group that denies ingress from the security group used by instances in the public subnet
D. Attach the security group to an Amazon RDS DB instance.
E. Create a security group that allows ingress from the security group used by instances in the private subnet
F. Attach the security group to an Amazon RDS DB instance.
G. Create a new peering connection between the public subnets and the private subnet
H. Create a different peering connection between the private subnets and the databasesubnets.

Answer: C

Explanation:
Security groups are stateful. All inbound traffic is blocked by default. If you create an inbound rule allowing traffic in, that traffic is automatically allowed back out again. You cannot block specific IP address using Security groups (instead use Network Access Control Lists).
"You can specify allow rules, but not deny rules." "When you first create a security group, it has no inbound rules. Therefore, no inbound traffic originating from another host to your instance is allowed until you add inbound rules to the security group." Source:
https://docs.aws.amazon.com/vpc/latest/userguide/VPC_SecurityGroups.html#VPCSecurit yGroups

NEW QUESTION 57
- (Topic 1)
A company is using a SQL database to store movie data that is publicly accessible. The database runs on an Amazon RDS Single-AZ DB instance A script runs queries at random intervals each day to record the number of new movies that have been added to the database. The script must report a final total during business hours The company's development team notices that the database performance is inadequate for development tasks when the script is running. A solutions architect must recommend a solution to resolve this issue. Which solution will meet this requirement with the LEAST operational overhead?

A. Modify the DB instance to be a Multi-AZ deployment
B. Create a read replica of the database Configure the script to query only the read replica
C. Instruct the development team to manually export the entries in the database at the end of each day
D. Use Amazon ElastiCache to cache the common queries that the script runs against the database

Answer: B

NEW QUESTION 61
- (Topic 1)
An application development team is designing a microservice that will convert large images to smaller, compressed images. When a user uploads an image through the web interface, the microservice should store the image in an Amazon S3 bucket, process and compress the image with an AWS Lambda function, and store the image in its compressed form in a different S3 bucket.
A solutions architect needs to design a solution that uses durable, stateless components to process the images automatically.
Which combination of actions will meet these requirements? (Choose two.)

A. Create an Amazon Simple Queue Service (Amazon SQS) queue Configure the S3 bucket to send a notification to the SQS queue when an image is uploaded to the S3 bucket
B. Configure the Lambda function to use the Amazon Simple Queue Service (Amazon SQS) queue as the invocation source When the SQS message is successfully processed, delete the message in the queue
C. Configure the Lambda function to monitor the S3 bucket for new uploads When an uploaded image is detected write the file name to a text file in memory and use the text file to keep track of the images that were processed
D. Launch an Amazon EC2 instance to monitor an Amazon Simple Queue Service(Amazon SQS) queue When items are added to the queue log the file name in a text file on the EC2 instance and invoke the Lambda function
E. Configure an Amazon EventBridge (Amazon CloudWatch Events) event to monitor the S3 bucket When an image is uploade
F. send an alert to an Amazon Simple Notification Service (Amazon SNS) topic with the application owner's email address for further processing

Answer: AB

Explanation:
? Creating an Amazon Simple Queue Service (SQS) queue and configuring the S3 bucket to send a notification to the SQS queue when an image is uploaded to the S3 bucket will ensure that the Lambda function is triggered in a stateless and durable manner.
? Configuring the Lambda function to use the SQS queue as the invocation source, and deleting the message in the queue after it is successfully processed will ensure that the Lambda function processes the image in a stateless and durable manner.
Amazon SQS is a fully managed message queuing service that enables you to decouple and scale microservices, distributed systems, and serverless applications. SQS eliminates the complexity and overhead associated with managing and operating-message oriented middleware, and empowers developers to focus on differentiating work. When new images are uploaded to the S3 bucket, SQS will trigger the Lambda function to process the image and compress it. Once the image is processed, the SQS message is deleted, ensuring that the Lambda function is stateless and durable.

NEW QUESTION 64
- (Topic 1)
A company is preparing to store confidential data in Amazon S3 For compliance reasons the data must be encrypted at rest Encryption key usage must be logged tor auditing purposes. Keys must be rotated every year.
Which solution meets these requirements and «the MOST operationally efferent?

A. Server-side encryption with customer-provided keys (SSE-C)
B. Server-side encryption with Amazon S3 managed keys (SSE-S3)
C. Server-side encryption with AWS KMS (SSE-KMS) customer master keys (CMKs) with manual rotation
D. Server-side encryption with AWS KMS (SSE-KMS) customer master keys (CMKs) with automate rotation

Answer: D

Explanation:
https://docs.aws.amazon.com/kms/latest/developerguide/rotate-keys.html When you enable automatic key rotation for a customer managed key, AWS KMS generates new cryptographic material for the KMS key every year. AWS KMS also saves the KMS key's older cryptographic material in perpetuity so it can be used to decrypt data that the KMS key encrypted.
Key rotation in AWS KMS is a cryptographic best practice that is designed to be transparent and easy to use. AWS KMS supports optional automatic key rotation only for customer managed CMKs. Enable and disable key rotation. Automatic key rotation is disabled by default on customer managed CMKs. When you enable (or re-enable) key rotation, AWS KMS automatically rotates the CMK 365 days after the enable date and every 365 days thereafter.

NEW QUESTION 66
- (Topic 1)
An application allows users at a company's headquarters to access product data. The product data is stored in an Amazon RDS MySQL DB instance. The operations team has isolated an application performance slowdown and wants to separate read traffic from write traffic. A solutions architect needs to optimize the application's performance quickly.
What should the solutions architect recommend?

A. Change the existing database to a Multi-AZ deploymen
B. Serve the read requests from the primary Availability Zone.
C. Change the existing database to a Multi-AZ deploymen
D. Serve the read requests from the secondary Availability Zone.
E. Create read replicas for the databas
F. Configure the read replicas with half of the compute and storage resources as the source database.
G. Create read replicas for the databas
H. Configure the read replicas with the same compute and storage resources as the source database.

Answer: D

Explanation:
https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_MySQL.Replication.R eadReplicas.html

NEW QUESTION 70
- (Topic 1)
A solutions architect is designing a VPC with public and private subnets. The VPC and subnets use IPv4 CIDR blocks. There is one public subnet and one private subnet in each of three Availability Zones (AZs) for high availability. An internet gateway is used to provide internet access for the public subnets. The private subnets require access to the internet to allow Amazon EC2 instances to download software updates.
What should the solutions architect do to enable Internet access for the private subnets?

A. Create three NAT gateways, one for each public subnet in each A
B. Create a private route table for each AZ that forwards non-VPC traffic to the NAT gateway in its AZ.
C. Create three NAT instances, one for each private subnet in each A
D. Create a private route table for each AZ that forwards non-VPC traffic to the NAT instance in its AZ.
E. Create a second internet gateway on one of the private subnet
F. Update the route table for the private subnets that forward non-VPC traffic to the private internet gateway.
G. Create an egress-only internet gateway on one of the public subnet
H. Update the route table for the private subnets that forward non-VPC traffic to the egress- only internet gateway.

Answer: A

Explanation:
https://aws.amazon.com/about-aws/whats-new/2018/03/introducing-amazon-vpc-nat-gateway-in-the-aws-govcloud-us-region/#:~:text=NAT%20Gateway%20is%20a%20highly,instances%20in%20a%20private
%20subnet.
https://docs.aws.amazon.com/vpc/latest/userguide/vpc-nat-comparison.html

NEW QUESTION 72
- (Topic 1)
A company has an automobile sales website that stores its listings in a database on Amazon RDS When an automobile is sold the listing needs to be removed from the website and the data must be sent to multiple target systems.
Which design should a solutions architect recommend?

A. Create an AWS Lambda function triggered when the database on Amazon RDS is updated to send the information to an Amazon Simple Queue Service (Amazon SQS> queue for the targets to consume
B. Create an AWS Lambda function triggered when the database on Amazon RDS is updated to send the information to an Amazon Simple Queue Service (Amazon SQS) FIFO queue for the targets to consume
C. Subscribe to an RDS event notification and send an Amazon Simple Queue Service (Amazon SQS) queue fanned out to multiple Amazon Simple Notification Service (Amazon SNS) topics Use AWS Lambda functions to update the targets
D. Subscribe to an RDS event notification and send an Amazon Simple Notification Service (Amazon SNS) topic fanned out to multiple Amazon Simple Queue Service (Amazon SQS) queues Use AWS Lambda functions to update the targets

Answer: D

Explanation:
https://docs.aws.amazon.com/lambda/latest/dg/services-rds.html https://docs.aws.amazon.com/lambda/latest/dg/with-sns.html

NEW QUESTION 73
- (Topic 1)
A company has a production web application in which users upload documents through a web interlace or a mobile app. According to a new regulatory requirement, new documents cannot be modified or deleted after they are stored.
What should a solutions architect do to meet this requirement?

A. Store the uploaded documents in an Amazon S3 bucket with S3 Versioning and S3 Object Lock enabled
B. Store the uploaded documents in an Amazon S3 bucke
C. Configure an S3 Lifecycle policy to archive the documents periodically.
D. Store the uploaded documents in an Amazon S3 bucket with S3 Versioning enabled Configure an ACL to restrict all access to read-only.
E. Store the uploaded documents on an Amazon Elastic File System (Amazon EFS) volum
F. Access the data by mounting the volume in read-only mode.

Answer: A

Explanation:
https://docs.aws.amazon.com/AmazonS3/latest/userguide/object-lock-overview.html

NEW QUESTION 75
- (Topic 1)
A company has an application that runs on Amazon EC2 instances and uses an Amazon
Aurora database. The EC2 instances connect to the database by using user names and passwords that are stored locally in a file. The company wants to minimize the operational overhead of credential management.
What should a solutions architect do to accomplish this goal?

A. Use AWS Secrets Manage
B. Turn on automatic rotation.
C. Use AWS Systems Manager Parameter Stor
D. Turn on automatic rotation.
E. Create an Amazon S3 bucket lo store objects that are encrypted with an AWS Key
F. Management Service (AWS KMS) encryption ke
G. Migrate the credential file to the S3 bucke
H. Point the application to the S3 bucket.
I. Create an encrypted Amazon Elastic Block Store (Amazon EBS) volume (or each EC2 instanc
J. Attach the new EBS volume to each EC2 instanc
K. Migrate the credential file to the new EBS volum
L. Point the application to the new EBS volume.

Answer: A

Explanation:
https://aws.amazon.com/cn/blogs/security/how-to-connect-to-aws-secrets-manager-service-within-a-virtual-private-cloud/
https://aws.amazon.com/blogs/security/rotate-amazon-rds-database-credentials-automatically-with-aws-secrets-manager/

NEW QUESTION 77
- (Topic 1)
A company uses AWS Organizations to manage multiple AWS accounts for different departments. The management account has an Amazon S3 bucket that contains project reports. The company wants to limit access to this S3 bucket to only users of accounts within the organization in AWS Organizations.
Which solution meets these requirements with the LEAST amount of operational overhead?

