



Amazon-Web-Services

Exam Questions DBS-C01

AWS Certified Database - Specialty

About ExamBible

Your Partner of IT Exam

Found in 1998

ExamBible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, ExamBible has its unique advantages that other companies could not achieve.

Our Advances

* 99.9% Uptime

All examinations will be up to date.

* 24/7 Quality Support

We will provide service round the clock.

* 100% Pass Rate

Our guarantee that you will pass the exam.

* Unique Gurantee

If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

NEW QUESTION 1

A software company uses an Amazon RDS for MySQL Multi-AZ DB instance as a data store for its critical applications. During an application upgrade process, a database specialist runs a custom SQL script that accidentally removes some of the default permissions of the master user. What is the MOST operationally efficient way to restore the default permissions of the master user?

- A. Modify the DB instance and set a new master user password.
- B. Use AWS Secrets Manager to modify the master user password and restart the DB instance.
- C. Create a new master user for the DB instance.
- D. Review the IAM user that owns the DB instance, and add missing permissions.

Answer: A

NEW QUESTION 2

A company's applications store data in Amazon Aurora MySQL DB clusters. The company has separate AWS accounts for its production, test, and development environments. To test new functionality in the test environment, the company's development team requires a copy of the production database four times a day. Which solution meets this requirement with the MOST operational efficiency?

- A. Take a manual snapshot in the production account
- B. Share the snapshot with the test account
- C. Restore the database from the snapshot.
- D. Take a manual snapshot in the production account
- E. Export the snapshot to Amazon S3. Copy the snapshot to an S3 bucket in the test account
- F. Restore the database from the snapshot.
- G. Share the Aurora DB cluster with the test account
- H. Create a snapshot of the production database in the test account
- I. Restore the database from the snapshot.
- J. Share the Aurora DB cluster with the test account
- K. Create a clone of the production database in the test account.

Answer: D

Explanation:

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Aurora.Managing.Clone.html>

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Aurora.Managing.Clone.html#Aurora.Managing.Clone>

NEW QUESTION 3

An application reads and writes data to an Amazon RDS for MySQL DB instance. A new reporting dashboard needs read-only access to the database. When the application and reports are both under heavy load, the database experiences performance degradation. A database specialist needs to improve the database performance.

What should the database specialist do to meet these requirements?

- A. Create a read replica of the DB instance
- B. Configure the reports to connect to the replication instance endpoint.
- C. Create a read replica of the DB instance
- D. Configure the application and reports to connect to the cluster endpoint.
- E. Enable Multi-AZ deployment
- F. Configure the reports to connect to the standby replica.
- G. Enable Multi-AZ deployment
- H. Configure the application and reports to connect to the cluster endpoint.

Answer: A

Explanation:

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_ReadRepl.html

NEW QUESTION 4

For the first time, a database professional is establishing a test graph database on Amazon Neptune. The database expert must input millions of rows of test observations from an Amazon S3.csv file. The database professional uploaded the data to the Neptune DB instance through a series of API calls.

Which sequence of actions enables the database professional to upload the data most quickly? (Select three.)

- A. Ensure Amazon Cognito returns the proper AWS STS tokens to authenticate the Neptune DB instance to the S3 bucket hosting the CSV file.
- B. Ensure the vertices and edges are specified in different .csv files with proper header column formatting.
- C. Use AWS DMS to move data from Amazon S3 to the Neptune Loader.
- D. Curl the S3 URI while inside the Neptune DB instance and then run the addVertex or addEdge commands.
- E. Ensure an IAM role for the Neptune DB instance is configured with the appropriate permissions to allow access to the file in the S3 bucket.
- F. Create an S3 VPC endpoint and issue an HTTP POST to the database's loader endpoint.

Answer: BEF

Explanation:

<https://docs.aws.amazon.com/neptune/latest/userguide/bulk-load-optimize.html>

NEW QUESTION 5

A company's database specialist implements an AWS Database Migration Service (AWS DMS) task for change data capture (CDC) to replicate data from an on-premises Oracle database to Amazon S3. When usage of the company's application increases, the database specialist notices multiple hours of latency with the CDC.

Which solutions will reduce this latency? (Choose two.)

- A. Configure the DMS task to run in full large binary object (LOB) mode.
- B. Configure the DMS task to run in limited large binary object (LOB) mode.
- C. Create a Multi-AZ replication instance.
- D. Load tables in parallel by creating multiple replication instances for sets of tables that participate in common transactions.
- E. Replicate tables in parallel by creating multiple DMS tasks for sets of tables that do not participate in common transactions.

Answer: BE

NEW QUESTION 6

A company is using an Amazon Aurora PostgreSQL DB cluster with an xlarge primary instance master and two large Aurora Replicas for high availability and read-only workload scaling. A failover event occurs and application performance is poor for several minutes. During this time, application servers in all Availability Zones are healthy and responding normally.

What should the company do to eliminate this application performance issue?

- A. Configure both of the Aurora Replicas to the same instance class as the primary DB instance
- B. Enable cache coherence on the DB cluster, set the primary DB instance failover priority to tier-0, and assign a failover priority of tier-1 to the replicas.
- C. Deploy an AWS Lambda function that calls the DescribeDBInstances action to establish which instance has failed, and then use the PromoteReadReplica operation to promote one Aurora Replica to be the primary DB instance
- D. Configure an Amazon RDS event subscription to send a notification to an Amazon SNS topic to which the Lambda function is subscribed.
- E. Configure one Aurora Replica to have the same instance class as the primary DB instance
- F. Implement Aurora PostgreSQL DB cluster cache management
- G. Set the failover priority to tier-0 for the primary DB instance and one replica with the same instance class
- H. Set the failover priority to tier-1 for the other replicas.
- I. Configure both Aurora Replicas to have the same instance class as the primary DB instance
- J. Implement Aurora PostgreSQL DB cluster cache management
- K. Set the failover priority to tier-0 for the primary DB instance and to tier-1 for the replicas.

Answer: C

Explanation:

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraPostgreSQL.cluster-cache-mgmt.htm>

<https://aws.amazon.com/blogs/database/introduction-to-aurora-postgresql-cluster-cache-management/>

"You can customize the order in which your Aurora Replicas are promoted to the primary instance after a failure by assigning each replica a priority. Priorities range from 0 for the first priority to 15 for the last priority. If the primary instance fails, Amazon RDS promotes the Aurora Replica with the better priority to the new primary instance. You can modify the priority of an Aurora Replica at any time. Modifying the priority doesn't trigger a failover. More than one Aurora Replica can share the same priority, resulting in promotion tiers. If two or more Aurora Replicas share the same priority, then Amazon RDS promotes the replica that is largest in size. If two or more Aurora Replicas share the same priority and size, then Amazon RDS promotes an arbitrary replica in the same promotion tier. "

Amazon Aurora with PostgreSQL compatibility now supports cluster cache management, providing a faster path to full performance if there's a failover. With cluster cache management, you designate a specific reader DB instance in your Aurora PostgreSQL cluster as the failover target. Cluster cache management keeps the data in the designated reader's cache synchronized with the data in the read-write instance's cache. If a failover occurs, the designated reader is promoted to be the new read-write instance, and workloads benefit immediately from the data in its cache.

NEW QUESTION 7

A Database Specialist is migrating a 2 TB Amazon RDS for Oracle DB instance to an RDS for PostgreSQL DB instance using AWS DMS. The source RDS Oracle DB instance is in a VPC in the us-east-1 Region. The target RDS for PostgreSQL DB instance is in a VPC in the us-west-2 Region.

Where should the AWS DMS replication instance be placed for the MOST optimal performance?

- A. In the same Region and VPC of the source DB instance
- B. In the same Region and VPC as the target DB instance
- C. In the same VPC and Availability Zone as the target DB instance
- D. In the same VPC and Availability Zone as the source DB instance

Answer: C

Explanation:

https://docs.aws.amazon.com/dms/latest/userguide/CHAP_ReplicationInstance.VPC.html#CHAP_ReplicationInstance.VPC In fact, all the configurations list on above url prefer the replication instance putting into target vpc region / subnet / az.

https://docs.aws.amazon.com/dms/latest/sbs/CHAP_SQLServer2Aurora.Steps.CreateReplicationInstance.html

NEW QUESTION 8

A company is running an Amazon RDS for MySQL Multi-AZ DB instance for a business-critical workload. RDS encryption for the DB instance is disabled. A recent security audit concluded that all business-critical applications must encrypt data at rest. The company has asked its database specialist to formulate a plan to accomplish this for the DB instance.

Which process should the database specialist recommend?

- A. Create an encrypted snapshot of the unencrypted DB instance
- B. Copy the encrypted snapshot to Amazon S3. Restore the DB instance from the encrypted snapshot using Amazon S3.
- C. Create a new RDS for MySQL DB instance with encryption enabled
- D. Restore the unencrypted snapshot to this DB instance.
- E. Create a snapshot of the unencrypted DB instance
- F. Create an encrypted copy of the snapshot
- G. Restore the DB instance from the encrypted snapshot.
- H. Temporarily shut down the unencrypted DB instance
- I. Enable AWS KMS encryption in the AWS Management Console using an AWS managed CMK
- J. Restart the DB instance in an encrypted state.

Answer: C

Explanation:

<https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Overview.Encryption.html#Overview.Encryption>.

NEW QUESTION 9

A company is running a business-critical application on premises by using Microsoft SQL Server. A database specialist is planning to migrate the instance with several databases to the AWS Cloud. The database specialist will use SQL Server Standard edition hosted on Amazon EC2 Windows instances. The solution must provide high availability and must avoid a single point of failure in the SQL Server deployment architecture.

Which solution will meet these requirements?

- A. Create Amazon RDS for SQL Server Multi-AZ DB instance
- B. Use Amazon S3 as a shared storage option to host the databases.
- C. Set up Always On Failover Cluster Instances as a single SQL Server instance
- D. Use Multi-AZ Amazon FSx for Windows File Server as a shared storage option to host the databases.
- E. Set up Always On availability groups to group one or more user databases that fail over together across multiple SQL Server instances
- F. Use Multi-AZ Amazon FSx for Windows File Server as a shared storage option to host the databases.
- G. Create an Application Load Balancer to distribute database traffic across multiple EC2 instances in multiple Availability Zones
- H. Use Amazon S3 as a shared storage option to host the databases.

