

SAP-C02 Dumps

AWS Certified Solutions Architect - Professional

<https://www.certleader.com/SAP-C02-dumps.html>



NEW QUESTION 1

- (Exam Topic 1)

A company wants to deploy an AWS WAF solution to manage AWS WAF rules across multiple AWS accounts. The accounts are managed under different OUs in AWS Organizations.

Administrators must be able to add or remove accounts or OUs from managed AWS WAF rule sets as needed. Administrators also must have the ability to automatically update and remediate noncompliant AWS WAF rules in all accounts

Which solution meets these requirements with the LEAST amount of operational overhead?

- A. Use AWS Firewall Manager to manage AWS WAF rules across accounts in the organization
- B. Use an AWS Systems Manager Parameter Store parameter to store account numbers and OUs to manage Update the parameter as needed to add or remove accounts or OUs Use an Amazon EventBridge (Amazon CloudWatch Events) rule to identify any changes to the parameter and to invoke an AWS Lambda function to update the security policy in the Firewall Manager administrative account
- C. Deploy an organization-wide AWS Config rule that requires all resources in the selected OUs to associate the AWS WAF rule
- D. Deploy automated remediation actions by using AWS Lambda to fix noncompliant resource
- E. Deploy AWS WAF rules by using an AWS CloudFormation stack set to target the same OUs where the AWS Config rule is applied.
- F. Create AWS WAF rules in the management account of the organization
- G. Use AWS Lambda environment variables to store account numbers and OUs to manage Update environment variables as needed to add or remove accounts or OUs Create cross-account IAM roles in member account
- H. Assume the roles by using AWS Security Token Service (AWS STS) in the Lambda function to create and update AWS WAF rules in the member accounts
- I. Use AWS Control Tower to manage AWS WAF rules across accounts in the organization
- J. Use AWS Key Management Service (AWS KMS) to store account numbers and OUs to manage Update AWS KMS as needed to add or remove accounts or OUs
- K. Create IAM users in member accounts Allow AWS Control Tower in the management account to use the access key and secret access key to create and update AWS WAF rules in the member accounts

Answer: B

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 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 4

- (Exam Topic 1)

A solutions architect has an operational workload deployed on Amazon EC2 instances in an Auto Scaling group. The VPC architecture spans two Availability Zones (AZ) with a subnet in each that the Auto Scaling group is targeting. The VPC is connected to an on-premises environment and connectivity cannot be interrupted. The maximum size of the Auto Scaling group is 20 instances in service. The VPC IPv4 addressing is as follows:

VPC CIDR: 10.0.0.0/23

AZ1 subnet CIDR: 10.0.0.0/24 AZ2 subnet CIDR: 10.0.1.0/24

Since deployment, a third AZ has become available in the Region. The solutions architect wants to adopt the new AZ without adding additional IPv4 address space and without service downtime.

Which solution will meet these requirements?

- A. Update the Auto Scaling group to use the AZ2 subnet only
- B. Delete and re-create the AZ1 subnet using half the previous address space
- C. Adjust the Auto Scaling group to also use the new AZ1 subnet
- D. When the instances are healthy, adjust the Auto Scaling group to use the AZ1 subnet only
- E. Remove the current AZ2 subnet
- F. Create a new AZ2 subnet using the second half of the address space from the original AZ1 subnet
- G. Create a new AZ3 subnet using half the original AZ2 subnet address space, then update the Auto Scaling group to target all three new subnets.
- H. Terminate the EC2 instances in the AZ1 subnet
- I. Delete and re-create the AZ1 subnet using half the address space
- J. Update the Auto Scaling group to use this new subnet
- K. Repeat this for the second AZ
- L. Define a new subnet in AZ3, then update the Auto Scaling group to target all three new subnets.
- M. Create a new VPC with the same IPv4 address space and define three subnets, with one for each AZ
- N. Update the existing Auto Scaling group to target the new subnets in the new VPC.
- O. Update the Auto Scaling group to use the AZ2 subnet only
- P. Update the AZ1 subnet to have half the previous address space
- Q. Adjust the Auto Scaling group to also use the AZ1 subnet again
- R. When the instances are healthy, adjust the Auto Scaling group to use the AZ1 subnet only
- S. Update the current AZ2 subnet and assign the second half of the address space from the original AZ1 subnet
- T. Create a new AZ3 subnet using half the original AZ2 subnet address space, then update the Auto Scaling group to target all three new subnets.

Answer: A

Explanation:

https://aws.amazon.com/premiumsupport/knowledge-center/vpc-ip-address-range/?nc1=h_ls

It's not possible to modify the IP address range of an existing virtual private cloud (VPC) or subnet. You must delete the VPC or subnet, and then create a new VPC or subnet with your preferred CIDR block.

NEW QUESTION 5

- (Exam Topic 1)

A company wants to retire its Oracle Solaris NFS storage arrays. The company requires rapid data migration over its internet network connection to a combination of destinations for Amazon S3, Amazon Elastic File System (Amazon EFS), and Amazon FSx for Windows File Server. The company also requires a full initial copy, as well as incremental transfers of changes until the retirement of the storage arrays. All data must be encrypted and checked for integrity.

What should a solutions architect recommend to meet these requirements?

- A. Configure CloudEndure
- B. Create a project and deploy the CloudEndure agent and token to the storage arrays
- C. Run the migration plan to start the transfer.
- D. Configure AWS DataSync
- E. Configure the DataSync agent and deploy it to the local network
- F. Create a transfer task and start the transfer.
- G. Configure the aws S3 sync command
- H. Configure the AWS client on the client side with credential
- I. Run the sync command to start the transfer.
- J. Configure AWS Transfer (or FTP)
- K. Configure the FTP client with credential
- L. Script the client to connect and sync to start the transfer.

Answer: B

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 travel company built a web application that uses Amazon Simple Email Service (Amazon SES) to send email notifications to users. The company needs to enable logging to help troubleshoot email delivery issues. The company also needs the ability to do searches that are based on recipient, subject, and time sent.

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

- A. Create an Amazon SES configuration set with Amazon Kinesis Data Firehose as the destination
- B. Choose to send logs to an Amazon S3 bucket.
- C. Enable AWS CloudTrail logging
- D. Specify an Amazon S3 bucket as the destination for the logs.
- E. Use Amazon Athena to query the logs in the Amazon S3 bucket for recipient, subject, and time sent.
- F. Create an Amazon CloudWatch log group
- G. Configure Amazon SES to send logs to the log group
- H. Use Amazon Athena to query the logs in Amazon CloudWatch for recipient, subject, and time sent.

Answer: AC

Explanation:

<https://docs.aws.amazon.com/ses/latest/dg/event-publishing-retrieving-firehose.html>

To enable you to track your email sending at a granular level, you can set up Amazon SES to publish email sending events to Amazon CloudWatch, Amazon Kinesis Data Firehose, or Amazon Simple Notification Service based on characteristics that you define.

<https://docs.aws.amazon.com/ses/latest/dg/monitor-using-event-publishing.html>

<https://aws.amazon.com/getting-started/hands-on/build-serverless-real-time-data-processing-app-lambda-kinesis>

NEW QUESTION 8

- (Exam Topic 1)

A company has 50 AWS accounts that are members of an organization in AWS Organizations. Each account contains multiple VPCs. The company wants to use AWS Transit Gateway to establish connectivity between the VPCs in each member account. Each time a new member account is created, the company wants to automate the process of creating a new VPC and a transit gateway attachment.

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

- A. From the management account, share the transit gateway with member accounts by using AWS Resource Access Manager
- B. From the management account, share the transit gateway with member accounts by using an AWS Organizations SCP
- C. Launch an AWS CloudFormation stack set from the management account that automatically creates a new VPC and a VPC transit gateway attachment in a member account
- D. Associate the attachment with the transit gateway in the management account by using the transit gateway ID.
- E. Launch an AWS CloudFormation stack set from the management account that automatically creates a new VPC and a peering transit gateway attachment in a member account
- F. Share the attachment with the transit gateway in the management account by using a transit gateway service-linked role.
- G. From the management account, share the transit gateway with member accounts by using AWS Service Catalog

Answer: AC

NEW QUESTION 9

- (Exam Topic 1)

A company that tracks medical devices in hospitals wants to migrate its existing storage solution to the AWS Cloud. The company equips all of its devices with sensors that collect location and usage information. This sensor data is sent in unpredictable patterns with large spikes. The data is stored in a MySQL database running on premises at each hospital. The company wants the cloud storage solution to scale with usage.

The company's analytics team uses the sensor data to calculate usage by device type and hospital. The team needs to keep analysis tools running locally while fetching data from the cloud. The team also needs to use existing Java application and SQL queries with as few changes as possible.

How should a solutions architect meet these requirements while ensuring the sensor data is secure?

- A. Store the data in an Amazon Aurora Serverless database
- B. Serve the data through a Network Load Balancer (NLB). Authenticate users using the NLB with credentials stored in AWS Secrets Manager.
- C. Store the data in an Amazon S3 bucket
- D. Serve the data through Amazon QuickSight using an IAM user authorized with AWS Identity and Access Management (IAM) with the S3 bucket as the data source.
- E. Store the data in an Amazon Aurora Serverless database
- F. Serve the data through the Aurora Data API using an IAM user authorized with AWS Identity and Access Management (IAM) and the AWS Secrets Manager ARN.
- G. Store the data in an Amazon S3 bucket
- H. Serve the data through Amazon Athena using AWS PrivateLink to secure the data in transit.

Answer: C

Explanation:

<https://aws.amazon.com/blogs/aws/new-data-api-for-amazon-aurora-serverless/> <https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/data-api.html>

<https://aws.amazon.com/blogs/aws/aws-privatelink-for-amazon-s3-now-available/> <https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/data-api.html#data-api.access>

The data is currently stored in a MySQL database running on-prem. Storing MySQL data in S3 doesn't sound good so B & D are out. Aurora Data API "enables the SQL HTTP endpoint, a connectionless Web Service API for running SQL queries against this database. When the SQL HTTP endpoint is enabled, you can also query your database from inside the RDS console (these features are free to use)."

NEW QUESTION 10

- (Exam Topic 1)

A company is running a data-intensive application on AWS. The application runs on a cluster of hundreds of Amazon EC2 instances. A shared file system also runs on several EC2 instances that store 200 TB of data. The application reads and modifies the data on the shared file system and generates a report. The job runs once monthly, reads a subset of the files from the shared file system, and takes about 72 hours to complete. The compute instances scale in an Auto Scaling group, but the instances that host the shared file system run continuously. The compute and storage instances are all in the same AWS Region.