A. Add the aws:PrincipalOrgID global condition key with a reference to the organization ID to the S3 bucket policy.
B. Create an organizational unit (OU) for each departmen
C. Add the aws:PrincipalOrgPaths global condition key to the S3 bucket policy.
D. Use AWS CloudTrail to monitor the CreateAccount, InviteAccountToOrganization, LeaveOrganization, and RemoveAccountFromOrganization event
E. Update the S3 bucket policy accordingly.
F. Tag each user that needs access to the S3 bucke
G. Add the aws:PrincipalTag global condition key to the S3 bucket policy.

Answer: A

Explanation:
https://aws.amazon.com/blogs/security/control-access-to-aws-resources-by-using-the-aws-organization-of-iam-principals/
The aws:PrincipalOrgID global key provides an alternative to listing all the account IDs for all AWS accounts in an organization. For example, the following Amazon S3 bucket policy allows members of any account in the XXX organization to add an object into the
examtopics bucket.

```
{"Version": "2020-09-10",
"Statement": {
"Sid": "AllowPutObject", "Effect": "Allow",
"Principal": "*",
"Action": "s3:PutObject",
"Resource": "arn:aws:s3:::examtopics/*", "Condition": {"StringEquals":
{"aws:PrincipalOrgID":["XXX"]}}}}
```

https://docs.aws.amazon.com/IAM/latest/UserGuide/reference_policies_condition- keys.html

NEW QUESTION 78
- (Topic 1)
A company is running a business-critical web application on Amazon EC2 instances behind an Application Load Balancer. The EC2 instances are in an Auto Scaling group. The application uses an Amazon Aurora PostgreSQL database that is deployed in a single Availability Zone. The company wants the application to be highly available with minimum downtime and minimum loss of data.
Which solution will meet these requirements with the LEAST operational effort?

A. Place the EC2 instances in different AWS Region
B. Use Amazon Route 53 health checks to redirect traffi
C. Use Aurora PostgreSQL Cross-Region Replication.
D. Configure the Auto Scaling group to use multiple Availability Zone
E. Configure the database as Multi-A
F. Configure an Amazon RDS Proxy instance for the database.
G. Configure the Auto Scaling group to use one Availability Zon
H. Generate hourly snapshots of the databas
I. Recover the database from the snapshots in the event of a failure.
J. Configure the Auto Scaling group to use multiple AWS Region
K. Write the data from the application to Amazon S3. Use S3 Event Notifications to launch an AWS Lambda function to write the data to the database.

Answer: B

Explanation:
To achieve high availability with minimum downtime and minimum loss of data, the Auto Scaling group should be configured to use multiple Availability Zones to ensure that there is no single point of failure. The database should be configured as Multi- AZ to enable automatic failover in case of an outage in the primary Availability Zone. Additionally, an Amazon RDS Proxy instance can be used to improve the scalability and availability of the database by reducing connection failures and improving failover times.

NEW QUESTION 83
- (Topic 1)
A company recently launched a variety of new workloads on Amazon EC2 instances in its AWS account. The company needs to create a strategy to access and administer the instances remotely and securely. The company needs to implement a repeatable process that works with native AWS services and follows the AWS Well-Architected Framework.
Which solution will meet these requirements with the LEAST operational overhead?

A. Use the EC2 serial console to directly access the terminal interface of each instance foradministration.
B. Attach the appropriate IAM role to each existing instance and new instanc
C. Use AWS Systems Manager Session Manager to establish a remote SSH session.
D. Create an administrative SSH key pai
E. Load the public key into each EC2 instanc
F. Deploy a bastion host in a public subnet to provide a tunnel for administration of each instance.
G. Establish an AWS Site-to-Site VPN connectio
H. Instruct administrators to use their local on-premises machines to connect directly to the instances by using SSH keys across the VPN tunnel.

Answer: B

Explanation:
https://docs.aws.amazon.com/systems-manager/latest/userguide/setup-launch-managed-instance.html

NEW QUESTION 88
- (Topic 1)
A hospital recently deployed a RESTful API with Amazon API Gateway and AWS Lambda The hospital uses API Gateway and Lambda to upload reports that are in PDF format and JPEG format The hospital needs to modify the Lambda code to identify protected health information (PHI) in the reports
Which solution will meet these requirements with the LEAST operational overhead?

A. Use existing Python libraries to extract the text from the reports and to identify the PHI from the extracted text.
B. Use Amazon Textract to extract the text from the reports Use Amazon SageMaker to identify the PHI from the extracted text.
C. Use Amazon Textract to extract the text from the reports Use Amazon Comprehend Medical to identify the PHI from the extracted text
D. Use Amazon Rekognition to extract the text from the reports Use Amazon Comprehend Medical to identify the PHI from the extracted text

Answer: C

Explanation:
To meet the requirements of the company to have access to both AWS and on-premises file storage with minimum latency, a hybrid cloud architecture can be used. One solution is to deploy and configure Amazon FSx for Windows File Server on AWS, which provides fully managed Windows file servers. The on-premises file data can be moved to the FSx File Gateway, which can act as a bridge between on-premises and AWS file storage. The cloud workloads can be configured to use FSx for Windows File Server on AWS, while the on-premises workloads can be configured to use the FSx File Gateway. This solution minimizes operational overhead and requires no significant changes to the
existing file access patterns. The connectivity between on-premises and AWS can be established using an AWS Site-to-Site VPN connection.
Reference:
AWS FSx for Windows File Server: https://aws.amazon.com/fsx/windows/ AWS FSx File Gateway: https://aws.amazon.com/fsx/file-gateway/
AWS Site-to-Site VPN: https://aws.amazon.com/vpn/site-to-site-vpn/

NEW QUESTION 92
- (Topic 2)
A company hosts a two-tier application on Amazon EC2 instances and Amazon RDS. The application's demand varies based on the time of day. The load is minimal after work hours and on weekends. The EC2 instances run in an EC2 Auto Scaling group that is configured with a minimum of two instances and a maximum of five instances. The application must be available at all times, but the company is concerned about overall cost.
Which solution meets the availability requirement MOST cost-effectively?

A. Use all EC2 Spot Instance

B. Stop the RDS database when it is not in use.
C. Purchase EC2 Instance Savings Plans to cover five EC2 instance
D. Purchase an RDS Reserved DB Instance
E. Purchase two EC2 Reserved Instances Use up to three additional EC2 Spot Instances as neede
F. Stop the RDS database when it is not in use.
G. Purchase EC2 Instance Savings Plans to cover two EC2 instance
H. Use up to three additional EC2 On-Demand Instances as neede
I. Purchase an RDS Reserved DB Instance.

Answer: C

Explanation:
This solution meets the requirements of a two-tier application that has a variable demand based on the time of day and must be available at all times, while minimizing the overall cost. EC2 Reserved Instances can provide significant savings compared to On-Demand Instances for the baseline level of usage, and they can guarantee capacity reservation when needed. EC2 Spot Instances can provide up to 90% savings compared to On- Demand Instances for any additional capacity that the application needs during peak hours. Spot Instances are suitable for stateless applications that can tolerate interruptions and can be replaced by other instances. Stopping the RDS database when it is not in use can reduce the cost of running the database tier.
Option A is incorrect because using all EC2 Spot Instances can affect the availability of the application if there are not enough spare capacity or if the Spot price exceeds the maximum price. Stopping the RDS database when it is not in use can reduce the cost of running the database tier, but it can also affect the availability of the application. Option B is incorrect because purchasing EC2 Instance Savings Plans to cover five EC2 instances can lock in a fixed amount of compute usage per hour, which may not match the actual usage pattern of the application. Purchasing an RDS Reserved DB Instance can provide savings for the database tier, but it does not allow stopping the database when it is not in use. Option D is incorrect because purchasing EC2 Instance Savings Plans to cover two EC2 instances can lock in a fixed amount of compute usage per hour, which may not match the
actual usage pattern of the application. Using up to three additional EC2 On-Demand Instances as needed can incur higher costs than using Spot Instances.
References:
? https://aws.amazon.com/ec2/pricing/reserved-instances/
? https://aws.amazon.com/ec2/spot/
? https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_StopInstance.html

NEW QUESTION 94
- (Topic 2)
A company has a Windows-based application that must be migrated to AWS. The application requires the use of a shared Windows file system attached to multiple Amazon EC2 Windows instances that are deployed across multiple Availability Zones.
What should a solutions architect do to meet this requirement?

A. Configure AWS Storage Gateway in volume gateway mod
B. Mount the volume to each Windows instance.
C. Configure Amazon FSx for Windows File Serve
D. Mount the Amazon FSx file system to each Windows instance.
E. Configure a file system by using Amazon Elastic File System (Amazon EFS). Mount the EFS file system to each Windows instance.
F. Configure an Amazon Elastic Block Store (Amazon EBS) volume with the required siz
G. Attach each EC2 instance to the volum
H. Mount the file system within the volume to each Windows instance.

Answer: B

Explanation:
This solution meets the requirement of migrating a Windows-based application that requires the use of a shared Windows file system attached to multiple Amazon EC2 Windows instances that are deployed across multiple Availability Zones. Amazon FSx for Windows File Server provides fully managed shared storage built on Windows Server, and delivers a wide range of data access, data management, and administrative capabilities. It supports the Server Message Block (SMB) protocol and can be mounted to EC2 Windows instances across multiple Availability Zones.
Option A is incorrect because AWS Storage Gateway in volume gateway mode provides cloud-backed storage volumes that can be mounted as iSCSI devices from on-premises application servers, but it does not support SMB protocol or EC2 Windows instances. Option C is incorrect because Amazon Elastic File System (Amazon EFS) provides a scalable and elastic NFS file system for Linux-based workloads, but it does not support SMB protocol or EC2 Windows instances. Option D is incorrect because Amazon Elastic Block Store (Amazon EBS) provides persistent block storage volumes for use with EC2 instances, but it does not support SMB protocol or attaching multiple instances to the same volume.
References:
? https://aws.amazon.com/fsx/windows/
? https://docs.aws.amazon.com/fsx/latest/WindowsGuide/using-file-shares.html

NEW QUESTION 96
- (Topic 2)
A company wants to direct its users to a backup static error page if the company's primary website is unavailable. The primary website's DNS records are hosted in Amazon Route 53. The domain is pointing to an Application Load Balancer (ALB). The company needs a solution that minimizes changes and infrastructure overhead.
Which solution will meet these requirements?

A. Update the Route 53 records to use a latency routing polic
B. Add a static error page that is hosted in an Amazon S3 bucket to the records so that the traffic is sent to the most responsive endpoints.
C. Set up a Route 53 active-passive failover configuratio
D. Direct traffic to a static error page that is hosted in an Amazon S3 bucket when Route 53 health checks determine that the ALB endpoint is unhealthy.
E. Set up a Route 53 active-active configuration with the ALB and an Amazon EC2 instance that hosts a static error page as endpoint
F. Configure Route 53 to send requests to the instance only if the health checks fail for the ALB.
G. Update the Route 53 records to use a multivalue answer routing polic
H. Create a health chec
I. Direct traffic to the website if the health check passe
J. Direct traffic to a static error page that is hosted in Amazon S3 if the health check does not pass.

