

SAP-C02 Dumps

AWS Certified Solutions Architect - Professional

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



NEW QUESTION 1

- (Exam Topic 1)

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

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

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

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

Answer: C

Explanation:

<https://docs.aws.amazon.com/fsx/latest/WindowsGuide/what-is.html> https://docs.aws.amazon.com/directoryservice/latest/admin-guide/ms_ad_join_instance.html

NEW QUESTION 2

- (Exam Topic 1)

A company runs an IoT platform on AWS IoT sensors in various locations send data to the company's Node.js API servers on Amazon EC2 instances running behind an Application Load Balancer. The data is stored in an Amazon RDS MySQL DB instance that uses a 4 TB General Purpose SSD volume.

The number of sensors the company has deployed in the field has increased over time and is expected to grow significantly. The API servers are consistently overloaded and RDS metrics show high write latency.

Which of the following steps together will resolve the issues permanently and enable growth as new sensors are provisioned, while keeping this platform cost-efficient? (Select TWO.)

- A. Resize the MySQL General Purpose SSD storage to 6 TB to improve the volume's IOPS
- B. Re-architect the database tier to use Amazon Aurora instead of an RDS MySQL DB instance and add read replicas
- C. Leverage Amazon Kinesis Data Streams and AWS Lambda to ingest and process the raw data
- D. Use AWS X-Ray to analyze and debug application issues and add more API servers to match the load
- E. Re-architect the database tier to use Amazon DynamoDB instead of an RDS MySQL DB instance

Answer: CE

Explanation:

➤ Option C is correct because leveraging Amazon Kinesis Data Streams and AWS Lambda to ingest and process the raw data resolves the issues permanently and enables growth as new sensors are provisioned. Amazon Kinesis Data Streams is a serverless streaming data service that simplifies the capture, processing, and storage of data streams at any scale. Kinesis Data Streams can handle any amount of streaming data and process data from hundreds of thousands of sources with very low latency. AWS Lambda is a serverless compute service that lets you run code without provisioning or managing servers. Lambda can be triggered by Kinesis Data Streams events and process the data records in real time. Lambda can also scale automatically based on the incoming data volume. By using Kinesis Data Streams and Lambda, the company can reduce the load on the API servers and improve the performance and scalability of the data ingestion and processing layer.

➤ Option E is correct because re-architecting the database tier to use Amazon DynamoDB instead of an RDS MySQL DB instance resolves the issues permanently and enables growth as new sensors are provisioned. Amazon DynamoDB is a fully managed key-value and document database that delivers single-digit millisecond performance at any scale. DynamoDB supports auto scaling, which automatically adjusts read and write capacity based on actual traffic patterns. DynamoDB also supports on-demand capacity mode, which instantly accommodates up to double the previous peak traffic on a table. By using DynamoDB instead of RDS MySQL DB instance, the company can eliminate high write latency and improve scalability and performance of the database tier.

References: 1: <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-volume-types.html> 2:

https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/CHAP_AuroraOverview.html 3:

<https://docs.aws.amazon.com/streams/latest/dev/introduction.html> : <https://docs.aws.amazon.com/lambda/latest/dg/welcome.html> :

<https://docs.aws.amazon.com/xray/latest/devguide/aws-xray.html> : <https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Introduction.html> :

NEW QUESTION 3

- (Exam Topic 1)

A company has 10 accounts that are part of an organization in AWS Organizations. AWS Config is configured in each account. All accounts belong to either the Prod OU or the NonProd OU.

The company has set up an Amazon EventBridge rule in each AWS account to notify an Amazon Simple Notification Service (Amazon SNS) topic when an Amazon EC2 security group inbound rule is created with 0.0.0.0/0 as the source. The company's security team is subscribed to the SNS topic.

For all accounts in the NonProd OU, the security team needs to remove the ability to create a security group inbound rule that includes 0.0.0.0/0 as the source. Which solution will meet this requirement with the LEAST operational overhead?

- A. Modify the EventBridge rule to invoke an AWS Lambda function to remove the security group inbound rule and to publish to the SNS topic. Deploy the updated rule to the NonProd OU.
- B. Add the vpc-sg-open-only-to-authorized-ports AWS Config managed rule to the NonProd OU.

- C. Configure an SCP to allow the ec2 AuthorizeSecurityGroupIngress action when the value of the aws:SourceIp condition key is not 0.0.0.0/0 Apply the SCP to the NonProd OU
- D. Configure an SCP to deny the ec2 AuthorizeSecurityGroupIngress action when the value of the aws:SourceIp condition key is 0.0.0.0/0 Apply the SCP to the NonProd OU

Answer: D

Explanation:

This solution will meet the requirement with the least operational overhead because it directly denies the creation of the security group inbound rule with 0.0.0.0/0 as the source, which is the exact requirement. Additionally, it does not require any additional steps or resources such as invoking a Lambda function or adding a Config rule.

An SCP (Service Control Policy) is a policy that you can use to set fine-grained permissions for your AWS accounts within your organization. You can use SCPs to set permissions for the root user of an account and to delegate permissions to IAM users and roles in the accounts. You can use SCPs to set permissions that allow or deny access to specific services, actions, and resources.

To implement this solution, you would need to create an SCP that denies the ec2:AuthorizeSecurityGroupIngress action when the value of the aws:SourceIp condition key is 0.0.0.0/0. This SCP would then be applied to the NonProd OU. This would ensure that any security group inbound rule that includes 0.0.0.0/0 as the source will be denied, thus meeting the requirement.

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

https://docs.aws.amazon.com/IAM/latest/UserGuide/access_policies_condition-keys.html

NEW QUESTION 4

- (Exam Topic 1)

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

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

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

Answer: AC

Explanation:

<https://aws.amazon.com/blogs/mt/self-service-vpcs-in-aws-control-tower-using-aws-service-catalog/> <https://docs.aws.amazon.com/vpc/latest/tgw/tgw-transit-gateways.html>

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-resource-ec2-transitgatewayattachme>

NEW QUESTION 5

- (Exam Topic 1)

A publishing company's design team updates the icons and other static assets that an e-commerce web application uses. The company serves the icons and assets from an Amazon S3 bucket that is hosted in the company's production account. The company also uses a development account that members of the design team can access.

After the design team tests the static assets in the development account, the design team needs to load the assets into the S3 bucket in the production account. A solutions architect must provide the design team with access to the production account without exposing other parts of the web application to the risk of unwanted changes.

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

- A. In the production account, create a new IAM policy that allows read and write access to the S3 bucket.
- B. In the development account, create a new IAM policy that allows read and write access to the S3 bucket.
- C. In the production account, create a role
- D. Attach the new policy to the role
- E. Define the development account as a trusted entity.
- F. In the development account, create a role
- G. Attach the new policy to the role
- H. Define the production account as a trusted entity.
- I. In the development account, create a group that contains all the IAM users of the design team
- J. Attach a different IAM policy to the group to allow the sts:AssumeRole action on the role in the production account.
- K. In the development account, create a group that contains all the IAM users of the design team
- L. Attach a different IAM policy to the group to allow the sts:AssumeRole action on the role in the development account.

Answer: ACE

Explanation:

> A. In the production account, create a new IAM policy that allows read and write access to the S3 bucket. The policy grants the necessary permissions to access the assets in the production S3 bucket.

> C. In the production account, create a role. Attach the new policy to the role. Define the development account as a trusted entity. By creating a role and attaching the policy, and then defining the development account as a trusted entity, the development account can assume the role and access the production S3 bucket with the read and write permissions.

