

Amazon-Web-Services

Exam Questions SAP-C02

AWS Certified Solutions Architect - Professional



NEW QUESTION 1

- (Exam Topic 1)

A company is serving files to its customers through an SFTP server that is accessible over the internet. The SFTP server is running on a single Amazon EC2 instance with an Elastic IP address attached. Customers connect to the SFTP server through its Elastic IP address and use SSH for authentication. The EC2 instance also has an attached security group that allows access from all customer IP addresses.

A solutions architect must implement a solution to improve availability, minimize the complexity of infrastructure management, and minimize the disruption to customers who access files. The solution must not change the way customers connect.

Which solution will meet these requirements?

- A. Disassociate the Elastic IP address from the EC2 instance. Create an Amazon S3 bucket to be used for SFTP file hosting. Create an AWS Transfer Family server. Configure the Transfer Family server with a publicly accessible endpoint.
- B. Associate the SFTP Elastic IP address with the new endpoint.
- C. Point the Transfer Family server to the S3 bucket. Sync all files from the SFTP server to the S3 bucket.
- D. Disassociate the Elastic IP address from the EC2 instance.
- E. Create an Amazon S3 bucket to be used for SFTP file hosting. Create an AWS Transfer Family server.
- F. Configure the Transfer Family server with a VPC-hosted internet-facing endpoint.
- G. internet-facing endpoint.
- H. Associate the SFTP Elastic IP address with the new endpoint.
- I. Attach the security group with customer IP addresses to the new endpoint.
- J. Point the Transfer Family server to the S3 bucket.
- K. Sync all files from the SFTP server to the S3 bucket.
- L. Disassociate the Elastic IP address from the EC2 instance.
- M. Create a new Amazon Elastic File System (Amazon EFS) file system to be used for SFTP file hosting.
- N. Create an AWS Fargate task definition to run an SFTP server.
- O. Specify the EFS file system as a mount in the task definition. Create a Fargate service by using the task definition, and place a Network Load Balancer (NLB) in front of the service. When configuring the service, attach the security group with customer IP addresses to the tasks that run the SFTP server. Associate the Elastic IP address with the NLB. Sync all files from the SFTP server to the S3 bucket.
- P. Disassociate the Elastic IP address from the EC2 instance. Create a multi-attach Amazon Elastic Block Store (Amazon EBS) volume to be used for SFTP file hosting. Create a Network Load Balancer (NLB) with the Elastic IP address attached. Create an Auto Scaling group with EC2 instances that run an SFTP server. Define in the Auto Scaling group that instances that are launched should attach the new multi-attach EBS volume. Configure the Auto Scaling group to automatically add instances behind the NLB. Configure the Auto Scaling group to use the security group that allows customer IP addresses for the EC2 instances that the Auto Scaling group launches. Sync all files from the SFTP server to the new multi-attach EBS volume.

Answer: B

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/aws-sftp-endpoint-type/>

<https://docs.aws.amazon.com/transfer/latest/userguide/create-server-in-vpc.html> <https://aws.amazon.com/premiumsupport/knowledge-center/aws-sftp-endpoint-type/>

NEW QUESTION 2

- (Exam Topic 1)

A company has a website that enables users to upload videos. Company policy states the uploaded videos must be analyzed for restricted content. An uploaded video is placed in Amazon S3, and a message is pushed to an Amazon SQS queue with the video's location. A backend application pulls this location from Amazon SQS and analyzes the video.

The video analysis is compute-intensive and occurs sporadically during the day. The website scales with demand. The video analysis application runs on a fixed number of instances. Peak demand occurs during the holidays, so the company must add instances to the application during this time. All instances used are currently on-demand Amazon EC2 T2 instances. The company wants to reduce the cost of the current solution.

Which of the following solutions is MOST cost-effective?

- A. Keep the website on T2 instance.
- B. Determine the minimum number of website instances required during off-peak times and use Spot Instances to cover them while using Reserved Instances to cover peak demand.
- C. Use Amazon EC2 R4 and Amazon EC2 R5 Reserved Instances in an Auto Scaling group for the video analysis application.
- D. Keep the website on T2 instance.
- E. Determine the minimum number of website instances required during off-peak times and use Reserved Instances to cover them while using On-Demand Instances to cover peak demand.
- F. Use Spot Fleet for the video analysis application comprised of Amazon EC2 C4 and Amazon EC2 C5 Spot Instances.
- G. Migrate the website to AWS Elastic Beanstalk and Amazon EC2 C4 instance.
- H. Determine the minimum number of website instances required during off-peak times and use On-Demand Instances to cover them while using Spot capacity to cover peak demand. Use Spot Fleet for the video analysis application comprised of C4 and Amazon EC2 C5 instances.
- I. Migrate the website to AWS Elastic Beanstalk and Amazon EC2 R4 instance.
- J. Determine the minimum number of website instances required during off-peak times and use Reserved Instances to cover them while using On-Demand Instances to cover peak demand. Use Spot Fleet for the video analysis application comprised of R4 and Amazon EC2 R5 instances.

Answer: B

NEW QUESTION 3

- (Exam Topic 1)

A company hosts a large on-premises MySQL database at its main office that supports an issue tracking system used by employees around the world. The company already uses AWS for some workloads and has created an Amazon Route 53 entry for the database endpoint that points to the on-premises database. Management is concerned about the database being a single point of failure and wants a solutions architect to migrate the database to AWS without any data loss or downtime.

Which set of actions should the solutions architect implement?

- A. Create an Amazon Aurora DB cluster.
- B. Use AWS Database Migration Service (AWS DMS) to do a full load from the on-premises database to Aurora.
- C. Update the Route 53 entry for the database to point to the Aurora cluster endpoint.
- D. and shut down the on-premises database.

- E. During nonbusiness hours, shut down the on-premises database and create a backup
- F. Restore this backup to an Amazon Aurora DB cluster
- G. When the restoration is complete, update the Route 53 entry for the database to point to the Aurora cluster endpoint, and shut down the on-premises database.
- H. Create an Amazon Aurora DB cluster
- I. Use AWS Database Migration Service (AWS DMS) to do a full load with continuous replication from the on-premises database to Aurora
- J. When the migration is complete, update the Route 53 entry for the database to point to the Aurora cluster endpoint, and shut down the on-premises database.
- K. Create a backup of the database and restore it to an Amazon Aurora multi-master cluster
- L. This Aurora cluster will be in a master-master replication configuration with the on-premises database
- M. Update the Route 53 entry for the database to point to the Aurora cluster endpoint
- N. and shut down the on-premises database.

Answer: C

Explanation:

“Around the world” eliminates possibility for the maintenance window at night. The other difference is ability to leverage continuous replication in MySQL to Aurora case.

NEW QUESTION 4

- (Exam Topic 1)

A company has a complex web application that leverages Amazon CloudFront for global scalability and performance. Over time, users report that the web application is slowing down.

The company's operations team reports that the CloudFront cache hit ratio has been dropping steadily. The cache metrics report indicates that query strings on some URLs are inconsistently ordered and are specified sometimes in mixed-case letters and sometimes in lowercase letters.

Which set of actions should the solutions architect take to increase the cache hit ratio as quickly as possible?

- A. Deploy a Lambda@Edge function to sort parameters by name and force them to be lowercase
- B. Select the CloudFront viewer request trigger to invoke the function.
- C. Update the CloudFront distribution to disable caching based on query string parameters.
- D. Deploy a reverse proxy after the load balancer to post-process the emitted URLs in the application to force the URL strings to be lowercase.
- E. Update the CloudFront distribution to specify casing-insensitive query string processing.

Answer: A

Explanation:

https://docs.amazonaws.cn/en_us/AmazonCloudFront/latest/DeveloperGuide/lambda-examples.html#lambda-ex Before CloudFront serves content from the cache it will trigger any Lambda function associated with the Viewer Request, in which we can normalize parameters.

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/lambda-examples.html#lambda-examp>

NEW QUESTION 5

- (Exam Topic 1)

A solutions architect is designing the data storage and retrieval architecture for a new application that a company will be launching soon. The application is designed to ingest millions of small records per minute from devices all around the world. Each record is less than 4 KB in size and needs to be stored in a durable location where it can be retrieved with low latency. The data is ephemeral and the company is required to store the data for 120 days only, after which the data can be deleted.

The solutions architect calculates that, during the course of a year, the storage requirements would be about 10-15 TB.

Which storage strategy is the MOST cost-effective and meets the design requirements?

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

Answer: B

Explanation:

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

NEW QUESTION 6

- (Exam Topic 1)

A company hosts a photography website on AWS that has global visitors. The website has experienced steady increases in traffic during the last 12 months, and users have reported a delay in displaying images. The company wants to configure Amazon CloudFront to deliver photos to visitors with minimal latency.

Which actions will achieve this goal? (Select TWO.)

- A. Set the Minimum TTL and Maximum TTL to 0 in the CloudFront distribution.
- B. Set the Minimum TTL and Maximum TTL to a high value in the CloudFront distribution.
- C. Set the CloudFront distribution to forward all headers, all cookies, and all query strings to the origin.
- D. Set up additional origin servers that are geographically closer to the requester
- E. Configure latency-based routing in Amazon Route 53.
- F. Select Price Class 100 on the CloudFront distribution.

Answer: BD

NEW QUESTION 7

- (Exam Topic 1)

A company wants to change its internal cloud billing strategy for each of its business units. Currently, the cloud governance team shares reports for overall cloud spending with the head of each business unit. The company uses AWS Organizations to manage the separate AWS accounts for each business unit. The existing tagging standard in Organizations includes the application, environment, and owner. The cloud governance team wants a centralized solution so each business unit receives monthly reports on its cloud spending. The solution should also send notifications for any cloud spending that exceeds a set threshold. Which solution is the MOST cost-effective way to meet these requirements?

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

Answer: B

Explanation:

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

NEW QUESTION 8

- (Exam Topic 1)

A company has an Amazon VPC that is divided into a public subnet and a private subnet. A web application runs in Amazon VPC, and each subnet has its own NACL. The public subnet has a CIDR of 10.0.0.0/24. An Application Load Balancer is deployed to the public subnet. The private subnet has a CIDR of 10.0.1.0/24. Amazon EC2 instances that run a web server on port 80 are launched into the private subnet. Only network traffic that is required for the Application Load Balancer to access the web application can be allowed to travel between the public and private subnets. What collection of rules should be written to ensure that the private subnet's NACL meets the requirement? (Select TWO.)