Answer: B

Explanation:

<https://docs.aws.amazon.com/prescriptive-guidance/latest/migration-sql-server/ec2-fci.html>

An FCI is generally preferable over an Always on availability group when: You're using SQL Server Standard edition instead of Enterprise edition.

NEW QUESTION 10

An IT consulting company wants to reduce costs when operating its development environment databases. The company's workflow creates multiple Amazon Aurora MySQL DB clusters for each development group. The Aurora DB clusters are only used for 8 hours a day. The DB clusters can then be deleted at the end of the development cycle, which lasts 2 weeks.

Which of the following provides the MOST cost-effective solution?

- A. Use AWS CloudFormation template
- B. Deploy a stack with the DB cluster for each development group. Delete the stack at the end of the development cycle.
- C. Use the Aurora DB cloning feature
- D. Deploy a single development and test Aurora DB instance, and create clone instances for the development group
- E. Delete the clones at the end of the development cycle.
- F. Use Aurora Replica
- G. From the master automatic pause compute capacity option, create replicas for each development group, and promote each replica to master
- H. Delete the replicas at the end of the development cycle.
- I. Use Aurora Serverless
- J. Restore current Aurora snapshot and deploy to a serverless cluster for each development group
- K. Enable the option to pause the compute capacity on the cluster and set an appropriate timeout.

Answer: B

Explanation:

Aurora Serverless is not compatible to all Aurora provisioned engine versions. However, you can do clones with most engine versions. Meanwhile, I also consider the performance while restoring snapshots to Aurora Serverless.

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/aurora-serverless.how-it-works.html#aurora>

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/aurora-serverless.html#aurora-serverless.us>

NEW QUESTION 10

A company is due for renewing its database license. The company wants to migrate its 80 TB transactional database system from on-premises to the AWS Cloud. The migration should incur the least possible downtime on the downstream database applications. The company's network infrastructure has limited network bandwidth that is shared with other applications.

Which solution should a database specialist use for a timely migration?

- A. Perform a full backup of the source database to AWS Snowball Edge appliances and ship them to be loaded to Amazon S3. Use AWS DMS to migrate change data capture (CDC) data from the source database to Amazon S3. Use a second AWS DMS task to migrate all the S3 data to the target database.
- B. Perform a full backup of the source database to AWS Snowball Edge appliances and ship them to be loaded to Amazon S3. Periodically perform incremental backups of the source database to be shipped in another Snowball Edge appliance to handle syncing change data capture (CDC) data from the source to the target database.
- C. Use AWS DMS to migrate the full load of the source database over a VPN tunnel using the internet for its primary connection
- D. Allow AWS DMS to handle syncing change data capture (CDC) data from the source to the target database.
- E. Use the AWS Schema Conversion Tool (AWS SCT) to migrate the full load of the source database over a VPN tunnel using the internet for its primary connection
- F. Allow AWS SCT to handle syncing change data capture (CDC) data from the source to the target database.

Answer: A

Explanation:

https://docs.aws.amazon.com/dms/latest/userguide/CHAP_Target.S3.html Using Amazon S3 as a target for AWS Database Migration Service

NEW QUESTION 14

A company is running its customer feedback application on Amazon Aurora MySQL. The company runs a report every day to extract customer feedback, and a team reads the feedback to determine if the customer comments are positive or negative. It sometimes takes days before the company can contact unhappy customers and take corrective measures. The company wants to use machine learning to automate this workflow.

Which solution meets this requirement with the LEAST amount of effort?

- A. Export the Aurora MySQL database to Amazon S3 by using AWS Database Migration Service (AWS DMS). Use Amazon Comprehend to run sentiment analysis on the exported files.
- B. Export the Aurora MySQL database to Amazon S3 by using AWS Database Migration Service (AWS DMS). Use Amazon SageMaker to run sentiment analysis on the exported files.
- C. Set up Aurora native integration with Amazon Comprehend

- D. Use SQL functions to extract sentiment analysis.
- E. Set up Aurora native integration with Amazon SageMaker
- F. Use SQL functions to extract sentiment analysis.

Answer: C

Explanation:

For details about using Aurora and Amazon Comprehend together, see Using Amazon Comprehend for sentiment detection. Aurora machine learning uses a highly optimized integration between the Aurora database and the AWS machine learning (ML) services SageMaker and Amazon Comprehend.
<https://www.stackovercloud.com/2019/11/27/new-for-amazon-aurora-use-machine-learning-directly-from-your>

NEW QUESTION 19

AWS CloudFormation stack including an Amazon RDS database instance was mistakenly removed, resulting in the loss of recent data. A Database Specialist must apply RDS parameters to the CloudFormation template in order to minimize the possibility of future inadvertent instance data loss. Which settings will satisfy this criterion? (Select three.)

- A. Set DeletionProtection to True
- B. Set MultiAZ to True
- C. Set TerminationProtection to True
- D. Set DeleteAutomatedBackups to False
- E. Set DeletionPolicy to Delete
- F. Set DeletionPolicy to Retain

Answer: ADF

Explanation:

A <https://aws.amazon.com/about-aws/whats-new/2018/09/amazon-rds-now-provides-database-deletion-protection/>
D https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_WorkingWithAutomatedBackups.html
F - <https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-attribute-deletionpolicy.html>

NEW QUESTION 21

A company is using Amazon with Aurora Replicas for read-only workload scaling. A Database Specialist needs to split up two read-only applications so each application always connects to a dedicated replica. The Database Specialist wants to implement load balancing and high availability for the read-only applications. Which solution meets these requirements?

- A. Use a specific instance endpoint for each replica and add the instance endpoint to each read-only application connection string.
- B. Use reader endpoints for both the read-only workload applications.
- C. Use a reader endpoint for one read-only application and use an instance endpoint for the other read-only application.
- D. Use custom endpoints for the two read-only applications.

Answer: D

Explanation:

<https://aws.amazon.com/about-aws/whats-new/2018/11/amazon-aurora-simplifies-workload-management-with-c>

NEW QUESTION 24

A database specialist has been entrusted by an ecommerce firm with designing a reporting dashboard that visualizes crucial business KPIs derived from the company's primary production database running on Amazon Aurora. The dashboard should be able to read data within 100 milliseconds after an update. The Database Specialist must conduct an audit of the Aurora DB cluster's present setup and provide a cost-effective alternative. The solution must support the unexpected read demand generated by the reporting dashboard without impairing the DB cluster's write availability and performance. Which solution satisfies these criteria?

- A. Turn on the serverless option in the DB cluster so it can automatically scale based on demand.
- B. Provision a clone of the existing DB cluster for the new Application team.
- C. Create a separate DB cluster for the new workload, refresh from the source DB cluster, and set up ongoing replication using AWS DMS change data capture (CDC).
- D. Add an automatic scaling policy to the DB cluster to add Aurora Replicas to the cluster based on CPU consumption.

Answer: D

NEW QUESTION 27

Recently, an ecommerce business transferred one of its SQL Server databases to an Amazon RDS for SQL Server Enterprise Edition database instance. The corporation anticipates an increase in read traffic as a result of an approaching sale. To accommodate the projected read load, a database professional must establish a read replica of the database instance. Which procedures should the database professional do prior to establishing the read replica? (Select two.)

- A. Identify a potential downtime window and stop the application calls to the source DB instance.
- B. Ensure that automatic backups are enabled for the source DB instance.
- C. Ensure that the source DB instance is a Multi-AZ deployment with Always ON Availability Groups.
- D. Ensure that the source DB instance is a Multi-AZ deployment with SQL Server Database Mirroring (DBM).
- E. Modify the read replica parameter group setting and set the value to 1.

Answer: BC

Explanation:

<https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/SQLServer.ReadReplicas.html>

NEW QUESTION 30

A database specialist is responsible for an Amazon RDS for MySQL DB instance with one read replica. The DB instance and the read replica are assigned to the default parameter group. The database team currently runs test queries against a read replica. The database team wants to create additional tables in the read replica that will only be accessible from the read replica to benefit the tests.

Which should the database specialist do to allow the database team to create the test tables?

- A. Contact AWS Support to disable read-only mode on the read replic
- B. Reboot the read replic
- C. Connect to the read replica and create the tables.
- D. Change the read_only parameter to false (read_only=0) in the default parameter group of the read replic
- E. Perform a reboot without failove
- F. Connect to the read replica and create the tables using the local_only MySQL option.
- G. Change the read_only parameter to false (read_only=0) in the default parameter grou
- H. Reboot the read replic
- I. Connect to the read replica and create the tables.
- J. Create a new DB parameter grou
- K. Change the read_only parameter to false (read_only=0). Associate the read replica with the new grou
- L. Reboot the read replic
- M. Connect to the read replica and create the tables.

Answer: D

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/rds-read-replica/>

NEW QUESTION 34

A corporation is transitioning from an IBM Informix database to an Amazon RDS for SQL Server Multi-AZ implementation with Always On Availability Groups (AGs). SQL Server Agent tasks are scheduled to execute at 5-minute intervals on the Always On AG listener to synchronize data between the Informix and SQL Server databases. After a successful failover to the backup node with minimum delay, users endure hours of stale data.

How can a database professional guarantee that consumers view the most current data after a failover?

- A. Set TTL to less than 30 seconds for cached DNS values on the Always On AG listener.
- B. Break up large transactions into multiple smaller transactions that complete in less than 5 minutes.
- C. Set the databases on the secondary node to read-only mode.
- D. Create the SQL Server Agent jobs on the secondary node from a script when the secondary node takes over after a failure.

Answer: D

Explanation:

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_SQLServerMultiAZ.html

If you have SQL Server Agent jobs, recreate them on the secondary. You do so because these jobs are stored in the msdb database, and you can't replicate this database by using Database Mirroring (DBM) or Always On Availability Groups (AGs). Create the jobs first in the original primary, then fail over, and create the same jobs in the new primary.