> E. In the development account, create a group that contains all the IAM users of the design team. Attach a different IAM policy to the group to allow the sts:AssumeRole action on the role in the production account. The IAM policy attached to the group allows the design team members to assume the role created in the production account, thereby giving them access to the production S3 bucket.

Step 1: Create a role in the Production Account; create the role in the Production account and specify the Development account as a trusted entity. You also limit the role permissions to only read and write access to the productionapp bucket. Anyone granted permission to use the role can read and write to the productionapp

bucket. Step 2: Grant access to the role Sign in as an administrator in the Development account and allow the AssumeRole action on the UpdateApp role in the Production account. So, recap, production account you create the policy for S3, and you set development account as a trusted entity. Then on the development account you allow the sts:assumeRole action on the role in production account. https://docs.aws.amazon.com/IAM/latest/UserGuide/tutorial_cross-account-with-roles.html

NEW QUESTION 6

- (Exam Topic 1)

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

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

Answer: C

Explanation:

This solution requires the least amount of effort as it only requires to update the API endpoint to private in API Gateway and create an interface VPC endpoint. Then create a resource policy and attach it to the API. This will make the API only accessible from the VPC and still keep the authentication mechanism intact.

Reference:

> <https://aws.amazon.com/api-gateway/features/>

NEW QUESTION 7

- (Exam Topic 1)

A company has created an OU in AWS Organizations for each of its engineering teams. Each OU owns multiple AWS accounts. The organization has hundreds of AWS accounts. A solutions architect must design a solution so that each OU can view a breakdown of usage costs across its AWS accounts. Which solution meets these requirements?

- A. Create an AWS Cost and Usage Report (CUR) for each OU by using AWS Resource Access Manager. Allow each team to visualize the CUR through an Amazon QuickSight dashboard.
- B. Create an AWS Cost and Usage Report (CUR) from the AWS Organizations management account. Allow each team to visualize the CUR through an Amazon QuickSight dashboard.
- C. Create an AWS Cost and Usage Report (CUR) in each AWS Organizations member account. Allow each team to visualize the CUR through an Amazon QuickSight dashboard.
- D. Create an AWS Cost and Usage Report (CUR) by using AWS Systems Manager. Allow each team to visualize the CUR through Systems Manager OpsCenter dashboards.

Answer: B

Explanation:

<https://docs.aws.amazon.com/cur/latest/userguide/billing-cur-limits.html>

NEW QUESTION 8

- (Exam Topic 1)

A company wants to migrate an application to Amazon EC2 from VMware Infrastructure that runs in an on-premises data center. A solutions architect must preserve the software and configuration settings during the migration. What should the solutions architect do to meet these requirements?

- A. Configure the AWS DataSync agent to start replicating the data store to Amazon FSx for Windows FileServer. Use the SMB share to host the VMware data store.
- B. Use VM Import/Export to move the VMs to Amazon EC2.
- C. Use the VMware vSphere client to export the application as an image in Open Virtualization Format (OVF) format. Create an Amazon S3 bucket to store the image in the destination AWS Region.
- D. Create and apply an IAM role for VM Import. Use the AWS CLI to run the EC2 import command.
- E. . Configure AWS Storage Gateway for files service to export a Common Internet File System (CIFS) share.
- F. Create a backup copy to the shared folder.
- G. Sign in to the AWS Management Console and create an AMI from the backup copy. Launch an EC2 instance that is based on the AMI.
- H. Create a managed-instance activation for a hybrid environment in AWS Systems Manager.
- I. Download and install Systems Manager Agent on the on-premises VM. Register the VM with Systems Manager to be a managed instance. Use AWS Backup to create a snapshot of the VM and create an AMI.
- J. Launch an EC2 instance that is based on the AMI.

Answer: D

Explanation:

<https://docs.aws.amazon.com/vm-import/latest/userguide/vmimport-image-import.html>

- Export an OVF Template

- Create / use an Amazon S3 bucket for storing the exported images. The bucket must be in the Region where you want to import your VMs.
- Create an IAM role named vmimport.
- You'll use AWS CLI to run the import commands. <https://aws.amazon.com/premiumsupport/knowledge-center/import-instances/>

NEW QUESTION 9

- (Exam Topic 1)

A company developed a pilot application by using AWS Elastic Beanstalk and Java. To save costs during development, the company's development team deployed the application into a single-instance environment. Recent tests indicate that the application consumes more CPU than expected. CPU utilization is regularly greater than 85%, which causes some performance bottlenecks.

A solutions architect must mitigate the performance issues before the company launches the application to production.

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

- A. Create a new Elastic Beanstalk applicatio
- B. Select a load-balanced environment typ
- C. Select all Availability Zone
- D. Add a scale-out rule that will run if the maximum CPU utilization is over 85% for 5 minutes.
- E. Create a second Elastic Beanstalk environmen
- F. Apply the traffic-splitting deployment polic
- G. Specify a percentage of incoming traffic to direct to the new environment in the average CPU utilization is over 85% for 5 minutes.
- H. Modify the existing environment's capacity configuration to use a load-balanced environment type. Select all Availability Zone
- I. Add a scale-out rule that will run if the average CPU utilization is over 85% for 5 minutes.
- J. Select the Rebuild environment action with the load balancing option Select an Availability Zones Add a scale-out rule that will run if the sum CPU utilization is over 85% for 5 minutes.

Answer: C

Explanation:

This solution will meet the requirements with the least operational overhead because it allows the company to modify the existing environment's capacity configuration, so it becomes a load-balanced environment type. By selecting all availability zones, the company can ensure that the application is running in multiple availability zones, which can help to improve the availability and scalability of the application. The company can also add a scale-out rule that will run if the average CPU utilization is over 85% for 5 minutes, which can help to mitigate the performance issues. This solution does not require creating new Elastic Beanstalk environments or rebuilding the existing one, which reduces the operational overhead.

You can refer to the AWS Elastic Beanstalk documentation for more information on how to use this service: <https://aws.amazon.com/elasticbeanstalk/> You can refer to the AWS documentation for more information on how to use autoscaling: <https://aws.amazon.com/autoscaling/>

NEW QUESTION 10

- (Exam Topic 1)

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

The development team and testing team work on business days during business hours, but the production environment operates 24 hours a day. 7 days a week.

The company wants to reduce costs AH resources are tagged with an environment tag with either development, testing, or production as the key. What should a solutions architect do to reduce costs with the LEAST operational effort?

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

Answer: B

Explanation:

Creating an Amazon EventBridge rule that runs every business day in the evening to stop instances and another rule that runs every business day in the morning to start instances based on the tag will reduce costs with the least operational effort. This approach allows for instances to be stopped during non-business hours when they are not in use, reducing the costs associated with running them. It also allows for instances to be started again in the morning when the development and testing teams need to use them.

NEW QUESTION 10

- (Exam Topic 1)

A company hosts a Git repository in an on-premises data center. The company uses webhooks to invoke functionality that runs in the AWS Cloud. The company hosts the webhook logic on a set of Amazon EC2 instances in an Auto Scaling group that the company set as a target for an Application Load Balancer (ALB). The Git server calls the ALB for the configured webhooks. The company wants to move the solution to a serverless architecture.

