Free Questions for SAP-C02 by vceexamstest

Shared by Foster on 09-08-2024

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

A delivery company is running a serverless solution in tneAWS Cloud The solution manages user data, delivery information and past purchase details The solution consists of several microservices The central user service stores sensitive data in an Amazon DynamoDB table Several of the other microservices store a copy of parts of the sensitive data in different storage services

The company needs the ability to delete user information upon request As soon as the central user service deletes a user every other microservice must also delete its copy of the data immediately

Which solution will meet these requirements?

Options:

A- Activate DynamoDB Streams on the DynamoDB table Create an AWS Lambda trigger for the DynamoDB stream that will post events about user deletion in an Amazon Simple Queue Service (Amazon SQS) queue Configure each microservice to poll the queue and delete the user from the DynamoDB table

B- Set up DynamoDB event notifications on the DynamoDB table Create an Amazon Simple Notification Service (Amazon SNS) topic as a target for the DynamoDB event notification Configure each microservice to subscribe to the SNS topic and to delete the user from the DynamoDB table

C- Configure the central user service to post an event on a custom Amazon EventBridge event bus when the company deletes a user Create an EventBndge rule for each microservice to match the user deletion event pattern and invoke logic in the microservice to delete

the user from the DynamoDB table

D- Configure the central user service to post a message on an Amazon Simple Queue Service (Amazon SQS) queue when the company deletes a user Configure each microservice to create an event filter on the SQS queue and to delete the user from the DynamoDB table

Answer:

C

Explanation:

Set Up EventBridge Event Bus:

Step 1: Open the Amazon EventBridge console and create a custom event bus. This bus will be used to handle user deletion events.

Step 2: Name the event bus appropriately (e.g., user-deletion-bus).

Post Events on User Deletion:

Step 1: Modify the central user service to post an event to the custom EventBridge event bus whenever a user is deleted.

Step 2: Ensure the event includes relevant details such as the user ID and any other necessary metadata.

Create EventBridge Rules for Microservices:

Step 1: For each microservice that needs to delete user data, create a new rule in EventBridge that triggers on the user deletion event.

Step 2: Define the event pattern to match the user deletion event. This pattern should include the event details posted by the central user service.

Invoke Microservice Logic:

Step 1: Configure the EventBridge rule to invoke a target, such as an AWS Lambda function, which contains the logic to delete the user data from the microservice's data store.

Step 2: Each microservice should have its Lambda function or equivalent logic to handle the deletion of user data upon receiving the event.

Using Amazon EventBridge ensures a scalable, reliable, and decoupled approach to handle the deletion of user data across multiple microservices. This setup allows each microservice to independently process user deletion events without direct dependencies on other services.

Reference

AWS EventBridge Documentation

DynamoDB Streams and AWS Lambda Triggers

Implementing the Transactional Outbox Pattern with EventBridge Pipes (AWS Documentation) (Amazon Web Services, Inc.) (Amazon Web Services, Inc.) (AWS Documentation) (AWS Cloud Community).

Question 2

Question Type: MultipleChoice

A company needs to use an AWS Transfer Family SFTP-enabled server with an Amazon S3 bucket to receive updates from a third-party data supplier. The data is encrypted with Pretty Good Privacy (PGP) encryption The company needs a solution that will automatically decrypt the data after the company receives the data

A solutions architect will use a Transfer Family managed workflow The company has created an 1AM service role by using an 1AM policy that allows access to AWS Secrets Manager and the S3 bucket The role's trust relationship allows the transfer amazonaws com service to assume the rote

What should the solutions architect do next to complete the solution for automatic decryption'?