A solutions architect needs to reduce costs by replacing the shared file system instances. The file system must provide high performance access to the needed data for the duration of the 72-hour run.

Which solution will provide the LARGEST overall cost reduction while meeting these requirements?

- A. Migrate the data from the existing shared file system to an Amazon S3 bucket that uses the S3 Intelligent-Tiering storage class

- B. Before the job runs each month, use Amazon FSx for Lustre to create a new file system with the data from Amazon S3 by using lazy loadin
- C. Use the new file system as the shared storage for the duration of the jo
- D. Delete the file system when the job is complete.
- E. Migrate the data from the existing shared file system to a large Amazon Elastic Block Store (Amazon EBS) volume with Multi-Attach enable
- F. Attach the EBS volume to each of the instances by using a user data script in the Auto Scaling group launch templat
- G. Use the EBS volume as the shared storage for the duration of the jo
- H. Detach the EBS volume when the job is complete.
- I. Migrate the data from the existing shared file system to an Amazon S3 bucket that uses the S3 Standard storage clas
- J. Before the job runs each month, use Amazon FSx for Lustre to create a new file system with the data from Amazon S3 by using batch loadin
- K. Use the new file system as the shared storage for the duration of the jo
- L. Delete the file system when the job is complete.
- M. Migrate the data from the existing shared file system to an Amazon S3 bucke
- N. Before the job runs each month, use AWS Storage Gateway to create a file gateway with the data from Amazon S3. Use the file gateway as the shared storage for the jo
- O. Delete the file gateway when the job is complete.

Answer: B

NEW QUESTION 10

- (Exam Topic 1)

A company has application services that have been containerized and deployed on multiple Amazon EC2 instances with public IPs. An Apache Kafka cluster has been deployed to the EC2 instances. A PostgreSQL database has been migrated to Amazon RDS lor PostgreSQL. The company expects a significant increase of orders on its platform when a new version of its flagship product is released.

What changes to the current architecture will reduce operational overhead and support the product release?

- A. Create an EC2 Auto Scaling group behind an Application Load Balance
- B. Create additional read replicas for the DB instanc
- C. Create Amazon Kinesis data streams and configure the application services lo use the data stream
- D. Store and serve static content directly from Amazon S3.
- E. Create an EC2 Auto Scaling group behind an Application Load Balance
- F. Deploy the DB instance in Multi-AZ mode and enable storage auto scalin
- G. Create Amazon Kinesis data streams and configure the application services to use the data stream
- H. Store and serve static content directly from Amazon S3.
- I. Deploy the application on a Kubernetes cluster created on the EC2 instances behind an Application Load Balance
- J. Deploy the DB instance in Multi-AZ mode and enable storage auto scalin
- K. Create an Amazon Managed Streaming for Apache Kafka cluster and configure the application services to use the cluste
- L. Store static content in Amazon S3 behind an Amazon CloudFront distribution.
- M. Deploy the application on Amazon Elastic Kubernetes Service (Amazon EKS) with AWS Fargate and enable auto scaling behind an Application Load Balance
- N. Create additional read replicas for the DB instanc
- O. Create an Amazon Managed Streaming for Apache Kafka cluster and configure the application services to use the cluste
- P. Store static content in Amazon S3 behind an Amazon CloudFront distribution.

Answer: D

Explanation:

Deploy the application on Amazon Elastic Kubernetes Service (Amazon EKS) with AWS Fargate and enable auto scaling behind an Application Load Balancer. Create additional read replicas for the DB instance. Create an Amazon Managed Streaming for Apache Kafka cluster and configure the application services to use the cluster. Store static content in Amazon S3 behind an Amazon CloudFront distribution.

NEW QUESTION 12

- (Exam Topic 1)

A company is using AWS Organizations lo manage multiple accounts. Due to regulatory requirements, the company wants to restrict specific member accounts to certain AWS Regions, where they are permitted to deploy resources. The resources in the accounts must be tagged, enforced based on a group standard, and centrally managed with minimal configuration.

What should a solutions architect do to meet these requirements?

- A. Create an AWS Config rule in the specific member accounts to limit Regions and apply a tag policy.
- B. From the AWS Billing and Cost Management console, in the master account, disable Regions for the specific member accounts and apply a tag policy on the root.
- C. Associate the specific member accounts with the roo
- D. Apply a tag policy and an SCP using conditions to limit Regions.
- E. Associate the specific member accounts with a new O
- F. Apply a tag policy and an SCP using conditions to limit Regions.

Answer: D

NEW QUESTION 15

- (Exam Topic 1)

To abide by industry regulations, a solutions architect must design a solution that will store a company's critical data in multiple public AWS Regions, including in the United States, where the company's headquarters is located. The solutions architect is required to provide access to the data stored in AWS to the company's global WAN network. The security team mandates that no traffic accessing this data should traverse the public internet.

How should the solutions architect design a highly available solution that meets the requirements and is cost-effective?

- A. Establish AWS Direct Connect connections from the company headquarters to all AWS Regions in use. Use the company WAN lo send traffic over to the headquarters and then to the respective DX connection to access the data.
- B. Establish two AWS Direct Connect connections from the company headquarters to an AWS Region. Use the company WAN to send traffic over a DX connectio
- C. Use inter-region VPC peering to access the data in other AWS Regions.
- D. Establish two AWS Direct Connect connections from the company headquarters to an AWS Region. Use the company WAN to send traffic over a DX connectio
- E. Use an AWS transit VPC solution to access data in other AWS Regions.
- F. Establish two AWS Direct Connect connections from the company headquarters to an AWS Region. Use the company WAN to send traffic over a DX connectio

G. Use Direct Connect Gateway to access data in other AWS Regions.

Answer: D

Explanation:

This feature also allows you to connect to any of the participating VPCs from any Direct Connect location, further reducing your costs for making using AWS services on a cross-region basis. <https://aws.amazon.com/blogs/aws/new-aws-direct-connect-gateway-inter-region-vpc-access/>
<https://docs.aws.amazon.com/whitepapers/latest/aws-vpc-connectivity-options/aws-direct-connect-aws-transit-g>

NEW QUESTION 18

- (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 de facto 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 has developed a single-page web application in JavaScript. The source code is stored in a single Amazon S3 bucket in the us-east-1 Region. The company serves the web application to a global user base through Amazon CloudFront.

The company wants to experiment with two versions of the website without informing application users. Each version of the website will reside in its own S3 bucket. The company wants to determine which version is most successful in marketing a new product.

The solution must send application users that are based in Europe to the new website design. The solution must send application users that are based in the United States to the current website design. However, some exceptions exist. The company needs to be able to redirect specific users to the new website design, regardless of the users' location.

Which solution meets these requirements?

- A. Configure two CloudFront distributions
- B. Configure a geolocation routing policy in Amazon Route 53 to route traffic to the appropriate CloudFront endpoint based on the location of clients.
- C. Configure a single CloudFront distribution
- D. Create a behavior with different paths for each version of the site
- E. Configure Lambda@Edge on the default path to generate redirects and send the client to the correct version of the website.
- F. Configure a single CloudFront distribution
- G. Configure an alternate domain name on the distribution. Configure two behaviors to route users to the different S3 origins based on the domain name that the client uses in the HTTP request.
- H. Configure a single CloudFront distribution with Lambda@Edge
- I. Use Lambda@Edge to send user requests to different origins based on request attributes.

Answer: A

NEW QUESTION 24

- (Exam Topic 1)

A company plans to migrate to AWS. A solutions architect uses AWS Application Discovery Service over the fleet and discovers that there is an Oracle data warehouse and several PostgreSQL databases. Which combination of migration patterns will reduce licensing costs and operational overhead? (Select TWO.)

- A. Lift and shift the Oracle data warehouse to Amazon EC2 using AWS DMS.
- B. Migrate the Oracle data warehouse to Amazon Redshift using AWS SCT and AWS QMS.
- C. Lift and shift the PostgreSQL databases to Amazon EC2 using AWS DMS.
- D. Migrate the PostgreSQL databases to Amazon RDS for PostgreSQL using AWS DMS
- E. Migrate the Oracle data warehouse to an Amazon EMR managed cluster using AWS DMS.

Answer: BD

Explanation:

<https://aws.amazon.com/getting-started/hands-on/migrate-oracle-to-amazon-redshift/> <https://docs.aws.amazon.com/prescriptive-guidance/latest/patterns/migrate-an-on-premises-postgresql-database>

NEW QUESTION 26

- (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 30

- (Exam Topic 1)

A solutions architect works for a government agency that has strict disaster recovery requirements All Amazon Elastic Block Store (Amazon EBS) snapshots are required to be saved in at least two additional AWS Regions. The agency also is required to maintain the lowest possible operational overhead.

Which solution meets these requirements?

- A. Configure a policy in Amazon Data Lifecycle Manager (Amazon DLMJ) to run once daily to copy the EBS snapshots to the additional Regions.
- B. Use Amazon EventBridge (Amazon CloudWatch Events) to schedule an AWS Lambda function to copy the EBS snapshots to the additional Regions.
- C. Set up AWS Backup to create the EBS snapshot
- D. Configure Amazon S3 cross-Region replication to copy the EBS snapshots to the additional Regions.
- E. Schedule Amazon EC2 Image Builder to run once daily to create an AMI and copy the AMI to the additional Regions.

Answer: B

NEW QUESTION 34

- (Exam Topic 1)

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

Which solution meets these requirements?

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

Answer: A

Explanation:

A Direct Connect gateway is a globally available resource. You can create the Direct Connect gateway in any Region and access it from all other Regions. The following describe scenarios where you can use a Direct Connect gateway.

<https://docs.aws.amazon.com/directconnect/latest/UserGuide/direct-connect-gateways-intro.html>

NEW QUESTION 38

- (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 (or 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 instanc
- B. Create an Amazon S3 bucket to be used for SFTP file hostin
- C. Create an AWS Transfer Family server Configure the Transfer Family server with a publicly accessible endpoint Associate the SFTP Elastic IP address with the new endpoint Point the Transfer Family server to the S3 bucke
- D. Sync all files from the SFTP server to the S3 bucket.
- E. Disassociate the Elastic IP address from the EC2 instanc
- F. Create an Amazon S3 bucket to be used for SFTP file hostin
- G. Create an AWS Transfer Family serve
- H. Configure the Transfer Family server with aVPC-hoste
- I. internet-facing endpoint
- J. Associate the SFTP Elastic IP address with the new endpoint
- K. Attach the security group with customer IP addresses to the new endpoint
- L. Point the Transfer Family server to the S3 bucket Sync all files from the SFTP server to the S3 bucket.
- M. Disassociate the Elastic IP address from the EC2 instanc
- N. Create a new Amazon Elastic File System{Amazon EFS) file system to be used for SFTP file hostin
- O. Create an AWS Fargate task definition to run an SFTP serve
- P. Specify the EFS file system as a mount in the task definitio
- Q. 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 serve
- R. Associate the Elastic IP address with the NL
- S. Sync all files from the SFTP server to the S3 bucket.
- T. Disassociate the Elastic IP address from the EC2 instanc
- . Create a multi-attach Amazon Elastic Block Store (Amazon EBS) volume to be used for SFTP file hostin
- . Create a Network Load Balancer (NLB) with the Elastic IP address attache
- . 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 newmulti-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 launch
- . Sync all files from the SFTP server to the new multi-attach EBS volume.