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

- A. For each webhook, create and configure an AWS Lambda function UR
- B. Update the Git servers to call the individual Lambda function URLs.
- C. Create an Amazon API Gateway HTTP AP
- D. Implement each webhook logic in a separate AWS Lambda functio
- E. Update the Git servers to call the API Gateway endpoint.
- F. Deploy the webhook logic to AWS App Runne
- G. Create an ALB, and set App Runner as the target. Update the Git servers to call the ALB endpoint.
- H. Containerize the webhook logi
- I. Create an Amazon Elastic Container Service (Amazon ECS) cluster, and run the webhook logic in AWS Fargat

- J. Create an Amazon API Gateway REST API, and set Fargate as the target.
K. Update the Git servers to call the API Gateway endpoint.

Answer: B

Explanation:

<https://aws.amazon.com/solutions/implementations/git-to-s3-using-webhooks/> <https://medium.com/mindorks/building-webhook-is-easy-using-aws-lambda-and-api-gateway-56f5e5c3a596>

NEW QUESTION 14

- (Exam Topic 1)

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

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

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

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

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

Answer: CE

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/private-hosted-zone-different-account/>

NEW QUESTION 18

- (Exam Topic 1)

A company with global offices has a single 1 Gbps AWS Direct Connect connection to a single AWS Region. The company's on-premises network uses the connection to communicate with the company's resources in the AWS Cloud. The connection has a single private virtual interface that connects to a single VPC.

A solutions architect must implement a solution that adds a redundant Direct Connect connection in the same Region. The solution also must provide connectivity to other Regions through the same pair of Direct Connect connections as the company expands into other Regions.

Which solution meets these requirements?

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

Answer: A

Explanation:

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

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

NEW QUESTION 23

- (Exam Topic 1)

A company has a multi-tier web application that runs on a fleet of Amazon EC2 instances behind an Application Load Balancer (ALB). The instances are in an Auto Scaling group. The ALB and the Auto Scaling group are replicated in a backup AWS Region. The minimum value and the maximum value for the Auto Scaling group are set to zero. An Amazon RDS Multi-AZ DB instance stores the application's data. The DB instance has a read replica in the backup Region. The application presents an endpoint to end users by using an Amazon Route 53 record.

The company needs to reduce its RTO to less than 15 minutes by giving the application the ability to automatically fail over to the backup Region. The company does not have a large enough budget for an active-active strategy.

What should a solutions architect recommend to meet these requirements?

- A. Reconfigure the application's Route 53 record with a latency-based routing policy that load balances traffic between the two ALBs.
- B. Create an AWS Lambda function in the backup Region to promote the read replica and modify the Auto Scaling group value.
- C. Create an Amazon CloudWatch alarm that is based on the HTTPCode_Target_5XX_Count metric for the ALB in the primary Region.
- D. Configure the CloudWatch alarm to invoke the Lambda function.
- E. Create an AWS Lambda function in the backup Region to promote the read replica and modify the Auto Scaling group value.
- F. Configure Route 53 with a health check that monitors the web application and sends an Amazon Simple Notification Service (Amazon SNS) notification to the Lambda function when the health check status is unhealthy.
- G. Update the application's Route 53 record with a failover policy that routes traffic to the ALB in the backup Region when a health check failure occurs.

- H. Configure the Auto Scaling group in the backup Region to have the same values as the Auto Scaling group in the primary Region
- I. Reconfigure the application's Route 53 record with a latency-based routing policy that load balances traffic between the two ALB
- J. Remove the read replica
- K. Replace the read replica with a standalone RDS DB instance
- L. Configure Cross-Region Replication between the RDS DB instances by using snapshots and Amazon S3.
- M. Configure an endpoint in AWS Global Accelerator with the two ALBs as equal weighted target
- N. Create an AWS Lambda function in the backup Region to promote the read replica and modify the Auto Scaling group value
- O. Create an Amazon CloudWatch alarm that is based on the HTTPCode_Target_5XX_Count metric for the ALB in the primary Region
- P. Configure the CloudWatch alarm to invoke the Lambda function.

Answer: B

Explanation:

an AWS Lambda function in the backup region to promote the read replica and modify the Auto Scaling group values, and then configuring Route 53 with a health check that monitors the web application and sends an Amazon SNS notification to the Lambda function when the health check status is unhealthy. Finally, the application's Route 53 record should be updated with a failover policy that routes traffic to the ALB in the backup region when a health check failure occurs. This approach provides automatic failover to the backup region when a health check failure occurs, reducing the RTO to less than 15 minutes. Additionally, this approach is cost-effective as it does not require an active-active strategy.

NEW QUESTION 27

- (Exam Topic 1)

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

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

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

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

Answer: B

Explanation:

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

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

NEW QUESTION 32

- (Exam Topic 1)

A video processing company wants to build a machine learning (ML) model by using 600 TB of compressed data that is stored as thousands of files in the company's on-premises network attached storage system. The company does not have the necessary compute resources on premises for ML experiments and wants to use AWS.

The company needs to complete the data transfer to AWS within 3 weeks. The data transfer will be a one-time transfer. The data must be encrypted in transit. The measured upload speed of the company's internet connection is 100 Mbps, and multiple departments share the connection.

Which solution will meet these requirements MOST cost-effectively?

- A. Order several AWS Snowball Edge Storage Optimized devices by using the AWS Management Console
- B. Configure the devices with a destination S3 bucket
- C. Copy the data to the device
- D. Ship the devices back to AWS.
- E. Set up a 10 Gbps AWS Direct Connect connection between the company location and the nearest AWS Region
- F. Transfer the data over a VPN connection into the Region to store the data in Amazon S3.
- G. Create a VPN connection between the on-premises network storage and the nearest AWS Region. Transfer the data over the VPN connection.
- H. Deploy an AWS Storage Gateway file gateway on premises
- I. Configure the file gateway with a destination S3 bucket
- J. Copy the data to the file gateway.

Answer: A

Explanation:

This solution will meet the requirements of the company as it provides a secure, cost-effective and fast way of transferring large data sets from on-premises to AWS. Snowball Edge devices encrypt the data during transfer, and the devices are shipped back to AWS for import into S3. This option is more cost effective than using Direct Connect or VPN connections as it does not require the company to pay for long-term dedicated connections.

NEW QUESTION 34

- (Exam Topic 1)

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

Which solution will meet these requirements?

- A. Deploy and configure the AWS Agentless Discovery Connector virtual appliance on the on-premises host
- B. Configure Data Exploration in AWS Migration Hub

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

Answer: D

Explanation:

➤ it covers all the requirements mentioned in the question, it will allow collecting the detailed metrics, including process information and it provides a way to query and analyze the data using Amazon Athena.

NEW QUESTION 37

- (Exam Topic 1)

A company's solutions architect is reviewing a web application that runs on AWS. The application references static assets in an Amazon S3 bucket in the us-east-1 Region. The company needs resiliency across multiple AWS Regions. The company already has created an S3 bucket in a second Region. Which solution will meet these requirements with the LEAST operational overhead?

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

Answer: C

Explanation:

https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/high_availability_origin_failover.html

NEW QUESTION 42

- (Exam Topic 1)

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

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

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

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

Answer: AD

Explanation:

Option A, Deploy the security tool on EC2 instances in a new Auto Scaling group in the existing VPC, allows the company to use its existing security tool while still running it within the AWS environment. This ensures that all packets coming in and out of the VPC are inspected by the security tool in real time. Option D, Provision a Gateway Load Balancer for each Availability Zone to redirect the traffic to the security tool, allows for high availability within an AWS Region. By provisioning a Gateway Load Balancer for each Availability Zone, the traffic is redirected to the security tool in the event of any failures or outages. This ensures that the security tool is always available to inspect the traffic, even in the event of a failure.

