



Amazon-Web-Services

Exam Questions SAP-C02

AWS Certified Solutions Architect - Professional

NEW QUESTION 1

- (Exam Topic 1)

A startup company recently migrated a large ecommerce website to AWS. The website has experienced a 70% increase in sales. Software engineers are using a private GitHub repository to manage code. The DevOps learn is using Jenkins for builds and unit testing. The engineers need to receive notifications for bad builds and zero downtime during deployments. The engineers also need to ensure any changes to production are seamless for users and can be rolled back in the event of a major issue.

The software engineers have decided to use AWS CodePipeline to manage their build and deployment process.

Which solution will meet these requirements?

- A. Use GitHub websockets to trigger the CodePipeline pipeline
- B. Use the Jenkins plugin for AWS CodeBuild to conduct unit testing
- C. Send alerts to an Amazon SNS topic for any bad build
- D. Deploy in an in-place
- E. all-at-once deployment configuration using AWS CodeDeploy.
- F. Use GitHub webhooks to trigger the CodePipeline pipeline
- G. Use the Jenkins plugin for AWS CodeBuild to conduct unit testing
- H. Send alerts to an Amazon SNS topic for any bad build
- I. Deploy in a blue/green deployment using AWS CodeDeploy.
- J. Use GitHub websockets to trigger the CodePipeline pipeline
- K. Use AWS X-Ray for unit testing and static code analysis
- L. Send alerts to an Amazon SNS topic for any bad build
- M. Deploy in a blue/green deployment using AWS CodeDeploy.
- N. Use GitHub webhooks to trigger the CodePipeline pipeline
- O. Use AWS X-Ray for unit testing and static code analysis
- P. Send alerts to an Amazon SNS topic for any bad build
- Q. Deploy in an in-place, all-at-once deployment configuration using AWS CodeDeploy.

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 company is running an application on several Amazon EC2 instances in an Auto Scaling group behind an Application Load Balancer. The load on the application varies throughout the day, and EC2 instances are scaled in and out on a regular basis. Log files from the EC2 instances are copied to a central Amazon S3 bucket every 15 minutes. The security team discovers that log files are missing from some of the terminated EC2 instances.

Which set of actions will ensure that log files are copied to the central S3 bucket from the terminated EC2 instances?

- A. Create a script to copy log files to Amazon S3, and store the script in a file on the EC2 instance
- B. Create an Auto Scaling lifecycle hook and an Amazon EventBridge (Amazon CloudWatch Events) rule to detect lifecycle events from the Auto Scaling group
- C. Invoke an AWS Lambda function on the autoscaling:EC2_INSTANCE_TERMINATING transition to send ABANDON to the Auto Scaling group to prevent termination, run the script to copy the log files, and terminate the instance using the AWS SDK.
- D. Create an AWS Systems Manager document with a script to copy log files to Amazon S3. Create an Auto Scaling lifecycle hook and an Amazon EventBridge (Amazon CloudWatch Events) rule to detect lifecycle events from the Auto Scaling group
- E. Invoke an AWS Lambda function on the autoscaling:EC2_INSTANCE_TERMINATING transition to call the AWS Systems Manager API SendCommand operation to run the document to copy the log files and send CONTINUE to the Auto Scaling group to terminate the instance.
- F. Change the log delivery rate to every 5 minutes
- G. Create a script to copy log files to Amazon S3, and add the script to EC2 instance user data
- H. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to detect EC2 instance termination
- I. Invoke an AWS Lambda function from the EventBridge (CloudWatch Events) rule that uses the AWS CLI to run the user-data script to copy the log files and terminate the instance.
- J. Create an AWS Systems Manager document with a script to copy log files to Amazon S3. Create an Auto Scaling lifecycle hook that publishes a message to an Amazon Simple Notification Service (Amazon SNS) topic
- K. From the SNS notification, call the AWS Systems Manager API SendCommand operation to run the document to copy the log files and send ABANDON to the

Auto Scaling group to terminate the instance.

Answer: B

Explanation:

<https://docs.aws.amazon.com/autoscaling/ec2/userguide/adding-lifecycle-hooks.html>

- Refer to Default Result section - If the instance is terminating, both abandon and continue allow the instance to terminate. However, abandon stops any remaining actions, such as other lifecycle hooks, and continue allows any other lifecycle hooks to complete.

<https://aws.amazon.com/blogs/infrastructure-and-automation/run-code-before-terminating-an-ec2-auto-scaling-i> <https://github.com/aws-samples/aws-lambda-lifecycle-hooks-function>

<https://github.com/aws-samples/aws-lambda-lifecycle-hooks-function/blob/master/cloudformation/template.yaml>

NEW QUESTION 4

- (Exam Topic 1)

A financial services company logs personally identifiable information 10 its application logs stored in Amazon S3. Due to regulatory compliance requirements, the log files must be encrypted at rest. The security team has mandated that the company's on-premises hardware security modules (HSMs) be used to generate the CMK material.

Which steps should the solutions architect take to meet these requirements?

- A. Create an AWS CloudHSM cluster
- B. Create a new CMK in AWS KMS using AWS_CloudHSM as the source (or the key material and an origin of AWS_CLOUDHSM)
- C. Enable automatic key rotation on the CMK with a duration of 1 year
- D. Configure a bucket policy on the logging bucket that disallows uploads of unencrypted data and requires that the encryption source be AWS KMS.
- E. Provision an AWS Direct Connect connection, ensuring there is no overlap of the RFC 1918 address space between on-premises hardware and the VPC
- F. Configure an AWS bucket policy on the logging bucket that requires all objects to be encrypted
- G. Configure the logging application to query the on-premises HSMs from the AWS environment for the encryption key material, and create a unique CMK for each logging event.
- H. Create a CMK in AWS KMS with no key material and an origin of EXTERNAL
- I. Import the key material generated from the on-premises HSMs into the CMK using the public key and import token provided by AWS
- J. Configure a bucket policy on the logging bucket that disallows uploads of non-encrypted data and requires that the encryption source be AWS KMS.
- K. Create a new CMK in AWS KMS with AWS-provided key material and an origin of AWS_KMS
- L. Disable this CMK
- M. and overwrite the key material with the key material from the on-premises HSM using the public key and import token provided by AWS
- N. Re-enable the CMK
- O. Enable automatic key rotation on the CMK with a duration of 1 year
- P. Configure a bucket policy on the logging bucket that disallows uploads of non-encrypted data and requires that the encryption source be AWS KMS.

Answer: C

Explanation:

<https://aws.amazon.com/blogs/security/how-to-byok-bring-your-own-key-to-aws-kms-for-less-than-15-00-a-year>

<https://docs.aws.amazon.com/kms/latest/developerguide/importing-keys-create-cmk.html>

NEW QUESTION 5

- (Exam Topic 1)

An enterprise runs 103 line-of-business applications on virtual machines in an on-premises data center. Many of the applications are simple PHP, Java, or Ruby web applications, are no longer actively developed, and serve little traffic.

Which approach should be used to migrate these applications to AWS with the LOWEST infrastructure costs?

- A. Deploy the applications to single-instance AWS Elastic Beanstalk environments without a load balancer.
- B. Use AWS SMS to create AMIs for each virtual machine and run them in Amazon EC2.
- C. Convert each application to a Docker image and deploy to a small Amazon ECS cluster behind an Application Load Balancer.
- D. Use VM Import/Export to create AMIs for each virtual machine and run them in single-instance AWS Elastic Beanstalk environments by configuring a custom image.

Answer: C

NEW QUESTION 6

- (Exam Topic 1)

A development team has created a new flight tracker application that provides near-real-time data to users. The application has a front end that consists of an Application Load Balancer (ALB) in front of two large Amazon EC2 instances in a single Availability Zone. Data is stored in a single Amazon RDS MySQL DB instance. An Amazon Route 53 DNS record points to the ALB.

