

Amazon-Web-Services

Exam Questions SAP-C02

AWS Certified Solutions Architect - Professional



NEW QUESTION 1

- (Exam Topic 1)

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

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

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

Answer: B

Explanation:

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

NEW QUESTION 2

- (Exam Topic 1)

An application is using an Amazon RDS for MySQL Multi-AZ DB instance in the us-east-1 Region. After a failover test, the application lost the connections to the database and could not re-establish the connections. After a restart of the application, the application re-established the connections.

A solutions architect must implement a solution so that the application can re-establish connections to the database without requiring a restart.

Which solution will meet these requirements?

- A. Create an Amazon Aurora MySQL Serverless v1 DB instance
- B. Migrate the RDS DB instance to the Aurora Serverless v1 DB instance
- C. Update the connection settings in the application to point to the Aurora reader endpoint.
- D. Create an RDS proxy
- E. Configure the existing RDS endpoint as a target
- F. Update the connection settings in the application to point to the RDS proxy endpoint.
- G. Create a two-node Amazon Aurora MySQL DB cluster
- H. Migrate the RDS DB instance to the Aurora DB cluster
- I. Create an RDS proxy
- J. Configure the existing RDS endpoint as a target
- K. Update the connection settings in the application to point to the RDS proxy endpoint.
- L. Create an Amazon S3 bucket
- M. Export the database to Amazon S3 by using AWS Database Migration Service (AWS DMS). Configure Amazon Athena to use the S3 bucket as a data store
- N. Install the latest Open Database Connectivity (ODBC) driver for the application
- O. Update the connection settings in the application to point to the Athena endpoint

Answer: B

Explanation:

Amazon RDS Proxy is a fully managed database proxy service for Amazon Relational Database Service (RDS) that makes applications more scalable, resilient, and secure. It allows applications to pool and share connections to an RDS database, which can help reduce database connection overhead, improve scalability, and provide automatic failover and high availability.

NEW QUESTION 3

- (Exam Topic 1)

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

What is the MOST cost-effective migration recommendation?

- A. Create a queue using Amazon SQS
- B. Configure the existing web server to publish to the new queue. When there are messages in the queue, invoke an AWS Lambda function to pull requests from the queue and process the file
- C. Store the processed files in an Amazon S3 bucket.
- D. Create a queue using Amazon SNS
- E. Configure the existing web server to publish to the new queue
- F. When there are messages in the queue, create a new Amazon EC2 instance to pull requests from the queue and process the file
- G. Store the processed files in Amazon EFS
- H. Shut down the EC2 instance after the task is complete.
- I. Create a queue using Amazon MQ
- J. Configure the existing web server to publish to the new queue. When there are messages in the queue, invoke an AWS Lambda function to pull requests from the queue and process the file
- K. Store the processed files in Amazon EFS.

- L. Create a queue using Amazon SQS
- M. Configure the existing web server to publish to the new queue
- N. Use Amazon EC2 instances in an EC2 Auto Scaling group to pull requests from the queue and process the file
- O. Scale the EC2 instances based on the SQS queue length
- P. Store the processed files in an Amazon S3 bucket.

Answer: D

Explanation:

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

NEW QUESTION 4

- (Exam Topic 1)

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

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

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

Which solution will meet these requirements?

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

Answer: D

Explanation:

Extend the system with an application tier that uses AWS Step Functions and AWS Lambda. Configure this tier to use Amazon Textract and Amazon Comprehend to perform optical character recognition (OCR) on the forms when forms are uploaded. Store the output in Amazon S3. Parse this output by extracting the data that is required within the application tier. Submit the data to the target system's API. This solution meets the requirements of accurate form extraction, minimal time to market, and minimal long-term operational overhead. Amazon Textract and Amazon Comprehend are fully managed and serverless services that can perform OCR and extract relevant data from the forms, which eliminates the need to develop custom libraries or train and host models. Using AWS Step Functions and Lambda allows for easy automation of the process and the ability to scale as needed.

NEW QUESTION 5

- (Exam Topic 1)

A company uses AWS Organizations for a multi-account setup in the AWS Cloud. The company uses AWS Control Tower for governance and uses AWS Transit Gateway for VPC connectivity across accounts.

In an AWS application account, the company's application team has deployed a web application that uses AWS Lambda and Amazon RDS. The company's database administrators have a separate DBA account and use the account to centrally manage all the databases across the organization. The database administrators use an Amazon EC2 instance that is deployed in the DBA account to access an RDS database that is deployed in the application account.

The application team has stored the database credentials as secrets in AWS Secrets Manager in the application account. The application team is manually sharing the secrets with the database administrators. The secrets are encrypted by the default AWS managed key for Secrets Manager in the application account. A solutions architect needs to implement a solution that gives the database administrators access to the database and eliminates the need to manually share the secrets.

Which solution will meet these requirements?

- A. Use AWS Resource Access Manager (AWS RAM) to share the secrets from the application account with the DBA account
- B. In the DBA account, create an IAM role that is named DBA-Admin
- C. Grant the role the required permissions to access the shared secret
- D. Attach the DBA-Admin role to the EC2 instance for access to the cross-account secrets.
- E. In the application account, create an IAM role that is named DBA-Secret
- F. Grant the role the required permissions to access the secret
- G. In the DBA account, create an IAM role that is named DBA-Admin
- H. Grant the DBA-Admin role the required permissions to assume the DBA-Secret role in the application account
- I. Attach the DBA-Admin role to the EC2 instance for access to the cross-account secrets.
- J. In the DBA account, create an IAM role that is named DBA-Admin
- K. Grant the role the required permissions to access the secrets and the default AWS managed key in the application account
- L. In the application account, attach resource-based policies to the key to allow access from the DBA account
- M. Attach the DBA-Admin role to the EC2 instance for access to the cross-account secrets.

- N. In the DBA account, create an IAM role that is named DBA-Admi
O. Grant the role the required permissions to access the secrets in the application account
P. Attach an SCP to the application account to allow access to the secrets from the DBA account
Q. Attach the DBA-Admin role to the EC2 instance for access to the cross-account secrets.

Answer: B

Explanation:

➤ Option B is correct because creating an IAM role in the application account that has permissions to access the secrets and creating an IAM role in the DBA account that has permissions to assume the role in the application account eliminates the need to manually share the secrets. This approach uses cross-account IAM roles to grant access to the secrets in the application account. The database administrators can assume the role in the application account from their EC2 instance in the DBA account and retrieve the secrets without having to store them locally or share them manually.
References: 1: <https://docs.aws.amazon.com/ram/latest/userguide/what-is.html> 2: https://docs.aws.amazon.com/IAM/latest/UserGuide/tutorial_cross-account-with-roles.html 3: <https://docs.aws.amazon.com/kms/latest/developerguide/concepts.html> : https://docs.aws.amazon.com/secretsmanager/latest/userguide/tutorials_basic.html : <https://docs.aws.amazon.com/IAM/latest/UserGuide/introduction.html>

NEW QUESTION 6

- (Exam Topic 1)

A company is hosting a three-tier web application in an on-premises environment. Due to a recent surge in traffic that resulted in downtime and a significant financial impact, company management has ordered that the application be moved to AWS. The application is written in .NET and has a dependency on a MySQL database. A solutions architect must design a scalable and highly available solution to meet the demand of 200,000 daily users. Which steps should the solutions architect take to design an appropriate solution?

- A. Use AWS Elastic Beanstalk to create a new application with a web server environment and an Amazon RDS MySQL Multi-AZ DB instance. The environment should launch a Network Load Balancer (NLB) in front of an Amazon EC2 Auto Scaling group in multiple Availability Zones. Use an Amazon Route 53 alias record to route traffic from the company's domain to the NLB.
B. Use AWS CloudFormation to launch a stack containing an Application Load Balancer (ALB) in front of an Amazon EC2 Auto Scaling group spanning three Availability Zones.
C. The stack should launch a Multi-AZ deployment of an Amazon Aurora MySQL DB cluster with a Retain deletion policy.
D. Use an Amazon Route 53 alias record to route traffic from the company's domain to the ALB.
E. Use AWS Elastic Beanstalk to create an automatically scaling web server environment that spans two separate Regions with an Application Load Balancer (ALB) in each Region.
F. Create a Multi-AZ deployment of an Amazon Aurora MySQL DB cluster with a cross-Region read replica. Use Amazon Route 53 with a geoproximity routing policy to route traffic between the two Regions.
G. Use AWS CloudFormation to launch a stack containing an Application Load Balancer (ALB) in front of an Amazon ECS cluster of Spot Instances spanning three Availability Zones. The stack should launch an Amazon RDS MySQL DB instance with a Snapshot deletion policy. Use an Amazon Route 53 alias record to route traffic from the company's domain to the ALB.