NEW QUESTION 44

- (Exam Topic 1)

A company is building an electronic document management system in which users upload their documents. The application stack is entirely serverless and runs on AWS in the eu-central-1 Region. The system includes a web application that uses an Amazon CloudFront distribution for delivery with Amazon S3 as the origin. The web application communicates with Amazon API Gateway Regional endpoints. The API Gateway APIs call AWS Lambda functions that store metadata in an Amazon Aurora Serverless database and put the documents into an S3 bucket.

The company is growing steadily and has completed a proof of concept with its largest customer. The company must improve latency outside of Europe.

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

- A. Enable S3 Transfer Acceleration on the S3 bucket
- B. Ensure that the web application uses the Transfer Acceleration signed URLs.
- C. Create an accelerator in AWS Global Accelerator
- D. Attach the accelerator to the CloudFront distribution.

- E. Change the API Gateway Regional endpoints to edge-optimized endpoints.
- F. Provision the entire stack in two other locations that are spread across the world.
- G. Use global databases on the Aurora Serverless cluster.
- H. Add an Amazon RDS proxy between the Lambda functions and the Aurora Serverless database.

Answer: AC

Explanation:

<https://aws.amazon.com/global-accelerator/faqs/>

NEW QUESTION 47

- (Exam Topic 1)

A company is running several workloads in a single AWS account. A new company policy states that engineers can provision only approved resources and that engineers must use AWS CloudFormation to provision these resources. A solutions architect needs to create a solution to enforce the new restriction on the IAM role that the engineers use for access.

What should the solutions architect do to create the solution?

- A. Upload AWS CloudFormation templates that contain approved resources to an Amazon S3 bucket. Update the IAM policy for the engineers' IAM role to only allow access to Amazon S3 and AWS CloudFormation.
- B. Use AWS CloudFormation templates to provision resources.
- C. Update the IAM policy for the engineers' IAM role with permissions to only allow provisioning of approved resources and AWS CloudFormation.
- D. Use AWS CloudFormation templates to create stacks with approved resources.
- E. Update the IAM policy for the engineers' IAM role with permissions to only allow AWS CloudFormation action.
- F. Create a new IAM policy with permission to provision approved resources, and assign the policy to a new IAM service role.
- G. Assign the IAM service role to AWS CloudFormation during stack creation.
- H. Provision resources in AWS CloudFormation stack.
- I. Update the IAM policy for the engineers' IAM role to only allow access to their own AWS CloudFormation stack.

Answer: B

Explanation:

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/security-best-practices.html#use-iam-to-c>

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/using-iam-servicerole.html>

NEW QUESTION 49

- (Exam Topic 1)

A digital marketing company has multiple AWS accounts that belong to various teams. The creative team uses an Amazon S3 bucket in its AWS account to securely store images and media files that are used as content for the company's marketing campaigns. The creative team wants to share the S3 bucket with the strategy team so that the strategy team can view the objects.

A solutions architect has created an IAM role that is named `strategy_reviewer` in the Strategy account. The solutions architect also has set up a custom AWS Key Management Service (AWS KMS) key in the Creative account and has associated the key with the S3 bucket. However, when users from the Strategy account assume the IAM role and try to access objects in the S3 bucket, they receive an `AccessDenied` error.

The solutions architect must ensure that users in the Strategy account can access the S3 bucket. The solution must provide these users with only the minimum permissions that they need.

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

- A. Create a bucket policy that includes read permissions for the S3 bucket.
- B. Set the principal of the bucket policy to the account ID of the Strategy account.
- C. Update the `strategy_reviewer` IAM role to grant full permissions for the S3 bucket and to grant decrypt permissions for the custom KMS key.
- D. Update the custom KMS key policy in the Creative account to grant decrypt permissions to the `strategy_reviewer` IAM role.
- E. Create a bucket policy that includes read permissions for the S3 bucket.
- F. Set the principal of the bucket policy to an anonymous user.
- G. Update the custom KMS key policy in the Creative account to grant encrypt permissions to the `strategy_reviewer` IAM role.
- H. Update the `strategy_reviewer` IAM role to grant read permissions for the S3 bucket and to grant decrypt permissions for the custom KMS key.

Answer: ACF

Explanation:

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

NEW QUESTION 51

- (Exam Topic 1)

A company runs a serverless application in a single AWS Region. The application accesses external URLs and extracts metadata from those sites. The company uses an Amazon Simple Notification Service (Amazon SNS) topic to publish URLs to an Amazon Simple Queue Service (Amazon SQS) queue. An AWS Lambda function uses the queue as an event source and processes the URLs from the queue. Results are saved to an Amazon S3 bucket.

The company wants to process each URL in other Regions to compare possible differences in site localization. URLs must be published from the existing Region. Results must be written to the existing S3 bucket in the current Region.

Which combination of changes will produce multi-Region deployment that meets these requirements? (Select TWO.)

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

Answer: AC

Explanation:

<https://docs.aws.amazon.com/sns/latest/dg/sns-cross-region-delivery.html>

NEW QUESTION 55

- (Exam Topic 1)

An adventure company has launched a new feature on its mobile app. Users can use the feature to upload their hiking and rapping photos and videos anytime. The photos and videos are stored in Amazon S3 Standard storage in an S3 bucket and are served through Amazon CloudFront.

The company needs to optimize the cost of the storage. A solutions architect discovers that most of the uploaded photos and videos are accessed infrequently after 30 days. However, some of the uploaded photos and videos are accessed frequently after 30 days. The solutions architect needs to implement a solution that maintains millisecond retrieval availability of the photos and videos at the lowest possible cost.

Which solution will meet these requirements?

- A. Configure S3 Intelligent-Tiering on the S3 bucket.
- B. Configure an S3 Lifecycle policy to transition image objects and video objects from S3 Standard to S3 Glacier Deep Archive after 30 days.
- C. Replace Amazon S3 with an Amazon Elastic File System (Amazon EFS) file system that is mounted on Amazon EC2 instances.
- D. Add a Cache-Control: max-age header to the S3 image objects and S3 video object
- E. Set the header to 30 days.

Answer: A

Explanation:

Amazon S3 Intelligent-Tiering is a storage class that automatically moves objects between two access tiers based on changing access patterns. Objects that are accessed frequently are stored in the frequent access tier and objects that are accessed infrequently are stored in the infrequent access tier. This allows for cost optimization without requiring manual intervention. This makes it an ideal solution for the scenario described, as it can automatically move objects that are infrequently accessed after 30 days to a lower-cost storage tier while still maintaining millisecond retrieval availability.

NEW QUESTION 56

- (Exam Topic 1)

A company has an on-premises website application that provides real estate information for potential renters and buyers. The website uses a Java backend and a NOSQL MongoDB database to store subscriber data.

The company needs to migrate the entire application to AWS with a similar structure. The application must be deployed for high availability, and the company cannot make changes to the application

Which solution will meet these requirements?

- A. use an Amazon Aurora DB cluster as the database for the subscriber data
- B. Deploy Amazon EC2 instances in an Auto Scaling group across multiple Availability Zones for the Java backend application.
- C. Use MongoDB on Amazon EC2 instances as the database for the subscriber data
- D. Deploy EC2 instances in an Auto Scaling group in a single Availability Zone for the Java backend application.
- E. Configure Amazon DocumentDB (with MongoDB compatibility) with appropriately sized instances in multiple Availability Zones as the database for the subscriber data
- F. Deploy Amazon EC2 instances in an Auto Scaling group across multiple Availability Zones for the Java backend application.
- G. Configure Amazon DocumentDB (with MongoDB compatibility) in on-demand capacity mode in multiple Availability Zones as the database for the subscriber data
- H. Deploy Amazon EC2 instances in an Auto Scaling group across multiple Availability Zones for the Java backend application.

