



Amazon

Exam Questions AWS-Certified-Solutions-Architect-Professional

Amazon AWS Certified Solutions Architect Professional

NEW QUESTION 1

- (Exam Topic 1)

A retail company is operating its ecommerce application on AWS. The application runs on Amazon EC2 instances behind an Application Load Balancer (ALB). The company uses an Amazon RDS DB instance as the database backend. Amazon CloudFront is configured with one origin that points to the ALB. Static content is cached. Amazon Route 53 is used to host all public zones.

After an update of the application, the ALB occasionally returns a 502 status code (Bad Gateway) error. The root cause is malformed HTTP headers that are returned to the ALB. The webpage returns successfully when a solutions architect reloads the webpage immediately after the error occurs.

While the company is working on the problem, the solutions architect needs to provide a custom error page instead of the standard ALB error page to visitors. Which combination of steps will meet this requirement with the LEAST amount of operational overhead? (Choose two.)

- A. Create an Amazon S3 bucket
- B. Configure the S3 bucket to host a static webpage
- C. Upload the custom error pages to Amazon S3.
- D. Create an Amazon CloudWatch alarm to invoke an AWS Lambda function if the ALB health check response Target.FailedHealthChecks is greater than 0. Configure the Lambda function to modify the forwarding rule at the ALB to point to a publicly accessible web server.
- E. Modify the existing Amazon Route 53 records by adding health check
- F. Configure a fallback target if the health check fail
- G. Modify DNS records to point to a publicly accessible webpage.
- H. Create an Amazon CloudWatch alarm to invoke an AWS Lambda function if the ALB health check response Elb.InternalError is greater than 0. Configure the Lambda function to modify the forwarding rule at the ALB to point to a public accessible web server.
- I. Add a custom error response by configuring a CloudFront custom error page
- J. Modify DNS records to point to a publicly accessible web page.

Answer: AE

Explanation:

"Save your custom error pages in a location that is accessible to CloudFront. We recommend that you store them in an Amazon S3 bucket, and that you don't store them in the same place as the rest of your website or application's content. If you store the custom error pages on the same origin as your website or application, and the origin starts to return 5xx errors, CloudFront can't get the custom error pages because the origin server is unavailable."

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/GeneratingCustomErrorResponses.htm>

NEW QUESTION 2

- (Exam Topic 1)

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

Which solution will meet these requirements?

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

Answer: C

Explanation:

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

NEW QUESTION 3

- (Exam Topic 1)

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

Which solution meets these requirements?

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

Answer: B

Explanation:

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

NEW QUESTION 4

- (Exam Topic 1)

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

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

Answer: D

NEW QUESTION 5

- (Exam Topic 1)

A company needs to store and process image data that will be uploaded from mobile devices using a custom mobile app. Usage peaks between 8 AM and 5 PM on weekdays, with thousands of uploads per minute. The app is rarely used at any other time. A user is notified when image processing is complete. Which combination of actions should a solutions architect take to ensure image processing can scale to handle the load? (Select THREE.)

- A. Upload files from the mobile software directly to Amazon S3. Use S3 event notifications to create a message in an Amazon MQ queue.
- B. Upload files from the mobile software directly to Amazon S3. Use S3 event notifications to create a message in an Amazon Simple Queue Service (Amazon SQS) standard queue.
- C. Invoke an AWS Lambda function to perform image processing when a message is available in the queue.
- D. Invoke an S3 Batch Operations job to perform image processing when a message is available in the queue.
- E. Send a push notification to the mobile app by using Amazon Simple Notification Service (Amazon SNS) when processing is complete.
- F. Send a push notification to the mobile app by using Amazon Simple Email Service (Amazon SES) when processing is complete.

Answer: BCE

Explanation:

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/batch-ops-basics.html>

NEW QUESTION 6

- (Exam Topic 1)

A developer reports receiving an Error 403: Access Denied message when they try to download an object from an Amazon S3 bucket. The S3 bucket is accessed using an S3 endpoint inside a VPC, and is encrypted with an AWS KMS key. A solutions architect has verified that the developer is assuming the correct IAM role in the account that allows the object to be downloaded. The S3 bucket policy and the NACL are also valid. Which additional step should the solutions architect take to troubleshoot this issue?

- A. Ensure that blocking all public access has not been enabled in the S3 bucket.
- B. Verify that the IAM role has permission to decrypt the referenced KMS key.
- C. Verify that the IAM role has the correct trust relationship configured.
- D. Check that local firewall rules are not preventing access to the S3 endpoint.

Answer: B

NEW QUESTION 7

- (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 8

- (Exam Topic 1)

A company has a new application that needs to run on five Amazon EC2 instances in a single AWS Region. The application requires high-throughput, low-latency

network connections between all of the EC2 instances where the application will run. There is no requirement for the application to be fault tolerant. Which solution will meet these requirements?

- A. Launch five new EC2 instances into a cluster placement group
- B. Ensure that the EC2 instance type supports enhanced networking.
- C. Launch five new EC2 instances into an Auto Scaling group in the same Availability Zone
- D. Attach an extra elastic network interface to each EC2 instance.
- E. Launch five new EC2 instances into a partition placement group
- F. Ensure that the EC2 instance type supports enhanced networking.
- G. Launch five new EC2 instances into a spread placement group
- H. Attach an extra elastic network interface to each EC2 instance.

Answer: A

Explanation:

When you launch EC2 instances in a cluster they benefit from performance and low latency. No redundancy though as per the question
<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/placement-groups.html>.

NEW QUESTION 9

- (Exam Topic 1)

A company has an application that generates reports and stores them in an Amazon S3 bucket. When a user accesses their report, the application generates a signed URL to allow the user to download the report. The company's security team has discovered that the files are public and that anyone can download them without authentication. The company has suspended the generation of new reports until the problem is resolved.

Which set of actions will immediately remediate the security issue without impacting the application's normal workflow?

- A. Create an AWS Lambda function that applies a deny all policy for users who are not authenticated. Create a scheduled event to invoke the Lambda function.
- B. Review the AWS Trusted Advisor bucket permissions check and implement the recommended actions.
- C. Run a script that puts a private ACL on all of the objects in the bucket.
- D. Use the Block Public Access feature in Amazon S3 to set the IgnorePublicAcls option to TRUE on the bucket.

Answer: D

Explanation:

The S3 bucket is allowing public access and this must be immediately disabled. Setting the IgnorePublicAcls option to TRUE causes Amazon S3 to ignore all public ACLs on a bucket and any objects that it contains. The other settings you can configure with the Block Public Access Feature are:

- o BlockPublicAcls – PUT bucket ACL and PUT objects requests are blocked if granting public access.
- o BlockPublicPolicy – Rejects requests to PUT a bucket policy if granting public access.
- o RestrictPublicBuckets – Restricts access to principles in the bucket owners' AWS account. <https://aws.amazon.com/s3/features/block-public-access/>

NEW QUESTION 10

- (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 10

- (Exam Topic 1)

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

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

Which strategy should a solutions architect use?

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

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

Answer: B

Explanation:

IP reputation rule groups allow you to block requests based on their source. Choose one or more of these rule groups if you want to reduce your exposure to BOTS!!!! traffic or exploitation attempts
 The Amazon IP reputation list rule group contains rules that are based on Amazon internal threat intelligence. This is useful if you would like to block IP addresses typically associated with bots or other threats. Inspects for a list of IP addresses that have been identified as bots by Amazon threat intelligence.

NEW QUESTION 13

- (Exam Topic 1)

A company is running an application on Amazon EC2 instances in three environments; development, testing, and production. The company uses AMIs to deploy the EC2 instances. The company builds the AMIs by using custom deployment scripts and infrastructure orchestration tools for each release in each environment. The company is receiving errors in its deployment process. Errors appear during operating system package downloads and during application code installation from a third-party Git hosting service. The company needs deployments to become more reliable across all environments. Which combination of steps will meet these requirements? (Select THREE).