Answer: B

Explanation:

This solution meets the requirements of directing users to a backup static error page if the primary website is unavailable, minimizing changes and infrastructure overhead. Route 53 active-passive failover configuration can route traffic to a primary resource when it is healthy or to a secondary resource when the primary resource is unhealthy. Route 53 health checks can monitor the health of the ALB endpoint and trigger the failover when needed. The static error page can be hosted in an S3 bucket that is configured as a website, which is a simple and cost-effective way to serve static content.
Option A is incorrect because using a latency routing policy can route traffic based on the lowest network latency for users, but it does not provide failover functionality. Option C is incorrect because using an active-active configuration with the ALB and an EC2 instance can increase the infrastructure overhead and complexity, and it does not guarantee that the EC2 instance will always be healthy. Option D is incorrect because using a multivalue answer routing policy can return multiple values for a query, but it does not provide failover functionality.
References:
? https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-policy- failover.html
? https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/dns-failover.html
? https://docs.aws.amazon.com/AmazonS3/latest/userguide/WebsiteHosting.html

NEW QUESTION 101
- (Topic 2)
An ecommerce company hosts its analytics application in the AWS Cloud. The application generates about 300 MB of data each month. The data is stored in JSON format. The company is evaluating a disaster recovery solution to back up the data. The data must be accessible in milliseconds if it is needed, and the data must be kept for 30 days.
Which solution meets these requirements MOST cost-effectively?

A. Amazon OpenSearch Service (Amazon Elasticsearch Service)
B. Amazon S3 Glacier
C. Amazon S3 Standard
D. Amazon RDS for PostgreSQL

Answer: C

Explanation:
This solution meets the requirements of a disaster recovery solution to back up the data that is generated by an analytics application, stored in JSON format, and must be accessible in milliseconds if it is needed. Amazon S3 Standard is a durable and scalable storage class for frequently accessed data. It can store any amount of data and provide high availability and performance. It can also support millisecond access time for data retrieval.
Option A is incorrect because Amazon OpenSearch Service (Amazon Elasticsearch Service) is a search and analytics service that can index and query data, but it is not a backup solution for data stored in JSON format. Option B is incorrect because Amazon S3 Glacier is a low-cost storage class for data archiving and long-term backup, but it does not support millisecond access time for data retrieval. Option D is incorrect because Amazon RDS for PostgreSQL is a relational database service that can store and query structured data, but it is not a backup solution for data stored in JSON format.
References:
? https://aws.amazon.com/s3/storage-classes/
? https://aws.amazon.com/s3/faqs/#Durability_and_data_protection

NEW QUESTION 102
- (Topic 2)
A company uses a popular content management system (CMS) for its corporate website. However, the required patching and maintenance are burdensome. The company is redesigning its website and wants anew solution. The website will be updated four times a year and does not need to have any dynamic content available. The solution must provide high scalability and enhanced security.
Which combination of changes will meet these requirements with the LEAST operational overhead? (Choose two.)

A. Deploy an AWS WAF web ACL in front of the website to provide HTTPS functionality
B. Create and deploy an AWS Lambda function to manage and serve the website content
C. Create the new website and an Amazon S3 bucket Deploy the website on the S3 bucket with static website hosting enabled
D. Create the new websit
E. Deploy the website by using an Auto Scaling group of Amazon EC2 instances behind an Application Load Balancer.

Answer: AD

Explanation:
A -> We can configure CloudFront to require HTTPS from clients (enhanced security) https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/using- https-viewers-to-cloudfront.html D -> storing static website on S3 provides scalability and less operational overhead, then configuration of Application LB and EC2 instances (hence E is out)

NEW QUESTION 107
- (Topic 2)
A security team wants to limit access to specific services or actions in all of the team's AWS accounts. All accounts belong to a large organization in AWS Organizations. The solution must be scalable and there must be a single point where permissions can be maintained.
What should a solutions architect do to accomplish this?

A. Create an ACL to provide access to the services or actions.
B. Create a security group to allow accounts and attach it to user groups.
C. Create cross-account roles in each account to deny access to the services or actions.
D. Create a service control policy in the root organizational unit to deny access to the services or actions.

Answer: D

Explanation:
Service control policies (SCPs) are one type of policy that you can use to manage your organization. SCPs offer central control over the maximum available permissions for all accounts in your organization, allowing you to ensure your accounts stay within your organization's access control guidelines. See https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scp.ht
ml.

NEW QUESTION 110

- (Topic 2)
An online retail company has more than 50 million active customers and receives more than 25,000 orders each day. The company collects purchase data for customers and stores this data in Amazon S3. Additional customer data is stored in Amazon RDS.
The company wants to make all the data available to various teams so that the teams can perform analytics. The solution must provide the ability to manage fine-grained permissions for the data and must minimize operational overhead.
Which solution will meet these requirements?

A. Migrate the purchase data to write directly to Amazon RD
B. Use RDS access controls to limit access.
C. Schedule an AWS Lambda function to periodically copy data from Amazon RDS to Amazon S3. Create an AWS Glue crawle
D. Use Amazon Athena to query the dat
E. Use S3 policies to limit access.
F. Create a data lake by using AWS Lake Formatio
G. Create an AWS Glue JDBC connection to Amazon RD
H. Register (he S3 bucket in Lake Formatio
I. Use Lake Formation access controls to limit access.
J. Create an Amazon Redshift cluste
K. Schedule an AWS Lambda function to periodically copy data from Amazon S3 and Amazon RDS to Amazon Redshif
L. Use Amazon Redshift access controls to limit access.

Answer: C

Explanation:
To make all the data available to various teams and minimize operational overhead, the company can create a data lake by using AWS Lake Formation. This will allow the company to centralize all the data in one place and use fine-grained access controls to manage access to the data. To meet the requirements of the company, the solutions architect can create a data lake by using AWS Lake Formation, create an AWS Glue JDBC connection to Amazon RDS, and register the S3 bucket in Lake Formation. The solutions architect can then use Lake Formation access controls to limit access to the data. This solution will provide the ability to manage fine-grained permissions for the data and minimize operational overhead.

NEW QUESTION 115
- (Topic 2)
A company has an ecommerce checkout workflow that writes an order to a database and calls a service to process the payment. Users are experiencing timeouts during the checkout process. When users resubmit the checkout form, multiple unique orders are created for the same desired transaction.
How should a solutions architect refactor this workflow to prevent the creation of multiple orders?

A. Configure the web application to send an order message to Amazon Kinesis Data Firehos
B. Set the payment service to retrieve the message from Kinesis Data Firehose and process the order.
C. Create a rule in AWS CloudTrail to invoke an AWS Lambda function based on the logged application path request Use Lambda to query the database, call the payment service, and pass in the order information.
D. Store the order in the databas
E. Send a message that includes the order number to Amazon Simple Notification Service (Amazon SNS). Set the payment service to pollAmazon SN
F. retrieve the message, and process the order.
G. Store the order in the databas
H. Send a message that includes the order number to an Amazon Simple Queue Service (Amazon SQS) FIFO queu
I. Set the payment service to retrieve the message and process the orde
J. Delete the message from the queue.

Answer: D

Explanation:
This approach ensures that the order creation and payment processing steps are separate and atomic. By sending the order information to an SQS FIFO queue, the payment service can process the order one at a time and in the order they were received. If the payment service is unable to process an order, it can be retried later, preventing the creation of multiple orders. The deletion of the message from the queue after it is processed will prevent the same message from being processed multiple times.

NEW QUESTION 117
- (Topic 2)
A company is building a containerized application on premises and decides to move the application to AWS. The application will have thousands of users soon after li is deployed. The company Is unsure how to manage the deployment of containers at scale. The company needs to deploy the containerized application in a highly available architecture that minimizes operational overhead.
Which solution will meet these requirements?

A. Store container images In an Amazon Elastic Container Registry (Amazon ECR) repositor
B. Use an Amazon Elastic Container Service (Amazon ECS) cluster with the AWS Fargate launch type to run the container
C. Use target tracking to scale automatically based on demand.
D. Store container images in an Amazon Elastic Container Registry (Amazon ECR) repositor
E. Use an Amazon Elastic Container Service (Amazon ECS) cluster with the Amazon EC2 launch type to run the container
F. Use target tracking to scale automatically based on demand.
G. Store container images in a repository that runs on an Amazon EC2 instanc
H. Run the containers on EC2 instances that are spread across multiple Availability Zone
I. Monitor the average CPU utilization in Amazon CloudWatc
J. Launch new EC2 instances as needed
K. Create an Amazon EC2 Amazon Machine Image (AMI) that contains the container image Launch EC2 Instances in an Auto Scaling group across multiple Availability Zone
L. Use an Amazon CloudWatch alarm to scale out EC2 instances when the average CPU utilization threshold is breached.

Answer: A

Explanation:
AWS Fargate is a serverless experience for user applications, allowing the user to concentrate on building applications instead of configuring and managing servers. Fargate also automates resource management, allowing users to easily scale their applications in response to demand.

NEW QUESTION 118
- (Topic 2)
A company is planning to move its data to an Amazon S3 bucket. The data must be encrypted when it is stored in the S3 bucket. Additionally, the encryption key must be automatically rotated every year.
Which solution will meet these requirements with the LEAST operational overhead?

A. Move the data to the S3 bucke
B. Use server-side encryption with Amazon S3 managed encryption keys (SSE-S3). Use the built-in key rotation behavior of SSE-S3 encryption keys.
C. Create an AWS Key Management Service {AWS KMS) customer managed ke
D. Enable automatic key rotatio
E. Set the S3 bucket's default encryption behavior to use the customer managed KMS ke
F. Move the data to the S3 bucket.
G. Create an AWS Key Management Service (AWS KMS) customer managed ke
H. Set the S3 bucket's default encryption behavior to use the customer managed KMS ke
I. Move the data to the S3 bucke
J. Manually rotate the KMS key every year.
K. Encrypt the data with customer key material before moving the data to the S3 bucke
L. Create an AWS Key Management Service (AWS KMS) key without key materia
M. Import the customer key material into the KMS ke
N. Enable automatic key rotation.

