Amazon-Web-Services

Exam Questions SAA-C03

AWS Certified Solutions Architect - Associate (SAA-C03)

NEW QUESTION 1
- (Topic 1)
A company is running a popular social media website. The website gives users the ability to upload images to share with other users. The company wants to make sure that the images do not contain inappropriate content. The company needs a solution that minimizes development effort.
What should a solutions architect do to meet these requirements?

A. Use Amazon Comprehend to detect inappropriate conten
B. Use human review for low- confidence predictions.
C. Use Amazon Rekognition to detect inappropriate conten
D. Use human review for low- confidence predictions.
E. Use Amazon SageMaker to detect inappropriate conten
F. Use ground truth to label low- confidence predictions.
G. Use AWS Fargate to deploy a custom machine learning model to detect inappropriate conten
H. Use ground truth to label low-confidence predictions.

Answer: B

Explanation:
https://docs.aws.amazon.com/rekognition/latest/dg/moderation.html?pg=ln&sec=ft https://docs.aws.amazon.com/rekognition/latest/dg/a2i-rekognition.html

NEW QUESTION 2
- (Topic 1)
A company runs an online marketplace web application on AWS. The application serves hundreds of thousands of users during peak hours. The company needs a scalable, near- real-time solution to share the details of millions of financial transactions with several other internal applications Transactions also need to be processed to remove sensitive data before being stored in a document database for low-latency retrieval.
What should a solutions architect recommend to meet these requirements?

A. Store the transactions data into Amazon DynamoDB Set up a rule in DynamoDB to remove sensitive data from every transaction upon write Use DynamoDB Streams to share the transactions data with other applications
B. Stream the transactions data into Amazon Kinesis Data Firehose to store data in Amazon DynamoDB and Amazon S3 Use AWS Lambda integration with Kinesis Data Firehose to remove sensitive dat
C. Other applications can consume the data stored in Amazon S3
D. Stream the transactions data into Amazon Kinesis Data Streams Use AWS Lambda integration to remove sensitive data from every transaction and then store the transactionsdata in Amazon DynamoDB Other applications can consume the transactions data off the Kinesis data stream.
E. Store the batched transactions data in Amazon S3 as file
F. Use AWS Lambda to process every file and remove sensitive data before updating the files in Amazon S3 The Lambda function then stores the data in Amazon DynamoDB Other applications can consume transaction files stored in Amazon S3.

Answer: C

Explanation:
The destination of your Kinesis Data Firehose delivery stream. Kinesis Data Firehose can send data records to various destinations, including Amazon Simple Storage Service (Amazon S3), Amazon Redshift, Amazon OpenSearch Service, and any HTTP endpoint that is owned by you or any of your third-party service providers. The following are the supported destinations:
* Amazon OpenSearch Service
* Amazon S3
* Datadog
* Dynatrace
* Honeycomb
* HTTP Endpoint
* Logic Monitor
* MongoDB Cloud
* New Relic
* Splunk
* Sumo Logic https://docs.aws.amazon.com/firehose/latest/dev/create-name.html
https://aws.amazon.com/kinesis/data-streams/
Amazon Kinesis Data Streams (KDS) is a massively scalable and durable real-time data streaming service. KDS can continuously capture gigabytes of data per second from hundreds of thousands of sources such as website clickstreams, database event streams, financial transactions, social media feeds, IT logs, and location-tracking events.

NEW QUESTION 3
- (Topic 1)
A company observes an increase in Amazon EC2 costs in its most recent bill The billing team notices unwanted vertical scaling of instance types for a couple of EC2 instances A solutions architect needs to create a graph comparing the last 2 months of EC2 costs and perform an in-depth analysis to identify the root cause of the vertical scaling
How should the solutions architect generate the information with the LEAST operational overhead?

A. Use AWS Budgets to create a budget report and compare EC2 costs based on instance types
B. Use Cost Explorer's granular filtering feature to perform an in-depth analysis of EC2 costs based on instance types
C. Use graphs from the AWS Billing and Cost Management dashboard to compare EC2 costs based on instance types for the last 2 months
D. Use AWS Cost and Usage Reports to create a report and send it to an Amazon S3 bucket Use Amazon QuickSight with Amazon S3 as a source to generate an interactive graph based on instance types.

Answer: B

Explanation:
AWS Cost Explorer is a tool that enables you to view and analyze your costs and usage. You can explore your usage and costs using the main graph, the Cost Explorer cost and usage reports, or the Cost Explorer RI reports. You can view data for up to the last 12 months, forecast how much you're likely to spend for the next 12 months, and get recommendations for what Reserved Instances to purchase. You can use Cost Explorer to identify areas that need further inquiry and see

trends that you can use to understand your
costs. https://docs.aws.amazon.com/cost-management/latest/userguide/ce-what-is.html

NEW QUESTION 4
- (Topic 1)
A global company hosts its web application on Amazon EC2 instances behind an Application Load Balancer (ALB). The web application has static data and dynamic data. The company stores its static data in an Amazon S3 bucket. The company wants to improve performance and reduce latency for the static data and dynamic data. The company is using its own domain name registered with Amazon Route 53.
What should a solutions architect do to meet these requirements?

A. Create an Amazon CloudFront distribution that has the S3 bucket and the ALB as origins Configure Route 53 to route traffic to the CloudFront distribution.
B. Create an Amazon CloudFront distribution that has the ALB as an origin Create an AWS Global Accelerator standard accelerator that has the S3 bucket as an endpoin
C. Configure Route 53 to route traffic to the CloudFront distribution.
D. Create an Amazon CloudFront distribution that has the S3 bucket as an origin Create an AWS Global Accelerator standard accelerator that has the ALB and the CloudFront distribution as endpoints Create a custom domain name that points to the accelerator DNS name Use the custom domain name as an endpoint for the web application.
E. Create an Amazon CloudFront distribution that has the ALB as an origin
F. Create an AWS Global Accelerator standard accelerator that has the S3 bucket as an endpoint Create two domain name
G. Point one domain name to the CloudFront DNS name for dynamic content, Point the other domain name to the accelerator DNS name for static content Use the domain names as endpoints for the web application.

Answer: C

Explanation:
Static content can be cached at Cloud front Edge locations from S3 and dynamic content EC2 behind the ALB whose performance can be improved by Global Accelerator whose one endpoint is ALB and other Cloud front. So with regards to custom domain name endpoint is web application is R53 alias records for the custom domain point to web application https://aws.amazon.com/blogs/networking-and-content-delivery/improving-availability-and-performance-for-application-load-balancers-using-one-click-integration- with-aws-global-accelerator/

NEW QUESTION 5
- (Topic 1)
A company is migrating a distributed application to AWS The application serves variable workloads The legacy platform consists of a primary server trial coordinates jobs across multiple compute nodes The company wants to modernize the application with a solution that maximizes resiliency and scalability.
How should a solutions architect design the architecture to meet these requirements?

A. Configure an Amazon Simple Queue Service (Amazon SQS) queue as a destination for the jobs Implement the compute nodes with Amazon EC2 instances that are managed in an Auto Scaling grou
B. Configure EC2 Auto Scaling to use scheduled scaling
C. Configure an Amazon Simple Queue Service (Amazon SQS) queue as a destination for the jobs Implement the compute nodes with Amazon EC2 Instances that are managed in an Auto Scaling group Configure EC2 Auto Scaling based on the size of the queue
D. Implement the primary server and the compute nodes with Amazon EC2 instances that are managed In an Auto Scaling grou
E. Configure AWS CloudTrail as a destination for the fobs Configure EC2 Auto Scaling based on the load on the primary server
F. implement the primary server and the compute nodes with Amazon EC2 instances that are managed in an Auto Scaling group Configure Amazon EventBridge (Amazon CloudWatch Events) as a destination for the jobs Configure EC2 Auto Scaling based on the load on the compute nodes

Answer: B

Explanation:
To maximize resiliency and scalability, the best solution is to use an Amazon SQS queue as a destination for the jobs. This decouples the primary server from the compute nodes, allowing them to scale independently. This also helps to prevent job loss in the event of a failure. Using an Auto Scaling group of Amazon EC2 instances for the compute nodes allows for automatic scaling based on the workload. In this case, it's recommended to configure the Auto Scaling group based on the size of the Amazon SQS queue, which is a better indicator of the actual workload than the load on the primary server or compute nodes. This approach ensures that the application can handle variable workloads, while also minimizing costs by automatically scaling up or down the compute nodes as needed.

NEW QUESTION 6
- (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 7
- (Topic 1)
A company is implementing a new business application. The application runs on two Amazon EC2 instances and uses an Amazon S3 bucket for document storage. A solutions architect needs to ensure that the EC2 instances can access the S3 bucket.
What should the solutions architect do to meet this requirement?