Answer: C

Explanation:

Using AWS CloudFormation to launch a stack with an Application Load Balancer (ALB) in front of an Amazon EC2 Auto Scaling group spanning three Availability Zones, a Multi-AZ deployment of an Amazon Aurora MySQL DB cluster with a Retain deletion policy, and an Amazon Route 53 alias record to route traffic from the company's domain to the ALB will ensure that

NEW QUESTION 7

- (Exam Topic 1)

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

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

Answer: ABE

Explanation:

<https://docs.aws.amazon.com/config/latest/developerguide/config-rule-multi-account-deployment.html>
<https://docs.aws.amazon.com/config/latest/developerguide/aggregate-data.html>
https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps_examples_tagging.htm

NEW QUESTION 8

- (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-transitgatewayattachment.html>

NEW QUESTION 9

- (Exam Topic 1)

A company is migrating some of its applications to AWS. The company wants to migrate and modernize the applications quickly after it finalizes networking and security strategies. The company has set up an AWS Direct Connection connection in a central network account. The company expects to have hundreds of AWS accounts and VPCs in the near future. The corporate network must be able to access the resources on AWS seamlessly and also must be able to communicate with all the VPCs. The company also wants to route its cloud resources to the internet through its on-premises data center.

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

- A. Create a Direct Connect gateway in the central account
- B. In each of the accounts, create an association proposal by using the Direct Connect gateway and the account ID for every virtual private gateway.
- C. Create a Direct Connect gateway and a transit gateway in the central network account
- D. Attach the transit gateway to the Direct Connect gateway by using a transit VIF.
- E. Provision an internet gateway
- F. Attach the internet gateway to subnet
- G. Allow internet traffic through the gateway.
- H. Share the transit gateway with other account
- I. Attach VPCs to the transit gateway.
- J. Provision VPC peering as necessary.
- K. Provision only private subnet
- L. Open the necessary route on the transit gateway and customer gateway to allow outbound internet traffic from AWS to flow through NAT services that run in the data center.

Answer: BDF

Explanation:

➤ Option A is incorrect because creating a Direct Connect gateway in the central account and creating an association proposal by using the Direct Connect gateway and the account ID for every virtual private gateway does not enable active-passive failover between the regions. A Direct Connect gateway is a globally available resource that enables you to connect your AWS Direct Connect connection over a private virtual interface (VIF) to one or more VPCs in any AWS Region. A virtual private gateway is the VPN concentrator on the Amazon side of a VPN connection. You can associate a Direct Connect gateway with either a transit gateway or a virtual private gateway. However, a Direct Connect gateway does not provide any load balancing or failover capabilities by itself.

➤ Option B is correct because creating a Direct Connect gateway and a transit gateway in the central network account and attaching the transit gateway to the Direct Connect gateway by using a transit VIF meets the requirement of enabling the corporate network to access the resources on AWS seamlessly and also to communicate with all the VPCs. A transit VIF is a type of private VIF that you can use to connect your AWS Direct Connect connection to a transit gateway or a Direct Connect gateway. A transit gateway is a network transit hub that you can use to interconnect your VPCs and on-premises networks. By using a transit VIF, you can route traffic between your on-premises network and multiple VPCs across different AWS accounts and Regions through a single connection.

➤ Option C is incorrect because provisioning an internet gateway, attaching the internet gateway to subnets, and allowing internet traffic through the gateway does not meet the requirement of routing cloud resources to the internet through its on-premises data center. An internet gateway is a horizontally scaled, redundant, and highly available VPC component that allows communication between your VPC and the internet. An internet gateway serves two purposes: to provide a target in your VPC route tables for internet-routable traffic, and to perform network address translation (NAT) for instances that have been assigned public IPv4 addresses. By using an internet gateway, you are routing cloud resources directly to the internet, not through your on-premises data center.

➤ Option D is correct because sharing the transit gateway with other accounts and attaching VPCs to the transit gateway meets the requirement of enabling the corporate network to access the resources on AWS seamlessly and also to communicate with all the VPCs. You can share your transit gateway with other AWS accounts within the same organization by using AWS Resource Access Manager (AWS RAM). This allows you to centrally manage connectivity from multiple accounts without having to create individual peering connections between VPCs or duplicate network appliances in each account. You can attach VPCs from different accounts and Regions to your shared transit gateway and enable routing between them.

➤ Option E is incorrect because provisioning VPC peering as necessary does not meet the requirement of enabling the corporate network to access the resources on AWS seamlessly and also to communicate with all the VPCs. VPC peering is a networking connection between two VPCs that enables you to route traffic between them using private IPv4 addresses or IPv6 addresses. You can create a VPC peering connection between your own VPCs, or with a VPC in another AWS account within a single Region. However, VPC peering does not allow you to route traffic from your on-premises network to your VPCs or between multiple Regions. You would need to create multiple VPN connections or Direct Connect connections for each VPC peering connection, which increases operational complexity and costs.

➤ Option F is correct because provisioning only private subnets, opening the necessary route on the transit gateway and customer gateway to allow outbound internet traffic from AWS to flow through NAT services that run in the data center meets the requirement of routing cloud resources to the internet through its on-premises data center. A private subnet is a subnet that's associated with a route table that has no route to an internet gateway. Instances in a private subnet can communicate with other instances in the same VPC but cannot access resources on the internet directly. To enable outbound internet access from instances in private subnets, you can use NAT devices such as NAT gateways or NAT instances that are deployed in public subnets. A public subnet is a subnet that's associated with a route table that has a route to an internet gateway. Alternatively, you can use your on-premises data center as a NAT device by configuring routes on your transit gateway and customer gateway that direct outbound internet traffic from your private subnets through your VPN connection or Direct Connect connection. This way, you can route cloud resources to the internet through your on-premises data center instead of using an internet gateway.

References: 1:

<https://docs.aws.amazon.com/directconnect/latest/UserGuide/direct-connect-gateways-intro.html> 2:

<https://docs.aws.amazon.com/directconnect/latest/UserGuide/direct-connect-transit-virtual-interfaces.html> 3: <https://docs.aws.amazon.com/vpc/latest/tgw/what-is-transit-gateway.html> : https://docs.aws.amazon.com/vpc/latest/userguide/VPC_Internet_Gateway.html : <https://docs.aws.amazon.com/vpc/latest/tgw/tgw-sharing.html> : <https://docs.aws.amazon.com/vpc/latest/peering/what-is-vpc-peering.html> : https://docs.aws.amazon.com/vpc/latest/userguide/VPC_Scenario2.html : https://docs.aws.amazon.com/vpc/latest/userguide/VPC_Scenario3.html : https://docs.aws.amazon.com/vpc/latest/userguide/VPC_NAT_Instance.html :

https://docs.aws.amazon.com/vpc/latest/userguide/VPC_NAT_Gateway.html

NEW QUESTION 10

- (Exam Topic 1)

A publishing company's design team updates the icons and other static assets that an ecommerce 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 10

- (Exam Topic 1)

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

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

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

Answer: B

Explanation:

<https://aws.amazon.com/aws-transfer-family/faqs/> <https://docs.aws.amazon.com/transfer/latest/userguide/what-is-aws-transfer-family.html>
https://aws.amazon.com/about-aws/whats-new/2018/11/aws-transfer-for-sftp-fully-managed-sftp-for-s3/?nc1=h_

NEW QUESTION 15

- (Exam Topic 1)

A company is running a critical application that uses an Amazon RDS for MySQL database to store data. The RDS DB instance is deployed in Multi-AZ mode. A recent RDS database failover test caused a 40-second outage to the application. A solutions architect needs to design a solution to reduce the outage time to less than 20 seconds.

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

- A. Use Amazon ElastiCache for Memcached in front of the database
- B. Use Amazon ElastiCache for Redis in front of the database.
- C. Use RDS Proxy in front of the database
- D. Migrate the database to Amazon Aurora MySQL
- E. Create an Amazon Aurora Replica
- F. Create an RDS for MySQL read replica

Answer: CDE

Explanation:

Migrate the database to Amazon Aurora MySQL. - Create an Amazon Aurora Replica. - Use RDS Proxy in front of the database. - These options are correct because they address the requirement of reducing the failover time to less than 20 seconds. Migrating to Amazon Aurora MySQL and creating an Aurora replica can reduce the failover time to less than 20 seconds. Aurora has a built-in, fault-tolerant storage system that can automatically detect and repair failures. Additionally, Aurora has a feature called "Aurora Global Database" which allows you to create read-only replicas across multiple AWS regions which can further help to reduce the failover time. Creating an Aurora replica can also help to reduce the failover time as it can take over as the primary DB instance in case of a failure. Using RDS proxy can also help to reduce the failover time as it can route the queries to the healthy DB instance, it also helps to balance the load across multiple DB instances.

NEW QUESTION 18

- (Exam Topic 1)

An enterprise company wants to allow its developers to purchase third-party software through AWS Marketplace. The company uses an AWS Organizations account structure with full features enabled, and has a shared services account in each organizational unit (OU) that will be used by procurement managers. The procurement team's policy indicates that developers should be able to obtain third-party software from an approved list only and use Private Marketplace in AWS Marketplace to achieve this requirement. The procurement team wants administration of Private Marketplace to be restricted to a role named procurement-manager-role, which could be assumed by procurement managers. Other IAM users, groups, roles, and account administrators in the company should be denied Private Marketplace administrative access.

