



Amazon

Exam Questions AWS-Certified-Solutions-Architect-Professional

Amazon AWS Certified Solutions Architect Professional

About ExamBible

Your Partner of IT Exam

Found in 1998

ExamBible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, ExamBible has its unique advantages that other companies could not achieve.

Our Advances

* 99.9% Uptime

All examinations will be up to date.

* 24/7 Quality Support

We will provide service round the clock.

* 100% Pass Rate

Our guarantee that you will pass the exam.

* Unique Gurantee

If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

NEW QUESTION 1

- (Exam Topic 1)

A company has an application that sells tickets online and experiences bursts of demand every 7 days. The application has a stateless presentation layer running on Amazon EC2, an Oracle database to store unstructured data catalog information, and a backend API layer. The front-end layer uses an Elastic Load Balancer to distribute the load across nine On-Demand Instances over three Availability Zones (AZs). The Oracle database is running on a single EC2 instance. The company is experiencing performance issues when running more than two concurrent campaigns. A solutions architect must design a solution that meets the following requirements:

- Address scalability issues.
- Increase the level of concurrency.
- Eliminate licensing costs.
- Improve reliability.

Which set of steps should the solutions architect take?

- A. Create an Auto Scaling group for the front end with a combination of On-Demand and Spot Instances to reduce cost
- B. Convert the Oracle database into a single Amazon RDS reserved DB instance.
- C. Create an Auto Scaling group for the front end with a combination of On-Demand and Spot Instances to reduce cost
- D. Create two additional copies of the database instance, then distribute the databases in separate AZs.
- E. Create an Auto Scaling group for the front end with a combination of On-Demand and Spot Instances to reduce cost
- F. Convert the tables in the Oracle database into Amazon DynamoDB tables.
- G. Convert the On-Demand Instances into Spot Instances to reduce costs for the front end
- H. Convert the tables in the Oracle database into Amazon DynamoDB tables.

Answer: C

Explanation:

Combination of On-Demand and Spot Instances + DynamoDB.

NEW QUESTION 2

- (Exam Topic 1)

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

Which solution will meet the company's requirements?

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

Answer: C

Explanation:

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

NEW QUESTION 3

- (Exam Topic 1)

A company has a project that is launching Amazon EC2 instances that are larger than required. The project's account cannot be part of the company's organization in AWS Organizations due to policy restrictions to keep this activity outside of corporate IT. The company wants to allow only the launch of t3.small EC2 instances by developers in the project's account. These EC2 instances must be restricted to the us-east-2 Region.

What should a solutions architect do to meet these requirements?

- A. Create a new developer account
- B. Move all EC2 instances, users, and assets into us-east-2. Add the account to the company's organization in AWS Organization
- C. Enforce a tagging policy that denotes Region affinity.
- D. Create an SCP that denies the launch of all EC2 instances except t3.small EC2 instances in us-east-2. Attach the SCP to the project's account.
- E. Create and purchase a t3.small EC2 Reserved Instance for each developer in us-east-2. Assign each developer a specific EC2 instance with their name as the tag.
- F. Create an IAM policy that allows the launch of only t3.small EC2 instances in us-east-2. Attach the policy to the roles and groups that the developers use in the project's account.

Answer: D

NEW QUESTION 4

- (Exam Topic 1)

A team collects and routes behavioral data for an entire company. The company runs a Multi-AZ VPC environment with public subnets, private subnets, and an internet gateway. Each public subnet also contains a NAT gateway. Most of the company's applications read from and write to Amazon Kinesis Data Streams. Most of the workloads are in private subnets.

A solutions architect must review the infrastructure. The solutions architect needs to reduce costs and maintain the function of the applications. The solutions architect uses Cost Explorer and notices that the cost in the EC2-Other category is consistently high. A further review shows that NatGateway-Bytes charges are increasing the cost in the EC2-Other category.

What should the solutions architect do to meet these requirements?

- A. Enable VPC Flow Log
- B. Use Amazon Athena to analyze the logs for traffic that can be removed
- C. Ensure that security groups are blocking traffic that is responsible for high costs.
- D. Add an interface VPC endpoint for Kinesis Data Streams to the VPC
- E. Ensure that applications have the correct IAM permissions to use the interface VPC endpoint.
- F. Enable VPC Flow Logs and Amazon Detective Review Detective findings for traffic that is not related to Kinesis Data Streams Configure security groups to block that traffic
- G. Add an interface VPC endpoint for Kinesis Data Streams to the VPC
- H. Ensure that the VPC endpoint policy allows traffic from the applications.