Options:

A- Store the PGP public key in Secrets Manager Add a nominal step in the Transfer Family managed workflow to decrypt files Configure PGP encryption parameters in the nominal step Associate the workflow with the Transfer Family server

B- Store the PGP private key in Secrets Manager Add an exception-handling step in the Transfer Family managed workflow to decrypt files Configure PGP encryption parameters in the exception handler Associate the workflow with the SFTP user

C- Store the PGP private key in Secrets Manager Add a nominal step in the Transfer Family managed workflow to decrypt files. Configure PGP decryption parameters in the nominal step Associate the workflow with the Transfer Family server

D- Store the PGP public key in Secrets Manager Add an exception-handling step in the Transfer Family managed workflow to decrypt files Configure PGP decryption parameters in the exception handler Associate the workflow with the SFTP user

Answer:

C

Explanation:

Store the PGP Private Key:

Step 1: In the AWS Management Console, navigate to AWS Secrets Manager.

Step 2: Store the PGP private key in Secrets Manager. Ensure the key is encrypted and properly secured.

Set Up the Transfer Family Managed Workflow:

Step 1: In the AWS Transfer Family console, create a new managed workflow.

Step 2: Add a nominal step to the workflow that includes the decryption of the files. Configure this step with the PGP decryption parameters, referencing the PGP private key stored in Secrets Manager.

Step 3: Associate this workflow with the Transfer Family SFTP server, ensuring that incoming files are automatically decrypted upon receipt.

This solution ensures that the data is securely decrypted as it is transferred from the SFTP server to the S3 bucket, automating the decryption process and leveraging AWS Secrets Manager for key management.

Reference

AWS Transfer Family Documentation

Using AWS Secrets Manager for Managing Secrets

AWS Transfer Family Managed Workflows

Question 3

Question Type: MultipleChoice

A software as a service (SaaS) company provides a media software solution to customers The solution is hosted on 50 VPCs across various AWS Regions and AWS accounts One of the VPCs is designated as a management VPC The compute resources in the VPCs work independently

The company has developed a new feature that requires all 50 VPCs to be able to communicate with each other. The new feature also requires one-way access from each customer's VPC to the company's management VPC The management VPC hosts a compute resource that validates licenses for the media software solution

The number of VPCs that the company will use to host the solution will continue to increase as the solution grows

Which combination of steps will provide the required VPC connectivity with the LEAST operational overhead" (Select TWO.)

Options:

A- Create a transit gateway Attach all the company's VPCs and relevant subnets to the transit gateway

B- Create VPC peering connections between all the company's VPCs

C- Create a Network Load Balancer (NLB) that points to the compute resource for license validation. Create an AWS PrivateLink endpoint service that is available to each customer's VPC Associate the endpoint service with the NLB

D- Create a VPN appliance in each customer's VPC Connect the company's management VPC to each customer's VPC by using AWS Site-to-Site VPN

E- Create a VPC peering connection between the company's management VPC and each customer's VPC

Answer:

A, C

Explanation:

Create a Transit Gateway:

Step 1: In the AWS Management Console, navigate to the VPC Dashboard.

Step 2: Select 'Transit Gateways' and click on 'Create Transit Gateway'.

Step 3: Configure the transit gateway by providing a name and setting the options for Amazon side ASN and VPN ECMP support as needed.

Step 4: Attach each of the company's VPCs and relevant subnets to the transit gateway. This centralizes the network management and simplifies the routing configurations, supporting scalable and flexible network architecture.

Set Up AWS PrivateLink:

Step 1: Create a Network Load Balancer (NLB) in the management VPC that points to the compute resource responsible for license validation.

Step 2: Create an AWS PrivateLink endpoint service pointing to this NLB.

Step 3: Allow each customer's VPC to create an interface endpoint to this PrivateLink service. This setup enables secure and private communication between the customer VPCs and the management VPC, ensuring one-way access from each customer's VPC to the management VPC for license validation.

This combination leverages the benefits of AWS Transit Gateway for scalable and centralized routing, and AWS PrivateLink for secure and private service access, meeting the requirement with minimal operational overhead.

Reference

Amazon VPC-to-Amazon VPC Connectivity Options

AWS PrivateLink - Building a Scalable and Secure Multi-VPC AWS Network Infrastructure

Connecting Your VPC to Other VPCs and Networks Using a Transit Gateway

Question 4

Question Type: MultipleChoice

A company creates an AWS Control Tower landing zone to manage and govern a multi-account AWS environment. The company's security team will deploy preventive controls and detective controls to monitor AWS services across all the accounts. The security team needs a centralized view of the security state of all the accounts.