What is the MOST efficient way to design an architecture to meet these requirements?

- A. Create an IAM role named procurement-manager-role in all AWS accounts in the organization. Add the PowerUserAccess managed policy to the role. Apply an inline policy to all IAM users and roles in every AWS account to deny permissions on the AWSPrivateMarketplaceAdminFullAccess managed policy.
- B. Create an IAM role named procurement-manager-role in all AWS accounts in the organization. Add the AdministratorAccess managed policy to the role. Define a permissions boundary with the AWSPrivateMarketplaceAdminFullAccess managed policy and attach it to all the developer roles.
- C. Create an IAM role named procurement-manager-role in all the shared services accounts in the organization. Add the AWSPrivateMarketplaceAdminFullAccess managed policy to the role. Create an organization root-level SCP to deny permissions to administer Private Marketplace to everyone except the role named procurement-manager-role. Create another organization root-level SCP to deny permissions to create an IAM role named procurement-manager-role to everyone in the organization.
- D. Create an IAM role named procurement-manager-role in all AWS accounts that will be used by developer.
- E. Add the AWSPrivateMarketplaceAdminFullAccess managed policy to the role.
- F. Create an SCP in Organizations to deny permissions to administer Private Marketplace to everyone except the role named procurement-manager-role.
- G. Apply the SCP to all the shared services accounts in the organization.

Answer: C

Explanation:

SCP to deny permissions to administer Private Marketplace to everyone except the role named procurement-manager-role.

<https://aws.amazon.com/blogs/awsmarketplace/controlling-access-to-a-well-architected-private-marketplace-usi>

This approach allows the procurement managers to assume the procurement-manager-role in shared services accounts, which have the AWSPrivateMarketplaceAdminFullAccess managed policy attached to it and can then manage the Private Marketplace. The organization root-level SCP denies the permission to administer Private Marketplace to everyone except the role named procurement-manager-role and another SCP denies the permission to create an IAM role named procurement-manager-role to everyone in the organization, ensuring that only the procurement team can assume the role and manage the Private Marketplace. This approach provides a centralized way to manage and restrict access to Private Marketplace while maintaining a high level of security.

NEW QUESTION 19

- (Exam Topic 1)

A company gives users the ability to upload images from a custom application. The upload process invokes an AWS Lambda function that processes and stores the image in an Amazon S3 bucket. The application invokes the Lambda function by using a specific function version ARN.

The Lambda function accepts image processing parameters by using environment variables. The company often adjusts the environment variables of the Lambda function to achieve optimal image processing output. The company tests different parameters and publishes a new function version with the updated environment variables after validating results. This update process also requires frequent changes to the custom application to invoke the new function version ARN. These changes cause interruptions for users.

A solutions architect needs to simplify this process to minimize disruption to users. Which solution will meet these requirements with the LEAST operational overhead?

- A. Directly modify the environment variables of the published Lambda function version.
- B. Use the LATEST version to test image processing parameters.
- C. Create an Amazon DynamoDB table to store the image processing parameter.
- D. Modify the Lambda function to retrieve the image processing parameters from the DynamoDB table.
- E. Directly code the image processing parameters within the Lambda function and remove the environment variable.
- F. Publish a new function version when the company updates the parameters.
- G. Create a Lambda function alias.
- H. Modify the client application to use the function alias ARN.
- I. Reconfigure the Lambda alias to point to new versions of the function when the company finishes testing.

Answer: D

Explanation:

A Lambda function alias allows you to point to a specific version of a function and also can be updated to point to a new version of the function without modifying the client application. This way, the company can test different versions of the function with different environment variables and, once the optimal parameters are found, update the alias to point to the new version, without the need to update the client application.

By using this approach, the company can simplify the process of updating the environment variables, minimize disruption to users, and reduce the operational overhead.

Reference:

AWS Lambda documentation: <https://aws.amazon.com/lambda/>

AWS Lambda Aliases documentation: <https://docs.aws.amazon.com/lambda/latest/dg/aliases-intro.html> AWS Lambda versioning and aliases documentation:

<https://aws.amazon.com/blogs/compute/versioning-aliases-in-aws-lambda/>

NEW QUESTION 21

- (Exam Topic 1)

A company has developed a web application. The company is hosting the application on a group of Amazon EC2 instances behind an Application Load Balancer.

The company wants to improve the security posture of the application and plans to use AWS WAF web ACLs. The solution must not adversely affect legitimate traffic to the application.

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

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

Answer: A

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/waf-analyze-count-action-rules/>

NEW QUESTION 23

- (Exam Topic 1)

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

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

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

Answer: B

Explanation:

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

NEW QUESTION 25

- (Exam Topic 1)

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

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

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

Answer: B

Explanation:

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

NEW QUESTION 26

- (Exam Topic 1)

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

Which solutions will meet these requirements?

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

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

Answer: C

Explanation:

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

NEW QUESTION 31

- (Exam Topic 1)