- A. Mirror the application code to an AWS CodeCommit Git repositior
- B. Use the repository to build EC2 AMIs.
- C. Produce multiple EC2 AMI
- D. one for each environment, for each release.
- E. Produce one EC2 AMI for each release for use across all environments.
- F. Mirror the application code to a third-party Git repository that uses Amazon S3 storag
- G. Use the repository for deployment.
- H. Replace the custom scripts and tools with AWS CodeBuil
- I. Update the infrastructure deployment process to use EC2 Image Builder.

Answer: ACE

NEW QUESTION 17

- (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 queu
- 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 queu
- D. Configure an AWS Fargate container application to automatically scale to a single instance when the SQS queue contains message
- E. Have the application process each record, and transform the record into JSON forma
- F. When the queue is empty, send the results to another S3 bucket for internal processing and scale down the AWS Fargate instance.
- G. Create an AWS Glue crawler and custom classifier based on the data feed formats and build a table definition to matc
- H. 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 JSO
- 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 matc
- 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 JSO
- 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 22

- (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 26

- (Exam Topic 1)

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

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

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

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

Answer: A

Explanation:

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

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

NEW QUESTION 31

- (Exam Topic 1)

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

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

Which solution will meet these requirements?

- A. Disassociate the Elastic IP address from the EC2 instance
- 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 internet-facing endpoint
- I. 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 34

- (Exam Topic 1)

A media company uses Amazon DynamoDB to store metadata for its catalog of movies that are available to stream. Each media item contains user-facing content that concludes a description of the media, a list of search tags, and similar data. In addition, media items include a list of Amazon S3 key names that relate to movie files. The company stores these movie files in a single S3 bucket that has versioning enabled. The company uses Amazon CloudFront to serve these movie files. The company has 100,000 media items, and each media item can have many different S3 objects that represent different encodings of the same media. S3 objects that belong to the same media item are grouped together under the same key prefix, which is a random unique ID. Because of an expiring contract with a media provider, the company must remove 2,000 media items. The company must completely delete all DynamoDB keys and movie files on Amazon S3 that are related to these media items within 36 hours. The company must ensure that the content cannot be recovered. Which combination of actions will meet these requirements? (Select TWO.)

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

Answer: CE

NEW QUESTION 36

- (Exam Topic 1)

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

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

Answer: D

NEW QUESTION 40

- (Exam Topic 1)

A travel company built a web application that uses Amazon Simple Email Service (Amazon SES) to send email notifications to users. The company needs to enable logging to help troubleshoot email delivery issues. The company also needs the ability to do searches that are based on recipient, subject, and time sent. Which combination of steps should a solutions architect take to meet these requirements? (Select TWO.)

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

Answer: AC

Explanation:

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

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

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

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

NEW QUESTION 44

- (Exam Topic 1)

A company is storing data in several Amazon DynamoDB tables. A solutions architect must use a serverless architecture to make the data accessible publicly through a simple API over HTTPS. The solution must scale automatically in response to demand. Which solutions meet these requirements? (Choose two.)

- A. Create an Amazon API Gateway REST API
- B. Configure this API with direct integrations to DynamoDB by using API Gateway's AWS integration type.
- C. Create an Amazon API Gateway HTTP API
- D. Configure this API with direct integrations to DynamoDB by using API Gateway's AWS integration type.
- E. Create an Amazon API Gateway HTTP API
- F. Configure this API with integrations to AWS Lambda functions that return data from the DynamoDB tables.
- G. Create an accelerator in AWS Global Accelerator

- H. Configure this accelerator with AWS Lambda@Edge function integrations that return data from the DynamoDB tables.
- I. Create a Network Load Balance
- J. Configure listener rules to forward requests to the appropriate AWS Lambda functions

Answer: CD

Explanation:

<https://docs.aws.amazon.com/apigateway/latest/developerguide/http-api-dynamo-db.html>

NEW QUESTION 49

- (Exam Topic 1)

A solutions architect is designing a publicly accessible web application that is on an Amazon CloudFront distribution with an Amazon S3 website endpoint as the origin. When the solution is deployed, the website returns an Error 403: Access Denied message.

Which steps should the solutions architect take to correct the issue? (Select TWO.)

- A. Remove the S3 block public access option from the S3 bucket.
- B. Remove the requester pays option from the S3 bucket.
- C. Remove the origin access identity (OAI) from the CloudFront distribution.
- D. Change the storage class from S3 Standard to S3 One Zone-Infrequent Access (S3 One Zone-IA).
- E. Disable S3 object versioning.

Answer: AB

Explanation:

See using S3 to host a static website with Cloudfront: <https://aws.amazon.com/premiumsupport/knowledge-center/cloudfront-serve-static-website/>

- Using a REST API endpoint as the origin, with access restricted by an origin access identity (OAI)
- Using a website endpoint as the origin, with anonymous (public) access allowed
- Using a website endpoint as the origin, with access restricted by a Referer header

NEW QUESTION 52

- (Exam Topic 1)

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

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

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

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

Answer: B

NEW QUESTION 56

- (Exam Topic 1)

A large payroll company recently merged with a small staffing company. The unified company now has multiple business units, each with its own existing AWS account.

A solutions architect must ensure that the company can centrally manage the billing and access policies for all the AWS accounts. The solutions architect configures AWS Organizations by sending an invitation to all member accounts of the company from a centralized management account.

What should the solutions architect do next to meet these requirements?

- A. Create the OrganizationAccountAccess IAM group in each member account
- B. Include the necessary IAM roles for each administrator.
- C. Create the OrganizationAccountAccessPolicy IAM policy in each member account
- D. Connect the member accounts to the management account by using cross-account access.
- E. Create the OrganizationAccountAccessRole IAM role in each member account
- F. Grant permission to the management account to assume the IAM role.
- G. Create the OrganizationAccountAccessRole IAM role in the management account Attach the Administrator Access AWS managed policy to the IAM role
- H. Assign the IAM role to the administrators in each member account.

Answer: C

NEW QUESTION 61

- (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 64

- (Exam Topic 1)

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

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

Answer: B

NEW QUESTION 65

- (Exam Topic 1)

A company uses a service to collect metadata from applications that the company hosts on premises. Consumer devices such as TVs and internet radios access the applications. Many older devices do not support certain HTTP headers and exhibit errors when these headers are present in responses. The company has configured an on-premises load balancer to remove the unsupported headers from responses sent to older devices, which the company identified by the User-Agent headers.

The company wants to migrate the service to AWS, adopt serverless technologies, and retain the ability to support the older devices. The company has already migrated the applications into a set of AWS Lambda functions.

Which solution will meet these requirements?

- A. Create an Amazon CloudFront distribution for the metadata servic
- B. Create an Application Load Balancer (ALB). Configure the CloudFront distribution to forward requests to the AL
- C. Configure the ALB to invoke the correct Lambda function for each type of reques
- D. Create a CloudFront function to remove the problematic headers based on the value of the User-Agent header.
- E. Create an Amazon API Gateway REST API for the metadata servic
- F. Configure API Gateway to invoke the correct Lambda function for each type of reques
- G. Modify the default gateway responses to remove the problematic headers based on the value of the User-Agent header.
- H. Create an Amazon API Gateway HTTP API for the metadata servic
- I. Configure API Gateway to invoke the correct Lambda function for each type of reques
- J. Create a response mapping template to remove the problematic headers based on the value of the User-Agen
- K. Associate the response data mapping withthe HTTP API.
- L. Create an Amazon CloudFront distribution for the metadata servic
- M. Create an Application Load Balancer (ALB). Configure the CloudFront distribution to forward requests to the AL
- N. Configure the ALB to invoke the correct Lambda function for each type of reques
- O. Create a Lambda@Edge function that will remove the problematic headers in response to viewer requests based on the value of theUser-Agent header.

