

Amazon

Exam Questions AWS-Certified-Solutions-Architect-Professional

Amazon AWS Certified Solutions Architect Professional



NEW QUESTION 1

- (Exam Topic 2)

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

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

- A. Enable Aurora Auto Scaling for Aurora Replica
- B. Use a Network Load Balancer with the least outstanding requests routing algorithm and sticky sessions enabled.
- C. Enable Aurora Auto Scaling for Aurora writer
- D. Use an Application Load Balancer with the round robin routing algorithm and sticky sessions enabled.
- E. 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 2

- (Exam Topic 2)

A company built an application based on AWS Lambda deployed in an AWS CloudFormation stack. The last production release of the web application introduced an issue that resulted in an outage lasting several minutes. A solutions architect must adjust the deployment process to support a canary release.

Which solution will meet these requirements?

- A. Create an alias for every new deployed version of the Lambda function
- B. Use the AWS CLI update-alias command with the routing-config parameter to distribute the load.
- C. Deploy the application into a new CloudFormation stack
- D. Use an Amazon Route 53 weighted routing policy to distribute the load.
- E. Create a version for every new deployed Lambda function
- F. Use the AWS CLI update-function-configuration command with the routing-config parameter to distribute the load.
- G. Configure AWS CodeDeploy and use CodeDeployDefault.OneAtATime in the Deployment configuration to distribute the load.

Answer: A

Explanation:

<https://aws.amazon.com/blogs/compute/implementing-canary-deployments-of-aws-lambda-functions-with-alias>

NEW QUESTION 3

- (Exam Topic 2)

A solutions architect is redesigning a three-tier application that a company hosts on premises. The application provides personalized recommendations based on user profiles. The company already has an AWS account and has configured a VPC to host the application.

The frontend is a Java-based application that runs in on-premises VMs. The company hosts a personalization model on a physical application server and uses TensorFlow to implement the model. The personalization model uses artificial intelligence and machine learning (AI/ML). The company stores user information in a Microsoft SQL Server database. The web application calls the personalization model, which reads the user profiles from the database and provides recommendations.

The company wants to migrate the redesigned application to AWS.

Which solution will meet this requirement with the LEAST operational overhead?

- A. Use AWS Server Migration Service (AWS SMS) to migrate the on-premises physical application server and the web application VMs to AWS
- B. Use AWS Database Migration Service (AWS DMS) to migrate the SQL Server database to Amazon RDS for SQL Server.
- C. Export the personalization model
- D. Store the model artifacts in Amazon S3. Deploy the model to Amazon SageMaker and create an endpoint
- E. Host the Java application in AWS Elastic Beanstalk
- F. Use AWS Database Migration Service (AWS DMS) to migrate the SQL Server database to Amazon RDS for SQL Server.
- G. Use AWS Application Migration Service to migrate the on-premises personalization model and VMs to Amazon EC2 instances in an Auto Scaling group
- H. Use AWS Database Migration Service (AWS DMS) to migrate the SQL Server database to an EC2 instance.
- I. Containerize the personalization model and the Java application
- J. Use Amazon Elastic Kubernetes Service (Amazon EKS) managed node groups to deploy the model and the application to Amazon EKS. Host the node groups in a VPC
- K. Use AWS Database Migration Service (AWS DMS) to migrate the SQL Server database to Amazon RDS for SQL Server.

Answer: B

Explanation:

Amazon SageMaker is a fully managed machine learning service that allows users to build, train, and deploy machine learning models quickly and easily¹. Users can export their existing TensorFlow models and store the model artifacts in Amazon S3, a highly scalable and durable object storage service². Users can then deploy the model to Amazon SageMaker and create an endpoint that can be invoked by the web application to provide recommendations³. This way, the solution can leverage the AI/ML capabilities of Amazon SageMaker without having to rewrite the personalization model.

AWS Elastic Beanstalk is a service that allows users to deploy and manage web applications without worrying about the infrastructure that runs those applications. Users can host their Java application in AWS Elastic Beanstalk and configure it to communicate with the Amazon SageMaker endpoint. This way, the solution can reduce the operational overhead of managing servers, load balancers, scaling, and application health monitoring.

AWS Database Migration Service (AWS DMS) is a service that helps users migrate databases to AWS quickly and securely. Users can use AWS DMS to migrate their SQL Server database to Amazon RDS for SQL Server, a fully managed relational database service that offers high availability, scalability, security, and compatibility. This way, the solution can reduce the operational overhead of managing database servers, backups, patches, and upgrades.

Option A is incorrect because using AWS Server Migration Service (AWS SMS) to migrate the on-premises physical application server and the web application VMs to AWS is not cost-effective or scalable. AWS SMS is a service that helps users migrate on-premises workloads to AWS. However, for this use case, migrating the physical application server and the web application VMs to AWS will not take advantage of the AI/ML capabilities of Amazon SageMaker or the managed services of AWS Elastic Beanstalk and Amazon RDS.

Option C is incorrect because using AWS Application Migration Service to migrate the on-premises personalization model and VMs to Amazon EC2 instances in Auto Scaling groups is not cost-effective or scalable. AWS Application Migration Service is a service that helps users migrate applications from on-premises or other clouds to AWS without making any changes to their applications. However, for this use case, migrating the personalization model and VMs to EC2 instances will not take advantage of the AI/ML capabilities of Amazon SageMaker or the managed services of AWS Elastic Beanstalk and Amazon RDS.

Option D is incorrect because containerizing the personalization model and the Java application and using Amazon Elastic Kubernetes Service (Amazon EKS) managed node groups to deploy them to Amazon EKS is not necessary or cost-effective. Amazon EKS is a service that allows users to run Kubernetes on AWS without needing to install, operate, and maintain their own Kubernetes control plane or nodes. However, for this use case, containerizing and deploying the personalization model and the Java application will not take advantage of the AI/ML capabilities of Amazon SageMaker or the managed services of AWS Elastic Beanstalk. Moreover, using S3 Glacier Deep Archive as a storage class for images will incur a high retrieval fee and latency for accessing them.

NEW QUESTION 4

- (Exam Topic 2)

A company wants to refactor its retail ordering web application that currently has a load-balanced Amazon EC2 instance fleet for web hosting, database API services, and business logic. The company needs to create a decoupled, scalable architecture with a mechanism for retaining failed orders while also minimizing operational costs.

Which solution will meet these requirements?

- A. Use Amazon S3 for web hosting with Amazon API Gateway for database API service
- B. Use Amazon Simple Queue Service (Amazon SQS) for order queuein
- C. Use Amazon Elastic Container Service (Amazon ECS) for business logic with Amazon SQS long polling for retaining failed orders.
- D. Use AWS Elastic Beanstalk for web hosting with Amazon API Gateway for database API service
- E. Use Amazon MQ for order queuein
- F. Use AWS Step Functionsfor business logic with Amazon S3 Glacier Deep Archive for retaining failed orders.
- G. Use Amazon S3 for web hosting with AWS AppSync for database API service
- H. Use Amazon Simple Queue Service (Amazon SQS) for order queuein
- I. Use AWS Lambda for business logic with an Amazon SQS dead-letter queue for retaining failed orders.
- J. Use Amazon Lightsail for web hosting with AWS AppSync for database API service
- K. Use Amazon Simple Email Service (Amazon SES) for order queuein
- L. UseAmazon Elastic Kubernetes Service (Amazon EKS) for business logic with Amazon OpenSearch Service for retaining failed orders.

Answer: C

Explanation:

•Use Amazon S3 for web hosting with AWS AppSync for database API services. Use Amazon Simple Queue Service (Amazon SQS) for order queuing. Use AWS Lambda for business logic with an Amazon SQS dead-letter queue for retaining failed orders.

This solution will allow you to:

- Host a static website on Amazon S3 without provisioning or managing servers¹.
- Use AWS AppSync to create a scalable GraphQL API that connects to your database and other data sources¹.
- Use Amazon SQS to decouple and scale your order processing microservices¹.
- Use AWS Lambda to run code for your business logic without provisioning or managing servers¹.
- Use an Amazon SQS dead-letter queue to retain messages that can't be processed by your Lambda function¹.

NEW QUESTION 5

- (Exam Topic 2)

A company has an application that runs on Amazon EC2 instances in an Amazon EC2 Auto Scaling group. The company uses AWS CodePipeline to deploy the application. The instances that run in the Auto Scaling group are constantly changing because of scaling events.

When the company deploys new application code versions, the company installs the AWS CodeDeploy agent on any new target EC2 instances and associates the instances with the CodeDeploy deployment group. The application is set to go live within the next 24 hours.

What should a solutions architect recommend to automate the application deployment process with the LEAST amount of operational overhead?

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

Answer: D

Explanation:

<https://docs.aws.amazon.com/codedeploy/latest/userguide/integrations-aws-auto-scaling.html>

NEW QUESTION 6

- (Exam Topic 2)

A company has an application in the AWS Cloud. The application runs on a fleet of 20 Amazon EC2 instances. The EC2 instances are persistent and store data on multiple attached Amazon Elastic Block Store (Amazon EBS) volumes.

The company must maintain backups in a separate AWS Region. The company must be able to recover the EC2 instances and their configuration within 1 business day, with loss of no more than 1 day's worth of data. The company has limited staff and needs a backup solution that optimizes operational efficiency and cost. The company already has created an AWS CloudFormation template that can deploy the required network configuration in a secondary Region.

Which solution will meet these requirements?

- A. Create a second CloudFormation template that can recreate the EC2 instances in the secondary Region. Run daily multivolume snapshots by using AWS Systems Manager Automation runbook
- B. Copy the snapshots to the secondary Region
- C. In the event of a failure, launch the CloudFormation templates, restore the EBS volumes from snapshots, and transfer usage to the secondary Region.
- D. Use Amazon Data Lifecycle Manager (Amazon DLM) to create daily multivolume snapshots of the EBS volume
- E. In the event of a failure, launch the CloudFormation template and use Amazon DLM to restore the EBS volumes and transfer usage to the secondary Region.
- F. Use AWS Backup to create a scheduled daily backup plan for the EC2 instance
- G. Configure the backup task to copy the backups to a vault in the secondary Region
- H. In the event of a failure, launch the CloudFormation template, restore the instance volumes and configurations from the backup vault, and transfer usage to the secondary Region.
- I. Deploy EC2 instances of the same size and configuration to the secondary Region
- J. Configure AWS DataSync daily to copy data from the primary Region to the secondary Region
- K. In the event of a failure, launch the CloudFormation template and transfer usage to the secondary Region.

Answer: C

Explanation:

Using AWS Backup to create a scheduled daily backup plan for the EC2 instances will enable taking snapshots of the EC2 instances and their attached EBS volumes¹. Configuring the backup task to copy the backups to a vault in the secondary Region will enable maintaining backups in a separate Region¹. In the event of a failure, launching the CloudFormation template will enable deploying the network configuration in the secondary Region². Restoring the instance volumes and configurations from the backup vault will enable recovering the EC2 instances and their data¹. Transferring usage to the secondary Region will enable resuming operations².