Answer: D

Explanation:

<https://docs.aws.amazon.com/vpc/latest/privatelink/vpc-endpoints-access.html> <https://aws.amazon.com/premiumsupport/knowledge-center/vpc-reduce-nat-gateway-transfer-costs/>

VPC endpoint policies enable you to control access by either attaching a policy to a VPC endpoint or by using additional fields in a policy that is attached to an IAM user, group, or role to restrict access to only occur via the specified VPC endpoint

NEW QUESTION 5

- (Exam Topic 1)

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

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

Which solution will satisfy these requirements with minimal operational overhead?

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

Answer: A

Explanation:

SMS can handle migrating the data volumes:

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

NEW QUESTION 6

- (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 7

- (Exam Topic 1)

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

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

- A. Create an EC2 Auto Scaling group behind an Application Load Balance
- B. Create additional read replicas for the DB instance
- C. Create Amazon Kinesis data streams and configure the application services to use the data stream

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

Answer: D

Explanation:

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

NEW QUESTION 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 many services running in its on-premises data center. The data center is connected to AWS using AWS Direct Connect (DX) and an IPsec VPN. The service data is sensitive and connectivity cannot traverse the internet. The company wants to expand into a new market segment and begin offering its services to other companies that are using AWS.

Which solution will meet these requirements?

- A. Create a VPC Endpoint Service that accepts TCP traffic, host it behind a Network Load Balancer, and make the service available over DX.
- B. Create a VPC Endpoint Service that accepts HTTP or HTTPS traffic, host it behind an Application Load Balancer, and make the service available over DX.
- C. Attach an internet gateway to the VPC
- D. and ensure that network access control and security group rules allow the relevant inbound and outbound traffic.
- E. Attach a NAT gateway to the VPC
- F. and ensure that network access control and security group rules allow the relevant inbound and outbound traffic.

Answer: A

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 13

- (Exam Topic 1)

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

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

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

Which solution meets these requirements?

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

Answer: C

NEW QUESTION 15

- (Exam Topic 1)

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

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

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

Answer: B

Explanation:

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

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

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

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

NEW QUESTION 16

- (Exam Topic 1)

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

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

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

Answer: BD

Explanation:

https://aws.amazon.com/documentdb/?nc1=h_ls

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

NEW QUESTION 21

- (Exam Topic 1)

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

What is the MOST cost-effective migration recommendation?

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

Answer: D

Explanation:

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

NEW QUESTION 23

- (Exam Topic 1)

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

Which solutions will meet these requirements?

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

J. Once complete, have the ETL job send the results to another S3 bucket for internal processing.

K. Create an AWS Glue crawler and custom classifier based upon the data feed formats and build a table definition to match

L. Perform an Amazon Athena query on file delivery to start an Amazon EMR ETL job to transform the entire record according to the processing and transformation requirement

M. Define the output format as JSON

N. Once complete, send the results to another S3 bucket for internal processing and scale down the EMR cluster.

Answer: C

Explanation:

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

Create an AWS Glue crawler and custom classifier based on the data feed formats and build a table definition to match. Trigger an AWS Lambda function on file delivery to start an AWS Glue ETL job to transform the entire record according to the processing and transformation requirements. Define the output format as JSON.

Once complete, have the ETL job send the results to another S3 bucket for internal processing. <https://docs.aws.amazon.com/glue/latest/dg/trigger-job.html>
https://d1.awsstatic.com/Products/product-name/diagrams/product-page-diagram_Glue_Event-driven-ETL-Pipel

NEW QUESTION 28

- (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 30

- (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 34

- (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 36

- (Exam Topic 1)

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

Which approach should the company take to achieve this goal?

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

Answer: A

NEW QUESTION 40

- (Exam Topic 1)

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

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

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

Answer: C

Explanation:

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

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

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

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

NEW QUESTION 44

- (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 49

- (Exam Topic 1)

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

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

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

Answer: B

Explanation:

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

NEW QUESTION 54

- (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 55

- (Exam Topic 1)

A company wants to move a web application to AWS. The application stores session information locally on each web server, which will make auto scaling difficult. As part of the migration, the application will be rewritten to decouple the session data from the web servers. The company requires low latency, scalability, and availability.

Which service will meet the requirements for storing the session information in the MOST cost-effective way?

- A. Amazon ElastiCache with the Memcached engine
- B. Amazon S3
- C. Amazon RDS MySQL
- D. Amazon ElastiCache with the Redis engine

Answer: D

Explanation:

<https://aws.amazon.com/caching/session-management/>

Building real-time apps across versatile use cases like gaming, geospatial service, caching, session stores, or queuing, with advanced data structures, replication, and point-in-time snapshot support. Memcached: Building a simple, scalable caching layer for your data-intensive apps. <https://aws.amazon.com/elasticache/>

NEW QUESTION 59

- (Exam Topic 1)

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

To meet these requirements, a solutions architect plans to use an S3 access point that is restricted to specific VPCs for each application.