Answer: B

Explanation:
SSE-S3 - is free and uses AWS owned CMKs (CMK = Customer Master Key). The encryption key is owned and managed by AWS, and is shared among many accounts. Its rotation is automatic with time that varies as shown in the table here. The time is not explicitly defined.
SSE-KMS - has two flavors:
AWS managed CMK. This is free CMK generated only for your account. You can only view it policies and audit usage, but not manage it. Rotation is automatic - once per 1095 days (3 years),
Customer managed CMK. This uses your own key that you create and can manage. Rotation is not enabled by default. But if you enable it, it will be automatically rotated every 1 year. This variant can also use an imported key material by you. If you create such key with an imported material, there is no automated rotation. Only manual rotation.
SSE-C - customer provided key. The encryption key is fully managed by you outside of AWS. AWS will not rotate it.
This solution meets the requirements of moving data to an Amazon S3 bucket, encrypting the data when it is stored in the S3 bucket, and automatically rotating the encryption key every year with the least operational overhead. AWS Key Management Service (AWS KMS) is a service that enables you to create and manage encryption keys for your data. A customer managed key is a symmetric encryption key that you create and manage in AWS KMS. You can enable automatic key rotation for a customer managed key, which means that AWS KMS generates new cryptographic material for the key every year. You can set the S3 bucket's default encryption behavior to use the customer managed KMS key, which means that any object that is uploaded to the bucket without specifying an encryption method will be encrypted with that key.
Option A is incorrect because using server-side encryption with Amazon S3 managed encryption keys (SSE-S3) does not allow you to control or manage the encryption keys. SSE-S3 uses a unique key for each object, and encrypts that key with a master key that is regularly rotated by S3. However, you cannot enable or disable key rotation for SSE-S3 keys, or specify the rotation interval. Option C is incorrect because manually rotating the KMS key every year can increase the operational overhead and complexity, and it may not meet the requirement of rotating the key every year if you forget or delay the rotation
process. Option D is incorrect because encrypting the data with customer key material before moving the data to the S3 bucket can increase the operational overhead and complexity, and it may not provide consistent encryption for all objects in the bucket. Creating a KMS key without key material and importing the customer key material into the KMS key can enable you to use your own source of random bits to generate your KMS keys, but it does not support automatic key rotation.
References:
? https://docs.aws.amazon.com/kms/latest/developerguide/concepts.html
? https://docs.aws.amazon.com/kms/latest/developerguide/rotate-keys.html
? https://docs.aws.amazon.com/AmazonS3/latest/userguide/bucket-encryption.html

NEW QUESTION 120
- (Topic 2)
A company is planning to build a high performance computing (HPC) workload as a service solution that Is hosted on AWS A group of 16 AmazonEC2Ltnux Instances requires the lowest possible latency for node-to-node communication. The instances also need a shared block device volume for high-performing storage.
Which solution will meet these requirements?

A. Use a duster placement grou
B. Attach a single Provisioned IOPS SSD Amazon Elastic Block Store (Amazon E BS) volume to all the instances by using Amazon EBS Multi-Attach
C. Use a cluster placement grou
D. Create shared 'lie systems across the instances by using Amazon Elastic File System (Amazon EFS)
E. Use a partition placement grou
F. Create shared tile systems across the instances by using Amazon Elastic File System (Amazon EFS).
G. Use a spread placement grou
H. Attach a single Provisioned IOPS SSD Amazon Elastic Block Store (Amazon EBS) volume to all the instances by using Amazon EBS Multi-Attach

Answer: A

Explanation:
1. lowest possible latency + node to node ==> cluster placement(must be within one AZ), so C, D out
* 2. For EBS Multi-Attach, up to 16 instances can be attached to a single volume==>we have 16 linux instance==>more close to A
* 3. "need a shared block device volume"==>EBS Multi-attach is Block Storage whereas EFS is File Storage==> B out
* 4. EFS automatically replicates data within and across 3 AZ==>we use cluster placement
so all EC2 are within one AZ.
* 5. EBS Multi-attach volumes can be used for clients within a single AZ. https://repost.aws/questions/QUK2RANw1QTKCwpDUwCCI72A/efs-vs-ebs-mult-attach

NEW QUESTION 125
- (Topic 2)

A company wants to move its application to a serverless solution. The serverless solution needs to analyze existing and new data by using SL. The company stores the data in an Amazon S3 bucket. The data requires encryption and must be replicated to a different AWS Region.
Which solution will meet these requirements with the LEAST operational overhead?

A. Create a new S3 bucke
B. Load the data into the new S3 bucke
C. Use S3 Cross-Region Replication (CRR) to replicate encrypted objects to an S3 bucket in another Regio
D. Use server-side encryption with AWS KMS multi-Region kays (SSE-KMS). Use Amazon Athena to query the data.
E. Create a new S3 bucke
F. Load the data into the new S3 bucke
G. Use S3 Cross-Region Replication (CRR) to replicate encrypted objects to an S3 bucket in another Regio
H. Use server-side encryption with AWS KMS multi-Region keys (SSE-KMS). Use Amazon RDS to query the data.
I. Load the data into the existing S3 bucke
J. Use S3 Cross-Region Replication (CRR) to replicate encrypted objects to an S3 bucket in another Regio
K. Use server-side encryption with Amazon S3managed encryption keys (SSE-S3). Use Amazon Athena to query the data.
L. Load the data into the existing S3 bucke
M. Use S3 Cross-Region Replication (CRR) to replicate encrypted objects to an S3 bucket in another Regio
N. Use server-side encryption with Amazon S3 managed encryption keys (SSE-S3). Use Amazon RDS to query the data.

Answer: A

Explanation:
This solution meets the requirements of a serverless solution, encryption, replication, and SQL analysis with the least operational overhead. Amazon Athena is a serverless interactive query service that can analyze data in S3 using standard SQL. S3 Cross-Region Replication (CRR) can replicate encrypted objects to an S3 bucket in another Region automatically. Server-side encryption with AWS KMS multi-Region keys (SSE-KMS) can encrypt the data at rest using keys that are replicated across multiple Regions. Creating a new S3 bucket can avoid potential conflicts with existing data or configurations.
Option B is incorrect because Amazon RDS is not a serverless solution and it cannot query data in S3 directly. Option C is incorrect because server-side encryption with Amazon S3 managed encryption keys (SSE-S3) does not use KMS keys and it does not support multi- Region replication. Option D is incorrect because Amazon RDS is not a serverless solution and it cannot query data in S3 directly. It is also incorrect for the same reason as option C. References:
? https://docs.aws.amazon.com/AmazonS3/latest/userguide/replication-walkthrough-4.html
?https://aws.amazon.com/blogs/storage/considering-four-different-replication-options-for-data-in-amazon-s3/
? https://docs.aws.amazon.com/AmazonS3/latest/userguide/UsingEncryption.html
? https://aws.amazon.com/athena/

NEW QUESTION 130
- (Topic 2)
A company has an AWS account used for software engineering. The AWS account has access to the company's on-premises data center through a pair of AWS Direct Connect connections. All non-VPC traffic routes to the virtual private gateway.
A development team recently created an AWS Lambda function through the console. The development team needs to allow the function to access a database that runs in a private subnet in the company's data center.
Which solution will meet these requirements?

A. Configure the Lambda function to run in the VPC with the appropriate security group.
B. Set up a VPN connection from AWS to the data cente
C. Route the traffic from the Lambda function through the VPN.
D. Update the route tables in the VPC to allow the Lambda function to access the on- premises data center through Direct Connect.
E. Create an Elastic IP addres
F. Configure the Lambda function to send traffic through theElastic IP address without an elastic network interface.

Answer: A

Explanation:
https://docs.aws.amazon.com/lambda/latest/dg/configuration-vpc.html#vpc-managing-eni

NEW QUESTION 131
- (Topic 2)
A company runs a production application on a fleet of Amazon EC2 instances. The application reads the data from an Amazon SQS queue and processes the messages in parallel. The message volume is unpredictable and often has intermittent traffic. This application should continually process messages without any downtime.
Which solution meets these requirements MOST cost-effectively?

A. Use Spot Instances exclusively to handle the maximum capacity required.
B. Use Reserved Instances exclusively to handle the maximum capacity required.
C. Use Reserved Instances for the baseline capacity and use Spot Instances to handle additional capacity.
D. Use Reserved Instances for the baseline capacity and use On-Demand Instances to handle additional capacity.

Answer: D

Explanation:
We recommend that you use On-Demand Instances for applications with short-term, irregular workloads that cannot be interrupted.
https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-on-demand-instances.html

NEW QUESTION 134
- (Topic 2)
A corporation has recruited a new cloud engineer who should not have access to the CompanyConfidential Amazon S3 bucket. The cloud engineer must have read and write permissions on an S3 bucket named AdminTools.
Which IAM policy will satisfy these criteria?
A.