- A. An inbound rule for port 80 from source 0.0.0.0/0
- B. An inbound rule for port 80 from source 10.0.0.0/24
- C. An outbound rule for port 80 to destination 0.0.0.0/0
- D. An outbound rule for port 80 to destination 10.0.0.0/24
- E. An outbound rule for ports 1024 through 65535 to destination 10.0.0.0/24

Answer: BE

Explanation:

Ephemeral ports are not covered in the syllabus so be careful that you don't confuse day to day best practice with what is required for the exam. Link to an explanation on Ephemeral ports here: <https://acloud.guru/forums/aws-certified-solutions-architect-associate/discussion/-KUbcwo4IXefMI7janaK/netw>

NEW QUESTION 9

- (Exam Topic 1)

A solution architect is designing an AWS account structure for a company that consists of multiple terms. All the team will work in the same AWS Region. The company needs a VPC that is connected to the on-premises network. The company expects less than 50 Mbps of total to and from the on-premises network. Which combination of steps will meet these requirements MOST cost-effectively? (Select TWO)

- A. Create an AWS CloudFormation template that provisions a VPC and the required subnet
- B. Deploy the template to each AWS account
- C. Create an AWS CloudFormation template that provisions a VPC and the required subnet
- D. Deploy the template to a shared services account
- E. Share the subnets by using AWS Resource Access Manager
- F. Use AWS Transit Gateway along with an AWS Site-to-Site VPN for connectivity to the on-premises network
- G. Share the transit gateway by using AWS Resource Access Manager
- H. Use AWS Site-to-Site VPN for connectivity to the on-premises network
- I. Use AWS Direct Connect for connectivity to the on-premises network.

Answer: BD

NEW QUESTION 10

- (Exam Topic 1)

A finance company is running its business-critical application on current-generation Linux EC2 instances. The application includes a self-managed MySQL database performing heavy I/O operations. The application is working fine to handle a moderate amount of traffic during the month. However, it slows down during the final three days of each month due to month-end reporting, even though the company is using Elastic Load Balancers and Auto Scaling within its infrastructure to meet the increased demand.

Which of the following actions would allow the database to handle the month-end load with the LEAST impact on performance?

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

before the end of the month and reverting back afterwards.

Answer: B

Explanation:

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

NEW QUESTION 10

- (Exam Topic 1)

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

What is the MOST cost-effective migration recommendation?

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

Answer: D

Explanation:

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

NEW QUESTION 12

- (Exam Topic 1)

A web application is hosted in a dedicated VPC that is connected to a company's on-premises data center over a Site-to-Site VPN connection. The application is accessible from the company network only. This is a temporary non-production application that is used during business hours. The workload is generally low with occasional surges.

The application has an Amazon Aurora MySQL provisioned database cluster on the backend. The VPC has an internet gateway and a NAT gateway attached. The web servers are in private subnets in an Auto Scaling group behind an Elastic Load Balancer. The web servers also upload data to an Amazon S3 bucket through the internet.

A solutions architect needs to reduce operational costs and simplify the architecture. Which strategy should the solutions architect use?

- A. Review the Auto Scaling group settings and ensure the scheduled actions are specified to operate the Amazon EC2 instances during business hours only
- B. Use 3-year scheduled Reserved Instances for the web server EC2 instance
- C. Detach the internet gateway and remove the NAT gateways from the VPC
- D. Use an Aurora Serverless database and set up a VPC endpoint for the S3 bucket.
- E. Review the Auto Scaling group settings and ensure the scheduled actions are specified to operate the Amazon EC2 instances during business hours only
- F. Detach the internet gateway and remove the NAT gateways from the VPC
- G. Use an Aurora Serverless database and set up a VPC endpoint for the S3 bucket, then update the network routing and security rules and policies related to the changes.
- H. Review the Auto Scaling group settings and ensure the scheduled actions are specified to operate the Amazon EC2 instances during business hours only
- I. Detach the internet gateway from the VPC, and use an Aurora Serverless database
- J. Set up a VPC endpoint for the S3 bucket, then update the network routing and security rules and policies related to the changes.
- K. Use 3-year scheduled Reserved Instances for the web server Amazon EC2 instance
- L. Remove the NAT gateways from the VPC, and set up a VPC endpoint for the S3 bucket
- M. Use Amazon
- N. CloudWatch and AWS Lambda to stop and start the Aurora DB cluster so it operates during business hours only
- O. Update the network routing and security rules and policies related to the changes.

Answer: B

Explanation:

The application is accessible from the company network only remove NAT and IGW, application - S3 with VPC endpoint. Non-Production application no need to go for Reserved instances

To build site-to-site vpn, you don't need internet gateway. Instead, customer gateway is needed.

<https://docs.aws.amazon.com/vpn/latest/s2svpn/SetUpVPNConnections.html#vpn-create-cgw>

NEW QUESTION 15

- (Exam Topic 1)

A financial services company receives a regular data feed from its credit card servicing partner. Approximately 5.1 records are sent every 15 minutes in plaintext, delivered over HTTPS directly into an Amazon S3 bucket with server-side encryption. This feed contains sensitive credit card primary account number (PAN) data. The company needs to automatically mask the PAN before sending the data to another S3 bucket for additional internal processing. The company also needs to

remove and merge specific fields, and then transform the record into JSON format. Additionally, extra feeds are likely to be added in the future, so any design needs to be easily expandable.

Which solutions will meet these requirements?

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

Answer: C

Explanation:

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

Create an AWS Glue crawler and custom classifier based on the data feed formats and build a table definition to match. Trigger an AWS Lambda function on file delivery to start an AWS Glue ETL job to transform the entire record according to the processing and transformation requirements. Define the output format as JSON. Once complete, have the ETL job send the results to another S3 bucket for internal processing.

<https://docs.aws.amazon.com/glue/latest/dg/trigger-job.html>

https://d1.awsstatic.com/Products/product-name/diagrams/product-page-diagram_Glue_Event-driven-ETL-Pipeline.png

NEW QUESTION 17

- (Exam Topic 1)

A media company uses Amazon DynamoDB to store metadata for its catalog of movies that are available to stream. Each media item contains user-facing content that concludes a description of the media, a list of search tags, and similar data. In addition, media items include a list of Amazon S3 key names that relate to movie files. The company stores these movie files in a single S3 bucket that has versioning enabled. The company uses Amazon CloudFront to serve these movie files.

The company has 100,000 media items, and each media item can have many different S3 objects that represent different encodings of the same media. S3 objects that belong to the same media item are grouped together under the same key prefix, which is a random unique ID.

Because of an expiring contract with a media provider, the company must remove 2,000 media items. The company must completely delete all DynamoDB keys and movie files on Amazon S3 that are related to these media items within 36 hours. The company must ensure that the content cannot be recovered.

Which combination of actions will meet these requirements? (Select TWO.)

- A. Configure the DynamoDB table with a TTL field
- B. Create and invoke an AWS Lambda function to perform a conditional update. Set the TTL field to the time of the contract's expiration on every affected media item.
- C. Configure an S3 Lifecycle object expiration rule that is based on the contract's expiration date
- D. Write a script to perform a conditional delete on all the affected DynamoDB records
- E. Temporarily suspend versioning on the S3 bucket
- F. Create and invoke an AWS Lambda function that deletes affected objects. Reactivate versioning when the operation is complete
- G. Write a script to delete objects from Amazon S3. Specify in each request a NoncurrentVersionExpiration property with a NoncurrentDays attribute set to 0.

Answer: CE

NEW QUESTION 20

- (Exam Topic 1)

A solutions architect is responsible for redesigning a legacy Java application to improve its availability, data durability, and scalability. Currently, the application runs on a single high-memory Amazon EC2 instance. It accepts HTTP requests from upstream clients, adds them to an in-memory queue, and responds with a 200 status. A separate application thread reads items from the queue, processes them, and persists the results to an Amazon RDS MySQL instance. The processing time for each item takes 90 seconds on average, most of which is spent waiting on external service calls, but the application is written to process multiple items in parallel.

Traffic to this service is unpredictable. During periods of high load, items may sit in the internal queue for over an hour while the application processes the backlog. In addition, the current system has issues with availability and data loss if the single application node fails.

Clients that access this service cannot be modified. They expect to receive a response to each HTTP request they send within 10 seconds before they will time out and retry the request.

Which approach would improve the availability and durability of the system while decreasing the processing latency and minimizing costs?

- A. Create an Amazon API Gateway REST API that uses Lambda proxy integration to pass requests to an AWS Lambda function
- B. Migrate the core processing code to a Lambda function and write a wrapper class that provides a handler method that converts the proxy events to the internal application data model and invokes the processing module.
- C. Create an Amazon API Gateway REST API that uses a service proxy to put items in an Amazon SQS queue
- D. Extract the core processing code from the existing application and update it to pull items from Amazon SQS instead of an in-memory queue
- E. Deploy the new processing application to smaller EC2 instances within an Auto Scaling group that scales dynamically based on the approximate number of messages in the Amazon SQS queue.
- F. Modify the application to use Amazon DynamoDB instead of Amazon RDS
- G. Configure Auto Scaling for the DynamoDB table
- H. Deploy the application within an Auto Scaling group with a scaling policy based on CPU utilization

- I. Back the in-memory queue with a memory-mapped file to an instance store volume and periodically write that file to Amazon S3.
- J. Update the application to use a Redis task queue instead of the in-memory queue
- K. Build a Docker container image for the application
- L. Create an Amazon ECS task definition that includes the application container and a separate container to host Redis
- M. Deploy the new task definition as an ECS service using AWS Fargate, and enable Auto Scaling.

Answer: B

Explanation:

The obvious challenges here are long workloads, scalability based on queue load, and reliability. Almost always the defacto answer to queue related workload is SQS. Since the workloads are very long (90 minutes) Lambdas cannot be used (15 mins max timeout). So, autoscaled smaller EC2 nodes that wait on external services to complete the task makes more sense. If the task fails, the message is returned to the queue and retried.

NEW QUESTION 22

- (Exam Topic 1)

A company provides a centralized Amazon EC2 application hosted in a single shared VPC. The centralized application must be accessible from client applications running in the VPCs of other business units. The centralized application front end is configured with a Network Load Balancer (NLB) for scalability.

Up to 10 business unit VPCs will need to be connected to the shared VPC. Some of the business unit VPC CIDR blocks overlap with the shared VPC, and some overlap with each other. Network connectivity to the centralized application in the shared VPC should be allowed from authorized business unit VPCs only.

Which network configuration should a solutions architect use to provide connectivity from the client applications in the business unit VPCs to the centralized application in the shared VPC?