Management wants the development team to improve the solution to achieve maximum reliability with the least amount of operational overhead.

Which set of actions should the team take?

- A. Create RDS MySQL read replica
- B. Deploy the application to multiple AWS Region
- C. Use a Route 53 latency-based routing policy to route to the application.
- D. Configure the DB instance as Multi-AZ
- E. Deploy the application to two additional EC2 instances in different Availability Zones behind an ALB.
- F. Replace the DB instance with Amazon DynamoDB global table
- G. Deploy the application in multiple AWS Region
- H. Use a Route 53 latency-based routing policy to route to the application.
- I. Replace the DB instance with Amazon Aurora with Aurora Replica
- J. Deploy the application to multiple smaller EC2 instances across multiple Availability Zones in an Auto Scaling group behind an ALB.

Answer: D

Explanation:

Multi AZ ASG + ALB + Aurora = Less overhead and automatic scaling

NEW QUESTION 7

- (Exam Topic 1)

A company has developed an application that is running Windows Server on VMware vSphere VMs that the company hosts on-premises. The application data is stored in a proprietary format that must be read through the application. The company manually provisioned the servers and the application.

As part of its disaster recovery plan, the company wants the ability to host its application on AWS temporarily if the company's on-premises environment becomes unavailable. The company wants the application to return to on-premises hosting after a disaster recovery event is complete. The RPO is 15 minutes.

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

- A. Configure AWS DataSync
- B. Replicate the data to Amazon Elastic Block Store (Amazon EBS) volumes. When the on-premises environment is unavailable, use AWS CloudFormation templates to provision Amazon EC2 instances and attach the EBS volumes.
- C. Configure CloudEndure Disaster Recovery. Replicate the data to replication Amazon EC2 instances that are attached to Amazon Elastic Block Store (Amazon EBS) volumes. When the on-premises environment is unavailable, use CloudEndure to launch EC2 instances that use the replicated volumes.
- D. Provision an AWS Storage Gateway. Use the gateway to replicate the data to Amazon S3 buckets.
- E. Recreate the data in an Amazon S3 bucket.
- F. When the on-premises environment is unavailable, use AWS Backup to restore the data to Amazon Elastic Block Store (Amazon EBS) volumes and launch Amazon EC2 instances from these EBS volumes.
- G. Provision an Amazon FSx for Windows File Server file system on AWS. Replicate the data to the file system. When the on-premises environment is unavailable, use AWS CloudFormation templates to provision Amazon EC2 instances and use AWS CloudFormation Init commands to mount the Amazon FSx file shares.

Answer: D

NEW QUESTION 8

- (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 resources. Deploy AWS WAF rules by using an AWS CloudFormation stack set to target the same OUs where the AWS Config rule is applied.
- E. Create AWS WAF rules in the management account of the organization. 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 accounts. 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.
- F. Use AWS Control Tower to manage AWS WAF rules across accounts in the organization. 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. 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 9

- (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 10

- (Exam Topic 1)

A company has a three-tier application running on AWS with a web server, an application server, and an Amazon RDS MySQL DB instance. A solutions architect is designing a disaster recovery (DR) solution with an RPO of 5 minutes.

Which solution will meet the company's requirements?

- A. Configure AWS Backup to perform cross-Region backups of all servers every 5 minutes.
- B. Reprovision the three tiers in the DR Region from the backups using AWS CloudFormation in the event of a disaster.
- C. Maintain another running copy of the web and application server stack in the DR Region using AWS CloudFormation. Detect a disaster and promote the DR Region stack to production.
- D. Configure cross-Region snapshots of the DB instance to the DR Region every 5 minutes.
- E. In the event of a disaster, restore the DB instance using the snapshot in the DR Region.
- F. Use Amazon EC2 Image Builder to create and copy AMIs of the web and application server to both the primary and DR Region.
- G. Create a cross-Region read replica of the DB instance in the DR Region.
- H. In the event of a disaster, promote the read replica to become the master and reprovision the servers with AWS CloudFormation using the AMIs.
- I. Create AMIs of the web and application servers in the DR Region.
- J. Use scheduled AWS Glue jobs to synchronize the DB instance with another DB instance in the DR Region.
- K. In the event of a disaster, switch to the DB instance in the DR Region and reprovision the servers with AWS CloudFormation using the AMIs.

Answer: C

Explanation:

deploying a brand new RDS instance will take >30 minutes. You will use EC2 Image builder to put the AMIs into the new region, but not use image builder to LAUNCH them.

NEW QUESTION 10

- (Exam Topic 1)

A solutions architect must analyze a company's Amazon EC2 Instances and Amazon Elastic Block Store (Amazon EBS) volumes to determine whether the company is using resources efficiently. The company is running several large, high-memory EC2 instances to host database clusters that are deployed in active/passive configurations. The utilization of these EC2 instances varies by the applications that use the databases, and the company has not identified a pattern. The solutions architect must analyze the environment and take action based on the findings. Which solution meets these requirements MOST cost-effectively?

- A. Create a dashboard by using AWS Systems Manager OpsCenter. Configure visualizations for Amazon CloudWatch metrics that are associated with the EC2 instances and their EBS volumes. Review the dashboard periodically and identify usage patterns. Rightsize the EC2 instances based on the peaks in the metrics.
- B. Turn on Amazon CloudWatch detailed monitoring for the EC2 instances and their EBS volumes. Create and review a dashboard that is based on the metrics. Identify usage patterns. Rightsize the EC2 instances based on the peaks in the metrics.
- C. Install the Amazon CloudWatch agent on each of the EC2 instances. Turn on AWS Compute Optimizer, and let it run for at least 12 hours. Review the recommendations from Compute Optimizer, and rightsize the EC2 instances as directed.
- D. Sign up for the AWS Enterprise Support plan. Turn on AWS Trusted Advisor. Wait 12 hours. Review the recommendations from Trusted Advisor, and rightsize the EC2 instances as directed.

Answer: C

Explanation:

(<https://aws.amazon.com/compute-optimizer/pricing/> , <https://aws.amazon.com/systems-manager/pricing/>). <https://aws.amazon.com/compute-optimizer/>

NEW QUESTION 11

- (Exam Topic 1)

A company has multiple AWS accounts as part of an organization created with AWS Organizations. Each account has a VPC in the us-east-2 Region and is used for either production or development workloads. Amazon EC2 instances across production accounts need to communicate with each other, and EC2 instances across development accounts need to communicate with each other, but production and development instances should not be able to communicate with each other.

To facilitate connectivity, the company created a common network account. The company used AWS Transit Gateway to create a transit gateway in the us-east-2 Region in the network account and shared the transit gateway with the entire organization by using AWS Resource Access Manager. Network administrators then attached VPCs in each account to the transit gateway, after which the EC2 instances were able to communicate across accounts. However, production and development accounts were also able to communicate with one another.

Which set of steps should a solutions architect take to ensure production traffic and development traffic are completely isolated?

- A. Modify the security groups assigned to development EC2 instances to block traffic from production EC2 instances.
- B. Modify the security groups assigned to production EC2 instances to block traffic from development EC2 instances.
- C. Create a tag on each VPC attachment with a value of either production or development, according to the type of account being attached.
- D. Using the Network Manager feature of AWS Transit Gateway, create policies that restrict traffic between VPCs based on the value of this tag.
- E. Create separate route tables for production and development traffic.
- F. Delete each account's association and route propagation to the default AWS Transit Gateway route table.
- G. Attach development VPCs to the development AWS Transit Gateway route table and production VPCs to the production route table, and enable automatic route propagation on each attachment.
- H. Create a tag on each VPC attachment with a value of either production or development, according to the type of account being attached.
- I. Modify the AWS Transit Gateway routing table to route production tagged attachments to one another and development tagged attachments to one another.