Which solution will meet these requirements'?

Options:

A- From the AWS Control Tower management account, use AWS CloudFormation StackSets to deploy an AWS Config conformance pack to all accounts in the organization

B- Enable Amazon Detective for the organization in AWS Organizations Designate one AWS account as the delegated administrator for Detective

C- From the AWS Control Tower management account, deploy an AWS CloudFormation stack set that uses the automatic deployment option to enable Amazon Detective for the organization

D- Enable AWS Security Hub for the organization in AWS Organizations Designate one AWS account as the delegated administrator for Security Hub

Answer:

D

Explanation:

Enable AWS Security Hub:

Navigate to the AWS Security Hub console in your management account and enable Security Hub. This process integrates Security Hub with AWS Control Tower, allowing you to manage and monitor security findings across all accounts within your organization.

Designate a Delegated Administrator:

In AWS Organizations, designate one of the AWS accounts as the delegated administrator for Security Hub. This account will have the responsibility to manage and oversee the security posture of all accounts within the organization.

Deploy Controls Across Accounts:

Use AWS Security Hub to automatically enable security controls across all AWS accounts in the organization. This provides a centralized view of the security state of all accounts and ensures continuous monitoring and compliance.

Utilize AWS Security Hub Features:

Leverage the capabilities of Security Hub to aggregate security alerts, run continuous security checks, and generate findings based on the AWS Foundational Security Best Practices. Security Hub integrates with other AWS services like AWS Config, Amazon GuardDuty, and AWS IAM Access Analyzer to enhance security monitoring and remediation.

By integrating AWS Security Hub with AWS Control Tower and using a delegated administrator account, you can achieve a centralized and comprehensive view of your organization's security posture, facilitating effective management and remediation of security issues.

Reference

AWS Security Hub now integrates with AWS Control Tower77

AWS Control Tower and Security Hub Integration76

AWS Security Hub Features79

Question 5

Question Type: MultipleChoice

A company has implemented a new security requirement According to the new requirement, the company must scan all traffic from corporate AWS instances in the company's VPC for violations of the company's security policies. As a result of these scans the company can block access to and from specific IP addresses.

To meet the new requirement, the company deploys a set of Amazon EC2 instances in private subnets to serve as transparent proxies The company installs approved proxy server software on these EC2 instances The company modifies the route tables on all subnets to use the corresponding EC2 instances with proxy software as the default route The company also creates security groups that are compliant with the security policies and assigns these security groups to the EC2 instances

Despite these configurations, the traffic of the EC2 instances in their private subnets is not being properly forwarded to the internet.

What should a solutions architect do to resolve this issue?

Options:

A- Disable source'destination checks on the EC2 instances that run the proxy software

B- Add a rule to the security group that is assigned to the proxy EC2 instances to allow all traffic between instances that have this security group Assign this security group to all EC2 instances in the VPC.

C- Change the VPC's DHCP options set Set the DNS server options to point to the addresses of the proxy EC2 instances

D- Assign one additional elastic network interface to each proxy EC2 instance Ensure that one of these network interfaces has a route to the private subnets Ensure that the other network interface has a route to the internet.

Answer:

A

Explanation:

Identify Proxy EC2 Instances:

Determine which EC2 instances in the private subnets are running the proxy server software.

Disable Source/Destination Checks:

For each of these EC2 instances, go to the AWS Management Console.

Navigate to the EC2 dashboard, select the instance, and choose 'Actions' > 'Networking' > 'Change Source/Dest. Check'.

Disable the source/destination check for these instances.

Disabling source/destination checks allows the EC2 instances to route traffic appropriately, enabling them to function as network appliances or proxies. This ensures that traffic from other instances in the private subnets can be routed through the proxy instances to the internet, meeting the company's security requirements.