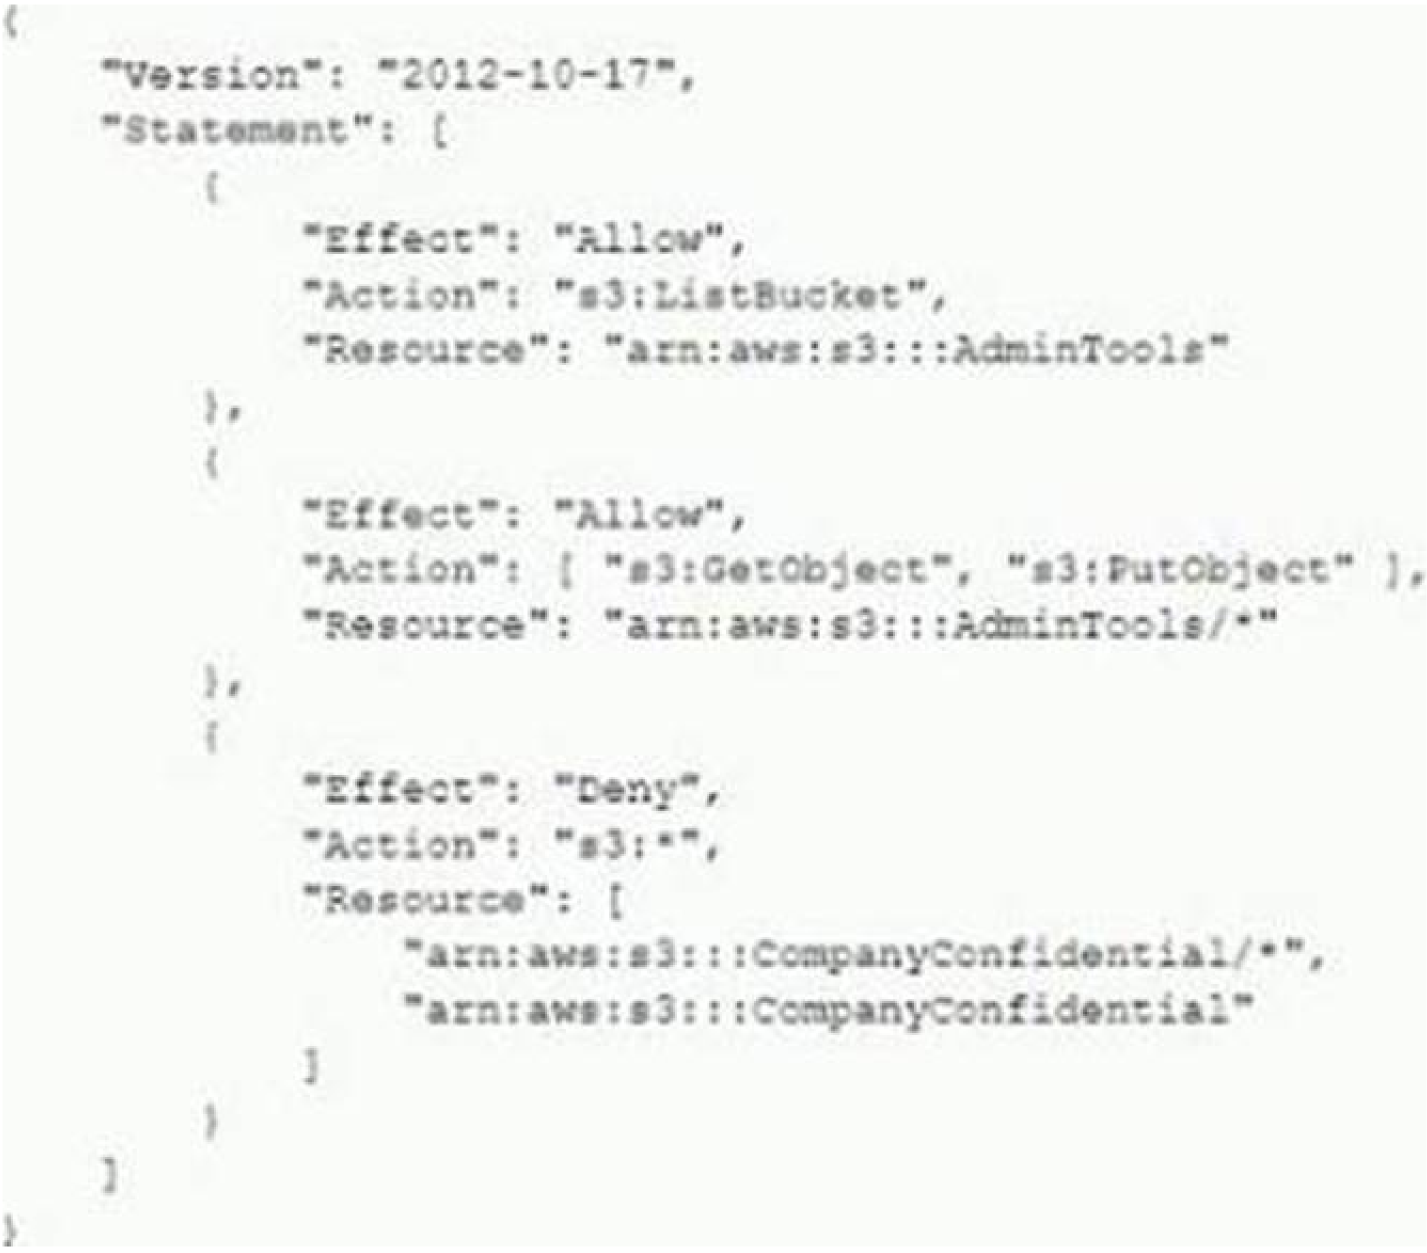

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": "s3:ListBucket",
            "Resource": "arn:aws:s3:::AdminTools"
        },
        {
            "Effect": "Allow",
            "Action": [ "s3:GetObject", "s3:PutObject" ],
            "Resource": "arn:aws:s3:::AdminTools/*"
        },
        {
            "Effect": "Deny",
            "Action": "s3:*",
            "Resource": [
                "arn:aws:s3:::CompanyConfidential/*",
                "arn:aws:s3:::CompanyConfidential"
            ]
        }
    ]
}
```

B.

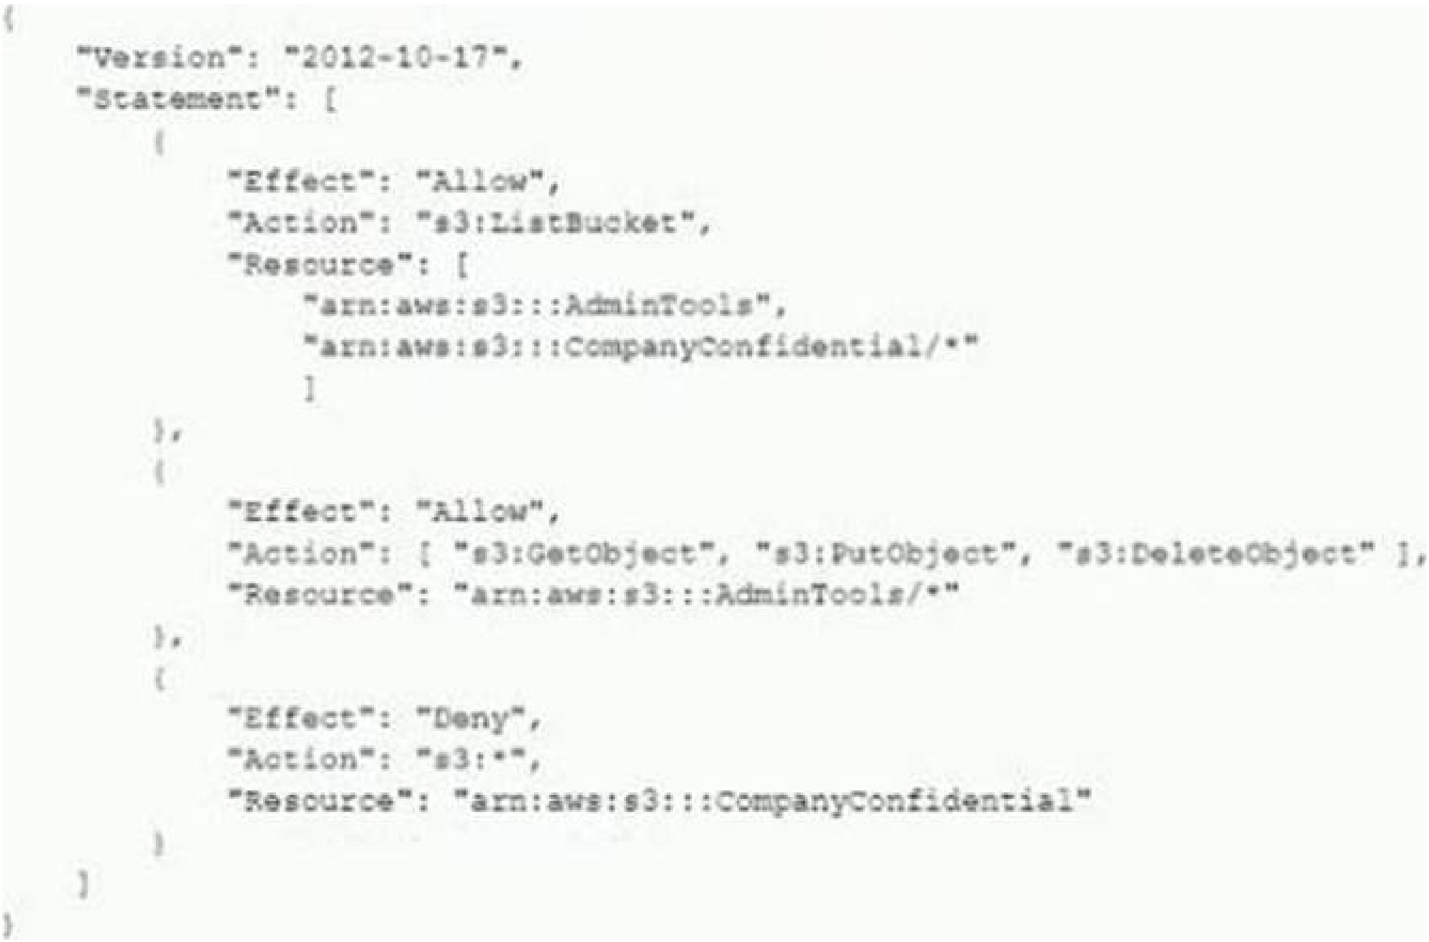

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": "s3:ListBucket",
            "Resource": [
                "arn:aws:s3:::AdminTools",
                "arn:aws:s3:::CompanyConfidential/*"
                ]
        },
        {
            "Effect": "Allow",
            "Action": [ "s3:GetObject", "s3:PutObject", "s3:DeleteObject" ],
            "Resource": "arn:aws:s3:::AdminTools/*"
        },
        {
            "Effect": "Deny",
            "Action": "s3:*",
            "Resource": "arn:aws:s3:::CompanyConfidential"
        }
    ]
}
```

C.