Answer: D

Explanation:

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

NEW QUESTION 70

- (Exam Topic 1)

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

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

J. Have the operators log in to running instances and go through a manual test plan to verify the application is running as expected.

Answer: B

Explanation:

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

NEW QUESTION 71

- (Exam Topic 1)

A company is hosting a single-page web application in the AWS Cloud. The company is using Amazon CloudFront to reach its goal audience. The CloudFront distribution has an Amazon S3 bucket that is configured as its origin. The static files for the web application are stored in this S3 bucket. The company has used a simple routing policy to configure an Amazon Route 53 A record. The record points to the CloudFront distribution. The company wants to use a canary deployment release strategy for new versions of the application. What should a solutions architect recommend to meet these requirements?

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

Answer: A

NEW QUESTION 72

- (Exam Topic 1)

A company is planning to set up a REST API application on AWS. The application team wants to set up a new identity store on AWS. The IT team does not want to maintain any infrastructure or servers for this deployment. What is the MOST operationally efficient solution that meets these requirements?

- A. Deploy the application as AWS Lambda function
- B. Set up Amazon API Gateway REST API endpoints for the application. Create a Lambda function, and configure a Lambda authorizer.
- C. Deploy the application in AWS AppSync, and configure AWS Lambda resolvers. Set up an Amazon Cognito user pool, and configure AWS AppSync to use the user pool for authorization.
- D. Deploy the application as AWS Lambda function.
- E. Set up Amazon API Gateway REST API endpoints for the application. Set up an Amazon Cognito user pool, and configure an Amazon Cognito authorizer.
- F. Deploy the application in Amazon Elastic Kubernetes Service (Amazon EKS) cluster.
- G. Set up an Application Load Balancer for the EKS pods. Set up an Amazon Cognito user pool and service pod for authentication.

Answer: C

NEW QUESTION 75

- (Exam Topic 1)

An online e-commerce business is running a workload on AWS. The application architecture includes a web tier, an application tier for business logic, and a database tier for user and transactional data management. The database server has a 100 GB memory requirement. The business requires cost-efficient disaster recovery for the application with an RTO of 5 minutes and an RPO of 1 hour. The business also has a regulatory requirement for out-of-region disaster recovery with a minimum distance between the primary and alternate sites of 250 miles. Which of the following options can the solutions architect design to create a comprehensive solution for this customer that meets the disaster recovery requirements?

- A. Back up the application and database data frequently and copy them to Amazon S3. Replicate the backups using S3 cross-region replication, and use AWS CloudFormation to instantiate infrastructure for disaster recovery and restore data from Amazon S3.
- B. Employ a pilot light environment in which the primary database is configured with mirroring to build a standby database on m4.large in the alternate region.
- C. Use AWS CloudFormation to instantiate the web servers, application servers, and load balancers in case of a disaster to bring the application up in the alternate region.
- D. Vertically resize the database to meet the full production demands, and use Amazon Route 53 to switch traffic to the alternate region.
- E. Use a scaled-down version of the fully functional production environment in the alternate region that includes one instance of the web server, one instance of the application server, and a replicated instance of the database server in standby mode.
- F. Place the web and the application tiers in an Auto Scaling group behind a load balancer, which can automatically scale when the load arrives to the application.
- G. Use Amazon Route 53 to switch traffic to the alternate region.
- H. Employ a multi-region solution with fully functional web
- I. application, and database tiers in both regions with equivalent capacity.
- J. Activate the primary database in one region only and the standby database in the other region.
- K. Use Amazon Route 53 to automatically switch traffic from one region to another using health check routing policies.

Answer: C

Explanation:

As RTO is in minutes

(<https://docs.aws.amazon.com/wellarchitected/latest/reliability-pillar/plan-for-disaster-recovery-dr.html>) Warm standby (RPO in seconds, RTO in minutes): Maintain a scaled-down version of a fully functional environment always running in the DR Region. Business-critical systems are fully duplicated and are always on, but with a scaled-down fleet. When the time comes for recovery, the system is scaled up quickly to handle the production load.

NEW QUESTION 80

- (Exam Topic 1)

A solutions architect has an operational workload deployed on Amazon EC2 instances in an Auto Scaling group. The VPC architecture spans two Availability

Zones (AZ) with a subnet in each that the Auto Scaling group is targeting. The VPC is connected to an on-premises environment and connectivity cannot be interrupted. The maximum size of the Auto Scaling group is 20 instances in service. The VPC IPv4 addressing is as follows:

VPC CIDR: 10.0.0.0/23

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

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

Which solution will meet these requirements?

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

Answer: A

Explanation:

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

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

NEW QUESTION 83

- (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 84

- (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 89

- (Exam Topic 1)

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

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

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

Answer: D

Explanation:

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

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

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

NEW QUESTION 93

- (Exam Topic 1)

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

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

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

Answer: C

Explanation:

Aurora Auto Scaling enables your Aurora DB cluster to handle sudden increases in connectivity or workload. When the connectivity or workload decreases, Aurora Auto Scaling removes unnecessary Aurora Replicas so that you don't pay for unused provisioned DB instances

NEW QUESTION 94

- (Exam Topic 1)

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

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

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

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

Answer: A

Explanation:

<https://aws.amazon.com/blogs/storage/new-enhancements-for-moving-data-between-amazon-fsx-for-lustre-and>

NEW QUESTION 98

- (Exam Topic 1)

A financial services company logs personally identifiable information in 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. Disable this CMK
- L. and overwrite the key material with the key material from the on-premises HSM using the public key and import token provided by AWS
- M. Re-enable the CMK
- N. Enable automatic key rotation on the CMK with a duration of 1 year
- O. 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 99

- (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 103

- (Exam Topic 1)

An education company is running a web application used by college students around the world. The application runs in an Amazon Elastic Container Service (Amazon ECS) cluster in an Auto Scaling group behind an Application Load Balancer (ALB). A system administrator detects a weekly spike in the number of failed login attempts, which overwhelm the application's authentication service. All the failed login attempts originate from about 500 different IP addresses that change each week. A solutions architect must prevent the failed login attempts from overwhelming the authentication service.

Which solution meets these requirements with the MOST operational efficiency?

- A. Use AWS Firewall Manager to create a security group and security group policy to deny access from the IP addresses.
- B. Create an AWS WAF web ACL with a rate-based rule, and set the rule action to Block
- C. Connect the web ACL to the ALB.
- D. Use AWS Firewall Manager to create a security group and security group policy to allow access only to specific CIDR ranges.
- E. Create an AWS WAF web ACL with an IP set match rule, and set the rule action to Block
- F. Connect the web ACL to the ALB.

Answer: B

Explanation:

<https://docs.aws.amazon.com/waf/latest/developerguide/waf-rule-statement-type-rate-based.html>

The IP set match statement inspects the IP address of a web request against a set of IP addresses and address ranges. Use this to allow or block web requests based on the IP addresses that the requests originate from. By default, AWS WAF uses the IP address from the web request origin, but you can configure the rule to use an HTTP header like X-Forwarded-For instead.