A. Create an IAM role that grants access to the S3 bucke
B. Attach the role to the EC2 instances.
C. Create an IAM policy that grants access to the S3 bucke
D. Attach the policy to the EC2 instances.
E. Create an IAM group that grants access to the S3 bucke
F. Attach the group to the EC2 instances.
G. Create an IAM user that grants access to the S3 bucke
H. Attach the user account to the EC2 instances.

Answer: A

Explanation:
https://aws.amazon.com/premiumsupport/knowledge-center/ec2-instance-access-s3- bucket/

NEW QUESTION 8
- (Topic 1)
A company uses NFS to store large video files in on-premises network attached storage. Each video file ranges in size from 1MB to 500 GB. The total storage is 70 TB and is no longer growing. The company decides to migrate the video files to Amazon S3. The company must migrate the video files as soon as possible while using the least possible network bandwidth.
Which solution will meet these requirements?

A. Create an S3 bucket Create an IAM role that has permissions to write to the S3 bucke
B. Use the AWS CLI to copy all files locally to the S3 bucket.
C. Create an AWS Snowball Edge jo
D. Receive a Snowball Edge device on premise
E. Use the Snowball Edge client to transfer data to the devic
F. Return the device so that AWS can import the data into Amazon S3.
G. Deploy an S3 File Gateway on premise
H. Create a public service endpoint to connect to the S3 File Gateway Create an S3 bucket Create a new NFS file share on the S3 File Gateway Point the new file share to the S3 bucke
I. Transfer the data from the existing NFS file share to the S3 File Gateway.
J. Set up an AWS Direct Connect connection between the on-premises network and AW
K. Deploy an S3 File Gateway on premise
L. Create a public virtual interlace (VIF) to connect to the S3 File Gatewa
M. Create an S3 bucke
N. Create a new NFS file share on the S3 File Gatewa
O. Point the new file share to the S3 bucke
P. Transfer the data from the existing NFS file share to the S3 File Gateway.

Answer: B

Explanation:
The basic difference between Snowball and Snowball Edge is the capacity they provide. Snowball provides a total of 50 TB or 80 TB, out of which 42 TB or 72 TB is available, while Amazon Snowball Edge provides 100 TB, out of which 83 TB is available.

NEW QUESTION 9
- (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 10
- (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 10
- (Topic 1)
A company wants to run its critical applications in containers to meet requirements tor scalability and availability The company prefers to focus on maintenance of the critical applications The company does not want to be responsible for provisioning and managing the underlying infrastructure that runs the containerized workload
What should a solutions architect do to meet those requirements?

A. Use Amazon EC2 Instances, and Install Docker on the Instances
B. Use Amazon Elastic Container Service (Amazon ECS) on Amazon EC2 worker nodes
C. Use Amazon Elastic Container Service (Amazon ECS) on AWS Fargate
D. Use Amazon EC2 instances from an Amazon Elastic Container Service (Amazon ECS)- op6mized Amazon Machine Image (AMI).

Answer: C

Explanation:
using AWS ECS on AWS Fargate since they requirements are for scalability and availability without having to provision and manage the underlying infrastructure to run the containerized workload. https://docs.aws.amazon.com/AmazonECS/latest/userguide/what-is-fargate.html

NEW QUESTION 14
- (Topic 1)
A company is building an ecommerce web application on AWS. The application sends information about new orders to an Amazon API Gateway REST API to process. The company wants to ensure that orders are processed in the order that they are received.
Which solution will meet these requirements?

A. Use an API Gateway integration to publish a message to an Amazon Simple Notification Service (Amazon SNS) topic when the application receives an orde
B. Subscribe an AWS Lambda function to the topic to perform processing.
C. Use an API Gateway integration to send a message to an Amazon Simple Queue Service (Amazon SQS) FIFO queue when the application receives an orde
D. Configure the SQS FIFO queue to invoke an AWS Lambda function for processing.
E. Use an API Gateway authorizer to block any requests while the application processes an order.
F. Use an API Gateway integration to send a message to an Amazon Simple Queue Service (Amazon SQS) standard queue when the application receives an orde
G. Configure the SQS standard queue to invoke an AWS Lambda function for processing.

Answer: B

Explanation:
To ensure that orders are processed in the order that they are received, the best solution is to use an Amazon SQS FIFO (First-In-First-Out) queue. This type of queue maintains the exact order in which messages are sent and received. In this case, the application can send information about new orders to an Amazon API Gateway REST API, which can then use an API Gateway integration to send a message to an Amazon SQS FIFO queue for processing. The queue can then be configured to invoke an AWS Lambda function to perform the necessary processing on each order. This ensures that orders are processed in the exact order in which they are received.

NEW QUESTION 18
- (Topic 1)
A company has a website hosted on AWS. The website is behind an Application Load Balancer (ALB) that is configured to handle HTTP and HTTPS separately. The company wants to forward all requests to the website so that the requests will use HTTPS.
What should a solutions architect do to meet this requirement?

A. Update the ALB's network ACL to accept only HTTPS traffic
B. Create a rule that replaces the HTTP in the URL with HTTPS.
C. Create a listener rule on the ALB to redirect HTTP traffic to HTTPS.
D. Replace the ALB with a Network Load Balancer configured to use Server Name Indication (SNI).

Answer: C

Explanation:
https://aws.amazon.com/premiumsupport/knowledge-center/elb-redirect-http-to-https- using-alb/
How can I redirect HTTP requests to HTTPS using an Application Load Balancer? Last updated: 2020-10-30 I want to redirect HTTP requests to HTTPS using Application Load Balancer listener rules. How can I do this? Resolution Reference: https://aws.amazon.com/premiumsupport/knowledge-center/elb-redirect-http-to-https- using-alb/

NEW QUESTION 22
- (Topic 1)
A company is launching a new application and will display application metrics on an Amazon CloudWatch dashboard. The company's product manager needs to access this dashboard periodically. The product manager does not have an AWS account. A solution architect must provide access to the product manager by following the principle of least privilege.
Which solution will meet these requirements?

A. Share the dashboard from the CloudWatch consol
B. Enter the product manager's email address, and complete the sharing step
C. Provide a shareable link for the dashboard to the product manager.
D. Create an IAM user specifically for the product manage
E. Attach the CloudWatch Read Only Access managed policy to the use
F. Share the new login credential with the product manage
G. Share the browser URL of the correct dashboard with the product manager.
H. Create an IAM user for the company's employees, Attach the View Only Access AWS managed policy to the IAM use
I. Share the new login credentials with the product manage
J. Ask the product manager to navigate to the CloudWatch console and locate the dashboard by name in the Dashboards section.
K. Deploy a bastion server in a public subne
L. When the product manager requires access to the dashboard, start the server and share the RDP credential
M. On the bastion server, ensure that the browser is configured to open the dashboard URL with cached AWS credentials that have appropriate permissions to view the dashboard.

Answer: B

Explanation:
To provide the product manager access to the Amazon CloudWatch dashboard while following the principle of least privilege, a solution architect should create an IAM user specifically for the product manager and attach the CloudWatch Read Only Access managed policy to the user. This policy allows the user to view the dashboard without being able to make any changes to it. The solution architect should then share the new login credential with the product manager and provide them with the browser URL of the correct dashboard.

NEW QUESTION 23
- (Topic 1)
A company is hosting a web application on AWS using a single Amazon EC2 instance that stores user-uploaded documents in an Amazon EBS volume. For better scalability and availability, the company duplicated the architecture and created a second EC2 instance and EBS volume in another Availability Zone placing both behind an Application Load Balancer After completing this change, users reported that, each time they refreshed the website, they could see one subset of their documents or the other, but never all of the documents at the same time.
What should a solutions architect propose to ensure users see all of their documents at once?

A. Copy the data so both EBS volumes contain all the documents.
B. Configure the Application Load Balancer to direct a user to the server with the documents
C. Copy the data from both EBS volumes to Amazon EFS Modify the application to save new documents to Amazon EFS
D. Configure the Application Load Balancer to send the request to both servers Return each document from the correct server.

Answer: C

Explanation:
https://docs.aws.amazon.com/efs/latest/ug/how-it-works.html#how-it-works-ec2

NEW QUESTION 25
- (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 28
- (Topic 1)
A company wants to migrate an on-premises data center to AWS. The data canter hosts an SFTP server that stores its data on an NFS-based file system. The server holds 200 GB of data that needs to be transferred. The server must be hosted on an Amazon EC2 instance that uses an Amazon Elastic File System (Amazon EFS) file system
When combination of steps should a solutions architect take to automate this task? (Select TWO)

A. Launch the EC2 instance into the same Avalability Zone as the EFS fie system
B. install an AWS DataSync agent m the on-premises data center
C. Create a secondary Amazon Elastic Block Store (Amazon EBS) volume on the EC2instance tor the data
D. Manually use an operating system copy command to push the data to the EC2 instance
E. Use AWS DataSync to create a suitable location configuration for the onprremises SFTP server

Answer: BE

Explanation:
 AWS DataSync is an online data movement and discovery service that simplifies data migration and helps users quickly, easily, and securely move their file or object data to, from, and between AWS storage services1. Users can use AWS DataSync to transfer data between on-premises and AWS storage services. To use AWS DataSync, users need to install an AWS DataSync agent in the on-premises data center. The agent is a software appliance that connects to the source or destination storage system and handles the data transfer to or from AWS over the network2. Users also need to use AWS DataSync to create a suitable location configuration for the on-premises SFTP server. A location is a logical representation of a storage system that contains files or objects that users want to transfer using DataSync. Users can create locations for NFS shares, SMB shares, HDFS file systems, self-managed object storage, Amazon S3 buckets, Amazon EFS file systems, Amazon FSx for Windows File Server file systems, Amazon FSx for Lustre file systems, Amazon FSx for OpenZFS file systems, Amazon FSx for NetApp ONTAP file systems, and AWS Snowcone devices3.