Which combination of steps should the solutions architect take to implement this solution? (Select TWO.)

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

Answer: AC

Explanation:

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

NEW QUESTION 64

- (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 service
- B. Create an Application Load Balancer (ALB). Configure the CloudFront distribution to forward requests to the ALB
- C. Configure the ALB to invoke the correct Lambda function for each type of request
- 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 service
- F. Configure API Gateway to invoke the correct Lambda function for each type of request
- 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 service
- I. Configure API Gateway to invoke the correct Lambda function for each type of request
- J. Create a response mapping template to remove the problematic headers based on the value of the User-Agent header
- K. Associate the response data mapping with the HTTP API.
- L. Create an Amazon CloudFront distribution for the metadata service
- M. Create an Application Load Balancer (ALB). Configure the CloudFront distribution to forward requests to the ALB
- N. Configure the ALB to invoke the correct Lambda function for each type of request
- O. Create a Lambda@Edge function that will remove the problematic headers in response to viewer requests based on the value of the User-Agent header.

Answer: D

Explanation:

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

NEW QUESTION 68

- (Exam Topic 1)

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

How should a solutions architect improve the CI/CD pipeline to reduce the likelihood that changes in the templates will cause downtime?

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

Answer: B

Explanation:

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

NEW QUESTION 73

- (Exam Topic 1)

A company stores sales transaction data in Amazon DynamoDB tables. To detect anomalous behaviors and respond quickly, all changes to the items stored in the DynamoDB tables must be logged within 30 minutes.

Which solution meets the requirements?

- A. Copy the DynamoDB tables into Apache Hive tables on Amazon EMR every hour and analyze them (or anomalous behavior)
- B. Send Amazon SNS notifications when anomalous behaviors are detected.
- C. Use AWS CloudTrail to capture all the APIs that change the DynamoDB table
- D. Send SNS notifications when anomalous behaviors are detected using CloudTrail event filtering.
- E. Use Amazon DynamoDB Streams to capture and send updates to AWS Lambda
- F. Create a Lambda function to output records to Amazon Kinesis Data Stream
- G. Analyze any anomalies with Amazon Kinesis Data Analytics
- H. Send SNS notifications when anomalous behaviors are detected.
- I. Use event patterns in Amazon CloudWatch Events to capture DynamoDB API call events with an AWS Lambda function as a target to analyze behavior
- J. Send SNS notifications when anomalous behaviors are detected.

Answer: C

Explanation:

[https://aws.amazon.com/blogs/database/dynamodb-streams-use-cases-and-design-patterns/#:~:text=DynamoDB DynamoDb Stream to capture DynamoDB update. And Kinesis Data Analytics for anomaly detection \(it uses AWS proprietary Random Cut Forest Algorithm\)](https://aws.amazon.com/blogs/database/dynamodb-streams-use-cases-and-design-patterns/#:~:text=DynamoDB%20Stream%20to%20capture%20DynamoDB%20update,And%20Kinesis%20Data%20Analytics%20for%20anomaly%20detection%20(it%20uses%20AWS%20proprietary%20Random%20Cut%20Forest%20Algorithm))

NEW QUESTION 74

- (Exam Topic 1)

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

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

Answer: A

NEW QUESTION 79

- (Exam Topic 1)

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

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

Which solution will meet these requirements?

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

Answer: B

NEW QUESTION 80

- (Exam Topic 1)

An AWS customer has a web application that runs on premises. The web application fetches data from a third-party API that is behind a firewall. The third party accepts only one public CIDR block in each client's allow list.

The customer wants to migrate their web application to the AWS Cloud. The application will be hosted on a set of Amazon EC2 instances behind an Application Load Balancer (ALB) in a VPC. The ALB is located in public subnets. The EC2 instances are located in private subnets. NAT gateways provide internet access to the private subnets.

How should a solutions architect ensure that the web application can continue to call the third-party API after the migration?

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

Answer: B

Explanation:

When EC2 instances reach third-party API through internet, their private IP addresses will be masked by NAT Gateway public IP address.

<https://aws.amazon.com/blogs/networking-and-content-delivery/introducing-bring-your-own-ip-byoip-for-amaz>

NEW QUESTION 82

- (Exam Topic 1)

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

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

Which of the following solutions is MOST cost-effective?

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

Answer: B

NEW QUESTION 83

- (Exam Topic 1)

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

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

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

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

Answer: B

Explanation:

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

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

NEW QUESTION 87

- (Exam Topic 1)

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

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

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

Answer: B

Explanation:

<https://aws.amazon.com/aws-transfer-family/faqs/> <https://docs.aws.amazon.com/transfer/latest/userguide/what-is-aws-transfer-family.html>

https://aws.amazon.com/about-aws/whats-new/2018/11/aws-transfer-for-sftp-fully-managed-sftp-for-s3/?nc1=h_

NEW QUESTION 89

- (Exam Topic 1)

A start up company hosts a fleet of Amazon EC2 instances in private subnets using the latest Amazon Linux 2 AMI. The company's engineers rely heavily on SSH access to the instances for troubleshooting.