NEW QUESTION 39

A major organization maintains a number of Amazon DB clusters. Each of these clusters is configured differently to meet certain needs. These configurations may be classified into wider groups based on the team and use case.

A database administrator wishes to streamline the process of storing and updating these settings. Additionally, the database administrator want to guarantee that changes to certain configuration categories are automatically implemented to all instances as necessary.

Which AWS service or functionality will assist in automating and achieving this goal?

- A. AWS Systems Manager Parameter Store
- B. DB parameter group
- C. AWS Config
- D. AWS Secrets Manager

Answer: B

Explanation:

Database parameters specify how the database is configured. For example, database parameters can specify the amount of resources, such as memory, to allocate to a database.

NEW QUESTION 42

A company uses Microsoft SQL Server on Amazon RDS in a Multi-AZ deployment as the database engine for its application. The company was recently acquired by another company. A database specialist must rename the database to follow a new naming standard.

Which combination of steps should the database specialist take to rename the database? (Choose two.)

- A. Turn off automatic snapshots for the DB instanc
- B. Rename the database with the rdsadmin.dbo.rds_modify_db_name stored procedur
- C. Turn on the automatic snapshots.
- D. Turn off Multi-AZ for the DB instanc
- E. Rename the database with the rdsadmin.dbo.rds_modify_db_name stored procedur
- F. Turn on Multi-AZ Mirroring.
- G. Delete all existing snapshots for the DB instanc
- H. Use the rdsadmin.dbo.rds_modify_db_name stored procedure.
- I. Update the application with the new database connection string.
- J. Update the DNS record for the DB instance.

Answer: BD

Explanation:

<https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Appendix.SQLServer.CommonDBATasks.Renami>

NEW QUESTION 46

A Database Specialist is designing a disaster recovery strategy for a production Amazon DynamoDB table. The table uses provisioned read/write capacity mode, global secondary indexes, and time to live (TTL). The Database Specialist has restored the latest backup to a new table. To prepare the new table with identical settings, which steps should be performed? (Choose two.)

- A. Re-create global secondary indexes in the new table
- B. Define IAM policies for access to the new table
- C. Define the TTL settings
- D. Encrypt the table from the AWS Management Console or use the update-table command
- E. Set the provisioned read and write capacity

Answer: BC

Explanation:

The following items need to be reconfigured after restoring the DynamoDB table.

- AutoScaling policy
- IAM policy
- CloudWatch settings
- Tags
- Stream settings
- TTL

https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/backuprestore_HowItWorks.html

NEW QUESTION 49

An worldwide gaming company's development team is experimenting with using Amazon DynamoDB to store in-game events for three mobile titles. Maximum concurrent users for the most popular game is 500,000, while the least popular game is 10,000. The typical event is 20 KB in size, while the average user session generates one event each second. Each event is assigned a millisecond time stamp and a globally unique identification.

The lead developer generated a single DynamoDB database with the following structure for the events:

- Partition key: game name
- Sort key: event identifier
- Local secondary index: player identifier
- Event time

In a small-scale development setting, the tests were successful. When the application was deployed to production, however, new events were not being added to the database, and the logs indicated DynamoDB failures with the `ItemCollectionSizeLimitExceededException` issue code.

Which design modification should a database professional offer to the development team?

- A. Use the player identifier as the partition ke
- B. Use the event time as the sort ke
- C. Add a global secondary index with the game name as the partition key and the event time as the sort key.
- D. Create two table
- E. Use the game name as the partition key in both table
- F. Use the event time as the sort key for the first tabl
- G. Use the player identifier as the sort key for the second table.
- H. Replace the sort key with a compound value consisting of the player identifier collated with the event time, separated by a das
- I. Add a local secondary index with the player identifier as the sort key.
- J. Create one table for each gam
- K. Use the player identifier as the partition ke
- L. Use the event time as the sort key.

Answer: D

NEW QUESTION 50

A company is using Amazon RDS for MySQL to redesign its business application. A Database Specialist has noticed that the Development team is restoring their MySQL database multiple times a day when Developers make mistakes in their schema updates. The Developers sometimes need to wait hours to the restores to complete.

Multiple team members are working on the project, making it difficult to find the correct restore point for each mistake.

Which approach should the Database Specialist take to reduce downtime?

- A. Deploy multiple read replicas and have the team members make changes to separate replica instances
- B. Migrate to Amazon RDS for SQL Server, take a snapshot, and restore from the snapshot
- C. Migrate to Amazon Aurora MySQL and enable the Aurora Backtrack feature
- D. Enable the Amazon RDS for MySQL Backtrack feature

Answer: C

Explanation:

"Amazon Aurora, a fully-managed relational database service in AWS, is now offering a backtrack feature. With Amazon Aurora with MySQL compatibility, users can backtrack, or "rewind", a database cluster to a specific point in time, without restoring data from a backup. The backtrack process allows a point in time to be specified with one second resolution, and the rewind process typically takes minutes. This new feature facilitates developers in undoing mistakes like deleting data inappropriately or dropping the wrong table."

NEW QUESTION 52

A bike rental company operates an application to track its bikes. The application receives location and condition data from bike sensors. The application also receives rental transaction data from the associated mobile app.

The application uses Amazon DynamoDB as its database layer. The company has configured DynamoDB with provisioned capacity set to 20% above the expected peak load of the application. On an average day, DynamoDB used 22 billion read capacity units (RCUs) and 60 billion write capacity units (WCUs). The application is running well. Usage changes smoothly over the course of the day and is generally shaped like a bell curve. The timing and magnitude of peaks vary based on the weather and season, but the general shape is consistent.

Which solution will provide the MOST cost optimization of the DynamoDB database layer?

- A. Change the DynamoDB tables to use on-demand capacity.
- B. Use AWS Auto Scaling and configure time-based scaling.
- C. Enable DynamoDB capacity-based auto scaling.
- D. Enable DynamoDB Accelerator (DAX).

Answer: C

Explanation:

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/AutoScaling.html>

NEW QUESTION 54

A company is deploying a solution in Amazon Aurora by migrating from an on-premises system. The IT department has established an AWS Direct Connect link from the company's data center. The company's Database Specialist has selected the option to require SSL/TLS for connectivity to prevent plaintext data from being set over the network. The migration appears to be working successfully, and the data can be queried from a desktop machine.

Two Data Analysts have been asked to query and validate the data in the new Aurora DB cluster. Both Analysts are unable to connect to Aurora. Their user names and passwords have been verified as valid and the Database Specialist can connect to the DB cluster using their accounts. The Database Specialist also verified that the security group configuration allows network from all corporate IP addresses.

What should the Database Specialist do to correct the Data Analysts' inability to connect?

- A. Restart the DB cluster to apply the SSL change.
- B. Instruct the Data Analysts to download the root certificate and use the SSL certificate on the connection string to connect.
- C. Add explicit mappings between the Data Analysts' IP addresses and the instance in the security group assigned to the DB cluster.
- D. Modify the Data Analysts' local client firewall to allow network traffic to AWS.

Answer: B

Explanation:

- To connect using SSL:
- Provide the SSLTrust certificate (can be downloaded from AWS)
- Provide SSL options when connecting to database
- Not using SSL on a DB that enforces SSL would result in error <https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/ssl-certificate-rotation-aurora-postgresql.ht>

NEW QUESTION 55

A business's production database is hosted on a single-node Amazon RDS for MySQL DB instance. The database instance is hosted in a United States AWS Region.

A week before a significant sales event, a fresh database maintenance update is released. The maintenance update has been designated as necessary. The firm want to minimize the database instance's downtime and requests that a database expert make the database instance highly accessible until the sales event concludes.

Which solution will satisfy these criteria?

- A. Defer the maintenance update until the sales event is over.
- B. Create a read replica with the latest updat
- C. Initiate a failover before the sales event.
- D. Create a read replica with the latest updat
- E. Transfer all read-only traffic to the read replica during the sales event.
- F. Convert the DB instance into a Multi-AZ deploymen
- G. Apply the maintenance update.

Answer: D

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/rds-required-maintenance/>

NEW QUESTION 59

A company is closing one of its remote data centers. This site runs a 100 TB on-premises data warehouse solution. The company plans to use the AWS Schema Conversion Tool (AWS SCT) and AWS DMS for the migration to AWS. The site network bandwidth is 500 Mbps. A Database Specialist wants to migrate the on-premises data using Amazon S3 as the data lake and Amazon Redshift as the data warehouse. This move must take place during a 2-week period when source systems are shut down for maintenance. The data should stay encrypted at rest and in transit.

Which approach has the least risk and the highest likelihood of a successful data transfer?

- A. Set up a VPN tunnel for encrypting data over the network from the data center to AW
- B. Leverage AWS SCT and apply the converted schema to Amazon Redshif
- C. Once complete, start an AWS DMS task tomove the data from the source to Amazon S3. Use AWS Glue to load the data from Amazon S3 to Amazon Redshift.
- D. Leverage AWS SCT and apply the converted schema to Amazon Redshif
- E. Start an AWS DMS task with two AWS Snowball Edge devices to copy data from on-premises to Amazon S3 with AWS KMS encryptio
- F. Use AWS DMS to finish copying data to Amazon Redshift.
- G. Leverage AWS SCT and apply the converted schema to Amazon Redshif
- H. Once complete, use a fleet of 10 TB dedicated encrypted drives using the AWS Import/Export feature to copy data from on-premises to Amazon S3 with AWS KMS encryptio
- I. Use AWS Glue to load the data to Amazon redshift.
- J. Set up a VPN tunnel for encrypting data over the network from the data center to AW
- K. Leverage a native database export feature to export the data and compress the file
- L. Use the aws S3 cp multi-port upload command to upload these files to Amazon S3 with AWS KMS encryptio

M. Once complete, load the data to Amazon Redshift using AWS Glue.

Answer: B

Explanation:

<https://aws.amazon.com/blogs/database/new-aws-dms-and-aws-snowball-integration-enables-mass-database-mi>

NEW QUESTION 60

A company's ecommerce website uses Amazon DynamoDB for purchase orders. Each order is made up of a Customer ID and an Order ID. The DynamoDB table uses the Customer ID as the partition key and the Order ID as the sort key.