A company is in the process of implementing AWS Organizations to constrain its developers to use only Amazon EC2, Amazon S3 and Amazon DynamoDB. The developers account resides in a dedicated organizational unit (OU). The solutions architect has implemented the following SCP on the developers account:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowEC2",
      "Effect": "Allow",
      "Action": "ec2:*",
      "Resource": "*"
    },
    {
      "Sid": "AllowDynamoDB",
      "Effect": "Allow",
      "Action": "dynamodb:*",
      "Resource": "*"
    },
    {
      "Sid": "AllowS3",
      "Effect": "Allow",
      "Action": "s3:*",
      "Resource": "*"
    }
  ]
}
```

When this policy is deployed, IAM users in the developers account are still able to use AWS services that are not listed in the policy. What should the solutions architect do to eliminate the developers' ability to use services outside the scope of this policy?

- A. Create an explicit deny statement for each AWS service that should be constrained
- B. Remove the Full AWS Access SCP from the developer account's OU
- C. Modify the Full AWS Access SCP to explicitly deny all services
- D. Add an explicit deny statement using a wildcard to the end of the SCP

Answer: B

Explanation:

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

NEW QUESTION 36

- (Exam Topic 1)

A company is using AWS Organizations to manage multiple AWS accounts. For security purposes, the company requires the creation of an Amazon Simple Notification Service (Amazon SNS) topic that enables integration with a third-party alerting system in all the Organizations member accounts.

A solutions architect used an AWS CloudFormation template to create the SNS topic and stack sets to automate the deployment of CloudFormation stacks. Trusted access has been enabled in Organizations.

What should the solutions architect do to deploy the CloudFormation StackSets in all AWS accounts?

- A. Create a stack set in the Organizations member account
- B. Use service-managed permission
- C. Set deployment options to deploy to an organization
- D. Use CloudFormation StackSets drift detection.
- E. Create stacks in the Organizations member account
- F. Use self-service permission

- G. Set deployment options to deploy to an organizatio
- H. Enable the CloudFormation StackSets automatic deployment.
- I. Create a stack set in the Organizations management account Use service-managed permission
- J. Set deployment options to deploy to the organizatio
- K. Enable CloudFormation StackSets automatic deployment.
- L. Create stacks in the Organizations management accoun
- M. Use service-managed permission
- N. Set deployment options to deploy to the organizatio
- O. Enable CloudFormation StackSets drift detection.

Answer: C

Explanation:

<https://aws.amazon.com/blogs/aws/use-cloudformation-stacksets-to-provision-resources-across-multiple-aws-ac>

NEW QUESTION 41

- (Exam Topic 1)

A company uses an on-premises data analytics platform. The system is highly available in a fully redundant configuration across 12 servers in the company's data center.

The system runs scheduled jobs, both hourly and daily, in addition to one-time requests from users. Scheduled jobs can take between 20 minutes and 2 hours to finish running and have tight SLAs. The scheduled jobs account for 65% of the system usage. User jobs typically finish running in less than 5 minutes and have no SLA. The user jobs account for 35% of system usage. During system failures, scheduled jobs must continue to meet SLAs. However, user jobs can be delayed.

A solutions architect needs to move the system to Amazon EC2 instances and adopt a consumption-based model to reduce costs with no long-term commitments.

The solution must maintain high availability and must not affect the SLAs.

Which solution will meet these requirements MOST cost-effectively?

- A. Split the 12 instances across two Availability Zones in the chosen AWS Regio
- B. Run two instances in each Availability Zone as On-Demand Instances with Capacity Reservation
- C. Run four instances in each Availability Zone as Spot Instances.
- D. Split the 12 instances across three Availability Zones in the chosen AWS Regio
- E. In one of the Availability Zones, run all four instances as On-Demand Instances with Capacity Reservation
- F. Run the remaining instances as Spot Instances.
- G. Split the 12 instances across three Availability Zones in the chosen AWS Regio
- H. Run two instances in each Availability Zone as On-Demand Instances with a Savings Pla
- I. Run two instances in each Availability Zone as Spot Instances.
- J. Split the 12 instances across three Availability Zones in the chosen AWS Regio
- K. Run three instances in each Availability Zone as On-Demand Instances with Capacity Reservation
- L. Run one instance in each Availability Zone as a Spot Instance.

Answer: D

Explanation:

By splitting the 12 instances across three Availability Zones, the system can maintain high availability and availability of resources in case of a failure. Option D also uses a combination of On-Demand Instances with Capacity Reservations and Spot Instances, which allows for scheduled jobs to be run on the On-Demand instances with guaranteed capacity, while also taking advantage of the cost savings from Spot Instances for the user jobs which have lower SLA requirements.

NEW QUESTION 42

- (Exam Topic 1)

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

Which solution will meet these requirements?

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

Answer: C

Explanation:

it's the one that handles failover while B (the one shown as the answer today) it almost the same but does not handle failover.

NEW QUESTION 46

- (Exam Topic 1)

A company is running a web application in the AWS Cloud. The application consists of dynamic content that is created on a set of Amazon EC2 instances. The EC2 instances run in an Auto Scaling group that is configured as a target group for an Application Load Balancer (ALB).

The company is using an Amazon CloudFront distribution to distribute the application globally. The CloudFront distribution uses the ALB as an origin. The

company uses Amazon Route 53 for DNS and has created an A record of www.example.com for the CloudFront distribution. A solutions architect must configure the application so that it is highly available and fault tolerant. Which solution meets these requirements?

- A. Provision a full, secondary application deployment in a different AWS Region
- B. Update the Route 53 A record to be a failover record
- C. Add both of the CloudFront distributions as values
- D. Create Route 53 health checks.
- E. Provision an ALB, an Auto Scaling group, and EC2 instances in a different AWS Region
- F. Update the CloudFront distribution, and create a second origin for the new AL
- G. Create an origin group for the two origins
- H. Configure one origin as primary and one origin as secondary.
- I. Provision an Auto Scaling group and EC2 instances in a different AWS Region
- J. Create a second target for the new Auto Scaling group in the AL
- K. Set up the failover routing algorithm on the ALB.
- L. Provision a full, secondary application deployment in a different AWS Region
- M. Create a second CloudFront distribution, and add the new application setup as an origin
- N. Create an AWS Global Accelerator accelerator
- O. Add both of the CloudFront distributions as endpoints.

Answer: B

Explanation:

[https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/DownloadDistS3AndCustomOrigins.h](https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/DownloadDistS3AndCustomOrigins.html)

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

You can set up CloudFront with origin failover for scenarios that require high availability. To get started, you create an origin group with two origins: a primary and a secondary. If the primary origin is unavailable, or returns specific HTTP response status codes that indicate a failure, CloudFront automatically switches to the secondary origin.

NEW QUESTION 49

- (Exam Topic 1)

A company that has multiple AWS accounts is using AWS Organizations. The company's AWS accounts host VPCs, Amazon EC2 instances, and containers. The company's compliance team has deployed a security tool in each VPC where the company has deployments. The security tools run on EC2 instances and send information to the AWS account that is dedicated for the compliance team. The company has tagged all the compliance-related resources with a key of "costCenter" and a value of "compliance".

The company wants to identify the cost of the security tools that are running on the EC2 instances so that the company can charge the compliance team's AWS account. The cost calculation must be as accurate as possible.

What should a solutions architect do to meet these requirements?

- A. In the management account of the organization, activate the costCenter user-defined tag
- B. Configure monthly AWS Cost and Usage Reports to save to an Amazon S3 bucket in the management account
- C. Use the tag breakdown in the report to obtain the total cost for the costCenter tagged resources.
- D. In the member accounts of the organization, activate the costCenter user-defined tag
- E. Configure monthly AWS Cost and Usage Reports to save to an Amazon S3 bucket in the management account
- F. Schedule a monthly AWS Lambda function to retrieve the reports and calculate the total cost for the costCenter tagged resources.
- G. In the member accounts of the organization activate the costCenter user-defined tag
- H. From the management account, schedule a monthly AWS Cost and Usage Report
- I. Use the tag breakdown in the report to calculate the total cost for the costCenter tagged resources.
- J. Create a custom report in the organization view in AWS Trusted Advisor
- K. Configure the report to generate a monthly billing summary for the costCenter tagged resources in the compliance team's AWS account.

Answer: A

Explanation:

<https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/custom-tags.html>

<https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/configurecostallocreport.html>

NEW QUESTION 50

- (Exam Topic 1)

A company has purchased appliances from different vendors. The appliances all have IoT sensors. The sensors send status information in the vendors' proprietary formats to a legacy application that parses the information into JSON. The parsing is simple, but each vendor has a unique format. Once daily, the application parses all the JSON records and stores the records in a relational database for analysis.

The company needs to design a new data analysis solution that can deliver faster and optimize costs. Which solution will meet these requirements?

- A. Connect the IoT sensors to AWS IoT Core
- B. Set a rule to invoke an AWS Lambda function to parse the information and save a .csv file to Amazon S3. Use AWS Glue to catalog the file
- C. Use Amazon Athena and Amazon QuickSight for analysis.
- D. Migrate the application server to AWS Fargate, which will receive the information from IoT sensors and parse the information into a relational format
- E. Save the parsed information to Amazon Redshift for analysis.
- F. Create an AWS Transfer for SFTP server
- G. Update the IoT sensor code to send the information as a .csv file through SFTP to the server
- H. Use AWS Glue to catalog the file
- I. Use Amazon Athena for analysis.
- J. Use AWS Snowball Edge to collect data from the IoT sensors directly to perform local analysis. Periodically collect the data into Amazon Redshift to perform global analysis.

Answer: A

Explanation:

➤ Connect the IoT sensors to AWS IoT Core. Set a rule to invoke an AWS Lambda function to parse the information and save a .csv file to Amazon S3. Use AWS Glue to catalog the files. Use Amazon Athena and Amazon QuickSight for analysis. This solution meets the requirement of faster analysis and cost