The company's existing architecture includes the following:

- A VPC with private and public subnets, and a NAT gateway
- Site-to-Site VPN for connectivity with the on-premises environment
- EC2 security groups with direct SSH access from the on-premises environment

The company needs to increase security controls around SSH access and provide auditing of commands executed by the engineers.

Which strategy should a solutions architect use?

- A. Install and configure EC2 Instance Connect on the fleet of EC2 instance
- B. Remove all security group rules attached to EC2 instances that allow inbound TCP on port 22. Advise the engineers to remotely access the instances by using the EC2 Instance Connect CLI.
- C. Update the EC2 security groups to only allow inbound TCP on port 22 to the IP addresses of the engineer's device
- D. Install the Amazon CloudWatch agent on all EC2 instances and send operating system audit logs to CloudWatch Logs.
- E. Update the EC2 security groups to only allow inbound TCP on port 22 to the IP addresses of the engineer's device
- F. Enable AWS Config for EC2 security group resource change
- G. Enable AWS Firewall Manager and apply a security group policy that automatically remediates changes to rules.
- H. Create an IAM role with the AmazonSSMManagedInstanceCore managed policy attached
- I. Attach the IAM role to all the EC2 instance
- J. Remove all security group rules attached to the EC2 instances that allow inbound TCP on port 22. Have the engineers install the AWS Systems Manager Session Manager plugin for their devices and remotely access the instances by using the start-session API call from Systems Manager.

Answer: D

Explanation:

Allows client machines to be able to connect to Session Manager using the AWS CLI instead of going through the AWS EC2 or AWS Server Manager console.

<https://docs.aws.amazon.com/systems-manager/latest/userguide/session-manager-working-with-install-plugin.html>

NEW QUESTION 92

- (Exam Topic 1)

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

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

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

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

Answer: ABE

NEW QUESTION 94

- (Exam Topic 1)

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

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

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

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

Answer: B

Explanation:

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

NEW QUESTION 99

- (Exam Topic 1)

A company is developing and hosting several projects in the AWS Cloud. The projects are developed across multiple AWS accounts under the same organization

in AWS Organizations. The company requires the cost for cloud infrastructure to be allocated to the owning project. The team responsible for all of the AWS accounts has discovered that several Amazon EC2 instances are lacking the Project tag used for cost allocation. Which actions should a solutions architect take to resolve the problem and prevent it from happening in the future? (Select THREE.)

- A. Create an AWS Config rule in each account to find resources with missing tags.
- B. Create an SCP in the organization with a deny action for ec2:RunInstances if the Project tag is missing.
- C. Use Amazon Inspector in the organization to find resources with missing tags.
- D. Create an IAM policy in each account with a deny action for ec2:RunInstances if the Project tag is missing.
- E. Create an AWS Config aggregator for the organization to collect a list of EC2 instances with the missing Project tag.
- F. Use AWS Security Hub to aggregate a list of EC2 instances with the missing Project tag.

Answer: CDE

NEW QUESTION 104

- (Exam Topic 1)

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

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

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

Answer: B

Explanation:

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

NEW QUESTION 108

- (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 112

- (Exam Topic 1)

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

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

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

Answer: C

Explanation:

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

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

NEW QUESTION 115

- (Exam Topic 1)

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

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

Which data migration strategy should the company use?

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

Answer: B

Explanation:

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

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

NEW QUESTION 118

- (Exam Topic 1)

A company wants to use a third-party software-as-a-service (SaaS) application. The third-party SaaS application is consumed through several API calls. The third-party SaaS application also runs on AWS inside a VPC.

The company will consume the third-party SaaS application from inside a VPC. The company has internal security policies that mandate the use of private connectivity that does not traverse the internet. No resources that run in the company VPC are allowed to be accessed from outside the company's VPC. All permissions must conform to the principles of least privilege.

Which solution meets these requirements?

- A. Create an AWS PrivateLink interface VPC endpoint
- B. Connect this endpoint to the endpoint service that the third-party SaaS application provide
- C. Create a security group to limit the access to the endpoint
- D. Associate the security group with the endpoint.
- E. Create an AWS Site-to-Site VPN connection between the third-party SaaS application and the company VP
- F. Configure network ACLs to limit access across the VPN tunnels.
- G. Create a VPC peering connection between the third-party SaaS application and the company VPC route tables by adding the needed routes for the peering connection.
- H. Create an AWS PrivateLink endpoint service
- I. Ask the third-party SaaS provider to create an interface VPC endpoint for this endpoint service
- J. Grant permissions for the endpoint service to the specific account of the third-party SaaS provider.