NEW QUESTION 31
- (Topic 1)
A company needs to keep user transaction data in an Amazon DynamoDB table. The company must retain the data for 7 years.
What is the MOST operationally efficient solution that meets these requirements?

A. Use DynamoDB point-in-time recovery to back up the table continuously.
B. Use AWS Backup to create backup schedules and retention policies for the table.
C. Create an on-demand backup of the table by using the DynamoDB consol
D. Store the backup in an Amazon S3 bucke
E. Set an S3 Lifecycle configuration for the S3 bucket.
F. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to invoke an AWS Lambda functio
G. Configure the Lambda function to back up the table and to store thebackup in an Amazon S3 bucke
H. Set an S3 Lifecycle configuration for the S3 bucket.

Answer: C

NEW QUESTION 35
- (Topic 1)
An ecommerce company wants to launch a one-deal-a-day website on AWS. Each day will feature exactly one product on sale for a period of 24 hours. The company wants to be able to handle millions of requests each hour with millisecond latency during peak hours.
Which solution will meet these requirements with the LEAST operational overhead?

A. Use Amazon S3 to host the full website in different S3 buckets Add Amazon CloudFront distributions Set the S3 buckets as origins for the distributions Store the order data in Amazon S3
B. Deploy the full website on Amazon EC2 instances that run in Auto Scaling groups across multiple Availability Zones Add an Application Load Balancer (ALB) to distribute the website traffic Add another ALB for the backend APIs Store the data in Amazon RDS for MySQL
C. Migrate the full application to run in containers Host the containers on Amazon Elastic Kubernetes Service (Amazon EKS) Use the Kubernetes Cluster Autoscaler to increase and decrease the number of pods to process bursts in traffic Store the data in Amazon RDS for MySQL
D. Use an Amazon S3 bucket to host the website's static content Deploy an Amazon CloudFront distributio
E. Set the S3 bucket as the origin Use Amazon API Gateway and AWS Lambda functions for the backend APIs Store the data in Amazon DynamoDB

Answer: D

Explanation:
 To launch a one-deal-a-day website on AWS with millisecond latency during peak hours and with the least operational overhead, the best option is to use an Amazon S3 bucket to host the website's static content, deploy an Amazon CloudFront distribution, set the S3 bucket as the origin, use Amazon API Gateway and AWS Lambda functions for the backend APIs, and store the data in Amazon DynamoDB. This option requires minimal operational overhead and can handle millions of requests each hour with millisecond latency during peak hours. Therefore, option D is the correct answer.
Reference: https://aws.amazon.com/blogs/compute/building-a-serverless-multi-player-game-with-aws-lambda-and-amazon-dynamodb/

NEW QUESTION 39
- (Topic 1)
A company is implementing a shared storage solution for a media application that is hosted m the AWS Cloud The company needs the ability to use SMB clients to access data The solution must he fully managed.
Which AWS solution meets these requirements?

A. Create an AWS Storage Gateway volume gatewa
B. Create a file share that uses the required client protocol Connect the application server to the file share.
C. Create an AWS Storage Gateway tape gateway Configure (apes to use Amazon S3 Connect the application server lo the tape gateway
D. Create an Amazon EC2 Windows instance Install and configure a Windows file share role on the instanc
E. Connect the application server to the file share.
F. Create an Amazon FSx for Windows File Server tile system Attach the fie system to the origin serve
G. Connect the application server to the file system

Answer: D

Explanation:
https://aws.amazon.com/fsx/lustre/
Amazon FSx has native support for Windows file system features and for the industry- standard Server Message Block (SMB) protocol to access file storage over a network. https://docs.aws.amazon.com/fsx/latest/WindowsGuide/what-is.html

NEW QUESTION 43
- (Topic 1)
A company runs multiple Windows workloads on AWS. The company's employees use Windows file shares that are hosted on two Amazon EC2 instances. The file shares synchronize data between themselves and maintain duplicate copies. The company wants a highly available and durable storage solution that preserves how users currently access the files.
What should a solutions architect do to meet these requirements?

A. Migrate all the data to Amazon S3 Set up IAM authentication for users to access files
B. Set up an Amazon S3 File Gatewa
C. Mount the S3 File Gateway on the existing EC2 Instances.
D. Extend the file share environment to Amazon FSx for Windows File Server with a Multi- AZ configuratio
E. Migrate all the data to FSx for Windows File Server.
F. Extend the file share environment to Amazon Elastic File System (Amazon EFS) with a Multi-AZ configuratio
G. Migrate all the data to Amazon EFS.

Answer: C

Explanation:
https://docs.aws.amazon.com/AWSEC2/latest/WindowsGuide/AmazonEFS.html Amazon FSx for Windows File Server provides fully managed Microsoft Windows file servers, backed by a fully native Windows file system. https://docs.aws.amazon.com/fsx/latest/WindowsGuide/what-is.html

NEW QUESTION 48
- (Topic 1)
A company recently launched Linux-based application instances on Amazon EC2 in a private subnet and launched a Linux-based bastion host on an Amazon EC2 instance in a public subnet of a VPC A solutions architect needs to connect from the on-premises
network, through the company's internet connection to the bastion host and to the application servers The solutions architect must make sure that the security groups of all the EC2 instances will allow that access
Which combination of steps should the solutions architect take to meet these requirements? (Select TWO)

A. Replace the current security group of the bastion host with one that only allows inbound access from the application instances
B. Replace the current security group of the bastion host with one that only allows inbound access from the internal IP range for the company
C. Replace the current security group of the bastion host with one that only allows inbound access from the external IP range for the company
D. Replace the current security group of the application instances with one that allows inbound SSH access from only the private IP address of the bastion host
E. Replace the current security group of the application instances with one that allows inbound SSH access from only the public IP address of the bastion host

Answer: CD

Explanation:
https://digitalcloud.training/ssh-into-ec2-in-private-subnet/

NEW QUESTION 53
- (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 56
- (Topic 1)
A solutions architect is designing the cloud architecture for a new application being deployed on AWS. The process should run in parallel while adding and removing application nodes as needed based on the number of jobs to be processed. The processor application is stateless. The solutions architect must ensure that the application is loosely coupled and the job items are durably stored.
Which design should the solutions architect use?

A. Create an Amazon SNS topic to send the jobs that need to be processed Create an Amazon Machine Image (AMI) that consists of the processor application Create a launch configuration that uses the AMI Create an Auto Scaling group using the launch configuration Set the scaling policy for the Auto Scaling group to add and remove nodes based on CPU usage
B. Create an Amazon SQS queue to hold the jobs that need to be processed Create an Amazon Machine image (AMI) that consists of the processor application Create a launch configuration that uses the AM' Create an Auto Scaling group using the launch configuration Set the scaling policy for the Auto Scaling group to add and remove nodes based on network usage
C. Create an Amazon SQS queue to hold the jobs that needs to be processed Create an Amazon Machine image (AMI) that consists of the processor application Create a launch template that uses the AMI Create an Auto Scaling group using the launch template Set the scaling policy for the Auto Scaling group to add and

remove nodes based on the number of items in the SQS queue
D. Create an Amazon SNS topic to send the jobs that need to be processed Create an Amazon Machine Image (AMI) that consists of the processor application Create a launch template that uses the AMI Create an Auto Scaling group using the launch template Set the scaling policy for the Auto Scaling group to add and remove nodes based on the number of messages published to the SNS topic

Answer: C

Explanation:
"Create an Amazon SQS queue to hold the jobs that needs to be processed. Create an Amazon EC2 Auto Scaling group for the compute application. Set the scaling policy for the Auto Scaling group to add and remove nodes based on the number of items in the SQS queue"
In this case we need to find a durable and loosely coupled solution for storing jobs. Amazon SQS is ideal for this use case and can be configured to use dynamic scaling based on the number of jobs waiting in the queue.To configure this scaling you can use the backlog per instance metric with the target value being the acceptable backlog per instance to maintain. You can calculate these numbers as follows: Backlog per instance: To calculate your backlog per instance, start with the ApproximateNumberOfMessages queue attribute to determine the length of the SQS queue

NEW QUESTION 58
- (Topic 1)
A company needs to review its AWS Cloud deployment to ensure that its Amazon S3 buckets do not have unauthorized configuration changes.
What should a solutions architect do to accomplish this goal?