NEW QUESTION 7

- (Exam Topic 2)

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

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

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

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

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

Answer: ABD

Explanation:

Cross region with AWS Backup:

<https://docs.aws.amazon.com/aws-backup/latest/devguide/cross-region-backup.html>

NEW QUESTION 8

- (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 accelerator
- 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 values
- 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: A

Explanation:

It describes a solution that uses an Application Load Balancer (ALB) and an Auto Scaling group for the MQTT broker. The ALB distributes incoming traffic across

the instances in the Auto Scaling group and allows for automatic scaling based on incoming traffic. The use of an alias record in Route 53 allows for easy updates to the DNS record without changing the IP address. This solution improves the reliability of the MQTT broker by allowing it to automatically scale based on incoming traffic, reducing the likelihood of lost data due to broker overload.

Reference: <https://aws.amazon.com/elasticloadbalancing/applicationloadbalancer/> <https://aws.amazon.com/autoscaling/> <https://aws.amazon.com/route53/>

NEW QUESTION 9

- (Exam Topic 2)

A company uses AWS Organizations to manage more than 1,000 AWS accounts. The company has created a new developer organization. There are 540 developer member accounts that must be moved to the new developer organization. All accounts are set up with all the required information so that each account can be operated as a standalone account.

Which combination of steps should a solutions architect take to move all of the developer accounts to the new developer organization? (Select THREE.)

- A. Call the MoveAccount operation in the Organizations API from the old organization's management account to migrate the developer accounts to the new developer organization.
- B. From the management account, remove each developer account from the old organization using the RemoveAccountFromOrganization operation in the Organizations API.
- C. From each developer account, remove the account from the old organization using the RemoveAccountFromOrganization operation in the Organizations API.
- D. Sign in to the new developer organization's management account and create a placeholder member account that acts as a target for the developer account migration.
- E. Call the InviteAccountToOrganization operation in the Organizations API from the new developer organization's management account to send invitations to the developer accounts.
- F. Have each developer sign in to their account and confirm to join the new developer organization.

Answer: BEF

Explanation:

"This operation can be called only from the organization's management account. Member accounts can remove themselves with LeaveOrganization instead."
https://docs.aws.amazon.com/organizations/latest/APIReference/API_RemoveAccountFromOrganization.html

NEW QUESTION 10

- (Exam Topic 2)

A company is building a call center by using Amazon Connect. The company's operations team is defining a disaster recovery (DR) strategy across AWS Regions. The contact center has dozens of contact flows, hundreds of users, and dozens of claimed phone numbers.

Which solution will provide DR with the LOWEST RTO?

- A. Create an AWS Lambda function to check the availability of the Amazon Connect instance and to send a notification to the operations team in case of unavailability.
- B. Create an Amazon EventBridge rule to invoke the Lambda function every 5 minutes.
- C. After notification, instruct the operations team to use the AWS Management Console to provision a new Amazon Connect instance in a second Region.
- D. Deploy the contact flows, users, and claimed phone numbers by using an AWS CloudFormation template.
- E. Provision a new Amazon Connect instance with all existing users in a second Region.
- F. Create an AWS Lambda function to check the availability of the Amazon Connect instance.
- G. Create an Amazon EventBridge rule to invoke the Lambda function every 5 minutes.
- H. In the event of an issue, configure the Lambda function to deploy an AWS CloudFormation template that provisions contact flows and claimed numbers in the second Region.
- I. Provision a new Amazon Connect instance with all existing contact flows and claimed phone numbers in a second Region.
- J. Create an Amazon Route 53 health check for the URL of the Amazon Connect instance.
- K. Create an Amazon CloudWatch alarm for failed health check.
- L. Create an AWS Lambda function to deploy an AWS CloudFormation template that provisions all users.
- M. Configure the alarm to invoke the Lambda function.
- N. Provision a new Amazon Connect instance with all existing users and contact flows in a second Region. Create an Amazon Route 53 health check for the URL of the Amazon Connect instance.
- O. Create an Amazon CloudWatch alarm for failed health check.
- P. Create an AWS Lambda function to deploy an AWS CloudFormation template that provisions claimed phone numbers.
- Q. Configure the alarm to invoke the Lambda function.

Answer: D

Explanation:

Option D provisions a new Amazon Connect instance with all existing users and contact flows in a second Region. It also sets up an Amazon Route 53 health check for the URL of the Amazon Connect instance, an Amazon CloudWatch alarm for failed health checks, and an AWS Lambda function to deploy an AWS CloudFormation template that provisions claimed phone numbers. This option allows for the fastest recovery time because all the necessary components are already provisioned and ready to go in the second Region. In the event of a disaster, the failed health check will trigger the AWS Lambda function to deploy the CloudFormation template to provision the claimed phone numbers, which is the only missing component.

NEW QUESTION 10

- (Exam Topic 2)

A company has developed a hybrid solution between its data center and AWS. The company uses Amazon VPC and Amazon EC2 instances that send application logs to Amazon CloudWatch. The EC2 instances read data from multiple relational databases that are hosted on premises.

The company wants to monitor which EC2 instances are connected to the databases in near-real time. The company already has a monitoring solution that uses Splunk on premises. A solutions architect needs to determine how to send networking traffic to Splunk.

How should the solutions architect meet these requirements?

- A. Enable VPC flow logs, and send them to CloudWatch.
- B. Create an AWS Lambda function to periodically export the CloudWatch logs to an Amazon S3 bucket by using the pre-defined export function.
- C. Generate ACCESS_KEY and SECRET_KEY AWS credentials.
- D. Configure Splunk to pull the logs from the S3 bucket by using those credentials.
- E. Create an Amazon Kinesis Data Firehose delivery stream with Splunk as the destination.
- F. Configure a pre-processing AWS Lambda function with a Kinesis Data Firehose stream processor that extracts individual log events from records sent by CloudWatch Logs subscription filter.

- G. Enable VPC flows logs, and send them to CloudWatc
- H. Create a CloudWatch Logs subscription that sends log events to the Kinesis Data Firehose delivery stream.
- I. Ask the company to log every request that is made to the databases along with the EC2 instance IP address
- J. Export the CloudWatch logs to an Amazon S3 bucket
- K. Use Amazon Athena to query the logs grouped by database name
- L. Export Athena results to another S3 bucket
- M. Invoke an AWS Lambda function to automatically send any new file that is put in the S3 bucket to Splunk.
- N. Send the CloudWatch logs to an Amazon Kinesis data stream with Amazon Kinesis Data Analytics for SQL Application
- O. Configure a 1 -minute sliding window to collect the event
- P. Create a SQL query that uses the anomaly detection template to monitor any networking traffic anomalies in near-real time
- Q. Send the result to an Amazon Kinesis Data Firehose delivery stream with Splunk as the destination.

Answer: B

Explanation:

<https://docs.aws.amazon.com/firehose/latest/dev/creating-the-stream-to-splunk.html>

NEW QUESTION 15

- (Exam Topic 2)

A company is migrating a document processing workload to AWS. The company has updated many applications to natively use the Amazon S3 API to store, retrieve, and modify documents that a processing server generates at a rate of approximately 5 documents every second. After the document processing is finished, customers can download the documents directly from Amazon S3.

During the migration, the company discovered that it could not immediately update the processing server that generates many documents to support the S3 API. The server runs on Linux and requires fast local access to the files that the server generates and modifies. When the server finishes processing, the files must be available to the public for download within 30 minutes.

Which solution will meet these requirements with the LEAST amount of effort?

- A. Migrate the application to an AWS Lambda function
- B. Use the AWS SDK for Java to generate, modify, and access the files that the company stores directly in Amazon S3.
- C. Set up an Amazon S3 File Gateway and configure a file share that is linked to the document store. Mount the file share on an Amazon EC2 instance by using NFS
- D. When changes occur in Amazon S3, initiate a RefreshCache API call to update the S3 File Gateway.
- E. Configure Amazon FSx for Lustre with an import and export policy
- F. Link the new file system to an S3 bucket
- G. Install the Lustre client and mount the document store to an Amazon EC2 instance by using NFS.
- H. Configure AWS DataSync to connect to an Amazon EC2 instance
- I. Configure a task to synchronize the generated files to and from Amazon S3.

Answer: C

Explanation:

The company should configure Amazon FSx for Lustre with an import and export policy. The company should link the new file system to an S3 bucket. The company should install the Lustre client and mount the document store to an Amazon EC2 instance by using NFS. This solution will meet the requirements with the least amount of effort because Amazon FSx for Lustre is a fully managed service that provides a high-performance file system optimized for fast processing of workloads such as machine learning, high performance computing, video processing, financial modeling, and electronic design automation¹. Amazon FSx for Lustre can be linked to an S3 bucket and can import data from and export data to the bucket². The import and export policy can be configured to automatically import new or changed objects from S3 and export new or changed files to S3. This will ensure that the files are available to the public for download within 30 minutes. Amazon FSx for Lustre supports NFS version 3.0 protocol for Linux clients.

The other options are not correct because:

- Migrating the application to an AWS Lambda function would require a lot of effort and may not be feasible for the existing server that generates many documents. Lambda functions have limitations on execution time, memory, disk space, and network bandwidth.
- Setting up an Amazon S3 File Gateway would not work because S3 File Gateway does not support write-back caching, which means that files written to the file share are uploaded to S3 immediately and are not available locally until they are downloaded again. This would not provide fast local access to the files that the server generates and modifies.
- Configuring AWS DataSync to connect to an Amazon EC2 instance would not meet the requirement of making the files available to the public for download within 30 minutes. DataSync is a service that transfers data between on-premises storage systems and AWS storage services over the internet or AWS Direct Connect. DataSync tasks can be scheduled to run at specific times or intervals, but they are not triggered by file changes.

References:

- <https://aws.amazon.com/fsx/lustre/>
- <https://docs.aws.amazon.com/fsx/latest/LustreGuide/create-fs-linked-data-repo.html>
- <https://docs.aws.amazon.com/fsx/latest/LustreGuide/import-export-data-repositories.html>
- <https://docs.aws.amazon.com/fsx/latest/LustreGuide/mounting-on-premises.html>
- <https://docs.aws.amazon.com/lambda/latest/dg/gettingstarted-limits.html>
- <https://docs.aws.amazon.com/storagegateway/latest/userguide/StorageGatewayConcepts.html>
- <https://docs.aws.amazon.com/datasync/latest/userguide/what-is-datasync.html>

NEW QUESTION 16

- (Exam Topic 2)

A company's interactive web application uses an Amazon CloudFront distribution to serve images from an Amazon S3 bucket. Occasionally, third-party tools ingest corrupted images into the S3 bucket. This image corruption causes a poor user experience in the application later. The company has successfully implemented and tested Python logic to detect corrupt images.

A solutions architect must recommend a solution to integrate the detection logic with minimal latency between the ingestion and serving.

Which solution will meet these requirements?

- A. Use a Lambda@Edge function that is invoked by a viewer-response event.
- B. Use a Lambda@Edge function that is invoked by an origin-response event.
- C. Use an S3 event notification that invokes an AWS Lambda function.

D. Use an S3 event notification that invokes an AWS Step Functions state machine.

Answer: B

Explanation:

This solution will allow the detection logic to be run as soon as the image is uploaded to the S3 bucket, before it is served to users via the CloudFront distribution. This way, the detection logic can quickly identify any corrupted images and prevent them from being served to users, minimizing latency between ingestion and serving.

Reference: AWS Lambda@Edge documentation:

<https://docs.aws.amazon.com/lambda/latest/dg/lambda-edge.html> You can use Lambda@Edge to run your code in response to CloudFront events, such as a viewer request, an origin request, a response, or an error.