optimization by using AWS IoT Core to collect data from the IoT sensors in real-time and then using AWS Glue and Amazon Athena for efficient data analysis. This solution involves connecting the IoT sensors to the AWS IoT Core, setting a rule to invoke an AWS Lambda function to parse the information, and saving a .csv file to Amazon S3. AWS Glue can be used to catalog the files and Amazon Athena and Amazon QuickSight can be used for analysis. This solution will enable faster and more cost-effective data analysis.

This solution is in line with the official Amazon Textbook and Resources for the AWS Certified Solutions Architect - Professional certification. In particular, the book states that: "AWS IoT Core can be used to ingest and process the data, AWS Lambda can be used to process and transform the data, and Amazon S3 can be used to store the data. AWS Glue can be used to catalog and access the data, Amazon Athena can be used to query the data, and Amazon QuickSight can be used to visualize the data." (Source: https://d1.awsstatic.com/training-and-certification/docs-sa-pro/AWS_Certified_Solutions_Architect_Professiona)

NEW QUESTION 53

- (Exam Topic 1)

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

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

Answer: A

Explanation:

<https://aws.amazon.com/blogs/compute/implementing-canary-deployments-of-aws-lambda-functions-with-alias->
<https://docs.aws.amazon.com/lambda/latest/dg/configuration-aliases.html>

NEW QUESTION 58

- (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 premise
- 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 62

- (Exam Topic 1)

A retail company has an on-premises data center in Europe. The company also has a multi-Region AWS presence that includes the eu-west-1 and us-east-1 Regions. The company wants to be able to route network traffic from its on-premises infrastructure into VPCs in either of those Regions. The company also needs to support traffic that is routed directly between VPCs in those Regions. No single points of failure can exist on the network.

The company already has created two 1 Gbps AWS Direct Connect connections from its on-premises data center. Each connection goes into a separate Direct Connect location in Europe for high availability. These two locations are named DX-A and DX-B, respectively. Each Region has a single AWS Transit Gateway that is configured to route all inter-VPC traffic within that Region.

Which solution will meet these requirements?

- A. Create a private VIF from the DX-A connection into a Direct Connect gateway
- B. Create a private VIF from the DX-B connection into the same Direct Connect gateway for high availability
- C. Associate both the eu-west-1 and us-east-1 transit gateways with the Direct Connect gateway
- D. Peer the transit gateways with each other to support cross-Region routing.
- E. Create a transit VIF from the DX-A connection into a Direct Connect gateway
- F. Associate the eu-west-1 transit gateway with this Direct Connect gateway
- G. Create a transit VIF from the DX-B connection into a separate Direct Connect gateway
- H. Associate the us-east-1 transit gateway with this separate Direct Connect gateway
- I. Peer the Direct Connect gateways with each other to support high availability and cross-Region routing.
- J. Create a transit VIF from the DX-A connection into a Direct Connect gateway
- K. Create a transit VIF from the DX-B connection into the same Direct Connect gateway for high availability
- L. Associate both the eu-west-1 and us-east-1 transit gateways with this Direct Connect gateway
- M. Configure the Direct Connect gateway to route traffic between the transit gateways.
- N. Create a transit VIF from the DX-A connection into a Direct Connect gateway

- O. Create a transit VIF from the DX-B connection into the same Direct Connect gateway for high availability
- P. Associate both the eu-west-1 and us-east-1 transit gateways with this Direct Connect gateway
- Q. Peer the transit gateways with each other to support cross-Region routing.

Answer: D

Explanation:

in this solution, two transit VIFs are created - one from the DX-A connection and one from the DX-B connection - into the same Direct Connect gateway for high availability. Both the eu-west-1 and us-east-1 transit gateways are then associated with this Direct Connect gateway. The transit gateways are then peered with each other to support cross-Region routing. This solution meets the requirements of the company by creating a highly available connection between the on-premises data center and the VPCs in both the eu-west-1 and us-east-1 regions, and by enabling direct traffic routing between VPCs in those regions.

NEW QUESTION 63

- (Exam Topic 1)

A company has a web application that allows users to upload short videos. The videos are stored on Amazon EBS volumes and analyzed by custom recognition software for categorization.

The website contains static content that has variable traffic with peaks in certain months. The architecture consists of Amazon EC2 instances running in an Auto Scaling group for the web application and EC2

instances running in an Auto Scaling group to process an Amazon SQS queue. The company wants to

re-architect the application to reduce operational overhead using AWS managed services where possible and remove dependencies on third-party software.

Which solution meets these requirements?

- A. Use Amazon ECS containers for the web application and Spot Instances for the Auto Scaling group that processes the SQS queue
- B. Replace the custom software with Amazon Rekognition to categorize the videos.
- C. Store the uploaded videos in Amazon EFS and mount the file system to the EC2 instances for the web application
- D. Process the SQS queue with an AWS Lambda function that calls the Amazon Rekognition API to categorize the videos.
- E. Host the web application in Amazon S3. Store the uploaded videos in Amazon S3. Use S3 event notifications to publish events to the SQS queue. Process the SQS queue with an AWS Lambda function that calls the Amazon Rekognition API to categorize the videos.
- F. Use AWS Elastic Beanstalk to launch EC2 instances in an Auto Scaling group for the web application and launch a worker environment to process the SQS queue. Replace the custom software with Amazon Rekognition to categorize the videos.

Answer: C

Explanation:

➤ Option C is correct because hosting the web application in Amazon S3, storing the uploaded videos in Amazon S3, and using S3 event notifications to publish events to the SQS queue reduces the operational overhead of managing EC2 instances and EBS volumes. Amazon S3 can serve static content such as HTML, CSS, JavaScript, and media files directly from S3 buckets. Amazon S3 can also trigger AWS Lambda functions through S3 event notifications when new objects are created or existing objects are updated or deleted. AWS Lambda can process the SQS queue with an AWS Lambda function that calls the Amazon Rekognition API to categorize the videos. This solution eliminates the need for custom recognition software and third-party dependencies.

References: 1: <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-spot-instances.html> 2:

<https://aws.amazon.com/efs/pricing/> 3:

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/WebsiteHosting.html> 4:

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/NotificationHowTo.html> 5:

<https://docs.aws.amazon.com/rekognition/latest/dg/what-is.html> 6: <https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/Welcome.html>

NEW QUESTION 68

- (Exam Topic 1)

A company is processing videos in the AWS Cloud by using Amazon EC2 instances in an Auto Scaling group. It takes 30 minutes to process a video. Several EC2 instances scale in and out depending on the number of videos in an Amazon Simple Queue Service (Amazon SQS) queue.

The company has configured the SQS queue with a redrive policy that specifies a target dead-letter queue and a maxReceiveCount of 1. The company has set the visibility timeout for the SQS queue to 1 hour. The company has set up an Amazon CloudWatch alarm to notify the development team when there are messages in the dead-letter queue.

Several times during the day, the development team receives notification that messages are in the dead-letter queue and that videos have not been processed properly. An investigation finds no errors in the application logs.

How can the company solve this problem?

- A. Turn on termination protection for the EC2 instances.
- B. Update the visibility timeout for the SQS queue to 3 hours.
- C. Configure scale-in protection for the instances during processing.
- D. Update the redrive policy and set maxReceiveCount to 0.

Answer: B

Explanation:

The best solution for this problem is to update the visibility timeout for the SQS queue to 3 hours. This is because when the visibility timeout is set to 1 hour, it means that if the EC2 instance doesn't process the message within an hour, it will be moved to the dead-letter queue. By increasing the visibility timeout to 3 hours, this should give the EC2 instance enough time to process the message before it gets moved to the dead-letter queue. Additionally, configuring scale-in protection for the EC2 instances during processing will help to ensure that the instances are not terminated while the messages are being processed.

NEW QUESTION 71

- (Exam Topic 1)

A company is planning to host a web application on AWS and works to load balance the traffic across a group of Amazon EC2 instances. One of the security requirements is to enable end-to-end encryption in transit between the client and the web server.

Which solution will meet this requirement?

- A. Place the EC2 instances behind an Application Load Balancer (ALB). Provision an SSL certificate using AWS Certificate Manager (ACM), and associate the SSL certificate with the ALB.
- B. Export the SSL certificate and install it on each EC2 instance.
- C. Configure the ALB to listen on port 443 and to forward traffic to port 443 on the instances.

- D. Associate the EC2 instances with a target group
- E. Provision an SSL certificate using AWS Certificate Manager (ACM). Create an Amazon CloudFront distribution and configure it to use the SSL certificate
- F. Set CloudFront to use the target group as the origin server
- G. Place the EC2 instances behind an Application Load Balancer (ALB). Provision an SSL certificate using AWS Certificate Manager (ACM), and associate the SSL certificate with the ALB
- H. Provision a third-party SSL certificate and install it on each EC2 instance
- I. Configure the ALB to listen on port 443 and to forward traffic to port 443 on the instances.
- J. Place the EC2 instances behind a Network Load Balancer (NLB). Provision a third-party SSL certificate and install it on the NLB and on each EC2 instance
- K. Configure the NLB to listen on port 443 and to forward traffic to port 443 on the instances.

Answer: A

Explanation:

➤ Option A is correct because placing the EC2 instances behind an Application Load Balancer (ALB) and associating an SSL certificate from AWS Certificate Manager (ACM) with the ALB enables encryption in transit between the client and the ALB. Exporting the SSL certificate and installing it on each EC2 instance enables encryption in transit between the ALB and the web server. Configuring the ALB to listen on port 443 and to forward traffic to port 443 on the instances ensures that HTTPS is used for both connections. This solution achieves end-to-end encryption in transit for the web application.

References: 1: <https://docs.aws.amazon.com/elasticloadbalancing/latest/application/introduction.html> 2: <https://docs.aws.amazon.com/acm/latest/userguide/acm-overview.html> 3: <https://docs.aws.amazon.com/elasticloadbalancing/latest/application/load-balancer-target-groups.html> : <https://aws.amazon.com/certificate-manager/faqs/> : <https://docs.aws.amazon.com/elasticloadbalancing/latest/network/introduction.html>

