DUMPSARENA

AWS Certified Solutions Architect - Professional

Amazon AWS SAP-C02

Version Demo

Total Demo Questions: 20

Total Premium Questions: 421

Buy Premium PDF

https://dumpsarena.co

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

Topic Break Down

Topic	No. of Questions
Topic 1, Exam Pool A	143
Topic 2, Exam Pool B	278
Total	421

QUESTION NO: 1

A company has used infrastructure as code (IaC) to provision a set of two Amazon EC2 instances. The instances have remained the same tor several years.

The company's business has grown rapidly in the past few months. In response, the company's operations team has implemented an Auto Scaling group to manage the sudden increases in traffic Company policy requires a monthly installation of security updates on all operating systems that are running.

The most recent security update required a reboot. As a result the Auto Scaling group terminated the instances and replaced them with new, unpatched instances.

Which combination of steps should a sol-tons architect recommend to avoid a recurrence of this issue? (Select TWO)

A. Modify the Auto Scaling group by setting the Update policy to target the oldest launch configuration for replacement.

B. Create a new Auto Scaling group before the next patch maintenance During the maintenance window patch both groups and reboot the instances.

C. Create an Elastic Load Balancer in front of the Auto Scaling group Configure monitoring to ensure that target group health checks return healthy after the Auto Scaling group replaces the terminated instances

D. Create automation scripts to patch an AMI. update the launch configuration, and invoke an Auto Scaling instance refresh.

E. Create an Elastic Load Balancer in front of the Auto Scaling group Configure termination protection on the instances.

ANSWER: A C

QUESTION NO: 2

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

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

A. Create a REST API in Amazon API Gateway and use AWS Lambda functions as the application layer.

B. Create an Application Load Balancer and migrate the Docker container to AWS Fargate.

C. Migrate the storage layer to Amazon DynamoD8.

D. Migrate the storage layer to Amazon DocumentD8 (with MongoDB compatibility).

E. Create an Application Load Balancer and move the storage layer to an EC2 Auto Scaling group.

ANSWER: B D

Explanation:

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

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

QUESTION NO: 3

A solutions architect needs to implement a client-side encryption mechanism for objects that will be stored in a new Amazon S3 bucket. The solutions architect created a CMK that is stored in AWS Key Management Service (AWS KMS) for this purpose.

The solutions architect created the following IAM policy and attached it to an IAM role:

During tests, me solutions architect was able to successfully get existing test objects m the S3 bucket However, attempts to upload a new object resulted in an error message. The error message stated that me action was forbidden.

Which action must me solutions architect add to the IAM policy to meet all the requirements?

A. Kms:GenerateDataKey

B. KmsGetKeyPolpcy

C. kmsGetPubKKey

D. kms:SKjn

ANSWER: A

Explanation:

https://aws.amazon.com/premiumsupport/knowledge-center/s3-access-denied-error-kms/

"An error occurred (AccessDenied) when calling the PutObject operation: Access Denied" This error message indicates that your IAM user or role needs permission for the kms:GenerateDataKey action.

QUESTION NO: 4

A company plans to migrate to AWS. A solutions architect uses AWS Application Discovery Service over the fleet and discovers that there is an Oracle data warehouse and several PostgreSQL databases. Which combination of migration patterns will reduce licensing costs and operational overhead? (Select TWO.)

A. Lift and shift the Oracle data warehouse to Amazon EC2 using AWS DMS.

B. Migrate the Oracle data warehouse to Amazon Redshift using AWS SCT and AWS QMS.

C. Lift and shift the PostgreSQL databases to Amazon EC2 using AWS DMS.

D. Migrate the PostgreSQL databases to Amazon RDS for PostgreSQL using AWS DMS

E. Migrate the Oracle data warehouse to an Amazon EMR managed cluster using AWS DMS.

ANSWER: B D

Explanation:

https://aws.amazon.com/getting-started/hands-on/migrate-oracle-to-amazon-redshift/

https://docs.aws.amazon.com/prescriptive-guidance/latest/patterns/migrate-an-on-premises-postgresql-database-to-amazon-rds-for-postgresql.html

QUESTION NO: 5

A solutions architect is building a web application that uses an Amazon RDS for PostgreSQL DB instance The DB instance is expected to receive many more reads than writes. The solutions architect needs to ensure that the large amount of read traffic can be accommodated and that the DB instance is highly available.