- A. Create an AWS Transit Gateway
- B. Attach the shared VPC and the authorized business unit VPCs to the transit gateway
- C. Create a single transit gateway route table and associate it with all of the attached VPC
- D. Allow automatic propagation of routes from the attachments into the route table
- E. Configure VPC routing tables to send traffic to the transit gateway.
- F. Create a VPC endpoint service using the centralized application NLB and enable the option to require endpoint acceptance
- G. Create a VPC endpoint in each of the business unit VPCs using the service name of the endpoint service
- H. Accept authorized endpoint requests from the endpoint service console.
- I. Create a VPC peering connection from each business unit VPC to the shared VPC
- J. Accept the VPC peering connections from the shared VPC console
- K. Configure VPC routing tables to send traffic to the VPC peering connection.
- L. Configure a virtual private gateway for the shared VPC and create customer gateways for each of the authorized business unit VPC
- M. Establish a Site-to-Site VPN connection from the business unit VPCs to the shared VPC
- N. Configure VPC routing tables to send traffic to the VPN connection.

Answer: B

Explanation:

Amazon Transit Gateway doesn't support routing between Amazon VPCs with overlapping CIDRs. If you attach a new Amazon VPC that has a CIDR which overlaps with an already attached Amazon VPC, Amazon Transit Gateway will not propagate the new Amazon VPC route into the Amazon Transit Gateway route table.

<https://docs.aws.amazon.com/elasticloadbalancing/latest/network/load-balancer-target-groups.html#client-ip-pre>

NEW QUESTION 25

- (Exam Topic 1)

A company needs to create and manage multiple AWS accounts for a number of departments from a central location. The security team requires read-only access to all accounts from its own AWS account. The company is using AWS Organizations and created an account for the security team.

How should a solutions architect meet these requirements?

- A. Use the OrganizationAccountAccessRole IAM role to create a new IAM policy with read-only access in each member account
- B. Establish a trust relationship between the IAM policy in each member account and the security account
- C. Ask the security team to use the IAM policy to gain access.
- D. Use the OrganizationAccountAccessRole IAM role to create a new IAM role with read-only access in each member account
- E. Establish a trust relationship between the IAM role in each member account and the security account
- F. Ask the security team to use the IAM role to gain access.
- G. Ask the security team to use AWS Security Token Service (AWS STS) to call the AssumeRole API for the OrganizationAccountAccessRole IAM role in the master account from the security account
- H. Use the generated temporary credentials to gain access.
- I. Ask the security team to use AWS Security Token Service (AWS STS) to call the AssumeRole API for the OrganizationAccountAccessRole IAM role in the member account from the security account
- J. Use the generated temporary credentials to gain access.

Answer: D

NEW QUESTION 30

- (Exam Topic 1)

A company is running an application distributed over several Amazon EC2 instances in an Auto Scaling group behind an Application Load Balancer. The security team requires that all application access attempts be made available for analysis. Information about the client IP address, connection type, and user agent must be included.

Which solution will meet these requirements?

- A. Enable EC2 detailed monitoring, and include network log
- B. Send all logs through Amazon Kinesis Data Firehose to an Amazon Elasticsearch Service (Amazon ES) cluster that the security team uses for analysis.
- C. Enable VPC Flow Logs for all EC2 instance network interfaces. Publish VPC Flow Logs to an Amazon S3 bucket. Have the security team use Amazon Athena to query and analyze the logs.
- D. Enable access logs for the Application Load Balancer, and publish the logs to an Amazon S3 bucket. Have the security team use Amazon Athena to query and analyze the logs.
- E. Enable Traffic Mirroring and specify all EC2 instance network interfaces as the source

F. Send all traffic information through Amazon Kinesis Data Firehose to an Amazon Elasticsearch Service (Amazon ES) cluster that the security team uses for analysis.

Answer: C

Explanation:

<https://docs.aws.amazon.com/elasticloadbalancing/latest/application/load-balancer-access-logs.html> <https://docs.aws.amazon.com/vpc/latest/mirroring/what-is-traffic-mirroring.html>

NEW QUESTION 31

- (Exam Topic 1)

A company maintains a restaurant review website. The website is a single-page application where files are stored in Amazon S3 and delivered using Amazon CloudFront. The company receives several fake postings every day that are manually removed.

The security team has identified that most of the fake posts are from bots with IP addresses that have a bad reputation within the same global region. The team needs to create a solution to help restrict the bots from accessing the website.

Which strategy should a solutions architect use?

- A. Use AWS Firewall Manager to control the CloudFront distribution security setting
- B. Create a geographical block rule and associate it with Firewall Manager.
- C. Associate an AWS WAF web ACL with the CloudFront distributio
- D. Select the managed Amazon IP reputation rule group for the web ACL with a deny action.
- E. Use AWS Firewall Manager to control the CloudFront distribution security setting
- F. Select the managed Amazon IP reputation rule group and associate it with Firewall Manager with a deny action.
- G. Associate an AWS WAF web ACL with the CloudFront distributio
- H. Create a rule group for the web ACL with a geographical match statement with a deny action.

Answer: B

Explanation:

IP reputation rule groups allow you to block requests based on their source. Choose one or more of these rule groups if you want to reduce your exposure to BOTS!!!! traffic or exploitation attempts

The Amazon IP reputation list rule group contains rules that are based on Amazon internal threat intelligence. This is useful if you would like to block IP addresses typically associated with bots or other threats. Inspects for a list of IP addresses that have been identified as bots by Amazon threat intelligence.

NEW QUESTION 32

- (Exam Topic 1)

A financial company is building a system to generate monthly, immutable bank account statements for its users. Statements are stored in Amazon S3. Users should have immediate access to their monthly statements for up to 2 years. Some users access their statements frequently, whereas others rarely access their statements. The company's security and compliance policy requires that the statements be retained for at least 7 years.

What is the MOST cost-effective solution to meet the company's needs?

- A. Create an S3 bucket with Object Lock disable
- B. Store statements in S3 Standar
- C. Define an S3 Lifecycle policy to transition the data to S3 Standard-Infrequent Access (S3 Standard-IA) after 30 day
- D. Define another S3 Lifecycle policy to move the data to S3 Glacier Deep Archive after 2 year
- E. Attach an S3 Glacier Vault Lock policy with deny delete permissions for archives less than 7 years old.
- F. Create an S3 bucket with versioning enable
- G. Store statements in S3 Intelligent-Tierin
- H. Usesame-Region replication to replicate objects to a backup S3 bucke
- I. Define an S3 Lifecycle policy for the backup S3 bucket to move the data to S3 Glacie
- J. Attach an S3 Glacier Vault Lock policy with deny delete permissions for archives less than 7 years old.
- K. Create an S3 bucket with Object Lock enable
- L. Store statements in S3 Intelligent-Tierin
- M. Enable compliance mode with a default retention period of 2 year
- N. Define an S3 Lifecycle policy to move the data to S3 Glacier after 2 year
- O. Attach an S3 Glacier Vault Lock policy with deny delete permissionsfor archives less than 7 years old.
- P. Create an S3 bucket with versioning disable
- Q. Store statements in S3 One Zone-Infrequent Access (S3 One Zone-IA). Define an S3 Lifecyde policy to move the data to S3 Glacier Deep Archive after 2 year
- R. Attach an S3 Glader Vault Lock policy with deny delete permissions for archives less than 7 years old.

Answer: C

Explanation:

<https://aws.amazon.com/about-aws/whats-new/2018/11/s3-object-lock/>

Create an S3 bucket with Object Lock enabled. Store statements in S3 Intelligent-Tiering. Enable compliance mode with a default retention period of 2 years.

Define an S3 Lifecycle policy to move the data to S3 Glacier after 2 years. Attach an S3 Glacier Vault Lock policy with deny delete permissions for archives less than 7 years old.

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

NEW QUESTION 34

- (Exam Topic 1)

A company is planning to set up a REST API application on AWS. The application team wants to set up a new identity store on AWS The IT team does not want to maintain any infrastructure or servers for this deployment.

What is the MOST operationally efficient solution that meets these requirements?

- A. Deploy the application as AWS Lambda function
- B. Set up Amazon API Gateway REST API endpoints for the application Create a Lambda function, and configure a Lambda authorizer
- C. Deploy the application in AWS AppSync, and configure AWS Lambda resolvers Set up an Amazon Cognito user pool, and configure AWS AppSync to use the user pool for authorization
- D. Deploy the application as AWS Lambda function

- E. Set up Amazon API Gateway REST API endpoints for the application Set up an Amazon Cognito user pool, and configure an Amazon Cognito authorizer
- F. Deploy the application in Amazon Elastic Kubernetes Service (Amazon EKS) cluster
- G. Set up an Application Load Balancer for the EKS pods Set up an Amazon Cognito user pool and service pod for authentication.

Answer: C

NEW QUESTION 35

- (Exam Topic 1)

A company hosts a web application that runs on a group of Amazon EC2 instances that are behind an Application Load Balancer (ALB) in a VPC. The company wants to analyze the network payloads to reverse-engineer a sophisticated attack of the application.

Which approach should the company take to achieve this goal?

- A. Enable VPC Flow Log
- B. Store the flow logs in an Amazon S3 bucket for analysis.
- C. Enable Traffic Mirroring on the network interface of the EC2 instance
- D. Send the mirrored traffic to a target for storage and analysis.
- E. Create an AWS WAF web ACL
- F. and associate it with the ALB
- G. Configure AWS WAF logging.
- H. Enable logging for the ALB
- I. Store the logs in an Amazon S3 bucket for analysis.

Answer: A

NEW QUESTION 36

- (Exam Topic 1)

A multimedia company needs to deliver its video-on-demand (VOD) content to its subscribers in a cost-effective way. The video files range in size from 1-15 GB and are typically viewed frequently for the first 6 months after creation, and then access decreases considerably. The company requires all video files to remain immediately available for subscribers. There are now roughly 30,000 files, and the company anticipates doubling that number over time.

What is the MOST cost-effective solution for delivering the company's VOD content?

- A. Store the video files in an Amazon S3 bucket using S3 Intelligent-Tiering
- B. Use Amazon CloudFront to deliver the content with the S3 bucket as the origin.
- C. Use AWS Elemental MediaConvert and store the adaptive bitrate video files in Amazon S3. Configure an AWS Elemental MediaPackage endpoint to deliver the content from Amazon S3.
- D. Store the video files in Amazon Elastic File System (Amazon EFS) Standard
- E. Enable EFS lifecycle management to move the video files to EFS Infrequent Access after 6 months
- F. Create an Amazon EC2 Auto Scaling group behind an Elastic Load Balancer to deliver the content from Amazon EFS.
- G. Store the video files in Amazon S3 Standard
- H. Create S3 Lifecycle rules to move the video files to S3 Standard-Infrequent Access (S3 Standard-IA) after 6 months and to S3 Glacier Deep Archive after 1 year
- I. Use Amazon CloudFront to deliver the content with the S3 bucket as the origin.

Answer: A

Explanation:

<https://d1.awsstatic.com/whitepapers/amazon-cloudfront-for-media.pdf> <https://aws.amazon.com/solutions/implementations/video-on-demand-on-aws/>

NEW QUESTION 37

- (Exam Topic 1)

A company uses AWS Transit Gateway for a hub-and-spoke model to manage network traffic between many VPCs. The company is developing a new service that must be able to send data at 100 Gbps. The company needs a faster connection to other VPCs in the same AWS Region.