Answer: C

Explanation:

On-demand capacity mode is the function of Dynamodb.

<https://aws.amazon.com/blogs/news/running-spiky-workloads-and-optimizing-costs-by-more-than-90-using-ama>

Amazon DocumentDB Elastic Clusters <https://aws.amazon.com/blogs/news/announcing-amazon-documentdb-elastic-clusters/>

Deploy Amazon EC2 instances in an Auto Scaling group across multiple Availability Zones for the Java backend application. This will provide high availability and scalability, while allowing the company to retain the same database structure as the original application.

NEW QUESTION 61

- (Exam Topic 1)

A health insurance company stores personally identifiable information (PII) in an Amazon S3 bucket. The company uses server-side encryption with S3 managed encryption keys (SSE-S3) to encrypt the objects. According to a new requirement, all current and future objects in the S3 bucket must be encrypted by keys that the company's security team manages. The S3 bucket does not have versioning enabled. Which solution will meet these requirements?

- A. In the S3 bucket properties, change the default encryption to SSE-S3 with a customer managed key
- B. Use the AWS CLI to re-upload all objects in the S3 bucket
- C. Set an S3 bucket policy to deny unencrypted PutObject requests.
- D. In the S3 bucket properties, change the default encryption to server-side encryption with AWS KMS managed encryption keys (SSE-KMS). Set an S3 bucket policy to deny unencrypted PutObject request
- E. Use the AWS CLI to re-upload all objects in the S3 bucket.
- F. In the S3 bucket properties, change the default encryption to server-side encryption with AWS KMS managed encryption keys (SSE-KMS). Set an S3 bucket policy to automatically encrypt objects on GetObject and PutObject requests.
- G. In the S3 bucket properties, change the default encryption to AES-256 with a customer managed key. Attach a policy to deny unencrypted PutObject requests to any entities that access the S3 bucket
- H. Use the AWS CLI to re-upload all objects in the S3 bucket.

Answer: D

Explanation:

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/ServerSideEncryptionCustomerKeys.html> Clearly says we need following header for SSE-C x-amz-server-side-encryption-customer-algorithm Use this header to specify the encryption algorithm. The header value must be AES256.

NEW QUESTION 62

- (Exam Topic 1)

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

A solutions architect needs to decouple the third-party service calls and ensure that all the calls are eventually completed. Which solution will meet these requirements?

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

Answer: A

Explanation:

Using an SQS queue to store events and invoke the Lambda function will decouple the third-party service calls and ensure that all the calls are eventually completed. SQS allows you to store messages in a queue and process them asynchronously, which eliminates the need for the application to wait for a response from the third-party service. The messages will be stored in the SQS queue until they are processed by the Lambda function, even if the Lambda function is currently unavailable or busy. This will ensure that all the calls are eventually completed, even if there are delays or errors.

AWS Step Functions state machines can also be used to pass events to the Lambda function, but it would require additional management and configuration to set up the state machine, which would increase operational overhead.

Amazon EventBridge rule can also be used to pass events to the Lambda function, but it would not provide the same level of decoupling and reliability as SQS.

Using Amazon Simple Notification Service (Amazon SNS) topic to store events and Invoke the Lambda function, is similar to SQS, but SNS is a publish-subscribe messaging service and SQS is a queue service. SNS is used for sending messages to multiple recipients, SQS is used for sending messages to a single recipient, so SQS is more appropriate for this use case.

References:

- > AWS SQS
- > AWS Step Functions
- > AWS EventBridge
- > AWS SNS

NEW QUESTION 65

- (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 67

- (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 selectio
- 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 70

- (Exam Topic 2)

A company has five development teams that have each created five AWS accounts to develop and host applications. To track spending, the development teams log in to each account every month, record the current cost from the AWS Billing and Cost Management console, and provide the information to the company's finance team.

The company has strict compliance requirements and needs to ensure that resources are created only in AWS Regions in the United States. However, some resources have been created in other Regions.

A solutions architect needs to implement a solution that gives the finance team the ability to track and consolidate expenditures for all the accounts. The solution also must ensure that the company can create resources only in Regions in the United States.

Which combination of steps will meet these requirements in the MOST operationally efficient way? (Select THREE.)

- A. Create a new account to serve as a management account
- B. Create an Amazon S3 bucket for the finance team. Use AWS Cost and Usage Reports to create monthly reports and to store the data in the finance team's S3 bucket.
- C. Create a new account to serve as a management account
- D. Deploy an organization in AWS Organizations with all features enable
- E. Invite all the existing accounts to the organization
- F. Ensure that each account accepts the invitation.
- G. Create an OU that includes all the development team
- H. Create an SCP that allows the creation of resources only in Regions that are in the United States
- I. Apply the SCP to the OU.
- J. Create an OU that includes all the development team
- K. Create an SCP that denies the creation of resources in Regions that are outside the United States
- L. Apply the SCP to the OU.
- M. Create an IAM role in the management account. Attach a policy that includes permissions to view the Billing and Cost Management console
- N. Allow the finance team users to assume the role
- O. Use AWS Cost Explorer and the Billing and Cost Management console to analyze cost.
- P. Create an IAM role in each AWS account
- Q. Attach a policy that includes permissions to view the Billing and Cost Management console
- R. Allow the finance team users to assume the role.

Answer: BCE

Explanation:

AWS Organizations is a service that enables you to consolidate multiple AWS accounts into an organization that you create and centrally manage. By creating a management account and inviting all the existing accounts to join the organization, the solutions architect can track and consolidate expenditures for all the accounts using AWS Cost Management tools such as AWS Cost Explorer and AWS Budgets. An organizational unit (OU) is a group of accounts within an organization that can be used to apply policies and simplify management. A service control policy (SCP) is a type of policy that you can use to manage permissions in your organization. By creating an OU that includes all the development teams and applying an SCP that allows the creation of resources only in Regions that are in the United States, the solutions architect can ensure that the company meets its compliance requirements and avoids unwanted charges from other Regions. An IAM role is an identity with permission policies that determine what the identity can and cannot do in AWS. By creating an IAM role in the management account and allowing the finance team users to assume it, the solutions architect can give them access to view the Billing and Cost Management console without sharing credentials or creating additional users. References:

- > https://docs.aws.amazon.com/organizations/latest/userguide/orgs_introduction.html
- > https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scp.html
- > https://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles.html
- > <https://docs.aws.amazon.com/aws-cost-management/latest/userguide/what-is-costmanagement.html>

NEW QUESTION 71

- (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 table
- 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 control). Establish trust with the marketing team's account
- D. In the marketing team's account, create an IAM role that has permissions to assume the IAM role in the finance team's account.
- E. Create a resource-based IAM policy that includes conditions for specific DynamoDB attributes (fine-grained access control). Attach the policy to the DynamoDB table
- F. In the marketing team's account, 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 DynamoDB table
- 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 73

- (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 77

- (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 81

- (Exam Topic 2)

A company has millions of objects in an Amazon S3 bucket. The objects are in the S3 Standard storage class. All the S3 objects are accessed frequently. The number of users and applications that access the objects is increasing rapidly. The objects are encrypted with server-side encryption with AWS KMS Keys (SSE-KMS).

A solutions architect reviews the company's monthly AWS invoice and notices that AWS KMS costs are increasing because of the high number of requests from Amazon S3. The solutions architect needs to optimize costs with minimal changes to the application.