Answer: C

Explanation:

<https://docs.aws.amazon.com/vpc/latest/tgw/vpc-tgw.pdf>

NEW QUESTION 12

- (Exam Topic 1)

A company is deploying a new cluster for big data analytics on AWS. The cluster will run across many Linux Amazon EC2 instances that are spread across multiple Availability Zones.

All of the nodes in the cluster must have read and write access to common underlying file storage. The file storage must be highly available, must be resilient, must be compatible with the Portable Operating System Interface (POSIX), and must accommodate high levels of throughput.

Which storage solution will meet these requirements?

- A. Provision an AWS Storage Gateway file gateway NFS file share that is attached to an Amazon S3 bucket.
- B. Mount the NFS file share on each EC2 instance in the cluster.
- C. Provision a new Amazon Elastic File System (Amazon EFS) file system that uses General Purpose performance mode.
- D. Mount the EFS file system on each EC2 instance in the cluster.
- E. Provision a new Amazon Elastic Block Store (Amazon EBS) volume that uses the io2 volume type. Attach the EBS volume to all of the EC2 instances in the cluster.
- F. Provision a new Amazon Elastic File System (Amazon EFS) file system that uses Max I/O performance mode.
- G. Mount the EFS file system on each EC2 instance in the cluster.

Answer: D

NEW QUESTION 14

- (Exam Topic 1)

A company is running a web application on Amazon EC2 instances in a production AWS account. The company requires all logs generated from the web application to be copied to a central AWS account (for analysis and archiving). The company's AWS accounts are currently managed independently. Logging agents

are configured on the EC2 instances to upload the log files to an Amazon S3 bucket in the central AWS account.

A solutions architect needs to provide access for a solution that will allow the production account to store log files in the central account. The central account also needs to have read access to the log files.

What should the solutions architect do to meet these requirements?

- A. Create a cross-account role in the central account
- B. Assume the role from the production account when the logs are being copied.
- C. Create a policy on the S3 bucket with the production account ID as the principal
- D. Allow S3 access from a delegated user.
- E. Create a policy on the S3 bucket with access from only the CIDR range of the EC2 instances in the production account
- F. Use the production account ID as the principal.
- G. Create a cross-account role in the production account
- H. Assume the role from the production account when the logs are being copied.

Answer: B

NEW QUESTION 15

- (Exam Topic 1)

A company is storing data on premises on a Windows file server. The company produces 5 GB of new data daily.

The company migrated part of its Windows-based workload to AWS and needs the data to be available on a file system in the cloud. The company already has established an AWS Direct Connect connection between the on-premises network and AWS.

Which data migration strategy should the company use?

- A. Use the file gateway option in AWS Storage Gateway to replace the existing Windows file server, and point the existing file share to the new file gateway.
- B. Use AWS DataSync to schedule a daily task to replicate data between the on-premises Windows file server and Amazon FSx.
- C. Use AWS Data Pipeline to schedule a daily task to replicate data between the on-premises Windows file server and Amazon Elastic File System (Amazon EFS).
- D. Use AWS DataSync to schedule a daily task to replicate data between the on-premises Windows file server and Amazon Elastic File System (Amazon EFS).

Answer: B

Explanation:

<https://aws.amazon.com/storagegateway/file/> <https://docs.aws.amazon.com/fsx/latest/WindowsGuide/migrate-files-to-fsx-datasync.html>

<https://docs.aws.amazon.com/systems-manager/latest/userguide/prereqs-operating-systems.html#prereqs-os-win>

NEW QUESTION 20

- (Exam Topic 1)

A large company with hundreds of AWS accounts has a newly established centralized internal process for purchasing new or modifying existing Reserved Instances. This process requires all business units that want to purchase or modify Reserved Instances to submit requests to a dedicated team for procurement or execution. Previously, business units would directly purchase or modify Reserved Instances in their own respective AWS accounts autonomously.

Which combination of steps should be taken to proactively enforce the new process in the MOST secure way possible? (Select TWO.)

- A. Ensure all AWS accounts are part of an AWS Organizations structure operating in all features mode.
- B. Use AWS Config to report on the attachment of an IAM policy that denies access to the ec2:PurchaseReservedInstancesOffering and ec2:ModifyReservedInstances actions.
- C. In each AWS account, create an IAM policy with a DENY rule to the ec2:PurchaseReservedInstancesOffering and ec2:ModifyReservedInstances actions.
- D. Create an SCP that contains a deny rule to the ec2:PurchaseReservedInstancesOffering and ec2:ModifyReservedInstances action
- E. Attach the SCP to each organizational unit (OU) of the AWS Organizations structure.
- F. Ensure that all AWS accounts are part of an AWS Organizations structure operating in consolidated billing features mode.

Answer: AD

Explanation:

https://docs.aws.amazon.com/organizations/latest/APIReference/API_EnableAllFeatures.html

https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scp-strategies.html

NEW QUESTION 21

- (Exam Topic 1)

An e-commerce company is revamping its IT infrastructure and is planning to use AWS services. The company's CIO has asked a solutions architect to design a simple, highly available, and loosely coupled order processing application. The application is responsible for receiving and processing orders before storing them in an Amazon DynamoDB table. The application has a sporadic traffic pattern and should be able to scale during marketing campaigns to process the orders with minimal delays.

Which of the following is the MOST reliable approach to meet the requirements?

- A. Receive the orders in an Amazon EC2-hosted database and use EC2 instances to process them.
- B. Receive the orders in an Amazon SQS queue and trigger an AWS Lambda function to process them.
- C. Receive the orders using the AWS Step Functions program and trigger an Amazon ECS container to process them.
- D. Receive the orders in Amazon Kinesis Data Streams and use Amazon EC2 instances to process them.

Answer: B

Explanation:

Q: How does Amazon Kinesis Data Streams differ from Amazon SQS?

Amazon Kinesis Data Streams enables real-time processing of streaming big data. It provides ordering of records, as well as the ability to read and/or replay records in the same order to multiple Amazon Kinesis Applications. The Amazon Kinesis Client Library (KCL) delivers all records for a given partition key to the same record processor, making it easier to build multiple applications reading from the same Amazon Kinesis data stream (for example, to perform counting, aggregation, and filtering).

<https://aws.amazon.com/kinesis/data-streams/faqs/>

<https://aws.amazon.com/blogs/big-data/unite-real-time-and-batch-analytics-using-the-big-data-lambda-architect>

NEW QUESTION 26

- (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 for 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 instance
- C. Create Amazon Kinesis data streams and configure the application services to 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 scaling
- 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 scaling
- K. Create an Amazon Managed Streaming for Apache Kafka cluster and configure the application services to use the cluster
- 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 instance
- O. Create an Amazon Managed Streaming for Apache Kafka cluster and configure the application services to use the cluster
- 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 30

- (Exam Topic 1)

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

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

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

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

Answer: B

Explanation:

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

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

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

NEW QUESTION 31

- (Exam Topic 1)

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

Which solutions will meet these requirements?

- A. Trigger an AWS Lambda function on file delivery that extracts each record and writes it to an Amazon SQS queue
- B. Trigger another Lambda function when new messages arrive in the SQS queue to process the records, writing the results to a temporary location in Amazon S3. Trigger a final Lambda function once the SQS queue is empty to transform the records into JSON format and send the results to another S3 bucket for internal processing.
- C. Trigger an AWS Lambda function on file delivery that extracts each record and writes it to an Amazon SQS queue
- D. Configure an AWS Fargate container application to

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

Answer: C

Explanation:

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

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

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

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

NEW QUESTION 34

- (Exam Topic 1)

A company has a photo sharing social networking application. To provide a consistent experience for users, the company performs some image processing on the photos uploaded by users before publishing on the application. The image processing is implemented using a set of Python libraries.