Which solution will meet these requirements?

- A. Establish VPC peering between the necessary VPCs
- B. Ensure that all route tables are updated as required.
- C. Attach an additional transit gateway to the VPC
- D. Update the route tables accordingly.
- E. Create AWS Site-to-Site VPN connections that use equal-cost multi-path (ECMP) routing between the necessary VPCs.
- F. Create an additional attachment from the necessary VPCs to the existing transit gateway.

Answer: D

NEW QUESTION 39

- (Exam Topic 1)

An ecommerce website running on AWS uses an Amazon RDS for MySQL DB instance with General Purpose SSD storage. The developers chose an appropriate instance type based on demand, and configured 100 GB of storage with a sufficient amount of free space.

The website was running smoothly for a few weeks until a marketing campaign launched. On the second day of the campaign, users reported long wait times and time outs. Amazon CloudWatch metrics indicated that both reads and writes to the DB instance were experiencing long response times. The CloudWatch metrics show 40% to 50% CPU and memory utilization, and sufficient free storage space is still available. The application server logs show no evidence of database connectivity issues.

What could be the root cause of the issue with the marketing campaign?

- A. It exhausted the I/O credit balance due to provisioning low disk storage during the setup phase.
- B. It caused the data in the tables to change frequently, requiring indexes to be rebuilt to optimize queries.
- C. It exhausted the maximum number of allowed connections to the database instance.
- D. It exhausted the network bandwidth available to the RDS for MySQL DB instance.

Answer: A

Explanation:

"When using General Purpose SSD storage, your DB instance receives an initial I/O credit balance of 5.4 million I/O credits. This initial credit balance is enough to sustain a burst performance of 3,000 IOPS for 30 minutes."

<https://aws.amazon.com/blogs/database/how-to-use-cloudwatch-metrics-to-decide-between-general-purpose-or>

NEW QUESTION 41

- (Exam Topic 1)

A company is hosting a single-page web application in the AWS Cloud. The company is using Amazon CloudFront to reach its goal audience. The CloudFront distribution has an Amazon S3 bucket that is configured as its origin. The static files for the web application are stored in this S3 bucket.

The company has used a simple routing policy to configure an Amazon Route 53 A record. The record points to the CloudFront distribution. The company wants to use a canary deployment release strategy for new versions of the application.

What should a solutions architect recommend to meet these requirements?

- A. Create a second CloudFront distribution for the new version of the application.
- B. Update the Route 53 record to use a weighted routing policy.
- C. Create a Lambda@Edge function.
- D. Configure the function to implement a weighting algorithm and rewrite the URL to direct users to a new version of the application.
- E. Create a second S3 bucket and a second CloudFront origin for the new S3 bucket. Create a CloudFront origin group that contains both origins. Configure origin weighting for the origin group.
- F. Create two Lambda@Edge functions.
- G. Use each function to serve one of the application versions. Set up a CloudFront weighted Lambda@Edge invocation policy.

Answer: A

NEW QUESTION 42

- (Exam Topic 2)

A company that develops consumer electronics with offices in Europe and Asia has 60 TB of software images stored on premises in Europe. The company wants to transfer the images to an Amazon S3 bucket in the ap-northeast-1 Region. New software images are created daily and must be encrypted in transit. The company needs a solution that does not require custom development to automatically transfer all existing and new software images to Amazon S3.

What is the next step in the transfer process?

- A. Deploy an AWS DataSync agent and configure a task to transfer the images to the S3 bucket.
- B. Configure Amazon Kinesis Data Firehose to transfer the images using S3 Transfer Acceleration.
- C. Use an AWS Snowball device to transfer the images with the S3 bucket as the target.
- D. Transfer the images over a Site-to-Site VPN connection using the S3 API with multipart upload.

Answer: A

NEW QUESTION 43

- (Exam Topic 2)

A company's solution architect is designing a disaster recovery (DR) solution for an application that runs on AWS. The application uses PostgreSQL 11.7 as its database. The company has an RPO of 30 seconds. The solutions architect must design a DR solution with the primary database in the us-east-1 Region and the database in the us-west-2 Region.

What should the solution architect do to meet these requirements with minimum application change?

- A. Migrate the database to Amazon RDS for PostgreSQL in us-east-1. Set up a read replica in us-west-2. Set the managed RPO for the RDS database to 30 seconds.
- B. Migrate the database to Amazon RDS for PostgreSQL in us-east-1. Set up a standby replica in an Availability Zone in us-west-2. Set the managed RPO for the RDS database to 30 seconds.
- C. Migrate the database to an Amazon Aurora PostgreSQL global database with the primary Region as us-east-1 and the secondary Region as us-west-2. Set the managed RPO for the Aurora database to 30 seconds.
- D. Migrate the database to Amazon DynamoDB in us-east-1. Set up global tables with replica tables that are created in us-west-2.

Answer: A

NEW QUESTION 47

- (Exam Topic 2)

A company is running a two-tier web-based application in an on-premises data center. The application layer consists of a single server running a stateful application. The application connects to a PostgreSQL database running on a separate server. The application's user base is expected to grow significantly, so the company is migrating the application and database to AWS. The solution will use Amazon Aurora PostgreSQL, Amazon EC2 Auto Scaling, and Elastic Load Balancing.

Which solution will provide a consistent user experience that will allow the application and database tiers to scale?

- A. Enable Aurora Auto Scaling for Aurora Replica.
- B. Use a Network Load Balancer with the least outstanding requests routing algorithm and sticky sessions enabled.
- C. Enable Aurora Auto Scaling for Aurora writer.
- D. Use an Application Load Balancer with the round robin routing algorithm and sticky sessions enabled.
- E. Aurora Auto Scaling for Aurora Replica.
- F. Use an Application Load Balancer with the round robin routing algorithm and sticky sessions enabled.
- G. Aurora Auto Scaling for Aurora writer.
- H. Use a Network Load Balancer with the least outstanding requests routing algorithm and sticky sessions enabled.

Answer: C

NEW QUESTION 49

- (Exam Topic 2)

A company is using an Amazon CloudFront distribution to distribute both static and dynamic content from a web application running behind an Application Load Balancer. The web application requires user authorization and session tracking for dynamic content. The CloudFront distribution has a single cache behavior.

configured to forward the Authorization, Host, and Agent HTTP allow list headers and a session cookie to the origin All other cache behavior settings are set to their default value

A valid ACM certificate is applied to the CloudFront distribution with a matching CNAME in the distribution settings The ACM certificate is also applied to the HTTPS listener for the Application Load Balancer The CloudFront origin protocol policy is set to HTTPS only Analysis of the cache statistics report shows that the miss rate for this distribution is very high

What can the solutions architect do to improve the cache hit rate for this distribution without causing the SSL/TLS handshake between CloudFront and the Application Load Balancer to fail?

- A. Create two cache behaviors for static and dynamic content Remove the user-Agent and Host HTTP headers from the allow list headers section on both of the cache behaviors Remove the session cookie from the allow list cookies section and the Authorization HTTP header from the allow list headers section for cache behavior configured for static content
- B. Remove the user-Agent and Authorization HTTP headers from the allow list headers section of the cache behaviour
- C. Then update the cache behaviour to use resigned cookies for authorization
- D. Remove the Host HTTP header from the allow list headers section and remove the session cookie from the allow list cookies section for the default cache behaviour Enable automatic object compression and use Lambda@Edge viewer request events for user authorization
- E. Create two cache behaviours for static and dynamic content Remove the User-Agent HTTP header from the allow list headers section on both of the cache behaviours Remove the session cookie from the allow list cookies section and the Authorization HTTP header from the allow list headers section for cache behaviour configured for static content

Answer: D

Explanation:

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/understanding-the-cache-key.html> Removing the host header will result in failed flow between CloudFront and ALB, because they have same certificate.

NEW QUESTION 53

- (Exam Topic 2)

A company wants to migrate its workloads from on premises to AWS. The workloads run on Linux and Windows. The company has a large on-premises intra structure that consists of physical machines and VMs that host numerous applications.

The company must capture details about the system configuration. system performance. running processure and network coi.net lions of its o. -premises ,on boards. The company also must divide the on-premises applications into groups for AWS migrations. The company needs recommendations for Amazon EC2 instance types so that the company can run its workloads on AWS in the most cost-effective manner.

Which combination of steps should a solutions architect take to meet these requirements? (Select THREE.)

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

Answer: BDF

NEW QUESTION 54

- (Exam Topic 2)

A company uses AWS Organizations with a single OU named Production to manage multiple accounts All accounts are members of the Production OU Administrators use deny list SCPs in the root of the organization to manage access to restricted services.

The company recently acquired a new business unit and invited the new unit's existing AWS account to the organization Once onboarded the administrators of the new business unit discovered that they are not able to update existing AWS Config rules to meet the company's policies.

Which option will allow administrators to make changes and continue to enforce the current policies without introducing additional long-term maintenance?

- A. Remove the organization's root SCPs that limit access to AWS Config Create AWS Service Catalog products for the company's standard AWS Config rules and deploy them throughout the organization, including the new account.
- B. Create a temporary OU named Onboarding for the new account Apply an SCP to the Onboarding OU to allow AWS Config actions Move the new account to the Production OU when adjustments to AWS Config are complete
- C. Convert the organization's root SCPs from deny list SCPs to allow list SCPs to allow the required services only Temporarily apply an SCP to the organization's root that allows AWS Config actions for principals only in the new account.
- D. Create a temporary OU named Onboarding for the new account Apply an SCP to the Onboarding OU to allow AWS Config action
- E. Move the organization's root SCP to the Production O
- F. Move the new account to the Production OU when adjustments to AWS Config are complete.

Answer: D

NEW QUESTION 58

- (Exam Topic 2)

A Solutions Architect is constructing a containerized.NET Core application for AWS Fargate. The application's backend needs a high-availability version of Microsoft SQL Server. All application levels must be extremely accessible. The credentials associated with the SQL Server connection string should not be saved to disk inside the.NET Core front-end containers.

Which tactics should the Solutions Architect use to achieve these objectives?