A. Turn on AWS Config with the appropriate rules.
B. Turn on AWS Trusted Advisor with the appropriate checks.
C. Turn on Amazon Inspector with the appropriate assessment template.
D. Turn on Amazon S3 server access loggin
E. Configure Amazon EventBridge (Amazon Cloud Watch Events).

Answer: A

Explanation:
To ensure that Amazon S3 buckets do not have unauthorized configuration changes, a solutions architect should turn on AWS Config with the appropriate rules. AWS Config is a service that allows users to audit and assess their AWS resource configurations for compliance with industry standards and internal policies. It provides a detailed view of the resources and their configurations, including information on how the resources are related to each other. By turning on AWS Config with the appropriate rules, users can identify and remediate unauthorized configuration changes to their Amazon S3 buckets.

NEW QUESTION 59
- (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 62
- (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 67
- (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 71
- (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 74
- (Topic 1)
A company needs to store its accounting records in Amazon S3. The records must be immediately accessible for 1 year and then must be archived for an additional 9 years. No one at the company, including administrative users and root users, can be able to delete the records during the entire 10-year period. The records must be stored with maximum resiliency.
Which solution will meet these requirements?

A. Store the records in S3 Glacier for the entire 10-year perio
B. Use an access control policy to deny deletion of the records for a period of 10 years.
C. Store the records by using S3 Intelligent-Tierin
D. Use an IAM policy to deny deletion of the record
E. After 10 years, change the IAM policy to allow deletion.
F. Use an S3 Lifecycle policy to transition the records from S3 Standard to S3 Glacier Deep Archive after 1 yea
G. Use S3 Object Lock in compliance mode for a period of 10 years.
H. Use an S3 Lifecycle policy to transition the records from S3 Standard to S3 One Zone-Infrequent Access (S3 One Zone-IA) after 1 yea
I. Use S3 Object Lock in governance mode for a period of 10 years.

Answer: C

Explanation:
To meet the requirements of immediately accessible records for 1 year and then archived for an additional 9 years with maximum resiliency, we can use S3 Lifecycle policy to transition records from S3 Standard to S3 Glacier Deep Archive after 1 year. And to ensure that the records cannot be deleted by anyone, including administrative and root users, we can use S3 Object Lock in compliance mode for a period of 10 years. Therefore, the correct answer is option C. Reference: https://docs.aws.amazon.com/AmazonS3/latest/userguide/object-lock.html

NEW QUESTION 75
- (Topic 1)
A company needs guaranteed Amazon EC2 capacity in three specific Availability Zones in a specific AWS Region for an upcoming event that will last 1 week.
What should the company do to guarantee the EC2 capacity?

A. Purchase Reserved instances that specify the Region needed
B. Create an On Demand Capacity Reservation that specifies the Region needed
C. Purchase Reserved instances that specify the Region and three Availability Zones needed
D. Create an On-Demand Capacity Reservation that specifies the Region and three Availability Zones needed

Answer: D

Explanation:
https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-capacity- reservations.html
Reserve instances: You will have to pay for the whole term (1 year or 3years) which is not cost effective

NEW QUESTION 77
- (Topic 1)
A company has several web servers that need to frequently access a common Amazon RDS MySQL Multi-AZ DB instance The company wants a secure method for the web servers to connect to the database while meeting a security requirement to rotate user credentials frequently.
Which solution meets these requirements?

A. Store the database user credentials in AWS Secrets Manager Grant the necessary IAM permissions to allow the web servers to access AWS Secrets Manager
B. Store the database user credentials in AWS Systems Manager OpsCenter Grant the necessary IAM permissions to allow the web servers to access OpsCenter
C. Store the database user credentials in a secure Amazon S3 bucket Grant the necessary IAM permissions to allow the web servers to retrieve credentials and access the databas
D. D.Store the database user credentials in files encrypted with AWS Key Management Service (AWS KMS) on the web server file syste
E. The web server should be able to decrypt the files and access the database

Answer: A

Explanation:
AWS Secrets Manager helps you protect secrets needed to access your applications, services, and IT resources. The service enables you to easily rotate, manage, and retrieve database credentials, API keys, and other secrets throughout their lifecycle.
https://docs.aws.amazon.com/secretsmanager/latest/userguide/intro.html
Secrets Manager enables you to replace hardcoded credentials in your code, including passwords, with an API call to Secrets Manager to retrieve the secret programmatically. This helps ensure the secret can't be compromised by someone examining your code, because the secret no longer exists in the code. Also, you can configure Secrets Manager to automatically rotate the secret for you according to a specified schedule. This enables you to replace long-term secrets with short-term ones, significantly reducing the risk of compromise.

NEW QUESTION 81
- (Topic 1)
A bicycle sharing company is developing a multi-tier architecture to track the location of its bicycles during peak operating hours The company wants to use these data points in its existing analytics platform A solutions architect must determine the most viable multi-tier option to support this architecture The data points must be accessible from the REST API.
Which action meets these requirements for storing and retrieving location data?

A. Use Amazon Athena with Amazon S3
B. Use Amazon API Gateway with AWS Lambda
C. Use Amazon QuickSight with Amazon Redshift.
D. Use Amazon API Gateway with Amazon Kinesis Data Analytics

Answer: D

Explanation:
https://aws.amazon.com/solutions/implementations/aws-streaming-data-solution-for- amazon-kinesis/

NEW QUESTION 84
- (Topic 1)
A company recently signed a contract with an AWS Managed Service Provider (MSP) Partner for help with an application migration initiative. A solutions architect needs to share an Amazon Machine Image (AMI) from an existing AWS account with the MSP Partner's AWS account. The AMI is backed by Amazon Elastic Block Store (Amazon EBS) and uses a customer managed customer master key (CMK) to encrypt EBS volume snapshots.
What is the MOST secure way for the solutions architect to share the AMI with the MSP Partner's AWS account?

A. Make the encrypted AMI and snapshots publicly availabl
B. Modify the CMK's key policy to allow the MSP Partner's AWS account to use the key
C. Modify the launchPermission property of the AM
D. Share the AMI with the MSP Partner's AWS account onl
E. Modify the CMK's key policy to allow the MSP Partner's AWS account to use the key.
F. Modify the launchPermission property of the AMI Share the AMI with the MSP Partner's AWS account onl
G. Modify the CMK's key policy to trust a new CMK that is owned by the MSP Partner for encryption.
H. Export the AMI from the source account to an Amazon S3 bucket in the MSP Partner's AWS accoun
I. Encrypt the S3 bucket with a CMK that is owned by the MSP Partner Copy and launch the AMI in the MSP Partner's AWS account.

Answer: B

Explanation:
Share the existing KMS key with the MSP external account because it has already been used to encrypt the AMI snapshot.
https://docs.aws.amazon.com/kms/latest/developerguide/key-policy-modifying-external- accounts.html

NEW QUESTION 85
- (Topic 1)
A company recently migrated to AWS and wants to implement a solution to protect the traffic that flows in and out of the production VPC. The company had an inspection server in its on-premises data center. The inspection server performed specific operations such as traffic flow inspection and traffic filtering. The company wants to have the same functionalities in the AWS Cloud.
Which solution will meet these requirements?

A. Use Amazon GuardDuty for traffic inspection and traffic filtering in the production VPC
B. Use Traffic Mirroring to mirror traffic from the production VPC for traffic inspection and filtering.

C. Use AWS Network Firewall to create the required rules for traffic inspection and traffic filtering for the production VPC.
D. Use AWS Firewall Manager to create the required rules for traffic inspection and traffic filtering for the production VPC.

Answer: C

Explanation:
AWS Network Firewall supports both inspection and filtering as required

NEW QUESTION 90
- (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 95
- (Topic 1)
A company is storing backup files by using Amazon S3 Standard storage. The files are accessed frequently for 1 month. However, the files are not accessed after 1 month. The company must keep the files indefinitely.
Which storage solution will meet these requirements MOST cost-effectively?

A. Configure S3 Intelligent-Tiering to automatically migrate objects.
B. Create an S3 Lifecycle configuration to transition objects from S3 Standard to S3 Glacier Deep Archive after 1 month.
C. Create an S3 Lifecycle configuration to transition objects from S3 Standard to S3 Standard-Infrequent Access (S3 Standard-IA) after 1 month.
D. Create an S3 Lifecycle configuration to transition objects from S3 Standard to S3 One Zone-Infrequent Access (S3 One Zone-IA) after 1 month.