<https://docs.aws.amazon.com/waf/latest/developerguide/waf-rule-statement-type-ipset-match.html>

<https://docs.aws.amazon.com/waf/latest/developerguide/waf-rule-statement-type-rate-based.html>

NEW QUESTION 105

- (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 mod
- 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 DynamoD
- E. Use an AWS KMS key to encrypt DynamoDB objects at rest.
- F. Deploy instances with Amazon EBS volumes attached to store this dat
- G. Use E8S 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 109

- (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 lo single-instance AWS Elastic Beanstalk environments without a load balancer.
- B. Use AWS SMS to create AMLs 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 AMLs for each virtual machine and run them in single-instance AWS Elastic Beanstalk environments by configuring a custom image.

Answer: C

NEW QUESTION 114

- (Exam Topic 1)

A company is building a hybrid solution between its existing on-premises systems and a new backend in AWS. The company has a management application to monitor the state of its current IT infrastructure and automate responses to issues. The company wants to incorporate the status of its consumed AWS services into the application. The application uses an HTTPS endpoint to receive updates.

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

- A. Configure AWS Systems Manager OpsCenter to ingest operational events from the on-premises systems Retire the on-premises management application and adopt OpsCenter as the hub
- B. Configure Amazon EventBridge (Amazon CloudWatch Events) to detect and react to changes for AWS Health events from the AWS Personal Health Dashboard Configure the EventBridge (CloudWatch Events) event to publish a message to an Amazon Simple Notification Service (Amazon SNS) topic and subscribe the topic to the HTTPS endpoint of the management application
- C. Modify the on-premises management application to call the AWS Health API to poll for status events of AWS services.
- D. Configure Amazon EventBridge (Amazon CloudWatch Events) to detect and react to changes for AWS Health events from the AWS Service Health Dashboard Configure the EventBridge (CloudWatch Events) event to publish a message to an Amazon Simple Notification Service (Amazon SNS) topic and subscribe the topic to an HTTPS endpoint for the management application with a topic filter corresponding to the services being used

Answer: A

Explanation:

ALB & NLB both supports IPs as targets. Questions is based on TCP traffic over VPN to on-premise. TCP is layer 4 and the , load balancer should be NLB. Then next questions does NLB supports loadbalcning traffic over VPN. And answer is YEs based on below URL.

<https://aws.amazon.com/about-aws/whats-new/2018/09/network-load-balancer-now-supports-aws-vpn/>

Target as IPs for NLB & ALB: <https://aws.amazon.com/elasticloadbalancing/faqs/?nc=sn&loc=5> <https://aws.amazon.com/elasticloadbalancing/application-load-balancer/>

NEW QUESTION 116

- (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 lo host database dusters 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 OpsConter Configure visualizations tor Amazon CloudWatch metrics that are associated with the EC2 instances and their EBS volumes Review the dashboard periodically and identify usage patterns Right size 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 Right size the FC? 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 right size 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 117

- (Exam Topic 2)

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

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

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

Answer: D

NEW QUESTION 119

- (Exam Topic 2)

A company is building a software-as-a-service (SaaS) solution on AWS. The company has deployed an Amazon API Gateway REST API with AWS Lambda integration in multiple AWS Regions and in the same production account.

The company offers tiered pricing that gives customers the ability to pay for the capacity to make a certain number of API calls per second. The premium tier offers up to 3,000 calls per second, and customers are identified by a unique API key. Several premium tier customers in various Regions report that they receive error responses of 429 Too Many Requests from multiple API methods during peak usage hours. Logs indicate that the Lambda function is never invoked.

What could be the cause of the error messages for these customers?

- A. The Lambda function reached its concurrency limit.
- B. The Lambda function its Region limit for concurrency.
- C. The company reached its API Gateway account limit for calls per second.
- D. The company reached its API Gateway default per-method limit for calls per second.

Answer: C

Explanation:

<https://docs.aws.amazon.com/apigateway/latest/developerguide/api-gateway-request-throttling.html#apig-reqs>

NEW QUESTION 120

- (Exam Topic 2)

A company plans to refactor a monolithic application into a modern application designed deployed on AWS. The CI/CD pipeline needs to be upgraded to support the modern design for the application with the following requirements

- It should allow changes to be released several times every hour.
- * It should be able to roll back the changes as quickly as possible Which design will meet these requirements?

- A. Deploy a CI-CD pipeline that incorporates AMIs to contain the application and their configurations Deploy the application by replacing Amazon EC2 instances
- B. Specify AWS Elastic Beanstalk to serve as a secondary environment as the deployment target for the CI/CD pipeline of the application
- C. To deploy swap the staging and production environment URLs.
- D. Use AWS Systems Manager to re-provision the infrastructure for each deployment Update the Amazon EC2 user data to pull the latest code artifact from Amazon S3 and use Amazon Route 53 weighted routing to point to the new environment
- E. Roll out application updates as part of an Auto Scaling event using prebuilt AMI
- F. Use new versions of the AMIs to add instances, and phase out all instances that use the previous AMI version with the configured termination policy during a deployment event.

Answer: B

Explanation:

It is the fastest when it comes to rollback and deploying changes every hour

NEW QUESTION 122

- (Exam Topic 2)

A company wants to send data from its on-premises systems to Amazon S3 buckets. The company created the S3 buckets in three different accounts. The company must send the data privately without the data traveling across the internet. The company has no existing dedicated connectivity to AWS Which combination of steps should a solutions architect take to meet these requirements? (Select TWO.)

- A. Establish a networking account in the AWS Cloud Create a private VPC in the networking account Set up an AWS Direct Connect connection with a private VIF between the on-premises environment and the private VPC
- B. Establish a networking account in the AWS Cloud Create a private VPC in the networking account Set up an AWS Direct Connect connection with a public VIF between the on-premises environment and the private VPC
- C. Create an Amazon S3 interface endpoint in the networking account
- D. Create an Amazon S3 gateway endpoint in the networking account
- E. Establish a networking account in the AWS Cloud
- F. Create a private VPC in the networking account Peer VPCs from the accounts that host the S3 buckets with the VPC in the network account

Answer: AD

NEW QUESTION 125

- (Exam Topic 2)

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

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

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

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

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

Answer: A

NEW QUESTION 127

- (Exam Topic 2)

A solutions architect needs to implement a client-side encryption mechanism for objects that will be stored in a new Amazon S3 bucket. The solutions architect created a CMK that is stored in AWS Key Management Service (AWS KMS) for this purpose.