Answer: B

Explanation:

<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 39

- (Exam Topic 1)

A medical company is running a REST API on a set of Amazon EC2 instances. The EC2 instances run in an Auto Scaling group behind an Application Load Balancer (ALB). The ALB runs in three public subnets, and the EC2 instances run in three private subnets. The company has deployed an Amazon CloudFront distribution that has the ALB as the only origin.

Which solution should a solutions architect recommend to enhance the origin security?

- A. Store a random string in AWS Secrets Manage
- B. Create an AWS Lambda (unction for automatic secret rotatio
- C. Configure CloudFront to inject the random string as a custom HTTP header for the origin reques
- D. Create an AWS WAF web ACL rule with a string match rule for the custom heade
- E. Associate the web ACL with the ALB.
- F. Create an AWS WAF web ACL rule with an IP match condition of the CloudFront service IP address range
- G. Associate the web ACL with the AL
- H. Move the ALB into the three private subnets.
- I. Store a random string in AWS Systems Manager Parameter Stor
- J. Configure Parameter Store automatic rotation for the strin
- K. Configure CloudFront to inject the random siring as a custom HTTP header for the origin reques
- L. Inspect the value of the custom HTTP header, and block access in the ALB.
- M. Configure AWS Shield Advance
- N. Create a security group policy to allow connections from CloudFront service IP address range
- O. Add the policy to AWS Shield Advanced, and attach the policy to the ALB.

Answer: D

Explanation:

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

it shows For Amazon EC2 Auto Scaling, there are two primary process types: Launch and Terminate. The Launch process adds a new Amazon EC2 instance to an Auto Scaling group, increasing its capacity. The Terminate process removes an Amazon EC2 instance from the group, decreasing its capacity. HealthCheck process for EC2 autoscaling is not a primary process! It is a process along with the following AddToLoadBalancer AlarmNotification AZRebalance HealthCheck InstanceRefresh ReplaceUnhealthy ScheduledActions From the requirements, Some EC2 instances are now being marked as unhealthy and are being terminated. Application is running at reduced capacity not because instances are marked unhealthy but because they are being terminated.

<https://docs.aws.amazon.com/autoscaling/ec2/userguide/as-suspend-resume-processes.html#choosing-suspend-r>

NEW QUESTION 44

- (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 48

- (Exam Topic 1)

A company runs a popular public-facing e-commerce website. Its user base is growing quickly from a local market to a national market. The website is hosted in an on-premises data center with web servers and a MySQL database. The company wants to migrate its workload to AWS. A solutions architect needs to create a solution to:

- Improve security
- Improve reliability
- Improve availability
- Reduce latency
- Reduce maintenance

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

- A. Use Amazon EC2 instances in two Availability Zones for the web servers in an Auto Scaling group behind an Application Load Balancer.
- B. Migrate the database to a Multi-AZ Amazon Aurora MySQL DB cluster.
- C. Use Amazon EC2 instances in two Availability Zones to host a highly available MySQL database cluster.
- D. Host static website content in Amazon S3. Use S3 Transfer Acceleration to reduce latency while serving webpage
- E. Use AWS WAF to improve website security.
- F. Host static website content in Amazon S3. Use Amazon CloudFront to reduce latency while serving webpage
- G. Use AWS WAF to improve website security
- H. Migrate the database to a single-AZ Amazon RDS for MySQL DB instance.

Answer: ABE

NEW QUESTION 50

- (Exam Topic 1)

A company is using AWS CodePipeline for the CI/CD of an application to an Amazon EC2 Auto Scaling group. All AWS resources are defined in AWS CloudFormation templates. The application artifacts are stored in an Amazon S3 bucket and deployed to the Auto Scaling group using instance user data scripts. As the application has become more complex, recent resource changes in the CloudFormation templates have caused unplanned downtime. How should a solutions architect improve the CI/CD pipeline to reduce the likelihood that changes in the templates will cause downtime?

- A. Adapt the deployment scripts to detect and report CloudFormation error conditions when performing deployment
- B. Write test plans for a testing team to execute in a non-production environment before approving the change for production.
- C. Implement automated testing using AWS CodeBuild in a test environment
- D. Use CloudFormation change sets to evaluate changes before deployment
- E. Use AWS CodeDeploy to leverage blue/green deployment patterns to allow evaluations and the ability to revert changes, if needed.
- F. Use plugins for the integrated development environment (IDE) to check the templates for errors, and use the AWS CLI to validate that the templates are correct
- G. Adapt the deployment code to check for error conditions and generate notifications on error
- H. Deploy to a test environment and execute a manual test plan before approving the change for production.
- I. Use AWS CodeDeploy and a blue/green deployment pattern with CloudFormation to replace the user data deployment script
- J. Have the operators log in to running instances and go through a manual test plan to verify the application is running as expected.

Answer: B

Explanation:

<https://aws.amazon.com/blogs/devops/performing-bluegreen-deployments-with-aws-codedeploy-and-auto-scaling/> When one adopts go infrastructure as code, we need to test the infrastructure code as well via automated testing, and revert to original if things are not performing correctly.

NEW QUESTION 55

- (Exam Topic 1)

A company runs a popular web application in an on-premises data center. The application receives four million views weekly. The company expects traffic to increase by 200% because of an advertisement that will be published soon.

The company needs to decrease the load on the origin before the increase of traffic occurs. The company does not have enough time to move the entire application to the AWS Cloud.

Which solution will meet these requirements?

- A. Create an Amazon CloudFront content delivery network (CDN). Enable query forwarding to the origin. Create a managed cache policy that includes query string
- B. Use an on-premises load balancer as the origin
- C. Offload the DNS querying to AWS to handle CloudFront CDN traffic.
- D. Create an Amazon CloudFront content delivery network (CDN) that uses a Real Time Messaging Protocol (RTMP) distribution
- E. Enable query forwarding to the origin
- F. Use an on-premises load balancer as the origin
- G. Offload the DNS querying to AWS to handle CloudFront CDN traffic.
- H. Create an accelerator in AWS Global Accelerator
- I. Add listeners for HTTP and HTTPS TCP ports. Create an endpoint group
- J. Create a Network Load Balancer (NLB), and attach it to the endpoint group
- K. Point the NLB to the on-premises server
- L. Offload the DNS querying to AWS to handle AWS Global Accelerator traffic.
- M. Create an accelerator in AWS Global Accelerator
- N. Add listeners for HTTP and HTTPS TCP ports. Create an endpoint group

- O. Create an Application Load Balancer (ALB), and attach it to the endpoint group
- P. Point the ALB to the on-premises server
- Q. Offload the DNS querying to AWS to handle AWS Global Accelerator traffic.

Answer: D

NEW QUESTION 56

- (Exam Topic 1)

A company is planning on hosting its ecommerce platform on AWS using a multi-tier web application designed for a NoSQL database. The company plans to use the us-west-2 Region as its primary Region. The company wants to ensure that copies of the application and data are available in a second Region, us-west-1, for disaster recovery. The company wants to keep the time to fail over as low as possible. Failing back to the primary Region should be possible without administrative interaction after the primary service is restored.

Which design should the solutions architect use?

- A. Use AWS CloudFormation StackSets to create the stacks in both Regions with Auto Scaling groups for the web and application tier
- B. Asynchronously replicate static content between Regions using Amazon S3 cross-Region replication
- C. Use an Amazon Route 53 DNS failover routing policy to direct users to the secondary site in us-west-1 in the event of an outage
- D. Use Amazon DynamoDB global tables for the database tier.
- E. Use AWS CloudFormation StackSets to create the stacks in both Regions with Auto Scaling groups for the web and application tier
- F. Asynchronously replicate static content between Regions using Amazon S3 cross-Region replication
- G. Use an Amazon Route 53 DNS failover routing policy to direct users to the secondary site in us-west-1 in the event of an outage
- H. Deploy an Amazon Aurora global database for the database tier.
- I. Use AWS Service Catalog to deploy the web and application servers in both Regions
- J. Asynchronously replicate static content between the two Regions using Amazon S3 cross-Region replication
- K. Use Amazon Route 53 health checks to identify a primary Region failure and update the public DNS entry listing to the secondary Region in the event of an outage
- L. Use Amazon RDS for MySQL with cross-Region replication for the database tier.
- M. Use AWS CloudFormation StackSets to create the stacks in both Regions using Auto Scaling groups for the web and application tier
- N. Asynchronously replicate static content between Regions using Amazon S3 cross-Region replication
- O. Use Amazon CloudFront with static files in Amazon S3, and multi-Region origins for the front-end web tier
- P. Use Amazon DynamoDB tables in each Region with scheduled backups to Amazon S3.

Answer: A

NEW QUESTION 60

- (Exam Topic 1)

A company has a policy that all Amazon EC2 instances that are running a database must exist within the same subnets in a shared VPC. Administrators must follow security compliance requirements and are not allowed to directly log in to the shared account. All company accounts are members of the same organization in AWS Organizations. The number of accounts will rapidly increase as the company grows.

A solutions architect uses AWS Resource Access Manager to create a resource share in the shared account. What is the MOST operationally efficient configuration to meet these requirements?

- A. Add the VPC to the resource share
- B. Add the account IDs as principals
- C. Add all subnets within the VPC to the resource share
- D. Add the account IDs as principals
- E. Add all subnets within the VPC to the resource share
- F. Add the organization as a principal.
- G. Add the VPC to the resource share
- H. Add the organization as a principal

Answer: C

Explanation:

<https://docs.aws.amazon.com/ram/latest/userguide/getting-started-sharing.html#getting-started-sharing-create> To restrict resource sharing to only principals in your organization, choose Allow sharing with principals in your organization only.