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

A. Create multiple read replicas and put them into an Auto Scaling group.

B. Create multiple read replicas in different Availability Zones.

C. Create an Amazon Route 53 hosted zone and a record set for each read replica with a TTL and a weighted routing policy.

D. Create an Application Load Balancer (ALB) and put the read replicas behind the ALB.

E. Configure an Amazon CloudWatch alarm to detect a failed read replica. Set the alarm to directly invoke an AWS Lambda function to delete its Route 53 record set.

F. Configure an Amazon Route 53 health check for each read replica using its endpoint

ANSWER: B C F

Explanation:

https://aws.amazon.com/premiumsupport/knowledge-center/requests-rds-read-replicas/

You can use Amazon Route 53 weighted record sets to distribute requests across your read replicas. Within a Route 53 hosted zone, create individual record sets for each DNS endpoint associated with your read replicas and give them the same weight. Then, direct requests to the endpoint of the record set. You can incorporate Route 53 health checks to be sure that Route 53 directs traffic away from unavailable read replicas

QUESTION NO: 6

A company wants to send data from its on-premises systems to Amazon S3 buckets. The company created the S3 buckets in three different accounts. The company must send the data privately without the data traveling across the internet. The company has no existing dedicated connectivity to AWS

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

A. Establish a networking account in the AWS Cloud Create a private VPC in the networking account Set up an AWS Direct Connect connection with a private VIF between the on-premises environment and the private VPC

B. Establish a networking account in the AWS Cloud Create a private VPC in the networking account Set up an AWS Direct Connect connection with a public VIF between the on-premises environment and the private VPC

C. Create an Amazon S3 interface endpoint in the networking account

D. Create an Amazon S3 gateway endpoint in the networking account

E. Establish a networking account in the AWS Cloud. Create a private VPC in the networking account Peer VPCs from the accounts that host the S3 buckets with the VPC in the network account

ANSWER: A C

Explanation:

https://aws.amazon.com/premiumsupport/knowledge-center/s3-bucket-access-direct-connect/

QUESTION NO: 7

A greeting card company recently advertised that customers could send cards to their favourite celebrities through the company's platform Since the advertisement was published, the platform has received constant traffic from 10.000 unique users each second.

The platform runs on m5.xlarge Amazon EC2 instances behind an Application Load Balancer (ALB) The instances run in an Auto Scaling group and use a custom AMI that is based on Amazon Linux. The platform uses a highly available Amazon Aurora MySQL DB cluster that uses primary and reader endpoints The platform also uses an Amazon ElastiCache for Redis cluster that uses its cluster endpoint

The platform generates a new process for each customer and holds open database connections to MySQL for the duration of each customer's session However, resource usage for the platform is low.

Many customers are reporting errors when they connect to the platform Logs show that connections to the Aurora database are failing Amazon CloudWatch metrics show that the CPU load is tow across the platform and that connections to the platform are successful through the ALB.

Which solution will remediate the errors MOST cost-effectively?

A. Set up an Amazon CloudFront distribution Set the ALB as the origin Move all customer traffic to the CloudFront distribution endpoint

B. Use Amazon RDS Proxy Reconfigure the database connections to use the proxy

C. Increase the number of reader nodes in the Aurora MySQL cluster

D. Increase the number of nodes in the ElastiCache for Redis cluster

ANSWER: B

Explanation:

https://aws.amazon.com/rds/proxy/faqs/

QUESTION NO: 8

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

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

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

A. Resize the MySQL General Purpose SSD storage to 6 TB to improve the volume's IOPS

B. Re-architect the database tier to use Amazon Aurora instead of an RDS MySQL DB instance and add read replicas

C. Leverage Amazon Kinesis Data Streams and AWS Lambda to ingest and process the raw data

D. Use AWS X-Ray to analyze and debug application issues and add more API servers to match the load

E. Re-architect the database tier to use Amazon DynamoDB instead of an RDS MySQL DB instance

ANSWER: C E

QUESTION NO: 9

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

Which solution meets the requirements?