Answer: B

Explanation:
The storage solution that will meet these requirements most cost-effectively is B: Create an S3 Lifecycle configuration to transition objects from S3 Standard to S3 Glacier Deep Archive after 1 month. Amazon S3 Glacier Deep Archive is a secure, durable, and extremely low-cost Amazon S3 storage class for long-term retention of data that is rarely accessed and for which retrieval times of several hours are acceptable. It is the lowest-cost storage option in Amazon S3, making it a cost-effective choice for storing backup files that are not accessed after 1 month. You can use an S3 Lifecycle configuration to automatically transition objects from S3 Standard to S3 Glacier Deep Archive after 1 month. This will minimize the storage costs for the backup files that are not accessed frequently.

NEW QUESTION 99
- (Topic 1)
A company uses Amazon S3 to store its confidential audit documents. The S3 bucket uses bucket policies to restrict access to audit team IAM user credentials according to the principle of least privilege. Company managers are worried about accidental deletion of documents in the S3 bucket and want a more secure solution.
What should a solutions architect do to secure the audit documents?

A. Enable the versioning and MFA Delete features on the S3 bucket.
B. Enable multi-factor authentication (MFA) on the IAM user credentials for each auditteam IAM user account.
C. Add an S3 Lifecycle policy to the audit team's IAM user accounts to deny the s3:DeleteObject action during audit dates.
D. Use AWS Key Management Service (AWS KMS) to encrypt the S3 bucket and restrict audit team IAM user accounts from accessing the KMS key.

Answer: A

NEW QUESTION 100
- (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 105

- (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 106
- (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 110
- (Topic 1)
An application runs on an Amazon EC2 instance in a VPC. The application processes logs that are stored in an Amazon S3 bucket. The EC2 instance needs to access the S3 bucket without connectivity to the internet.
Which solution will provide private network connectivity to Amazon S3?

A. Create a gateway VPC endpoint to the S3 bucket.
B. Stream the logs to Amazon CloudWatch Log
C. Export the logs to the S3 bucket.
D. Create an instance profile on Amazon EC2 to allow S3 access.
E. Create an Amazon API Gateway API with a private link to access the S3 endpoint.

Answer: A

Explanation:
VPC endpoint allows you to connect to AWS services using a private network instead of using the public Internet

NEW QUESTION 112
- (Topic 1)
A company has thousands of edge devices that collectively generate 1 TB of status alerts each day. Each alert is approximately 2 KB in size. A solutions architect needs to implement a solution to ingest and store the alerts for future analysis.
The company wants a highly available solution. However, the company needs to minimize costs and does not want to manage additional infrastructure. Ad ditionally, the company wants to keep 14 days of data available for immediate analysis and archive any data older than 14 days.
What is the MOST operationally efficient solution that meets these requirements?

A. Create an Amazon Kinesis Data Firehose delivery stream to ingest the alerts Configure the Kinesis Data Firehose stream to deliver the alerts to an Amazon S3 bucket Set up an S3 Lifecycle configuration to transition data to Amazon S3 Glacier after 14 days
B. Launch Amazon EC2 instances across two Availability Zones and place them behind an Elastic Load Balancer to ingest the alerts Create a script on the EC2 instances that will store tne alerts m an Amazon S3 bucket Set up an S3 Lifecycle configuration to transition data to Amazon S3 Glacier after 14 days
C. Create an Amazon Kinesis Data Firehose delivery stream to ingest the alerts Configure the Kinesis Data Firehose stream to deliver the alerts to an Amazon Elasticsearch Service (Amazon ES) duster Set up the Amazon ES cluster to take manual snapshots every day and delete data from the duster that is older than 14 days
D. Create an Amazon Simple Queue Service (Amazon SQS i standard queue to ingest the alerts and set the message retention period to 14 days Configure consumers to poll the SQS queue check the age of the message and analyze the message data as needed If the message is 14 days old the consumer should copy the message to an Amazon S3 bucketand delete the message from the SQS queue

Answer: A

Explanation:
https://aws.amazon.com/kinesis/data-firehose/features/?nc=sn&loc=2#:~:text=into%20Amazon%20S3%2C%20Amazon%20Redshift%2C%20Amazon%20OpenS
earch%20Service%2C%20Kinesis,Delivery%20streams

NEW QUESTION 114
- (Topic 1)
A company is designing an application. The application uses an AWS Lambda function to receive information through Amazon API Gateway and to store the information in an Amazon Aurora PostgreSQL database.
During the proof-of-concept stage, the company has to increase the Lambda quotas significantly to handle the high volumes of data that the company needs to load into the database. A solutions architect must recommend a new design to improve scalability and minimize the configuration effort.
Which solution will meet these requirements?

A. Refactor the Lambda function code to Apache Tomcat code that runs on Amazon EC2 instance
B. Connect the database by using native Java Database Connectivity (JDBC) drivers.
C. Change the platform from Aurora to Amazon DynamoD
D. Provision a DynamoDB Accelerator (DAX) cluste
E. Use the DAX client SDK to point the existing DynamoDB API calls at the DAX cluster.
F. Set up two Lambda function
G. Configure one function to receive the informatio
H. Configure the other function to load the information into the databas
I. Integrate the Lambda functions by using Amazon Simple Notification Service (Amazon SNS).
J. Set up two Lambda function
K. Configure one function to receive the informatio
L. Configure the other function to load the information into the databas
M. Integrate the Lambda functions by using an Amazon Simple Queue Service (Amazon SQS) queue.

Answer: B

Explanation:
 bottlenecks can be avoided with queues (SQS).

NEW QUESTION 116
- (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 121
- (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 124
- (Topic 1)
A company wants to move a multi-tiered application from on premises to the AWS Cloud to improve the application's performance. The application consists of application tiers that communicate with each other by way of RESTful services. Transactions are dropped when one tier becomes overloaded. A solutions architect must design a solution that resolves these issues and modernizes the application.
Which solution meets these requirements and is the MOST operationally efficient?

A. Use Amazon API Gateway and direct transactions to the AWS Lambda functions as the application laye
B. Use Amazon Simple Queue Service (Amazon SQS) as the communication layer between application services.
C. Use Amazon CloudWatch metrics to analyze the application performance history to determine the server's peak utilization during the performance failure
D. Increase the size of the application server's Amazon EC2 instances to meet the peak requirements.
E. Use Amazon Simple Notification Service (Amazon SNS) to handle the messaging between application servers running on Amazon EC2 in an Auto Scaling grou
F. Use Amazon CloudWatch to monitor the SNS queue length and scale up and down as required.
G. Use Amazon Simple Queue Service (Amazon SQS) to handle the messaging between application servers running on Amazon EC2 in an Auto Scaling grou
H. Use Amazon CloudWatch to monitor the SQS queue length and scale up when communication failures are detected.

Answer: A

Explanation:
https://aws.amazon.com/getting-started/hands-on/build-serverless-web-app-lambda-apigateway-s3-dynamodb-cognito/module-4/
Build a Serverless Web Application with AWS Lambda, Amazon API Gateway, AWS Amplify, Amazon DynamoDB, and Amazon Cognito. This example showed similar setup as question: Build a Serverless Web Application with AWS Lambda, Amazon API Gateway, AWS Amplify, Amazon DynamoDB, and Amazon Cognito

NEW QUESTION 126
- (Topic 1)
A company wants to reduce the cost of its existing three-tier web architecture. The web, application, and database servers are running on Amazon EC2 instances for the development, test, and production environments. The EC2 instances average 30% CPU utilization during peak hours and 10% CPU utilization during non-peak hours.
The production EC2 instances run 24 hours a day. The development and test EC2 instances run for at least 8 hours each day. The company plans to implement automation to stop the development and test EC2 instances when they are not in use.
Which EC2 instance purchasing solution will meet the company's requirements MOST cost-effectively?

A. Use Spot Instances for the production EC2 instance
B. Use Reserved Instances for the development and test EC2 instances.
C. Use Reserved Instances for the production EC2 instance
D. Use On-Demand Instances for the development and test EC2 instances.
E. Use Spot blocks for the production EC2 instance
F. Use Reserved Instances for the development and test EC2 instances.
G. Use On-Demand Instances for the production EC2 instance
H. Use Spot blocks for the development and test EC2 instances.

Answer: B

NEW QUESTION 131
- (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 136
- (Topic 2)
A medical records company is hosting an application on Amazon EC2 instances. The application processes customer data files that are stored on Amazon S3. The EC2 instances are hosted in public subnets. The EC2 instances access Amazon S3 over the internet, but they do not require any other network access.
A new requirement mandates that the network traffic for file transfers take a private route and not be sent over the internet.
Which change to the network architecture should a solutions architect recommend to meet this requirement?