<https://docs.aws.amazon.com/ram/latest/userguide/ram-ug.pdf>

NEW QUESTION 64

- (Exam Topic 1)

A company is migrating its three-tier web application from on-premises to the AWS Cloud. The company has the following requirements for the migration process:

- Ingest machine images from the on-premises environment.
- Synchronize changes from the on-premises environment to the AWS environment until the production cutover.
- Minimize downtime when executing the production cutover.
- Migrate the virtual machines' root volumes and data volumes.

Which solution will satisfy these requirements with minimal operational overhead?

- A. Use AWS Server Migration Service (SMS) to create and launch a replication job for each tier of the application
- B. Launch instances from the AMIs created by AWS SMS
- C. After initial testing, perform a final replication and create new instances from the updated AMIs.
- D. Create an AWS CLI Import/Export script to migrate each virtual machine
- E. Schedule the script to run incrementally to maintain changes in the application
- F. Launch instances from the AMIs created by VM Import/Export
- G. Once testing is done, rerun the script to do a final import and launch the instances from the AMIs.
- H. Use AWS Server Migration Service (SMS) to upload the operating system volume
- I. Use the AWS CLI import-snap command for the data volume
- J. Launch instances from the AMIs created by AWS SMS and attach the data volumes to the instances
- K. After initial testing, perform a final replication, launch new instances from the replicated AMI
- L. and attach the data volumes to the instances.
- M. Use AWS Application Discovery Service and AWS Migration Hub to group the virtual machines as an application

- N. Use the AWS CLI VM Import/Export script to import the virtual machines as AMI
- O. Schedule the script to run incrementally to maintain changes in the applicatio
- P. Launch instances from the AMI
- Q. After initial testing, perform a final virtual machine import and launch new instances from the AMIs.

Answer: A

Explanation:

SMS can handle migrating the data volumes:

<https://aws.amazon.com/about-aws/whats-new/2018/09/aws-server-migration-service-adds-support-for-migratin>

NEW QUESTION 68

- (Exam Topic 1)

A company is running a web application with On-Demand Amazon EC2 instances in Auto Scaling groups that scale dynamically based on custom metrics. After extensive testing, the company determines that the m5.2xlarge instance size is optimal for the workload. Application data is stored in db.r4.4xlarge Amazon RDS instances that are confirmed to be optimal. The traffic to the web application spikes randomly during the day.

What other cost-optimization methods should the company implement to further reduce costs without impacting the reliability of the application?

- A. Double the instance count in the Auto Scaling groups and reduce the instance size to m5.large
- B. Reserve capacity for the RDS database and the minimum number of EC2 instances that are constantly running.
- C. Reduce the RDS instance size to db.r4.xlarge and add five equivalentⁿ sized read replicas to provide reliability.
- D. Reserve capacity for all EC2 instances and leverage Spot Instance pricing for the RDS database.

Answer: B

Explanation:

People are being confused by the term 'reserve capacity'. This is not the same as an on-demand capacity reservation. This article by AWS clearly states that by 'reserving capacity' you are reserving the instances and reducing your costs. See <https://aws.amazon.com/aws-cost-management/aws-cost-optimization/reserved-instances/>

NEW QUESTION 72

- (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 VPC
- 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 77

- (Exam Topic 1)

A finance company hosts a data lake in Amazon S3. The company receives financial data records over SFTP each night from several third parties. The company runs its own SFTP server on an Amazon EC2 instance in a public subnet of a VPC. After the files are uploaded, they are moved to the data lake by a cron job that runs on the same instance. The SFTP server is reachable on DNS `sftp.example.com` through the use of Amazon Route 53.

What should a solutions architect do to improve the reliability and scalability of the SFTP solution?

- A. Move the EC2 instance into an Auto Scaling group
- B. Place the EC2 instance behind an Application Load Balancer (ALB). Update the DNS record `sftp.example.com` in Route 53 to point to the ALB.
- C. Migrate the SFTP server to AWS Transfer for SFT
- D. Update the DNS record `sftp.example.com` in Route 53 to point to the server endpoint hostname.
- E. Migrate the SFTP server to a file gateway in AWS Storage Gateway
- F. Update the DNS record `sftp.example.com` in Route 53 to point to the file gateway endpoint.
- G. Place the EC2 instance behind a Network Load Balancer (NLB). Update the DNS record `sftp.example.com` in Route 53 to point to the NLB.

Answer: B

NEW QUESTION 78

- (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 81

- (Exam Topic 1)

A solutions architect at a large company needs to set up network security for outbound traffic to the internet from all AWS accounts within an organization m AWS Organizations The organization has more than 100 AWS accounts, and the accounts route to each other by using a centralized AWS Transit Gateway. Each account has both an internet gateway and a NAT gateway for outbound traffic to the internet The company deploys resources only into a single AWS Region The company needs the ability to add centrally managed rule-based filtering on all outbound traffic to the internet for all AWS accounts in the organization The peak load of outbound traffic will not exceed 25 Gbps in each Availability Zone

Which solution meets these requirements?

- A. Create a new VPC for outbound traffic to the internet Connect the existing transit gateway to the new VPC Configure a new NAT gateway Create an Auto Scaling group of Amazon EC2 Instances that run an open-source internet proxy for rule-based filtering across all Availability Zones in the Region Modify all default routes to point to the proxy's Auto Scaling group
- B. Create a new VPC for outbound traffic to the internet Connect the existing transit gateway to the new VPC Configure a new NAT gateway Use an AWS Network Firewall firewall for rule-based filtering Create Network Firewall endpoints in each Availability Zone Modify all default routes to point to the Network Firewall endpoints
- C. Create an AWS Network Firewall firewall for rule-based filtering in each AWS account Modify all default routes to point to the Network Firewall firewalls in each account.
- D. In each AWS account, create an Auto Scaling group of network-optimized Amazon EC2 instances that run an open-source internet proxy for rule-based filtering Modify all default routes to point to the proxy's Auto Scaling group.

Answer: B

Explanation:

<https://aws.amazon.com/blogs/networking-and-content-delivery/deployment-models-for-aws-network-firewall/>

<https://aws.amazon.com/blogs/networking-and-content-delivery/deploy-centralized-traffic-filtering-using-aws-n>

NEW QUESTION 84

- (Exam Topic 1)

A company needs to implement a patching process for its servers. The on-premises servers and Amazon EC2 instances use a variety of tools to perform patching. Management requires a single report showing the patch status of all the servers and instances.

Which set of actions should a solutions architect take to meet these requirements?

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

Answer: A

Explanation:

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

NEW QUESTION 86

- (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 CloudFrontorigin 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 87

- (Exam Topic 2)

A life sciences company is using a combination of open source tools to manage data analysis workflows and Docker containers running on servers in its on-premises data center to process genomics data Sequencing data is generated and stored on a local storage area network (SAN), and then the data is processed. The research and development teams are running into capacity issues and have decided to re-architect their genomics analysis platform on AWS to scale based on workload demands and reduce the turnaround time from weeks to days

The company has a high-speed AWS Direct Connect connection Sequencers will generate around 200 GB of data for each genome, and individual jobs can take several hours to process the data with ideal compute capacity. The end result will be stored in Amazon S3. The company is expecting 10-15 job requests each day

Which solution meets these requirements?

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

Answer: C

NEW QUESTION 89

- (Exam Topic 2)

A company is migrating its marketing website and content management system from an on-premises data center to AWS. The company wants the AWS application to be deployed in a VPC with Amazon EC2 instances used for the web servers and an Amazon RDS instance for the database. The company has a runbook document that describes the installation process of the on-premises system. The company would like to base the AWS system on the processes referenced in the runbook document. The runbook document describes the installation and configuration of the operating systems, network settings, the website, and content management system software on the servers. After the migration is complete, the company wants to be able to make changes quickly to take advantage of other AWS features.

How can the application and environment be deployed and automated in AWS, while allowing for future changes?

- A. Update the runbook to describe how to create the VPC.
- B. Create the EC2 instances and the RDS instance for the application by using the AWS Console. Make sure that the rest of the steps in the runbook are updated to reflect any changes that may come from the AWS migration.
- C. Write a Python script that uses the AWS API to create the VPC.
- D. Create the EC2 instances and the RDS instance for the application. Write shell scripts that implement the rest of the steps in the runbook. Have the Python script copy and run the shell scripts on the newly created instances to complete the installation.
- E. Write an AWS CloudFormation template that creates the VPC, the EC2 instances, and the RDS instance for the application. Ensure that the rest of the steps in the runbook are updated to reflect any changes that may come from the AWS migration.
- F. Write an AWS CloudFormation template that creates the VPC, the EC2 instances, and the RDS instance for the application. Include EC2 user data in the AWS CloudFormation template to install and configure the software.

Answer: D

NEW QUESTION 91

- (Exam Topic 2)

A company hosts a blog post application on AWS using Amazon API Gateway, Amazon DynamoDB, and AWS Lambda. The application currently does not use API keys to authorize requests. The API model is as follows:

GET /posts/{postId} to get post details

GET /users/{userId} to get user details

GET /comments/{commentId} to get comments details

The company has noticed users are actively discussing topics in the comments section, and the company wants to increase user engagement by making the comments appear in real time.

Which design should be used to reduce comment latency and improve user experience?

- A. Use edge-optimized API with Amazon CloudFront to cache API responses.
- B. Modify the blog application code to request GET/commentsV{commentId} every 10 seconds.
- C. Use AWS AppSync and leverage WebSockets to deliver comments.
- D. Change the concurrency limit of the Lambda functions to lower the API response time.

Answer: C

NEW QUESTION 96

- (Exam Topic 2)

A company wants to allow its marketing team to perform SQL queries on customer records to identify market segments. The data is spread across hundreds of files. The records must be encrypted in transit and at rest. The team manager must have the ability to manage users and groups but no team members should have access to services or resources not required for the SQL queries. Additionally, administrators need to audit the queries made and receive notifications when a query violates rules defined by the security team.

AWS Organizations has been used to create a new account and an AWS IAM user with administrator permissions for the team manager. Which design meets these requirements?