A. Copy the DynamoDB tables into Apache Hive tables on Amazon EMR every hour and analyze them (or anomalous behaviors. Send Amazon SNS notifications when anomalous behaviors are detected.

B. Use AWS CloudTrail to capture all the APIs that change the DynamoDB tables. Send SNS notifications when anomalous behaviors are detected using CloudTrail event filtering.

C. Use Amazon DynamoDB Streams to capture and send updates to AWS Lambda. Create a Lambda function to output records lo Amazon Kinesis Data Streams. Analyze any anomalies with Amazon Kinesis Data Analytics. Send SNS notifications when anomalous behaviors are detected.

D. Use event patterns in Amazon CloudWatch Events to capture DynamoDB API call events with an AWS Lambda (unction as a target to analyze behavior. Send SNS notifications when anomalous behaviors are detected.

ANSWER: C

Explanation:

https://aws.amazon.com/blogs/database/dynamodb-streams-use-cases-and-design-patterns/#:~:text=DynamoDB%20Streams%20is%20a%20powerful,for%20up%20to%2024%20hours.

DynamoDb Stream to capture DynamoDB update. And Kinesis Data Analytics for anomaly detection (it uses AWS proprietary Random Cut Forest Algorithm)

QUESTION NO: 10

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

Which approach should the company take to achieve this goal?

A. Enable VPC Flow Logs. Store the flow logs in an Amazon S3 bucket for analysis.

B. Enable Traffic Mirroring on the network interface of the EC2 instances. Send the mirrored traffic lo a target for storage and analysis.

C. Create an AWS WAF web ACL. and associate it with the ALB. Configure AWS WAF logging.

D. Enable logging for the ALB. Store the logs in an Amazon S3 bucket for analysis.

ANSWER: B

Explanation:

Traffic Mirroring allows to copy network traffic from a network interface to a destination network interface, Amazon EC2 instance or Amazon S3 bucket. The company can use Traffic Mirroring to analyze network payloads, detect sophisticated attacks and reverse-engineer the same.

Reference: AWS Certified Solutions Architect Professional Official Text Book, Chapter 9: Networking and Content Delivery, section: VPC Traffic Mirroring

QUESTION NO: 11

A new application is running on Amazon Elastic Container Service (Amazon ECS) with AWS Fargate The application uses an Amazon Aurora MySQL database The application and the database run m the same subnets of a VPC with distinct security groups that are configured.

The password (or the database is stored m AWS Secrets Manager and is passed to the application through the D8_PASSWORD environment variable The hostname of the database is passed to the application through the DB_HOST environment variable The application Is failing to access the database.

Which combination of actions should a solutions architect take to resolve this error? (Select THREE)

A. Ensure that the container has the environment variable with name "DB_PASSWORD" specified with a "ValueFrom" and the ARN of the secret

B. Ensure that the container has the environment variable with name *D8_PASSWORD" specified with a "ValueFrom" and the secret name of the secret.

C. Ensure that the Fargate service security group allows inbound network traffic from the Aurora MySQL database on the MySQL TCP port 3306.

D. Ensure that the Aurora MySQL database security group allows inbound network traffic from the Fargate service on the MySQL TCP port 3306.

E. Ensure that the container has the environment variable with name "D8_HOST" specified with the hostname of a DB instance endpoint.

F. Ensure that the container has the environment variable with name "DB_HOST" specified with the hostname of the OB duster endpoint.

ANSWER: A D E

QUESTION NO: 12

A company is planning to migrate its on-premises data analysis application to AWS. The application is hosted across a fleet of servers and requires consistent system time.