Answer: A

Explanation:

Reference architecture - <https://docs.aws.amazon.com/vpc/latest/privatelink/privatelink-access-saas.html> Note from documentation that Interface Endpoint is at client side

NEW QUESTION 123

- (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 125

- (Exam Topic 1)

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

Which solution is the MOST cost-effective way to meet these requirements?

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

Answer: B

Explanation:

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

NEW QUESTION 126

- (Exam Topic 1)

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

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

Which of the following solutions should the solutions architect recommend?

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

Answer: B

Explanation:

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

NEW QUESTION 130

- (Exam Topic 1)

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

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

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

Answer: C

NEW QUESTION 135

- (Exam Topic 1)

A company needs to architect a hybrid DNS solution. This solution will use an Amazon Route 53 private hosted zone for the domain cloud.example.com for the resources stored within VPCs.

The company has the following DNS resolution requirements:

- On-premises systems should be able to resolve and connect to cloud.example.com.
- All VPCs should be able to resolve cloud.example.com.

There is already an AWS Direct Connect connection between the on-premises corporate network and AWS Transit Gateway. Which architecture should the company use to meet these requirements with the HIGHEST performance?

- A. Associate the private hosted zone to all the VPC
- B. Create a Route 53 inbound resolver in the shared services VPC
- C. Attach all VPCs to the transit gateway and create forwarding rules in the on-premises DNS server for cloud.example.com that point to the inbound resolver.
- D. Associate the private hosted zone to all the VPC
- E. Deploy an Amazon EC2 conditional forwarder in the shared services VPC
- F. Attach all VPCs to the transit gateway and create forwarding rules in the on-premises DNS server for cloud.example.com that point to the conditional forwarder.
- G. Associate the private hosted zone to the shared services VPC
- H. Create a Route 53 outbound resolver in the shared services VPC
- I. Attach all VPCs to the transit gateway and create forwarding rules in the on-premises DNS server for cloud.example.com that point to the outbound resolver.
- J. Associate the private hosted zone to the shared services VPC
- K. Create a Route 53 inbound resolver in the shared services VPC

L. Attach the shared services VPC to the transit gateway and create forwarding rules in the on-premises DNS server for cloud.example.com that point to the inbound resolver.

Answer: D

Explanation:

<https://aws.amazon.com/blogs/networking-and-content-delivery/centralized-dns-management-of-hybrid-cloud-w>

NEW QUESTION 138

- (Exam Topic 1)

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

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

Which set of actions should the team take?

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

Answer: D

Explanation:

Multi AZ ASG + ALB + Aurora = Less over head and automatic scaling

NEW QUESTION 140

- (Exam Topic 1)

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

- A. Create a dashboard by using AWS Systems Manager OpsCenter. Configure visualizations for Amazon CloudWatch metrics that are associated with the EC2 instances and their EBS volumes. Review the dashboard periodically and identify usage patterns. 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 EC2 instances based on the peaks in the metrics.
- C. Install the Amazon CloudWatch agent on each of the EC2 instances. Turn on AWS Compute Optimizer, and let it run for at least 12 hours. Review the recommendations from Compute Optimizer, and 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 141

- (Exam Topic 2)

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

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

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

Answer: D

NEW QUESTION 143

- (Exam Topic 2)

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

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

- A. Deploy the AWS Server Migration Service (AWS SMS) connector using the OVA image on the VMware cluster to collect configuration and utilization data from

the VMs

- B. Use the AWS Migration Portfolio Assessment (MPA) tool to connect to each of the VMs to collect the configuration and utilization data.
- C. Install the AWS Application Discovery Service on each of the VMs to collect the configuration and utilization data
- D. Register the on-premises VMs with the AWS Migration Hub to collect configuration and utilization data

Answer: A

NEW QUESTION 148

- (Exam Topic 2)

A company wants to improve cost awareness for its Amazon EMR platform. The company has allocated budgets for each team's Amazon EMR usage. When a budgetary threshold is reached, a notification should be sent by email to the budget office's distribution list. Teams should be able to view their EMR cluster expenses to date. A solutions architect needs to create a solution that ensures this policy is proactively and centrally enforced in a multi-account environment. Which combination of steps should the solutions architect take to meet these requirements? (Select TWO.)