The solutions architect created the following IAM policy and attached it to an IAM role:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "DownloadUpload",
      "Action": [
        "s3:GetObject",
        "s3:GetObjectVersion",
        "s3:PutObject",
        "s3:PutObjectAcl"
      ],
      "Effect": "Allow",
      "Resource": "arn:aws:s3:::BucketName/*"
    },
    {
      "Sid": "KMSAccess",
      "Action": [
        "kms:Decrypt",
        "kms:Encrypt"
      ],
      "Effect": "Allow",
      "Resource": "arn:aws:kms:Region:Account:key/Key ID"
    }
  ]
}
```

During tests, the solutions architect was able to successfully get existing test objects in the S3 bucket. However, attempts to upload a new object resulted in an error message. The error message stated that the action was forbidden.

Which action must the solutions architect add to the IAM policy to meet all the requirements?

- A. kms:GenerateDataKey
- B. kms:GetKeyPolicy
- C. kms:GetPublicKey
- D. kms:SKJn

Answer: A

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/s3-access-denied-error-kms/>

"An error occurred (AccessDenied) when calling the PutObject operation: Access Denied" This error message indicates that your IAM user or role needs permission for the kms:GenerateDataKey action.

NEW QUESTION 130

- (Exam Topic 2)

A company runs a content management application on a single Windows Amazon EC2 instance in a development environment. The application reads and writes static content to a 2 TB Amazon Elastic Block Store (Amazon EBS) volume that is attached to the instance as the root device. The company plans to deploy this application in production as a highly available and fault-tolerant solution that runs on at least three EC2 instances across multiple Availability Zones.

A solutions architect must design a solution that joins all the instances that run the application to an Active Directory domain. The solution also must implement Windows ACLs to control access to file contents. The application always must maintain exactly the same content on all running instances at any given point in time.

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

- A. Create an Amazon Elastic File System (Amazon EFS) file share
- B. Create an Auto Scaling group that extends across three Availability Zones and maintains a minimum size of three instance
- C. Implement a user data script to install the application, join the instance to the AD domain, and mount the EFS file share.
- D. Create a new AMI from the current EC2 instance that is running
- E. Create an Amazon FSx for Lustre file system
- F. Create an Auto Scaling group that extends across three Availability Zones and maintains a minimum size of three instance
- G. Implement a user data script to join the instance to the AD domain and mount the FSx for Lustre file system.
- H. Create an Amazon FSx for Windows File Server file system
- I. Create an Auto Scaling group that extends across three Availability Zones and maintains a minimum size of three instance
- J. Implement a user data script to install the application and mount the FSx for Windows File Server file system
- K. Perform a seamless domain join to join the instance to the AD domain.
- L. Create a new AMI from the current EC2 instance that is running
- M. Create an Amazon Elastic File System (Amazon EFS) file system
- N. Create an Auto Scaling group that extends across three Availability Zones and maintains a minimum size of three instance
- O. Perform a seamless domain join to join the instance to the AD domain.

Answer: B

NEW QUESTION 133

- (Exam Topic 2)

A solutions architect wants to make sure that only AWS users or roles with suitable permissions can access a new Amazon API Gateway endpoint. The solutions architect wants an end-to-end view of each request to analyze the latency of the request and create service maps. How can the solutions architect design the API Gateway access control and perform request inspections?

- A. For the API Gateway method, set the authorization to AWS_IAM. Then, give the IAM user or role execute-api:Invoke permission on the REST API resource. Enable the API caller to sign requests with AWS Signature when accessing the endpoint. Use AWS X-Ray to trace and analyze user requests to API Gateway.
- B. For the API Gateway resource, set CORS to enabled and only return the company's domain in Access-Control-Allow-Origin headers. Then give the IAM user or role execute-api:Invoke permission on the REST API resource. Use Amazon CloudWatch to trace and analyze user requests to API Gateway.
- C. Create an AWS Lambda function as the custom authorizer. Ask the API client to pass the key and secret when making the call, and then use Lambda to validate the key/secret pair against the IAM system. Use AWS X-Ray to trace and analyze user requests to API Gateway.
- D. Create a client certificate for API Gateway. Distribute the certificate to the AWS users and roles that need to access the endpoint. Enable the API caller to pass the client certificate when accessing the endpoint.
- E. Use Amazon CloudWatch to trace and analyze user requests to API Gateway.

Answer: A

NEW QUESTION 134

- (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 behavior.
- C. Then update the cache behavior 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 behavior. Enable automatic object compression and use Lambda@Edge viewer request events for user authorization.
- E. Create two cache behaviors for static and dynamic content. Remove the User-Agent HTTP header 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.

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 138

- (Exam Topic 2)

During an audit, a security team discovered that a development team was putting IAM user secret access keys in their code and then committing it to an AWS CodeCommit repository. The security team wants to automatically find and remediate instances of this security vulnerability. Which solution will ensure that the credentials are appropriately secured automatically?

- A. Run a script nightly using AWS Systems Manager Run Command to search for credentials on the development instances. If found, use AWS Secrets Manager to rotate the credentials.
- B. Use a scheduled AWS Lambda function to download and scan the application code from CodeCommit. If credentials are found, generate new credentials and store them in AWS KMS.

- C. Configure Amazon Macie to scan for credentials in CodeCommit repositories. If credentials are found, trigger an AWS Lambda function to disable the credentials and notify the user.
- D. Configure a CodeCommit trigger to invoke an AWS Lambda function to scan new code submissions for credentials. If credentials are found, disable them in AWS IAM and notify the user.

Answer: A

NEW QUESTION 139

- (Exam Topic 2)

A company has a latency-sensitive trading platform that uses Amazon DynamoDB as a storage backend. The company configured the DynamoDB table to use on-demand capacity mode. A solutions architect needs to design a solution to improve the performance of the trading platform. The new solution must ensure high availability for the trading platform.

Which solution will meet these requirements with the LEAST latency?

- A. Create a two-node DynamoDB Accelerator (DAX) cluster. Configure an application to read and write data by using DAX.
- B. Create a three-node DynamoDB Accelerator (DAX) cluster.
- C. Configure an application to read data by using DAX and to write data directly to the DynamoDB table.
- D. Create a three-node DynamoDB Accelerator (DAX) cluster.
- E. Configure an application to read data directly from the DynamoDB table and to write data by using DAX.
- F. Create a single-node DynamoDB Accelerator (DAX) cluster.
- G. Configure an application to read data by using DAX and to write data directly to the DynamoDB table.

Answer: A

NEW QUESTION 142

- (Exam Topic 2)

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

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

Answer: A

NEW QUESTION 147

- (Exam Topic 2)

A company is running an application in the AWS Cloud. Recent application metrics show inconsistent response times and a significant increase in error rates. Calls to third-party services are causing the delays. Currently, the application calls third-party services synchronously by directly invoking an AWS Lambda function. A solutions architect needs to decouple the third-party service calls and ensure that all the calls are eventually completed.

Which solution will meet these requirements?

- A. Use an Amazon Simple Queue Service (Amazon SQS) queue to store events and invoke the Lambda function.
- B. Use an AWS Step Functions state machine to pass events to the Lambda function.
- C. Use an Amazon EventBridge rule to pass events to the Lambda function.
- D. Use an Amazon Simple Notification Service (Amazon SNS) topic to store events and invoke the Lambda function.

Answer: A

NEW QUESTION 149

- (Exam Topic 2)

A company has developed APIs that use Amazon API Gateway with Regional endpoints. The APIs call AWS Lambda functions that use API Gateway authentication mechanisms. After a design review, a solutions architect identifies a set of APIs that do not require public access. The solutions architect must design a solution to make the set of APIs accessible only from a VPC. All APIs need to be called with an authenticated user. Which solution will meet these requirements with the LEAST amount of effort?