The company has established an AWS Direct Connect connection from its on-premises data center to AWS. The company has a high-precision stratum-0 atomic dock network appliance that acts as an NTP source for all on-premises servers.

After the migration to AWS is complete, the clock on all Amazon EC2 instances that host the application must be synchronized with the on-premises atomic clock network appliance.

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

A. Configure a DHCP options set with the on-premises NTP server address Assign the options set to the VPC. Ensure that NTP traffic is allowed between AWS and the on-premises networks.

B. Create a custom AMI to use the Amazon Time Sync Service at 169.254.169.123 Use this AMI for the application Use AWS Config to audit the NTP configuration.

C. Deploy a third-party time server from the AWS Marketplace. Configure the time server to synchronize with the on-premises atomic clock network appliance. Ensure that NTP traffic is allowed inbound in the network ACLs for the VPC that contains the third-party server.
D Create an IPsec VPN tunnel from the on-premises atomic clock network appliance to the VPC to encrypt the traffic over the Direct Connect connection. Configure the VPC route tables to direct NTP traffic over the tunnel.

ANSWER: B

Explanation:

This AMI will run a cron job that periodically synchronizes the time on the Amazon EC2 instances with the Amazon Time Sync Service. There is no need to worry about configuring DHCP options sets, configuring network ACLs, setting up third-party time servers, or setting up IPsec VPN tunnels. Additionally, using AWS Config to audit the NTP configuration ensures that the NTP service is running correctly on the instances.

QUESTION NO: 13

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

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

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

A. Create an S3 access point for each application in the AWS account that owns the S3 bucket. Configure each access point to be accessible only from the application's VPC. Update the bucket policy to require access from an access point.

B. Create an interface endpoint for Amazon S3 in each application's VPC. Configure the endpoint policy to allow access to an S3 access point. Create a VPC gateway attachment for the S3 endpoint.

C. Create a gateway endpoint lor Amazon S3 in each application's VPConfigure the endpoint policy to allow access to an S3 access point. Specify the route table that is used to access the access point.

D. Create an S3 access point for each application in each AWS account and attach the access points to the S3 bucket. Configure each access point to be accessible only from the application's VPC. Update the bucket policy to require access from an access point.

E. Create a gateway endpoint for Amazon S3 in the data lake's VPC. Attach an endpoint policy to allow access to the S3 bucket. Specify the route table that is used to access the bucket.

ANSWER: A C

Explanation:

https://joe.blog.freemansoft.com/2020/04/protect-data-in-cloud-with-s3-access.html

https://aws.amazon.com/s3/features/access-points/

https://aws.amazon.com/s3/features/access-points/

& https://aws.amazon.com/blogs/storage/managing-amazon-s3-access-with-vpc-endpoints-and-s3-access-points/

QUESTION NO: 14

A company has a platform that contains an Amazon S3 bucket for user content. The S3 bucket has thousands of terabytes of objects, all in the S3 Standard storage class. The company has an RTO of 6 hours The company must replicate the data from its primary AWS Region to a replication S3 bucket in another Region

The user content S3 bucket contains user-uploaded files such as videos and photos. The user content S3 bucket has an unpredictable access pattern. The number of users is increasing quickly, and the company wants to create an S3 Lifecycle policy to reduce storage costs

Which combination of steps will meet these requirements MOST cost-effectively'? (Select TWO)

A. Move the objects in the user content S3 bucket to S3 Intelligent-Tiering immediately

B. Move the objects in the user content S3 bucket to S3 Intelligent-Tiering after 30 days

C. Move the objects in the replication S3 bucket to S3 Standard-Infrequent Access (S3 Standard-IA) after 30 days and to S3 Glacier after 90 days

D. Move the objects in the replication S3 bucket to S3 One Zone-Infrequent Access (S3 One Zone-IA) after 30 days and to S3 Glacier Deep Archive after 90 days

E. Move the objects in the replication S3 bucket to S3 Standard-infrequent Access (S3 Standard-IA) after 30 days and to S3 Glacier Deep Archive after 180 days

ANSWER: A C

Explanation:

S3 Intelligent-Tiering is an ideal storage class for unpredictable access patterns, as it automatically and cost-effectively moves objects between the S3 Standard and S3 Infrequent Access storage classes based on actual access patterns. This can significantly reduce storage costs compared to storing all objects in S3 Standard.