- A. Apply a service control policy (SCP) that allows access to IAM, Amazon RDS, and AWS CloudTrail. Load customer records in Amazon RDS MySQL and train users to run queries using the AWS CLI.
- B. Stream the query logs to Amazon CloudWatch Logs from the RDS database instance. Use a subscription filter with AWS Lambda functions to audit and alarm on queries against personal data.
- C. Apply a service control policy (SCP) that denies access to all services except IAM, Amazon Athena, Amazon S3, and AWS CloudTrail. Store customer record files in Amazon S3 and train users to run queries using the CLI via Athena. Analyze CloudTrail events to audit and alarm on queries against personal data.
- D. Apply a service control policy (SCP) that denies access to all services except IAM, Amazon DynamoDB, and AWS CloudTrail. Store customer records in DynamoDB and train users to run queries using the AWS CLI. Enable DynamoDB streams to track the queries that are issued and use an AWS Lambda function for real-time monitoring and alerting.
- E. Apply a service control policy (SCP) that allows access to IAM, Amazon Athena, Amazon S3, and AWS CloudTrail. Store customer records as files in Amazon S3 and train users to leverage the Amazon S3 Select feature and run queries using the AWS CLI. Enable S3 object-level logging and analyze CloudTrail events to audit and alarm on queries against personal data.

Answer: B

NEW QUESTION 100

- (Exam Topic 2)

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

Which strategy should the solutions architect use?

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

Answer: C

NEW QUESTION 101

- (Exam Topic 2)

A company wants to migrate its website from an on-premises data center onto AWS. At the same time, it wants to migrate the website to a containerized microservice-based architecture to improve the availability and cost efficiency. The company's security policy states that privileges and network permissions must be configured according to best practice, using least privilege.

A solutions architect must create a containerized architecture that meets the security requirements and has deployed the application to an Amazon ECS cluster. What steps are required after the deployment to meet the requirements? (Select TWO.)

- A. Create tasks using the bridge network mode.
- B. Create tasks using the awsvpc network mode.
- C. Apply security groups to Amazon EC2 instances and use IAM roles for EC2 instances to access other resources.
- D. Apply security groups to the tasks, and pass IAM credentials into the container at launch time to access other resources.
- E. Apply security groups to the tasks; and use IAM roles for tasks to access other resources.

Answer: BE

NEW QUESTION 103

- (Exam Topic 2)

A company is running a three-tier web application in an on-premises data center. The frontend is served by an Apache web server, the middle tier is a monolithic Java application, and the storage tier is a PostgreSQL database.

During a recent marketing promotion, customers could not place orders through the application because the application crashed. An analysis showed that all three tiers were overloaded. The application became unresponsive, and the database reached its capacity limit because of read operations. The company already has several similar promotions scheduled in the near future.

A solutions architect must develop a plan for migration to AWS to resolve these issues. The solution must maximize scalability and must minimize operational effort.

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

- A. Refactor the frontend so that static assets can be hosted on Amazon S3. Use Amazon CloudFront to serve the frontend to customer.
- B. Connect the frontend to the Java application.
- C. Rehost the Apache web server of the frontend on Amazon EC2 instances that are in an Auto Scaling group.
- D. Use a load balancer in front of the Auto Scaling group.
- E. Use Amazon Elastic File System (Amazon EFS) to host the static assets that the Apache web server needs.
- F. Rehost the Java application in an AWS Elastic Beanstalk environment that includes auto scaling.
- G. Refactor the Java application.
- H. Develop a Docker container to run the Java application.
- I. Use AWS Fargate to host the container.
- J. Use AWS Database Migration Service (AWS DMS) to replatform the PostgreSQL database to an Amazon Aurora PostgreSQL database.
- K. Use Aurora Auto Scaling for read replicas.
- L. Rehost the PostgreSQL database on an Amazon EC2 instance that has twice as much memory as the on-premises server.

Answer: BCF

NEW QUESTION 104

- (Exam Topic 2)

A large company has many business units. Each business unit has multiple AWS accounts for different purposes. The CIO of the company sees that each business unit has data that would be useful to share with other parts of the company. In total, there are about 10 PB of data that needs to be shared with users in 1,000 AWS accounts. The data is proprietary, so some of it should only be available to users with specific job types. Some of the data is used for throughput-intensive workloads such as simulations. The number of accounts changes frequently because of new initiatives, acquisitions, and divestitures.

A solutions architect has been asked to design a system that will allow for sharing data for use in AWS with all of the employees in the company.

Which approach will allow for secure data sharing in a scalable way?

- A. Store the data in a single Amazon S3 bucket. Create an IAM role for every combination of job type and business unit that allows for appropriate read/write access based on object prefixes in the S3 bucket. The roles should have trust policies that allow the business unit's AWS accounts to assume their roles. Use IAM in each business unit's AWS account to prevent them from assuming roles for a different job type. Users get credentials to access the data by using AssumeRole from their business unit's AWS account. Users can then use those credentials with an S3 client.
- B. Store the data in a single Amazon S3 bucket. Write a bucket policy that uses conditions to grant read and write access where appropriate based on each user's business unit and job type.
- C. Determine the business unit with the AWS account accessing the bucket and the job type with a prefix in the IAM user's name. Users can access data by using IAM credentials from their business unit's AWS account with an S3 client.
- D. Store the data in a series of Amazon S3 buckets. Create an application running on Amazon EC2 that is integrated with the company's identity provider (IdP) that authenticates users and allows them to download or upload data through the application. The application uses the business unit and job type information in the IdP to control what users can upload and download through the application. The users can access the data through the application's API.
- E. Store the data in a series of Amazon S3 buckets. Create an AWS STS token vending machine that is integrated with the company's identity provider (IdP). When a user logs in, have the token vending machine attach an IAM policy that assumes the role that limits the user's access and/or upload only the data the user is authorized to access. Users can get credentials by authenticating to the token vending machine's website or API and then use those credentials with an S3 client.
- F. D

Answer: E

NEW QUESTION 107

- (Exam Topic 2)

A development team is deploying new APIs as serverless applications within a company. The team is currently using the AWS Management Console to provision Amazon API Gateway, AWS Lambda, and Amazon DynamoDB resources. A solutions architect has been tasked with automating the future deployments of these serverless APIs.

How can this be accomplished?

- A. Use AWS CloudFormation with a Lambda-backed custom resource to provision API Gateway. Use the `MyDynamoDB::Table` and `AWS::Lambda::Function` resources to create the Amazon DynamoDB table and Lambda functions. Write a script to automate the deployment of the CloudFormation template.
- B. Use the AWS Serverless Application Model to define the resources. Upload a YAML template and application files to the code repository. Use AWS CodePipeline to connect to the code repository and to create an action to build using AWS CodeBuild.
- C. Use the AWS CloudFormation deployment provider in CodePipeline to deploy the solution.
- D. Use AWS CloudFormation to define the serverless application.
- E. Implement versioning on the Lambda functions and create aliases to point to the version.
- F. When deploying, configure weights to implement shifting traffic to the newest version, and gradually update the weights as traffic moves over.
- G. Commit the application code to the AWS CodeCommit code repository.
- H. Use AWS CodePipeline and connect to the CodeCommit code repository. Use AWS CodeBuild to build and deploy the Lambda functions using AWS CodeDeploy. Specify the deployment preference type in CodeDeploy to gradually shift traffic over to the new version.

Answer: B

NEW QUESTION 112

- (Exam Topic 2)

A company uses multiple AWS accounts in a single AWS Region. A solutions architect is designing a solution to consolidate logs generated by Elastic Load Balancers (ELBs) in the AppDev, AppTest, and AppProd accounts. The logs should be stored in an existing Amazon S3 bucket named `s3-elb-logs` in the central AWS account. The central account is used for log consolidation only and does not have ELBs deployed. ELB logs must be encrypted at rest. Which combination of steps should the solutions architect take to build the solution? (Select TWO)

- A. Update the S3 bucket policy for the `s3-elb-logs` bucket to allow the `s3:PutBucketLogging` action for the central AWS account ID.
- B. Update the S3 bucket policy for the `s3-elb-logs` bucket to allow the `s3:PutObject` and `s3:DeleteObject` actions for the AppDev, AppTest, and AppProd account IDs.
- C. Update the S3 bucket policy for the `s3-elb-logs` bucket to allow the `s3:PutObject` action for the AppDev, AppTest, and AppProd account IDs.
- D. Enable access logging for the ELB.
- E. Set the S3 location to the `s3-elb-logs` bucket.
- F. Enable Amazon S3 default encryption using server-side encryption with S3 managed encryption keys (SSE-S3) for the `s3-elb-logs` S3 bucket.

Answer: AE

NEW QUESTION 115

- (Exam Topic 2)

A company is planning to migrate an application from on-premises to the AWS Cloud. The company will begin the migration by moving the application's underlying data storage to AWS. The application data is stored on a shared file system on-premises, and the application servers connect to the shared file system through SMB.

A solutions architect must implement a solution that uses an Amazon S3 bucket for shared storage. Until the application is fully migrated and code is rewritten to use native Amazon S3 APIs, the application must continue to have access to the data through SMB. The solutions architect must migrate the application data to AWS to its new location while still allowing the on-premises application to access the data.

Which solution will meet these requirements?

- A. Create a new Amazon FSx for Windows File System file system. Configure AWS DataSync with one location for the on-premises file share and one location for the new Amazon FSx file system. Create a new DataSync task to copy the data from the on-premises file share location to the Amazon FSx file system.
- B. Create an S3 bucket for the application.
- C. Copy the data from the on-premises storage to the S3 bucket.
- D. Deploy an AWS Server Migration Service (AWS SMS) VM to the on-premises environment.
- E. Use AWS SMS to migrate the file storage server from on-premises to an Amazon EC2 instance.
- F. Create an S3 bucket for the application.
- G. Deploy a new AWS Storage Gateway File gateway on on-premises.
- H. Create a new file share that stores data in the S3 bucket and is associated with the file gateway.
- I. Copy the data from the on-premises storage to the new file gateway endpoint.

Answer: A

NEW QUESTION 120

- (Exam Topic 2)

A company is using a lift-and-shift strategy to migrate applications from several on-premises Windows servers to AWS. The Windows servers will be hosted on Amazon EC2 instances in the `us-east-1` Region.

The company's security policy allows the installation of migration tools on servers. The migration data must be encrypted in transit and encrypted at rest. The applications are business critical. The company needs to minimize the cutover window and minimize the downtime that results from the migration. The company wants to use Amazon CloudWatch and AWS CloudTrail for monitoring.

Which solution will meet these requirements?

- A. Use AWS Application Migration Service (CloudEndure Migration) to migrate the Windows servers to AWS.
- B. Create a Replication Settings template.
- C. Install the AWS Replication Agent on the source servers.
- D. Use AWS DataSync to migrate the Windows servers to AWS.
- E. Install the DataSync agent on the source server.
- F. Configure a blueprint for the target server.
- G. Begin the replication process.
- H. Use AWS Server Migration Service (AWS SMS) to migrate the Windows servers to AWS.