Reference

Amazon EC2 User Guide on Source/Destination Checks

Question 6

Question Type: **MultipleChoice**

A company has developed a new release of a popular video game and wants to make it available for public download The new release package is approximately 5 GB in size. The company provides downloads for existing releases from a Linux-based publicly facing FTP site hosted in an on-premises data center The company expects the new release will be downloaded by users worldwide The company wants a solution that provides improved download performance and low transfer costs regardless of a user's location

Which solutions will meet these requirements'?

Options:

A- Store the game files on Amazon EBS volumes mounted on Amazon EC2 instances within an Auto Scaling group Configure an FTP service on the EC2 instances Use an Application Load Balancer in front of the Auto Scaling group. Publish the game download URL for users to download the package

B- Store the game files on Amazon EFS volumes that are attached to Amazon EC2 instances within an Auto Scaling group Configure an FTP service on each of the EC2 instances Use an Application Load Balancer in front of the Auto Scaling group Publish the game download URL for users to download the package

C- Configure Amazon Route 53 and an Amazon S3 bucket for website hosting Upload the game files to the S3 bucket Use Amazon CloudFront for the website Publish the game download URL for users to download the package

D- Configure Amazon Route 53 and an Amazon S3 bucket for website hosting Upload the game files to the S3 bucket Set Requester Pays for the S3 bucket Publish the game download URL for users to download the package

Answer:

C

Explanation:

Create an S3 Bucket:

Navigate to Amazon S3 in the AWS Management Console and create a new S3 bucket to store the game files. Enable static website hosting on this bucket.

Upload Game Files:

Upload the 5 GB game release package to the S3 bucket. Ensure that the files are publicly accessible if required for download.

Configure Amazon Route 53:

Set up a new domain or subdomain in Amazon Route 53 and point it to the S3 bucket. This allows users to access the game files using a custom URL.

Use Amazon CloudFront:

Create a CloudFront distribution with the S3 bucket as the origin. CloudFront is a content delivery network (CDN) that caches content at edge locations worldwide, improving download performance and reducing latency for users regardless of their location.

Publish the Download URL:

Use the CloudFront distribution URL as the download link for users to access the game files. CloudFront will handle the efficient distribution and caching of the content.

This solution leverages the scalability of Amazon S3 and the performance benefits of CloudFront to provide an optimal download experience for users globally while minimizing costs.

Reference

Amazon CloudFront Documentation

Amazon S3 Static Website Hosting

Question 7

Question Type: MultipleChoice

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

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

Which solution will meet these requirements?

Options:

A- Create a new Amazon FSx for Windows File Server file system Configure AWS DataSync with one location for the on-premises file share and one location for the new Amazon FSx file system Create a new DataSync task to copy the data from the on-premises file share location to the Amazon FSx file system

B- Create an S3 bucket for the application Copy the data from the on-premises storage to the S3 bucket

C- Deploy an AWS Server Migration Service (AWS SMS) VM to the on-premises environment Use AWS SMS to migrate the file storage server from on premises to an Amazon EC2 instance

D- Create an S3 bucket for the application Deploy a new AWS Storage Gateway file gateway on an on-premises VM Create a new file share that stores data in the S3 bucket and is associated with the file gateway Copy the data from the on-premises storage to the new file gateway endpoint

Answer:

D

Explanation:

Create an S3 Bucket:

Log in to the AWS Management Console and navigate to Amazon S3.

Create a new S3 bucket that will serve as the destination for the application data.

Deploy AWS Storage Gateway:

Download and deploy the AWS Storage Gateway virtual machine (VM) on your on-premises environment. This VM can be deployed on VMware ESXi, Microsoft Hyper-V, or Linux KVM.

Configure the File Gateway:

Configure the deployed Storage Gateway as a file gateway. This will enable it to present Amazon S3 buckets as SMB file shares to your on-premises applications.

Create a New File Share:

Within the Storage Gateway configuration, create a new file share that is associated with the S3 bucket you created earlier. This file share will use the SMB protocol, allowing your on-premises applications to access the S3 bucket as if it were a local SMB file share.

Copy Data to the File Gateway:

Use your preferred method (such as robocopy, rsync, or similar tools) to copy data from the on-premises storage to the newly created file gateway endpoint. This data will be stored in the S3 bucket, maintaining accessibility through SMB.