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

- A. Create a new S3 bucket that has server-side encryption with customer-provided keys (SSE-C) as the encryption type
- B. Copy the existing objects to the new S3 bucket
- C. Specify SSE-C.
- D. Create a new S3 bucket that has server-side encryption with Amazon S3 managed keys (SSE-S3) as the encryption type
- E. Use S3 Batch Operations to copy the existing objects to the new S3 bucket
- F. Specify SSE-S3.

- G. Use AWS CloudHSM to store the encryption key
- H. Create a new S3 bucket
- I. Use S3 Batch Operations to copy the existing objects to the new S3 bucket
- J. Encrypt the objects by using the keys from CloudHSM.
- K. Use the S3 Intelligent-Tiering storage class for the S3 bucket
- L. Create an S3 Intelligent-Tiering archive configuration to transition objects that are not accessed for 90 days to S3 Glacier Deep Archive.

Answer: B

Explanation:

To reduce the volume of Amazon S3 calls to AWS KMS, use Amazon S3 bucket keys, which are protected encryption keys that are reused for a limited time in Amazon S3. Bucket keys can reduce costs for AWS KMS requests by up to 99%. You can configure a bucket key for all objects in an Amazon S3 bucket, or for a specific object in an Amazon S3 bucket. https://docs.aws.amazon.com/fr_fr/kms/latest/developerguide/services-s3.html

NEW QUESTION 85

- (Exam Topic 2)

A company is deploying a new web-based application and needs a storage solution for the Linux application servers. The company wants to create a single location for updates to application data for all instances. The active dataset will be up to 100 GB in size. A solutions architect has determined that peak operations will occur for 3 hours daily and will require a total of 225 MiBps of read throughput.

The solutions architect must design a Multi-AZ solution that makes a copy of the data available in another AWS Region for disaster recovery (DR). The DR copy has an RPO of less than 1 hour.

Which solution will meet these requirements?

- A. Deploy a new Amazon Elastic File System (Amazon EFS) Multi-AZ file system
- B. Configure the file system for 75 MiBps of provisioned throughput
- C. Implement replication to a file system in the DR Region.
- D. Deploy a new Amazon FSx for Lustre file system
- E. Configure Bursting Throughput mode for the file system
- F. Use AWS Backup to back up the file system to the DR Region.
- G. Deploy a General Purpose SSD (gp3) Amazon Elastic Block Store (Amazon EBS) volume with 225 MiBps of throughput
- H. Enable Multi-Attach for the EBS volume
- I. Use AWS Elastic Disaster Recovery to replicate the EBS volume to the DR Region.
- J. Deploy an Amazon FSx for OpenZFS file system in both the production Region and the DR Region. Create an AWS DataSync scheduled task to replicate the data from the production file system to the DR file system every 10 minutes.

Answer: A

Explanation:

The company should deploy a new Amazon Elastic File System (Amazon EFS) Multi-AZ file system. The company should configure the file system for 75 MiBps of provisioned throughput. The company should implement replication to a file system in the DR Region. This solution will meet the requirements because Amazon EFS is a serverless, fully elastic file storage service that lets you share file data without provisioning or managing storage capacity and performance. Amazon EFS is built to scale on demand to petabytes without disrupting applications, growing and shrinking automatically as you add and remove files¹. By deploying a new Amazon EFS Multi-AZ file system, the company can create a single location for updates to application data for all instances. A Multi-AZ file system replicates data across multiple Availability Zones (AZs) within a Region, providing high availability and durability². By configuring the file system for 75 MiBps of provisioned throughput, the company can ensure that it meets the peak operations requirement of 225 MiBps of read throughput. Provisioned throughput is a feature that enables you to specify a level of throughput that the file system can drive independent of the file system's size or burst credit balance³. By implementing replication to a file system in the DR Region, the company can make a copy of the data available in another AWS Region for disaster recovery. Replication is a feature that enables you to replicate data from one EFS file system to another EFS file system across AWS Regions. The replication process has an RPO of less than 1 hour.

The other options are not correct because:

➤ Deploying a new Amazon FSx for Lustre file system would not provide a single location for updates to application data for all instances. Amazon FSx for Lustre is a fully managed service that provides cost-effective, high-performance storage for compute workloads. However, it does not support concurrent write access from multiple instances. Using AWS Backup to back up the file system to the DR Region would not provide real-time replication of data. AWS Backup is a service that enables you to centralize and automate data protection across AWS services. However, it does not support continuous data replication or cross-Region disaster recovery.

➤ Deploying a General Purpose SSD (gp3) Amazon Elastic Block Store (Amazon EBS) volume with 225 MiBps of throughput would not provide a single location for updates to application data for all instances. Amazon EBS is a service that provides persistent block storage volumes for use with Amazon EC2 instances. However, it does not support concurrent access from multiple instances, unless Multi-Attach is enabled. Enabling Multi-Attach for the EBS volume would not provide Multi-AZ resilience or cross-Region replication. Multi-Attach is a feature that enables you to attach an EBS volume to multiple EC2 instances within the same Availability Zone. Using AWS Elastic Disaster Recovery to replicate the EBS volume to the DR Region would not provide real-time replication of data. AWS Elastic Disaster Recovery (AWS DRS) is a service that enables you to orchestrate and automate disaster recovery workflows across AWS Regions. However, it does not support continuous data replication or sub-hour RPOs.

➤ Deploying an Amazon FSx for OpenZFS file system in both the production Region and the DR Region would not be as simple or cost-effective as using Amazon EFS. Amazon FSx for OpenZFS is a fully managed service that provides high-performance storage with strong data consistency and advanced data management features for Linux workloads. However, it requires more configuration and management than Amazon EFS, which is serverless and fully elastic. Creating an AWS DataSync scheduled task to replicate the data from the production file system to the DR file system every 10 minutes would not provide real-time replication of data. AWS DataSync is a service that enables you to transfer data between on-premises storage and AWS services, or between AWS services. However, it does not support continuous data replication or sub-minute RPOs.

References:

- <https://aws.amazon.com/efs/>
- <https://docs.aws.amazon.com/efs/latest/ug/how-it-works.html#how-it-works-azs>
- <https://docs.aws.amazon.com/efs/latest/ug/performance.html#provisioned-throughput>
- <https://docs.aws.amazon.com/efs/latest/ug/replication.html>
- <https://aws.amazon.com/fsx/lustre/>
- <https://aws.amazon.com/backup/>
- <https://aws.amazon.com/ebs/>
- <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-volumes-multi.html>

NEW QUESTION 89

- (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 subne
- 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 NL
- F. In the case of a failed health check, redeploy the NLB in different subnets.
- G. Deploy a single NAT gateway in a public subne
- H. Assign the Elastic IP address to the NAT gateway.Use Amazon CloudWatch with a custom metric tomonitor the NAT gatewa
- I. If the NAT gateway is unhealthy, invoke an AWS Lambda function to create a new NAT gateway in a different subne
- J. Assign the Elastic IP address to the new NAT gateway.
- K. Assign the Elastic IP address to the AL
- L. Create an Amazon Route 53 simple record with the Elastic IP address as the valu
- M. Create a Route 53 health chec
- 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 91

- (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 applicatio
- 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 AP
- 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 CloudFron
- 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 94

- (Exam Topic 2)

A company has a critical application in which the data tier is deployed in a single AWS Region. The data tier uses an Amazon DynamoDB table and an Amazon Aurora MySQL DB cluster. The current Aurora MySQL engine version supports a global database. The application tier is already deployed in two Regions. Company policy states that critical applications must have application tier components and data tier components deployed across two Regions. The RTO and RPO must be no more than a few minutes each. A solutions architect must recommend a solution to make the data tier compliant with company policy. Which combination of steps will meet these requirements? (Choose two.)