A. Create a NAT gatewa
B. Configure the route table for the public subnets to send traffic to Amazon S3 through the NAT gateway.
C. Configure the security group for the EC2 instances to restrict outbound traffic so that only traffic to the S3 prefix list is permitted.
D. Move the EC2 instances to private subnet
E. Create a VPC endpoint for Amazon S3, and link the endpoint to the route table for the private subnets
F. Remove the internet gateway from the VP
G. Set up an AWS Direct Connect connection, and route traffic to Amazon S3 over the Direct Connect connection.

Answer: C

Explanation:
 To meet the new requirement of transferring files over a private route, the EC2 instances should be moved to private subnets, which do not have direct access to the internet. This ensures that the traffic for file transfers does not go over the internet. To enable the EC2 instances to access Amazon S3, a VPC endpoint for Amazon S3 can be created. VPC endpoints allow resources within a VPC to communicate with resources in other services without the traffic being sent over the internet. By linking the VPC endpoint to the route table for the private subnets, the EC2 instances can access Amazon S3 over a
private connection within the VPC.

NEW QUESTION 138
- (Topic 2)
A company needs to save the results from a medical trial to an Amazon S3 repository. The repository must allow a few scientists to add new files and must restrict all other users to read-only access. No users can have the ability to modify or delete any files in the repository. The company must keep every file in the repository for a minimum of 1 year after its creation date.
Which solution will meet these requirements?

A. Use S3 Object Lock In governance mode with a legal hold of 1 year
B. Use S3 Object Lock in compliance mode with a retention period of 365 days.
C. Use an IAM role to restrict all users from deleting or changing objects in the S3 bucket Use an S3 bucket policy to only allow the IAM role
D. Configure the S3 bucket to invoke an AWS Lambda function every tune an object is added Configure the function to track the hash of the saved object to that modified objects can be marked accordingly

Answer: B

Explanation:
 In compliance mode, a protected object version can't be overwritten or deleted by any user, including the root user in your AWS account. When an object is locked in compliance mode, its retention mode can't be changed, and its retention period can't be shortened. Compliance mode helps ensure that an object version can't be overwritten or deleted for the duration of the retention period. In governance mode, users can't overwrite or delete an object version or alter its lock settings unless they have special permissions. With governance mode, you protect objects against being deleted by most users, but you can still grant some users permission to alter the retention settings or delete the object if necessary. In Governance mode, Objects can be deleted by some users with special permissions, this is against the requirement.
Compliance:
- Object versions can't be overwritten or deleted by any user, including the root user
- Objects retention modes can't be changed, and retention periods can't be shortened
Governance:
- Most users can't overwrite or delete an object version or alter its lock settings
- Some users have special permissions to change the retention or delete the object

NEW QUESTION 139
- (Topic 2)
A company wants to migrate its on-premises data center to AWS. According to the company's compliance requirements, the company can use only the ap-northeast-3 Region. Company administrators are not permitted to connect VPCs to the internet.
Which solutions will meet these requirements? (Choose two.)

A. Use AWS Control Tower to implement data residency guardrails to deny internet access and deny access to all AWS Regions except ap-northeast-3.
B. Use rules in AWS WAF to prevent internet acces
C. Deny access to all AWS Regions except ap-northeast-3 in the AWS account settings.
D. Use AWS Organizations to configure service control policies (SCPS) that prevent VPCs from gaining internet acces
E. Deny access to all AWS Regions except ap-northeast-3.
F. Create an outbound rule for the network ACL in each VPC to deny all traffic from 0.0.0.0/0. Create an IAM policy for each user to prevent the use of any AWS Region other than ap-northeast-3.
G. Use AWS Config to activate managed rules to detect and alert for internet gateways and to detect and alert for new resources deployed outside of ap-northeast-3.

Answer: AC

Explanation:

https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps_examples_vpc.html#example_vpc_2

NEW QUESTION 144
- (Topic 2)
A solutions architect needs to help a company optimize the cost of running an application on AWS. The application will use Amazon EC2 instances, AWS Fargate, and AWS Lambda for compute within the architecture.

The EC2 instances will run the data ingestion layer of the application. EC2 usage will be sporadic and unpredictable. Workloads that run on EC2 instances can be interrupted at any time. The application front end will run on Fargate, and Lambda will serve the API layer. The front-end utilization and API layer utilization will be predictable over the course of the next year.

Which combination of purchasing options will provide the MOST cost-effective solution for hosting this application? (Choose two.)

A. Use Spot Instances for the data ingestion layer
B. Use On-Demand Instances for the data ingestion layer
C. Purchase a 1-year Compute Savings Plan for the front end and API layer.
D. Purchase 1-year All Upfront Reserved instances for the data ingestion layer.
E. Purchase a 1-year EC2 instance Savings Plan for the front end and API layer.

Answer: AC

Explanation:
EC2 instance Savings Plan saves 72% while Compute Savings Plans saves 66%. But according to link, it says "Compute Savings Plans provide the most flexibility and help to reduce your costs by up to 66%. These plans automatically apply to EC2 instance usage regardless of instance family, size, AZ, region, OS or tenancy, and also apply to Fargate and Lambda usage." EC2 instance Savings Plans are not applied to Fargate or Lambda

NEW QUESTION 145
- (Topic 2)
A company wants to use the AWS Cloud to make an existing application highly available and resilient. The current version of the application resides in the company's data center. The application recently experienced data loss after a database server crashed because of an unexpected power outage.
The company needs a solution that avoids any single points of failure. The solution must give the application the ability to scale to meet user demand.
Which solution will meet these requirements?

A. Deploy the application servers by using Amazon EC2 instances in an Auto Scaling group across multiple Availability Zone
B. Use an Amazon RDS DB instance in a Multi-AZ configuration.
C. Deploy the application servers by using Amazon EC2 instances in an Auto Scaling group in a single Availability Zon
D. Deploy the databaseon an EC2 instanc
E. Enable EC2 Auto Recovery.
F. Deploy the application servers by using Amazon EC2 instances in an Auto Scalinggroup across multiple Availability Zone
G. Use an Amazon RDS DB instance with a read replica in a single Availability Zon
H. Promote the read replica to replace the primary DB instance if the primary DB instance fails.
I. Deploy the application servers by using Amazon EC2 instances in an Auto Scaling group across multiple Availability Zones Deploy the primary and secondary database servers on EC2 instances across multiple Availability Zones Use Amazon Elastic Block Store (Amazon EBS) Multi-Attach to create shared storage between the instances.

Answer: A

Explanation:
Deploy the application servers by using Amazon EC2 instances in an Auto Scaling group across multiple Availability Zones. Use an Amazon RDS DB instance in a Multi-AZ configuration. To make an existing application highly available and resilient while avoiding any single points of failure and giving the application the ability to scale to meet user demand, the best solution would be to deploy the application servers using Amazon EC2 instances in an Auto Scaling group across multiple Availability Zones and use an Amazon RDS DB instance in a Multi-AZ configuration. By using an Amazon RDS DB instance in a Multi-AZ configuration, the database is automatically replicated across multiple Availability Zones, ensuring that the database is highly available and can withstand the failure of a single Availability Zone. This provides fault tolerance and avoids any single points of failure.

NEW QUESTION 150
- (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 152
- (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 a new 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 154
- (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 157
- (Topic 2)
A company runs an application using Amazon ECS. The application creates esi/ed versions of an original image and then makes Amazon S3 API calls to store the resized images in Amazon S3.
How can a solutions architect ensure that the application has permission to access Amazon S3?

A. Update the S3 role in AWS IAM to allow read/write access from Amazon ECS, and then relaunch the container.
B. Create an IAM role with S3 permissions, and then specify that role as the taskRoleAm in the task definition.
C. Create a security group that allows access from Amazon ECS to Amazon S3, and update the launch configuration used by the ECS cluster.
D. Create an IAM user with S3 permissions, and then relaunch the Amazon EC2 instancesfor the ECS cluster while logged in as this account.

Answer: B

Explanation:
https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-resource-ecs-taskdefinition.html

NEW QUESTION 160
- (Topic 2)
A company is concerned about the security of its public web application due to recent web attacks. The application uses an Application Load Balancer (ALB). A solutions architect must reduce the risk of DDoS attacks against the application.
What should the solutions architect do to meet this requirement?

A. Add an Amazon Inspector agent to the ALB.
B. Configure Amazon Macie to prevent attacks.
C. Enable AWS Shield Advanced to prevent attacks.
D. Configure Amazon GuardDuty to monitor the ALB.

Answer: C

Explanation:
AWS Shield Advanced provides expanded DDoS attack protection for your Amazon EC2 instances, Elastic Load Balancing load balancers, CloudFront distributions, Route 53 hosted zones, and AWS Global Accelerator standard accelerators. https://docs.aws.amazon.com/waf/latest/developerguide/what-is-aws-waf.html

NEW QUESTION 161
- (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 163
- (Topic 2)
An entertainment company is using Amazon DynamoDB to store media metadata. The application is read intensive and experiencing delays. The company does not have staff to handle additional operational overhead and needs to improve the performance efficiency of DynamoDB without reconfiguring the application.
What should a solutions architect recommend to meet this requirement?