NEW QUESTION 18

- (Exam Topic 2)

A company uses AWS Organizations with a single OU named Production to manage multiple accounts All accounts are members of the Production OU

Administrators use deny list SCPs in the root of the organization to manage access to restricted services.

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

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

A. Remove the organization's root SCPs that limit access to AWS Config Create AWS Service Catalog products for the company's standard AWS Config rules and deploy them throughout the organization, including the new account.

B. Create a temporary OU named Onboarding for the new account Apply an SCP to the Onboarding OU to allow AWS Config actions Move the new account to the Production OU when adjustments to AWS Config are complete

C. Convert the organization's root SCPs from deny list SCPs to allow list SCPs to allow the required services only Temporarily apply an SCP to the organization's root that allows AWS Config actions for principals only in the new account.

D. Create a temporary OU named Onboarding for the new account Apply an SCP to the Onboarding OU to allow AWS Config action

E. Move the organization's root SCP to the Production O

F. Move the new account to the Production OU when adjustments to AWS Config are complete.

Answer: D

Explanation:

An SCP at a lower level can't add a permission after it is blocked by an SCP at a higher level. SCPs can only filter; they never add permissions. SO you need to create a new OU for the new account assign an SCP, and move the root SCP to Production OU. Then move the new account to production OU when AWS config is done.

NEW QUESTION 22

- (Exam Topic 2)

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

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

Which solution will meet these requirements MOST cost-effectively?

A. Use S3 Select to query the dat

B. Create an S3 Lifecycle policy to transition data that is more than 1 year old to S3 Glacier Deep Archive.

C. Use Amazon Redshift Spectrum to query the dat

D. Create an S3 Lifecycle policy to transition data that is more than 1 year old to S3 Glacier Deep Archive.

E. Use an AWS Glue Data Catalog and Amazon Athena to query the dat

F. Create an S3 Lifecycle policy to transition data that is more than 1 year old to S3 Glacier Deep Archive.

G. Use Amazon Redshift Spectrum to query the dat

H. Create an S3 Lifecycle policy to transition data that is more than 1 year old to S3 Intelligent-Tiering.

Answer: C

Explanation:

Generally, unstructured data should be converted structured data before querying them. AWS Glue can do that.

<https://docs.aws.amazon.com/glue/latest/dg/schema-relationalize.html> <https://docs.aws.amazon.com/athena/latest/ug/glue-athena.html>

NEW QUESTION 23

- (Exam Topic 2)

A solutions architect at a large company needs to set up network security for outbound traffic to the internet from all AWS accounts within an organization in AWS Organizations. The organization has more than 100 AWS accounts, and the accounts route to each other by using a centralized AWS Transit Gateway. Each account has both an internet gateway and a NAT gateway for outbound traffic to the internet The company deploys resources only into a single AWS Region.

The company needs the ability to add centrally managed rule-based filtering on all outbound traffic to the internet for all AWS accounts in the organization. The peak load of outbound traffic will not exceed 25 Gbps in each Availability Zone.

Which solution meets these requirements?

A. Create a new VPC for outbound traffic to the interne

B. Connect the existing transit gateway to the new VP

C. Configure a new NAT gatewa

D. Create an Auto Scaling group of Amazon EC2 instances that run an open-source internet proxy for rule-based filtering across all Availability Zones in the Regio

E. Modify all default routes to point to the proxy's Auto Scaling group.

F. Create a new VPC for outbound traffic to the interne

G. Connect the existing transit gateway to the new VP

H. Configure a new NAT gatewa

I. Use an AWSNetwork Firewall firewall for rule-based filterin

J. Create Network Firewall endpoints in each Availability Zon

K. Modify all default routes to point to the Network Firewall endpoints.

- L. Create an AWS Network Firewall firewall for rule-based filtering in each AWS account
- M. Modify all default routes to point to the Network Firewall firewalls in each account.
- N. In each AWS account, create an Auto Scaling group of network-optimized Amazon EC2 instances that run an open-source internet proxy for rule-based filtering
- O. Modify all default routes to point to the proxy's Auto Scaling group.

Answer: B

Explanation:

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

NEW QUESTION 28

- (Exam Topic 2)

A company is implementing a serverless architecture by using AWS Lambda functions that need to access a Microsoft SQL Server DB instance on Amazon RDS. The company has separate environments for development and production, including a clone of the database system.

The company's developers are allowed to access the credentials for the development database. However, the credentials for the production database must be encrypted with a key that only members of the IT security team's IAM user group can access. This key must be rotated on a regular basis.

What should a solutions architect do in the production environment to meet these requirements?

- A. Store the database credentials in AWS Systems Manager Parameter Store by using a SecureString parameter that is encrypted by an AWS Key Management Service (AWS KMS) customer managed key
- B. Attach a role to each Lambda function to provide access to the SecureString parameter
- C. Restrict access to the SecureString parameter and the customer managed key so that only the IT security team can access the parameter and the key.
- D. Encrypt the database credentials by using the AWS Key Management Service (AWS KMS) default Lambda key
- E. Store the credentials in the environment variables of each Lambda function
- F. Load the credentials from the environment variables in the Lambda code
- G. Restrict access to the KMS key so that only the IT security team can access the key.
- H. Store the database credentials in the environment variables of each Lambda function
- I. Encrypt the environment variables by using an AWS Key Management Service (AWS KMS) customer managed key
- J. Restrict access to the customer managed key so that only the IT security team can access the key.
- K. Store the database credentials in AWS Secrets Manager as a secret that is associated with an AWS Key Management Service (AWS KMS) customer managed key
- L. Attach a role to each Lambda function to provide access to the secret
- M. Restrict access to the secret and the customer managed key so that only the IT security team can access the secret and the key.

Answer: D

Explanation:

Storing the database credentials in AWS Secrets Manager as a secret that is associated with an AWS Key Management Service (AWS KMS) customer managed key will enable encrypting and managing the credentials securely. AWS Secrets Manager helps you to securely encrypt, store, and retrieve credentials for your databases and other services. Attaching a role to each Lambda function to provide access to the secret will enable retrieving the credentials programmatically. Restricting access to the secret and the customer managed key so that only members of the IT security team's IAM user group can access them will enable meeting the security requirements.

NEW QUESTION 32

- (Exam Topic 2)

A company has a few AWS accounts for development and wants to move its production application to AWS. The company needs to enforce Amazon Elastic Block Store (Amazon EBS) encryption at rest current production accounts and future production accounts only. The company needs a solution that includes built-in blueprints and guardrails.

Which combination of steps will meet these requirements? (Choose three.)

- A. Use AWS CloudFormation StackSets to deploy AWS Config rules on production accounts.
- B. Create a new AWS Control Tower landing zone in an existing developer account
- C. Create OUs for account
- D. Add production and development accounts to production and development OUs, respectively.
- E. Create a new AWS Control Tower landing zone in the company's management account
- F. Add production and development accounts to production and development OU
- G. respectively.
- H. Invite existing accounts to join the organization in AWS Organization
- I. Create SCPs to ensure compliance.
- J. Create a guardrail from the management account to detect EBS encryption.
- K. Create a guardrail for the production OU to detect EBS encryption.

Answer: CDF

Explanation:

<https://docs.aws.amazon.com/controltower/latest/userguide/controls.html> <https://docs.aws.amazon.com/controltower/latest/userguide/strongly-recommended-controls.html#ebs-enable-en> AWS is now transitioning the previous term 'guardrail' new term 'control'.

NEW QUESTION 35

- (Exam Topic 2)

A software-as-a-service (SaaS) provider exposes APIs through an Application Load Balancer (ALB). The ALB connects to an Amazon Elastic Kubernetes Service (Amazon EKS) cluster that is deployed in the

us-east-1 Region. The exposed APIs contain usage of a few non-standard REST methods: LINK, UNLINK, LOCK, and UNLOCK.

Users outside the United States are reporting long and inconsistent response times for these APIs. A solutions architect needs to resolve this problem with a solution that minimizes operational overhead.

Which solution meets these requirements?

- A. Add an Amazon CloudFront distribution
- B. Configure the ALB as the origin.
- C. Add an Amazon API Gateway edge-optimized API endpoint to expose the API

- D. Configure the ALB as the target.
- E. Add an accelerator in AWS Global Accelerator.
- F. Configure the ALB as the origin.
- G. Deploy the APIs to two additional AWS Regions: eu-west-1 and ap-southeast-2. Add latency-based routing records in Amazon Route 53.

Answer: C

Explanation:

Adding an accelerator in AWS Global Accelerator will enable improving the performance of the APIs for local and global users¹. AWS Global Accelerator is a service that uses the AWS global network to route traffic to the optimal regional endpoint based on health, client location, and policies¹. Configuring the ALB as the origin will enable connecting the accelerator to the ALB that exposes the APIs². AWS Global Accelerator supports non-standard REST methods such as LINK, UNLINK, LOCK, and UNLOCK³.

NEW QUESTION 40

- (Exam Topic 2)

A solutions architect is planning to migrate critical Microsoft SQL Server databases to AWS. Because the databases are legacy systems, the solutions architect will move the databases to a modern data architecture. The solutions architect must migrate the databases with near-zero downtime. Which solution will meet these requirements?

- A. Use AWS Application Migration Service and the AWS Schema Conversion Tool (AWS SCT). Perform an In-place upgrade before the migration.
- B. Export the migrated data to Amazon Aurora Serverless after cutover.
- C. Repoint the applications to Amazon Aurora.
- D. Use AWS Database Migration Service (AWS DMS) to Rehost the databases.
- E. Set Amazon S3 as a target. Set up change data capture (CDC) replication.
- F. When the source and destination are fully synchronized, load the data from Amazon S3 into an Amazon RDS for Microsoft SQL Server DB Instance.
- G. Use native database high availability tools. Connect the source system to an Amazon RDS for Microsoft SQL Server DB instance. Configure replication accordingly.
- H. When data replication is finished, transition the workload to an Amazon RDS for Microsoft SQL Server DB instance.
- I. Use AWS Application Migration Service.
- J. Rehost the database server on Amazon EC2. When data replication is finished, detach the database and move the database to an Amazon RDS for Microsoft SQL Server DB instance.
- K. Reattach the database and then cut over all networking.

Answer: B

Explanation:

AWS DMS can migrate data from a source database to a target database in AWS, using change data capture (CDC) to replicate ongoing changes and keep the databases in sync. Setting Amazon S3 as a target allows storing the migrated data in a durable and cost-effective storage service. When the source and destination are fully synchronized, the data can be loaded from Amazon S3 into an Amazon RDS for Microsoft SQL Server DB instance, which is a managed database service that simplifies database administration tasks. References:

- > https://docs.aws.amazon.com/dms/latest/userguide/CHAP_Source.SQLServer.html
- > https://docs.aws.amazon.com/dms/latest/userguide/CHAP_Target.S3.html
- > https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_SQLServer.html

NEW QUESTION 42

- (Exam Topic 2)

A solutions architect must create a business case for migration of a company's on-premises data center to the AWS Cloud. The solutions architect will use a configuration management database (CMDB) export of all the company's servers to create the case. Which solution will meet these requirements MOST cost-effectively?