- A. Create an internal Application Load Balancer (ALB). Create a target group.
- B. Select the Lambda function to call.
- C. Use the ALB DNS name to call the API from the VPC.
- D. Remove the DNS entry that is associated with the API in API Gateway.
- E. Create a hosted zone in Amazon Route 53. Create a CNAME record in the hosted zone.
- F. Update the API in API Gateway with the CNAME record.
- G. Use the CNAME record to call the API from the VPC.
- H. Update the API endpoint from Regional to private in API Gateway.
- I. Create an interface VPC endpoint in the VPC.
- J. Create a resource policy, and attach it to the API.
- K. Use the VPC endpoint to call the API from the VPC.
- L. Deploy the Lambda functions inside the VPC.
- M. Provision an EC2 instance, and install an Apache server. From the Apache server, call the Lambda function.
- N. Use the internal CNAME record of the EC2 instance to call the API from the VPC.

Answer: D

NEW QUESTION 151

- (Exam Topic 2)

A company has a serverless multi-tenant content management system on AWS. The architecture contains a web-based front end that interacts with an Amazon API Gateway API that uses a custom AWS Lambda authorizer. The authorizer authenticates a user to its tenant ID and encodes the information in a JSON Web Token (JWT) token. After authentication, each API call through API Gateway targets a Lambda function that interacts with a single Amazon DynamoDB table to fulfill requests.

To comply with security standards, the company needs a stronger isolation between tenants. The company will have hundreds of customers within the first year. Which solution will meet these requirements with the LEAST operational?

- A. Create a DynamoDB table for each tenant by using the tenant ID in the table name
- B. Create a service that uses the JWT token to retrieve the appropriate Lambda execution role that is tenant-specific
- C. Attach IAM policies to the execution role to allow access only to the DynamoDB table for the tenant.
- D. Add tenant ID information to the partition key of the DynamoDB table
- E. Create a service that uses the JWT token to retrieve the appropriate Lambda execution role that is tenant-specific
- F. Attach IAM policies to the execution role to allow access to items in the table only when the key matches the tenant ID.
- G. Create a separate AWS account for each tenant of the application
- H. Use dedicated infrastructure for each tenant
- I. Ensure that no cross-account network connectivity exists.
- J. Add tenant ID as a sort key in every DynamoDB table
- K. Add logic to each Lambda function to use the tenant ID that comes from the JWT token as the sort key in every operation on the DynamoDB table.

Answer: B

NEW QUESTION 154

- (Exam Topic 2)

A company is using AWS CloudFormation to deploy its infrastructure. The company is concerned that if a production CloudFormation stack is deleted, important data stored in Amazon RDS databases or Amazon EBS volumes might also be deleted.

How can the company prevent users from accidentally deleting data in this way?

- A. Modify the CloudFormation templates to add a DeletionPolicy attribute to RDS and EBS resources.
- B. Configure a stack policy that disallows the deletion of RDS and EBS resources.
- C. Modify IAM policies to deny deleting RDS and EBS resources that are tagged with an "aws:cloudformation:stack-name" tag.
- D. Use AWS Config rules to prevent deleting RDS and EBS resources.

Answer: A

NEW QUESTION 155

- (Exam Topic 2)

A financial company is planning to migrate its web application from on-premises to AWS. The company uses a third-party security tool to monitor the inbound traffic to the application. The company has used the security tool for the last 15 years, and the tool has no cloud solutions available from its vendor. The company's security team is concerned about how to integrate the security tool with AWS technology.

The company plans to deploy the application migration to AWS on Amazon EC2 instances. The EC2 instances will run in an Auto Scaling group in a dedicated VPC. The company needs to use the security tool to inspect all packets that come in and out of the VPC. This inspection must occur in real time and must not affect the application's performance. A solutions architect must design a target architecture on AWS that is highly available within an AWS Region.

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

- A. Deploy the security tool on EC2 instances in a new Auto Scaling group in the existing VPC.
- B. Deploy the web application behind a Network Load Balancer.
- C. Deploy an Application Load Balancer in front of the security tool instances.
- D. Provision a Gateway Load Balancer for each Availability Zone to redirect the traffic to the security tool.
- E. Provision a transit gateway to facilitate communication between VPCs.

Answer: AD

NEW QUESTION 156

- (Exam Topic 2)

A company manages hundreds of AWS accounts centrally in an organization in AWS Organizations. The company recently started to allow product teams to create and manage their own S3 access points in their accounts. The S3 access points can be accessed only within VPCs, not on the internet.

What is the MOST operationally efficient way to enforce this requirement?

- A. Set the S3 access point resource policy to deny the s3:CreateAccessPoint action unless the s3:AccessPointNetworkOrigin condition key evaluates to VPC.
- B. Create an SCP at the root level in the organization to deny the s3:CreateAccessPoint action unless the s3:AccessPointNetworkOrigin condition key evaluates to VPC.
- C. Use AWS CloudFormation StackSets to create a new IAM policy in each AWS account that allows the s3:CreateAccessPoint action only if the s3:AccessPointNetworkOrigin condition key evaluates to VPC.
- D. Set the S3 bucket policy to deny the s3:CreateAccessPoint action unless the s3:AccessPointNetworkOrigin condition key evaluates to VPC.

Answer: A

NEW QUESTION 158

- (Exam Topic 2)

A company is launching a web-based application in multiple regions around the world. The application consists of both static content stored in a private Amazon S3 bucket and dynamic ECS containers behind an Application Load Balancer (ALB). The company requires that the static and dynamic application content be accessible through Amazon CloudFront only.

Which combination of steps should a solutions architect recommend to restrict direct content access to CloudFront? (Select THREE)

- A. Create a web ACL in AWS WAF with a rule to validate the presence of a custom header and associate the web ACL with the ALB
- B. Create a web ACL in AWS WAF with a rule to validate the presence of a custom header and associate the web ACL with the CloudFront distribution
- C. Configure CloudFront to add a custom header to origin requests

- D. Configure the ALB to add a custom header to HTTP requests
- E. Update the S3 bucket ACL to allow access from the CloudFront distribution only
- F. Create a CloudFront Origin Access Identity (OAI) and add it to the CloudFront distribution Update the S3 bucket policy to allow access to the OAI only

Answer: ACF

NEW QUESTION 161

- (Exam Topic 2)

A company is planning to migrate 1,000 on-premises servers to AWS. The servers run on several VMware clusters in the company's data center. As part of the migration plan, the company wants to gather server metrics such as CPU details, RAM usage, operating system information, and running processes. The company then wants to query and analyze the data.

Which solution will meet these requirements?

- A. Deploy and configure the AWS Agentless Discovery Connector virtual appliance on the on-premises host
- B. Configure Data Exploration in AWS Migration Hub
- C. Use AWS Glue to perform an ETL job against the data
- D. Query the data by using Amazon S3 Select.
- E. Export only the VM performance information from the on-premises host
- F. Directly import the required data into AWS Migration Hub
- G. Update any missing information in Migration Hub
- H. Query the data by using Amazon QuickSight.
- I. Create a script to automatically gather the server information from the on-premises host
- J. Use the AWS CLI to run the put-resource-attributes command to store the detailed server data in AWS Migration Hub
- K. Query the data directly in the Migration Hub console.
- L. Deploy the AWS Application Discovery Agent to each on-premises server
- M. Configure Data Exploration in AWS Migration Hub
- N. Use Amazon Athena to run predefined queries against the data in Amazon S3.

Answer: C

NEW QUESTION 165

- (Exam Topic 2)

A retail company is running an application that stores invoice files in an Amazon S3 bucket and metadata about the files in an Amazon DynamoDB table. The application software runs in both us-east-1 and eu-west-1. The S3 bucket and DynamoDB table are in us-east-1. The company wants to protect itself from data corruption and loss of connectivity to either Region.

Which option meets these requirements?

- A. Create a DynamoDB global table to replicate data between us-east-1 and eu-west-1. Enable continuous backup on the DynamoDB table in us-east-1. Enable versioning on the S3 bucket
- B. Create an AWS Lambda function triggered by Amazon CloudWatch Events to make regular backups of the DynamoDB table. Set up S3 cross-region replication from us-east-1 to eu-west-1. Set up MFA delete on the S3 bucket in us-east-1.
- C. Create a DynamoDB global table to replicate data between us-east-1 and eu-west-1. Enable versioning on the S3 bucket. Implement strict ACLs on the S3 bucket