A. Use Amazon ElastiCache for Redis.
B. Use Amazon DynamoDB Accelerator (DAX).
C. Replicate data by using DynamoDB global tables.
D. Use Amazon ElastiCache for Memcached with Auto Discovery enabled.

Answer: B

Explanation:
https://aws.amazon.com/dynamodb/dax/

NEW QUESTION 164
- (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 167
- (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 170
- (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 175
- (Topic 2)
A gaming company hosts a browser-based application on AWS. The users of the application consume a large number of videos and images that are stored in Amazon S3. This content is the same for all users.
The application has increased in popularity, and millions of users worldwide are accessing these media files. The company wants to provide the files to the users while reducing the load on the origin.
Which solution meets these requirements MOST cost-effectively?

A. Deploy an AWS Global Accelerator accelerator in front of the web servers.
B. Deploy an Amazon CloudFront web distribution in front of the S3 bucket.
C. Deploy an Amazon ElastiCache for Redis instance in front of the web servers.
D. Deploy an Amazon ElastiCache for Memcached instance in front of the web servers.

Answer: B

Explanation:
ElastiCache, enhances the performance of web applications by quickly retrieving information from fully-managed in-memory data stores. It utilizes Memcached and Redis, and manages to considerably reduce the time your applications would, otherwise, take to read data from disk-based databases. Amazon CloudFront supports dynamic content from HTTP and WebSocket protocols, which are based on the Transmission Control Protocol (TCP) protocol. Common use cases include dynamic API calls, web pages and web applications, as well as an application's static files such as audio and images. It also supports on-demand media streaming over HTTP. AWS Global Accelerator supports both User Datagram Protocol (UDP) and TCP-based protocols. It is commonly used for non- HTTP use cases, such as gaming, IoT and voice over IP. It is also good for HTTP use cases that need static IP addresses or fast regional failover

NEW QUESTION 177
- (Topic 2)
A hospital wants to create digital copies for its large collection of historical written records. The hospital will continue to add hundreds of new documents each day. The hospital's data team will scan the documents and will upload the documents to the AWS Cloud.
A solutions architect must implement a solution to analyze the documents, extract the medical information, and store the documents so that an application can run SQL queries on the data. The solution must maximize scalability and operational efficiency.
Which combination of steps should the solutions architect take to meet these requirements? (Select TWO.)

A. Write the document information to an Amazon EC2 instance that runs a MySQL database.
B. Write the document information to an Amazon S3 bucke
C. Use Amazon Athena to query the data.
D. Create an Auto Scaling group of Amazon EC2 instances to run a custom application that processes the scanned files and extracts the medical information.
E. Create an AWS Lambda function that runs when new documents are uploade
F. Use Amazon Rekognition to convert the documents to raw tex
G. Use Amazon Transcribe Medical to detect and extract relevant medical information from the text.
H. Create an AWS Lambda function that runs when new documents are uploade
I. Use Amazon Textract to convert the documents to raw tex
J. Use Amazon Comprehend Medical to detect and extract relevant medical information from the text.

Answer: BE

Explanation:
This solution meets the requirements of creating digital copies for a large collection of historical written records, analyzing the documents, extracting the medical information, and storing the documents so that an application can run SQL queries on the data. Writing the document information to an Amazon S3 bucket can provide scalable and durable storage for the scanned files. Using Amazon Athena to query the data can provide serverless and interactive SQL analysis on data stored in S3. Creating an AWS Lambda function that runs when new documents are uploaded can provide event-driven and serverless processing of the scanned files. Using Amazon Textract to convert the documents to raw text can provide
accurate optical character recognition (OCR) and extraction of structured data such as tables and forms from documents using artificial intelligence (AI). Using Amazon Comprehend Medical to detect and extract relevant medical information from the text can provide natural language processing (NLP) service that uses machine learning that has been pre-trained to understand and extract health data from medical text.
Option A is incorrect because writing the document information to an Amazon EC2 instance that runs a MySQL database can increase the infrastructure overhead and complexity, and it may not be able to handle large volumes of data. Option C is incorrect because creating an Auto Scaling group of Amazon EC2 instances to run a custom application that processes the scanned files and extracts the medical information can increase the infrastructure overhead and complexity, and it may not be able to leverage existing AI and NLP services such as Textract and Comprehend Medical. Option D is incorrect because using Amazon Rekognition to convert the documents to raw text can provide image and video analysis, but it does not support OCR or extraction of structured data from documents. Using Amazon Transcribe Medical to detect and extract relevant medical information from the text can provide speech-to-text transcription service for medical conversations, but it does not support text analysis or extraction of health data from medical text.
References:
? https://aws.amazon.com/s3/
? https://aws.amazon.com/athena/
? https://aws.amazon.com/lambda/
? https://aws.amazon.com/textract/
? https://aws.amazon.com/comprehend/medical/

NEW QUESTION 182

- (Topic 2)
An ecommerce company has an order-processing application that uses Amazon API Gateway and an AWS Lambda function. The application stores data in an Amazon Aurora PostgreSQL database. During a recent sales event, a sudden surge in customer orders occurred. Some customers experienced timeouts and the application did not process the orders of those customers A solutions architect determined that the CPU utilization and memory utilization were high on the database because of a large number of open connections The solutions architect needs to prevent the timeout errors while making the least possible changes to the application.

Which solution will meet these requirements?

A. Configure provisioned concurrency for the Lambda function Modify the database to be a global database in multiple AWS Regions
B. Use Amazon RDS Proxy to create a proxy for the database Modify the Lambda function to use the RDS Proxy endpoint instead of the database endpoint
C. Create a read replica for the database in a different AWS Region Use query string parameters in API Gateway to route traffic to the read replica
D. Migrate the data from Aurora PostgreSQL to Amazon DynamoDB by using AWS Database Migration Service (AWS DMS| Modify the Lambda function to use the OynamoDB table

Answer: B

Explanation:
Many applications, including those built on modern serverless architectures, can have a large number of open connections to the database server and may open and close database connections at a high rate, exhausting database memory and compute resources. Amazon RDS Proxy allows applications to pool and share connections established with the database, improving database efficiency and application scalability. https://aws.amazon.com/id/rds/proxy/

NEW QUESTION 187
- (Topic 2)
A solutions architect is optimizing a website for an upcoming musical event. Videos of the performances will be streamed in real time and then will be available on demand. The event is expected to attract a global online audience.
Which service will improve the performance of both the real-lime and on-demand streaming?

A. Amazon CloudFront
B. AWS Global Accelerator
C. Amazon Route 53
D. Amazon S3 Transfer Acceleration

Answer: A

Explanation:
You can use CloudFront to deliver video on demand (VOD) or live streaming video using any HTTP origin. One way you can set up video workflows in the cloud is by using CloudFront together with AWS Media Services. https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/on-demand- streaming-video.html

NEW QUESTION 190
- (Topic 2)
A company's web application is running on Amazon EC2 instances behind an Application Load Balancer. The company recently changed its policy, which now requires the application to be accessed from one specific country only.
Which configuration will meet this requirement?

A. Configure the security group for the EC2 instances.
B. Configure the security group on the Application Load Balancer.
C. Configure AWS WAF on the Application Load Balancer in a VPC.
D. Configure the network ACL for the subnet that contains the EC2 instances.

Answer: C

Explanation:
https://aws.amazon.com/about-aws/whats-new/2017/10/aws-waf-now-supports- geographic-match/

NEW QUESTION 195
- (Topic 2)
Organizers for a global event want to put daily reports online as static HTML pages. The pages are expected to generate millions of views from users around the world. The files are stored In an Amazon S3 bucket. A solutions architect has been asked to design an efficient and effective solution.
Which action should the solutions architect take to accomplish this?

A. Generate presigned URLs for the files.
B. Use cross-Region replication to all Regions.
C. Use the geoproximtty feature of Amazon Route 53.
D. Use Amazon CloudFront with the S3 bucket as its origin.

Answer: D

Explanation:
Amazon CloudFront is a content delivery network (CDN) that speeds up the delivery of static and dynamic web content, such as HTML pages, images, and videos. By using CloudFront, the HTML pages will be served to users from the edge location that is closest to them, resulting in faster delivery and a better user experience. CloudFront can also handle the high traffic and large number of requests expected for the global event, ensuring that the HTML pages are available and accessible to users around the world.

NEW QUESTION 197
- (Topic 2)
A company hosts a website analytics application on a single Amazon EC2 On-Demand Instance. The analytics software is written in PHP and uses a MySQL database. The analytics software, the web server that provides PHP, and the database server are all hosted on the EC2 instance. The application is showing signs of performance degradation during busy times and is presenting 5xx errors. The company needs to make the application scale seamlessly.

Which solution will meet these requirements MOST cost-effectively?