- A. Use AWS Well-Architected Tool to import the CMDB data to perform an analysis and generate recommendations.
- B. Use Migration Evaluator to perform an analysis.
- C. Use the data import template to upload the data from the CMDB export.
- D. Implement resource matching rule.
- E. Use the CMDB export and the AWS Price List Bulk API to query CMDB data against AWS services in bulk.
- F. Use AWS Application Discovery Service to import the CMDB data to perform an analysis.

Answer: B

Explanation:

<https://aws.amazon.com/blogs/architecture/accelerating-your-migration-to-aws/> Build a business case with AWS Migration Evaluator The foundation for a successful migration starts with a defined business objective (for example, growth or new offerings). In order to enable the business drivers, the established business case must then be aligned to a technical capability (increased security and elasticity). AWS Migration Evaluator (formerly known as TSO Logic) can help you meet these objectives. To get started, you can choose to upload exports from third-party tools such as Configuration Management Database (CMDB) or install a collector agent to monitor. You will receive an assessment after data collection, which includes a projected cost estimate and savings of running your on-premises workloads in the AWS Cloud. This estimate will provide a summary of the projected costs to re-host on AWS based on usage patterns. It will show the breakdown of costs by infrastructure and software licenses. With this information, you can make the business case and plan next steps.

NEW QUESTION 47

- (Exam Topic 2)

A company is updating an application that customers use to make online orders. The number of attacks on the application by bad actors has increased recently. The company will host the updated application on an Amazon Elastic Container Service (Amazon ECS) cluster. The company will use Amazon DynamoDB to store application data. A public Application Load Balancer (ALB) will provide end users with access to the application. The company must prevent prevent attacks and ensure business continuity with minimal service interruptions during an ongoing attack. Which combination of steps will meet these requirements MOST cost-effectively? (Select TWO.)

- A. Create an Amazon CloudFront distribution with the ALB as the origin.
- B. Add a custom header and random value on the CloudFront domain.
- C. Configure the ALB to conditionally forward traffic if the header and value match.

- D. Deploy the application in two AWS Region
- E. Configure Amazon Route 53 to route to both Regions with equal weight.
- F. Configure auto scaling for Amazon ECS task
- G. Create a DynamoDB Accelerator (DAX) cluster.
- H. Configure Amazon ElastiCache to reduce overhead on DynamoDB.
- I. Deploy an AWS WAF web ACL that includes an appropriate rule group
- J. Associate the web ACL with the Amazon CloudFront distribution.

Answer: AE

Explanation:

The company should create an Amazon CloudFront distribution with the ALB as the origin. The company should add a custom header and random value on the CloudFront domain. The company should configure the ALB to conditionally forward traffic if the header and value match. The company should also deploy an AWS WAF web ACL that includes an appropriate rule group. The company should associate the web ACL with the Amazon CloudFront distribution. This solution will meet the requirements most cost-effectively because Amazon CloudFront is a fast content delivery network (CDN) service that securely delivers data, videos, applications, and APIs to customers globally with low latency, high transfer speeds, all within a developer-friendly environment¹. By creating an Amazon CloudFront distribution with the ALB as the origin, the company can improve the performance and availability of its application by caching static content at edge locations closer to end users. By adding a custom header and random value on the CloudFront domain, the company can prevent direct access to the ALB and ensure that only requests from CloudFront are forwarded to the ECS tasks. By configuring the ALB to conditionally forward traffic if the header and value match, the company can implement origin access identity (OAI) for its ALB origin. OAI is a feature that enables you to restrict access to your content by requiring users to access your content through CloudFront URLs². By deploying an AWS WAF web ACL that includes an appropriate rule group, the company can prevent attacks and ensure business continuity with minimal service interruptions during an ongoing attack. AWS WAF is a web application firewall that lets you monitor and control web requests that are forwarded to your web applications. You can use AWS WAF to define customizable web security rules that control which traffic can access your web applications and which traffic should be blocked³. By associating the web ACL with the Amazon CloudFront distribution, the company can apply the web security rules to all requests that are forwarded by CloudFront.

The other options are not correct because:

- Deploying the application in two AWS Regions and configuring Amazon Route 53 to route to both Regions with equal weight would not prevent attacks or ensure business continuity. Amazon Route 53 is a highly available and scalable cloud Domain Name System (DNS) web service that routes end users to Internet applications by translating names like `www.example.com` into numeric IP addresses⁴. However, routing traffic to multiple Regions would not protect against attacks or provide failover in case of an outage. It would also increase operational complexity and costs compared to using CloudFront and AWS WAF.
- Configuring auto scaling for Amazon ECS tasks and creating a DynamoDB Accelerator (DAX) cluster would not prevent attacks or ensure business continuity. Auto scaling is a feature that enables you to automatically adjust your ECS tasks based on demand or a schedule. DynamoDB Accelerator (DAX) is a fully managed, highly available, in-memory cache for DynamoDB that delivers up to a 10x performance improvement. However, these features would not protect against attacks or provide failover in case of an outage. They would also increase operational complexity and costs compared to using CloudFront and AWS WAF.
- Configuring Amazon ElastiCache to reduce overhead on DynamoDB would not prevent attacks or ensure business continuity. Amazon ElastiCache is a fully managed in-memory data store service that makes it easy to deploy, operate, and scale popular open-source compatible in-memory data stores. However, this service would not protect against attacks or provide failover in case of an outage. It would also increase operational complexity and costs compared to using CloudFront and AWS WAF.

References:

- <https://aws.amazon.com/cloudfront/>
- <https://aws.amazon.com/waf/>
- <https://aws.amazon.com/route53/>
- <https://aws.amazon.com/dynamodb/dax/>
- <https://aws.amazon.com/elasticache/>

NEW QUESTION 52

- (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, primary, 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, primary, 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 primary 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, including all instances, when the processing is complete
- I. Purchase Compute Savings Plans to cover the On-Demand Instance usage
- J. Launch the primary and core nodes on On-Demand Instance
- K. Launch the task nodes on Spot Instances in an instance fleet
- L. Terminate only the task node instances when the processing is complete
- M. Purchase Compute Savings Plans to cover the On-Demand Instance usage.

Answer: A

Explanation:

Amazon EC2 Spot Instances offer spare compute capacity at steep discounts compared to On-Demand prices. Spot Instances can be interrupted by EC2 with two minutes of notification when EC2 needs the capacity back. Amazon EMR can handle Spot interruptions gracefully by decommissioning the nodes and redistributing the tasks to other nodes. By launching all nodes on Spot Instances in an instance fleet, the solutions architect can minimize the compute costs of the EMR cluster. An instance fleet is a collection of EC2 instances with different types and sizes that EMR automatically provisions to meet a defined target capacity. By terminating the cluster when the processing is completed, the solutions architect can avoid paying for idle resources. References:

- <https://docs.aws.amazon.com/emr/latest/ManagementGuide/emr-managed-scaling.html>
-

<https://docs.aws.amazon.com/emr/latest/ManagementGuide/emr-instance-fleet.html>

➤ <https://aws.amazon.com/blogs/big-data/optimizing-amazon-emr-for-resilience-and-cost-with-capacity-opt>

NEW QUESTION 55

- (Exam Topic 2)

A company is running an application on Amazon EC2 instances in the AWS Cloud. The application is using a MongoDB database with a replica set as its data tier. The MongoDB database is installed on systems in the company's on-premises data center and is accessible through an AWS Direct Connect connection to the data center environment.

A solutions architect must migrate the on-premises MongoDB database to Amazon DocumentDB (with MongoDB compatibility).

Which strategy should the solutions architect choose to perform this migration?

- A. Create a fleet of EC2 instance
- B. Install MongoDB Community Edition on the EC2 instances, and create a databas
- C. Configure continuous synchronous replication with the database that is running in the on-premises data center.
- D. Create an AWS Database Migration Service (AWS DMS) replication instanc
- E. Create a source endpoint for the on-premises MongoDB database by using change data capture (CDC). Create a target endpoint for the Amazon DocumentDB databas
- F. Create and run a DMS migration task.
- G. Create a data migration pipeline by using AWS Data Pipelin
- H. Define data nodes for the on-premises MongoDB database and the Amazon DocumentDB databas
- I. Create a scheduled task to run the data pipeline.
- J. Create a source endpoint for the on-premises MongoDB database by using AWS Glue crawlers. Configure continuous asynchronous replication between the MongoDB database and the Amazon DocumentDB database.

Answer: B

Explanation:

<https://aws.amazon.com/getting-started/hands-on/move-to-managed/migrate-mongodb-to-documentdb/>

NEW QUESTION 59

- (Exam Topic 2)

A company wants to optimize AWS data-transfer costs and compute costs across developer accounts within the company's organization in AWS Organizations. Developers can configure VPCs and launch Amazon EC2 instances in a single AWS Region. The EC2 instances retrieve approximately 1 TB of data each day from Amazon S3.

The developer activity leads to excessive monthly data-transfer charges and NAT gateway processing charges between EC2 instances and S3 buckets, along with high compute costs. The company wants to proactively enforce approved architectural patterns for any EC2 instance and VPC infrastructure that developers deploy within the AWS accounts. The company does not want this enforcement to negatively affect the speed at which the developers can perform their tasks.

Which solution will meet these requirements MOST cost-effectively?

- A. Create SCPs to prevent developers from launching unapproved EC2 instance types. Provide the developers with an AWS CloudFormation template to deploy an approved VPC configuration with S3 interface endpoints. Scope the developers' IAM permissions so that the developers can launch VPC resources only with CloudFormation.
- B. Create a daily forecasted budget with AWS Budgets to monitor EC2 compute costs and S3 data-transfer costs across the developer accounts. When the forecasted cost is 75% of the actual budget cost, send an alert to the developer teams. If the actual budget cost is 100%, create a budget action to terminate the developers' EC2 instances and VPC infrastructure.
- C. Create an AWS Service Catalog portfolio that users can use to create an approved VPC configuration with S3 gateway endpoints and approved EC2 instances. Share the portfolio with the developer accounts. Configure an AWS Service Catalog launch constraint to use an approved IAM role. Scope the developers' IAM permissions to allow access only to AWS Service Catalog.
- D. Create and deploy AWS Config rules to monitor the compliance of EC2 and VPC resources in the developer AWS accounts. If developers launch unapproved EC2 instances or if developers create VPCs without S3 gateway endpoints, perform a remediation action to terminate the unapproved resources.

Answer: C

Explanation:

This solution allows developers to quickly launch resources using pre-approved configurations and instance types, while also ensuring that the resources launched comply with the company's architectural patterns. This can help reduce data transfer and compute costs associated with the resources. Using AWS Service Catalog also allows the company to control access to the approved configurations and resources through the use of IAM roles, while also allowing developers to quickly provision resources without negatively affecting their ability to perform their tasks.

Reference:

AWS Service Catalog: <https://aws.amazon.com/service-catalog/> AWS Service Catalog Constraints:

<https://docs.aws.amazon.com/servicecatalog/latest/adminguide/constraints.html>

AWS Service Catalog Launch Constraints: <https://docs.aws.amazon.com/servicecatalog/latest/adminguide/launch-constraints.html>

NEW QUESTION 61

- (Exam Topic 2)

A company uses an AWS CodeCommit repository. The company must store a backup copy of the data that is in the repository in a second AWS Region.

Which solution will meet these requirements?