- D. Create a DynamoDB global table to replicate data between us-east-1 and eu-west-1. Enable continuous backup on the DynamoDB table in us-east-1. Set up S3 cross-region replication from us-east-1 to eu-west-1.

Answer: B

NEW QUESTION 167

- (Exam Topic 2)

An external audit of a company's serverless application reveals IAM policies that grant too many permissions. These policies are attached to the company's AWS Lambda execution roles. Hundreds of the company's Lambda functions have broad access permissions, such as full access to Amazon S3 buckets and Amazon DynamoDB tables. The company wants each function to have only the minimum permissions that the function needs to complete its task.

A solutions architect must determine which permissions each Lambda function needs.

What should the solutions architect do to meet this requirement with the LEAST amount of effort?

- A. Set up Amazon CodeGuru to profile the Lambda functions and search for AWS API calls
- B. Create an inventory of the required API calls and resources for each Lambda function
- C. Create new IAM access policies for each Lambda function
- D. Review the new policies to ensure that they meet the company's business requirements.
- E. Turn on AWS CloudTrail logging for the AWS account
- F. Use AWS Identity and Access Management Access Analyzer to generate IAM access policies based on the activity recorded in the CloudTrail log. Review the generated policies to ensure that they meet the company's business requirements.
- G. Turn on AWS CloudTrail logging for the AWS account
- H. Create a script to parse the CloudTrail log, search for AWS API calls by Lambda execution role, and create a summary report
- I. Review the report
- J. Create IAM access policies that provide more restrictive permissions for each Lambda function.
- K. Turn on AWS CloudTrail logging for the AWS account
- L. Export the CloudTrail logs to Amazon S3. Use Amazon EMR to process the CloudTrail logs in Amazon S3 and produce a report of API calls and resources used by each execution role
- M. Create a new IAM access policy for each role
- N. Export the generated roles to an S3 bucket
- O. Review the generated policies to ensure that they meet the company's business requirements.

Answer: B

Explanation:

IAM Access Analyzer helps you identify the resources in your organization and accounts, such as Amazon S3 buckets or IAM roles, shared with an external entity. This lets you identify unintended access to your resources and data, which is a security risk. IAM Access Analyzer identifies resources shared with external principals by using logic-based reasoning to analyze the resource-based policies in your AWS environment.

<https://docs.aws.amazon.com/IAM/latest/UserGuide/what-is-access-analyzer.html>

NEW QUESTION 169

- (Exam Topic 2)

A company runs many workloads on AWS and uses AWS Organizations to manage its accounts. The workloads are hosted on Amazon EC2, AWS Fargate, and AWS Lambda. Some of the workloads have unpredictable demand. Accounts record high usage in some months and low usage in other months. The company wants to optimize its compute costs over the next 3 years. A solutions architect obtains a 6-month average for each of the accounts across the organization to calculate usage.

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

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

Answer: A

NEW QUESTION 174

- (Exam Topic 2)

A company is migrating an on-premises content management system (CMS) to AWS Fargate. The company uses the CMS for blog posts that include text, images, and videos. The company has observed that traffic to blog posts drops by more than 80% after the posts are more than 30 days old. The CMS runs on multiple VMs and stores application state on disk. This application state is shared across all instances across multiple Availability Zones. Images and other media are stored on a separate NFS file share. The company needs to reduce the costs of the existing solution while minimizing the impact on performance.

Which combination of steps will meet these requirements MOST cost-effectively? (Select TWO.)

- A. Store media in an Amazon S3 Standard bucket. Create an S3 Lifecycle configuration that transitions objects that are older than 30 days to the S3 Standard-Infrequent Access (S3 Standard-IA) storage class.
- B. Store media on an Amazon Elastic File System (Amazon EFS) volume. Attach the EFS volume to all Fargate instances.
- C. Store application state on an Amazon Elastic File System (Amazon EFS) volume. Attach the EFS volume to all Fargate instances.
- D. Store application state on an Amazon Elastic Block Store (Amazon EBS) volume. Attach the EBS volume to all Fargate instances.
- E. Store media in an Amazon S3 Standard bucket. Create an S3 Lifecycle configuration that transitions objects that are older than 30 days to the S3 Glacier storage class.

Answer: AC

NEW QUESTION 177

- (Exam Topic 2)

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

What should the solutions architect do to meet these requirements?

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

Answer: C

Explanation:

<https://aws.amazon.com/blogs/aws/new-use-aws-cloudformation-stacksets-for-multiple-accounts-in-an-aws-orga>

NEW QUESTION 178

- (Exam Topic 2)

A solutions architect needs to review the design of an Amazon EMR cluster that is using the EMR File System (EMRFS). The cluster performs tasks that are critical to business needs. The cluster is running Amazon EC2 On-Demand Instances at all times for all task, master, and core nodes. The EMR tasks run each morning, starting at 1:00 AM, and take 6 hours to finish running. The amount of time to complete the processing is not a priority because the data is not referenced until late in the day.

The solutions architect must review the architecture and suggest a solution to minimize the compute costs. Which solution should the solutions architect recommend to meet these requirements?

- A. Launch all task, master, and core nodes on Spot Instances in an instance fleet.
- B. Terminate the cluster, including all instances, when the processing is completed.
- C. Launch the master and core nodes on On-Demand Instance.
- D. Launch the task nodes on Spot Instances in an instance fleet.
- E. Terminate the cluster, including all instances, when the processing is complete.
- F. Purchase Compute Savings Plans to cover the On-Demand Instance usage.
- G. Continue to launch all nodes on On-Demand Instance.
- H. Terminate the cluster.
- I. Including all instances, when the processing is complete.
- J. Purchase Compute Savings Plans to cover the On-Demand Instance usage.
- K. Launch the master and core nodes on On-Demand Instance.
- L. Launch the task nodes on Spot Instances in an instance fleet.
- M. Terminate only the task node Instances when the processing is completed. Purchase Compute Savings Plans to cover the On-Demand Instance usage.

Answer: B

NEW QUESTION 179

- (Exam Topic 2)

A company operates quick-service restaurants. The restaurants follow a predictable model with high sales traffic for 4 hours daily. Sales traffic is lower outside of those peak hours.

The point of sale and management platform is deployed in the AWS Cloud and has a backend that is based on Amazon DynamoDB. The database table uses provisioned throughput mode with 100,000 RCUs and 80,000 WCUs to match known peak resource consumption.

The company wants to reduce its DynamoDB cost and minimize the operational overhead for the IT staff. Which solution meets these requirements MOST cost-effectively?

- A. Reduce the provisioned RCUs and WCUs
- B. Change the DynamoDB table to use on-demand capacity
- C. Enable DynamoDB auto scaling for the table.
- D. Purchase 1-year reserved capacity that is sufficient to cover the peak load for 4 hours each day.

Answer: C

NEW QUESTION 184

- (Exam Topic 2)

A company wants to migrate its data analytics environment from on-premises to AWS. The environment consists of two simple Node.js applications. One of the applications collects sensor data and loads it into a MySQL database. The other application aggregates the data into reports. When the aggregation jobs run, some of the load jobs fail to run correctly.

The company must resolve the data loading issue. The company also needs the migration to occur without interruptions or changes for the company's customers. What should a solutions architect do to meet these requirements?

- A. Set up an Amazon Aurora MySQL database as a replication target for the on-premises database. Create an Aurora Replica for the Aurora MySQL database, and move the aggregation jobs to run against the Aurora Replica. Set up collection endpoints as AWS Lambda functions behind a Network Load Balancer (NLB), and use Amazon RDS Proxy to write to the Aurora MySQL database. When the databases are synced, disable the replication job and restart the Aurora Replica as the primary instance.
- B. Point the collector DNS record to the NLB.
- C. Set up an Amazon Aurora MySQL database. Use AWS Database Migration Service (AWS DMS) to perform continuous data replication from the on-premises database to Aurora. Move the aggregation jobs to run against the Aurora MySQL database. Set up collection endpoints behind an Application Load Balancer (ALB) as Amazon EC2 instances in an Auto Scaling group. When the databases are synced, point the collector DNS record to the ALB. Disable the AWS DMS sync task after the cutover from on-premises to AWS.