A. Migrate the database to an Amazon RDS for MySQL DB instanc
B. Create an AMI of the web applicatio
C. Use the AMI to launch a second EC2 On-Demand Instanc
D. Use an Application Load Balancer to distribute the load to each EC2 instance.
E. Migrate the database to an Amazon RDS for MySQL DB instanc
F. Create an AMI of the web applicatio
G. Use the AMI to launch a second EC2 On-Demand Instanc
H. Use Amazon Route 53 weighted routing to distribute the load across the two EC2 instances.
I. Migrate the database to an Amazon Aurora MySQL DB instanc
J. Create an AWS Lambda function to stop the EC2 instance and change the instance typ
K. Create an Amazon CloudWatch alarm to invoke the Lambda function when CPU utilization surpasses 75%.
L. Migrate the database to an Amazon Aurora MySQL DB instanc
M. Create an AMI of the web applicatio
N. Apply the AMI to a launch templat
O. Create an Auto Scaling group with the launch template Configure the launch template to use a Spot Flee
P. Attach an Application Load Balancer to the Auto Scaling group.

Answer: D

Explanation:
Migrate the database to Amazon Aurora MySQL - this will let the DB scale on it's own; it'll scale automatically without needing adjustment. Create AMI of the web app and using a launch template - this will make the creating of any future instances of the app seamless. They can then be added to the auto scaling group which will save them money as it will scale up and down based on demand. Using a spot fleet to launch instances- This solves the "MOST cost-effective" portion of the question as spot instances come at a huge discount at the cost of being terminated at any time Amazon deems fit. I think this is why there's a bit of disagreement on this. While it's the most cost effective, it would be a terrible choice if amazon were to terminate that spot instance during a busy period.

NEW QUESTION 198
- (Topic 2)
A company runs a stateless web application in production on a group of Amazon EC2 On- Demand Instances behind an Application Load Balancer. The application experiences heavy usage during an 8-hour period each business day. Application usage is moderate and steady overnight Application usage is low during weekends.
The company wants to minimize its EC2 costs without affecting the availability of the application.
Which solution will meet these requirements?

A. Use Spot Instances for the entire workload.
B. Use Reserved instances for the baseline level of usage Use Spot Instances for any additional capacity that the application needs.
C. Use On-Demand Instances for the baseline level of usag
D. Use Spot Instances for any additional capacity that the application needs
E. Use Dedicated Instances for the baseline level of usag
F. Use On-Demand Instances for any additional capacity that the application needs

Answer: B

Explanation:
Reserved is cheaper than on demand the company has. And it's meet the availabilty (HA) requirement as to spot instance that can be disrupted at any time.
PRICING BELOW. On- Demand: 0% There's no commitment from you. You pay the most with this option. Reserved : 40%-60%1-year or 3-year commitment from you. You save money from that commitment. Spot 50%-90% Ridiculously inexpensive because there's no commitment from the AWS side.

NEW QUESTION 200
- (Topic 2)
A solutions architect is designing a customer-facing application for a company. The application's database will have a clearly defined access pattern throughout the year and will have a variable number of reads and writes that depend on the time of year. The company must retain audit records for the database for 7 days. The recovery point objective (RPO) must be less than 5 hours.
Which solution meets these requirements?

A. Use Amazon DynamoDB with auto scaling Use on-demand backups and Amazon DynamoDB Streams
B. Use Amazon Redshif
C. Configure concurrency scalin
D. Activate audit loggin
E. Perform database snapshots every 4 hours.
F. Use Amazon RDS with Provisioned IOPS Activate the database auditing parameter Perform database snapshots every 5 hours
G. Use Amazon Aurora MySQL with auto scalin
H. Activate the database auditing parameter

Answer: A

Explanation:
This solution meets the requirements of a customer-facing application that has a clearly defined access pattern throughout the year and a variable number of reads and writes that depend on the time of year. Amazon DynamoDB is a fully managed NoSQL database service that can handle any level of request traffic and data size. DynamoDB auto scaling can automatically adjust the provisioned read and write capacity based on the actual workload. DynamoDB on-demand backups can create full backups of the tables for data protection and archival purposes. DynamoDB Streams can capture a time-ordered sequence of item-level modifications in the tables for audit purposes.
Option B is incorrect because Amazon Redshift is a data warehouse service that is designed for analytical workloads, not for customer-facing applications. Option C is incorrect because Amazon RDS with Provisioned IOPS can provide consistent performance for relational databases, but it may not be able to handle unpredictable spikes in traffic and data size. Option D is incorrect because Amazon Aurora MySQL with auto scaling can provide high performance and availability for relational databases, but it does not support audit logging as a parameter.
References:
? https://aws.amazon.com/dynamodb/
? https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/AutoScalin g.html

? https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/BackupRe store.html
? https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Streams.ht ml

NEW QUESTION 202
- (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 203
- (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 206
- (Topic 3)
A company has migrated an application to Amazon EC2 Linux instances. One of these EC2 instances runs several 1-hour tasks on a schedule. These tasks were written by different teams and have no common programming language. The company is concerned about performance and scalability while these tasks run on a single instance. A solutions architect needs to implement a solution to resolve these concerns.
Which solution will meet these requirements with the LEAST operational overhead?

A. Use AWS Batch to run the tasks as job
B. Schedule the jobs by using Amazon EventBridge (Amazon CloudWatch Events).
C. Convert the EC2 instance to a containe
D. Use AWS App Runner to create the containeron demand to run the tasks as jobs.
E. Copy the tasks into AWS Lambda function
F. Schedule the Lambda functions by using Amazon EventBridge (Amazon CloudWatch Events).
G. Create an Amazon Machine Image (AMI) of the EC2 instance that runs the task
H. Create an Auto Scaling group with the AMI to run multiple copies of the instance.

Answer: A

Explanation:
AWS Batch is a fully managed service that enables users to run batch jobs on AWS. It can handle different types of tasks written in different languages and run them on EC2 instances. It also integrates with Amazon EventBridge (Amazon CloudWatch Events) to schedule jobs based on time or event triggers. This solution will meet the requirements of performance, scalability and low operational overhead12.
* B. Convert the EC2 instance to a container. Use AWS App Runner to create the container on demand to run the tasks as jobs. This solution will not meet the requirement of low operational overhead, as it involves converting the EC2 instance to a container and using AWS App Runner, which is a service that automatically builds and deploys web applications and load balances traffic2. This is not necessary for running batch jobs.
* C. Copy the tasks into AWS Lambda functions. Schedule the Lambda functions by using Amazon EventBridge (Amazon CloudWatch Events). This solution will not meet the requirement of performance, as AWS Lambda has a limit of 15 minutes for execution time and 10 GB for memory allocation3. These limits may not be sufficient for running 1-hour tasks.
* D. Create an Amazon Machine Image (AMI) of the EC2 instance that runs the tasks. Create an Auto Scaling group with the AMI to run multiple copies of the instance. This solution will not meet the requirement of low operational overhead, as it involves creating and maintaining AMIs and Auto Scaling groups, which are additional resources that need to be configured and managed2.
Reference URL: https://docs.aws.amazon.com/whitepapers/latest/aws-overview/compute- services.html

NEW QUESTION 210
- (Topic 3)
A media company hosts its website on AWS. The website application's architecture includes a fleet of Amazon EC2 instances behind an Application Load Balancer (ALB) and a database that is hosted on Amazon Aurora The company's cyber security teem reports that the application is vulnerable to SOL injection.
How should the company resolve this issue?

A. Use AWS WAF in front of the ALB Associate the appropriate web ACLs with AWS WAF.
B. Create an ALB listener rule to reply to SQL injection with a fixed response
C. Subscribe to AWS Shield Advanced to block all SQL injection attempts automatically.
D. Set up Amazon Inspector to block all SOL injection attempts automatically

Answer: A

Explanation:
https://aws.amazon.com/premiumsupport/knowledge-center/waf-block-common-
attacks/#:~:text=To%20protect%20your%20applications%20against,%2C%20query%20stri
ng%2C%20or%20URI. --
------------------------------ Protect against SQL injection and cross-site scripting To protect your applications against SQL injection and cross-site scripting (XSS)
attacks, use the built- in SQL injection and cross-site scripting engines. Remember that attacks can be performed on different parts of the HTTP request, such as
the HTTP header, query string, or URI. Configure the AWS WAF rules to inspect different parts of the HTTP request against the built-in mitigation engines.

NEW QUESTION 215
- (Topic 3)
A company is running a multi-tier recommerce 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 219
......

Thank You for Trying Our Product

We offer two products:

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

2nd - Questons and Answers in PDF Format

SAA-C03 Practice Exam Features:

* SAA-C03 Questions and Answers Updated Frequently

* SAA-C03 Practice Questions Verified by Expert Senior Certified Staff

* SAA-C03 Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* SAA-C03 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

100% Actual & Verified — Instant Download, Please Click
Order The SAA-C03 Practice Test Here