Ensure Secure and Efficient Data Transfer:

AWS Storage Gateway ensures that all data in transit is encrypted using TLS, providing secure data transfer to AWS. It also provides local caching for frequently accessed data, improving access performance for on-premises applications.

This approach allows your existing on-premises applications to continue accessing data via SMB while leveraging the scalability and durability of Amazon S3.

Reference

AWS Storage Gateway Overview67.

AWS DataSync and Storage Gateway Hybrid Architecture66.

AWS S3 File Gateway Details68.

Question 8

Question Type: MultipleChoice

A solutions architect is creating an AWS CloudFormation template from an existing manually created non-production AWS environment The CloudFormation template can be destroyed and recreated as needed The environment contains an Amazon EC2 instance The EC2 instance has an instance profile that the EC2 instance uses to assume a role in a parent account

The solutions architect recreates the role in a CloudFormation template and uses the same role name When the CloudFormation template is launched in the child account, the EC2 instance can no longer assume the role in the parent account because of insufficient permissions

What should the solutions architect do to resolve this issue?

Options:

A- In the parent account edit the trust policy for the role that the EC2 instance needs to assume Ensure that the target role ARN in the existing statement that allows the sts AssumeRole action is correct Save the trust policy

B- In the parent account edit the trust policy for the role that the EC2 instance needs to assume Add a statement that allows the sts AssumeRole action for the root principal of the child account Save the trust policy

C- Update the CloudFormation stack again Specify only the CAPABILITY_NAMED_IAM capability

D- Update the CloudFormation stack again Specify the CAPABIUTYJAM capability and the CAPABILITY_NAMEDJAM capability

Answer:

A

Explanation:

Edit the Trust Policy:

Go to the IAM console in the parent account and locate the role that the EC2 instance needs to assume.

Edit the trust policy of the role to ensure that it correctly allows the sts

action for the role ARN in the child account.

Update the Role ARN:

Verify that the target role ARN specified in the trust policy matches the role ARN created by the CloudFormation stack in the child account.

If necessary, update the ARN to reflect the correct role in the child account.

Save and Test:

Save the updated trust policy and ensure there are no syntax errors.

Test the setup by attempting to assume the role from the EC2 instance in the child account. Verify that the instance can successfully assume the role and perform the required actions.

This ensures that the EC2 instance in the child account can assume the role in the parent account, resolving the permission issue.

Reference

AWS IAM Documentation on Trust Policies51.

Question 9

Question Type: MultipleChoice

A solutions architect has deployed a web application that serves users across two AWS Regions under a custom domain The application uses Amazon Route 53 latency-based routing The solutions architect has associated weighted record sets with a pair of web servers in separate Availability Zones for each Region

The solutions architect runs a disaster recovery scenario When all the web servers in one Region are stopped. Route 53 does not automatically redirect users to the other Region

Which of the following are possible root causes of this issue1? (Select TWO)

Options:

A- The weight for the Region where the web servers were stopped is higher than the weight for the other Region.

B- One of the web servers in the secondary Region did not pass its HTTP health check

C- Latency resource record sets cannot be used in combination with weighted resource record sets

D- The setting to evaluate target health is not turned on for the latency alias resource record set that is associated with the domain in the Region where the web servers were stopped.

E- An HTTP health check has not been set up for one or more of the weighted resource record sets associated with the stopped web servers

Answer:

D, E

Explanation:

Evaluate Target Health Setting:

Ensure that the 'Evaluate Target Health' setting is enabled for the latency alias resource record sets in Route 53. This setting helps Route 53 determine the health of the resources associated with the alias record and redirect traffic appropriately.

HTTP Health Checks:

Configure HTTP health checks for all weighted resource record sets. Health checks monitor the availability and performance of the web servers, allowing Route 53 to reroute traffic to healthy servers in case of a failure.

Verify that the health checks are correctly set up and associated with the resource record sets. This ensures that Route 53 can detect server failures and redirect traffic to the servers in the other Region.