NEW QUESTION 72

- (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 76

- (Exam Topic 1)

A large company is running a popular web application. The application runs on several Amazon EC2 Linux Instances in an Auto Scaling group in a private subnet. An Application Load Balancer is targeting the Instances in the Auto Scaling group in the private subnet. AWS Systems Manager Session Manager is configured, and AWS Systems Manager Agent is running on all the EC2 instances. The company recently released a new version of the application. Some EC2 instances are now being marked as unhealthy and are being terminated. As a result, the application is running at reduced capacity. A solutions architect tries to determine the root cause by analyzing Amazon CloudWatch logs that are collected from the application, but the logs are inconclusive. How should the solutions architect gain access to an EC2 instance to troubleshoot the issue?

- A. Suspend the Auto Scaling group's HealthCheck scaling process
- B. Use Session Manager to log in to an instance that is marked as unhealthy
- C. Enable EC2 instance termination protection. Use Session Manager to log in to an instance that is marked as unhealthy.
- D. Set the termination policy to OldestInstance on the Auto Scaling group
- E. Use Session Manager to log in to an instance that is marked as unhealthy
- F. Suspend the Auto Scaling group's Terminate process
- G. Use Session Manager to log in to an instance that is marked as unhealthy

Answer: D

Explanation:

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

NEW QUESTION 81

- (Exam Topic 1)

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

- A. Create an API Gateway endpoint in the us-west-2 Region to direct traffic to the Lambda function in us-east-1. Configure Amazon Route 53 to use a failover routing policy to route traffic for the two API Gateway endpoints.
- B. Create an Amazon Simple Queue Service (Amazon SQS) queue

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

Answer: B

Explanation:

This solution allows for deploying the Lambda function and API Gateway endpoint to another region, providing a failover option in case of any issues in the primary region. Using Route 53's failover routing policy allows for automatic routing of traffic to the healthy endpoint, ensuring that the application is available even in case of issues in one region. This solution provides a cost-effective and simple way to implement failover while minimizing operational overhead.

NEW QUESTION 85

- (Exam Topic 1)

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

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

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

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

Answer: B

Explanation:

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

NEW QUESTION 89

- (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-service-role.html>

NEW QUESTION 93

- (Exam Topic 1)

A company needs to implement a patching process for its servers. The on-premises servers and Amazon EC2 instances use a variety of tools to perform patching. Management requires a single report showing the patch status of all the servers and instances.

Which set of actions should a solutions architect take to meet these requirements?

- A. Use AWS Systems Manager to manage patches on the on-premises servers and EC2 instances
- B. Use Systems Manager to generate patch compliance reports.
- C. Use AWS OpsWorks to manage patches on the on-premises servers and EC2 instances
- D. Use Amazon QuickSight integration with OpsWorks to generate patch compliance reports.
- E. Use an Amazon EventBridge (Amazon CloudWatch Events) rule to apply patches by scheduling an AWS Systems Manager patch remediation job
- F. Use Amazon Inspector to generate patch compliance reports.
- G. Use AWS OpsWorks to manage patches on the on-premises servers and EC2 instances
- H. Use AWS X-Ray to post the patch status to AWS Systems Manager OpsCenter to generate patch compliance reports.

Answer: A

Explanation:

<https://docs.aws.amazon.com/systems-manager/latest/userguide/systems-manager-patch.html>

NEW QUESTION 96

- (Exam Topic 1)

The company needs to determine which costs on the monthly AWS bill are attributable to each application or team. The company also must be able to create reports to compare costs from the last 12 months and to help forecast costs for the next 12 months. A solutions architect must recommend an AWS Billing and Cost Management solution that provides these cost reports.

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

- A. Activate the user-defined cost allocation tags that represent the application and the team.
- B. Activate the AWS generated cost allocation tags that represent the application and the team.
- C. Create a cost category for each application in Billing and Cost Management.
- D. Activate IAM access to Billing and Cost Management.
- E. Create a cost budget.
- F. Enable Cost Explorer.

Answer: ACF

Explanation:

<https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/manage-cost-categories.html> <https://aws.amazon.com/premiumsupport/knowledge-center/cost-explorer-analyze-spending-and-usage/> <https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/manage-cost-categories.html>
<https://docs.aws.amazon.com/cost-management/latest/userguide/ce-enable.html>

The best combination of actions to meet the company's requirements is Options A, C, and F.

Option A involves activating the user-defined cost allocation tags that represent the application and the team. This will allow the company to assign costs to different applications or teams, and will allow them to be tracked in the monthly AWS bill.

Option C involves creating a cost category for each application in Billing and Cost Management. This will allow the company to easily identify and compare costs across different applications and teams.

Option F involves enabling Cost Explorer. This will allow the company to view the costs of their AWS resources over the last 12 months and to create forecasts for the next 12 months.

These recommendations are in line with the official Amazon Textbook and Resources for the AWS Certified Solutions Architect - Professional certification. In particular, the book states that "You can use cost allocation tags to group your costs by application, team, or other categories" (Source:

https://d1.awsstatic.com/training-and-certification/docs-sa-pro/AWS_Certified_Solutions_Architect_Professiona Additionally, the book states that "Cost Explorer enables you to view the costs of your AWS resources over the last 12 months and to create forecasts for the next 12 months" (Source:

https://d1.awsstatic.com/training-and-certification/docs-sa-pro/AWS_Certified_Solutions_Architect_Professiona

NEW QUESTION 98

- (Exam Topic 1)

A company has hundreds of AWS accounts. The company recently implemented a centralized internal process for purchasing new Reserved Instances and modifying existing Reserved Instances. This process requires all business units that want to purchase or modify Reserved Instances to submit requests to a dedicated team for procurement. Previously, business units directly purchased or modified Reserved Instances in their own respective AWS accounts autonomously.

A solutions architect needs to enforce the new process in the most secure way possible.

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

- A. Ensure that all AWS accounts are part of an organization in AWS Organizations with all features enabled.
- B. Use AWS Config to report on the attachment of an IAM policy that denies access to the ec2:PurchaseReservedInstancesOffering action and the ec2:ModifyReservedInstances action.
- C. In each AWS account, create an IAM policy that denies the ec2:PurchaseReservedInstancesOffering action and the ec2:ModifyReservedInstances action.
- D. Create an SCP that denies the ec2:PurchaseReservedInstancesOffering action and the ec2:ModifyReservedInstances action.
- E. Attach the SCP to each OU of the organization.
- F. Ensure that all AWS accounts are part of an organization in AWS Organizations that uses the consolidated billing feature.

Answer: AD

Explanation:

All features – The default feature set that is available to AWS Organizations. It includes all the functionality of consolidated billing, plus advanced features that give you more control over accounts in your organization. For example, when all features are enabled the management account of the organization has full control over what member accounts can do. The management account can apply SCPs to restrict the services and actions that users (including the root user) and roles in an account can access. https://docs.aws.amazon.com/organizations/latest/userguide/orgs_getting-started_concepts.html#feature-set

NEW QUESTION 102

- (Exam Topic 1)

A company has a legacy monolithic application that is critical to the company's business. The company hosts the application on an Amazon EC2 instance that runs Amazon Linux 2. The company's application team receives a directive from the legal department to back up the data from the instance's encrypted Amazon Elastic Block Store (Amazon EBS) volume to an Amazon S3 bucket. The application team does not have the administrative SSH key pair for the instance. The application must continue to serve the users.

Which solution will meet these requirements?

- A. Attach a role to the instance with permission to write to Amazon S3. Use the AWS Systems Manager Session Manager option to gain access to the instance and run commands to copy data into Amazon S3.
- B. Create an image of the instance with the reboot option turned on.
- C. Launch a new EC2 instance from the image.
- D. Attach a role to the new instance with permission to write to Amazon S3. Run a command to copy data into Amazon S3.
- E. Take a snapshot of the EBS volume by using Amazon Data Lifecycle Manager (Amazon DLM). Copy the data to Amazon S3.
- F. Create an image of the instance.
- G. Launch a new EC2 instance from the image.
- H. Attach a role to the new instance with permission to write to Amazon S3. Run a command to copy data into Amazon S3.

Answer: C

Explanation:

Taking a snapshot of the EBS volume using Amazon Data Lifecycle Manager (DLM) will meet the requirements because it allows you to create a backup of the volume without the need to access the instance or its SSH key pair. Additionally, DLM allows you to schedule the backups to occur at specific intervals and also enables you to copy the snapshots to an S3 bucket. This approach will not impact the running application as the backup is performed on the EBS volume level.

NEW QUESTION 107

- (Exam Topic 1)

A company used Amazon EC2 instances to deploy a web fleet to host a blog site. The EC2 instances are behind an Application Load Balancer (ALB) and are configured in an Auto Scaling group. The web application stores all blog content on an Amazon EFS volume.

The company recently added a feature for bloggers to add video to their posts, attracting 10 times the previous user traffic. At peak times of day, users report buffering and timeout issues while attempting to reach the site or watch videos.

Which is the MOST cost-efficient and scalable deployment that will resolve the issues for users?

- A. Reconfigure Amazon EFS to enable maximum I/O.
- B. Update the blog site to use instance store volumes for storage.
- C. Copy the site contents to the volumes at launch and to Amazon S3 at shutdown.
- D. Configure an Amazon CloudFront distribution.
- E. Point the distribution to an S3 bucket, and migrate the videos from EFS to Amazon S3.
- F. Set up an Amazon CloudFront distribution for all site contents, and point the distribution at the ALB.

Answer: C

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/cloudfront-https-connection-fails/> Using an Amazon S3 bucket

Using a MediaStore container or a MediaPackage channel Using an Application Load Balancer