To meet a new requirement, the company also wants the ability to query the table by using a third attribute named Invoice ID. Queries using the Invoice ID must be strongly consistent. A database specialist must provide this capability with optimal performance and minimal overhead.

What should the database administrator do to meet these requirements?

- A. Add a global secondary index on Invoice ID to the existing table.
- B. Add a local secondary index on Invoice ID to the existing table.
- C. Recreate the table by using the latest snapshot while adding a local secondary index on Invoice ID.
- D. Use the partition key and a FilterExpression parameter with a filter on Invoice ID for all queries.

Answer: C

Explanation:

as Local secondary index can only be created while creating the Dynamodb table. and query needs to use third attribute on top of primary and sort key, so Local Secondary index has primary and sort key as well as the third attribute. Global secondary index can be created without primary and sort key

NEW QUESTION 64

A company is running a two-tier ecommerce application in one AWS account. The web server is deployed using an Amazon RDS for MySQL Multi-AZ DB instance. A Developer mistakenly deleted the database in the production environment. The database has been restored, but this resulted in hours of downtime and lost revenue.

Which combination of changes in existing IAM policies should a Database Specialist make to prevent an error like this from happening in the future? (Choose three.)

- A. Grant least privilege to groups, users, and roles
- B. Allow all users to restore a database from a backup that will reduce the overall downtime to restore the database
- C. Enable multi-factor authentication for sensitive operations to access sensitive resources and API operations
- D. Use policy conditions to restrict access to selective IP addresses
- E. Use AccessList Controls policy type to restrict users for database instance deletion
- F. Enable AWS CloudTrail logging and Enhanced Monitoring

Answer: ACD

Explanation:

<https://aws.amazon.com/blogs/database/using-iam-multifactor-authentication-with-amazon-rds/>

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/security_iam_id-based-policy.htmlhttps://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/security_iam_id-based-policy.html

NEW QUESTION 66

A retail company with its main office in New York and another office in Tokyo plans to build a database solution on AWS. The company's main workload consists of a mission-critical application that updates its application data in a data store. The team at the Tokyo office is building dashboards with complex analytical queries using the application data. The dashboards will be used to make buying decisions, so they need to have access to the application data in less than 1 second.

Which solution meets these requirements?

- A. Use an Amazon RDS DB instance deployed in the us-east-1 Region with a read replica instance in the ap-northeast-1 Region
- B. Create an Amazon ElastiCache cluster in the ap-northeast-1 Region to cache application data from the replica to generate the dashboards.
- C. Use an Amazon DynamoDB global table in the us-east-1 Region with replication into the ap-northeast-1 Region
- D. Use Amazon QuickSight for displaying dashboard results.
- E. Use an Amazon RDS for MySQL DB instance deployed in the us-east-1 Region with a read replica instance in the ap-northeast-1 Region
- F. Have the dashboard application read from the read replica.
- G. Use an Amazon Aurora global database
- H. Deploy the writer instance in the us-east-1 Region and the replica in the ap-northeast-1 Region
- I. Have the dashboard application read from the replica in the ap-northeast-1 Region.

Answer: D

Explanation:

<https://aws.amazon.com/blogs/database/aurora-postgresql-disaster-recovery-solutions-using-amazon-aurora-global>

NEW QUESTION 69

A gaming company wants to deploy a game in multiple Regions. The company plans to save local high scores in Amazon DynamoDB tables in each Region. A Database Specialist needs to design a solution to automate the deployment of the database with identical configurations in additional Regions, as needed. The solution should also automate configuration changes across all Regions.

Which solution would meet these requirements and deploy the DynamoDB tables?

- A. Create an AWS CLI command to deploy the DynamoDB table to all the Regions and save it for future deployments.
- B. Create an AWS CloudFormation template and deploy the template to all the Regions.
- C. Create an AWS CloudFormation template and use a stack set to deploy the template to all the Regions.
- D. Create DynamoDB tables using the AWS Management Console in all the Regions and create a step-by-step guide for future deployments.

Answer: C

Explanation:

<https://aws.amazon.com/blogs/aws/use-cloudformation-stacksets-to-provision-resources-across-multiple-aws-ac>
<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/stacksets-concepts.html>

NEW QUESTION 74

An information management services company is storing JSON documents on premises. The company is using a MongoDB 3.6 database but wants to migrate to AWS. The solution must be compatible, scalable, and fully managed. The solution also must result in as little downtime as possible during the migration. Which solution meets these requirements?

- A. Create an AWS Database Migration Service (AWS DMS) replication instance, a source endpoint for MongoDB, and a target endpoint of Amazon DocumentDB (with MongoDB compatibility).
- B. Create an AWS Database Migration Service (AWS DMS) replication instance, a source endpoint for MongoDB, and a target endpoint of a MongoDB image that is hosted on Amazon EC2
- C. Use the mongodump and mongorestore tools to migrate the data from the source MongoDB deployment to Amazon DocumentDB (with MongoDB compatibility).
- D. Use the mongodump and mongorestore tools to migrate the data from the source MongoDB deployment to a MongoDB image that is hosted on Amazon EC2.

Answer: A

Explanation:

<https://docs.aws.amazon.com/documentdb/latest/developerguide/docdb-migration.html#docdb-migration-approa>

NEW QUESTION 75

A company is using Amazon Redshift as its data warehouse solution. The Redshift cluster handles the following types of workloads:

- *Real-time inserts through Amazon Kinesis Data Firehose
- *Bulk inserts through COPY commands from Amazon S3
- *Analytics through SQL queries

Recently, the cluster has started to experience performance issues.

Which combination of actions should a database specialist take to improve the cluster's performance? (Choose three.)

- A. Modify the Kinesis Data Firehose delivery stream to stream the data to Amazon S3 with a high buffer size and to load the data into Amazon Redshift by using the COPY command.
- B. Stream real-time data into Redshift temporary tables before loading the data into permanent tables.
- C. For bulk inserts, split input files on Amazon S3 into multiple files to match the number of slices on Amazon Redshift
- D. Then use the COPY command to load data into Amazon Redshift.
- E. For bulk inserts, use the parallel parameter in the COPY command to enable multi-threading.
- F. Optimize analytics SQL queries to use sort keys.
- G. Avoid using temporary tables in analytics SQL queries.

Answer: BCE

Explanation:

<https://aws.amazon.com/blogs/big-data/top-10-performance-tuning-techniques-for-amazon-redshift/> Tip #6: Improving the efficiency of temporary tables
Tip #9: Maintaining efficient data loads

Amazon Redshift best practices suggest using the COPY command to perform data loads of file-based data. Tip #3: Sort key recommendation

Sorting a table on an appropriate sort key can accelerate query performance, especially queries with range-restricted predicates, by requiring fewer table blocks to be read from disk.

NEW QUESTION 77

A company uses Amazon Aurora for secure financial transactions. The data must always be encrypted at rest and in transit to meet compliance requirements. Which combination of actions should a database specialist take to meet these requirements? (Choose two.)

- A. Create an Aurora Replica with encryption enabled using AWS Key Management Service (AWS KMS). Then promote the replica to master.
- B. Use SSL/TLS to secure the in-transit connection between the financial application and the Aurora DB cluster.
- C. Modify the existing Aurora DB cluster and enable encryption using an AWS Key Management Service (AWS KMS) encryption key
- D. Apply the changes immediately.
- E. Take a snapshot of the Aurora DB cluster and encrypt the snapshot using an AWS Key Management Service (AWS KMS) encryption key
- F. Restore the snapshot to a new DB cluster and update the financial application database endpoints.
- G. Use AWS Key Management Service (AWS KMS) to secure the in-transit connection between the financial application and the Aurora DB cluster.

Answer: AB

Explanation:

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/aurora-replicas-adding.html>

NEW QUESTION 82

A company is about to launch a new product, and test databases must be re-created from production data. The company runs its production databases on an Amazon Aurora MySQL DB cluster. A Database Specialist needs to deploy a solution to create these test databases as quickly as possible with the least amount of administrative effort.

What should the Database Specialist do to meet these requirements?

- A. Restore a snapshot from the production cluster into test clusters
- B. Create logical dumps of the production cluster and restore them into new test clusters
- C. Use database cloning to create clones of the production cluster
- D. Add an additional read replica to the production cluster and use that node for testing

Answer: C

Explanation:

<https://aws.amazon.com/getting-started/hands-on/aurora-cloning-backtracking/>

"Cloning an Aurora cluster is extremely useful if you want to assess the impact of changes to your database, or if you need to perform workload-intensive operations—such as exporting data or running analytical queries, or simply if you want to use a copy of your production database in a development or testing environment. You can make multiple clones of your Aurora DB cluster. You can even create additional clones from other clones, with the constraint that the clone databases must be created in the same region as the source databases.

NEW QUESTION 84

Application developers have reported that an application is running slower as more users are added. The application database is running on an Amazon Aurora DB cluster with an Aurora Replica. The application is written to take advantage of read scaling through reader endpoints. A database specialist looks at the performance metrics of the database and determines that, as new users were added to the database, the primary instance CPU utilization steadily increased while the Aurora Replica CPU utilization remained steady.

How can the database specialist improve database performance while ensuring minimal downtime?

- A. Modify the Aurora DB cluster to add more replicas until the overall load stabilize
- B. Then, reduce the number of replicas once the application meets service level objectives.
- C. Modify the primary instance to a larger instance size that offers more CPU capacity.
- D. Modify a replica to a larger instance size that has more CPU capacity
- E. Then, promote the modified replica.
- F. Restore the Aurora DB cluster to one that has an instance size with more CPU capacity
- G. Then, swap the names of the old and new DB clusters.

Answer: C

NEW QUESTION 87

A ride-hailing application stores bookings in a persistent Amazon RDS for MySQL DB instance. This program is very popular, and the corporation anticipates a tenfold rise in the application's user base over the next several months. The application receives a higher volume of traffic in the morning and evening.

This application is divided into two sections:

An internal booking component that takes online reservations in response to concurrent user queries. A component of a third-party customer relationship management (CRM) system that customer service professionals utilize. Booking data is accessed using queries in the CRM.