- D. Set up an Amazon Aurora MySQL database. Use AWS Database Migration Service (AWS DMS) to perform continuous data replication from the on-premises database to Aurora. Create an Aurora Replica for the Aurora MySQL database and move the aggregation jobs to run against the Aurora Replica. Set up collection endpoints as AWS Lambda functions behind an Application Load Balancer (ALB) and use Amazon RDS Proxy to write to the Aurora MySQL database. When the databases are synced, point the collector DNS record to the ALB. Disable the AWS DMS sync task after the cutover from on-premises to AWS.
- E. Set up an Amazon Aurora MySQL database. Create an Aurora Replica for the Aurora MySQL database and move the aggregation jobs to run against the Aurora Replica. Set up collection endpoints as an Amazon Kinesis data stream. Use Amazon Kinesis Data Firehose to replicate the data to the Aurora MySQL database. When the databases are synced, disable the replication job and restart the Aurora Replica as the primary instance. Point the collector DNS record to the Kinesis data stream.

Answer: C

Explanation:

Set up an Amazon Aurora MySQL database. Use AWS Database Migration Service (AWS DMS) to perform continuous data replication from the on-premises database to Aurora. Create an Aurora Replica for the Aurora MySQL database, and move the aggregation jobs to run against the Aurora Replica. Set up collection endpoints as AWS Lambda functions behind an Application Load Balancer (ALB), and use Amazon RDS Proxy to write to the Aurora MySQL database. When the databases are synced, point the collector DNS record to the ALB. Disable the AWS DMS sync task after the cutover from on-premises to AWS.

Amazon RDS Proxy allows applications to pool and share connections established with the database, improving database efficiency and application scalability. With RDS Proxy, failover times for Aurora and RDS databases are reduced by up to 66%.

NEW QUESTION 187

- (Exam Topic 2)

A company is running an application in the AWS Cloud. The application collects and stores a large amount of unstructured data in an Amazon S3 bucket. The S3 bucket contains several terabytes of data and uses the S3 Standard storage class. The data increases in size by several gigabytes every day.

The company needs to query and analyze the data. The company does not access data that is more than 1 year old. However, the company must retain all the data indefinitely for compliance reasons.

Which solution will meet these requirements MOST cost-effectively?

- A. Use S3 Select to query the data.
- B. Create an S3 Lifecycle policy to transition data that is more than 1 year old to S3 Glacier Deep Archive.
- C. Use Amazon Redshift Spectrum to query the data.
- D. Create an S3 Lifecycle policy to transition data that is more than 1 year old to S3 Glacier Deep Archive.
- E. Use an AWS Glue Data Catalog and Amazon Athena to query the data.
- F. Create an S3 Lifecycle policy to transition data that is more than 1 year old to S3 Glacier Deep Archive.
- G. Use Amazon Redshift Spectrum to query the data.
- H. Create an S3 Lifecycle policy to transition data that is more than 1 year old to S3 Intelligent-Tiering.

Answer: A

NEW QUESTION 188

- (Exam Topic 2)

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

What should a solution architect do to meet these requirements?

- A. Create an AWS Site-to-Site VPN connection. Configure integration between a VPN and AD DS.
- B. Use an Amazon Workspaces client with MFA support enabled to establish a VPN connection.

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

Answer: B

NEW QUESTION 190

- (Exam Topic 2)

A company has multiple business units Each business unit has its own AWS account and runs a single website within that account. The company also has a single logging account. Logs from each business unit website are aggregated into a single Amazon S3 bucket in the logging account. The S3 bucket policy provides each business unit with access to write data into the bucket and requires data to be encrypted.

The company needs to encrypt logs uploaded into the bucket using a Single AWS Key Management Service (AWS KMS) CMK The CMK that protects the data must be rotated once every 365 days

Which strategy is the MOST operationally efficient for the company to use to meet these requirements?

- A. Create a customer managed CMK ri the logging account Update the CMK key policy to provide access to the logging account only Manually rotate the CMK every 365 days.
- B. Create a customer managed CMK in the logging account
- C. Update the CMK key policy to provide access to the logging account and business unit account
- D. Enable automatic rotation of the CMK
- E. Use an AWS managed CMK m the toggging account
- F. Update the CMK key policy to provide access to the logging account and business unit accounts Manually rotate the CMK every 365 days.
- G. Use an AWS managed CMK in the toggging account Update the CMK key policy to provide access to the toggging account onl
- H. Enable automatic rotation of the CMK.

Answer: A

NEW QUESTION 192

- (Exam Topic 2)

A video streaming company recently launched a mobile app for video sharing. The app uploads various files to an Amazon S3 bucket in the us-east-1 Region. The files range in size from 1 GB to 1 0 GB.

Users who access the app from Australia have experienced uploads that take long periods of time Sometimes the files fail to completely upload for these users . A solutions architect must improve the app' performance for these uploads

Which solutions will meet these requirements? (Select TWO.)

- A. Enable S3 Transfer Acceleration on the S3 bucket Configure the app to use the Transfer Acceleration endpoint for uploads
- B. Configure an S3 bucket in each Region to receive the upload
- C. Use S3 Cross-Region Replication to copy the files to the distribution S3 bucket.
- D. Set up Amazon Route 53 with latency-based routing to route the uploads to the nearest S3 bucket Region.
- E. Configure the app to break the video files into chunks Use a multipart upload to transfer files to Amazon S3.
- F. Modify the app to add random prefixes to the files before uploading

Answer: AD

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/s3-upload-large-files/>

NEW QUESTION 194

- (Exam Topic 2)

A company has migrated its forms-processing application to AWS. When users interact with the application, they upload scanned forms as files through a web application. A database stores user metadata and references to files that are stored in Amazon S3. The web application runs on Amazon EC2 instances and an Amazon RDS for PostgreSQL database.

When forms are uploaded, the application sends notifications to a team through Amazon Simple Notification Service (Amazon SNS). A team member then logs in and processes each form. The team member performs data validation on the form and extracts relevant data before entering the information into another system that uses an API.

A solutions architect needs to automate the manual processing of the forms. The solution must provide accurate form extraction, minimize time to market, and minimize long-term operational overhead.

Which solution will meet these requirements?

- A. Develop custom libraries to perform optical character recognition (OCR) on the form
- B. Deploy the libraries to an Amazon Elastic Kubernetes Service (Amazon EKS) cluster as an application tie
- C. Use this tier to process the forms when forms are uploade
- D. Store the output in Amazon S3. Parse this output by extracting the data into an Amazon DynamoDB tabl
- E. Submit the data to the target system's AP
- F. Host the new application tier on EC2 instances.
- G. Extend the system with an application tier that uses AWS Step Functions and AWS Lambd
- H. Configure this tier to use artificial intelligence and machine learning (AI/ML) models that are trained and hosted on an EC2 instance to perform optical character recognition (OCR) on the forms when forms are uploade
- I. Store the output in Amazon S3. Parse this output by extracting the data that is required within the application tie
- J. Submit the data to the target system's API.
- K. Host a new application tier on EC2 instance
- L. Use this tier to call endpoints that host artificial intelligence and machine learning (AI/ML) models that are trained and hosted in Amazon SageMaker to perform optical character recognition (OCR) on the form
- M. Store the output in Amazon ElastiCach
- N. Parse this output by extracting the data that is required within the application tie
- O. Submit the data to the target system's API.
- P. Extend the system with an application tier that uses AWS Step Functions and AWS Lambd

- Q. Configure this tier to use Amazon Textract and Amazon Comprehend to perform optical character recognition (OCR) on the forms when forms are uploaded
- R. Store the output in Amazon S3. Parse this output by extracting the data that is required within the application tier
- S. Submit the data to the target system's API.

Answer: D

NEW QUESTION 198

.....

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

AWS-Certified-Solutions-Architect-Professional Practice Exam Features:

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

100% Actual & Verified — Instant Download, Please Click
[Order The AWS-Certified-Solutions-Architect-Professional Practice Test Here](#)