Using a Lambda function URL

Using Amazon EC2 (or another custom origin)

Using CloudFront origin groups <https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/restrict-access-to-load-balancer.html>

NEW QUESTION 108

- (Exam Topic 1)

A company has an organization in AWS Organizations. The company is using AWS Control Tower to deploy a landing zone for the organization. The company wants to implement governance and policy enforcement. The company must implement a policy that will detect Amazon RDS DB instances that are not encrypted at rest in the company's production OU.

Which solution will meet this requirement?

- A. Turn on mandatory guardrails in AWS Control Tower.
- B. Apply the mandatory guardrails to the production OU.
- C. Enable the appropriate guardrail from the list of strongly recommended guardrails in AWS Control Tower.
- D. Apply the guardrail to the production OU.
- E. Use AWS Config to create a new mandatory guardrail.
- F. Apply the rule to all accounts in the production OU.
- G. Create a custom SCP in AWS Control Tower.
- H. Apply the SCP to the production OU.

Answer: B

Explanation:

AWS Control Tower provides a set of "strongly recommended guardrails" that can be enabled to implement governance and policy enforcement. One of these guardrails is "Encrypt Amazon RDS instances" which will detect RDS DB instances that are not encrypted at rest. By enabling this guardrail and applying it to the production OU, the company will be able to enforce encryption for RDS instances in the production environment.

NEW QUESTION 113

- (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 DocumentDB.

<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 116

- (Exam Topic 1)

A company is storing data in several Amazon DynamoDB tables. A solutions architect must use a serverless architecture to make the data accessible publicly through a simple API over HTTPS. The solution must scale automatically in response to demand.

Which solutions meet these requirements? (Choose two.)

- A. Create an Amazon API Gateway REST AP
- B. Configure this API with direct integrations to DynamoDB by using API Gateway's AWS integration type.
- C. Create an Amazon API Gateway HTTP AP
- D. Configure this API with direct integrations to Dynamo DB by using API Gateway's AWS integration type.
- E. Create an Amazon API Gateway HTTP AP
- F. Configure this API with integrations to AWS Lambda functions that return data from the DynamoDB tables.
- G. Create an accelerator in AWS Global Accelerato
- H. Configure this accelerator with AWS Lambda@Edge function integrations that return data from the DynamoDB tables.
- I. Create a Network Load Balance
- J. Configure listener rules to forward requests to the appropriate AWS Lambda functions

Answer: AC

Explanation:

<https://docs.aws.amazon.com/apigateway/latest/developerguide/api-gateway-overview-developer-experience.htm>

NEW QUESTION 118

- (Exam Topic 1)

A company runs a proprietary stateless ETL application on an Amazon EC2 Linux instance. The application is a Linux binary, and the source code cannot be modified. The application is single-threaded, uses 2 GB of RAM. and is highly CPU intensive The application is scheduled to run every 4 hours and runs for up to 20 minutes A solutions architect wants to revise the architecture for the solution.

Which strategy should the solutions architect use?

- A. Use AWS Lambda to run the applicatio
- B. Use Amazon CloudWatch Logs to invoke the Lambda function every 4 hours.
- C. Use AWS Batch to run the applicatio
- D. Use an AWS Step Functions state machine to invoke the AWS Batch job every 4 hours.
- E. Use AWS Fargate to run the applicatio
- F. Use Amazon EventBridge (Amazon CloudWatch Events) to invoke the Fargate task every 4 hours.
- G. Use Amazon EC2 Spot Instances to run the applicatio
- H. Use AWS CodeDeploy to deploy and run the application every 4 hours.

Answer: C

Explanation:

step function could run a scheduled task when triggered by eventbrige, but why would you add that layer of complexity just to run aws batch when you could directly invoke it through eventbridge. The link provided - <https://aws.amazon.com/pt/blogs/compute/orchestrating-high-performance-computing-with-aws-step-functions-> makes sense only for HPC, this is a single instance that needs to be run

NEW QUESTION 121

- (Exam Topic 1)

A company runs a new application as a static website in Amazon S3. The company has deployed the application to a production AWS account and uses Amazon CloudFront to deliver the website. The website calls an Amazon API Gateway REST API. An AWS Lambda function backs each API method.

The company wants to create a CSV report every 2 weeks to show each API Lambda function's recommended configured memory, recommended cost, and the price difference between current configurations and the recommendations. The company will store the reports in an S3 bucket.

Which solution will meet these requirements with the LEAST development time?

- A. Create a Lambda function that extracts metrics data for each API Lambda function from Amazon CloudWatch Logs for the 2-week period_ Collate the data into tabular forma
- B. Store the data as a _csvfile in an S3 bucke
- C. Create an Amazon Eventaridge rule to schedule the Lambda function to run every 2 weeks.
- D. Opt in to AWS Compute Optimize
- E. Create a Lambda function that calls the ExportLambdaFunctionRecommendatlons operatio
- F. Export the _csv file to an S3 bucke
- G. Create an Amazon Eventaridge rule to schedule the Lambda function to run every 2 weeks.
- H. Opt in to AWS Compute Optimize
- I. Set up enhanced infrastructure metric
- J. Within the Compute Optimizer console, schedule a job to export the Lambda recommendations to a _csvfile_ Store the file in an S3 bucket every 2 weeks.
- K. Purchase the AWS Business Support plan for the production accoun
- L. Opt in to AWS Compute Optimizer for AWS Trusted Advisor check
- M. In the Trusted Advisor console, schedule a job to export the cost optimization checks to a _csvfile_ Store the file in an S3 bucket every 2 weeks.

Answer: B

Explanation:

https://docs.aws.amazon.com/compute-optimizer/latest/APIReference/API_ExportLambdaFunctionRecommend

NEW QUESTION 123

- (Exam Topic 1)

A company has an organization that has many AWS accounts in AWS Organizations. A solutions architect must improve how the company manages common security group rules for the AWS accounts in the organization.

The company has a common set of IP CIDR ranges in an allow list in each AWS account to allow access to and from the company's on-premises network.

Developers within each account are responsible for adding new IP CIDR ranges to their security groups. The security team has its own AWS account. Currently, the security team notifies the owners of the other AWS accounts when changes are made to the allow list. The solutions architect must design a solution that distributes the common set of CIDR ranges across all accounts. Which solution meets these requirements with the LEAST amount of operational overhead?

- A. Set up an Amazon Simple Notification Service (Amazon SNS) topic in the security team's AWS account
- B. Deploy an AWS Lambda function in each AWS account
- C. Configure the Lambda function to run every time an SNS topic receives a message
- D. Configure the Lambda function to take an IP address as input and add it to a list of security groups in the account
- E. Instruct the security team to distribute changes by publishing messages to its SNS topic.
- F. Create new customer-managed prefix lists in each AWS account within the organization
- G. Populate the prefix lists in each account with all internal CIDR range
- H. Notify the owner of each AWS account to allow the new customer-managed prefix list IDs in their accounts in their security group
- I. Instruct the security team to share updates with each AWS account owner.
- J. Create a new customer-managed prefix list in the security team's AWS account
- K. Populate the customer-managed prefix list with all internal CIDR range
- L. Share the customer-managed prefix list with the organization by using AWS Resource Access Manager
- M. Notify the owner of each AWS account to allow the new customer-managed prefix list ID in their security groups.
- N. Create an IAM role in each account in the organization
- O. Grant permissions to update security groups. Deploy an AWS Lambda function in the security team's AWS account
- P. Configure the Lambda function to take a list of internal IP addresses as input, assume a role in each organization account, and add the list of IP addresses to the security groups in each account.

Answer: C

Explanation:

Create a new customer-managed prefix list in the security team's AWS account. Populate the customer-managed prefix list with all internal CIDR ranges. Share the customer-managed prefix list with the organization by using AWS Resource Access Manager. Notify the owner of each AWS account to allow the new customer-managed prefix list ID in their security groups. This solution meets the requirements with the least amount of operational overhead as it requires the security team to create and maintain a single customer-managed prefix list, and share it with the organization using AWS Resource Access Manager. The owners of each AWS account are then responsible for allowing the prefix list in their security groups, which eliminates the need for the security team to manually notify each account owner when changes are made. This solution also eliminates the need for a separate AWS Lambda function in each account, reducing the overall complexity of the solution.

NEW QUESTION 126

- (Exam Topic 1)

A company recently completed the migration from an on-premises data center to the AWS Cloud by using a replatforming strategy. One of the migrated servers is running a legacy Simple Mail Transfer Protocol (SMTP) service that a critical application relies upon. The application sends outbound email messages to the company's customers. The legacy SMTP server does not support TLS encryption and uses TCP port 25. The application can use SMTP only. The company decides to use Amazon Simple Email Service (Amazon SES) and to decommission the legacy SMTP server. The company has created and validated the SES domain. The company has lifted the SES limits. What should the company do to modify the application to send email messages from Amazon SES?

- A. Configure the application to connect to Amazon SES by using TLS Wrapper
- B. Create an IAM role that has ses:SendEmail and ses:SendRawEmail permission
- C. Attach the IAM role to an Amazon EC2 instance.
- D. Configure the application to connect to Amazon SES by using STARTTLS
- E. Obtain Amazon SES SMTP credential
- F. Use the credentials to authenticate with Amazon SES.
- G. Configure the application to use the SES API to send email message
- H. Create an IAM role that has ses:SendEmail and ses:SendRawEmail permission
- I. Use the IAM role as a service role for Amazon SES.
- J. Configure the application to use AWS SDKs to send email message
- K. Create an IAM user for Amazon SES
- L. Generate API access key
- M. Use the access keys to authenticate with Amazon SES.

Answer: B

Explanation:

To set up a STARTTLS connection, the SMTP client connects to the Amazon SES SMTP endpoint on port 25, 587, or 2587, issues an EHLO command, and waits for the server to announce that it supports the STARTTLS SMTP extension. The client then issues the STARTTLS command, initiating TLS negotiation. When negotiation is complete, the client issues an EHLO command over the new encrypted connection, and the SMTP session proceeds normally. To set up a TLS Wrapper connection, the SMTP client connects to the Amazon SES SMTP endpoint on port 465 or 2465. The server presents its certificate, the client issues an EHLO command, and the SMTP session proceeds normally.
<https://docs.aws.amazon.com/ses/latest/dg/smtp-connect.html>

NEW QUESTION 127