The current architecture is as follows:

- The image processing Python code runs in a single Amazon EC2 instance and stores the processed images in an Amazon S3 bucket named ImageBucket.
- The front-end application, hosted in another bucket, loads the images from ImageBucket to display to users. With plans for global expansion, the company wants to implement changes in its existing architecture to be able to scale for increased demand on the application and reduce management complexity as the application scales.

Which combination of changes should a solutions architect make? (Select TWO.)

- A. Place the image processing EC2 instance into an Auto Scaling group.
- B. Use AWS Lambda to run the image processing tasks.
- C. Use Amazon Rekognition for image processing.
- D. Use Amazon CloudFront in front of ImageBucket.
- E. Deploy the applications in an Amazon ECS cluster and apply Service Auto Scaling.

Answer: BD

Explanation:

<https://prismatic.io/blog/why-we-moved-from-lambda-to-ecs/>

NEW QUESTION 39

- (Exam Topic 1)

A company is running a containerized application in the AWS Cloud. The application is running by using Amazon Elastic Container Service (Amazon ECS) on a set of Amazon EC2 instances. The EC2 instances run in an Auto Scaling group.

The company uses Amazon Elastic Container Registry (Amazon ECR) to store its container images. When a new image version is uploaded, the new image version receives a unique tag.

The company needs a solution that inspects new image versions for common vulnerabilities and exposures. The solution must automatically delete new image tags that have Critical or High severity findings. The solution also must notify the development team when such a deletion occurs.

Which solution meets these requirements?

- A. Configure scan on push on the repository.
- B. Use Amazon EventBridge (Amazon CloudWatch Events) to invoke an AWS Step Functions state machine when a scan is complete for images that have Critical or High severity findings. Use the Step Functions state machine to delete the image tag for those images and to notify the development team through Amazon Simple Notification Service (Amazon SNS).
- C. Configure scan on push on the repository. Configure scan results to be pushed to an Amazon Simple Queue Service (Amazon SQS) queue. Invoke an AWS Lambda function when a new message is added to the SQS queue. Use the Lambda function to delete the image tag for images that have Critical or High severity findings.
- D. Notify the development team by using Amazon Simple Email Service (Amazon SES).
- E. Schedule an AWS Lambda function to start a manual image scan every hour. Configure Amazon EventBridge (Amazon CloudWatch Events) to invoke another Lambda function when a scan is complete.
- F. Use the second Lambda function to delete the image tag for images that have Critical or High severity findings.
- G. Notify the development team by using Amazon Simple Notification Service (Amazon SNS).
- H. Configure periodic image scan on the repository. Configure scan results to be added to an Amazon Simple Queue Service (Amazon SQS) queue. Invoke an AWS Step Functions state machine when a new message is added to the SQS queue. Use the Step Functions state machine to delete the image tag for images that have Critical or High severity findings.
- I. Notify the development team by using Amazon Simple Email Service (Amazon SES).

Answer: C

NEW QUESTION 41

- (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 45

- (Exam Topic 1)

A solutions architect is evaluating the reliability of a recently migrated application running on AWS. The front end is hosted on Amazon S3 and accelerated by Amazon CloudFront. The application layer is running in a stateless Docker container on an Amazon EC2 On-Demand Instance with an Elastic IP address. The storage layer is a MongoDB database running on an EC2 Reserved Instance in the same Availability Zone as the application layer.

Which combination of steps should the solutions architect take to eliminate single points of failure with minimal application code changes? (Select TWO.)

- A. Create a REST API in Amazon API Gateway and use AWS Lambda functions as the application layer.
- B. Create an Application Load Balancer and migrate the Docker container to AWS Fargate.
- C. Migrate the storage layer to Amazon DynamoD8.
- D. Migrate the storage layer to Amazon DocumentD8 (with MongoDB compatibility).
- E. Create an Application Load Balancer and move the storage layer to an EC2 Auto Scaling group.

Answer: BD

Explanation:

https://aws.amazon.com/documentdb/?nc1=h_ls
<https://aws.amazon.com/blogs/containers/using-alb-ingress-controller-with-amazon-eks-on-fargate/>

NEW QUESTION 50

- (Exam Topic 1)

A company runs an application that gives users the ability to search for videos and related information by using keywords that are curated from content providers. The application data is stored in an on-premises Oracle database that is 800 GB in size.

The company wants to migrate the data to an Amazon Aurora MySQL DB instance. A solutions architect plans to use the AWS Schema Conversion Tool and AWS Database Migration Service (AWS DMS) for the migration. During the migration, the existing database must serve ongoing requests. The migration must be completed with minimum downtime. Which solution will meet these requirements?

- A. Create primary key indexes, secondary indexes, and referential integrity constraints in the target database before starting the migration process
- B. Use AWS DMS to run the conversion report for Oracle to Aurora MySQL
- C. Remediate any issues Then use AWS DMS to migrate the data
- D. Use the M5 or CS DMS replication instance type for ongoing replication
- E. Turn off automatic backups and logging of the target database until the migration and cutover processes are complete

Answer: B

Explanation:

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Aurora.Managing.Backups.html>

NEW QUESTION 55

- (Exam Topic 1)

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

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

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

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

Answer: B

Explanation:

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

table.

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

NEW QUESTION 59

- (Exam Topic 1)

A company manages an on-premises JavaScript front-end web application. The application is hosted on two servers secured with a corporate Active Directory. The application calls a set of Java-based microservices on an application server and stores data in a clustered MySQL database. The application is heavily used during the day on weekdays. It is lightly used during the evenings and weekends.

Daytime traffic to the application has increased rapidly, and reliability has diminished as a result. The company wants to migrate the application to AWS with a solution that eliminates the need for server maintenance, with an API to securely connect to the microservices.

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

- A. Host the web application on Amazon S3. Use Amazon Cognito identity pools (federated identities) with SAML for authentication and authorization.
- B. Host the web application on Amazon EC2 with Auto Scaling
- C. Use Amazon Cognito federation and Login with Amazon for authentication and authorization.
- D. Create an API layer with Amazon API Gateway
- E. Rehost the microservices on AWS Fargate containers.
- F. Create an API layer with Amazon API Gateway
- G. Rehost the microservices on Amazon Elastic Container Service (Amazon ECS) containers.
- H. Replatform the database to Amazon RDS for MySQL.
- I. Replatform the database to Amazon Aurora MySQL Serverless.

Answer: ACE

NEW QUESTION 60

- (Exam Topic 1)

A company is migrating an application to AWS. It wants to use fully managed services as much as possible during the migration. The company needs to store large, important documents within the application with the following requirements:

- * 1. The data must be highly durable and available.
- * 2. The data must always be encrypted at rest and in transit.
- * 3. The encryption key must be managed by the company and rotated periodically.

Which of the following solutions should the solutions architect recommend?

- A. Deploy the storage gateway to AWS in file gateway mode
- B. Use Amazon EBS volume encryption using an AWS KMS key to encrypt the storage gateway volumes.
- C. Use Amazon S3 with a bucket policy to enforce HTTPS for connections to the bucket and to enforce server-side encryption and AWS KMS for object encryption.
- D. Use Amazon DynamoDB with SSL to connect to DynamoDB
- E. Use an AWS KMS key to encrypt DynamoDB objects at rest.
- F. Deploy instances with Amazon EBS volumes attached to store this data
- G. Use EBS volume encryption using an AWS KMS key to encrypt the data.