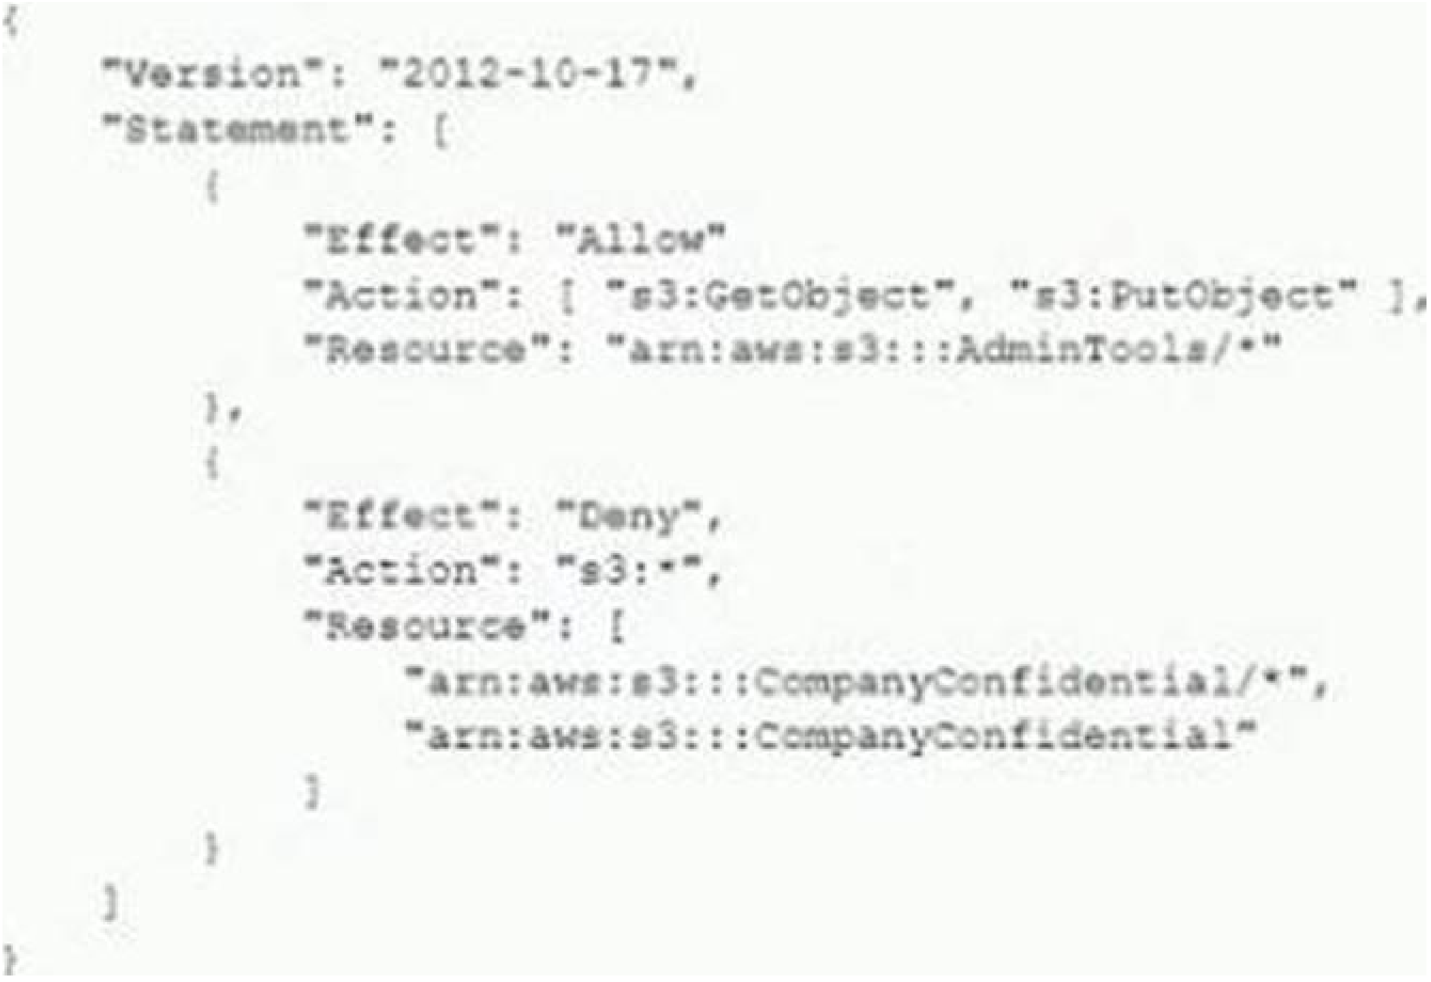

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow"
            "Action": [ "s3:GetObject", "s3:PutObject" ],
            "Resource": "arn:aws:s3:::AdminTools/*"
        },
        {
            "Effect": "Deny",
            "Action": "s3:*",
            "Resource": [
                "arn:aws:s3:::CompanyConfidential/*",
                "arn:aws:s3:::CompanyConfidential"
            ]
        }
    ]
}
```

D.

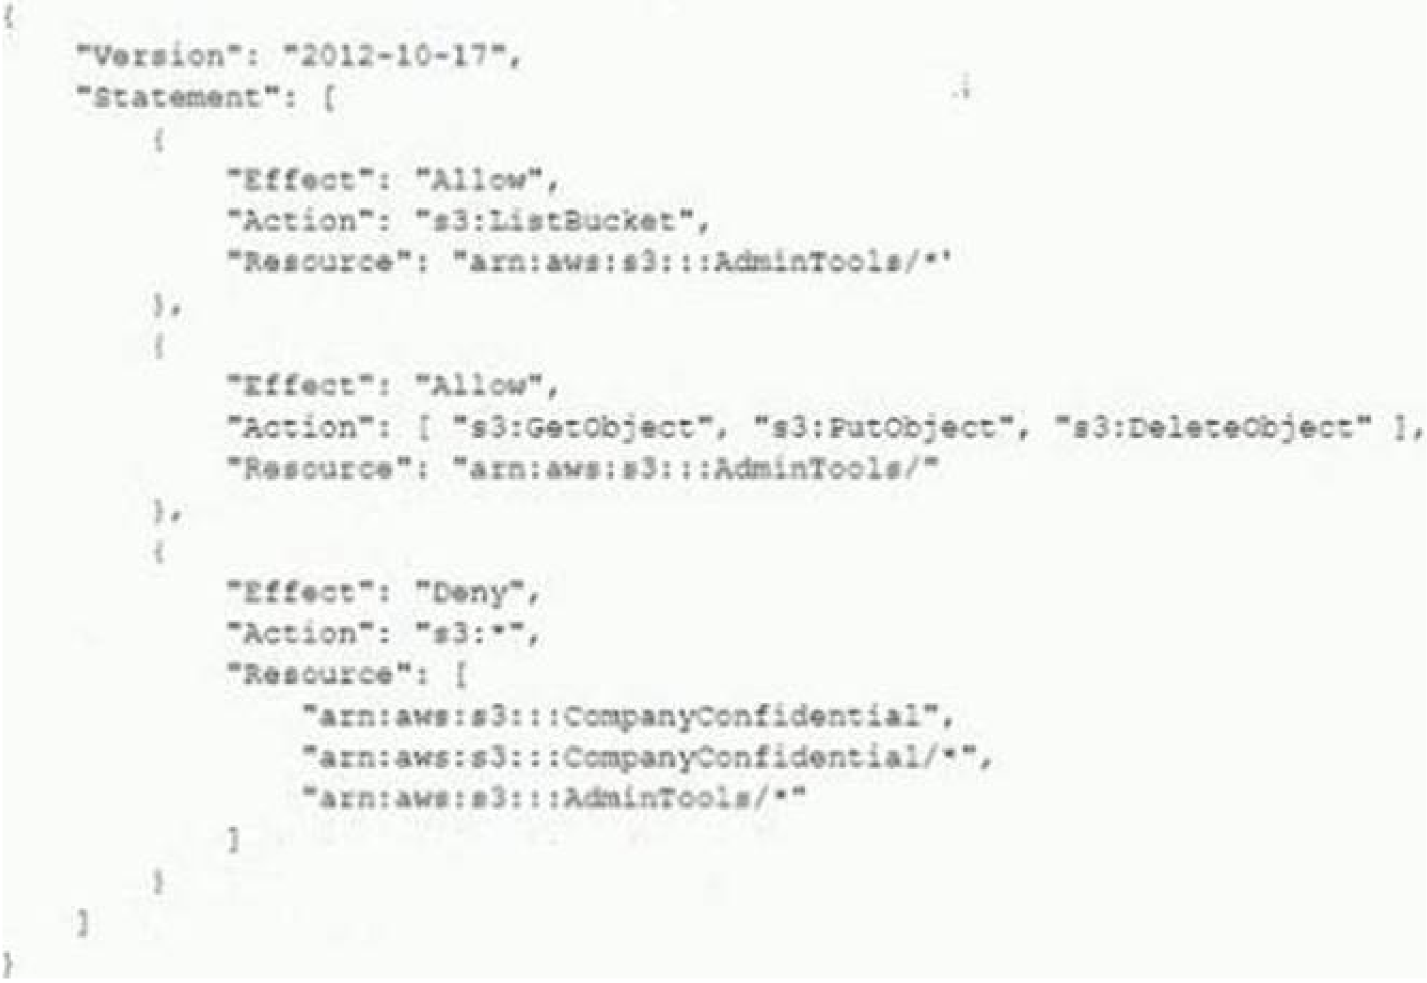

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": "s3:ListBucket",
            "Resource": "arn:aws:s3:::AdminTools/*'
        },
        {
            "Effect": "Allow",
            "Action": [ "s3:GetObject", "s3:PutObject", "s3:DeleteObject" ],
            "Resource": "arn:aws:s3:::AdminTools/"
        },
        {
            "Effect": "Deny",
            "Action": "s3:*",
            "Resource": [
                "arn:aws:s3:::CompanyConfidential",
                "arn:aws:s3:::CompanyConfidential/*",
                "arn:aws:s3:::AdminTools/*"
            ]
        }
    ]
}
```

A.

Answer: A

Explanation:
https://docs.amazonaws.cn/en_us/IAM/latest/UserGuide/reference_policies_examples_s3_ rw-bucket.html
The policy is separated into two parts because the ListBucket action requires permissions on the bucket while the other actions require permissions on the objects in the bucket. You must use two different Amazon Resource Names (ARNs) to specify bucket-level and object-level permissions. The first Resource element specifies arn:aws:s3:::AdminTools for the ListBucket action so that applications can list all objects in the AdminTools bucket.

NEW QUESTION 135
- (Topic 2)
A company has an event-driven application that invokes AWS Lambda functions up to 800 times each minute with varying runtimes. The Lambda functions access data that is stored in an Amazon Aurora MySQL OB cluster. The company is noticing connection timeouts as user activity increases The database shows no signs of being overloaded. CPU. memory, and disk access metrics are all low.
Which solution will resolve this issue with the LEAST operational overhead?

A. Adjust the size of the Aurora MySQL nodes to handle more connection
B. Configure retry logic in the Lambda functions for attempts to connect to the database
C. Set up Amazon ElastiCache tor Redls to cache commonly read items from the databas
D. Configure the Lambda functions to connect to ElastiCache for reads.
E. Add an Aurora Replica as a reader nod
F. Configure the Lambda functions to connect to the reader endpoint of the OB cluster rather than lo the writer endpoint.

G. Use Amazon ROS Proxy to create a prox
H. Set the DB cluster as the target database Configure the Lambda functions lo connect to the proxy rather than to the DB cluster.

Answer: D

Explanation:
1. database shows no signs of being overloaded. CPU, memory, and disk access metrics are all low==>A and C out. We cannot only add nodes instance or add read replica, because database workload is totally fine, very low. 2. "least operational overhead"==>B out, because b need to configure lambda. 3. ROS proxy: Shares infrequently used connections; High availability with failover; Drives increased efficiency==>proxy can leverage failover to redirect traffic from timeout rds instance to
healthy rds instance. So D is right.

NEW QUESTION 140
- (Topic 2)
A company runs a web-based portal that provides users with global breaking news, local alerts, and weather updates. The portal delivers each user a personalized view by using mixture of static and dynamic content. Content is served over HTTPS through an API server running on an Amazon EC2 instance behind an Application Load Balancer (ALB). The company wants the portal to provide this content to its users across the world as quickly as possible.
How should a solutions architect design the application to ensure the LEAST amount of latency for all users?

A. Deploy the application stack in a single AWS Regio
B. Use Amazon CloudFront to serve all static and dynamic content by specifying the ALB as an origin.
C. Deploy the application stack in two AWS Region
D. Use an Amazon Route 53 latency routing policy to serve all content from the ALB in the closest Region.
E. Deploy the application stack in a single AWS Regio
F. Use Amazon CloudFront to serve the static conten
G. Serve the dynamic content directly from the ALB.
H. Deploy the application stack in two AWS Region
I. Use an Amazon Route 53 geolocation routing policy to serve all content from the ALB in the closest Region.

Answer: A

Explanation:
https://aws.amazon.com/blogs/networking-and-content-delivery/deliver-your-apps-dynamic-content-using-amazon-cloudfront-getting-started-template/

NEW QUESTION 142
- (Topic 2)
A company has two applications: a sender application that sends messages with payloads to be processed and a processing application intended to receive the messages with payloads. The company wants to implement an AWS service to handle messages between the two applications. The sender application can send about 1.000 messages each hour. The messages may take up to 2 days to be processed. If the messages fail to process, they must be retained so that they do not impact the processing of any remaining messages.
Which solution meets these requirements and is the MOST operationally efficient?

A. Set up an Amazon EC2 instance running a Redis databas
B. Configure both applications to use the instanc
C. Store, process, and delete the messages, respectively.
D. Use an Amazon Kinesis data stream to receive the messages from the sender applicatio
E. Integrate the processing application with the Kinesis Client Library (KCL).
F. Integrate the sender and processor applications with an Amazon Simple Queue Service(Amazon SQS) queu
G. Configure a dead-letter queue to collect the messages that failed to process.
H. Subscribe the processing application to an Amazon Simple Notification Service (Amazon SNS) topic to receive notifications to proces
I. Integrate the sender application to write to the SNS topic.

Answer: C

Explanation:
https://aws.amazon.com/blogs/compute/building-loosely-coupled-scalable-c-applications-with-amazon-sqs-and-amazon-sns/
https://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/sqs-dead-letter-queues.html

NEW QUESTION 146
- (Topic 2)
A business's backup data totals 700 terabytes (TB) and is kept in network attached storage (NAS) at its data center. This backup data must be available in the event of occasional regulatory inquiries and preserved for a period of seven years. The organization has chosen to relocate its backup data from its on-premises data center to Amazon Web Services (AWS). Within one month, the migration must be completed. The company's public internet connection provides 500 Mbps of dedicated capacity for data transport.
What should a solutions architect do to ensure that data is migrated and stored at the LOWEST possible cost?