By enabling the 'Evaluate Target Health' setting and configuring HTTP health checks, Route 53 can effectively manage traffic during failover scenarios, ensuring high availability and reliability.

Reference

AWS Route 53 Documentation on Latency-Based Routing50.

AWS Architecture Blog on Cross-Account and Cross-Region Setup49.

Question 10

Question Type: **MultipleChoice**

A company has a web application that uses Amazon API Gateway. AWS Lambda and Amazon DynamoDB A recent marketing campaign has increased demand Monitoring software reports that many requests have significantly longer response times than before the marketing campaign

A solutions architect enabled Amazon CloudWatch Logs for API Gateway and noticed that errors are occurring on 20% of the requests. In CloudWatch. the Lambda function. Throttles metric represents 1% of the requests and the Errors metric represents 10% of the requests Application logs indicate that, when errors occur there is a call to DynamoDB

What change should the solutions architect make to improve the current response times as the web application becomes more popular'?

Options:

A- Increase the concurrency limit of the Lambda function

B- Implement DynamoDB auto scaling on the table

C- Increase the API Gateway throttle limit

D- Re-create the DynamoDB table with a better-partitioned primary index.

Answer:

B

Explanation:

Enable DynamoDB Auto Scaling:

Navigate to the DynamoDB console and select the table experiencing high demand.

Go to the 'Capacity' tab and enable auto scaling for both read and write capacity units. Auto scaling adjusts the provisioned throughput capacity automatically in response to actual traffic patterns, ensuring the table can handle the increased load.

Configure Auto Scaling Policies:

Set the minimum and maximum capacity units to define the range within which auto scaling can adjust the provisioned throughput.

Specify target utilization percentages for read and write operations, typically around 70%, to maintain a balance between performance and cost.

Monitor and Adjust:

Use Amazon CloudWatch to monitor the auto scaling activity and ensure it is effectively handling the increased demand.

Adjust the auto scaling settings if necessary to better match the traffic patterns and application requirements.

By enabling DynamoDB auto scaling, you ensure that the database can handle the fluctuating traffic volumes without manual intervention, improving response times and reducing errors.

Reference

AWS Compute Blog on Using API Gateway as a Proxy for DynamoDB60.

AWS Database Blog on DynamoDB Accelerator (DAX)59.

Question 11

Question Type: **MultipleChoice**

A company has an application that analyzes and stores image data on premises The application receives millions of new image files every day Files are an average of 1 MB in size The files are analyzed in batches of 1 GB When the application analyzes a batch the application zips the images together The application then archives the images as a single file in an on-premises NFS server for long-term storage

The company has a Microsoft Hyper-V environment on premises and has compute capacity available The company does not have storage capacity and wants to archive the images on AWS The company needs the ability to retrieve archived data within t week of a request.

The company has a 10 Gbps AWS Direct Connect connection between its on-premises data center and AWS. The company needs to set bandwidth limits and schedule archived images to be copied to AWS dunng non-business hours.

Which solution will meet these requirements MOST cost-effectively?

Options:

A- Deploy an AWS DataSync agent on a new GPU-based Amazon EC2 instance Configure the DataSync agent to copy the batch of files from the NFS on-premises server to Amazon S3 Glacier Instant Retrieval After the successful copy delete the data from the on-premises storage

B- Deploy an AWS DataSync agent as a Hyper-V VM on premises Configure the DataSync agent to copy the batch of files from the NFS on-premises server to Amazon S3 Glacier Deep Archive After the successful copy delete the data from the on-premises storage

C- Deploy an AWS DataSync agent on a new general purpose Amazon EC2 instance Configure the DataSync agent to copy the batch of files from the NFS on-premises server to Amazon S3 Standard After the successful copy deletes the data from the on-premises storage Create an S3 Lifecycle rule to transition objects from S3 Standard to S3 Glacier Deep Archive after 1 day

D- Deploy an AWS Storage Gateway Tape Gateway on premises in the Hyper-V environment Connect the Tape Gateway to AWS Use automatic tape creation Specify an Amazon S3 Glacier Deep Archive pool Eject the tape after the batch of images is copied

Answer:

B

Explanation:

Deploy DataSync Agent:

Install the AWS DataSync agent as a VM in your Hyper-V environment. This agent facilitates the data transfer between your on-premises storage and AWS.