- I. Install the SMS Connector on the source server
- J. Replicate the source servers to AW
- K. Convert the replicated volumes to AMIs to launch EC2 instances.
- L. Use AWS Migration Hub to migrate the Windows servers to AW
- M. Create a project in Migration Hub. Track the progress of server migration by using the built-in dashboard.

Answer: A

NEW QUESTION 121

- (Exam Topic 2)

A solutions architect has been assigned to migrate a 50 TB Oracle data warehouse that contains sales data from on-premises to Amazon Redshift. Major updates to the sales data occur on the final calendar day of the month. For the remainder of the month, the data warehouse only receives minor daily updates and is primarily used for reading and reporting. Because of this, the migration process must start on the first day of the month and must be complete before the next set of updates occur. This provides approximately 30 days to complete the migration and ensure that the minor daily changes have been synchronized with the Amazon Redshift data warehouse. Because the migration cannot impact normal business network operations, the bandwidth allocated to the migration for moving data over the internet is 50 Mbps. The company wants to keep data migration costs low.

Which steps will allow the solutions architect to perform the migration within the specified timeline?

- A. Install Oracle database software on an Amazon EC2 instance. Configure VPN connectivity between AWS and the company's data center. Configure the Oracle database running on Amazon EC2 to join the Oracle Real Application Clusters (RAC). When the Oracle database on Amazon EC2 finishes synchronizing, create an AWS DMS ongoing replication task to migrate the data from the Oracle database on Amazon EC2 to Amazon Redshift. Verify the data migration is complete and perform the cut over to Amazon Redshift.
- B. Create an AWS Snowball import job. Export a backup of the Oracle data warehouse. Copy the exported data to the Snowball device. Return the Snowball device to AWS. Create an Amazon RDS for Oracle database and restore the backup file to that RDS instance. Create an AWS DMS task to migrate the data from the RDS for Oracle database to Amazon Redshift. Copy daily incremental backups from Oracle in the data center to the RDS for Oracle database over the internet. Verify the data migration is complete and perform the cut over to Amazon Redshift.
- C. Install Oracle database software on an Amazon EC2 instance. To minimize the migration time, configure VPN connectivity between AWS and the company's data center by provisioning a 1 Gbps AWS Direct Connect connection. Configure the Oracle database running on Amazon EC2 to be a read replica of the data center Oracle database. Start the synchronization process between the company's on-premises data center and the Oracle database on Amazon EC2. When the Oracle database on Amazon EC2 is synchronized with the on-premises database, create an AWS DMS ongoing replication task from the Oracle database read replica that is running on Amazon EC2 to Amazon Redshift. Verify the data migration is complete and perform the cut over to Amazon Redshift.
- D. Create an AWS Snowball import job.
- E. Configure a server in the company's data center with an extraction agent.
- F. Use AWS SCT to manage the extraction agent and convert the Oracle schema to an Amazon Redshift schema.
- G. Create a new project in AWS SCT using the registered data extraction agent.
- H. Create a local task and an AWS DMS task in AWS SCT with replication of ongoing changes.
- I. Copy data to the Snowball device and return the Snowball device to AWS.
- J. Allow AWS DMS to copy data from Amazon S3 to Amazon Redshift.
- K. Verify that the data migration is complete and perform the cut over to Amazon Redshift.

Answer: D

Explanation:

Create an AWS Snowball import job. Configure a server in the company's data center with an extraction agent. Use AWS SCT to manage the extraction agent and convert the Oracle schema to an Amazon Redshift schema. Create a new project in AWS SCT using the registered data extraction agent. Create a local task and an AWS DMS task in AWS SCT with replication of ongoing changes. Copy data to the Snowball device and return the Snowball device to AWS. Allow AWS DMS to copy data from Amazon S3 to Amazon Redshift. Verify that the data migration is complete and perform the cut over to Amazon Redshift.

<https://aws.amazon.com/getting-started/hands-on/migrate-oracle-to-amazon-redshift/>

NEW QUESTION 126

- (Exam Topic 2)

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

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

Answer: B

NEW QUESTION 131

- (Exam Topic 2)

A company needs to create a centralized logging architecture for all of its AWS accounts. The architecture should provide near-real-time data analysis for all AWS CloudTrail logs and VPC Flow logs across all AWS accounts. The company plans to use Amazon Elasticsearch Service (Amazon ES) to perform log analyses in the logging account.

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

- A. Configure CloudTrail and VPC Flow Logs in each AWS account to send data to a centralized Amazon S3 bucket in the logging account.
- B. Create an AWS Lambda function to load data from the S3 bucket to Amazon ES in the logging account.
- C. Configure CloudTrail and VPC Flow Logs to send data to a log group in Amazon CloudWatch Logs in each AWS account. Configure a CloudWatch subscription filter in each AWS account to send data to Amazon Kinesis Data Firehose in the logging account. Load data from Kinesis Data Firehose into Amazon ES in the logging account.
- D. Configure CloudTrail and VPC Flow Logs to send data to a separate Amazon S3 bucket in each AWS account.
- E. Create an AWS Lambda function triggered by S3 events to copy the data to a centralized logging bucket.
- F. Create another Lambda function to load data from the S3 bucket to Amazon ES in the logging account.
- G. Configure CloudTrail and VPC Flow Logs to send data to a log group in Amazon CloudWatch Logs in each AWS account. Create AWS Lambda functions in each AWS account to subscribe to the log groups and stream the data to an Amazon S3 bucket in the logging account.

H. Create another Lambda function to load data from the S3 bucket to Amazon ES in the logging account.

Answer: A

NEW QUESTION 132

- (Exam Topic 2)

A company is running multiple workloads in the AWS Cloud. The company has separate units for software development. The company uses AWS Organizations and federation with SAML to give permissions to developers to manage resources in their AWS accounts. The development units each deploy their production workloads into a common production account.

Recently, an incident occurred in the production account in which members of a development unit terminated an EC2 instance that belonged to a different development unit. A solutions architect must create a solution that prevents a similar incident from happening in the future. The solution also must allow developers the possibility to manage the instances used for their workloads.

Which strategy will meet these requirements?

- A. Create separate OUs in AWS Organizations for each development unit. Assign the created OUs to the company AWS accounts. Create separate SCPs with a deny action and a StringNotEquals condition for the DevelopmentUnit resource tag that matches the development unit name. Assign the SCP to the corresponding OU.
- B. Pass an attribute for DevelopmentUnit as an AWS Security Token Service (AWS STS) session tag during SAML federation. Update the IAM policy for the developers' assumed IAM role with a deny action and a StringNotEquals condition for the DevelopmentUnit resource tag and aws:PrincipalTag/DevelopmentUnit.
- C. Pass an attribute for DevelopmentUnit as an AWS Security Token Service (AWS STS) session tag during SAML federation. Create an SCP with an allow action and a StringEquals condition for the DevelopmentUnit resource tag and aws:PrincipalTag/DevelopmentUnit. Assign the SCP to the root OU.
- D. Create separate IAM policies for each development unit. For every IAM policy, add an allow action and a StringEquals condition for the DevelopmentUnit resource tag and the development unit name. During SAML federation, use AWS Security Token Service (AWS STS) to assign the IAM policy and match the development unit name to the assumed IAM role.

Answer: A

NEW QUESTION 134

- (Exam Topic 2)

A company owns a chain of travel agencies and is running an application in the AWS Cloud. Company employees use the application to search (or information) about travel destinations. Destination content is updated four times each year.

Two fixed Amazon EC2 instances serve the application. The company uses an Amazon Route 53 public hosted zone with a multivalue record of travel.example.com that returns the Elastic IP addresses for the EC2 instances. The application uses Amazon DynamoDB as its primary data store. The company uses a self-hosted Redis instance as a caching solution.

During content updates, the load on the EC2 instances and the caching solution increases drastically. This increased load has led to downtime on several occasions. A solutions architect must update the application so that the application is highly available and can handle the load that is generated by the content updates.

Which solution will meet these requirements?

- A. Set up DynamoDB Accelerator (DAX) as in-memory cache.
- B. Update the application to use DAX.
- C. Create an Auto Scaling group for the EC2 instance.
- D. Create an Application Load Balancer (ALB). Set the Auto Scaling group as a target for the ALB.
- E. Update the Route 53 record to use a simple routing policy that targets the ALB's DNS alias.
- F. Configure scheduled scaling for the EC2 instances before the content updates.
- G. Set up Amazon ElastiCache for Redis.
- H. Update the application to use ElastiCache.
- I. Create an Auto Scaling group for the EC2 instance.
- J. Create an Amazon CloudFront distribution.
- K. and set the Auto Scaling group as an origin for the distribution.
- L. Update the Route 53 record to use a simple routing policy that targets the CloudFront distribution's DNS alias. Manually scale up EC2 instances before the content updates.
- M. Set up Amazon ElastiCache for Memcached.
- N. Update the application to use ElastiCache.
- O. Create an Auto Scaling group for the EC2 instances. Create an Application Load Balancer (ALB). Set the Auto Scaling group as a target for the ALB.
- P. Update the Route 53 record to use a simple routing policy that targets the ALB's DNS alias.
- Q. Configure scheduled scaling for the application before the content updates.
- R. Set up DynamoDB Accelerator (DAX) as in-memory cache.
- S. Update the application to use DAX.
- T. Create an Auto Scaling group for the EC2 instance.
- . Create an Amazon CloudFront distribution, and set the Auto Scaling group as an origin for the distribution.
- . Update the Route 53 record to use a simple routing policy that targets the CloudFront distribution's DNS alias.
- . Manually scale up EC2 instances before the content updates.

Answer: B

NEW QUESTION 137

- (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 leaderboard data in a different table.
- C. Use Apache HiveQL JOIN statements to build the leaderboard.
- D. Keep the leaderboard data in the RDS database 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 138

- (Exam Topic 2)

A solutions architect is migrating an existing workload to AWS Fargate. The task can only run in a private subnet within the VPC where there is no direct connectivity from outside the system to the application. When the Fargate task is launched, the task fails with the following error:

```
CannotPullContainerError: API error (500): Get https://111122223333.dkr.ecr.us-east-1.amazonaws.com/v2/: net/http: request canceled  
While waiting for connection
```

How should the solutions architect correct this error?