A. Order AWS Snowball devices to transfer the dat
B. Use a lifecycle policy to transition the files to Amazon S3 Glacier Deep Archive.
C. Deploy a VPN connection between the data center and Amazon VP
D. Use the AWS CLI to copy the data from on premises to Amazon S3 Glacier.
E. Provision a 500 Mbps AWS Direct Connect connection and transfer the data to Amazon S3. Use a lifecycle policy to transition the files to Amazon S3 Glacier Deep Archive.
F. Use AWS DataSync to transfer the data and deploy a DataSync agent on premise
G. Use the DataSync task to copy files from the on-premises NAS storage to Amazon S3 Glacier.

Answer: A

Explanation:
https://www.omnicalculator.com/other/data-transfer

NEW QUESTION 150
- (Topic 2)
A company runs its ecommerce application on AWS. Every new order is published as a message in a RabbitMQ queue that runs on an Amazon EC2 instance in a single Availability Zone. These messages are processed by a different application that runs on a separate EC2 instance. This application stores the details in a PostgreSQL database on another EC2 instance. All the EC2 instances are in the same Availability Zone.
The company needs to redesign its architecture to provide the highest availability with the least operational overhead.
What should a solutions architect do to meet these requirements?

A. Migrate the queue to a redundant pair (active/standby) of RabbitMQ instances on Amazon M
B. Create a Multi-AZ Auto Scaling group (or EC2 instances that host the applicatio
C. Create another Multi-AZAuto Scaling group for EC2 instances that host the PostgreSQL database.
D. Migrate the queue to a redundant pair (active/standby) of RabbitMQ instances on Amazon M
E. Create a Multi-AZ Auto Scaling group for EC2 instances that host the applicatio
F. Migrate the database to run on a Multi-AZ deployment of Amazon RDS for PostgreSQL.
G. Create a Multi-AZ Auto Scaling group for EC2 instances that host the RabbitMQ queu
H. Create another Multi-AZ Auto Scaling group for EC2 instances that host the application.Migrate the database to run on a Multi-AZ deployment of Amazon RDS fqjPostgreSQL.
I. Create a Multi-AZ Auto Scaling group for EC2 instances that host the RabbitMQ queue.Create another Multi-AZ Auto Scaling group for EC2 instances that host the applicatio
J. Create a third Multi-AZ Auto Scaling group for EC2 instances that host the PostgreSQL database.

Answer: B

Explanation:
Migrating to Amazon MQ reduces the overhead on the queue management. C and D are dismissed. Deciding between A and B means deciding to go for an AutoScaling group for EC2 or an RDS for Postgress (both multi- AZ). The RDS option has less operational impact, as provide as a service the tools and software required. Consider for instance, the effort to add an additional node like a read replica, to the DB. https://docs.aws.amazon.com/amazon-mq/latest/developer-guide/active-standby-broker- deployment.html https://aws.amazon.com/rds/postgresql/

NEW QUESTION 152
- (Topic 2)
A company wants to migrate its existing on-premises monolithic application to AWS.
The company wants to keep as much of the front- end code and the backend code as possible. However, the company wants to break the application into smaller applications. A different team will manage each application. The company needs a highly scalable solution that minimizes operational overhead.
Which solution will meet these requirements?

A. Host the application on AWS Lambda Integrate the application with Amazon API Gateway.
B. Host the application with AWS Amplif
C. Connect the application to an Amazon API Gateway API that is integrated with AWS Lambda.
D. Host the application on Amazon EC2 instance
E. Set up an Application Load Balancer with EC2 instances in an Auto Scaling group as targets.
F. Host the application on Amazon Elastic Container Service (Amazon ECS) Set up an Application Load Balancer with Amazon ECS as the target.

Answer: D

Explanation:
https://aws.amazon.com/blogs/compute/microservice-delivery-with-amazon-ecs-and-application-load-balancers/

NEW QUESTION 155
- (Topic 2)
A solutions architect needs to securely store a database user name and password that an application uses to access an Amazon RDS DB instance. The application that accesses the database runs on an Amazon EC2 instance. The solutions architect wants to create a secure parameter in AWS Systems Manager Parameter Store.
What should the solutions architect do to meet this requirement?

A. Create an IAM role that has read access to the Parameter Store paramete
B. Allow Decrypt access to an AWS Key Management Service (AWS KMS) key that is used to encrypt the paramete
C. Assign this IAM role to the EC2 instance.
D. Create an IAM policy that allows read access to the Parameter Store paramete
E. Allow Decrypt access to an AWS Key Management Service (AWS KMS) key that is used to encrypt the paramete
F. Assign this IAM policy to the EC2 instance.
G. Create an IAM trust relationship between the Parameter Store parameter and the EC2 instanc
H. Specify Amazon RDS as a principal in the trust policy.
I. Create an IAM trust relationship between the DB instance and the EC2 instanc
J. Specify Systems Manager as a principal in the trust policy.

Answer: B

Explanation:
https://docs.aws.amazon.com/IAM/latest/UserGuide/reference_aws-services-that-work- with-iam.html

NEW QUESTION 156
- (Topic 2)
A gaming company is designing a highly available architecture. The application runs on a modified Linux kernel and supports only UDP-based traffic. The company needs the front- end tier to provide the best possible user experience. That tier must have low latency, route traffic to the nearest edge location, and provide static IP addresses for entry into the application endpoints.
What should a solutions architect do to meet these requirements?

A. Configure Amazon Route 53 to forward requests to an Application Load Balance
B. Use AWS Lambda for the application in AWS Application Auto Scaling.

C. Configure Amazon CloudFront to forward requests to a Network Load Balance
D. Use AWS Lambda for the application in an AWS Application Auto Scaling group.
E. Configure AWS Global Accelerator to forward requests to a Network Load Balance
F. Use Amazon EC2 instances for the application in an EC2 Auto Scaling group.
G. Configure Amazon API Gateway to forward requests to an Application Load Balance
H. Use Amazon EC2 instances for the application in an EC2 Auto Scaling group.

Answer: C

Explanation:
AWS Global Accelerator and Amazon CloudFront are separate services that use the AWS global network and its edge locations around the world. CloudFront improves performance for both cacheable content (such as images and videos) and dynamic content (such as API acceleration and dynamic site delivery). Global Accelerator improves performance for a wide range of applications over TCP or UDP by proxying packets at the edge to applications running in one or more AWS Regions. Global Accelerator is a good fit for non- HTTP use cases, such as gaming (UDP), IoT (MQTT), or Voice over IP, as well as for
HTTP use cases that specifically require static IP addresses or deterministic, fast regional failover. Both services integrate with AWS Shield for DDoS protection.

NEW QUESTION 158
- (Topic 2)
A company recently started using Amazon Aurora as the data store for its global ecommerce application When large reports are run developers report that the ecommerce application is performing poorly After reviewing metrics in Amazon CloudWatch, a solutions architect finds that the ReadlOPS and CPUUtilization metrics are spiking when monthly reports run.
What is the MOST cost-effective solution?

A. Migrate the monthly reporting to Amazon Redshift.
B. Migrate the monthly reporting to an Aurora Replica
C. Migrate the Aurora database to a larger instance class
D. Increase the Provisioned IOPS on the Aurora instance

Answer: B

Explanation:
https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Aurora.Replication.htm l
#Aurora.Replication.Replicas Aurora Replicas have two main purposes. You can issue queries to them to scale the read operations for your application. You typically do so by connecting to the reader endpoint of the cluster. That way, Aurora can spread the load for read-only connections across as many Aurora Replicas as you have in the cluster. Aurora Replicas also help to increase availability. If the writer instance in a cluster becomes unavailable, Aurora automatically promotes one of the reader instances to take its place as the new writer. https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Aurora.Overview.html

NEW QUESTION 162
- (Topic 2)
A company uses a three-tier web application to provide training to new employees. The application is accessed for only 12 hours every day. The company is using an Amazon RDS for MySQL DB instance to store information and wants to minimize costs.
What should a solutions architect do to meet these requirements?

A. Configure an IAM policy for AWS Systems Manager Session Manage
B. Create an IAM role for the polic
C. Update the trust relationship of the rol
D. Set up automatic start and stop for the DB instance.
E. Create an Amazon ElastiCache for Redis cache cluster that gives users the ability to access the data from the cache when the DB instance is stoppe
F. Invalidate the cache after the DB instance is started.
G. Launch an Amazon EC2 instanc
H. Create an IAM role that grants access to Amazon RD
I. Attach the role to the EC2 instanc
J. Configure a cron job to start and stop the EC2 instance on the desired schedule.
K. Create AWS Lambda functions to start and stop the DB instanc
L. Create Amazon EventBridge (Amazon CloudWatch Events) scheduled rules to invoke the Lambda function
M. Configure the Lambda functions as event targets for the rules

Answer: D

Explanation:
In a typical development environment, dev and test databases are mostly utilized for 8 hours a day and sit idle when not in use. However, the databases are billed for the compute and storage costs during this idle time. To reduce the overall cost, Amazon RDS allows instances to be stopped temporarily. While the instance is stopped, you’re charged for storage and backups, but not for the DB instance hours. Please note that a stopped instance will automatically be started after 7 days. This post presents a solution using AWS Lambda and Amazon EventBridge that allows you to schedule a Lambda function to stop and start the idle databases with specific tags to save on compute costs. The second post presents a solution that accomplishes stop and start of the idle Amazon RDS databases using AWS Systems Manager.

NEW QUESTION 165
- (Topic 2)
A company wants to migrate its MySQL database from on premises to AWS. The company recently experienced a database outage that significantly impacted the business. To ensure this does not happen again, the company wants a reliable database solution on AWS that minimizes data loss and stores every transaction on at least two nodes.
Which solution meets these requirements?

A. Create an Amazon RDS DB instance with synchronous replication to three nodes in three Availability Zones.
B. Create an Amazon RDS MySQL DB instance with Multi-AZ functionality enabled to synchronously replicate the data.
C. Create an Amazon RDS MySQL DB instance and then create a read replica in a separate AWS Region that synchronously replicates the data.
D. Create an Amazon EC2 instance with a MySQL engine installed that triggers an AWS Lambda function to synchronously replicate the data to an Amazon RDS MySQL DB instance.

Answer: B

Explanation:
Q: What does Amazon RDS manage on my behalf?
Amazon RDS manages the work involved in setting up a relational database: from provisioning the infrastructure capacity you request to installing the database software. Once your database is up and running, Amazon RDS automates common administrative tasks such as performing backups and patching the software that powers your database. With optional Multi-AZ deployments, Amazon RDS also manages synchronous data replication across Availability Zones with automatic failover. https://aws.amazon.com/rds/faqs/

NEW QUESTION 168
- (Topic 2)
A company needs to move data from an Amazon EC2 instance to an Amazon S3 bucket. The company must ensure that no API calls and no data are routed through public internet routes. Only the EC2 instance can have access to upload data to the S3 bucket.
Which solution will meet these requirements?