S3 Standard-Infrequent Access and S3 Glacier are ideal for long-term storage and data archiving, as they are much cheaper than S3 Standard but still provide a fast retrieval time. By moving the replication bucket objects to S3 Standard-IA after 30 days and to S3 Glacier after 90 days, the company can ensure that their data is stored securely and cost-effectively while still meeting their RTO of 6 hours.

References: [1] Amazon Web Services Certified Solutions Architect Professional Official Text Book, Chapter 6: Storage (https://aws.amazon.com/certification/certification-prep/architect/) [2] AWS Documentation: Amazon S3 Storage Classes (https://docs.aws.amazon.com/AmazonS3/latest/dev/storage-class-intro.html)

QUESTION NO: 15

A company wants to host a new global website that consists of static content. A solutions architect is working on a solution that uses Amazon CloudFront with an origin access identity

During testing, the solutions architect receives 404 errors from the S3 bucket. Error messages appear only for attempts to access paths that end with a forward slash. such as example.com/path/. These requests should return the existing S3 object path/index.html. Any potential solution must not prevent CloudFront from caching the content.

What should the solutions architect do to resolve this problem?

A. Change the CloudFront origin to an Amazon API Gateway proxy endpoint. Rewrite the S3 request URL by using an AWS Lambda function.

B. Change the CloudFront origin to an Amazon API Gateway endpoint. Rewrite the S3 request URL in an AWS service integration.

C. Change the CloudFront configuration to use an AWS Lambda@Edge function that is invoked by a viewer request event to rewrite the S3 request URL.

D. Change the CloudFront configuration to use an AWS Lambda@Edge function that is invoked by an origin request event to rewrite the S3 request URL.

ANSWER: C

QUESTION NO: 16

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. and creating several additional read replicas to handle the load during end of month

C. Using Amazon CioudWatch with AWS Lambda to change the type. size, or IOPS of Amazon EBS volumes in the cluster based on a specific CloudWatch metric

D. 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.

QUESTION NO: 17

An e-commerce company is revamping its IT infrastructure and is planning to use AWS services. The company's CIO has asked a solutions architect to design a simple, highly available, and loosely coupled order processing application. The application is responsible (or receiving and processing orders before storing them in an Amazon DynamoDB table. The application has a sporadic traffic pattern and should be able to scale during markeling campaigns to process the orders with minimal delays.

Which of the following is the MOST reliable approach to meet the requirements?

A. Receive the orders in an Amazon EC2-hosted database and use EC2 instances to process them.

B. Receive the orders in an Amazon SOS queue and trigger an AWS Lambda function lo process them.

C. Receive the orders using the AWS Step Functions program and trigger an Amazon ECS container lo process them.

D. Receive the orders in Amazon Kinesis Data Streams and use Amazon EC2 instances to process them.

ANSWER: B

Explanation:

Q: How does Amazon Kinesis Data Streams differ from Amazon SQS?

Amazon Kinesis Data Streams enables real-time processing of streaming big data. It provides ordering of records, as well as the ability to read and/or replay records in the same order to multiple Amazon Kinesis Applications. The Amazon Kinesis Client Library (KCL) delivers all records for a given partition key to the same record processor, making it easier to build multiple applications reading from the same Amazon Kinesis data stream (for example, to perform counting, aggregation, and filtering).

https://aws.amazon.com/kinesis/data-streams/faqs/

https://aws.amazon.com/blogs/big-data/unite-real-time-and-batch-analytics-using-the-big-data-lambda-architecture-without-servers/

QUESTION NO: 18

A company recently deployed an application on AWS. The application uses Amazon DynamoDB. The company measured the application load and configured the RCUs and WCUs on the DynamoDB table to match the expected peak load. The peak load occurs once a week for a 4-hour period and is double the average load. The application load is close to the average load tor the rest of the week. The access pattern includes many more writes to the table than reads of the table.