- A. Set up SQL Server to run in Fargate with Service Auto Scalin
- B. Create an Amazon ECS task execution role that allows the Fargate task definition to get the secret value for the credentials to SQL Server running in Fargat
- C. Specify the ARN of the secret in AWS Secrets Manager in the secrets section of the Fargate task definition so the sensitive data can be injected into the containers as environment variables on startup for reading into the application to construct the connection strin
- D. Set up the .NET Core service using Service Auto Scaling behind an Application Load Balancer in multiple Availability Zones.
- E. Create a Multi-AZ deployment of SQL Server on Amazon RD
- F. Create a secret in AWS Secrets Manager for the credentials to the RDS databas
- G. Create an Amazon ECS task execution role that allows the Fargate task definition to get the secret value for the credentials to the RDS database in Secrets

Manage

- H. Specify the ARN of the secret in Secrets Manager in the secrets section of the Fargate task definition so the sensitive data can be injected into the containers as environment variables on startup for reading into the application to construct the connection string
- I. Set up the .NET Core service in Fargate using Service Auto Scaling behind an Application Load Balancer in multiple Availability Zones.
- J. Create an Auto Scaling group to run SQL Server on Amazon EC2. Create a secret in AWS Secrets Manager for the credentials to SQL Server running on EC2. Create an Amazon ECS task execution role that allows the Fargate task definition to get the secret value for the credentials to SQL Server on EC2. Specify the ARN of the secret in Secrets Manager in the secrets section of the Fargate task definition so the sensitive data can be injected into the containers as environment variables on startup for reading into the application to construct the connection string
- K. Set up the .NET Core service using Service Auto Scaling behind an Application Load Balancer in multiple Availability Zones.
- L. Create a Multi-AZ deployment of SQL Server on Amazon RD
- M. Create a secret in AWS Secrets Manager for the credentials to the RDS database
- N. Create non-persistent empty storage for the .NET Core containers in the Fargate task definition to store the sensitive information
- O. Create an Amazon ECS task execution role that allows the Fargate task definition to get the secret value for the credentials to the RDS database in Secrets Manager
- P. Specify the ARN of the secret in Secrets Manager in the secrets section of the Fargate task definition so the sensitive data can be written to the non-persistent empty storage on startup for reading into the application to construct the connection string
- Q. Set up the .NET Core service using Service Auto Scaling behind an Application Load Balancer in multiple Availability Zones.

Answer: B

Explanation:

Secrets Manager natively supports SQL Server on RDS. No real need to create additional 'ephemeral storage' to fetch credentials, as these can be injected to containers as environment variables. <https://aws.amazon.com/premiumsupport/knowledge-center/ecs-data-security-container-task/>

NEW QUESTION 62

- (Exam Topic 2)

A solutions architect uses AWS Organizations to manage several AWS accounts for a company. The full Organizations feature set is activated for the organization. All production AWS accounts exist under an OU that is named "production". Systems operators have full administrative privileges within these accounts by using IAM roles.

The company wants to ensure that security groups in all production accounts do not allow inbound traffic for TCP port 22. All noncompliant security groups must be remediated immediately, and no new rules that allow port 22 can be created.

Which solution will meet these requirements?

- A. Write an SCP that denies the CreateSecurityGroup action with a condition of (ec2:ingress rule with value 22. Apply the SCP to the 'production' OU.
- B. Configure an AWS CloudTrail trail for all accounts. Send CloudTrail logs to an Amazon S3 bucket. In the Organizations management account, configure an AWS Lambda function on the management account with permissions to assume a role in all production accounts to describe and modify security groups.
- C. Configure an AWS Lambda function on the management account with permissions to assume a role in all production accounts to describe and modify security groups. Configure the Lambda function to analyze each CloudTrail event for noncompliant security group actions and to automatically remediate any issues.
- D. Configure Amazon S3 to invoke the Lambda function on every PutObject event on the S3 bucket. Configure the Lambda function to analyze each CloudTrail event for noncompliant security group actions and to automatically remediate any issues.
- E. Create an Amazon EventBridge (Amazon CloudWatch Events) event bus in the Organizations management account.
- F. Create an AWS CloudFormation template to deploy configurations that send CreateSecurityGroup events to the event bus from all production accounts. Configure an AWS Lambda function in the management account with permissions to assume a role in all production accounts to describe and modify security groups. Configure the event bus to invoke the Lambda function. Configure the Lambda function to analyze each event for noncompliant security group actions and to automatically remediate any issues.
- G. Create an AWS CloudFormation template to turn on AWS Config. Activate the INCOMING_SSH_DISABLED AWS Config managed rule. Deploy an AWS Lambda function that will run based on AWS Config findings and will remediate noncompliant resources. Deploy the CloudFormation template by using a StackSet that is assigned to the "production" OU.
- H. Apply an SCP to the OU to deny modification of the resources that the CloudFormation template provisions.

Answer: D

NEW QUESTION 66

- (Exam Topic 2)

A company is migrating its data center from on-premises to the AWS Cloud. The migration will take several months to complete. The company will use Amazon Route 53 for private DNS zones.

During the migration, the company must keep its AWS services pointed at the VPC's Route 53 Resolver for DNS. The company also must maintain the ability to resolve addresses from its on-premises DNS server. A solutions architect must set up DNS so that Amazon EC2 instances can use native Route 53 endpoints to resolve on-premises DNS queries.

Which configuration will meet these requirements?

- A. Configure the VPC DHCP options set to point to on-premises DNS server IP addresses.
- B. Ensure that security groups for EC2 instances allow outbound access to port 53 on those DNS server IP addresses.
- C. Launch an EC2 instance that has DNS BIND installed and configure it.
- D. Ensure that the security groups that are attached to the EC2 instance can access the on-premises DNS server IP address on port 53. Configure BIND to forward DNS queries to on-premises DNS server IP addresses. Configure each migrated EC2 instance's DNS settings to point to the BIND server IP address.
- E. Create a new outbound endpoint in Route 53 and attach the endpoint to the VPC.
- F. Ensure that the security groups that are attached to the endpoint can access the on-premises DNS server IP address on port 53. Create a new Route 53 Resolver rule that routes on-premises designated traffic to the on-premises DNS server.
- G. Create a new private DNS zone in Route 53 with the same domain name as the on-premises domain. Create a single wildcard record with the on-premises DNS server IP address as the record's address.

Answer: A

NEW QUESTION 69

- (Exam Topic 2)

A software development company has multiple engineers who are working remotely. The company is running Active Directory Domain Services (AD DS) on an Amazon EC2 instance. The company's security policy states that all internal, nonpublic services that are deployed in a VPC must be accessible through a VPN. Multi-factor authentication (MFA) must be used for access to a VPN.

What should a solution architect do to meet these requirements?

- A. Create an AWS Site-to-Site VPN connection Configure integration between a VPN and AD D
- B. Use an Amazon Workspaces client with MFA support enabled to establish a VPN connection.
- C. Create an AWS Client VPN endpoint Create an AD Connector directory for integration with AD DS Enable MFA for AD Connector Use AWS Client VPN to establish a VPN connection.
- D. Create multiple AWS Site-to-Site VPN connections by using AWS VPN CloudHub Configure integration between AWS VPN CloudHub and AD DS Use AWS Cop4ot to establish a VPN connection.
- E. Create an Amazon WorkLink endpoint Configure integration between Amazon WorkLink and AD D
- F. Enable MFA in Amazon WorkLink Use AWS Client VPN to establish a VPN connection.

Answer: B

NEW QUESTION 72

- (Exam Topic 2)

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

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

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

Answer: ABE

NEW QUESTION 75

- (Exam Topic 2)

A digital marketing company has multiple AWS accounts that belong to various teams. The creative team uses an Amazon S3 bucket in its AWS account to securely store images and media files that are used as content for the company's marketing campaigns. The creative team wants to share the S3 bucket with the strategy team so that the strategy team can view the objects.

A solutions architect has created an IAM role that is named strategy_reviewer in the Strategy account. The solutions architect also has set up a custom AWS Key Management Service (AWS KMS) key in the Creative account and has associated the key with the S3 bucket. However, when users from the Strategy account assume the IAM role and try to access objects in the S3 bucket, they receive an Account.

The solutions architect must ensure that users in the Strategy account can access the S3 bucket. The solution must provide these users with only the minimum permissions that they need.

Which combination of steps should the solutions architect take to meet these requirements? (Select THREE.)

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

Answer: ACE

NEW QUESTION 80

- (Exam Topic 2)

A gaming company created a game leaderboard by using a Multi-AZ deployment of an Amazon RDS database. The number of users is growing, and the queries to get individual player rankings are getting slower over time. The company expects a surge in users for an upcoming version and wants to optimize the design for scalability and performance.

Which solution will meet these requirements?

- A. Migrate the database to Amazon DynamoDB
- B. Store the leader different table
- C. Use Apache HiveQL JOIN statements to build the leaderboard
- D. Keep the leaderboard data in the RDS DB instance
- E. Provision a Multi-AZ deployment of an Amazon ElastiCache for Redis cluster.
- F. Stream the leaderboard data by using Amazon Kinesis Data Firehose with an Amazon S3 bucket as the destination
- G. Query the S3 bucket by using Amazon Athena for the leaderboard.
- H. Add a read-only replica to the RDS DB instance
- I. Add an RDS Proxy database proxy.

Answer: C

NEW QUESTION 82

- (Exam Topic 2)

A company is creating a sequel for a popular online game. A large number of users from all over the world will play the game within the first week after launch. Currently, the game consists of the following components deployed in a single AWS Region:

- Amazon S3 bucket that stores game assets
- Amazon DynamoDB table that stores player scores

A solutions architect needs to design a Region solution that will reduce latency improve reliability, and require the least effort to implement

What should the solutions architect do to meet these requirements?

- A. Create an Amazon CloudFront distribution to serve assets from the S3 bucket Configure S3Cross-Region Replication Create a new DynamoDB table in a new Region Use the new table as a replica target for DynamoDB global tables.
- B. Create an Amazon CloudFront distribution to serve assets from the S3 bucket
- C. Configure S3Same-Region Replication
- D. Create a new DynamoDB table in a new Region
- E. Configure asynchronous replication between the DynamoDB tables by using AWS Database Migration Service (AWS DMS) with change data capture (CDC)
- F. Create another S3 bucket in a new Region and configure S3 Cross-Region Replication between the buckets Create an Amazon CloudFront distribution and configure origin failover with two origins accessing the S3 buckets in each Region
- G. Configure DynamoDB global tables by enabling Amazon DynamoDB Streams, and add a replica table in a new Region.
- H. Create another S3 bucket in the same Region, and configure S3 Same-Region Replication between the buckets- Create an Amazon CloudFront distribution and configure origin failover with two origins accessing the S3 buckets Create a new DynamoDB table in a new Region Use the new table as a replica target for DynamoDB global tables.

Answer: B

NEW QUESTION 84

- (Exam Topic 2)

A company is planning to migrate an application from on premises to AWS. The application currently uses an Oracle database and the company can tolerate a brief downtime of 1 hour when performing the switch to the new infrastructure As part of the migration, the database engine will be changed to MySQL. A solutions architect needs to determine which AWS services can be used to perform the migration while minimizing the amount of work and time required. Which of the following will meet the requirements?