- A. Update the AWS CloudFormation template to include the AWS Budgets Budget resource with the NotificationsWithSubscribers property
- B. Implement Amazon CloudWatch dashboards for Amazon EMR usage
- C. Create an EMR bootstrap action that runs at startup that calls the Cost Explorer API to set the budget on the cluster with the GetCostForecast and NotificationsWithSubscribers actions
- D. Create an AWS Service Catalog portfolio for each team
- E. Add each team's Amazon EMR cluster as an AWS CloudFormation template to their Service Catalog portfolio as a Product
- F. Create an Amazon CloudWatch metric for billing. Create a custom alert when costs exceed the budgetary threshold.

Answer: BE

NEW QUESTION 149

- (Exam Topic 2)

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

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

Which solution will meet these requirements?

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

Answer: A

NEW QUESTION 152

- (Exam Topic 2)

A company has an organization in AWS Organizations that has a large number of AWS accounts. One of the AWS accounts is designated as a transit account and has a transit gateway that is shared with all of the other AWS accounts. AWS Site-to-Site VPN connections are configured between all of the company's global offices and the transit account. The company has AWS Config enabled on all of its accounts.

The company's networking team needs to centrally manage a list of internal IP address ranges that belong to the global offices. Developers will reference this list to gain access to applications securely.

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

- A. Create a JSON file that is hosted in Amazon S3 and that lists all of the internal IP address ranges. Configure an Amazon Simple Notification Service (Amazon SNS) topic in each of the accounts that can be involved when the JSON file is updated.
- B. Subscribe an AWS Lambda function to the SNS topic to update all relevant security group rules with the updated IP address ranges.
- C. Create a new AWS Config managed rule that contains all of the internal IP address ranges. Use the rule to check the security groups in each of the accounts to ensure compliance with the list of IP address ranges.
- D. Configure the rule to automatically remediate any noncompliant security group that is detected.
- E. In the transit account, create a VPC prefix list with all of the internal IP address ranges.
- F. Use AWS Resource Access Manager to share the prefix list with all of the other accounts.
- G. Use the shared prefix list to configure security group rules in the other accounts.
- H. In the transit account, create a security group with all of the internal IP address ranges.
- I. Configure the security groups in the other accounts to reference the transit account's security group by using a nested security group reference of `*-<transit-account-id>/.sg-1a2b3c4d`.

Answer: C

NEW QUESTION 157

- (Exam Topic 2)

A video processing company has an application that downloads images from an Amazon S3 bucket, processes the images, stores a transformed image in a second S3 bucket, and updates metadata about the image in an Amazon DynamoDB table. The application is written in Node.js and runs by using an AWS Lambda function. The Lambda function is invoked when a new image is uploaded to Amazon S3.

The application ran without incident for a while. However, the size of the images has grown significantly. The Lambda function is now failing frequently with timeout errors. The function timeout is set to its maximum value. A solutions architect needs to refactor the application's architecture to prevent invocation failures. The company does not want to manage the underlying infrastructure.

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

- A. Modify the application deployment by building a Docker image that contains the application code. Publish the image to Amazon Elastic Container Registry (Amazon ECR).
- B. Create a new Amazon Elastic Container Service (Amazon ECS) task definition with a compatibility type of AWS Fargat
- C. Configure the task definition to use the new image in Amazon Elastic Container Registry (Amazon ECR). Adjust the Lambda function to invoke an ECS task by using the ECS task definition when a new file arrives in Amazon S3.
- D. Create an AWS Step Functions state machine with a Parallel state to invoke the Lambda function. Increase the provisioned concurrency of the Lambda function.
- E. Create a new Amazon Elastic Container Service (Amazon ECS) task definition with a compatibility type of Amazon EC2. Configure the task definition to use the new image in Amazon Elastic Container Registry (Amazon ECR). Adjust the Lambda function to invoke an ECS task by using the ECS task definition when a new file arrives in Amazon S3.
- F. Modify the application to store images on Amazon Elastic File System (Amazon EFS) and to store metadata on an Amazon RDS DB instance
- G. Adjust the Lambda function to mount the EFS file share.

Answer: DE

NEW QUESTION 161

- (Exam Topic 2)

A company is hosting a critical application on a single Amazon EC2 instance. The application uses an Amazon ElastiCache for Redis single-node cluster for an in-memory data store. The application uses an Amazon RDS for MariaDB DB instance for a relational database. For the application to function, each piece of the infrastructure must be healthy and must be in an active state.

A solutions architect needs to improve the application's architecture so that the infrastructure can automatically recover from failure with the least possible downtime.