- (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 129

- (Exam Topic 1)

A solutions architect is auditing the security setup of an AWS Lambda function for a company. The Lambda function retrieves the latest changes from an Amazon Aurora database. The Lambda function and the database run in the same VPC. Lambda environment variables are providing the database credentials to the Lambda function.

The Lambda function aggregates data and makes the data available in an Amazon S3 bucket that is configured for server-side encryption with AWS KMS managed encryption keys (SSE-KMS). The data must not travel across the internet. If any database credentials become compromised, the company needs a solution that minimizes the impact of the compromise.

What should the solutions architect recommend to meet these requirements?

- A. Enable IAM database authentication on the Aurora DB cluster
- B. Change the IAM role for the Lambda function to allow the function to access the database by using IAM database authentication
- C. Deploy a gateway VPC endpoint for Amazon S3 in the VPC.
- D. Enable IAM database authentication on the Aurora DB cluster
- E. Change the IAM role for the Lambda function to allow the function to access the database by using IAM database authentication
- F. Enforce HTTPS on the connection to Amazon S3 during data transfers.
- G. Save the database credentials in AWS Systems Manager Parameter Store
- H. Set up password rotation on the credentials in Parameter Store
- I. Change the IAM role for the Lambda function to allow the function to access Parameter Store
- J. Modify the Lambda function to retrieve the credentials from Parameter Store
- K. Deploy a gateway VPC endpoint for Amazon S3 in the VPC.
- L. Save the database credentials in AWS Secrets Manager
- M. Set up password rotation on the credentials in Secrets Manager
- N. Change the IAM role for the Lambda function to allow the function to access Secrets Manager
- O. Modify the Lambda function to retrieve the credentials from Secrets Manager
- P. Enforce HTTPS on the connection to Amazon S3 during data transfers.

Answer: A

Explanation:

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/UsingWithRDS.IAMDBAuth.html>

NEW QUESTION 130

- (Exam Topic 1)

A solutions architect is investigating an issue in which a company cannot establish new sessions in Amazon Workspaces. An initial analysis indicates that the issue involves user profiles. The Amazon Workspaces environment is configured to use Amazon FSx for Windows File Server as the profile share storage. The FSx for Windows File Server file system is configured with 10 TB of storage.

The solutions architect discovers that the file system has reached its maximum capacity. The solutions architect must ensure that users can regain access. The solution also must prevent the problem from occurring again.

Which solution will meet these requirements?

- A. Remove old user profiles to create space
- B. Migrate the user profiles to an Amazon FSx for Lustre file system.
- C. Increase capacity by using the update-file-system command
- D. Implement an Amazon CloudWatch metric that monitors free space
- E. Use Amazon EventBridge to invoke an AWS Lambda function to increase capacity as required.
- F. Monitor the file system by using the FreeStorageCapacity metric in Amazon CloudWatch
- G. Use AWS Step Functions to increase the capacity as required.
- H. Remove old user profiles to create space
- I. Create an additional FSx for Windows File Server file system. Update the user profile redirection for 50% of the users to use the new file system.

Answer: B

Explanation:

➤ It can prevent the issue from happening again by monitoring the file system with the FreeStorageCapacity metric in Amazon CloudWatch and using Amazon EventBridge to invoke an AWS Lambda function to increase the capacity as required. This ensures that the file system always has enough free space to store user profiles and avoids reaching maximum capacity.

NEW QUESTION 131

- (Exam Topic 2)

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

groups and a single instance type.

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

Which solution will meet this requirement?

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

Answer: B

Explanation:

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

NEW QUESTION 134

- (Exam Topic 2)

A company runs a processing engine in the AWS Cloud. The engine processes environmental data from logistics centers to calculate a sustainability index. The company has millions of devices in logistics centers that are spread across Europe. The devices send information to the processing engine through a RESTful API. The API experiences unpredictable bursts of traffic. The company must implement a solution to process all data that the devices send to the processing engine. Data loss is unacceptable.

Which solution will meet these requirements?

- A. Create an Application Load Balancer (ALB) for the RESTful API. Create an Amazon Simple Queue Service (Amazon SQS) queue. Create a listener and a target group for the ALB. Add the SQS queue as the target. Use a container that runs in Amazon Elastic Container Service (Amazon ECS) with the Fargate launch type to process messages in the queue.
- B. Create an Amazon API Gateway HTTP API that implements the RESTful API. Create an Amazon Simple Queue Service (Amazon SQS) queue. Create an API Gateway service integration with the SQS queue. Create an AWS Lambda function to process messages in the SQS queue.
- C. Create an Amazon API Gateway REST API that implements the RESTful API. Create a fleet of Amazon EC2 instances in an Auto Scaling group. Create an API Gateway Auto Scaling group proxy integration. Use the EC2 instances to process incoming data.
- D. Create an Amazon CloudFront distribution for the RESTful API. Create a data stream in Amazon Kinesis Data Streams. Set the data stream as the origin for the distribution. Create an AWS Lambda function to consume and process data in the data stream.

Answer: A

Explanation:

It will use the ALB to handle the unpredictable bursts of traffic and route it to the SQS queue. The SQS queue will act as a buffer to store incoming data temporarily, and the container running in Amazon ECS with the Fargate launch type will process messages in the queue. This approach will ensure that all data is processed and prevent data loss.

NEW QUESTION 139

- (Exam Topic 2)

A company runs its sales reporting application in an AWS Region in the United States. The application uses an Amazon API Gateway Regional API and AWS Lambda functions to generate on-demand reports from data in an Amazon RDS for MySQL database. The frontend of the application is hosted on Amazon S3 and is accessed by users through an Amazon CloudFront distribution. The company is using Amazon Route 53 as the DNS service for the domain. Route 53 is configured with a simple routing policy to route traffic to the API Gateway API.

In the next 6 months, the company plans to expand operations to Europe. More than 90% of the database traffic is read-only traffic. The company has already deployed an API Gateway API and Lambda functions in the new Region.

A solutions architect must design a solution that minimizes latency for users who download reports. Which solution will meet these requirements?

- A. Use an AWS Database Migration Service (AWS DMS) task with full load to replicate the primary database in the original Region to the database in the new Region.
- B. Change the Route 53 record to latency-based routing to connect to the API Gateway API.
- C. Use an AWS Database Migration Service (AWS DMS) task with full load plus change data capture (CDC) to replicate the primary database in the original Region to the database in the new Region.
- D. Change the Route 53 record to geolocation routing to connect to the API Gateway API.
- E. Configure a cross-Region read replica for the RDS database in the new Region.
- F. Change the Route 53 record to latency-based routing to connect to the API Gateway API.
- G. Configure a cross-Region read replica for the RDS database in the new Region.
- H. Change the Route 53 record to geolocation routing to connect to the API.

Answer: C

Explanation:

The company should configure a cross-Region read replica for the RDS database in the new Region. The company should change the Route 53 record to latency-based routing to connect to the API Gateway API. This solution will meet the requirements because a cross-Region read replica is a feature that enables you to create a MariaDB, MySQL, Oracle, PostgreSQL, or SQL Server read replica in a different Region from the source DB instance. You can use cross-Region read replicas to improve availability and disaster recovery, scale out globally, or migrate an existing database to a new Region¹. By creating a cross-Region read replica for the RDS database in the new Region, the company can have a standby copy of its primary database that can serve read-only traffic from users in Europe. A latency-based routing policy is a feature that enables you to route traffic based on the latency between your users and your resources. You can use latency-based routing to route traffic to the resource that provides the best latency². By changing the Route 53 record to latency-based routing, the company can minimize latency for users who download reports by connecting them to the API Gateway API in the Region that provides the best response time.

The other options are not correct because:

➤ Using AWS Database Migration Service (AWS DMS) to replicate the primary database in the original Region to the database in the new Region would not be as cost-effective or simple as using a cross-Region read replica. AWS DMS is a service that enables you to migrate relational databases, data warehouses, NoSQL databases, and other types of data stores. You can use AWS DMS to perform one-time migrations or continuous data replication with high availability and consolidate databases into a petabyte-scale data warehouse³. However, AWS DMS requires more configuration and management than creating a cross-Region read replica, which is fully

managed by Amazon RDS. AWS DMS also incurs additional charges for replication instances and tasks.

➤ Creating an Amazon API Gateway Data API service integration with Amazon Redshift would not help with disaster recovery or minimizing latency. The Data API is a feature that enables you to query your Amazon Redshift cluster using HTTP requests, without needing a persistent connection or a SQL client. It is useful for building applications that interact with Amazon Redshift, but not for replicating or recovering data from an RDS database.

➤ Creating an AWS Data Exchange datashare by connecting AWS Data Exchange to the Redshift cluster would not help with disaster recovery or minimizing latency. AWS Data Exchange is a service that makes it easy for AWS customers to exchange data in the cloud. You can use AWS Data Exchange to subscribe to a diverse selection of third-party data products or offer your own data products to other AWS customers. A datashare is a feature that enables you to share live and secure access to your Amazon Redshift data across your accounts or with third parties without copying or moving the underlying data. It is useful for sharing query results and views with other users, but not for replicating or recovering data from an RDS database.

References:

➤ <https://aws.amazon.com/dms/>

➤ <https://docs.aws.amazon.com/redshift/latest/mgmt/data-api.html>

➤ <https://aws.amazon.com/data-exchange/>

➤ <https://docs.aws.amazon.com/redshift/latest/dg/datashare-overview.html>

NEW QUESTION 142

- (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 145

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questions and Answers in PDF Format

SAP-C02 Practice Exam Features:

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

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