- A. Configure AWS Elastic Disaster Recovery to replicate the CodeCommit repository data to the second Region.
- B. Use AWS Backup to back up the CodeCommit repository on an hourly schedule. Create a cross-Region copy in the second Region.
- C. Create an Amazon EventBridge rule to invoke AWS CodeBuild when the company pushes code to the repository. Use CodeBuild to clone the repository. Create a zip file of the content. Copy the file to an S3 bucket in the second Region.
- D. Create an AWS Step Functions workflow on an hourly schedule to take a snapshot of the CodeCommit repository. Configure the workflow to copy the snapshot to an S3 bucket in the second Region.

Answer: B

Explanation:

AWS Backup is a fully managed service that makes it easy to centralize and automate the creation, retention, and restoration of backups across AWS services. It

provides a way to schedule automatic backups for CodeCommit repositories on an hourly basis. Additionally, it also supports cross-Region replication, which allows you to copy the backups to a second Region for disaster recovery.

By using AWS Backup, the company can set up an automatic and regular backup schedule for the CodeCommit repository, ensuring that the data is regularly backed up and stored in a second Region. This can provide a way to recover quickly from any disaster event that might occur.

Reference:

AWS Backup documentation: <https://aws.amazon.com/backup/> AWS Backup for AWS CodeCommit documentation:

<https://aws.amazon.com/about-aws/whats-new/2020/07/aws-backup-now-supports-aws-codecommit-repositorie>

NEW QUESTION 66

- (Exam Topic 2)

A company has an on-premises Microsoft SQL Server database that writes a nightly 200 GB export to a local drive. The company wants to move the backups to more robust cloud storage on Amazon S3. The company has set up a 10 Gbps AWS Direct Connect connection between the on-premises data center and AWS. Which solution meets these requirements MOST cost-effectively?

- A. Create a new S3 bucket
- B. Deploy an AWS Storage Gateway file gateway within the VPC that is connected to the Direct Connect connection
- C. Create a new SMB file share
- D. Write nightly database exports to the new SMB file share.
- E. Create an Amazon FSx for Windows File Server Single-AZ file system within the VPC that is connected to the Direct Connect connection
- F. Create a new SMB file share
- G. Write nightly database exports to an SMB file share on the Amazon FSx file system
- H. Enable nightly backups.
- I. Create an Amazon FSx for Windows File Server Multi-AZ file system within the VPC that is connected to the Direct Connect connection
- J. Create a new SMB file share
- K. Write nightly database exports to an SMB file share on the Amazon FSx file system
- L. Enable nightly backups.
- M. Create a new S3 bucket
- N. Deploy an AWS Storage Gateway volume gateway within the VPC that is connected to the Direct Connect connection
- O. Create a new SMB file share
- P. Write nightly database exports to the new SMB file share on the volume gateway, and automate copies of this data to an S3 bucket.

Answer: A

Explanation:

<https://docs.aws.amazon.com/filegateway/latest/files3/CreatingAnSMBFileShare.html>

NEW QUESTION 71

- (Exam Topic 2)

A company's public API runs as tasks on Amazon Elastic Container Service (Amazon ECS). The tasks run on AWS Fargate behind an Application Load Balancer (ALB) and are configured with Service Auto Scaling for the tasks based on CPU utilization. This service has been running well for several months.

Recently, API performance slowed down and made the application unusable. The company discovered that a significant number of SQL injection attacks had occurred against the API and that the API service had scaled to its maximum amount.

A solutions architect needs to implement a solution that prevents SQL injection attacks from reaching the ECS API service. The solution must allow legitimate traffic through and must maximize operational efficiency. Which solution meets these requirements?

- A. Create a new AWS WAF web ACL to monitor the HTTP requests and HTTPS requests that are forwarded to the ALB in front of the ECS tasks.
- B. Create a new AWS WAF Bot Control implementation
- C. Add a rule in the AWS WAF Bot Control managed rule group to monitor traffic and allow only legitimate traffic to the ALB in front of the ECS tasks.
- D. Create a new AWS WAF web ACL
- E. Add a new rule that blocks requests that match the SQL database rule group
- F. Set the web ACL to allow all other traffic that does not match those rules
- G. Attach the web ACL to the ALB in front of the ECS tasks.
- H. Create a new AWS WAF web ACL
- I. Create a new empty IP set in AWS IAM
- J. Add a new rule to the web ACL to block requests that originate from IP addresses in the new IP set
- K. Create an AWS Lambda function that scrapes the API logs for IP addresses that send SQL injection attacks, and add those IP addresses to the IP set
- L. Attach the web ACL to the ALB in front of the ECS tasks.

Answer: C

Explanation:

The company should create a new AWS WAF web ACL. The company should add a new rule that blocks requests that match the SQL database rule group. The company should set the web ACL to allow all other traffic that does not match those rules. The company should attach the web ACL to the ALB in front of the ECS tasks. This solution will meet the requirements because AWS WAF is a web application firewall that lets you monitor and control web requests that are forwarded to your web applications. You can use AWS WAF to define customizable web security rules that control which traffic can access your web applications and which traffic should be blocked¹. By creating a new AWS WAF web ACL, the company can create a collection of rules that define the conditions for allowing or blocking web requests. By adding a new rule that blocks requests that match the SQL database rule group, the company can prevent SQL injection attacks from reaching the ECS API service. The SQL database rule group is a managed rule group provided by AWS that contains rules to protect against common SQL injection attack patterns². By setting the web ACL to allow all other traffic that does not match those rules, the company can ensure that legitimate traffic can access the API service. By attaching the web ACL to the ALB in front of the ECS tasks, the company can apply the web security rules to all requests that are forwarded by the load balancer.

The other options are not correct because:

➤ Creating a new AWS WAF Bot Control implementation would not prevent SQL injection attacks from reaching the ECS API service. AWS WAF Bot Control is a feature that gives you visibility and control over common and pervasive bot traffic that can consume excess resources, skew metrics, cause downtime, or perform other undesired activities. However, it does not protect against SQL injection attacks, which are malicious attempts to execute unauthorized SQL statements against your database³.

➤ Creating a new AWS WAF web ACL to monitor the HTTP requests and HTTPS requests that are forwarded to the ALB in front of the ECS tasks would not prevent SQL injection attacks from reaching the ECS API service. Monitoring mode is a feature that enables you to evaluate how your rules would perform without actually blocking any requests. However, this mode does not provide any protection against attacks, as it only logs and counts requests that match your rules⁴.

➤ Creating a new AWS WAF web ACL and creating a new empty IP set in AWS WAF would not prevent SQL injection attacks from reaching the ECS API service. An IP set is a feature that enables you to specify a list of IP addresses or CIDR blocks that you want to allow or block based on their source IP address. However, this approach would not be effective or efficient against SQL injection attacks, as it would require constantly updating the IP set with new IP addresses of attackers, and it would not block attackers who use proxies or VPNs.

References:

- <https://aws.amazon.com/waf/>
- <https://docs.aws.amazon.com/waf/latest/developerguide/waf-bot-control.html>
- <https://docs.aws.amazon.com/waf/latest/developerguide/web-acl-monitoring-mode.html>
- <https://docs.aws.amazon.com/waf/latest/developerguide/waf-ip-sets.html>

NEW QUESTION 75

- (Exam Topic 2)

A company runs an application in an on-premises data center. The application gives users the ability to upload media files. The files persist in a file server. The web application has many users. The application server is overutilized, which causes data uploads to fail occasionally. The company frequently adds new storage to the file server. The company wants to resolve these challenges by migrating the application to AWS.

Users from across the United States and Canada access the application. Only authenticated users should have the ability to access the application to upload files. The company will consider a solution that refactors the application, and the company needs to accelerate application development.

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

- A. Use AWS Application Migration Service to migrate the application server to Amazon EC2 instances. Create an Auto Scaling group for the EC2 instance
- B. Use an Application Load Balancer to distribute the request
- C. Modify the application to use Amazon S3 to persist the file
- D. Use Amazon Cognito to authenticate users.
- E. Use AWS Application Migration Service to migrate the application server to Amazon EC2 instances. Create an Auto Scaling group for the EC2 instance
- F. Use an Application Load Balancer to distribute the request
- G. Set up AWS IAM Identity Center (AWS Single Sign-On) to give users the ability to sign in to the application
- H. Modify the application to use Amazon S3 to persist the files.
- I. Create a static website for uploads of media file
- J. Store the static assets in Amazon S3. Use AWS AppSync to create an API
- K. Use AWS Lambda resolvers to upload the media files to Amazon S3. Use Amazon Cognito to authenticate users.
- L. Use AWS Amplify to create a static website for uploads of media file
- M. Use Amplify Hosting to serve the website through Amazon CloudFront
- N. Use Amazon S3 to store the uploaded media file
- O. Use Amazon Cognito to authenticate users.

Answer: D

Explanation:

The company should use AWS Amplify to create a static website for uploads of media files. The company should use Amplify Hosting to serve the website through Amazon CloudFront. The company should use Amazon S3 to store the uploaded media files. The company should use Amazon Cognito to authenticate users.

This solution will meet the requirements with the least operational overhead because AWS Amplify is a complete solution that lets frontend web and mobile developers easily build, ship, and host full-stack applications on AWS, with the flexibility to leverage the breadth of AWS services as use cases evolve. No cloud expertise needed¹. By using AWS Amplify, the company can refactor the application to a serverless architecture that reduces operational complexity and costs.

AWS Amplify offers the following features and benefits:

- Amplify Studio: A visual interface that enables you to build and deploy a full-stack app quickly, including frontend UI and backend.
- Amplify CLI: A local toolchain that enables you to configure and manage an app backend with just a few commands.
- Amplify Libraries: Open-source client libraries that enable you to build cloud-powered mobile and web apps.
- Amplify UI Components: Open-source design system with cloud-connected components for building feature-rich apps fast.
- Amplify Hosting: Fully managed CI/CD and hosting for fast, secure, and reliable static and server-side rendered apps.

By using AWS Amplify to create a static website for uploads of media files, the company can leverage Amplify Studio to visually build a pixel-perfect UI and connect it to a cloud backend in clicks. By using Amplify Hosting to serve the website through Amazon CloudFront, the company can easily deploy its web app or website to the fast, secure, and reliable AWS content delivery network (CDN), with hundreds of points of presence globally. By using Amazon S3 to store the uploaded media files, the company can benefit from a highly scalable, durable, and cost-effective object storage service that can handle any amount of data². By using Amazon Cognito to authenticate users, the company can add user sign-up, sign-in, and access control to its web app with a fully managed service that scales to support millions of users³.

The other options are not correct because:

➤ Using AWS Application Migration Service to migrate the application server to Amazon EC2 instances would not refactor the application or accelerate development. AWS Application Migration Service (AWS MGN) is a service that enables you to migrate physical servers, virtual machines (VMs), or cloud servers from any source infrastructure to AWS without requiring agents or specialized tools. However, this would not address the challenges of overutilization and data uploads failures. It would also not reduce operational overhead or costs compared to a serverless architecture.

➤ Creating a static website for uploads of media files and using AWS AppSync to create an API would not be as simple or fast as using AWS Amplify. AWS AppSync is a service that enables you to create flexible APIs for securely accessing, manipulating, and combining data from one or more data sources. However, this would require more configuration and management than using Amplify Studio and Amplify Hosting. It would also not provide authentication features like Amazon Cognito.