Answer: B

Explanation:

Use Amazon S3 with a bucket policy to enforce HTTPS for connections to the bucket and to enforce server-side encryption and AWS KMS for object encryption.

NEW QUESTION 64

- (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 instance
- B. Create an Amazon S3 bucket to be used for SFTP file hosting
- 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 bucket.
- D. Sync all files from the SFTP server to the S3 bucket.
- E. Disassociate the Elastic IP address from the EC2 instance
- F. Create an Amazon S3 bucket to be used for SFTP file hosting
- G. Create an AWS Transfer Family server
- H. Configure the Transfer Family server with a VPC-hosted
- 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 instance
- N. Create a new Amazon Elastic File System (Amazon EFS) file system to be used for SFTP file hosting
- O. Create an AWS Fargate task definition to run an SFTP server
- P. Specify the EFS file system as a mount in the task definition
- 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 server
- R. Associate the Elastic IP address with the NLB
- S. Sync all files from the SFTP server to the S3 bucket.
- T. Disassociate the Elastic IP address from the EC2 instance
- . Create a multi-attach Amazon Elastic Block Store (Amazon EBS) volume to be used for SFTP file hosting
- . Create a Network Load Balancer (NLB) with the Elastic IP address attached

. Create an Auto Scaling group with EC2 instances that run an SFTP server Define in the Auto Scaling group that instances that are launched should attach the new multi-attach EBS volume Configure the Auto Scaling group to automatically add instances behind the NLB Configure the Auto Scaling group to use the security group that allows customer IP addresses for the EC2 instances that the Auto Scaling group launches
 . Sync all files from the SFTP server to the new multi-attach EBS volume.

Answer: B

Explanation:

<https://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 69

- (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 71

- (Exam Topic 1)

A solutions architect is building a web application that uses an Amazon RDS for PostgreSQL DB instance. The DB instance is expected to receive many more reads than writes. The solutions architect needs to ensure that the large amount of read traffic can be accommodated and that the DB instance is highly available. Which steps should the solutions architect take to meet these requirements? (Select THREE.)

- A. Create multiple read replicas and put them into an Auto Scaling group
- B. Create multiple read replicas in different Availability Zones.
- C. Create an Amazon Route 53 hosted zone and a record set for each read replica with a TTL and a weighted routing policy
- D. Create an Application Load Balancer (ALB) and put the read replicas behind the ALB.
- E. Configure an Amazon CloudWatch alarm to detect a failed read replica. Set the alarm to directly invoke an AWS Lambda function to delete its Route 53 record set.
- F. Configure an Amazon Route 53 health check for each read replica using its endpoint

Answer: BCF

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/requests-rds-read-replicas/>

You can use Amazon Route 53 weighted record sets to distribute requests across your read replicas. Within a Route 53 hosted zone, create individual record sets for each DNS endpoint associated with your read replicas and give them the same weight. Then, direct requests to the endpoint of the record set. You can incorporate Route 53 health checks to be sure that Route 53 directs traffic away from unavailable read replicas

NEW QUESTION 75

- (Exam Topic 1)

A company is moving a business-critical multi-tier application to AWS. The architecture consists of a desktop client application and server infrastructure. The server infrastructure resides in an on-premises data center that frequently fails to maintain the application uptime SLA of 99.95%. A solutions architect must re-architect the application to ensure that it can meet or exceed the SLA.

The application contains a PostgreSQL database running on a single virtual machine. The business logic and presentation layers are load balanced between multiple virtual machines. Remote users complain about slow load times while using this latency-sensitive application.

Which of the following will meet the availability requirements with little change to the application while improving user experience and minimizing costs?

- A. Migrate the database to a PostgreSQL database in Amazon EC2. Host the application and presentation layers in automatically scaled Amazon ECS containers behind an Application Load Balance
- B. Allocate an Amazon Workspaces Workspace for each end user to improve the user experience.
- C. Migrate the database to an Amazon RDS Aurora PostgreSQL configuration
- D. Host the application and presentation layers in an Auto Scaling configuration on Amazon EC2 instances behind an Application Load Balance
- E. Use Amazon AppStream 2.0 to improve the user experience.
- F. Migrate the database to an Amazon RDS PostgreSQL Multi-AZ configuration
- G. Host the application and presentation layers in automatically scaled AWS Fargate containers behind a Network Load Balance
- H. Use Amazon ElastiCache to improve the user experience.
- I. Migrate the database to an Amazon Redshift cluster with at least two nodes
- J. Combine and host the application and presentation layers in automatically scaled Amazon ECS containers behind an Application Load Balance
- K. Use Amazon CloudFront to improve the user experience.

Answer: B

Explanation:

Aurora would improve availability that can replicate to multiple AZ (6 copies). Auto scaling would improve the performance together with a ALB. AppStream is like Citrix that deliver hosted Apps to users.

NEW QUESTION 80

- (Exam Topic 1)

A company wants to migrate its corporate data center from on premises to the AWS Cloud. The data center includes physical servers and VMs that use VMware and Hyper-V. An administrator needs to select the correct services to collect data (or the initial migration discovery process). The data format should be supported by AWS Migration Hub. The company also needs the ability to generate reports from the data. Which solution meets these requirements?

- A. Use the AWS Agentless Discovery Connector for data collection on physical servers and all VM
- B. Store the collected data in Amazon S3. Query the data with S3 Select
- C. Generate reports by using Kibana hosted on Amazon EC2.
- D. Use the AWS Application Discovery Service agent for data collection on physical servers and all VMs. Store the collected data in Amazon Elastic File System (Amazon EFS). Query the data and generate reports with Amazon Athena.
- E. Use the AWS Application Discovery Service agent for data collection on physical servers and Hyper-
- F. Use the AWS Agentless Discovery Connector for data collection on VMwar
- G. Store the collected data in Amazon S3. Query the data with Amazon Athen
- H. Generate reports by using Amazon QuickSight.
- I. Use the AWS Systems Manager agent for data collection on physical server
- J. Use the AWS Agentless Discovery Connector for data collection on all VM
- K. Store, query, and generate reports from the collected data by using Amazon Redshift.

Answer: C

Explanation:

<https://docs.aws.amazon.com/application-discovery/latest/userguide/discovery-agent.html> <https://docs.aws.amazon.com/application-discovery/latest/userguide/discovery-connector.html>

NEW QUESTION 81

- (Exam Topic 1)

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

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

Answer: B

Explanation:

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

NEW QUESTION 86

- (Exam Topic 1)

A company has a data lake in Amazon S3 that needs to be accessed by hundreds of applications across many AWS accounts. The company's information security policy states that the S3 bucket must not be accessed over the public internet and that each application should have the minimum permissions necessary to function.

To meet these requirements, a solutions architect plans to use an S3 access point that is restricted to specific VPCs for each application. Which combination of steps should the solutions architect take to implement this solution? (Select TWO.)

- A. Create an S3 access point for each application in the AWS account that owns the S3 bucket
- B. Configure each access point to be accessible only from the application's VPC
- C. Update the bucket policy to require access from an access point.
- D. Create an interface endpoint for Amazon S3 in each application's VPC
- E. Configure the endpoint policy to allow access to an S3 access point
- F. Create a VPC gateway attachment for the S3 endpoint.
- G. Create a gateway endpoint for Amazon S3 in each application's VPC
- H. Configure the endpoint policy to allow access to an S3 access point
- I. Specify the route table that is used to access the access point.
- J. Create an S3 access point for each application in each AWS account and attach the access points to the S3 bucket
- K. Configure each access point to be accessible only from the application's VPC
- L. Update the bucket policy to require access from an access point.
- M. Create a gateway endpoint for Amazon S3 in the data lake's VPC
- N. Attach an endpoint policy to allow access to the S3 bucket
- O. Specify the route table that is used to access the bucket.