- A. Ensure the task is set to ENABLED for the auto-assign public IP setting when launching the task
- B. Ensure the task is set to DISABLED (or the auto-assign public IP setting when launching the task) Configure a NAT gateway in the public subnet in the VPC to route requests to the internet
- C. Ensure the task is set to DISABLED for the auto-assign public IP setting when launching the task. Configure a NAT gateway in the private subnet in the VPC to route requests to the internet
- D. Ensure the network mode is set to bridge in the Fargate task definition

Answer: B

NEW QUESTION 140

- (Exam Topic 2)

A company runs a software-as-a-service (SaaS) application on AWS. The application consists of AWS Lambda function and an Amazon RDS for MySQL Multi-AZ database. During market events, the application has a much higher workload than normal. Users notice slow response times during the peak periods because of many database connections. The company needs to improve the scalable performance and availability of the database.

Which solution meets these requirements?

- A. Create an Amazon CloudWatch alarm action that triggers a Lambda function to add an Amazon RDS for MySQL read replica when resource utilization hits a threshold.
- B. Migrate the database to Amazon Aurora and add a read replica. Add a database connection pool outside of the Lambda handler function.
- C. Migrate the database to Amazon Aurora and add a read replica.
- D. Use Amazon Route 53 weighted records.
- E. Migrate the database to Amazon Aurora and add an Aurora Replica.
- F. Configure Amazon RDS Proxy to manage database connection pools.

Answer: D

NEW QUESTION 144

- (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 147

- (Exam Topic 2)

A company is configuring connectivity to a multi-account AWS environment to support application workloads that serve users in a single geographic region. The workloads depend on a highly available, on-premises legacy system deployed across two locations. It is critical for the AWS workloads to maintain connectivity to the legacy system, and a minimum of 5 Gbps of bandwidth is required. All application workloads within AWS must have connectivity with one another.

Which solution will meet these requirements?

- A. Configure multiple AWS Direct Connect (DX) 10 Gbps dedicated connections from a DX partner for each on-premises location. Create private virtual interfaces on each connection for each AWS account VPC. Associate each private virtual interface with a virtual private gateway attached to each VPC.
- B. Configure multiple AWS Direct Connect (DX) 10 Gbps dedicated connections from two DX partners for each on-premises location. Create and attach a virtual private gateway for each AWS account VPC.
- C. Create a DX gateway in a central network account and associate it with the virtual private gateways. Create a public virtual interface on each DX connection and associate the interface with the DX gateway.
- D. Configure multiple AWS Direct Connect (DX) 10 Gbps dedicated connections from two DX partners for each on-premises location. Create a transit gateway and a DX gateway in a central network account.
- E. Create a transit virtual interface for each DX interface and associate them with the DX gateway.
- F. Create a gateway association between the DX gateway and the transit gateway.

- G. Configure multiple AWS Direct Connect (DX) 10 Gbps dedicated connections from a DX partner for each on-premises location Create and attach a virtual private gateway for each AWS account VP
- H. Create a transit gateway in a central network account and associate It with the virtual private gateways Create a transit virtual interface on each DX connection and attach the interface to the transit gateway.

Answer: B

NEW QUESTION 149

- (Exam Topic 2)

A company runs applications on Amazon EC2 instances. The company plans to begin using an Auto Scaling group for the instances. As part of this transition, a solutions architect must ensure that Amazon CloudWatch Logs automatically collects logs from all new instances The new Auto Scaling group will use a launch template that includes the Amazon Linux 2 AMI and no key pair

Which solution meets these requirements?

- A. Create an Amazon CloudWatch agent configuration for the workload Store the CloudWatch agent configuration in an Amazon S3 bucket Write an EC2 user data script to fetch the configuration He from Amazon S3. Configure the cloudWatch agent on the instance during Initial boot.
- B. Create an Amazon CloudWatch agent configuration for the workload In AWS Systems Manager Parameter Store Create a Systems Manager document that Installs and configures the CloudWatch agent by using the configuration Create an Amazon EventBridge (Amazon CloudWatch Events) rule on the default event bus with a Systems Manager Run Command target that runs the document whenever an instance enters the running state.
- C. Create an Amazon CloudWatch agent configuration for the workload Create an AWS Lambda function to Install and configure CloudWatch agent by using AWS Systems Manager Session Manage
- D. Include the agent configuration inside the Lambda package Create an AWS Config custom rule to identify changes to the EC2 instances and invoke the Lambda function
- E. Create an Amazon CloudWatch agent configuration for the workloa
- F. Save the CloudWatch agent configuration as pan of an AWS Lambda deployment packag
- G. Use AWS CloudTrail to capture EC2 tagging events and initiate agent installatio
- H. Use AWS CodeBuild to configure the CloudWatch agent on the instances that run the workload.

Answer: B

NEW QUESTION 150

- (Exam Topic 2)

A company's CI SO has asked a solutions architect to re-engineer the company's current CI/CD practices to make sure patch deployments to its application can happen as quickly as possible with minimal downtime if vulnerabilities are discovered The company must also be able to quickly roll back a change in case of errors.

The web application is deployed in a fleet of Amazon EC2 instances behind an Application Load Balancer The company is currently using GitHub to host the application source code. and has configured an AWS CodeBuild project to build the application The company also intends to use AWS CodePipeline to trigger builds from GitHub commits using the existing CodeBuild project.

What CI/CD configuration meets all of the requirements?

- A. Configure CodePipeline with a deploy stage using AWS CodeDeploy configured for in-place deployment Monitor the newly deployed code, and, if there are any issues, push another code update
- B. Configure CodePipeline with a deploy stage using AWS CodeDeploy configured for blue/green deployments Monitor the newly deployed code and if there are any issues, trigger a manual rollback using CodeDeploy
- C. Configure CodePipeline with a deploy stage using AWS CloudFormation to create a pipeline for test and production stacks Monitor the newly deployed code, and, if there are any issues, push another code update
- D. Configure the CodePipeline with a deploy stage using AWS OpsWorks and m-place deployments Monitor the newly deployed code an
- E. if there are any issues, push another code update

Answer: B

NEW QUESTION 151

- (Exam Topic 2)

A company runs many workloads on AWS and uses AWS Organizations to manage its accounts. The workloads are hosted on Amazon EC2, AWS Fargate, and AWS Lambda. Some of the workloads have unpredictable demand. Accounts record high usage in some months and low usage in other months.

The company wants to optimize its compute costs over the next 3 years. A solutions architect obtains a 6-month average for each of the accounts across the organization to calculate usage.

Which solution will provide the MOST cost savings for all the organization's compute usage?

- A. Purchase Reserved Instances for the organization to match the size and number of the most common EC2 instances from the member accounts.
- B. Purchase a Compute Savings Plan for the organization from the management account by using the recommendation at the management account level.
- C. Purchase Reserved Instances for each member account that had high EC2 usage according to the data from the last 6 months.
- D. Purchase an EC2 Instance Savings Plan for each member account from the management account based on EC2 usage data from the last 6 months.

Answer: A

NEW QUESTION 155

- (Exam Topic 2)

A company wants to migrate its on-premises data center to the AWS Cloud. This includes thousands of virtualized Linux and Microsoft Windows servers SAN storage, Java and PHP applications with MySQL, and Oracle databases. There are many dependent services hosted either in the same data center or externally. The technical documentation is incomplete and outdated A solutions architect needs to understand the current environment and estimate the cloud resource costs after the migration

Which tools or services should the solutions architect use to plan the cloud migration? (Select THREE.)

- A. AWS Application Discovery Service
- B. AWS SMS
- C. AWS X-Ray
- D. AWS Cloud Adoption Readiness Tool (CART)
- E. Amazon Inspector

F. AWS Migration Hub

Answer: ADF

NEW QUESTION 160

- (Exam Topic 2)

A company is using multiple AWS accounts. The DNS records are stored in a private hosted zone for Amazon Route 53 in Account A. The company's applications and databases are running in Account B.

A solutions architect will deploy a two-net application in a new VPC. To simplify the configuration, the db.example.com CNAME record set for the Amazon RDS endpoint was created in a private hosted zone for Amazon Route 53.

During deployment, the application failed to start. Troubleshooting revealed that db.example.com is not resolvable on the Amazon EC2 instance. The solutions architect confirmed that the record set was created correctly in Route 53.

Which combination of steps should the solutions architect take to resolve this issue? (Select TWO)

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

Answer: CE

NEW QUESTION 164

- (Exam Topic 2)

A company runs its application in the eu-west-1 Region and has one account for each of its environments: development, testing, and production. All the environments are running 24 hours a day, 7 days a week by using stateful Amazon EC2 instances and Amazon RDS for MySQL databases. The databases are between 500 GB and 800 GB in size.

The development team and testing team work on business days during business hours, but the production environment operates 24 hours a day, 7 days a week. The company wants to reduce costs. All resources are tagged with an environment tag with either development, testing, or production as the key.

What should a solutions architect do to reduce costs with the LEAST operational effort?

- A. Create an Amazon EventBridge (Amazon CloudWatch Events) rule that runs once every day. Configure the rule to invoke one AWS Lambda function that starts or stops instances based on the tag, day, and time.
- B. Create an Amazon EventBridge (Amazon CloudWatch Events) rule that runs every business day in the evening.
- C. Configure the rule to invoke an AWS Lambda function that stops instances based on the tag. Create a second EventBridge (CloudWatch Events) rule that runs every business day in the morning. Configure the second rule to invoke another Lambda function that starts instances based on the tag.
- D. Create an Amazon EventBridge (Amazon CloudWatch Events) rule that runs every business day in the evening. Configure the rule to invoke an AWS Lambda function that terminates instances based on the tag. Create a second EventBridge (CloudWatch Events) rule that runs every business day in the morning. Configure the second rule to invoke another Lambda function that restores the instances from their last backup based on the tag.
- E. Create an Amazon EventBridge (Amazon CloudWatch Events) rule that runs every hour. Configure the rule to invoke one AWS Lambda function that terminates or restores instances from their last backup based on the tag, day, and time.
- F. day, and time

Answer: C

NEW QUESTION 165

- (Exam Topic 2)

A company is running an application in the AWS Cloud. The application uses AWS Lambda functions and Amazon Elastic Container Service (Amazon ECS) containers that run with AWS Fargate technology as its primary compute. The load on the application is irregular. The application experiences long periods of no usage, followed by sudden and significant increases and decreases in traffic. The application is write-heavy and stores data in an Amazon Aurora MySQL database. The database runs on an Amazon RDS memory optimized D8 instance that is not able to handle the load.

What is the MOST cost-effective way for the company to handle the sudden and significant changes in traffic?

- A. Add additional read replicas to the database.
- B. Purchase Instance Savings Plans and RDS Reserved Instances.
- C. Migrate the database to an Aurora multi-master DB cluster.
- D. Purchase Instance Savings Plans.
- E. Migrate the database to an Aurora global database. Purchase Compute Savings Plans and RDS Reserved Instances.
- F. Migrate the database to Aurora Serverless v1. Purchase Compute Savings Plans.