- A. Use AWS SCT to generate the schema scripts and apply them on the target prior to migration Use AWS DMS to analyse the current schema and provide a recommendation for the optimal database engine Then, use AWS DMS to migrate to the recommended engine Use AWS SCT to identify what embedded SQL code in the application can be converted and what has to be done manually
- B. Use AWS SCT to generate the schema scripts and apply them on the target prior to migration
- C. Use AWS DMS to begin moving data from the on-premises database to AWS
- D. After the initial copy continue to use AWS DMS to keep the databases in sync until cutting over to the new database Use AWS SCT to identify what embedded SQL code in the application can be converted and what has to be done manually.
- E. Use AWS DMS to help identify the best target deployment between installing the database engine on Amazon EC2 directly or moving to Amazon RDS
- F. Then, use AWS DMS to migrate to the platform
- G. Use AWS Application Discovery Service to identify what embedded SQL code in the application can be converted and what has to be done manually.
- H. Use AWS DMS to begin moving data from the on-premises database to AWS After the initial copy, continue to use AWS DMS to keep the databases in sync until cutting over to the new database use AWS Application Discovery Service to identify what embedded SQL code in the application can be converted and what has to be done manually

Answer: B

NEW QUESTION 86

- (Exam Topic 2)

A company has an organization that has many AWS accounts in AWS Organizations A solutions architect must improve how the company manages common security group rules for the AWS accounts in the organization.

The company has a common set of IP CIDR ranges in an allow list in each AWS account to allow access to and from the company's on-premises network Developers within each account are responsible for adding new IP CIDR ranges to their security groups. The security team has its own AWS account. Currently, the security team notifies the owners of the other AWS accounts when changes are made to the allow list.

The solutions architect must design a solution that distributes the common set of CIDR ranges across all accounts Which solution meets these requirements with the LEAST amount of operational overhead.

- A. Set up an Amazon Simple Notification Service (Amazon SNS) topic in the security team's AWS account Deploy an AWS Lambda function in each AWS account Configure the Lambda function to run every time an SNS topic receives a message Configure the Lambda function to take an IP address as input and add it to a list of security groups in the account Instruct the security team to distribute changes by publishing messages to its SNS topic
- B. Create new customer-managed prefix lists in each AWS account within the organization Populate the prefix lists in each account with all internal CIDR ranges Notify the owner of each AWS account to allow the new customer-managed prefix list IDs in their accounts in their security groups Instruct the security team to share updates with each AWS account owner.
- C. Create a new customer-managed prefix list in the security team's AWS account Populate the customer-managed prefix list with all internal CIDR ranges
- D. Share the customer-managed prefix list.... organization by using AWS Resource Access Manager Notify the owner of each AWS account to allow the new customer-managed prefix list ID in their security groups

Answer: A

NEW QUESTION 90

- (Exam Topic 2)

A company's site reliability engineer is performing a review of Amazon FSx for Windows File Server deployments within an account that the company acquired Company policy states that all Amazon FSx file systems must be configured to be highly available across Availability Zones.

During the review, the site reliability engineer discovers that one of the Amazon FSx file systems used a deployment type of Single-AZ 2 A solutions architect needs to minimize downtime while aligning this Amazon FSx file system with company policy.

What should the solutions architect do to meet these requirements?

- A. Reconfigure the deployment type to Multi-AZ for this Amazon FSx file system
- B. Create a new Amazon FSx file system with a deployment type of Multi-AZ
- C. Use AWS DataSync to transfer data to the new Amazon FSx file system
- D. Point users to the new location
- E. Create a second Amazon FSx file system with a deployment type of Single-AZ 2. Use AWS DataSync to keep the data in sync
- F. Switch users to the second Amazon FSx file system in the event of failure
- G. Use the AWS Management Console to take a backup of the Amazon FSx file system Create a new Amazon FSx file system with a deployment type of Multi-AZ Restore the backup to the new Amazon FSx file system
- H. Point users to the new location.

Answer: B

NEW QUESTION 91

- (Exam Topic 2)

A new startup is running a serverless application using AWS Lambda as the primary source of compute. New versions of the application must be made available to a subset of users before deploying changes to all users. Developers should also have the ability to stop the deployment and have access to an easy rollback mechanism. A solutions architect decides to use AWS CodeDeploy to deploy changes when a new version is available.

Which CodeDeploy configuration should the solutions architect use?

- A. A blue/green deployment
- B. A linear deployment
- C. A canary deployment
- D. An all-at-once deployment

Answer: C

NEW QUESTION 96

- (Exam Topic 2)

A news company wants to implement an AWS Lambda function that calls an external API to receive new press releases every 10 minutes. The API provider is planning to use an IP address allow list to protect the API, so the news company needs to provide any public IP addresses that access the API. The company's current architecture includes a VPC with an internet gateway and a NAT gateway. A solutions architect must implement a static IP address for the Lambda function.

Which combination of steps should the solutions architect take to meet these requirements? (Select TWO.)

- A. Use the Elastic IP address that is associated with the NAT gateway for the IP address allow list.
- B. Assign an Elastic IP address to the Lambda function.
- C. Use the Lambda function's Elastic IP address for the IP address allow list.
- D. Configure the Lambda function to launch in the private subnet of the VPC.
- E. Configure the Lambda function to launch in the public subnet of the VPC.
- F. Create a transit gateway.
- G. Attach the VPC and the Lambda function to the transit gateway.

Answer: AC

NEW QUESTION 101

- (Exam Topic 2)

A company is hosting a three-tier web application in an on-premises environment. Due to a recent surge in traffic that resulted in downtime and a significant financial impact, company management has ordered that the application be moved to AWS. The application is written in .NET and has a dependency on a MySQL database. A solutions architect must design a scalable and highly available solution to meet the demand of 200,000 daily users.

Which steps should the solutions architect take to design an appropriate solution?

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

Answer: B

NEW QUESTION 103

- (Exam Topic 2)

A company processes environmental data. The company has set up sensors to provide a continuous stream of data from different areas in a city. The data is available in JSON format.

The company wants to use an AWS solution to send the data to a database that does not require fixed schemas for storage. The data must be sent in real time.

Which solution will meet these requirements?

- A. Use Amazon Kinesis Data Firehose to send the data to Amazon Redshift.
- B. Use Amazon Kinesis Data Streams to send the data to Amazon DynamoDB.
- C. Use Amazon Managed Streaming for Apache Kafka (Amazon MSK) to send the data to Amazon Aurora.
- D. Use Amazon Kinesis Data Firehose to send the data to Amazon Keyspaces (for Apache Cassandra).

Answer: B

NEW QUESTION 106

- (Exam Topic 2)

A retail company runs a business-critical web service on an Amazon Elastic Container Service (Amazon ECS) cluster that runs on Amazon EC2 instances. The web service receives POST requests from end users and writes data to a MySQL database that runs on a separate EC2 instance. The company needs to ensure that data loss does not occur.

The current code deployment process includes manual updates of the ECS service. During a recent deployment, end users encountered intermittent 502 Bad Gateway errors in response to valid web requests.

The company wants to implement a reliable solution to prevent this issue from recurring. The company also wants to automate code deployments. The solution must be highly available and must optimize

cost-effectiveness

- A. Run the web service on an ECS cluster that has a Fargate launch type Use AWS CodePipeline and AWS CodeDeploy to perform a blue/green deployment with validation testing to update the ECS service.
- B. Migrate the MySQL database to run on an Amazon RDS for MySQL Multi-AZ DB instance that uses Provisioned IOPS SSD (io2) storage
- C. Configure an Amazon Simple Queue Service (Amazon SQS) queue as an event source to receive the POST requests from the web service Configure an AWS Lambda function to poll the queue Write the data to the database.
- D. Run the web service on an ECS cluster that has a Fargate launch type Use AWS CodePipeline and AWS CodeDeploy to perform a canary deployment to update the ECS service.

Answer: CD

NEW QUESTION 108

- (Exam Topic 2)

A company's solution architect is designing a disaster recovery (DR) solution for an application that runs on AWS. The application uses PostgreSQL 11.7 as its database. The company has an RPO of 30 seconds. The solution architect must design a DR solution with the primary database in the us-east-1 Region and the database in the us-west-2 Region.

What should the solution architect do to meet these requirements with minimum application change?

- A. Migrate the database to Amazon RDS for PostgreSQL in us-east-1. Set up a read replica up a read replica in us-west-2. Set the managed RPO for the RDS database to 30 seconds.
- B. Migrate the database to Amazon for PostgreSQL in us-east-1. Set up a standby replica in an Availability Zone in us-west-2, Set the managed RPO for the RDS database to 30 seconds.
- C. Migrate the database to an Amazon Aurora PostgreSQL global database with the primary Region as us-east-1 and the secondary Region as us-west-2. Set the managed RPO for the Aurora database to 30 seconds.
- D. Migrate the database to Amazon DynamoDB in us-east-1. Set up global tables with replica tables that are created in us-west-2.

Answer: A

NEW QUESTION 111

- (Exam Topic 2)

A company has its cloud infrastructure on AWS A solutions architect needs to define the infrastructure as code. The infrastructure is currently deployed in one AWS Region. The company's business expansion plan includes deployments in multiple Regions across multiple AWS accounts

What should the solutions architect do to meet these requirements?

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

Answer: B

NEW QUESTION 112

- (Exam Topic 2)

A company is planning to migrate its on-premises data analysis application to AWS. The application is hosted across a fleet of servers and requires consistent system time.

The company has established an AWS Direct Connect connection from its on-premises data center to AWS. The company has a high-precision stratum-0 atomic clock network appliance that acts as an NTP source for all on-premises servers.

After the migration to AWS is complete, the clock on all Amazon EC2 instances that host the application must be synchronized with the on-premises atomic clock network appliance.

Which solution will meet these requirements with the LEAST administrative overhead?

- A. Configure a DHCP options set with the on-premises NTP server address Assign the options set to the VPC
- B. Ensure that NTP traffic is allowed between AWS and the on-premises networks.
- C. Create a custom AMI to use the Amazon Time Sync Service at 169.254.169.123 Use this AMI for the application Use AWS Config to audit the NTP configuration.
- D. Deploy a third-party time server from the AWS Marketplace
- E. Configure the time server to synchronize with the on-premises atomic clock network appliance
- F. Ensure that NTP traffic is allowed inbound in the network ACLs for the VPC that contains the third-party server.
- G. Create an IPsec VPN tunnel from the on-premises atomic clock network appliance to the VPC to encrypt the traffic over the Direct Connect connection
- H. Configure the VPC route tables to direct NTP traffic over the tunnel.

Answer: B

NEW QUESTION 114

- (Exam Topic 2)

What should the solutions architect do to meet this requirement?