Answer: AC

Explanation:

<https://joe.blog.freemansoft.com/2020/04/protect-data-in-cloud-with-s3-access.html> <https://aws.amazon.com/s3/features/access-points/>
<https://aws.amazon.com/s3/features/access-points/>
&

<https://aws.amazon.com/blogs/storage/managing-amazon-s3-access-with-vpc-endpoints-and-s3-access-points/>

NEW QUESTION 88

- (Exam Topic 1)

A company requires that all internal application connectivity use private IP addresses. To facilitate this policy, a solutions architect has created interface endpoints to connect to AWS public services. Upon testing, the solutions architect notices that the service names are resolving to public IP addresses, and that internal services cannot connect to the interface endpoints.

Which step should the solutions architect take to resolve this issue?

- A. Update the subnet route table with a route to the interface endpoint.
- B. Enable the private DNS option on the VPC attributes.
- C. Configure the security group on the interface endpoint to allow connectivity to the AWS services.
- D. Configure an Amazon Route 53 private hosted zone with a conditional forwarder for the internal application.

Answer: C

Explanation:

<https://docs.aws.amazon.com/vpc/latest/privatelink/vpce-interface.html>

NEW QUESTION 93

- (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 96

- (Exam Topic 1)

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

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

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

Answer: A

Explanation:

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

NEW QUESTION 101

- (Exam Topic 1)

A company has an internal application running on AWS that is used to track and process shipments in the company's warehouse. Currently, after the system receives an order, it emails the staff the information needed to ship a package. Once the package is shipped, the staff replies to the email and the order is marked as shipped.

The company wants to stop using email in the application and move to a serverless application model. Which architecture solution meets these requirements?

- A. Use AWS Batch to configure the different tasks required to ship a package

- B. Have AWS Batch trigger an AWS Lambda function that creates and prints a shipping label
- C. Once that label is scanned
- D. as it leaves the warehouse, have another Lambda function move the process to the next step in the AWS Batch job.
- E. When a new order is created, store the order information in Amazon SQS
- F. Have AWS Lambda check the queue every 5 minutes and process any needed work
- G. When an order needs to be shipped, have Lambda print the label in the warehouse
- H. Once the label has been scanned, as it leaves the warehouse, have an Amazon EC2 instance update Amazon S3.
- I. Update the application to store new order information in Amazon DynamoDB
- J. When a new order is created, trigger an AWS Step Functions workflow, mark the orders as "in progress," and print a package label to the warehouse
- K. Once the label has been scanned and fulfilled, the application will trigger an AWS Lambda function that will mark the order as shipped and complete the workflow.
- L. Store new order information in Amazon EFS
- M. Have instances pull the new information from the NFS and send that information to printers in the warehouse
- N. Once the label has been scanned, as it leaves the warehouse, have Amazon API Gateway call the instances to remove the order information from Amazon EFS.

Answer: C

NEW QUESTION 106

- (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 SFTP
- 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 108

- (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. Deploy the EC2 instances and the RDS instance for the application by using the AWS CloudFormation 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. Deploy 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 110

- (Exam Topic 2)

A company has developed a new billing application that will be released in two weeks. Developers are testing the application running on 10 EC2 instances managed by an Auto Scaling group in subnet `172.31.0.0/24` within VPC A with CIDR block `172.31.0.0/16`. The developers noticed connection timeout errors in the application logs while connecting to an Oracle database running on an Amazon EC2 instance in the same region within VPC B with CIDR block `172.50.0.0/16`. The IP of the database instance is hard-coded in the application instances.

Which recommendations should a solutions architect present to the developers to solve the problem in a secure way with minimal maintenance and overhead?

- A. Disable the `SrcDestCheck` attribute for all instances running the application and Oracle Database. Change the default route of VPC A to point to the ENI of the Oracle Database that has an IP address assigned within the range of `172.50.0.0/16`
- B. Create and attach internet gateways for both VPCs
- C. Configure default routes to the internet gateways for both VPCs
- D. Assign an Elastic IP for each Amazon EC2 instance in VPC A
- E. Create a VPC peering connection between the two VPCs and add a route to the routing table of VPC A that points to the IP address range of `172.50.0.0/16`
- F. Create an additional Amazon EC2 instance for each VPC as a customer gateway; create one virtual private gateway (VGW) for each VPC
- G. Configure an end-to-end VPC, and advertise the routes for `172.50.0.0/16`

Answer: C

NEW QUESTION 112

- (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 RD
- B. and AWS CloudTrail Load customer records in Amazon RDS MySQL and train users to run queries using the AWS CL
- C. 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
- D. 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 tram users to run queries using the CLI via Athena Analyze CloudTrail events to audit and alarm on queries against personal data
- E. Apply a service control policy (SCP) that denies access to all services except IAM Amazon DynamoD
- F. 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
- G. 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 115

- (Exam Topic 2)

A company that runs applications on AWS recently subscribed to a new software-as-a-service (SaaS) data vendor. The vendor provides the data by way of a REST API that the vendor hosts in its AWS environment. The vendor offers multiple options for connectivity to the API and is working with the company to find the best way to connect.

The company's AWS account does not allow outbound internet access from its AWS environment. The vendor's services run on AWS in the same AWS Region as the company's applications.

A solutions architect must implement connectivity to the vendor's API so that the API is highly available in the company's VPC.

Which solution will meet these requirements?

- A. Connect to the vendor's public API address for the data service.
- B. Connect to the vendor by way of a VPC peering connection between the vendor's VPC and the company's VPC
- C. Connect to the vendor by way of a VPC endpoint service that uses AWS PrivateLink
- D. Connect to a public bastion host that the vendor provides. Tunnel the API traffic.

Answer: C

NEW QUESTION 120

- (Exam Topic 2)

A company is using an Amazon CloudFront distribution to distribute both static and dynamic content from a web application running behind an Application Load Balancer. The web application requires user authorization and session tracking for dynamic content. The CloudFront distribution has a single cache behavior configured to forward the Authorization, Host, and Agent HTTP allow list headers and a session cookie to the origin. All other cache behavior settings are set to their default value.

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

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

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

Answer: D

Explanation:

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

NEW QUESTION 124

- (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 128

- (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 132

- (Exam Topic 2)

A company has more than 10,000 sensors that send data to an on-premises Apache Kafka server by using the Message Queuing Telemetry Transport (MQTT) protocol. The on-premises Kafka server transforms the data and then stores the results as objects in an Amazon S3 bucket.

Recently, the Kafka server crashed. The company lost sensor data while the server was being restored. A solutions architect must create a new design on AWS that is highly available and scalable to prevent a similar occurrence.

Which solution will meet these requirements?

- A. Launch two Amazon EC2 instances to host the Kafka server in an active/standby configuration across two Availability Zones.
- B. Create a domain name in Amazon Route 53. Create a Route 53 failover policy. Route the sensors to send the data to the domain name.
- C. Migrate the on-premises Kafka server to Amazon Managed Streaming for Apache Kafka (Amazon MSK). Create a Network Load Balancer (NLB) that points to the Amazon MSK broker.
- D. Enable NLB health checks. Route the sensors to send the data to the NLB.
- E. Deploy AWS IoT Core, and connect it to an Amazon Kinesis Data Firehose delivery stream. Use an AWS Lambda function to handle data transformation. Route the sensors to send the data to AWS IoT Core.
- F. Deploy AWS IoT Core, and launch an Amazon EC2 instance to host the Kafka server. Configure AWS IoT Core to send the data to the EC2 instance. Route the sensors to send the data to AWS IoT Core.