- A. Add another Region to the Aurora MySQL DB cluster
- B. Add another Region to each table in the Aurora MySQL DB cluster
- C. Set up scheduled cross-Region backups for the DynamoDB table and the Aurora MySQL DB cluster
- D. Convert the existing DynamoDB table to a global table by adding another Region to its configuration
- E. Use Amazon Route 53 Application Recovery Controller to automate database backup and recovery to the secondary Region

Answer: AD

Explanation:

The company should use Amazon Aurora global database and Amazon DynamoDB global table to deploy the data tier components across two Regions. Amazon Aurora global database is a feature that allows a single Aurora database to span multiple AWS Regions, enabling low-latency global reads and fast recovery from Region-wide outages¹. Amazon DynamoDB global table is a feature that allows a single DynamoDB table to span multiple AWS Regions, enabling low-latency global reads and writes and fast recovery from Region-wide outages².

References:

- <https://aws.amazon.com/rds/aurora/global-database/>
- https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/globaltables_HowItWorks.html
- <https://aws.amazon.com/route53/application-recovery-controller/>

NEW QUESTION 99

- (Exam Topic 2)

A company has several AWS accounts. A development team is building an automation framework for cloud governance and remediation processes. The automation framework uses AWS Lambda functions in a centralized account. A solutions architect must implement a least privilege permissions policy that allows the Lambda functions to run in each of the company's AWS accounts. Which combination of steps will meet these requirements? (Choose two.)

- A. In the centralized account, create an IAM role that has the Lambda service as a trusted entity
- B. Add an inline policy to assume the roles of the other AWS accounts.
- C. In the other AWS accounts, create an IAM role that has minimal permission
- D. Add the centralized account's Lambda IAM role as a trusted entity.
- E. In the centralized account, create an IAM role that has roles of the other accounts as trusted entities. Provide minimal permissions.
- F. In the other AWS accounts, create an IAM role that has permissions to assume the role of the centralized account
- G. Add the Lambda service as a trusted entity.
- H. In the other AWS accounts, create an IAM role that has minimal permission
- I. Add the Lambda service as a trusted entity.

Answer: AB

Explanation:

<https://medium.com/@it.melnichenko/invoke-a-lambda-across-multiple-aws-accounts-8c094b2e70be>

NEW QUESTION 100

- (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 102

- (Exam Topic 2)

A company has multiple business units that each have separate accounts on AWS. Each business unit manages its own network with several VPCs that have CIDR ranges that overlap. The company's marketing team has created a new internal application and wants to make the application accessible to all the other business units. The solution must use private IP addresses only.

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

- A. Instruct each business unit to add a unique secondary CIDR range to the business unit's VP
- B. Peer the VPCs and use a private NAT gateway in the secondary range to route traffic to the marketing team.
- C. Create an Amazon EC2 instance to serve as a virtual appliance in the marketing account's VP
- D. Create an AWS Site-to-Site VPN connection between the marketing team and each business unit's VP
- E. Perform NAT where necessary.
- F. Create an AWS PrivateLink endpoint service to share the marketing applicatio
- G. Grant permission to specific AWS accounts to connect to the servic
- H. Create interface VPC endpoints in other accounts to access the application by using private IP addresses.
- I. Create a Network Load Balancer (NLB) in front of the marketing application in a private subne
- J. Create an API Gateway AP
- K. Use the Amazon API Gateway private integration to connect the API to the NL
- L. Activate IAM authorization for the AP
- M. Grant access to the accounts of the other business units.

Answer: C

Explanation:

With AWS PrivateLink, the marketing team can create an endpoint service to share their internal application with other accounts securely using private IP addresses. They can grant permission to specific AWS accounts to connect to the service and create interface VPC endpoints in the other accounts to access the application by using private IP addresses. This option does not require any changes to the network of the other business units, and it does not require peering or NATing. This solution is both scalable and secure.

<https://aws.amazon.com/blogs/networking-and-content-delivery/connecting-networks-with-overlapping-ip-range>

NEW QUESTION 103

- (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 104

- (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 106

- (Exam Topic 2)

A company has a new application that needs to run on five Amazon EC2 instances in a single AWS Region. The application requires high-throughput, low-latency network connections between all to the EC2 instances where the application will run. There is no requirement for the application to be fault tolerant. Which solution will meet these requirements?

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

Answer: A

Explanation:

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

NEW QUESTION 111

- (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 112

- (Exam Topic 2)

A large company runs workloads in VPCs that are deployed across hundreds of AWS accounts. Each VPC consists to public subnets and private subnets that span across multiple Availability Zones. NAT gateways are deployed in the public subnets and allow outbound connectivity to the internet from the private subnets. A solutions architect is working on a hub-and-spoke design. All private subnets in the spoke VPCs must route traffic to the internet through an egress VPC. The solutions architect already has deployed a NAT gateway in an egress VPC in a central AWS account.

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

- A. Create peering connections between the egress VPC and the spoke VPC
- B. Configure the required routing to allow access to the internet.
- C. Create a transit gateway, and share it with the existing AWS account
- D. Attach existing VPCs to the transit gateway Configure the required routing to allow access to the internet.
- E. Create a transit gateway in every account
- F. Attach the NAT gateway to the transit gateway
- G. Configure the required routing to allow access to the internet.
- H. Create an AWS PrivateLink connection between the egress VPC and the spoke VPC
- I. Configure the required routing to allow access to the internet

Answer: B

Explanation:

<https://d1.awsstatic.com/architecture-diagrams/ArchitectureDiagrams/NAT-gateway-centralized-egress-ra.pdf?d>

NEW QUESTION 113

- (Exam Topic 2)

A company is running an application that uses an Amazon ElastiCache for Redis cluster as a caching layer. A recent security audit revealed that the company has configured encryption at rest for ElastiCache. However, the company did not configure ElastiCache to use encryption in transit. Additionally, users can access the cache without authentication.

A solutions architect must make changes to require user authentication and to ensure that the company is using end-to-end encryption.

Which solution will meet these requirements?

- A. Create an AUTH token. Store the token in AWS System Manager Parameter Store, as an encrypted parameter. Create a new cluster with AUTH and configure encryption in transit. Update the application to retrieve the AUTH token from Parameter Store when necessary and to use the AUTH token for authentication.
- B. Create an AUTH token. Store the token in AWS Secrets Manager. Configure the existing cluster to use the AUTH token and configure encryption in transit. Update the application to retrieve the AUTH token from Secrets Manager when necessary and to use the AUTH token for authentication.
- C. Create an SSL certificate. Store the certificate in AWS Secrets Manager. Create a new cluster and configure encryption in transit. Update the application to retrieve the SSL certificate from Secrets Manager when necessary and to use the certificate for authentication.
- D. Create an SSL certificate. Store the certificate in AWS Systems Manager Parameter Store, as an encrypted advanced parameter. Update the existing cluster to configure encryption in transit. Update the application to retrieve the SSL certificate from Parameter Store when necessary and to use the certificate for authentication.

Answer: B

Explanation:

Creating an AUTH token and storing it in AWS Secrets Manager and configuring the existing cluster to use the AUTH token and configure encryption in transit, and updating the application to retrieve the AUTH token from Secrets Manager when necessary and to use the AUTH token for authentication, would meet the requirements for user authentication and end-to-end encryption.