- A. / Use Amazon CloudWatch to monitor the Sample Count statistic for each service in the ECS cluster Set an alarm for when the math expression sample Notification SERVICE_QUOTA(service)*100 is greater than 80 Notify the development team by using Amazon Simple Notification Service (Amazon SNS)
- B. Use Amazon CloudWatch to monitor service quotas that are published under the AWS-Usage metric namespace Set an alarm for when the math expression metric SERVICE_QUOTA(metric)*100 is greater than 80 Notify the development team by using Amazon Simple Notification Service (Amazon SNS).
- C. Create an AWS Lambda function to poll detailed metrics from the ECS cluster
- D. When the number running Fargate tasks is greater than 80. invoke Amazon Simple Email Service (Amazon SES) to notify the development team
- E. Create an AWS Config rule to evaluate whether the Fargate SERVICE_QUOTA is greater than 80. Use Amazon Simple Email Service (Amazon SES) to notify the development team when the AWS Config rule is not compliant.

Answer: B

NEW QUESTION 115

- (Exam Topic 2)

A company that designs multiplayer online games wants to expand its user base outside of Europe. The company transfers a significant amount of UDP traffic to Keep all the live and interactive sessions of the games The company has plans for rapid expansion and wants to build its architecture to provide an optimized online experience to its users

Which architecture will meet these requirements with the LOWEST latency for users"

- A. Set up a Multi-AZ environment in a single AWS Region Use Amazon CloudFront to cache user sessions
- B. Set up environments in multiple AWS Regions Create an accelerator in AWS Global Accelerator, and add endpoints from different Regions to it
- C. Set up environments in multiple AWS Regions Use Amazon Route 53. and select latency-based routing
- D. Set up a Multi-AZ environment in a single AWS Region
- E. Use AWS Lambda@Edge to update sessions closer to the users

Answer: B

NEW QUESTION 117

- (Exam Topic 2)

A company has a web application that securely uploads pictures and videos to an Amazon S3 bucket The company requires that only authenticated users are allowed to post content T.he application generates a presigned URL that is used to upload objects through a browser interface. Most users are reporting slow upload times for objects larger than 100 MB

What can a solutions architect do to improve the performance of these uploads while ensuring only authenticated users are allowed to post content?

- A. Set up an Amazon API Gateway with an edge-optimized API endpoint that has a resource as an S3 service proxy Configure the PUT method for this resource to expose the S3 Putobject operation Secure the API Gateway using a cognito_user_pools authorizer Have the browser interface use API Gateway instead of the presigned URL to upload objects
- B. Set up an Amazon API Gateway with a regional API endpoint that has a resource as an S3 service proxy Configure the PUT method for this resource to expose the S3 Putobject operation Secure the API Gateway using an AWS Lambda authorizer Have the browser interface use API Gateway instead of the presigned URL to upload objects
- C. Enable an S3 Transfer Acceleration endpoint on the S3 bucket Use the endpoint when generating the presigned URL Have the browser interface upload the objects to this URL using the S3 multipart upload API
- D. Configure an Amazon CloudFront distribution for the destination S3 bucket Enable PUT and POST methods for the CloudFront cache behavior Update the CloudFront origin to use an origin access identity (OAI) Give the OAI user s 3: Putobject permissions in the bucket policy Have the browser interface upload objects using the CloudFront distribution

Answer: D

NEW QUESTION 118

- (Exam Topic 2)

A company has several applications running in an on-premises data center. The data center runs a mix of Windows and Linux VMs managed by VMware vCenter. A solutions architect needs to create a plan to migrate the applications to AWS However, the solutions architect discovers that the documentation for the applications is not up to date and that there are no complete infrastructure diagrams The company's developers lack time to discuss their applications and current usage with the solutions architect

What should the solutions architect do to gather the required information?

- A. Deploy the AWS Server Migration Service (AWS SMS) connector using the OVA image on the VMware cluster to collect configuration and utilization data from the VMs
- B. Use the AWS Migration Portfolio Assessment (MPA) tool to connect to each of the VMs to collect the configuration and utilization data.
- C. Install the AWS Application Discovery Service on each of the VMs to collect the configuration and utilization data
- D. Register the on-premises VMs with the AWS Migration Hub to collect configuration and utilization data

Answer: A

NEW QUESTION 123

- (Exam Topic 2)

A company wants to deploy an API to AWS. The company plans to run the API on AWS Fargate behind a load balancer. The API requires the use of header-based routing and must be accessible from on-premises networks through an AWS Direct Connect connection and a private VIF.

The company needs to add the client IP addresses that connect to the API to an allow list in AWS. The company also needs to add the IP addresses of the API to the allow list. The company's security team will allow /27 CIDR ranges to be added to the allow list. The solution must minimize complexity and operational overhead.

Which solution will meet these requirements?

- A. Create a new Network Load Balancer (NLB) in the same subnets as the Fargate task deployments. Create a security group that includes only the client IP addresses that need access to the AP
- B. Attach the new security group to the Fargate task
- C. Provide the security team with the NLB's IP addresses for the allow list.
- D. Create two new /27 subnet
- E. Create a new Application Load Balancer (ALB) that extends across the new subnet
- F. Create a security group that includes only the client IP addresses that need access to the AP
- G. Attach the security group to the AL
- H. Provide the security team with the new subnet IP ranges for the allow list.
- I. Create two new '27 subnet
- J. Create a new Network Load Balancer (NLB) that extends across the new subnet
- K. Create a new Application Load Balancer (ALB) within the new subnet
- L. Create a security group that includes only the client IP addresses that need access to the AP
- M. Attach the security group to the AL
- N. Add the ALB's IP addresses as targets behind the NL
- O. Provide the security team with the NLB's IP addresses for the allow list.
- P. Create a new Application Load Balancer (ALB) in the same subnets as the Fargate task deployments. Create a security group that includes only the client IP addresses that need access to the AP

- Q. Attach the security group to the ALB
R. Provide the security team with the ALB's IP addresses for the allow list.

Answer: A

NEW QUESTION 127

- (Exam Topic 2)

An enterprise company wants to allow its developers to purchase third-party software through AWS Marketplace. The company uses an AWS Organizations account structure with full features enabled, and has a shared services account in each organizational unit (OU) that will be used by procurement managers. The procurement team's policy indicates that developers should be able to obtain third-party software from an approved list only and use Private Marketplace in AWS Marketplace to achieve this requirement. The procurement team wants administration of Private Marketplace to be restricted to a role named procurement-manager-role, which could be assumed by procurement managers. Other IAM users, groups, roles, and account administrators in the company should be denied Private Marketplace administrative access.

What is the MOST efficient way to design an architecture to meet these requirements?

- A. Create an IAM role named procurement-manager-role in all AWS accounts in the organization. Add the PowerUserAccess managed policy to the role. Apply an inline policy to all IAM users and roles in every AWS account to deny permissions on the AWSPRivateMarketplaceAdminFullAccess managed policy.
- B. Create an IAM role named procurement-manager-role in all AWS accounts in the organization. Add the AdministratorAccess managed policy to the role. Define a permissions boundary with the AWSPRivateMarketplaceAdminFullAccess managed policy and attach it to all the developer roles.
- C. Create an IAM role named procurement-manager-role in all the shared services accounts in the organization. Add the AWSPRivateMarketplaceAdminFullAccess managed policy to the role. Create an organization root-level SCP to deny permissions to administer Private Marketplace to everyone except the role named procurement-manager-role. Create another organization root-level SCP to deny permissions to create an IAM role named procurement-manager-role to everyone in the organization.
- D. Create an IAM role named procurement-manager-role in the AWS accounts that will be used by developers. Add the AWSPRivateMarketplaceAdminFullAccess managed policy to the role.
- E. Create...Organizations to deny permissions to administer Private Marketplace to everyone except the role named procurement-manager-role. Apply the SCP to all the shared services accounts in the.....

Answer: C

NEW QUESTION 130

- (Exam Topic 2)

A company is creating a sequel for a popular online game. A large number of users from all over the world will play the game within the first week after launch. Currently, the game consists of the following components deployed in a single AWS Region:

- Amazon S3 bucket that stores game assets
- Amazon DynamoDB table that stores player scores

A solutions architect needs to design a multi-Region solution that will reduce latency, improve reliability, and require the least effort to implement.

What should the solutions architect do to meet these requirements?

- A. Create an Amazon CloudFront distribution to serve assets from the S3 bucket. Configure S3 Cross-Region Replication. Create a new DynamoDB table in a new Region. Use the new table as a replica target for DynamoDB global tables.
- B. Create an Amazon CloudFront distribution to serve assets from the S3 bucket.
- C. Configure S3 Same-Region Replication.
- D. Create a new DynamoDB table in a new Region.
- E. Configure asynchronous replication between the DynamoDB tables by using AWS Database Migration Service (AWS DMS) with change data capture (CDC).
- F. Create another S3 bucket in a new Region and configure S3 Cross-Region Replication between the buckets. Create an Amazon CloudFront distribution and configure origin failover with two origins accessing the S3 buckets in each Region.
- G. Configure DynamoDB global tables by enabling Amazon DynamoDB Streams, and add a replica table in a new Region.
- H. Create another S3 bucket in the same Region, and configure S3 Same-Region Replication between the buckets. Create an Amazon CloudFront distribution and configure origin failover with two origins accessing the S3 buckets. Create a new DynamoDB table in a new Region. Use the new table as a replica target for DynamoDB global tables.

Answer: C

NEW QUESTION 131

- (Exam Topic 2)

An auction website enables users to bid on collectible items. The auction rules require that each bid is processed only once and in the order it was received. The current implementation is based on a fleet of Amazon EC2 web servers that write bid records into Amazon Kinesis Data Streams. A single large instance has a cron job that runs the bid processor, which reads incoming bids from Kinesis Data Streams and processes each bid. The auction site is growing in popularity, but users are complaining that some bids are not registering.

Troubleshooting indicates that the bid processor is too slow during peak demand hours, sometimes crashes while processing, and occasionally loses track of which record is being processed.

What changes should be made to make the bid processing more reliable?

- A. Refactor the web application to use the Amazon Kinesis Producer Library (KPL) when posting bids to Kinesis Data Streams. Refactor the bid processor to flag each record in Kinesis Data Streams as being unread, processing, and processed. At the start of each bid processing run, scan Kinesis Data Streams for unprocessed records.
- B. Refactor the web application to post each incoming bid to an Amazon SNS topic in place of Kinesis Data Streams. Configure the SNS topic to trigger an AWS Lambda function that
- C. processes each bid as soon as a user submits it.
- D. Refactor the web application to post each incoming bid to an Amazon SQS FIFO queue in place of Kinesis Data Streams. Refactor the bid processor to continuously consume the SQS queue. Place the bid processing EC2 instance in an Auto Scaling group with a minimum and a maximum size of 1.
- E. Switch the EC2 instance type from t2.large to a larger general compute instance type. Put the bid processor EC2 instances in an Auto Scaling group that scales out the number of EC2 instances running the bid processor based on the incomingRecords metric in Kinesis Data Streams.