To manage this workload effectively, a database professional must create a cost-effective database system. Which solution satisfies these criteria?

- A. Use Amazon ElastiCache for Redis to accept the booking
- B. Associate an AWS Lambda function to capture changes and push the booking data to the RDS for MySQL DB instance used by the CRM.
- C. Use Amazon DynamoDB to accept the booking
- D. Enable DynamoDB Streams and associate an AWS Lambda function to capture changes and push the booking data to an Amazon SQS queue
- E. This triggers another Lambda function that pulls data from Amazon SQS and writes it to the RDS for MySQL DB instance used by the CRM.
- F. Use Amazon ElastiCache for Redis to accept the booking
- G. Associate an AWS Lambda function to capture changes and push the booking data to an Amazon Redshift database used by the CRM.
- H. Use Amazon DynamoDB to accept the booking
- I. Enable DynamoDB Streams and associate an AWS Lambda function to capture changes and push the booking data to Amazon Athena, which is used by the CRM.

Answer: B

Explanation:

"AWS Lambda function to capture changes" capture changes to what? ElastiCache? The main use of ElastiCache is to cache frequently read data. Also "the company expects a tenfold increase in the user base" and "correspond to simultaneous requests from users"

NEW QUESTION 92

A company is going to use an Amazon Aurora PostgreSQL DB cluster for an application backend. The DB cluster contains some tables with sensitive data. A Database Specialist needs to control the access privileges at the table level.

How can the Database Specialist meet these requirements?

- A. Use AWS IAM database authentication and restrict access to the tables using an IAM policy.
- B. Configure the rules in a NACL to restrict outbound traffic from the Aurora DB cluster.
- C. Execute GRANT and REVOKE commands that restrict access to the tables containing sensitive data.
- D. Define access privileges to the tables containing sensitive data in the pg_hba.conf file.

Answer: C

NEW QUESTION 96

A database specialist wants to ensure that an Amazon Aurora DB cluster is always automatically upgraded to the most recent minor version available. Noticing that there is a new minor version available, the database specialist has issues an AWS CLI command to enable automatic minor version updates. The command runs successfully, but checking the Aurora DB cluster indicates that no update to the Aurora version has been made.

What might account for this? (Choose two.)

- A. The new minor version has not yet been designated as preferred and requires a manual upgrade.
- B. Configuring automatic upgrades using the AWS CLI is not supported
- C. This must be enabled expressly using the AWS Management Console.
- D. Applying minor version upgrades requires sufficient free space.
- E. The AWS CLI command did not include an apply-immediately parameter.
- F. Aurora has detected a breaking change in the new minor version and has automatically rejected the upgrade.

Answer: AD

Explanation:

"When Amazon RDS designates a minor engine version as the preferred minor engine version, each database that meets both of the following conditions is upgraded to the minor engine version automatically"

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_UpgradeDBInstance.Upgrading.html

Call the modify-db-instance Amazon CLI command. Specify the name of your DB instance for the --db-instance-identifier option and true for the --auto-minor-version-upgrade option. Optionally, specify the --apply-immediately option to immediately enable this setting for your DB instance. Run a separate modify-db-instance command for each DB instance in the cluster.

https://docs.amazonaws.cn/en_us/AmazonRDS/latest/AuroraUserGuide/AuroraMySQL.Updates.Patching.html#

NEW QUESTION 99

A database specialist needs to delete user data and sensor data 1 year after it was loaded in an Amazon DynamoDB table. TTL is enabled on one of the attributes. The database specialist monitors TTL rates on the Amazon CloudWatch metrics for the table and observes that items are not being deleted as expected.

What is the MOST likely reason that the items are not being deleted?

- A. The TTL attribute's value is set as a Number data type.
- B. The TTL attribute's value is set as a Binary data type.
- C. The TTL attribute's value is a timestamp in the Unix epoch time format in seconds.
- D. The TTL attribute's value is set with an expiration of 1 year.

Answer: B

Explanation:

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/TTL.html#time-to-live-ttl-before-you-sta>

NEW QUESTION 100

A company wants to improve its ecommerce website on AWS. A database specialist decides to add Amazon ElastiCache for Redis in the implementation stack to ease the workload off the database and shorten the website response times. The database specialist must also ensure the ecommerce website is highly available within the company's AWS Region.

How should the database specialist deploy ElastiCache to meet this requirement?

- A. Launch an ElastiCache for Redis cluster using the AWS CLI with the -cluster-enabled switch.
- B. Launch an ElastiCache for Redis cluster and select read replicas in different Availability Zones.
- C. Launch two ElastiCache for Redis clusters in two different Availability Zone
- D. Configure Redis streams to replicate the cache from the primary cluster to another.
- E. Launch an ElastiCache cluster in the primary Availability Zone and restore the cluster's snapshot to a different Availability Zone during disaster recovery.

Answer: B

Explanation:

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

You can enable Multi-AZ only on Redis (cluster mode disabled) clusters that have at least one available read replica. Clusters without read replicas do not provide high availability or fault tolerance.

NEW QUESTION 103

A company uses an Amazon RDS for PostgreSQL DB instance for its customer relationship management (CRM) system. New compliance requirements specify that the database must be encrypted at rest.

Which action will meet these requirements?

- A. Create an encrypted copy of manual snapshot of the DB instanc
- B. Restore a new DB instance from the encrypted snapshot.
- C. Modify the DB instance and enable encryption.
- D. Restore a DB instance from the most recent automated snapshot and enable encryption.
- E. Create an encrypted read replica of the DB instanc
- F. Promote the read replica to a standalone instance.

Answer: A

Explanation:

<https://docs.aws.amazon.com/prescriptive-guidance/latest/patterns/encrypt-an-existing-amazon-rds-for-postgresq> You can enable encryption for an Amazon RDS DB instance when you create it, but not after it's created.

However, you can add encryption to an unencrypted DB instance by creating a snapshot of your DB instance, and then creating an encrypted copy of that snapshot. You can then restore a DB instance from the encrypted snapshot to get an encrypted copy of your original DB instance. The pattern uses AWS Database Migration Service (AWS DMS) to migrate data and AWS Key Management Service (AWS KMS) for encryption.

NEW QUESTION 108

A financial services company has an application deployed on AWS that uses an Amazon Aurora PostgreSQL DB cluster. A recent audit showed that no log files contained database administrator activity. A database specialist needs to recommend a solution to provide database access and activity logs. The solution should use the least amount of effort and have a minimal impact on performance.

Which solution should the database specialist recommend?

- A. Enable Aurora Database Activity Streams on the database in synchronous mod
- B. Connect the Amazon Kinesis data stream to Kinesis Data Firehos
- C. Set the Kinesis Data Firehose destination to an Amazon S3 bucket.
- D. Create an AWS CloudTrail trail in the Region where the database run
- E. Associate the database activity logs with the trail.
- F. Enable Aurora Database Activity Streams on the database in asynchronous mod
- G. Connect the Amazon Kinesis data stream to Kinesis Data Firehos
- H. Set the Firehose destination to an Amazon S3 bucket.
- I. Allow connections to the DB cluster through a bastion host onl
- J. Restrict database access to the bastion host and application server

K. Push the bastion host logs to Amazon CloudWatch Logs using the CloudWatch Logs agent.

Answer: C

Explanation:

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/DBActivityStreams.Overview.html>

NEW QUESTION 110

A bank intends to utilize Amazon RDS to host a MySQL database instance. The database should be able to handle high-volume read requests with extremely few repeated queries.

Which solution satisfies these criteria?

- A. Create an Amazon ElastiCache cluster
- B. Use a write-through strategy to populate the cache.
- C. Create an Amazon ElastiCache cluster
- D. Use a lazy loading strategy to populate the cache.
- E. Change the DB instance to Multi-AZ with a standby instance in another AWS Region.
- F. Create a read replica of the DB instance
- G. Use the read replica to distribute the read traffic.

Answer: D

NEW QUESTION 111

A manufacturing company's website uses an Amazon Aurora PostgreSQL DB cluster.

Which configurations will result in the LEAST application downtime during a failover? (Choose three.)

- A. Use the provided read and write Aurora endpoints to establish a connection to the Aurora DB cluster.
- B. Create an Amazon CloudWatch alert triggering a restore in another Availability Zone when the primary Aurora DB cluster is unreachable.
- C. Edit and enable Aurora DB cluster cache management in parameter groups.
- D. Set TCP keepalive parameters to a high value.
- E. Set JDBC connection string timeout variables to a low value.
- F. Set Java DNS caching timeouts to a high value.

Answer: ABC

NEW QUESTION 116

A Database Specialist is setting up a new Amazon Aurora DB cluster with one primary instance and three Aurora Replicas for a highly intensive, business-critical application. The Aurora DB cluster has one medium-sized primary instance, one large-sized replica, and two medium-sized replicas. The Database Specialist did not assign a promotion tier to the replicas.

In the event of a primary failure, what will occur?

- A. Aurora will promote an Aurora Replica that is of the same size as the primary instance
- B. Aurora will promote an arbitrary Aurora Replica
- C. Aurora will promote the largest-sized Aurora Replica
- D. Aurora will not promote an Aurora Replica

Answer: C

Explanation:

Priority: If you don't select a value, the default is tier-1. This priority determines the order in which Aurora

https://docs.amazonaws.cn/en_us/AmazonRDS/latest/AuroraUserGuide/aurora-replicas-adding.html

More than one Aurora Replica can share the same priority, resulting in promotion tiers. If two or more Aurora Replicas share the same priority, then Amazon RDS promotes the replica that is largest in size. If two or more Aurora Replicas share the same priority and size, then Amazon RDS promotes an arbitrary replica in the same promotion tier.

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Aurora.Managing.Backups.html#Aurora.M> If two or more Aurora Replicas share the same priority, then Amazon RDS promotes the replica that is largest in size. If two or more Aurora Replicas share the same priority and size, then Amazon RDS promotes an

arbitrary replica in the same promotion tier. <https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Concepts.AuroraHighAvailability.html>

NEW QUESTION 121

A vehicle insurance company needs to choose a highly available database to track vehicle owners and their insurance details. The persisted data should be immutable in the database, including the complete and sequenced history of changes over time with all the owners and insurance transfer details for a vehicle. The data should be easily verifiable for the data lineage of an insurance claim. Which approach meets these requirements with MINIMAL effort?