Configure Source and Destination:

Set up the source location to point to your on-premises NFS server where the image batches are stored.

Configure the destination location to be an Amazon S3 bucket with the Glacier Deep Archive storage class. This storage class is cost-effective for long-term storage with retrieval times of up to 12 hours.

Create DataSync Tasks:

Create and configure DataSync tasks to manage the data transfer. Schedule these tasks to run during non-business hours to minimize bandwidth usage during peak times. The tasks will handle the copying of data batches from the NFS server to the S3 bucket.

Set Bandwidth Limits:

In the DataSync configuration, set bandwidth limits to control the amount of data being transferred at any given time. This ensures that your network's performance is not adversely affected during business hours.

Delete On-Premises Data:

After successfully copying the data to S3 Glacier Deep Archive, configure the DataSync task to delete the data from your on-premises NFS server. This helps manage storage capacity on-premises and ensures data is securely archived on AWS.

This approach leverages AWS DataSync for efficient, secure, and automated data transfer, and S3 Glacier Deep Archive for cost-effective long-term storage.

Reference

AWS DataSync Overview41.

AWS Storage Blog on DataSync Migration40.

Amazon S3 Transfer Acceleration Documentation42.

Question 12

Question Type: **MultipleChoice**

A company uses AWS Organizations to manage its development environment. Each development team at the company has its own AWS account Each account has a single VPC and CIDR blocks that do not overlap.

The company has an Amazon Aurora DB cluster in a shared services account All the development teams need to work with live data from the DB cluster

Which solution will provide the required connectivity to the DB cluster with the LEAST operational overhead?

Options:

A- Create an AWS Resource Access Manager (AWS RAM) resource share tor the DB cluster. Share the DB cluster with all the development accounts

B- Create a transit gateway in the shared services account Create an AWS Resource Access Manager (AWS RAM) resource share for the transit gateway Share the transit gateway with all the development accounts Instruct the developers to accept the resource share Configure networking.

C- Create an Application Load Balancer (ALB) that points to the IP address of the DB cluster Create an AWS PrivateLink endpoint service that uses the ALB Add permissions to allow each development account to connect to the endpoint service

D- Create an AWS Site-to-Site VPN connection in the shared services account Configure networking Use AWS Marketplace VPN software in each development account to connect to the Site-to-Site VPN connection

Answer:

B

Explanation:

Create a Transit Gateway:

In the shared services account, create a new AWS Transit Gateway. This serves as a central hub to connect multiple VPCs, simplifying the network topology and management.

Configure Transit Gateway Attachments:

Attach the VPC containing the Aurora DB cluster to the transit gateway. This allows the shared services VPC to communicate through the transit gateway.

Create Resource Share with AWS RAM:

Use AWS Resource Access Manager (AWS RAM) to create a resource share for the transit gateway. Share this resource with all development accounts. AWS RAM allows you to securely share your AWS resources across AWS accounts without needing to duplicate them.

Accept Resource Shares in Development Accounts:

Instruct each development team to log into their respective AWS accounts and accept the transit gateway resource share. This step is crucial for enabling cross-account access to the shared transit gateway.

Configure VPC Attachments in Development Accounts:

Each development account needs to attach their VPC to the shared transit gateway. This allows their VPCs to route traffic through the transit gateway to the Aurora DB cluster in the shared services account.

Update Route Tables:

Update the route tables in each VPC to direct traffic intended for the Aurora DB cluster through the transit gateway. This ensures that network traffic is properly routed between the development VPCs and the shared services VPC.

Using a transit gateway simplifies the network management and reduces operational overhead by providing a scalable and efficient way to interconnect multiple VPCs across different AWS accounts.

Reference

AWS Database Blog on RDS Proxy for Cross-Account Access48.

AWS Architecture Blog on Cross-Account and Cross-Region Aurora Setup49.

DEV Community on Managing Multiple AWS Accounts with Organizations51.