Answer: C

Explanation:

<https://aws.amazon.com/sqs/faqs/#:~:text=A%20single%20Amazon%20SQS%20message,20%2C000%20for%2>

NEW QUESTION 133

- (Exam Topic 2)

A company that uses AWS Organizations is creating several new AWS accounts. The company is setting up controls to properly allocate AWS costs to business units. The company must implement a solution to ensure that all resources include a tag that has a key of costcenter and a value from a predefined list of business units. The solution must send a notification each time a resource tag does not meet these criteria. The solution must not prevent the creation of resources. Which solution will meet these requirements with the LEAST operational overhead?

- A. Create an IAM policy for all actions that create AWS resource
- B. Add a condition to the policy that aws:RequestTag/costcenter must exist and must contain a valid business unit value
- C. Create an Amazon EventBridge (Amazon CloudWatch Events) rule that monitors IAM service events and Amazon EC2 service events for noncompliant tag policies
- D. Configure the rule to send notifications through Amazon Simple Notification Service (Amazon SNS).
- E. Create an IAM policy for all actions that create AWS resource
- F. Add a condition to the policy that aws:ResourceTag/costcenter must exist and must contain a valid business unit value
- G. Create an Amazon EventBridge (Amazon CloudWatch Events) rule that monitors IAM service events and Amazon EC2 service events for noncompliant tag policies
- H. Configure the rule to send notifications through Amazon Simple Notification Service (Amazon SNS).
- I. Create an organization tag policy that ensures that all resources have the costcenter tag with a valid business unit value
- J. Do not select the option to prevent operations when tags are noncompliant
- K. Create an Amazon EventBridge (Amazon CloudWatch Events) rule that monitors all events for noncompliant tag policies
- L. Configure the rule to send notifications through Amazon Simple Notification Service (Amazon SNS).
- M. Create an organization tag policy that ensures that all resources have the costcenter tag with a valid business unit value
- N. Select the option to prevent operations when tags are noncompliant
- O. Create an Amazon EventBridge (Amazon CloudWatch Events) rule that monitors all events for noncompliant tag policies
- P. Configure the rule to send notifications through Amazon Simple Notification Service (Amazon SNS).

Answer: B

NEW QUESTION 136

- (Exam Topic 2)

A company has an application. Once a month, the application creates a compressed file that contains every object within an Amazon S3 bucket. The total size of the objects before compression is 1 TB.

The application runs by using a scheduled cron job on an Amazon EC2 instance that has a 5 TB Amazon Elastic Block Store (Amazon EBS) volume attached. The application downloads all the files from the source S3 bucket to the EBS volume, compresses the file, and uploads the file to a target S3 bucket. Every invocation of the application takes 2 hours from start to finish.

Which combination of actions should a solutions architect take to OPTIMIZE costs for this application? (Select TWO.)

- A. Migrate the application to run as an AWS Lambda function. Use Amazon EventBridge (Amazon CloudWatch Events) to schedule the Lambda function to run once each month.
- B. Configure the application to download the source files by using streams. Direct the streams into a compression library. Direct the output of the compression library into a target object in Amazon S3.
- C. Configure the application to download the source files from Amazon S3 and save the files to local storage. Compress the files and upload them to Amazon S3.
- D. Configure the application to run as a container in AWS Fargate. Use Amazon EventBridge (Amazon CloudWatch Events) to schedule the task to run once each month.
- E. Provision an Amazon Elastic File System (Amazon EFS) file system. Attach the file system to the AWS Lambda function.

Answer: CD

NEW QUESTION 137

- (Exam Topic 2)

A company is using an existing orchestration tool to manage thousands of Amazon EC2 instances. A recent penetration test found a vulnerability in the company's software stack. This vulnerability has prompted the company to perform a full evaluation of its current production environment. The analysts determined that the following vulnerabilities exist within the environment:

- Operating systems with outdated libraries and known vulnerabilities are being used in production
- Relational databases hosted and managed by the company are running unsupported versions with known vulnerabilities
- Data stored in databases is not encrypted.

The solutions architect intends to use AWS Config to continuously audit and assess the compliance of the company's AWS resource configurations with the company's policies and guidelines. What additional steps will enable the company to secure its environments and track resources while adhering to best practices?

- A. Use AWS Application Discovery Service to evaluate all running EC2 instances. Use the AWS CLI to modify each instance, and use EC2 user data to install the AWS Systems Manager Agent during boot. Schedule patching to run as a Systems Manager Maintenance Window task.
- B. Migrate all relational databases to Amazon RDS and enable AWS KMS encryption.
- C. Create an AWS CloudFormation template for the EC2 instances. Use EC2 user data in the CloudFormation template to install the AWS Systems Manager Agent, and enable AWS KMS encryption on all Amazon EBS volumes.
- D. Have CloudFormation replace all running instances.
- E. Use Systems Manager Patch Manager to establish a patch baseline and deploy a Systems Manager Maintenance Windows task to run AWS-RunPatchBaseline using the patch baseline.
- F. Install the AWS Systems Manager Agent on all existing instances using the company's current orchestration tool. Use the Systems Manager Run Command to run a list of commands to upgrade software on each instance using operating system-specific tools.
- G. Enable AWS KMS encryption on all Amazon EBS volumes.
- H. Install the AWS Systems Manager Agent on all existing instances using the company's current orchestration tool.
- I. Migrate all relational databases to Amazon RDS and enable AWS KMS encryption. Use Systems Manager Patch Manager to establish a patch baseline and deploy a Systems Manager Maintenance Windows task to run AWS-RunPatchBaseline using the patch baseline.

Answer: D

NEW QUESTION 140

- (Exam Topic 2)

A large company has a business-critical application that runs in a single AWS Region. The application consists of multiple Amazon EC2 instances and an Amazon RDS Multi-AZ DB instance. The EC2 instances run in an Amazon EC2 Auto Scaling group across multiple Availability Zones.

A solutions architect is implementing a disaster recovery (DR) plan for the application. The solutions architect has created a pilot light application deployment in a

new Region, which is referred to as the DR Region The DR environment has an Auto Scaling group with a single EC2 instance and a read replica of the RDS DB instance

The solutions architect must automate a failover from the primary application environment to the pilot light environment in the DR Region

Which solution meets these requirements with the MOST operational efficiency"

- A. Publish an application availability metric to Amazon CloudWatch in the DR Region from the application environment in the primary Region Create a CloudWatch alarm in the DR Region that is invoked when the application availability metric stops being delivered Configure the CloudWatch alarm to send a notification to an Amazon Simple Notification Service (Amazon SNS) topic in the DR Region Add an email subscription to the SNS topic that sends messages to the application owner upon notification, instruct a systems operator to sign in to the AWS Management Console and initiate failover operations for the application
- B. Create a cron task that runs every 5 minutes by using one of the application's EC2 instances in the primary Region Configure the cron task to check whether the application is available Upon failure, the cron task notifies a systems operator and attempts to restart the application services
- C. Create a cron task that runs every 5 minutes by using one of the application's EC2 instances in the primary Region Configure the cron task to check whether the application is available Upon failure, the cron task modifies the DR environment by promoting the read replica and by adding EC2 instances to the Auto Scaling group
- D. Publish an application availability metric to Amazon CloudWatch in the DR Region from the application environment in the primary Region Create a CloudWatch alarm in the DR Region that is invoked when the application availability metric stops being delivered Configure the CloudWatch alarm to send a notification to an Amazon Simple Notification Service (Amazon SNS) topic in the DR Region Use an AWS Lambda function that is invoked by Amazon SNS in the DR Region to promote the read replica and to add EC2 instances to the Auto Scaling group

Answer: D

NEW QUESTION 142

- (Exam Topic 2)

A company is developing a gene reporting device that will collect genomic information to assist researchers with collecting large samples of data from a diverse population. The device will push 8 KB of genomic data every second to a data platform that will need to process and analyze the data and provide information back to researchers The data platform must meet the following requirements:

- Provide near-real-time analytics of the inbound genomic data
- Ensure the data is flexible, parallel, and durable
- Deliver results of processing to a data warehouse

Which strategy should a solutions architect use to meet these requirements?

- A. Use Amazon Kinesis Data Firehose to collect the inbound sensor data analyze the data with Kinesis client
- B. and save the results to an Amazon RDS instance
- C. Use Amazon Kinesis Data Streams to collect the inbound sensor data analyze the data with Kinesis clients and save the results to an Amazon Redshift duster using Amazon EMR
- D. Use Amazon S3 to collect the inbound device data analyze the data from Amazon SOS with Kinesis and save the results to an Amazon Redshift duster
- E. Use an Amazon API Gateway to put requests into an Amazon SQS queue analyze the data with an AWS Lambda function and save the results » an Amazon Redshift duster using Amazon EMR

Answer: A

NEW QUESTION 144

- (Exam Topic 2)

A retail company has a small ecommerce web application that uses an Amazon RDS for PostgreSQL DB instance The DB instance is deployed with the Multi-AZ option turned on.

Application usage recently increased exponentially and users experienced frequent HTTP 503 errors Users reported the errors, and the company's reputation suffered The company could not identify a definitive root cause.

The company wants to improve its operational readiness and receive alerts before users notice an incident The company also wants to collect enough information to determine the root cause of any future incident.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Turn on Enhanced Monitoring for the DB instance Modify the corresponding parameter group to turn on query logging for all the slow queries Create Amazon CloudWatch alarms Set the alarms to appropriate thresholds that are based on performance metrics in CloudWatch
- B. Turn on Enhanced Monitoring and Performance Insights for the DB instance Create Amazon CloudWatch alarms Set the alarms to appropriate thresholds that are based on performance metrics in CloudWatch
- C. Turn on log exports to Amazon CloudWatch for the PostgreSQL logs on the DB instance Analyze the logs by using Amazon Elasticsearch Service (Amazon ES) and Kibana Create a dashboard in Kibana Configure alerts that are based on the metrics that are collected
- D. Turn on Performance Insights for the DB instance Modify the corresponding parameter group to turn on query logging for all the slow queries Create Amazon CloudWatch alarms Set the alarms to appropriate thresholds that are based on performance metrics in CloudWatch

Answer: A

NEW QUESTION 147

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questions and Answers in PDF Format

SAP-C02 Practice Exam Features:

- * SAP-C02 Questions and Answers Updated Frequently
- * SAP-C02 Practice Questions Verified by Expert Senior Certified Staff
- * SAP-C02 Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- * SAP-C02 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The SAP-C02 Practice Test Here](#)