- A. Create a blockchain to store the insurance detail
- B. Validate the data using a hash function to verify the data lineage of an insurance claim.
- C. Create an Amazon DynamoDB table to store the insurance detail
- D. Validate the data using AWS DMS validation by moving the data to Amazon S3 to verify the data lineage of an insurance claim.
- E. Create an Amazon QLDB ledger to store the insurance detail
- F. Validate the data by choosing the ledger name in the digest request to verify the data lineage of an insurance claim.
- G. Create an Amazon Aurora database to store the insurance detail
- H. Validate the data using AWS DMS validation by moving the data to Amazon S3 to verify the data lineage of an insurance claim.

Answer: C

NEW QUESTION 124

A company has two separate AWS accounts: one for the business unit and another for corporate analytics. The company wants to replicate the business unit data stored in Amazon RDS for MySQL in us-east-1 to its corporate analytics Amazon Redshift environment in us-west-1. The company wants to use AWS DMS with

Amazon RDS as the source endpoint and Amazon Redshift as the target endpoint.
Which action will allow AVS DMS to perform the replication?

- A. Configure the AWS DMS replication instance in the same account and Region as Amazon Redshift.
- B. Configure the AWS DMS replication instance in the same account as Amazon Redshift and in the same Region as Amazon RDS.
- C. Configure the AWS DMS replication instance in its own account and in the same Region as AmazonRedshift.
- D. Configure the AWS DMS replication instance in the same account and Region as Amazon RDS.

Answer: A

Explanation:

https://docs.aws.amazon.com/dms/latest/userguide/CHAP_Target.Redshift.html

NEW QUESTION 125

A significant automotive manufacturer is switching a mission-critical finance application's database to Amazon DynamoDB. According to the company's risk and compliance policy, any update to the database must be documented as a log entry for auditing purposes. Each minute, the system anticipates about 500,000 log entries. Log entries should be kept in Apache Parquet files in batches of at least 100,000 records per file.
How could a database professional approach these needs while using DynamoDB?

- A. Enable Amazon DynamoDB Streams on the tabl
- B. Create an AWS Lambda function triggered by the strea
- C. Write the log entries to an Amazon S3 object.
- D. Create a backup plan in AWS Backup to back up the DynamoDB table once a da
- E. Create an AWS Lambda function that restores the backup in another table and compares both tables for change
- F. Generate the log entries and write them to an Amazon S3 object.
- G. Enable AWS CloudTrail logs on the tabl
- H. Create an AWS Lambda function that reads the log files once an hour and filters DynamoDB API action
- I. Write the filtered log files to Amazon S3.
- J. Enable Amazon DynamoDB Streams on the tabl
- K. Create an AWS Lambda function triggered by the strea
- L. Write the log entries to an Amazon Kinesis Data Firehose delivery stream with buffering and Amazon S3 as the destination.

Answer: D

NEW QUESTION 127

A retail company manages a web application that stores data in an Amazon DynamoDB table. The company is undergoing account consolidation efforts. A database engineer needs to migrate the DynamoDB table from the current AWS account to a new AWS account.
Which strategy meets these requirements with the LEAST amount of administrative work?

- A. Use AWS Glue to crawl the data in the DynamoDB tabl
- B. Create a job using an available blueprint to export the data to Amazon S3. Import the data from the S3 file to a DynamoDB table in the new account.
- C. Create an AWS Lambda function to scan the items of the DynamoDB table in the current account and write to a file in Amazon S3. Create another Lambda function to read the S3 file and restore the items of a DynamoDB table in the new account.
- D. Use AWS Data Pipeline in the current account to export the data from the DynamoDB table to a file in Amazon S3. Use Data Pipeline to import the data from the S3 file to a DynamoDB table in the new account.
- E. Configure Amazon DynamoDB Streams for the DynamoDB table in the current accoun
- F. Create an AWS Lambda function to read from the stream and write to a file in Amazon S3. Create another Lambda functionto read the S3 file and restore the items to a DynamoDB table in the new account.

Answer: C

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/dynamodb-cross-account-migration/> <https://aws.amazon.com/premiumsupport/knowledge-center/data-pipeline-account-access-dynamodb-s3/>

NEW QUESTION 132

A manufacturing company has an. inventory system that stores information in an Amazon Aurora MySQL DB cluster. The database tables are partitioned. The database size has grown to 3 TB. Users run one-time queries by using a SQL client. Queries that use an equijoin to join large tables are taking a long time to run.
Which action will improve query performance with the LEAST operational effort?

- A. Migrate the database to a new Amazon Redshift data warehouse.
- B. Enable hash joins on the database by setting the variable optimizer_switch to hash_join=on.
- C. Take a snapshot of the DB cluste
- D. Create a new DB instance by using the snapshot, and enable parallel query mode.
- E. Add an Aurora read replica.

Answer: B

Explanation:

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraMySQL.BestPractices.html>

NEW QUESTION 133

A media company is using Amazon RDS for PostgreSQL to store user data. The RDS DB instance currently has a publicly accessible setting enabled and is hosted in a public subnet. Following a recent AWS Well- Architected Framework review, a Database Specialist was given new security requirements. Only certain on-premises corporate network IPs should connect to the DB instance. Connectivity is allowed from the corporate network only.
Which combination of steps does the Database Specialist need to take to meet these new requirements? (Choose three.)

- A. Modify the pg_hba.conf fil
- B. Add the required corporate network IPs and remove the unwanted IPs.

- C. Modify the associated security group.
- D. Add the required corporate network IPs and remove the unwanted IPs.
- E. Move the DB instance to a private subnet using AWS DMS.
- F. Enable VPC peering between the application host running on the corporate network and the VPC associated with the DB instance.
- G. Disable the publicly accessible setting.
- H. Connect to the DB instance using private IPs and a VPN.

Answer: BEF

Explanation:

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_VPC.WorkingWithRDSInstanceinaVPC.html

NEW QUESTION 135

The Amazon CloudWatch metric for FreeLocalStorage on an Amazon Aurora MySQL DB instance shows that the amount of local storage is below 10 MB. A database engineer must increase the local storage available in the Aurora DB instance. How should the database engineer meet this requirement?

- A. Modify the DB instance to use an instance class that provides more local SSD storage.
- B. Modify the Aurora DB cluster to enable automatic volume resizing.
- C. Increase the local storage by upgrading the database engine version.
- D. Modify the DB instance and configure the required storage volume in the configuration section.

Answer: A

Explanation:

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Aurora.AuroraMySQL.Monitoring.Metrics.html>. Unlike for other DB engines, for Aurora DB instances this metric reports the amount of storage available to each DB instance. This value depends on the DB instance class (for pricing information, see the Amazon RDS product page). You can increase the amount of free storage space for an instance by choosing a larger DB instance class for your instance."

NEW QUESTION 139

A company is using an Amazon ElastiCache for Redis cluster to host its online shopping website. Shoppers receive the following error when the website's application queries the cluster:

```
OOM command not allowed when used memory > 'maxmemory'
```

Which solutions will resolve this memory issues with the LEAST amount of effort? (Choose three.)

- A. Reduce the TTL value for keys on the node.
- B. Choose a larger node type.
- C. Test different values in the parameter group for the maxmemory-policy parameter to find the ideal value to use.
- D. Increase the number of nodes.
- E. Monitor the EngineCPUUtilization Amazon CloudWatch metric.
- F. Create an AWS Lambda function to delete keys on nodes when a threshold is reached.
- G. Increase the TTL value for keys on the node.

Answer: ABC

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/oom-command-not-allowed-redis/>

NEW QUESTION 140

A company with 500,000 employees needs to supply its employee list to an application used by human resources. Every 30 minutes, the data is exported using the LDAP service to load into a new Amazon DynamoDB table. The data model has a base table with Employee ID for the partition key and a global secondary index with Organization ID as the partition key.

While importing the data, a database specialist receives ProvisionedThroughputExceededException errors.

After increasing the provisioned write capacity units

(WCUs) to 50,000, the specialist receives the same errors. Amazon CloudWatch metrics show a consumption of 1,500 WCUs.

What should the database specialist do to address the issue?

- A. Change the data model to avoid hot partitions in the global secondary index.
- B. Enable auto scaling for the table to automatically increase write capacity during bulk imports.
- C. Modify the table to use on-demand capacity instead of provisioned capacity.
- D. Increase the number of retries on the bulk loading application.

Answer: A

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/dynamodb-table-throttled/>

NEW QUESTION 145

A company is planning to close for several days. A Database Specialist needs to stop all applications along with the DB instances to ensure employees do not have access to the systems during this time. All databases are running on Amazon RDS for MySQL.

The Database Specialist wrote and executed a script to stop all the DB instances. When reviewing the logs, the Database Specialist found that Amazon RDS DB instances with read replicas did not stop.

How should the Database Specialist edit the script to fix this issue?

- A. Stop the source instances before stopping their read replicas
- B. Delete each read replica before stopping its corresponding source instance
- C. Stop the read replicas before stopping their source instances
- D. Use the AWS CLI to stop each read replica and source instance at the same time

Answer: B

Explanation:

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_StopInstance.html

"The following are some limitations to stopping and starting a DB instance: You can't stop a DB instance that has a read replica, or that is a read replica." So if you can't stop a DB with a read replica, you have to delete the read replica first to then stop it???

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_MySQL.Replication.ReadReplicas.html#U

NEW QUESTION 147

A large retail company recently migrated its three-tier ecommerce applications to AWS. The company's backend database is hosted on Amazon Aurora PostgreSQL. During peak times, users complain about longer page load times. A database specialist reviewed Amazon RDS Performance Insights and found a spike in IO:XactSync wait events. The SQL attached to the wait events are all single INSERT statements.

How should this issue be resolved?