Answer: D

NEW QUESTION 170

- (Exam Topic 2)

A company deploys a new web application. As part of the setup, the company configures AWS WAF to log to Amazon S3 through Amazon Kinesis Data Firehose. The company develops an Amazon Athena query that runs once daily to return AWS WAF log data from the previous 24 hours. The volume of daily logs is constant. However, over time, the same query is taking more time to run.

A solutions architect needs to design a solution to prevent the query time from continuing to increase. The solution must minimize operational overhead. Which solution will meet these requirements?

- A. Create an AWS Lambda function that consolidates each day's AWS WAF logs into one log file.
- B. Reduce the amount of data scanned by configuring AWS WAF to send logs to a different S3 bucket each day.
- C. Update the Kinesis Data Firehose configuration to partition the data in Amazon S3 by date and time. Create external tables for Amazon Redshift. Configure Amazon Redshift Spectrum to query the data source.
- D. Modify the Kinesis Data Firehose configuration and Athena table definition to partition the data by date and time.
- E. Change the Athena query to view the relevant partitions.

Answer: D

NEW QUESTION 172

- (Exam Topic 2)

A company implements a containerized application by using Amazon Elastic Container Service (Amazon ECS) and Amazon API Gateway. The application data is stored in Amazon Aurora databases and Amazon DynamoDB databases. The company automates infrastructure provisioning by using AWS CloudFormation. The company automates application deployment by using AWS CodePipeline.

A solutions architect needs to implement a disaster recovery (DR) strategy that meets an RPO of 2 hours and an RTO of 4 hours.

Which solution will meet these requirements MOST cost-effectively?

- A. Set up an Aurora global database and DynamoDB global tables to replicate the databases to a secondary AWS Region
- B. In the primary Region and in the secondary Region, configure an API Gateway API with a Regional endpoint
- C. Implement Amazon CloudFront with origin failover to route traffic to the secondary Region during a DR scenario.
- D. Use AWS Database Migration Service (AWS DMS), Amazon EventBridge (Amazon CloudWatch Events), and AWS Lambda to replicate the Aurora databases to a secondary AWS Region
- E. Use DynamoDB Stream
- F. EventBridge (CloudWatch Events), and Lambda to replicate the DynamoDB databases to the secondary Region
- G. In the primary Region and In the secondary Region, configure an API Gateway API with a Regional endpoint
- H. Implement Amazon Route 53 failover routing to switch traffic from the primary Region to the secondary Region.
- I. Use AWS Backup to create backups of the Aurora databases and the DynamoDB databases in a secondary AWS Region
- J. In the primary Region and in the secondary Region, configure an API Gateway API with a Regional endpoint
- K. Implement Amazon Route 53 failover routing to switch traffic from the primary Region to the secondary Region
- L. Set up an Aurora global database and DynamoDB global tables to replicate the databases to a secondary AWS Region
- M. In the primary Region and in the secondary Region, configure an API Gateway API with a Regional endpoint
- N. Implement Amazon Route 53 failover routing to switch traffic from the primary Region to the secondary Region.

Answer: C

NEW QUESTION 177

- (Exam Topic 2)

An ecommerce company runs its infrastructure on AWS. The company exposes its APIs to its web and mobile clients through an Application Load Balancer (ALB) in front of an Amazon Elastic Kubernetes Service (Amazon EKS) cluster. The EKS cluster runs thousands of pods that provide the APIs.

After extending delivery to a new continent, the company adds an Amazon CloudFront distribution and sets the ALB as the origin. The company also adds AWS WAF to its architecture.

After implementation of the new architecture, API calls are significantly slower. However, there is a sudden increase in HTTP status code 504 (Gateway Timeout) errors and HTTP status code 502 (Bad Gateway) errors. This increase in errors seems to be for a specific domain. Which factors could be a cause of these errors? (Select TWO.)

- A. AWS WAF is blocking suspicious requests.
- B. The origin is not properly configured in CloudFront.
- C. There is an SSL/TLS handshake issue between CloudFront and the origin.
- D. EKS Kubernetes pods are being cycled.
- E. Some pods are taking more than 30 seconds to answer API calls.

Answer: AE

NEW QUESTION 181

- (Exam Topic 2)

A company has used infrastructure as code (IaC) to provision a set of two Amazon EC2 instances. The instances have remained the same for several years.

The company's business has grown rapidly in the past few months. In response, the company's operations team has implemented an Auto Scaling group to manage the sudden increases in traffic. Company policy requires a monthly installation of security updates on all operating systems that are running.

The most recent security update required a reboot. As a result the Auto Scaling group terminated the instances and replaced them with new, unpatched instances.

Which combination of steps should a solutions architect recommend to avoid a recurrence of this issue? (Select TWO)

- A. Modify the Auto Scaling group by setting the Update policy to target the oldest launch configuration for replacement.
- B. Create a new Auto Scaling group before the next patch maintenance. During the maintenance window patch both groups and reboot the instances.
- C. Create an Elastic Load Balancer in front of the Auto Scaling group. Configure monitoring to ensure that target group health checks return healthy after the Auto Scaling group replaces the terminated instances.
- D. Create automation scripts to patch an AMI.
- E. Update the launch configuration, and invoke an Auto Scaling instance refresh.
- F. Create an Elastic Load Balancer in front of the Auto Scaling group. Configure termination protection on the instances.

Answer: AC

NEW QUESTION 184

- (Exam Topic 2)

A media company has a 30-TB repository of digital news videos. These videos are stored on tape in an on-premises tape library and referenced by a Media Asset Management (MAM) system. The company wants to enrich the metadata for these videos in an automated fashion and put them into a searchable catalog by using a MAM feature. The company must be able to search based on information in the video such as objects, scenery, items, or people's faces. A catalog is available that contains faces of people who have appeared in the videos that include an image of each person. The company would like to migrate these videos to AWS.

The company has a high-speed AWS Direct Connect connection with AWS and would like to move the MAM solution video content directly from its current file system.

How can these requirements be met by using the LEAST amount of ongoing management overhead and causing MINIMAL disruption to the existing system?

- A. Set up an AWS Storage Gateway file gateway appliance on-premise.
- B. Use the MAM solution to extract the videos from the current archive and push them into the file gateway. Use the catalog of faces to build a collection in Amazon Rekognition. Build an AWS Lambda function that invokes the Rekognition Javascript SDK to have Rekognition pull the video from the Amazon S3 files backing the file gateway, retrieve the required metadata, and push the metadata into the MAM solution.
- C. Set up an AWS Storage Gateway tape gateway appliance on-premises. Use the MAM solution to extract the videos from the current archive and push them into the tape gateway. Use the catalog of faces to build a collection in Amazon Rekognition. Build an AWS Lambda function that invokes the Rekognition Javascript SDK.

to have Amazon Rekognition process the video in the tape gateway retrieve the required metadata, and push the metadata into the MAM solution

D. Configure a video ingestion stream by using Amazon Kinesis Video Streams Use the catalog of faces to build a collection in Amazon Rekognition Stream the videos from the MAM solution into Kinesis Video Streams Configure Amazon Rekognition to process the streamed videos Then, use a stream consumer to retrieve the required metadata and push the metadata into the MAM solution Configure the stream to store the videos in Amazon S3

E. Set up an Amazon EC2 instance that runs the OpenCV libranes Copy the videos, images, and facecatalog from the on-premises library into an Amazon EBS volumemounted on this EC2 instance Process the videos to retrieve the required metadata, and push the metadata into the MAM solution, while also copying the video files to an Amazon S3 bucket

Answer: C

NEW QUESTION 187

- (Exam Topic 2)

A software company is using three AWS accounts for each of its 10 development teams The company has developed an AWS CloudFormation standard VPC template that includes three NAT gateways The template is added to each account for each team The company is concerned that network costs will increase each time a new development team is added A solutions architect must maintain the reliability of the company's solutions and minimize operational complexity What should the solutions architect do to reduce the network costs while meeting these requirements?

- A. Create a single VPC with three NAT gateways in a shared services account Configure each account VPC with a default route through a transit gateway to the NAT gateway in the shared services account VPC Remove all NAT gateways from the standard VPC template
- B. Create a single VPC with three NAT gateways in a shared services account Configure each account VPC with a default route through a VPC peering connection to the NAT gateway in the shared services account VPC Remove all NAT gateways from the standard VPC template
- C. Remove two NAT gateways from the standard VPC template Rely on the NAT gateway SLA to cover reliability for the remaining NAT gateway.
- D. Create a single VPC with three NAT gateways in a shared services account Configure a Site-to-Site VPN connection from each account to the shared services account Remove all NAT gateways from the standard VPC template

Answer: A

NEW QUESTION 189

- (Exam Topic 2)

A company has an application that runs on Amazon EC2 instances in an Amazon EC2 Auto Scaling group. The company uses AWS CodePipeline to deploy the application. The instances that run in the Auto Scaling group are constantly changing because of scaling events When the company deploys new application code versions the company Installs the AWS CodeDeploy agent on any new target EC2 instances and associates the instances with the CodeDeploy deployment group The application is set to go live within the next 24 hours What should a solutions architect recommend to automate the application deployment process with the LEAST amount of operational overhead?

- A. Configure Amazon EventBridge (Amazon CloudWatch Events) to invoke an AWS Lambda function when a new EC2 instance is launched into the Auto Scaling group
- B. Code the Lambda function to associate the EC2 instances with the CodeDeploy deployment group.
- C. Write a script to suspend Amazon EC2 Auto Scaling operations before the deployment of new code. When the deployment is complete, create a new AMI and configure the Auto Scaling group's launch template to use the new AMI for new launches
- D. Resume Amazon EC2 Auto Scaling operations
- E. Create a new AWS CodeBuild project that creates a new AMI that contains the new code Configure CodeBuild to update the Auto Scaling group's launch template to the new AMI Run an Amazon EC2 Auto Scaling instance refresh operation.
- F. Create a new AMI that has the CodeDeploy agent installed Configure the Auto Scaling group's launch template to use the new AMI Associate the CodeDeploy deployment group with the Auto Scaling group instead of the EC2 instances.

Answer: C

NEW QUESTION 190

- (Exam Topic 2)

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

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

Answer: CD

NEW QUESTION 195

- (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 API
- 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 AL
- R. Provide the security team with the ALB's IP addresses for the allow list.

Answer: A

NEW QUESTION 198

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your SAP-C02 Exam with Our Prep Materials Via below:

<https://www.certleader.com/SAP-C02-dumps.html>