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

- A. Use an Elastic Load Balancer to distribute traffic across multiple EC2 instances
- B. Ensure that the EC2 instances are part of an Auto Scaling group that has a minimum capacity of two instances.
- C. Use an Elastic Load Balancer to distribute traffic across multiple EC2 instances. Ensure that the EC2 instances are configured in unlimited mode.
- D. Modify the DB instance to create a read replica in the same Availability Zone
- E. Promote the read replica to be the primary DB instance in failure scenarios.
- F. Modify the DB instance to create a Multi-AZ deployment that extends across two Availability Zones.
- G. Create a replication group for the ElastiCache for Redis cluster
- H. Configure the cluster to use an Auto Scaling group that has a minimum capacity of two instances.
- I. Create a replication group for the ElastiCache for Redis cluster
- J. Enable Multi-AZ on the cluster.

Answer: ADE

NEW QUESTION 164

- (Exam Topic 2)

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

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

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

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

Answer: E

NEW QUESTION 167

- (Exam Topic 2)

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

GET /posts/{postId} to get post details
GET /users/{userId} to get user details

GET /comments/{commentId} to get comments details

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

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

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

Answer: C

NEW QUESTION 172

- (Exam Topic 2)

A company with several AWS accounts is using AWS Organizations and service control policies (SCPs). An Administrator created the following SCP and has attached it to an organizational unit (OU) that contains AWS account 1111-1111-1111:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowsAllActions",
      "Effect": "Allow",
      "Action": "*",
      "Resource": "*"
    },
    {
      "Sid": "DenyCloudTrail",
      "Effect": "Deny",
      "Action": "cloudtrail:*",
      "Resource": "*"
    }
  ]
}
```

Developers working in account 1111-1111-1111 complain that they cannot create Amazon S3 buckets. How should the Administrator address this problem?

- A. Add s3:CreateBucket with Allow effect to the SCP.
- B. Remove the account from the OU, and attach the SCP directly to account 1111-1111-1111.
- C. Instruct the Developers to add Amazon S3 permissions to their IAM entities.
- D. Remove the SCP from account 1111-1111-1111.

Answer: C

Explanation:

is incorrect - https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps.html

"SCPs are similar to AWS Identity and Access Management (IAM) permission policies and use almost the same syntax. However, an SCP never grants permissions."

SCPs alone are not sufficient to granting permissions to the accounts in your organization. No permissions are granted by an SCP. An SCP defines a guardrail, or sets limits, on the actions that the account's administrator can delegate to the IAM users and roles in the affected accounts. The administrator must still attach identity-based or resource-based policies to IAM users or roles, or to the resources in your accounts to actually grant permissions. The effective permissions are the logical intersection between what is allowed by the SCP and what is allowed by the IAM and resource-based policies.

NEW QUESTION 175

- (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 176

- (Exam Topic 2)

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

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

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

- A. Create a new RDS for PostgreSQL DB instance in the target account Use the AWS Schema Conversion Tool (AWS SCT) to migrate the database schema from the source database to the target database.
- B. Use the AWS Schema Conversion Tool (AWS SCT) to create a new RDS for PostgreSQL DB instance in the target account with the schema and initial data from the source database
- C. Configure VPC peering between the VPCs in the two AWS accounts to provide connectivity to both DB instances from the target account
- D. Configure the security groups that are attached to each DB instance to allow traffic on the database port from the VPC in the target account

- E. Temporarily allow the source DB instance to be publicly accessible to provide connectivity from the VPC in the target account Configure the security groups that are attached to each DB instance to allow traffic on the database port from the VPC in the target account.
- F. Use AWS Database Migration Service (AWS DMS) in the target account to perform a full load plus change data capture (CDC) migration from the source database to the target database When the migration is complete, change the CNAME record to point to the target DB instance endpoint
- G. Use AWS Database Migration Service (AWS DMS) in the target account to perform a change data capture (CDC) migration from the source database to the target database When the migration is complete change the CNAME record to point to the target DB instance endpoint

Answer: BCE

NEW QUESTION 181

- (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 184

- (Exam Topic 2)

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

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

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

Answer: D

NEW QUESTION 188

- (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 effort?

- 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 192

- (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 194

- (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 198

- (Exam Topic 2)

A company has an asynchronous HTTP application that is hosted as an AWS Lambda function. A public Amazon API Gateway endpoint invokes the Lambda function. The Lambda function and the API Gateway endpoint reside in the us-east-1 Region. A solutions architect needs to redesign the application to support failover to another AWS Region.