A. Create an interface VPC endpoint for Amazon S3 in the subnet where the EC2 instance is locate
B. Attach a resource policy to the S3 bucket to only allow the EC2 instance's IAM role for access.
C. Create a gateway VPC endpoint for Amazon S3 in the Availability Zone where the EC2 instance is locate
D. Attach appropriate security groups to the endpoin
E. Attach a resource policy lo the S3 bucket to only allow the EC2 instance's IAM role for access.
F. Run the nslookup tool from inside the EC2 instance to obtain the private IP address of the S3 bucket's service API endpoin
G. Create a route in the VPC route table to provide theEC2 instance with access to the S3 bucke
H. Attach a resource policy to the S3 bucket to only allow the EC2 instance's IAM role for access.
I. Use the AWS provided, publicly available ip-ranges.json tile to obtain the private IP address of the S3 bucket's service API endpoin
J. Create a route in the VPC route table to provide the EC2 instance with access to the S3 bucke
K. Attach a resource policy to the S3 bucket to only allow the EC2 instance's IAM role for access.

Answer: A

Explanation:
(https://aws.amazon.com/blogs/security/how-to-restrict-amazon-s3-bucket-access-to-a-specific-iam-role/)

NEW QUESTION 172
- (Topic 2)
A company's website provides users with downloadable historical performance reports. The website needs a solution that will scale to meet the company's website demands globally. The solution should be cost-effective, limit the provisioning of infrastructure resources, and provide the fastest possible response time.
Which combination should a solutions architect recommend to meet these requirements?

A. Amazon CloudFront and Amazon S3
B. AWS Lambda and Amazon DynamoDB
C. Application Load Balancer with Amazon EC2 Auto Scaling
D. Amazon Route 53 with internal Application Load Balancers

Answer: A

Explanation:
Cloudfront for rapid response and s3 to minimize infrastructure.

NEW QUESTION 174
- (Topic 2)
A company is running a multi-tier web application on premises. The web application is containerized and runs on a number of Linux hosts connected to a PostgreSQL database that contains user records. The operational overhead of maintaining the infrastructure and capacity planning is limiting the company's growth. A solutions architect must improve the application's infrastructure.
Which combination of actions should the solutions architect take to accomplish this? (Choose two.)

A. Migrate the PostgreSQL database to Amazon Aurora
B. Migrate the web application to be hosted on Amazon EC2 instances.
C. Set up an Amazon CloudFront distribution for the web application content.
D. Set up Amazon ElastiCache between the web application and the PostgreSQL database.
E. Migrate the web application to be hosted on AWS Fargate with Amazon Elastic Container Service (Amazon ECS).

Answer: AE

Explanation:
Amazon Aurora is a fully managed, scalable, and highly available relational database service that is compatible with PostgreSQL. Migrating the database to Amazon Aurora would reduce the operational overhead of maintaining the database infrastructure and allow the company to focus on building and scaling the application. AWS Fargate is a fully managed container orchestration service that enables users to run containers without the need to manage the underlying EC2 instances. By using AWS Fargate with Amazon Elastic Container Service (Amazon ECS), the solutions architect can improve the scalability and efficiency of the web application and reduce the operational overhead of maintaining the underlying infrastructure.

NEW QUESTION 178
- (Topic 2)
A large media company hosts a web application on AWS. The company wants to start caching confidential media files so that users around the world will have reliable access to the files. The content is stored in Amazon S3 buckets. The company must deliver the content quickly, regardless of where the requests originate geographically.
Which solution will meet these requirements?

A. Use AWS DataSync to connect the S3 buckets to the web application.
B. Deploy AWS Global Accelerator to connect the S3 buckets to the web application.

C. Deploy Amazon CloudFront to connect the S3 buckets to CloudFront edge servers.
D. Use Amazon Simple Queue Service (Amazon SQS) to connect the S3 buckets to the web application.

Answer: C

Explanation:
CloudFront uses a local cache to provide the response, AWS Global accelerator proxies requests and connects to the application all the time for the response.
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/private-content-restricting-access-to-s3.html#private-content-granting-permissions-to-oai

NEW QUESTION 180
- (Topic 2)
A company is running several business applications in three separate VPCs within me us- east-1 Region. The applications must be able to communicate between VPCs. The applications also must be able to consistently send hundreds to gigabytes of data each day to a latency-sensitive application that runs in a single on-premises data center.
A solutions architect needs to design a network connectivity solution that maximizes cost- effectiveness
Which solution moots those requirements?

A. Configure three AWS Site-to-Site VPN connections from the data center to AWS Establish connectivity by configuring one VPN connection for each VPC
B. Launch a third-party virtual network appliance in each VPC Establish an iPsec VPN tunnel between the Data center and each virtual appliance
C. Set up three AWS Direct Connect connections from the data center to a Direct Connect gateway in us-east-1 Establish connectivity by configuring each VPC to use one of the Direct Connect connections
D. Set up one AWS Direct Connect connection from the data center to AW
E. Create a transit gateway, and attach each VPC to the transit gatewa
F. Establish connectivity between the Direct Connect connection and the transit gateway.

Answer: D

Explanation:
https://docs.aws.amazon.com/whitepapers/latest/aws-vpc-connectivity-options/aws-direct-connect-aws-transit-gateway.html

NEW QUESTION 185
- (Topic 2)
A company is migrating an application from on-premises servers to Amazon EC2 instances. As part of the migration design requirements, a solutions architect must implement infrastructure metric alarms. The company does not need to take action if CPU utilization increases to more than 50% for a short burst of time. However, if the CPU
utilization increases to more than 50% and read IOPS on the disk are high at the same time, the company needs to act as soon as possible. The solutions architect also must reduce false alarms.
What should the solutions architect do to meet these requirements?

A. Create Amazon CloudWatch composite alarms where possible.
B. Create Amazon CloudWatch dashboards to visualize the metrics and react to issues quickly.
C. Create Amazon CloudWatch Synthetics canaries to monitor the application and raise an alarm.
D. Create single Amazon CloudWatch metric alarms with multiple metric thresholds where possible.

Answer: A

Explanation:
Composite alarms determine their states by monitoring the states of other alarms. You can **use composite alarms to reduce alarm noise**. For example, you can create a composite alarm where the underlying metric alarms go into ALARM when they meet specific conditions. You then can set up your composite alarm to go into ALARM and send you notifications when the underlying metric alarms go into ALARM by configuring the underlying metric alarms never to take actions. Currently, composite alarms can take the following actions: https://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/Create_Composite_Al arm.html

NEW QUESTION 189
- (Topic 3)
A company is migrating an old application to AWS The application runs a batch job every hour and is CPU intensive The batch job takes 15 minutes on average with an on-premises server The server has 64 virtual CPU (vCPU) and 512 GiB of memory
Which solution will run the batch job within 15 minutes with the LEAST operational overhead?

A. Use AWS Lambda with functional scaling
B. Use Amazon Elastic Container Service (Amazon ECS) with AWS Fargate
C. Use Amazon Lightsail with AWS Auto Scaling
D. Use AWS Batch on Amazon EC2

Answer: D

Explanation:
Use AWS Batch on Amazon EC2. AWS Batch is a fully managed batch processing service that can be used to easily run batch jobs on Amazon EC2 instances. It can scale the number of instances to match the workload, allowing the batch job to be completed in the desired time frame with minimal operational overhead.
Using AWS Lambda with Amazon API Gateway - AWS Lambda https://docs.aws.amazon.com/lambda/latest/dg/services-apigateway.html
AWS Lambda FAQs https://aws.amazon.com/lambda/faqs/

NEW QUESTION 191
- (Topic 3)
A company is running a multi-tier recommence web application in the AWS Cloud. The application runs on Amazon EC2 instances with an Amazon RDS for MySQL Multi-AZ OB instance. Amazon ROS is configured with the latest generation DB instance with 2.000 GB of storage In a General Purpose SSD (gp3) Amazon Elastic Block Store (Amazon EBSl volume. The database performance affects the application during periods high demand.
A database administrator analyzes the logs in Amazon CloudWatch Logs and discovers that the application performance always degrades when the number of read and write IOPS is higher than 20.000.
What should a solutions architect do to improve the application performance?

A. Replace the volume with a magnetic volume.
B. Increase the number of IOPS on the gp3 volume.
C. Replace the volume with a Provisioned IOPS SSD (Io2) volume.
D. Replace the 2.000 GB gp3 volume with two 1.000 GB gp3 volumes

Answer: C

Explanation:
https://aws.amazon.com/ebs/features/ Amazon EBS provides a range of options that allow you to optimize storage performance and cost for your workload. These options are divided into two major categories: SSD-backed storage for transactional workloads, such as databases and boot volumes (performance depends primarily on IOPS), and HDD-backed storage for throughput intensive workloads, such as MapReduce and log processing (performance depends primarily on MB/s).

NEW QUESTION 192
- (Topic 3)
A gaming company is moving its public scoreboard from a data center to the AWS Cloud. The company uses Amazon EC2 Windows Server instances behind an Application Load Balancer to host its dynamic application. The company needs a highly available storage solution for the application. The application consists of static files and dynamic server-side code.
Which combination of steps should a solutions architect take to meet these requirements? (Select TWO.)

A. Store the static files on Amazon S3. Use Amazon CloudFront to cache objects at the edge.
B. Store the static files on Amazon S3. Use Amazon ElastiCache to cache objects at the edge.
C. Store the server-side code on Amazon Elastic File System (Amazon EFS). Mount the EFS volume on each EC2 instance to share the files.
D. Store the server-side code on Amazon FSx for Windows File Serve
E. Mount the FSx for Windows File Server volume on each EC2 instance to share the files.
F. Store the server-side code on a General Purpose SSD (gp2) Amazon Elastic Block Store (Amazon EBS) volum
G. Mount the EBS volume on each EC2 instance to share the files.

Answer: AD

Explanation:
A because Elasticache, despite being ideal for leaderboards per Amazon, doesn't cache at edge locations. D because FSx has higher performance for low latency needs. https://www.techtarget.com/searchaws/tip/Amazon-FSx-vs-EFS-Compare-the-AWS-file- services "FSx is built for high performance and submillisecond latency using solid-state drive storage volumes. This design enables users to select storage capacity and latency independently. Thus, even a subterabyte file system can have 256 Mbps or higher throughput and support volumes up to 64 TB."
Amazon S3 is an object storage service that can store static files such as images, videos, documents, etc. Amazon EFS is a file storage service that can store files in a hierarchical structure and supports NFS protocol. Amazon FSx for Windows File Server is a file storage service that can store files in a hierarchical structure and supports SMB protocol. Amazon EBS is a block storage service that can store data in fixed-size blocks and attach to EC2 instances.
Based on these definitions, the combination of steps that should be taken to meet the requirements are:
* A. Store the static files on Amazon S3. Use Amazon CloudFront to cache objects at the edge. D. Store the server-side code on Amazon FSx for Windows File Server. Mount the FSx for Windows File Server volume on each EC2 instance to share the files.

NEW QUESTION 196
......