Answer: A

NEW QUESTION 134

- (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 `MfS::DynamoDB::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 138

- (Exam Topic 2)

A company uses AWS Organizations with a single OU named Production to manage multiple accounts. All accounts are members of the Production OU. Administrators use deny list SCPs in the root of the organization to manage access to restricted services. The company recently acquired a new business unit and invited the new unit's existing AWS account to the organization. Once onboarded, the administrators of the new business unit discovered that they are not able to update existing AWS Config rules to meet the company's policies.

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

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

Answer: D

NEW QUESTION 141

- (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. 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 144

- (Exam Topic 2)

A company has developed a web application. The company is hosting the application on a group of Amazon EC2 instances behind an Application Load Balancer. The company wants to improve the security posture of the application and plans to use AWS WAF web ACLs. The solution must not adversely affect legitimate traffic to the application.

How should a solutions architect configure the web ACLs to meet these requirements?

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

Answer: B

NEW QUESTION 147

- (Exam Topic 2)

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

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

- A. Set up SQL Server to run in Fargate with Service Auto Scaling.
- B. Create an Amazon ECS task execution role that allows the Fargate task definition to get the secret value for the credentials to SQL Server running in Fargate.
- C. Specify the ARN of the secret in AWS Secrets Manager in the secrets section of the Fargate task definition so the sensitive data can be injected into the containers as environment variables on startup for reading into the application to construct the connection string.

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

Answer: B

Explanation:

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

NEW QUESTION 148

- (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 (CloudEnsure Migration) to migrate the Windows servers to AW
- B. Create a Replication Settings templat
- C. Install the AWS Replication Agent on the source servers
- D. Use AWS DataSync to migrate the Windows servers to AW
- 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 AW
- 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 150

- (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 jo
- E. Configure a server in the company's data center with an extraction agen
- F. Use AWS SCT to manage the extraction agent and convert the Oracle schema to an Amazon Redshift schem
- G. Create a new project in AWS SCT using the registered data extraction agen
- H. Create a local task and an AWS DMS task in AWS SCT with replication of ongoing change

- I. Copy data to the Snowball device and return the Snowball device to AW
- 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 155

- (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 156

- (Exam Topic 2)

A company is running a workload that consists of thousands of Amazon EC2 instances. The workload is running in a VPC that contains several public subnets and private subnets. The public subnets have a route for 0.0.0.0/0 to an existing internet gateway. The private subnets have a route for 0.0.0.0/0 to an existing NAT gateway.

A solutions architect needs to migrate the entire fleet of EC2 instances to use IPv6. The EC2 instances that are in private subnets must not be accessible from the public internet.

What should the solutions architect do to meet these requirements?

- A. Update the existing VPC and associate a custom IPv6 CIDR block with the VPC and all subnets. Update all the VPC route tables and add a route for /0 to the internet gateway.
- B. Update the existing VPC and associate an Amazon-provided IPv6 CIDR block with the VPC and all subnets. Update the VPC route tables for all private subnets, and add a route for /0 to the NAT gateway.
- C. Update the existing VPC and associate an Amazon-provided IPv6 CIDR block with the VPC and all subnets. Update the VPC route tables for all private subnets, and add a route for /0 to the egress-only internet gateway.
- D. Update the existing VPC and associate a custom IPv6 CIDR block with the VPC and all subnets. Create a new NAT gateway, and enable IPv6 support. Update the VPC route tables for all private subnets and add a route for /0 to the IPv6-enabled NAT gateway.
- E. Update the existing VPC and associate a custom IPv6 CIDR block with the VPC and all subnets. Update all the VPC route tables and add a route for /0 to the internet gateway.
- F. Update the existing VPC and associate an Amazon-provided IPv6 CIDR block with the VPC and all subnets. Update the VPC route tables for all private subnets, and add a route for /0 to the NAT gateway.

Answer: C

NEW QUESTION 157

- (Exam Topic 2)

A company recently started hosting new application workloads in the AWS Cloud. The company is using Amazon EC2 instances, Amazon Elastic File System (Amazon EFS) file systems, and Amazon RDS DB instances.

To meet regulatory and business requirements, the company must make the following changes for data backups:

- Backups must be retained based on custom daily, weekly, and monthly requirements.
- Backups must be replicated to at least one other AWS Region immediately after capture.
- The backup solution must provide a single source of backup status across the AWS environment.
- The backup solution must send immediate notifications upon failure of any resource backup.

Which combination of steps will meet these requirements with the LEAST amount of operational overhead? (Select THREE.)

- A. Create an AWS Backup plan with a backup rule for each of the retention requirements.
- B. Configure an AWS Backup plan to copy backups to another Region.
- C. Create an AWS Lambda function to replicate backups to another Region and send notification if a failure occurs.
- D. Add an Amazon Simple Notification Service (Amazon SNS) topic to the backup plan to send a notification for finished jobs that have any status except BACKUP_JOB_COMPLETE.
- E. Create an Amazon Data Lifecycle Manager (Amazon DLM) snapshot lifecycle policy for each of the retention requirements.
- F. Set up RDS snapshots on each database.

Answer: BDE

NEW QUESTION 160

- (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 StrmgEquals condition for the DevelopmentUnit resource tag and aws Principal Tag '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 162

- (Exam Topic 2)

A company is planning to migrate an Amazon RDS for Oracle database to an RDS for PostgreSQL DB instance in another AWS account A solutions architect needs to design a migration strategy that will require no downtime and that will minimize the amount of time necessary to complete the migration The migration strategy must replicate all existing data and any new data that is created during the migration The target database must be identical to the source database at completion of the migration process

All applications currently use an Amazon Route 53 CNAME record as their endpoint for communication with the RDS for Oracle DB instance The RDS for Oracle DB instance is in a private subnet

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

- A. Create a new RDS for PostgreSQL DB instance in the target account Use the AWS Schema Conversion Tool (AWS SCT) to migrate the database schema from the source database to the target database.
- B. Use the AWS Schema Conversion Tool (AWS SCT) to create a new RDS for PostgreSQL DB instance in the target account with the schema and initial data from the source database
- C. Configure VPC peering between the VPCs in the two AWS accounts to provide connectivity to both DB instances from the target account
- D. Configure the security groups that are attached to each DB instance to allow traffic on the database port from the VPC in the target account
- E. Temporarily allow the source DB instance to be publicly accessible to provide connectivity from the VPC in the target account Configure the security groups that are attached to each DB instance to allow traffic on the database port from the VPC in the target account.
- F. Use AWS Database Migration Service (AWS DMS) in the target account to perform a full load plus change data capture (CDC) migration from the source database to the target database When the migration is complete, change the CNAME record to point to the target DB instance endpoint
- G. Use AWS Database Migration Service (AWS DMS) in the target account to perform a change data capture (CDC) migration from the source database to the target database When the migration is complete change the CNAME record to point to the target DB instance endpoint

Answer: BCE

NEW QUESTION 163

- (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 cach
- B. Update the application to use DA
- 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 AL
- E. Update the Route 53 record to use a simple routing policy that targets the ALB's DNS alia
- F. Configure scheduled scaling for the EC2 instances before the content updates.
- G. Set up Amazon ElastiCache for Redi
- H. Update the application to use ElastiCach
- I. Create an Auto Scaling group for the EC2 instance
- J. Create an AmazonCloudFront distnbutio
- K. and set the Auto Scaling group as an origin for the distributio
- 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 Memcache
- N. Update the application to use ElastiCach
- 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 AL
- P. Update the Route 53 record to use a simple routing policy that targets the ALB's DNS alia
- Q. Configure scheduled scaling for the application before the content updates.
- R. Set up DynamoDB Accelerator (DAX) as in-memory cach
- S. Update the application to use DA
- 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 distributio
- . Update the Route 53 record to use a simple routing policy that targets the CloudFront distribution's DNS alia
- . Manually scale up EC2 instances before the content updates.