A solutions architect needs to implement a solution to minimize the cost of the table.

Which solution will meet these requirements?

A. Use AWS Application Auto Scaling to increase capacity during the peak period. Purchase reserved RCUs and WCUs to match the average load.

B. Configure on-demand capacity mode for the table.

C. Configure DynamoDB Accelerator (DAX) in front of the table. Reduce the provisioned read capacity to match the new peak load on the table.

D. Configure DynamoDB Accelerator (DAX) in front of the table. Configure on-demand capacity mode for the table.

ANSWER: A

Explanation:

This solution meets the requirements by using Application Auto Scaling to automatically increase capacity during the peak period, which will handle the double the average load. And by purchasing reserved RCUs and WCUs to match the average load, it will minimize the cost of the table for the rest of the week when the load is close to the average.

QUESTION NO: 19

A company runs an application in the cloud that consists of a database and a website Users can post data to the website, have the data processed, and have the data sent back to them in an email. Data is stored in a MySQL database running on an Amazon EC2 instance The database is running in a VPC with two private subnets The website is running on Apache Tomcat in a single EC2 instance in a different VPC with one public subnet There is a single VPC peering connection between the database and website VPC.

The website has suffered several outages during the last month due to high traffic

Which actions should a solutions architect take to increase the reliability of the application? (Select THREE)

A. Place the Tomcat server in an Auto Scaling group with multiple EC2 instances behind an Application Load Balancer

B. Provision an additional VPC peering connection

C. Migrate the MySQL database to Amazon Aurora with one Aurora Replica

D. Provision two NAT gateways in the database VPC

E. Move the Tomcat server to the database VPC

F. Create an additional public subnet in a different Availability Zone in the website VPC

ANSWER: A C F

QUESTION NO: 20

A group of research institutions and hospitals are in a partnership to study 2 PBs of genomic data. The institute that owns the data stores it in an Amazon S3 bucket and updates it regularly. The institute would like to give all of the organizations in the partnership read access to the data. All members of the partnership are extremety cost-conscious, and the institute that owns the account with the S3 bucket is concerned about covering the costs tor requests and data transfers from Amazon S3.

Which solution allows for secure datasharing without causing the institute that owns the bucket to assume all the costs for S3 requests and data transfers'?

A. Ensure that all organizations in the partnership have AWS accounts. In the account with the S3 bucket, create a cross-account role for each account in the partnership that allows read access to the data. Have the organizations assume and use that read role when accessing the data.

B. Ensure that all organizations in the partnership have AWS accounts. Create a bucket policy on the bucket that owns the data The policy should allow the accounts in the partnership read access to the bucket. Enable Requester Pays on the bucket. Have the organizations use their AWS credentials when accessing the data.

C. Ensure that all organizations in the partnership have AWS accounts. Configure buckets in each of the accounts with a bucket policy that allows the institute that owns the data the ability to write to the bucket Periodically sync the data from the institute's account to the other organizations. Have the organizations use their AWS credentials when accessing the data using their accounts

D. Ensure that all organizations in the partnership have AWS accounts. In the account with the S3 bucket, create a cross-account role for each account in the partnership that allows read access to the data. Enable Requester Pays on the bucket. Have the organizations assume and use that read role when accessing the data.

ANSWER: B

Explanation:

In general, bucket owners pay for all Amazon S3 storage and data transfer costs associated with their bucket. A bucket owner, however, can configure a bucket to be a Requester Pays bucket. With Requester Pays buckets, the requester instead of the bucket owner pays the cost of the request and the data download from the bucket. The bucket owner always pays the cost of storing data. If you enable Requester Pays on a bucket, anonymous access to that bucket is not allowed.

https://docs.aws.amazon.com/AmazonS3/latest/userguide/RequesterPaysExamples.html