- A. Modify the application to commit transactions in batches
- B. Add a new Aurora Replica to the Aurora DB cluster.
- C. Add an Amazon ElastiCache for Redis cluster and change the application to write through.
- D. Change the Aurora DB cluster storage to Provisioned IOPS (PIOPS).

Answer: A

Explanation:

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraPostgreSQL.Reference.html> "This wait most often arises when there is a very high rate of commit activity on the system. You can

sometimes alleviate this wait by modifying applications to commit transactions in batches. "

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/apg-waits.xactsync.html>

NEW QUESTION 149

A ride-hailing application uses an Amazon RDS for MySQL DB instance as persistent storage for bookings. This application is very popular and the company expects a tenfold increase in the user base in next few months. The application experiences more traffic during the morning and evening hours.

This application has two parts:

- An in-house booking component that accepts online bookings that directly correspond to simultaneous requests from users.
- A third-party customer relationship management (CRM) component used by customer care representatives. The CRM uses queries to access booking data.

A database specialist needs to design a cost-effective database solution to handle this workload. Which solution meets these requirements?

- A. Use Amazon ElastiCache for Redis to accept the booking
- B. Associate an AWS Lambda function to capture changes and push the booking data to the RDS for MySQL DB instance used by the CRM.
- C. Use Amazon DynamoDB to accept the booking
- D. Enable DynamoDB Streams and associate an AWS Lambda function to capture changes and push the booking data to an Amazon SQS queue
- E. This triggers another Lambda function that pulls data from Amazon SQS and writes it to the RDS for MySQL DB instance used by the CRM.
- F. Use Amazon ElastiCache for Redis to accept the booking
- G. Associate an AWS Lambda function to capture changes and push the booking data to an Amazon Redshift database used by the CRM.
- H. Use Amazon DynamoDB to accept the booking
- I. Enable DynamoDB Streams and associate an AWS Lambda function to capture changes and push the booking data to Amazon Athena, which is used by the CRM.

Answer: D

NEW QUESTION 151

A large financial services company requires that all data be encrypted in transit. A Developer is attempting to connect to an Amazon RDS DB instance using the company VPC for the first time with credentials provided by a Database Specialist. Other members of the Development team can connect, but this user is consistently receiving an error indicating a communications link failure. The Developer asked the Database Specialist to reset the password a number of times, but the error persists.

Which step should be taken to troubleshoot this issue?

- A. Ensure that the database option group for the RDS DB instance allows ingress from the Developer machine's IP address
- B. Ensure that the RDS DB instance's subnet group includes a public subnet to allow the Developer to connect
- C. Ensure that the RDS DB instance has not reached its maximum connections limit
- D. Ensure that the connection is using SSL and is addressing the port where the RDS DB instance is listening for encrypted connections

Answer: D

Explanation:

<https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/SQLServer.Concepts.General.SSL.Using.html>

NEW QUESTION 152

A company is using an Amazon RDS for MySQL DB instance for its internal applications. A security audit shows that the DB instance is not encrypted at rest. The company's application team needs to encrypt the DB instance.

What should the team do to meet this requirement?

- A. Stop the DB instance and modify it to enable encryption
- B. Apply this setting immediately without waiting for the next scheduled RDS maintenance window.
- C. Stop the DB instance and create an encrypted snapshot
- D. Restore the encrypted snapshot to a new encrypted DB instance
- E. Delete the original DB instance, and update the applications to point to the new encrypted DB instance.
- F. Stop the DB instance and create a snapshot
- G. Copy the snapshot into another encrypted snapshot
- H. Restore the encrypted snapshot to a new encrypted DB instance

- I. Delete the original DB instance, and update the applications to point to the new encrypted DB instance.
- J. Create an encrypted read replica of the DB instance.
- K. Promote the read replica to master.
- L. Delete the original DB instance, and update the applications to point to the new encrypted DB instance.

Answer: C

NEW QUESTION 154

A company has a production environment running on Amazon RDS for SQL Server with an in-house web application as the front end. During the last application maintenance window, new functionality was added to the web application to enhance the reporting capabilities for management. Since the update, the application is slow to respond to some reporting queries.

How should the company identify the source of the problem?

- A. Install and configure Amazon CloudWatch Application Insights for Microsoft .NET and Microsoft SQL Server.
- B. Use a CloudWatch dashboard to identify the root cause.
- C. Enable RDS Performance Insights and determine which query is creating the problem.
- D. Request changes to the query to address the problem.
- E. Use AWS X-Ray deployed with Amazon RDS to track query system traces.
- F. Create a support request and work with AWS Support to identify the source of the issue.

Answer: B

Explanation:

Amazon RDS Performance Insights is a database performance tuning and monitoring feature that helps you quickly assess the load on your database, and determine when and where to take action. Performance Insights allows non-experts to detect performance problems with an easy-to-understand dashboard that visualizes database load. <https://aws.amazon.com/rds/performance-insights/>

NEW QUESTION 156

A gaming company is designing a mobile gaming app that will be accessed by many users across the globe. The company wants to have replication and full support for multi-master writes. The company also wants to ensure low latency and consistent performance for app users.

Which solution meets these requirements?

- A. Use Amazon DynamoDB global tables for storage and enable DynamoDB automatic scaling.
- B. Use Amazon Aurora for storage and enable cross-Region Aurora Replicas.
- C. Use Amazon Aurora for storage and cache the user content with Amazon ElastiCache.
- D. Use Amazon Neptune for storage.

Answer: A

NEW QUESTION 161

A company has a database fleet that includes an Amazon RDS for MySQL DB instance. During an audit, the company discovered that the data that is stored on the DB instance is unencrypted.

A database specialist must enable encryption for the DB instance. The database specialist also must encrypt all connections to the DB instance.

Which combination of actions should the database specialist take to meet these requirements? (Choose three.)

- A. In the RDS console, choose "Enable encryption" to encrypt the DB instance by using an AWS Key Management Service (AWS KMS) key.
- B. Encrypt the read replica of the unencrypted DB instance by using an AWS Key Management Service (AWS KMS) key.
- C. Fail over the read replica to the primary DB instance.
- D. Create a snapshot of the unencrypted DB instance.
- E. Encrypt the snapshot by using an AWS Key Management Service (AWS KMS) key.
- F. Restore the DB instance from the encrypted snapshot.
- G. Delete the original DB instance.
- H. Require SSL connections for applicable database user accounts.
- I. Use SSL/TLS from the application to encrypt a connection to the DB instance.
- J. Enable SSH encryption on the DB instance.

Answer: ACE

Explanation:

<https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Overview.Encryption.html#Overview.Encryption>.

NEW QUESTION 165

A company uses Amazon DynamoDB as the data store for its ecommerce website. The website receives little to no traffic at night, and the majority of the traffic occurs during the day. The traffic growth during peak hours is gradual and predictable on a daily basis, but it can be orders of magnitude higher than during off-peak hours.

The company initially provisioned capacity based on its average volume during the day without accounting for the variability in traffic patterns. However, the website is experiencing a significant amount of throttling during peak hours. The company wants to reduce the amount of throttling while minimizing costs.

What should a database specialist do to meet these requirements?

- A. Use reserved capacity.
- B. Set it to the capacity levels required for peak daytime throughput.
- C. Use provisioned capacity.
- D. Set it to the capacity levels required for peak daytime throughput.
- E. Use provisioned capacity.
- F. Create an AWS Application Auto Scaling policy to update capacity based on consumption.
- G. Use on-demand capacity.

Answer: C

Explanation:

On-demand mode is a good option if any of the following are true: You create new tables with unknown workloads. You have unpredictable application traffic. You prefer the ease of paying for only what you use.

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/HowItWorks.ReadWriteCapacityMode.h>

Amazon DynamoDB auto scaling uses the AWS Application Auto Scaling service to dynamically adjust provisioned throughput capacity on your behalf
<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/AutoScaling.html>

NEW QUESTION 166

A company is migrating a mission-critical 2-TB Oracle database from on premises to Amazon Aurora. The cost for the database migration must be kept to a minimum, and both the on-premises Oracle database and the Aurora DB cluster must remain open for write traffic until the company is ready to completely cut over to Aurora.

Which combination of actions should a database specialist take to accomplish this migration as quickly as possible? (Choose two.)

- A. Use the AWS Schema Conversion Tool (AWS SCT) to convert the source database schem
- B. Then restore the converted schema to the target Aurora DB cluster.
- C. Use Oracle's Data Pump tool to export a copy of the source database schema and manually edit the schema in a text editor to make it compatible with Aurora.
- D. Create an AWS DMS task to migrate data from the Oracle database to the Aurora DB cluste
- E. Select the migration type to replicate ongoing changes to keep the source and target databases in sync until the company is ready to move all user traffic to the Aurora DB cluster.
- F. Create an AWS DMS task to migrate data from the Oracle database to the Aurora DB cluste
- G. Once the initial load is complete, create an AWS Kinesis Data Firehose stream to perform change data capture (CDC) until the company is ready to move all user traffic to the Aurora DB cluster.
- H. Create an AWS Glue job and related resources to migrate data from the Oracle database to the Aurora DB cluster
- I. Once the initial load is complete, create an AWS DMS task to perform change data capture (CDC) until the company is ready to move all user traffic to the Aurora DB cluster.

Answer: AC

NEW QUESTION 170

A company requires near-real-time notifications when changes are made to Amazon RDS DB security groups. Which solution will meet this requirement with the LEAST operational overhead?

- A. Configure an RDS event notification subscription for DB security group events.
- B. Create an AWS Lambda function that monitors DB security group change
- C. Create an Amazon Simple Notification Service (Amazon SNS) topic for notification.
- D. Turn on AWS CloudTrai
- E. Configure notifications for the detection of changes to DB security groups.
- F. Configure an Amazon CloudWatch alarm for RDS metrics about changes to DB security groups.

Answer: A

Explanation:

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_Events.Messages.html#USER_Events.Mes

NEW QUESTION 174

A Database Specialist is constructing a new Amazon Neptune DB cluster and tries to load data from Amazon S3 using the Neptune bulk loader API. The Database Specialist is confronted with the following error message:

€Unable to establish a connection to the s3 endpoint. The source URL is s3://mybucket/graphdata/ and the region code is us-east-1. Kindly confirm your Configuration S3.