➤ Setting up AWS IAM Identity Center (AWS Single Sign-On) to give users the ability to sign in to the application would not be as suitable as using Amazon Cognito. AWS Single Sign-On (AWS SSO) is a service that enables you to centrally manage SSO access and user permissions across multiple AWS accounts and business applications. However, this service is designed for enterprise customers who need to manage access for employees or partners across multiple resources. It is not intended for authenticating end users of web or mobile apps.

References:

- <https://aws.amazon.com/amplify/>
- <https://aws.amazon.com/s3/>
- <https://aws.amazon.com/cognito/>
- <https://aws.amazon.com/mgn/>
- <https://aws.amazon.com/appsync/>
- <https://aws.amazon.com/single-sign-on/>

NEW QUESTION 77

- (Exam Topic 2)

A company has migrated a legacy application to the AWS Cloud. The application runs on three Amazon EC2 instances that are spread across three Availability Zones. One EC2 instance is in each Availability Zone. The EC2 instances are running in three private subnets of the VPC and are set up as targets for an Application Load Balancer (ALB) that is associated with three public subnets.

The application needs to communicate with on-premises systems. Only traffic from IP addresses in the company's IP address range are allowed to access the on-premises systems. The company's security team is bringing only one IP address from its internal IP address range to the cloud. The company has added this IP address to the allow list for the company firewall. The company also has created an Elastic IP address for this IP address.

A solutions architect needs to create a solution that gives the application the ability to communicate with the on-premises systems. The solution also must be able to mitigate failures automatically.

Which solution will meet these requirements?

- A. Deploy three NAT gateways, one in each public subnet
- B. Assign the Elastic IP address to the NAT gateway
- C. Turn on health checks for the NAT gateway
- D. If a NAT gateway fails a health check, recreate the NAT gateway and assign the Elastic IP address to the new NAT gateway.
- E. Replace the ALB with a Network Load Balancer (NLB). Assign the Elastic IP address to the NLB. Turn on health checks for the NLB.
- F. In the case of a failed health check, redeploy the NLB in different subnets.
- G. Deploy a single NAT gateway in a public subnet
- H. Assign the Elastic IP address to the NAT gateway. Use Amazon CloudWatch with a custom metric to monitor the NAT gateway.
- I. If the NAT gateway is unhealthy, invoke an AWS Lambda function to create a new NAT gateway in a different subnet
- J. Assign the Elastic IP address to the new NAT gateway.
- K. Assign the Elastic IP address to the ALB
- L. Create an Amazon Route 53 simple record with the Elastic IP address as the value
- M. Create a Route 53 health check
- N. In the case of a failed health check, recreate the ALB in different subnets.

Answer: C

Explanation:

to connect out from the private subnet you need an NAT gateway and since only one Elastic IP whitelisted on firewall its one NATGateway at time and if AZ failure happens Lambda creates a new NATGATEWAY in a different AZ using the Same Elastic IP ,dont be tempted to select D since application that needs to connect is on a private subnet whose outbound connections use the NATGateway Elastic IP

NEW QUESTION 82

- (Exam Topic 2)

A company is running a web application in a VPC. The web application runs on a group of Amazon EC2 instances behind an Application Load Balancer (ALB). The ALB is using AWS WAF.

An external customer needs to connect to the web application. The company must provide IP addresses to all external customers.

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

- A. Replace the ALB with a Network Load Balancer (NLB). Assign an Elastic IP address to the NLB.
- B. Allocate an Elastic IP address
- C. Assign the Elastic IP address to the ALB. Provide the Elastic IP address to the customer.
- D. Create an AWS Global Accelerator standard accelerator
- E. Specify the ALB as the accelerator's endpoint. Provide the accelerator's IP addresses to the customer.
- F. Configure an Amazon CloudFront distribution
- G. Set the ALB as the origin
- H. Ping the distribution's DNS name to determine the distribution's public IP address
- I. Provide the IP address to the customer.

Answer: C

Explanation:

<https://docs.aws.amazon.com/global-accelerator/latest/dg/about-accelerators.alb-accelerator.html> Option A is wrong. AWS WAF does not support associating with NLB.

<https://docs.aws.amazon.com/waf/latest/developerguide/waf-chapter.html> Option B is wrong. An ALB does not support an Elastic IP address.

<https://aws.amazon.com/elasticloadbalancing/features/>

NEW QUESTION 84

- (Exam Topic 2)

A company uses a Grafana data visualization solution that runs on a single Amazon EC2 instance to monitor the health of the company's AWS workloads. The company has invested time and effort to create dashboards that the company wants to preserve. The dashboards need to be highly available and cannot be down for longer than 10 minutes. The company needs to minimize ongoing maintenance.

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

- A. Migrate to Amazon CloudWatch dashboard
- B. Recreate the dashboards to match the existing Grafana dashboard
- C. Use automatic dashboards where possible.
- D. Create an Amazon Managed Grafana workspace
- E. Configure a new Amazon CloudWatch data source. Export dashboards from the existing Grafana instance
- F. Import the dashboards into the new workspace.
- G. Create an AMI that has Grafana pre-installed
- H. Store the existing dashboards in Amazon Elastic File System (Amazon EFS). Create an Auto Scaling group that uses the new AMI
- I. Set the Auto Scaling group's minimum, desired, and maximum number of instances to one
- J. Create an Application Load Balancer that serves at least two Availability Zones.
- K. Configure AWS Backup to back up the EC2 instance that runs Grafana once each hour
- L. Restore the EC2 instance from the most recent snapshot in an alternate Availability Zone when required.

Answer: C

Explanation:

By creating an AMI that has Grafana pre-installed and storing the existing dashboards in Amazon Elastic File System (Amazon EFS) it allows for faster and more efficient scaling, and by creating an Auto Scaling group that uses the new AMI and setting the Auto Scaling group's minimum, desired, and maximum number of instances to one and creating an Application Load Balancer that serves at least two Availability Zones, it ensures high availability and minimized downtime.

NEW QUESTION 86

- (Exam Topic 2)

A company runs an intranet application on premises. The company wants to configure a cloud backup of the application. The company has selected AWS Elastic Disaster Recovery for this solution.

The company requires that replication traffic does not travel through the public internet. The application also must not be accessible from the internet. The company does not want this solution to consume all available network bandwidth because other applications require bandwidth.

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

- A. Create a VPC that has at least two private subnets, two NAT gateways, and a virtual private gateway.
- B. Create a VPC that has at least two public subnets, a virtual private gateway, and an internet gateway.
- C. Create an AWS Site-to-Site VPN connection between the on-premises network and the target AWS network.
- D. Create an AWS Direct Connect connection and a Direct Connect gateway between the on-premises network and the target AWS network.
- E. During configuration of the replication servers, select the option to use private IP addresses for data replication.
- F. During configuration of the launch settings for the target servers, select the option to ensure that the Recovery instance's private IP address matches the source server's private IP address.

Answer: BDE

Explanation:

AWS Elastic Disaster Recovery (AWS DRS) is a service that minimizes downtime and data loss with fast, reliable recovery of on-premises and cloud-based applications using affordable storage, minimal compute, and point-in-time recovery¹. Users can set up AWS DRS on their source servers to initiate secure data replication to a staging area subnet in their AWS account, in the AWS Region they select. Users can then launch recovery instances on AWS within minutes, using the most up-to-date server state or a previous point in time.

To configure a cloud backup of the application with AWS DRS, users need to create a VPC that has at least two public subnets, a virtual private gateway, and an internet gateway. A VPC is a logically isolated section of the AWS Cloud where users can launch AWS resources in a virtual network that they define². A public subnet is a subnet that has a route to an internet gateway³. A virtual private gateway is the VPN concentrator on the Amazon side of the Site-to-Site VPN connection⁴. An internet gateway is a horizontally scaled, redundant, and highly available VPC component that allows communication between instances in the VPC and the internet. Users need to create at least two public subnets for redundancy and high availability. Users need to create a virtual private gateway and attach it to the VPC to enable VPN connectivity between the on-premises network and the target AWS network. Users need to create an internet gateway and attach it to the VPC to enable internet access for the replication servers.

To ensure that replication traffic does not travel through the public internet, users need to create an AWS Direct Connect connection and a Direct Connect gateway between the on-premises network and the target AWS network. AWS Direct Connect is a service that establishes a dedicated network connection from an on-premises network to one or more VPCs. A Direct Connect gateway is a globally available resource that allows users to connect multiple VPCs across different Regions to their on-premises networks using one or more Direct Connect connections. Users need to create an AWS Direct Connect connection between their on-premises network and an AWS Region. Users need to create a Direct Connect gateway and associate it with their VPC and their Direct Connect connection.

To ensure that the application is not accessible from the internet, users need to select the option to use private IP addresses for data replication during configuration of the replication servers. This option configures the replication servers with private IP addresses only, without assigning any public IP addresses or Elastic IP addresses. This way, the replication servers can only communicate with other resources within the VPC or through VPN connections.

Option A is incorrect because creating a VPC that has at least two private subnets, two NAT gateways, and a virtual private gateway is not necessary or cost-effective. A private subnet is a subnet that does not have a route to an internet gateway³. A NAT gateway is a highly available, managed Network Address Translation (NAT) service that enables instances in a private subnet to connect to the internet or other AWS services, but prevents the internet from initiating connections with those instances. Users do not need to create private subnets or NAT gateways for this use case, as they can use public subnets with private IP addresses for data replication.

Option C is incorrect because creating an AWS Site-to-Site VPN connection between the on-premises network and the target AWS network will not ensure that replication traffic does not travel through the public internet. A Site-to-Site VPN connection consists of two VPN tunnels between an on-premises customer gateway device and a virtual private gateway in your VPC⁴. The VPN tunnels are encrypted using IPsec protocols, but they still use public IP addresses for communication. Users need to use AWS Direct Connect instead of Site-to-Site VPN for this use case.

Option F is incorrect because selecting the option to ensure that the Recovery instance's private IP address matches the source server's private IP address during configuration of the launch settings for the target servers will not ensure that the application is not accessible from the internet. This option configures the Recovery instance with an identical private IP address as its source server when launched in drills or recovery mode. However, this option does not prevent assigning public IP addresses or Elastic IP addresses to the Recovery instance. Users need to select the option to use private IP addresses for data replication instead.

NEW QUESTION 87

- (Exam Topic 2)

A company has many separate AWS accounts and uses no central billing or management. Each AWS account hosts services for different departments in the company. The company has a Microsoft Azure Active Directory that is deployed.

A solution architect needs to centralize billing and management of the company's AWS accounts. The company wants to start using identify federation instead of manual user management. The company also wants to use temporary credentials instead of long-lived access keys.

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

- A. Create a new AWS account to serve as a management account
- B. Deploy an organization in AWS Organization
- C. Invite each existing AWS account to join the organization
- D. Ensure that each account accepts the invitation.
- E. Configure each AWS Account's email address to be aws+<account id>@example.com so that account management email messages and invoices are sent to the same place.
- F. Deploy AWS IAM Identity Center (AWS Single Sign-On) in the management account
- G. Connect IAM Identity Center to the Azure Active Director
- H. Configure IAM Identity Center for automatic synchronization of users and groups.
- I. Deploy an AWS Managed Microsoft AD directory in the management account
- J. Share the directory with all other accounts in the organization by using AWS Resource Access Manager (AWS RAM).
- K. Create AWS IAM Identity Center (AWS Single Sign-On) permission set
- L. Attach the permission sets to the appropriate IAM Identity Center groups and AWS accounts.
- M. Configure AWS Identity and Access Management (IAM) in each AWS account to use AWS Managed Microsoft AD for authentication and authorization.