Which solution will meet these requirements?

- A. Create an API Gateway endpoint in the us-west-2 Region to direct traffic to the Lambda function in us-east-1. Configure Amazon Route 53 to use a failover routing policy to route traffic for the two API Gateway endpoints.
- B. Create an Amazon Simple Queue Service (Amazon SQS) queue.
- C. Configure API Gateway to direct traffic to the SQS queue instead of to the Lambda function.
- D. Configure the Lambda function to pull messages from the queue for processing.
- E. Deploy the Lambda function to the us-west-2 Region.
- F. Create an API Gateway endpoint in us-west-2 to direct traffic to the Lambda function in us-west-2. Configure AWS Global Accelerator and an Application Load Balancer to manage traffic across the two API Gateway endpoints.
- G. Deploy the Lambda function and an API Gateway endpoint to the us-west-2 Region.
- H. Configure Amazon Route 53 to use a failover routing policy to route traffic for the two API Gateway endpoints.

Answer: B

NEW QUESTION 203

- (Exam Topic 2)

A company runs an IoT application in the AWS Cloud. The company has millions of sensors that collect data from houses in the United States. The sensors use the MQTT protocol to connect and send data to a custom MQTT broker. The MQTT broker stores the data on a single Amazon EC2 instance. The sensors connect to the broker through the domain named `iot.example.com`. The company uses Amazon Route 53 as its DNS service. The company stores the data in Amazon DynamoDB.

On several occasions, the amount of data has overloaded the MQTT broker and has resulted in lost sensor data. The company must improve the reliability of the solution.

Which solution will meet these requirements?

- A. Create an Application Load Balancer (ALB) and an Auto Scaling group for the MQTT broker.
- B. Use the Auto Scaling group as the target for the ALB.
- C. Update the DNS record in Route 53 to an alias record.
- D. Point the alias record to the ALB.
- E. Use the MQTT broker to store the data.
- F. Set up AWS IoT Core to receive the sensor data.
- G. Create and configure a custom domain to connect to AWS IoT Core.
- H. Update the DNS record in Route 53 to point to the AWS IoT Core Data-ATS endpoint.
- I. Configure an AWS IoT rule to store the data.
- J. Create a Network Load Balancer (NLB). Set the MQTT broker as the target.
- K. Create an AWS Global Accelerator endpoint group.
- L. Set the NLB as the endpoint for the accelerator.

- M. Update the DNS record in Route 53 to a multivalued answer record
- N. Set the Global Accelerator IP addresses as value
- O. Use the MQTT broker to store the data.
- P. Set up AWS IoT Greengrass to receive the sensor data
- Q. Update the DNS record in Route 53 to point to the AWS IoT Greengrass endpoint
- R. Configure an AWS IoT rule to invoke an AWS Lambda function to store the data.

Answer: C

NEW QUESTION 208

- (Exam Topic 2)

A solutions architect needs to provide AWS Cost and Usage Report data from a company's AWS Organizations management account. The company already has an Amazon S3 bucket to store the reports. The reports must be automatically ingested into a database that can be visualized with other tools. Which combination of steps should the solutions architect take to meet these requirements? (Select THREE.)

- A. Create an Amazon EventBridge (Amazon CloudWatch Events) rule that a new object creation in the S3 bucket will trigger
- B. Create an AWS Cost and Usage Report configuration to deliver the data into the S3 bucket
- C. Configure an AWS Glue crawler that a new object creation in the S3 bucket will trigger.
- D. Create an AWS Lambda function that a new object creation in the S3 bucket will trigger
- E. Create an AWS Glue crawler that the AWS Lambda function will trigger to crawl objects in the S3 bucket
- F. Create an AWS Glue crawler that the Amazon EventBridge (Amazon CloudWatch Events) rule will trigger to crawl objects in the S3 bucket

Answer: BDF

NEW QUESTION 210

- (Exam Topic 2)

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

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

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

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

Answer: A

NEW QUESTION 214

- (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. Terminate 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 217

- (Exam Topic 2)

A company has automated the nightly retraining of its machine learning models by using AWS Step Functions. The workflow consists of multiple steps that use AWS Lambda. Each step can fail for various reasons, and any failure causes a failure of the overall workflow.

A review reveals that the retraining has failed multiple nights in a row without the company noticing the failure. A solutions architect needs to improve the workflow so that notifications are sent for all types of failures in the retraining process.

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

- A. Create an Amazon Simple Notification Service (Amazon SNS) topic with a subscription of type "Email" that targets the team's mailing list.
- B. Create a task named "Email" that forwards the input arguments to the SNS topic

- C. Add a Catch field to all Tas
- D. Ma
- E. and Parallel states that have a statement of "ErrorEquals": ["states.all"] and "Next": "Email".
- F. Add a new email address to Amazon Simple Email Service (Amazon SES). Verify the email address.
- G. Create a task named "Email" that forwards the input arguments to the SES email address
- H. Add a Catch field to all Task, Map, and Parallel states that have a statement of "ErrorEquals": ["states.Bun time"] and "Next": "Email".

Answer: BCD

NEW QUESTION 218

.....

Relate Links

100% Pass Your AWS-Certified-Solutions-Architect-Professional Exam with ExamBible Prep Materials

<https://www.exambible.com/AWS-Certified-Solutions-Architect-Professional-exam/>

Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - <https://www.exambible.com/>