AWS Secrets Manager is a service that enables you to easily rotate, manage, and retrieve database credentials, API keys, and other secrets throughout their lifecycle. Secrets Manager also enables you to encrypt the data and ensure that only authorized users and applications can access it.

By configuring the existing cluster to use the AUTH token and encryption in transit, all data will be encrypted as it is sent over the network, providing additional security for the data stored in ElastiCache.

Additionally, by updating the application to retrieve the AUTH token from Secrets Manager when necessary and to use the AUTH token for authentication, it ensures that only authorized users and applications can access the cache.

Reference:

AWS Secrets Manager documentation: <https://aws.amazon.com/secrets-manager/> Encryption in transit for ElastiCache:

<https://docs.aws.amazon.com/AmazonElastiCache/latest/red-ug/encryption.html>

Authentication and Authorization for ElastiCache: <https://docs.aws.amazon.com/AmazonElastiCache/latest/red-ug/accessing-elasticache.html>

NEW QUESTION 114

- (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 117

- (Exam Topic 2)

A solutions architect needs to define a reference architecture for a solution for three-tier applications with web, application, and NoSQL data layers. The reference architecture must meet the following requirements:

- High availability within an AWS Region
- Able to fail over in 1 minute to another AWS Region for disaster recovery
- Provide the most efficient solution while minimizing the impact on the user experience Which combination of steps will meet these requirements? (Select THREE.)

- A. Use an Amazon Route 53 weighted routing policy set to 100/0 across the two selected Region
- B. Set Time to Live (TTL) to 1 hour.
- C. Use an Amazon Route 53 failover routing policy for failover from the primary Region to the disaster recovery Region
- D. Set Time to Live (TTL) to 30 seconds.
- E. Use a global table within Amazon DynamoDB so data can be accessed in the two selected Regions.
- F. Back up data from an Amazon DynamoDB table in the primary Region every 60 minutes and then write the data to Amazon S3. Use S3 Cross-Region replication to copy the data from the primary Region to the disaster recovery Region
- G. Have a script import the data into DynamoDB in a disaster recovery scenario.
- H. Implement a hot standby model using Auto Scaling groups for the web and application layers across multiple Availability Zones in the Region
- I. Use zonal Reserved Instances for the minimum number of servers and On-Demand Instances for any additional resources.
- J. Use Auto Scaling groups for the web and application layers across multiple Availability Zones in the Region
- K. Use Spot Instances for the required resources.

Answer: BCE

Explanation:

The requirements can be achieved by using an Amazon DynamoDB database with a global table. DynamoDB is a NoSQL database so it fits the requirements. A global table also allows both reads and writes to occur in both Regions. For the web and application tiers Auto Scaling groups should be configured. Due to the 1-minute RTO these must be configured in an active/passive state. The best pricing model to lower price but ensure resources are available when needed is to use a combination of zonal reserved instances and on-demand instances. To failover between the Regions, a Route 53 failover routing policy can be configured with a TTL configured on the record of 30 seconds. This will mean clients must resolve against Route 53 every 30 seconds to get the latest record. In a failover scenario the clients would be redirected to the secondary site if the primary site is unhealthy.

NEW QUESTION 121

- (Exam Topic 2)

A company uses a load balancer to distribute traffic to Amazon EC2 instances in a single Availability Zone. The company is concerned about security and wants a solutions architect to re-architect the solution to meet the following requirements:

- Inbound requests must be filtered for common vulnerability attacks.
- Rejected requests must be sent to a third-party auditing application.
- All resources should be highly available. Which solution meets these requirements?

- A. Configure a Multi-AZ Auto Scaling group using the application's AM
- B. Create an Application Load Balancer (ALB) and select the previously created Auto Scaling group as the target
- C. Use Amazon Inspector to monitor traffic to the ALB and EC2 instance
- D. Create a web ACL in WA
- E. Create an AWS WAF using the web ACL and AL
- F. Use an AWS Lambda function to frequently push the Amazon Inspector report to the third-party auditing application.
- G. Configure an Application Load Balancer (ALB) and add the EC2 instances as targets Create a web ACL in WA
- H. Create an AWS WAF using the web ACL and ALB name and enable logging with Amazon CloudWatch Log
- I. Use an AWS Lambda function to frequently push the logs to the third-party auditing application.
- J. Configure an Application Load Balancer (ALB) along with a target group adding the EC2 instances as target
- K. Create an Amazon Kinesis Data Firehose with the destination of the third-party auditing applicatio

- L. Create a web ACL in WA
- M. Create an AWS WAF using the web ACL and ALB then enable logging by selecting the Kinesis Data Firehose as the destination
- N. Subscribe to AWS Managed Rules in AWS Marketplace, choosing the WAF as the subscriber.
- O. Configure a Multi-AZ Auto Scaling group using the application's AM
- P. Create an Application Load Balancer (ALB) and select the previously created Auto Scaling group as the target
- Q. Create an Amazon Kinesis Data Firehose with a destination of the third-party auditing application
- R. Create a web ACL in WA
- S. Create an AWS WAF using the WebACL and ALB then enable logging by selecting the Kinesis Data Firehose as the destination
- T. Subscribe to AWS Managed Rules in AWS Marketplace, choosing the WAF as the subscriber.

Answer: D

Explanation:

<https://docs.aws.amazon.com/waf/latest/developerguide/marketplace-managed-rule-groups.html>

NEW QUESTION 124

- (Exam Topic 2)

A company has deployed its database on an Amazon RDS for MySQL DB instance in the us-east-1 Region. The company needs to make its data available to customers in Europe. The customers in Europe must have access to the same data as customers in the United States (US) and will not tolerate high application latency or stale data. The customers in Europe and the customers in the US need to write to the database. Both groups of customers need to see updates from the other group in real time.

Which solution will meet these requirements?

- A. Create an Amazon Aurora MySQL replica of the RDS for MySQL DB instance
- B. Pause application writes to the RDS DB instance
- C. Promote the Aurora Replica to a standalone DB instance
- D. Reconfigure the application to use the Aurora database and resume writes
- E. Add eu-west-1 as a secondary Region to the DB instance
- F. Enable write forwarding on the DB instance
- G. Deploy the application in eu-west-1. Configure the application to use the Aurora MySQL endpoint in eu-west-1.
- H. Add a cross-Region replica in eu-west-1 for the RDS for MySQL DB instance
- I. Configure the replica to replicate write queries back to the primary DB instance
- J. Deploy the application in eu-west-1. Configure the application to use the RDS for MySQL endpoint in eu-west-1.
- K. Copy the most recent snapshot from the RDS for MySQL DB instance to eu-west-1. Create a new RDS for MySQL DB instance in eu-west-1 from the snapshot
- L. Configure MySQL logical replication from us-east-1 to eu-west-1. Enable write forwarding on the DB instance
- M. Deploy the application in eu-west-1. Configure the application to use the RDS for MySQL endpoint in eu-west-1.
- N. Convert the RDS for MySQL DB instance to an Amazon Aurora MySQL DB instance
- O. Add eu-west-1 as a secondary Region to the DB instance
- P. Enable write forwarding on the DB instance
- Q. Deploy the application in eu-west-1. Configure the application to use the Aurora MySQL endpoint in eu-west-1.

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 needed. 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 data. 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 users.

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 127

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

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

* One year free update

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

* Trusted by Millions

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

* Shop Securely

All transactions are protected by VeriSign!

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

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