Answer: ACE

NEW QUESTION 89

- (Exam Topic 2)

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 VP
- 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 VP
- 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 VP
- H. Create a Route 53 outbound resolver in the shared services VP
- 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 VP
- K. Create a Route 53 inbound resolver in the shared services VP
- 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: A

Explanation:

Amazon Route 53 Resolver is a managed DNS resolver service from Route 53 that helps to create conditional forwarding rules to redirect query traffic¹. By associating the private hosted zone to all the VPCs, the solutions architect can enable DNS resolution for cloud.example.com within the VPCs. By creating a Route 53 inbound resolver in the shared services VPC, the solutions architect can enable DNS resolution for cloud.example.com from on-premises systems. By attaching all VPCs to the transit gateway, the solutions architect can enable connectivity between the VPCs and the on-premises network through AWS Direct Connect. By creating forwarding rules in the on-premises DNS server for cloud.example.com that point to the inbound resolver, the solutions architect can direct DNS queries for cloud.example.com to the Route 53 Resolver endpoint in AWS. This solution will provide the highest performance as it leverages Route 53 Resolver's optimized routing and caching capabilities.

References: 1: <https://aws.amazon.com/route53/resolver/>

NEW QUESTION 94

- (Exam Topic 2)

A company operates an on-premises software-as-a-service (SaaS) solution that ingests several files daily. The company provides multiple public SFTP endpoints to its customers to facilitate the file transfers. The customers add the SFTP endpoint IP addresses to their firewall allow list for outbound traffic. Changes to the SFTP endpoint IP addresses are not permitted.

The company wants to migrate the SaaS solution to AWS and decrease the operational overhead of the file transfer service.

Which solution meets these requirements?

- A. Register the customer-owned block of IP addresses in the company's AWS account
- B. Create Elastic IP addresses from the address pool and assign them to an AWS Transfer for SFTP endpoint
- C. Use AWS Transfer to store the files in Amazon S3.
- D. Add a subnet containing the customer-owned block of IP addresses to a VPC. Create Elastic IP addresses from the address pool and assign them to an Application Load Balancer (ALB). Launch EC2 instances hosting FTP services in an Auto Scaling group behind the ALB.
- E. Store the files in attached Amazon Elastic Block Store (Amazon EBS) volumes.
- F. Register the customer-owned block of IP addresses with Amazon Route 53. Create alias records in Route 53 that point to a Network Load Balancer (NLB). Launch EC2 instances hosting FTP services in an Auto Scaling group behind the NLB.
- G. Store the files in Amazon S3.
- H. Register the customer-owned block of IP addresses in the company's AWS account
- I. Create Elastic IP addresses from the address pool and assign them to an Amazon S3 VPC endpoint
- J. Enable SFTP support on the S3 bucket.

Answer: A

Explanation:

Bring your own IP addresses (BYOIP) You can bring part or all of your publicly routable IPv4 or IPv6 address range from your on-premises network to your AWS account. You continue to own the address range, but AWS advertises it on the internet by default. After you bring the address range to AWS, it appears in your AWS account as an address pool. <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-byoip.html> AWS Transfer for SFTP enables you to easily move your file transfer workloads that use the Secure Shell File Transfer Protocol (SFTP) to AWS without needing to modify your applications or manage any SFTP servers. <https://aws.amazon.com/about-aws/whats-new/2018/11/aws-transfer-for-sftp-fully-managed-sftp-for-s3/>

NEW QUESTION 98

- (Exam Topic 2)

A company is migrating a document processing workload to AWS. The company has updated many applications to natively use the Amazon S3 API to store, retrieve, and modify documents that a processing server generates at a rate of approximately 5 documents every second. After the document processing is finished, customers can download the documents directly from Amazon S3.

During the migration, the company discovered that it could not immediately update the processing server that generates many documents to support the S3 API. The server runs on Linux and requires fast local access to the files that the server generates and modifies. When the server finishes processing, the files must be available to the public for download within 30 minutes.

Which solution will meet these requirements with the LEAST amount of effort?

- A. Migrate the application to an AWS Lambda function
- B. Use the AWS SDK for Java to generate, modify, and access the files that the company stores directly in Amazon S3.

- C. Set up an Amazon S3 File Gateway and configure a file share that is linked to the document store. Mount the file share on an Amazon EC2 instance by using NFS.
- D. When changes occur in Amazon S3, initiate a RefreshCache API call to update the S3 File Gateway.
- E. Configure Amazon FSx for Lustre with an import and export policy.
- F. Link the new file system to an S3 bucket.
- G. Install the Lustre client and mount the document store to an Amazon EC2 instance by using NFS.
- H. Configure AWS DataSync to connect to an Amazon EC2 instance.
- I. Configure a task to synchronize the generated files to and from Amazon S3.

Answer: C

Explanation:

Amazon FSx for Lustre is a fully managed service that provides cost-effective, high-performance, scalable storage for compute workloads. Powered by Lustre, the world's most popular high-performance file system, FSx for Lustre offers shared storage with sub-ms latencies, up to terabytes per second of throughput, and millions of IOPS. FSx for Lustre file systems can also be linked to Amazon Simple Storage Service (S3) buckets, allowing you to access and process data concurrently from both a high-performance file system and from the S3 API.

NEW QUESTION 99

- (Exam Topic 2)

A solutions architect wants to cost-optimize and appropriately size Amazon EC2 instances in a single AWS account. The solutions architect wants to ensure that the instances are optimized based on CPU, memory, and network metrics.

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

- A. Purchase AWS Business Support or AWS Enterprise Support for the account.
- B. Turn on AWS Trusted Advisor and review any "Low Utilization Amazon EC2 Instances" recommendations.
- C. Install the Amazon CloudWatch agent and configure memory metric collection on the EC2 instances.
- D. Configure AWS Compute Optimizer in the AWS account to receive findings and optimization recommendations.
- E. Create an EC2 Instance Savings Plan for the AWS Regions, instance families, and operating systems of interest.

Answer: BD

Explanation:

AWS Trusted Advisor is a service that provides real-time guidance to help users provision their resources following AWS best practices¹. One of the Trusted Advisor checks is "Low Utilization Amazon EC2 Instances", which identifies EC2 instances that appear to be underutilized based on CPU, network I/O, and disk I/O metrics¹. This check can help users optimize the cost and size of their EC2 instances by recommending smaller or more appropriate instance types.

AWS Compute Optimizer is a service that analyzes the configuration and utilization metrics of AWS resources and generates optimization recommendations to reduce the cost and improve the performance of workloads². Compute Optimizer supports four types of AWS resources: EC2 instances, EBS volumes, ECS services on AWS Fargate, and Lambda functions². For EC2 instances, Compute Optimizer evaluates the vCPUs, memory, storage, and other specifications, as well as the CPU utilization, network in and out, disk read and write, and other utilization metrics of currently running instances³. It then recommends optimal instance types based on price-performance trade-offs.

Option A is incorrect because purchasing AWS Business Support or AWS Enterprise Support for the account will not directly help with cost-optimization and sizing of EC2 instances. However, these support plans do provide access to more Trusted Advisor checks than the basic support plan¹.

Option C is incorrect because installing the Amazon CloudWatch agent and configuring memory metric collection on the EC2 instances will not provide any optimization recommendations by itself. However, memory metrics can be used by Compute Optimizer to enhance its recommendations if enabled³.

Option E is incorrect because creating an EC2 Instance Savings Plan for the AWS Regions, instance families, and operating systems of interest will not help with cost-optimization and sizing of EC2 instances. Savings Plans are a flexible pricing model that offer lower prices on Amazon EC2 usage in exchange for a commitment to a consistent amount of usage for a 1- or 3-year term⁴. Savings Plans do not affect the configuration or utilization of EC2 instances.

NEW QUESTION 104

- (Exam Topic 2)

A company has built a high performance computing (HPC) cluster in AWS for a tightly coupled workload that generates a large number of shared files stored in Amazon EFS. The cluster was performing well when the number of Amazon EC2 instances in the cluster was 100. However, when the company increased the cluster size to 1,000 EC2 instances, overall performance was well below expectations.

Which collection of design choices should a solutions architect make to achieve the maximum performance from the HPC cluster? (Select THREE.)

- A. Ensure the HPC cluster is launched within a single Availability Zone.
- B. Launch the EC2 instances and attach elastic network interfaces in multiples of four.
- C. Select EC2 Instance types with an Elastic Fabric Adapter (EFA) enabled.
- D. Ensure the cluster is launched across multiple Availability Zones.
- E. Replace Amazon EFS with multiple Amazon EBS volumes in a RAID array.
- F. Replace Amazon EFS with Amazon FSx for Lustre.

Answer: ACF

Explanation:

* A. High performance computing (HPC) workload cluster should be in a single AZ.

* C. Elastic Fabric Adapter (EFA) is a network device that you can attach to your Amazon EC2 instances to accelerate High Performance Computing (HPC)

* F. Amazon FSx for Lustre - Use it for workloads where speed matters, such as machine learning, high performance computing (HPC), video processing, and financial modeling.

Cluster – packs instances close together inside an Availability Zone. This strategy enables workloads to achieve the low-latency network performance necessary for tightly-coupled node-to-node communication that is typical of HPC applications.

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/placement-groups.html>

NEW QUESTION 106

- (Exam Topic 2)

A company uses AWS Organizations for a multi-account setup in the AWS Cloud. The company's finance team has a data processing application that uses AWS Lambda and Amazon DynamoDB. The company's marketing team wants to access the data that is stored in the DynamoDB table.

The DynamoDB table contains confidential data. The marketing team can have access to only specific attributes of data in the DynamoDB table. The finance team and the marketing team have separate AWS accounts.

What should a solutions architect do to provide the marketing team with the appropriate access to the DynamoDB table?

- A. Create an SCP to grant the marketing team's AWS account access to the specific attributes of the DynamoDB tabl
- B. Attach the SCP to the OU of the finance team.
- C. Create an IAM role in the finance team's account by using IAM policy conditions for specific DynamoDB attributes (fine-grained access con-trol). Establish trust with the marketing team's account
- D. In the mar-keting team's account, create an IAM role that has permissions to as-sume the IAM role in the finance team's account.
- E. Create a resource-based IAM policy that includes conditions for spe-cific DynamoDB attributes (fine-grained access control). Attach the policy to the DynamoDB tabl
- F. In the marketing team'saccount, create an IAM role that has permissions to access the DynamoDB table in the finance team's account.
- G. Create an IAM role in the finance team's account to access the Dyna-moDB tabl
- H. Use an IAM permissions boundary to limit the access to the specific attribute
- I. In the marketing team's account, create an IAM role that has permissions to assume the IAM role in the finance team's account.