Answer: B

NEW QUESTION 165

- (Exam Topic 2)

A financial services company loaded millions of historical stock trades into an Amazon DynamoDB table The table uses on-demand capacity mode Once each day

at midnight, a few million new records are loaded into the table Application read activity against the table happens in bursts throughout the day, and a limited set of keys are repeatedly looked up. The company needs to reduce costs associated with DynamoDB. Which strategy should a solutions architect recommend to meet this requirement?

- A. Deploy an Amazon ElastiCache cluster in front of the DynamoDB table.
- B. Deploy DynamoDB Accelerator (DAX) Configure DynamoDB auto scaling Purchase Savings Plans in Cost Explorer
- C. Use provisioned capacity mode Purchase Savings Plans in Cost Explorer
- D. Deploy DynamoDB Accelerator (DAX) Use provisioned capacity mode Configure DynamoDB auto scaling

Answer: D

NEW QUESTION 170

- (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 DynamoD
- B. Store the leader different table
- C. Use Apache HiveQL JOIN statements to build the leaderboard
- D. Keep the leaderboard data in the RDS DB instanc
- 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 destinatio
- G. Query the S3 bucket by using Amazon Athena for the leaderboard.
- H. Add a read-only replica to the RDS DB instanc
- I. Add an RDS Proxy database proxy.

Answer: C

NEW QUESTION 175

- (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 180

- (Exam Topic 2)

A car rental company has built a serverless REST API to provide data to its mobile app. The app consists of an Amazon API Gateway API with a Regional endpoint, AWS Lambda functions and an Amazon Aurora MySQL Serverless DB cluster The company recently opened the API to mobile apps of partners A significant increase in the number of requests resulted causing sporadic database memory errors Analysis of the API traffic indicates that clients are making multiple HTTP GET requests for the same queries in a short period of time Traffic is concentrated during business hours, with spikes around holidays and other events

The company needs to improve its ability to support the additional usage while minimizing the increase in costs associated with the solution.

Which strategy meets these requirements?

- A. Convert the API Gateway Regional endpoint to an edge-optimized endpoint Enable caching in the production stage.
- B. Implement an Amazon ElastiCache for Redis cache to store the results of the database calls Modify the Lambda functions to use the cache
- C. Modify the Aurora Serverless DB cluster configuration to increase the maximum amount of available memory
- D. Enable throttling in the API Gateway production stage Set the rate and burst values to limit the incoming calls

Answer: A

NEW QUESTION 183

- (Exam Topic 2)

A company's CI/CO 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 187

- (Exam Topic 2)

A company is migrating an application to the AWS Cloud. The application runs in an on-premises data center and writes thousands of images into a mounted NFS file system each night After the company migrates the application, the company will host the application on an Amazon EC2 instance with a mounted Amazon Elastic File System (Amazon EFS) file system.

The company has established an AWS Direct Connect connection to AWS Before the migration cutover. a solutions architect must build a process that will replicate the newly created on-premises images to the EFS file system

What is the MOST operationally efficient way to replicate the images?

- A. Configure a periodic process to run the aws s3 sync command from the on-premises file system to Amazon S3 Configure an AWS Lambda function to process event notifications from Amazon S3 and copy the images from Amazon S3 to the EFS file system
- B. Deploy an AWS Storage Gateway file gateway with an NFS mount poin
- C. Mount the file gateway file system on the on-premises serve
- D. Configure a process to periodically copy the images to the mount point
- E. Deploy an AWS DataSync agent to an on-premises server that has access to the NFS file system Send data over the Direct Connect connection to an S3 bucket by using a public VIF Configure an AWS Lambda function to process event notifications from Amazon S3 and copy the images from Amazon S3 to the EFS file system
- F. Deploy an AWS DataSync agent to an on-premises server that has access to the NFS file system Send data over the Direct Connect connection to an AWS PrivateLink interface VPC endpoint for Amazon EFS by using a private VIF Configure a DataSync scheduled task to send the images to the EFS file system every 24 hours.

Answer: A

NEW QUESTION 191

- (Exam Topic 2)

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

Which CodeDeploy configuration should the solutions architect use?

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

Answer: C

NEW QUESTION 194

- (Exam Topic 2)

A company's solutions architect is reviewing a web application that runs on AWS. The application references static assets in an Amazon S3 bucket in the us-east-1 Region. The company needs resiliency across multiple AWS Regions. The company already has created an S3 bucket in a second Region.

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

- A. Configure the application to write each object to both S3 bucket
- B. Set up an Amazon Route 53 public hosted zone with a record set by using a weighted routing policy for each S3 bucke
- C. Configure the application to reference the objects by using the Route 53 DNS name.
- D. Create an AWS Lambda function to copy objects from the S3 bucket in us-east-1 to the S3 bucket in the second Regio
- E. Invoke the Lambda function each time an object is written to the S3 bucket in us-east-1. Set up an Amazon CloudFront distribution with an origin group that contains the two S3 buckets as origins.
- F. Configure replication on the S3 bucket in us-east-1 to replicate objects to the S3 bucket in the second Region Set up an Amazon CloudFront distribution with an origin group that contains the two S3 bucketsas origins.
- G. Configure replication on the S3 bucket in us-east-1 to replicate objects to the S3 bucket in the second Regio
- H. If failover is required, update the application code to load S3 objects from the S3 bucket in the second Region.

Answer: D

NEW QUESTION 197

- (Exam Topic 2)

A financial services company in North America plans to release a new online web application to its customers on AWS . The company will launch the application in the us-east-1 Region on Amazon EC2 instances. The application must be highly available and must dynamically scale to meet user traffic. The company also wants to implement a disaster recovery environment for the application in the us-west-1 Region by using active-passive failover.

Which solution will meet these requirements?

- A. Create a VPC in us-east-1 and a VPC in us-west-1 Configure VPC peering In the us-east-1 VP
- B. create an Application Load Balancer (ALB) that extends across multiple Availability Zones in both VPCs Create an Auto Scaling group that deploys the EC2 instances across the multiple Availability Zones in both VPCs Place the Auto Scaling group behind the ALB.
- C. Create a VPC in us-east-1 and a VPC in us-west-1. In the us-east-1 VP
- D. create an Application Load Balancer (ALB) that extends across multiple Availability Zones in that VP
- E. Create an Auto Scaling group that deploys the EC2 instances across the multiple Availability Zones in the us-east-1 VPC Place the Auto Scaling group behind the ALB Set up the same configuration in the us-west-1 VP
- F. Create an Amazon Route 53 hosted zone Create separate records for each ALB Enable health checks to ensure high availability between Regions.
- G. Create a VPC in us-east-1 and a VPC in us-west-1 In the us-east-1 VP

- H. create an Application Load Balancer (ALB) that extends across multiple Availability Zones in that VPC Create an Auto Scaling group that deploys the EC2 instances across the multiple Availability Zones in the us-east-1 VPC Place the Auto Scaling group behind the ALB Set up the same configuration in the us-west-1 VPC Create an Amazon Route 53 hosted zon
- I. Create separate records for each ALB Enable health checks and configure a failover routing policy for each record.
- J. Create a VPC in us-east-1 and a VPC in us-west-1 Configure VPC peering In the us-east-1 VP
- K. create an Application Load Balancer (ALB) that extends across multiple Availability Zones in Create an Auto Scaling group that deploys the EC2 instances across the multiple Availability Zones in both VPCs Place the Auto Scaling group behind the ALB Create an Amazon Route 53 host.. Create a record for the ALB.

Answer: C

NEW QUESTION 200

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questions and Answers in PDF Format

SAP-C02 Practice Exam Features:

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

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