Answer: C

Explanation:

The company should create a resource-based IAM policy that includes conditions for specific DynamoDB attributes (fine-grained access control). The company should attach the policy to the DynamoDB table. In the marketing team's account, the company should create an IAM role that has permissions to access the DynamoDB table in the finance team's account. This solution will meet the requirements because a resource-based IAM policy is a policy that you attach to an AWS resource (such as a DynamoDB table) to control who can access that resource and what actions they can perform on it. You can use IAM policy conditions to specify fine-grained access control for DynamoDB items and attributes. For example, you can allow or deny access to specific attributes of all items in a table by matching on attribute names¹. By creating a resource-based policy that allows access to only specific attributes of the DynamoDB table and attaching it to the table, the company can restrict access to confidential data. By creating an IAM role in the marketing team's account that has permissions to access the DynamoDB table in the finance team's account, the company can enable cross-account access. The other options are not correct because:

- Creating an SCP to grant the marketing team's AWS account access to the specific attributes of the DynamoDB table would not work because SCPs are policies that you can use with AWS Organizations to manage permissions in your organization's accounts. SCPs do not grant permissions; instead, they specify the maximum permissions that identities in an account can have². SCPs cannot be used to specify fine-grained access control for DynamoDB items and attributes.
- Creating an IAM role in the finance team's account by using IAM policy conditions for specific DynamoDB attributes and establishing trust with the marketing team's account would not work because IAM roles are identities that you can create in your account that have specific permissions. You can use an IAM role to delegate access to users, applications, or services that don't normally have access to your AWS resources³. However, creating an IAM role in the finance team's account would not restrict access to specific attributes of the DynamoDB table; it would only allow cross-account access. The company would still need a resource-based policy attached to the table to enforce fine-grained access control.
- Creating an IAM role in the finance team's account to access the DynamoDB table and using an IAM permissions boundary to limit the access to the specific attributes would not work because IAM permissions boundaries are policies that you use to delegate permissions management to other users. You can use permissions boundaries to limit the maximum permissions that an identity-based policy can grant to an IAM entity (user or role)⁴. Permissions boundaries cannot be used to specify fine-grained access control for DynamoDB items and attributes.

References:

- <https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/specifying-conditions.html>
- https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps.html
- https://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles.html
- https://docs.aws.amazon.com/IAM/latest/UserGuide/access_policies_boundaries.html

NEW QUESTION 108

- (Exam Topic 2)

A company wants to containerize a multi-tier web application and move the application from an on-premises data center to AWS. The application includes web, application, and database tiers. The company needs to make the application fault tolerant and scalable. Some frequently accessed data must always be available across application servers. Frontend web servers need session persistence and must scale to meet increases in traffic. Which solution will meet these requirements with the LEAST ongoing operational overhead?

- A. Run the application on Amazon Elastic Container Service (Amazon ECS) on AWS Fargat
- B. Use Amazon Elastic File System (Amazon EFS) for data that is frequently accessed between the web and application tier
- C. Store the frontend web server session data in Amazon Simple Queue Service (Amazon SOS).
- D. Run the application on Amazon Elastic Container Service (Amazon ECS) on Amazon EC2. Use Amazon ElastiCache for Redis to cache frontend web server session dat
- E. Use Amazon Elastic Block Store (Amazon EBS) with Multi-Attach on EC2 instances that are distributed across multiple Availability Zones.
- F. Run the application on Amazon Elastic Kubernetes Service (Amazon EKS). Configure Amazon EKS to use managed node group
- G. Use ReplicaSets to run the web servers and application
- H. Create an Amazon Elastic File System (Amazon EFS) Me syste
- I. Mount the EFS file system across all EKS pods to store frontend web server session data.
- J. Deploy the application on Amazon Elastic Kubernetes Service (Amazon EKS) Configure Amazon EKS to use managed node group
- K. Run the web servers and application as Kubernetes deployments in the EKS cluste
- L. Store the frontend web server session data in an Amazon DynamoDB tabl
- M. Create an Amazon Elastic File System (Amazon EFS) volume that all applications will mount at the time of deployment.

Answer: D

Explanation:

Deploying the application on Amazon EKS with managed node groups simplifies the operational overhead of managing the Kubernetes cluster. Running the web servers and application as Kubernetes deployments ensures that the desired number of pods are always running and can scale up or down as needed. Storing the frontend web server session data in an Amazon DynamoDB table provides a fast, scalable, and durable storage option that can be accessed across multiple Availability Zones. Creating an Amazon EFS volume that all applications will mount at the time of deployment allows the application to share data that is frequently accessed between the web and application tiers. References:

- <https://docs.aws.amazon.com/eks/latest/userguide/managed-node-groups.html>
- <https://docs.aws.amazon.com/eks/latest/userguide/deployments.html>
- <https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Introduction.html>
- <https://docs.aws.amazon.com/efs/latest/ug/mounting-fs.html>

NEW QUESTION 113

- (Exam Topic 2)

A solutions architect is designing a solution to process events. The solution must have the ability to scale in and out based on the number of events that the solution receives. If a processing error occurs, the event must move into a separate queue for review.

Which solution will meet these requirements?

- A. Send event details to an Amazon Simple Notification Service (Amazon SNS) topic
- B. Configure an AWS Lambda function as a subscriber to the SNS topic to process the event
- C. Add an on-failure destination to the function
- D. Set an Amazon Simple Queue Service (Amazon SQS) queue as the target.
- E. Publish events to an Amazon Simple Queue Service (Amazon SQS) queue
- F. Create an Amazon EC2 Auto Scaling group
- G. Configure the Auto Scaling group to scale in and out based on the ApproximateAgeOfOldestMessage metric of the queue
- H. Configure the application to write failed messages to a dead-letter queue.
- I. Write events to an Amazon DynamoDB table
- J. Configure a DynamoDB stream for the table
- K. Configure the stream to invoke an AWS Lambda function
- L. Configure the Lambda function to process the events.
- M. Publish events to an Amazon EventBridge event bus
- N. Create and run an application on an Amazon EC2 instance with an Auto Scaling group that is behind an Application Load Balancer (ALB). Set the ALB as the event bus target
- O. Configure the event bus to retry event
- P. Write messages to a dead-letter queue if the application cannot process the messages.

Answer: A

Explanation:

Amazon Simple Notification Service (Amazon SNS) is a fully managed pub/sub messaging service that enables users to send messages to multiple subscribers¹. Users can send event details to an Amazon SNS topic and configure an AWS Lambda function as a subscriber to the SNS topic to process the events. Lambda is a serverless compute service that runs code in response to events and automatically manages the underlying compute resources². Users can add an on-failure destination to the function and set an Amazon Simple Queue Service (Amazon SQS) queue as the target. Amazon SQS is a fully managed message queuing service that enables users to decouple and scale microservices, distributed systems, and serverless applications³. This way, if a processing error occurs, the event will move into the separate queue for review.

Option B is incorrect because publishing events to an Amazon SQS queue and creating an Amazon EC2 Auto Scaling group will not have the ability to scale in and out based on the number of events that the solution receives. Amazon EC2 is a web service that provides secure, resizable compute capacity in the cloud. Auto Scaling is a feature that helps users maintain application availability and allows them to scale their EC2 capacity up or down automatically according to conditions they define. However, for this use case, using SQS and EC2 will not take advantage of the serverless capabilities of Lambda and SNS.

Option C is incorrect because writing events to an Amazon DynamoDB table and configuring a DynamoDB stream for the table will not have the ability to move events into a separate queue for review if a processing error occurs. Amazon DynamoDB is a fully managed key-value and document database that delivers single-digit millisecond performance at any scale. DynamoDB Streams is a feature that captures data modification events in DynamoDB tables. Users can configure the stream to invoke a Lambda function, but they cannot configure an on-failure destination for the function.

Option D is incorrect because publishing events to an Amazon EventBridge event bus and setting an Application Load Balancer (ALB) as the event bus target will not have the ability to move events into a separate queue for review if a processing error occurs. Amazon EventBridge is a serverless event bus service that makes it easy to connect applications with data from a variety of sources. An ALB is a load balancer that distributes incoming application traffic across multiple targets, such as EC2 instances, containers, IP addresses, Lambda functions, and virtual appliances. Users can configure EventBridge to retry events, but they cannot configure an on-failure destination for the ALB.

NEW QUESTION 117

- (Exam Topic 2)

A company provides auction services for artwork and has users across North America and Europe. The company hosts its application in Amazon EC2 instances in the us-east-1 Region. Artists upload photos of their work as large-size, high-resolution image files from their mobile phones to a centralized Amazon S3 bucket created in the us-east-1 Region. The users in Europe are reporting slow performance for their image uploads.

How can a solutions architect improve the performance of the image upload process?

- A. Redeploy the application to use S3 multipart uploads.
- B. Create an Amazon CloudFront distribution and point to the application as a custom origin
- C. Configure the buckets to use S3 Transfer Acceleration.
- D. Create an Auto Scaling group for the EC2 instances and create a scaling policy.

Answer: C

Explanation:

Transfer acceleration. S3 Transfer Acceleration utilizes the Amazon CloudFront global network of edge locations to accelerate the transfer of data to and from S3 buckets. By enabling S3 Transfer Acceleration on the centralized S3 bucket, the users in Europe will experience faster uploads as their data will be routed through the closest CloudFront edge location.

NEW QUESTION 119

- (Exam Topic 2)

A company is running a compute workload by using Amazon EC2 Spot Instances that are in an Auto Scaling group. The launch template uses two placement groups and a single instance type.

Recently, a monitoring system reported Auto Scaling instance launch failures that correlated with longer wait times for system users. The company needs to improve the overall reliability of the workload.

Which solution will meet this requirement?

- A. Replace the launch template with a launch configuration to use an Auto Scaling group that uses attribute-based instance type selection.
- B. Create a new launch template version that uses attribute-based instance type selection
- C. Configure the Auto Scaling group to use the new launch template version.
- D. Update the launch template Auto Scaling group to increase the number of placement groups.
- E. Update the launch template to use a larger instance type.

Answer: B

Explanation:

<https://docs.aws.amazon.com/autoscaling/ec2/userguide/create-asg-instance-type-requirements.html#use-attribut>

NEW QUESTION 122

- (Exam Topic 2)

A company needs to audit the security posture of a newly acquired AWS account. The company's data security team requires a notification only when an Amazon S3 bucket becomes publicly exposed. The company has already established an Amazon Simple Notification Service (Amazon SNS) topic that has the data security team's email address subscribed.

Which solution will meet these requirements?

- A. Create an S3 event notification on all S3 buckets for the isPublic even
- B. Select the SNS topic as the target for the event notifications.
- C. Create an analyzer in AWS Identity and Access Management Access Analyze
- D. Create an Amazon EventBridge rule for the event type "Access Analyzer Finding" with a filter for "isPublic: true." Select the SNS topic as the EventBridge rule target.
- E. Create an Amazon EventBridge rule for the event type "Bucket-Level API Call via CloudTrail" with a filter for "PutBucketPolicy." Select the SNS topic as the EventBridge rule target.
- F. Activate AWS Config and add the cloudtrail-s3-dataevents-enabled rul
- G. Create an Amazon EventBridge rule for the event type "Config Rules Re-evaluation Status" with a filter for "NON_COMPLIANT." Select the SNS topic as the EventBridge rule target.

Answer: B

Explanation:

Access Analyzer is to assess the access policy. https://docs.aws.amazon.com/ja_jp/AmazonS3/latest/userguide/access-control-block-public-access.html

NEW QUESTION 126

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questions and Answers in PDF Format

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

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

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