Which of the following activities should the Database Specialist take to resolve the issue? (Select two.)

- A. Check that Amazon S3 has an IAM role granting read access to Neptune
- B. Check that an Amazon S3 VPC endpoint exists
- C. Check that a Neptune VPC endpoint exists
- D. Check that Amazon EC2 has an IAM role granting read access to Amazon S3
- E. Check that Neptune has an IAM role granting read access to Amazon S3

Answer: BE

Explanation:

<https://docs.aws.amazon.com/neptune/latest/userguide/bulk-load-tutorial-IAM.html> <https://docs.aws.amazon.com/neptune/latest/userguide/bulk-load-data.html>

“An IAM role for the Neptune DB instance to assume that has an IAM policy that allows access to the data files in the S3 bucket. The policy must grant Read and List permissions.” “An Amazon S3 VPC endpoint. For more information, see the Creating an Amazon S3 VPC Endpoint section.”

NEW QUESTION 178

A business is transferring its on-premises database workloads to the Amazon Web Services (AWS) Cloud. A database professional migrating an Oracle database with a huge table to Amazon RDS has picked AWS DMS. The database professional observes that AWS DMS is consuming considerable time migrating the data. Which activities would increase the pace of data migration? (Select three.)

- A. Create multiple AWS DMS tasks to migrate the large table.
- B. Configure the AWS DMS replication instance with Multi-AZ.
- C. Increase the capacity of the AWS DMS replication server.
- D. Establish an AWS Direct Connect connection between the on-premises data center and AWS.
- E. Enable an Amazon RDS Multi-AZ configuration.
- F. Enable full large binary object (LOB) mode to migrate all LOB data for all large tables.

Answer: ACD

Explanation:

https://docs.aws.amazon.com/dms/latest/userguide/CHAP_Tasks.LOBSupport.html

NEW QUESTION 181

A company hosts a 2 TB Oracle database in its on-premises data center. A database specialist is migrating the database from on premises to an Amazon Aurora PostgreSQL database on AWS.

The database specialist identifies a problem that relates to compatibility Oracle stores metadata in its data dictionary in uppercase, but PostgreSQL stores the metadata in lowercase. The database specialist must resolve this problem to complete the migration.

What is the MOST operationally efficient solution that meets these requirements?

- A. Override the default uppercase format of Oracle schema by encasing object names in quotation marks during creation.
- B. Use AWS Database Migration Service (AWS DMS) mapping rules with rule-action as convert-lowercase.
- C. Use the AWS Schema Conversion Tool conversion agent to convert the metadata from uppercase to lowercase.
- D. Use an AWS Glue job that is attached to an AWS Database Migration Service (AWS DMS) replication task to convert the metadata from uppercase to lowercase.

Answer: B

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/dms-mapping-oracle-postgresql/>

NEW QUESTION 183

A database professional maintains a fleet of Amazon RDS database instances that are configured to utilize the default database parameter group. A database expert must connect a custom parameter group with certain database instances.

When will the instances be allocated to this new parameter group once the database specialist performs this change?

- A. Instantaneously after the change is made to the parameter group
- B. In the next scheduled maintenance window of the DB instances
- C. After the DB instances are manually rebooted
- D. Within 24 hours after the change is made to the parameter group

Answer: C

Explanation:

When you associate a new DB parameter group with a DB instance, the modified static and dynamic parameters are applied only after the DB instance is rebooted.

NEW QUESTION 188

A database specialist must create nightly backups of an Amazon DynamoDB table in a mission-critical workload as part of a disaster recovery strategy.

Which backup methodology should the database specialist use to MINIMIZE management overhead?

- A. Install the AWS CLI on an Amazon EC2 instanc
- B. Write a CLI command that creates a backup of theDynamoDB tabl
- C. Create a scheduled job or task that executes the command on a nightly basis.
- D. Create an AWS Lambda function that creates a backup of the DynamoDB tabl
- E. Create an Amazon CloudWatch Events rule that executes the Lambda function on a nightly basis.
- F. Create a backup plan using AWS Backup, specify a backup frequency of every 24 hours, and give the plan a nightly backup window.
- G. Configure DynamoDB backup and restore for an on-demand backup frequency of every 24 hours.

Answer: C

Explanation:

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/CreateBackup.html#:~:text=If%20you%2>

https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/backuprestore_HowItWorks.html

NEW QUESTION 190

A company's database specialist is building an Amazon RDS for Microsoft SQL Server DB instance to store hundreds of records in CSV format. A customer service tool uploads the records to an Amazon S3 bucket.

An employee who previously worked at the company already created a custom stored procedure to map the necessary CSV fields to the database tables. The database specialist needs to implement a solution that reuses this previous work and minimizes operational overhead.

Which solution will meet these requirements?

- A. Create an Amazon S3 event to invoke an AWS Lambda functio
- B. Configure the Lambda function to parse the .csv file and use a SQL client library to run INSERT statements to load the data into the tables.
- C. Write a custom .NET app that is hosted on Amazon EC2. Configure the .NET app to load the .csv file and call the custom stored procedure to insert the data into the tables.
- D. Download the .csv file from Amazon S3 to the RDS D drive by using an AWS msdb stored procedure.Call the custom stored procedure to insert the data from the RDS D drive into the tables.
- E. Create an Amazon S3 event to invoke AWS Step Functions to parse the .csv file and call the custom stored procedure to insert the data into the tables.

Answer: C

Explanation:

Step 1: Download S3 Files

Amazon RDS for SQL Server comes with several custom stored procedures and functions. These are located in the msdb database. The stored procedure to download files from S3 is "rds_download_from_s3". The syntax for this stored procedure is shown here:

```
exec msdb.dbo.rds_download_from_s3
@s3_arn_of_file='arn:aws:s3:::<bucket_name>/<file_name>',
@rds_file_path='D:\S3\<custom_folder_name>\<file_name>',
@overwrite_file=1;
```

NEW QUESTION 195

A global company is developing an application across multiple AWS Regions. The company needs a database solution with low latency in each Region and automatic disaster recovery. The database must be deployed in an active-active configuration with automatic data synchronization between Regions. Which solution will meet these requirements with the LOWEST latency?

- A. Amazon RDS with cross-Region read replicas
- B. Amazon DynamoDB global tables
- C. Amazon Aurora global database
- D. Amazon Athena and Amazon S3 with S3 Cross Region Replication

Answer: B

NEW QUESTION 200

A company is launching a new Amazon RDS for MySQL Multi-AZ DB instance to be used as a data store for a custom-built application. After a series of tests with point-in-time recovery disabled, the company decides that it must have point-in-time recovery reenabled before using the DB instance to store production data. What should a database specialist do so that point-in-time recovery can be successful?

- A. Enable binary logging in the DB parameter group used by the DB instance.
- B. Modify the DB instance and enable audit logs to be pushed to Amazon CloudWatch Logs.
- C. Modify the DB instance and configure a backup retention period
- D. Set up a scheduled job to create manual DB instance snapshots.

Answer: C

Explanation:

You can restore a DB instance to a specific point in time (PITR), creating a new DB instance. To support PITR, your DB instances must have backup retention set to a nonzero value. <https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/custom-backup-sqlserver.html>
<https://aws.amazon.com/blogs/database/setting-up-a-binlog-server-for-amazon-rds-mysql-and-mariadb-using-m> "After you run the command, it's okay to enable backup retention on the RDS instance by using the AWS CLI or the console. Enabling backup retention also enables binary logging."
<https://aws.amazon.com/blogs/storage/point-in-time-recovery-and-continuous-backup-for-amazon-rds-with-aws>

NEW QUESTION 201

A stock market analysis firm maintains two locations: one in the us-east-1 Region and another in the eu-west-2 Region. The business want to build an AWS database solution capable of providing rapid and accurate updates.

Dashboards with advanced analytical queries are used to present data in the eu-west-2 office. Because the corporation will use these dashboards to make purchasing choices, they must have less than a second to obtain application data.

Which solution satisfies these criteria and gives the MOST CURRENT dashboard?

- A. Deploy an Amazon RDS DB instance in us-east-1 with a read replica instance in eu-west-2. Create an Amazon ElastiCache cluster in eu-west-2 to cache data from the read replica to generate the dashboards.
- B. Use an Amazon DynamoDB global table in us-east-1 with replication into eu-west-2. Use multi-active replication to ensure that updates are quickly propagated to eu-west-2.
- C. Use an Amazon Aurora global databas
- D. Deploy the primary DB cluster in us-east-1. Deploy the secondary DB cluster in eu-west-2. Configure the dashboard application to read from the secondary cluster.
- E. Deploy an Amazon RDS for MySQL DB instance in us-east-1 with a read replica instance in eu-west-2. Configure the dashboard application to read from the read replica.

Answer: C

Explanation:

Amazon Aurora global databases span multiple AWS Regions, enabling low latency global reads and providing fast recovery from the rare outage that might affect an entire AWS Region. An Aurora global database has a primary DB cluster in one Region, and up to five secondary DB clusters in different Regions. <https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/aurora-global-database.html>

NEW QUESTION 202

A company is hosting critical business data in an Amazon Redshift cluster. Due to the sensitive nature of the data, the cluster is encrypted at rest using AWS KMS.

As a part of disaster recovery requirements, the company needs to copy the Amazon Redshift snapshots to another Region.

Which steps should be taken in the AWS Management Console to meet the disaster recovery requirements?

- A. Create a new KMS customer master key in the source Regio
- B. Switch to the destination Region, enable Amazon Redshift cross-Region snapshots, and use the KMS key of the source Region.
- C. Create a new IAM role with access to the KMS ke
- D. Enable Amazon Redshift cross-Region replication using the new IAM role, and use the KMS key of the source Region.
- E. Enable Amazon Redshift cross-Region snapshots in the source Region, and create a snapshot copy grant and use a KMS key in the destination Region.
- F. Create a new KMS customer master key in the destination Region and create a new IAM role with access to the new KMS ke
- G. Enable Amazon Redshift cross-Region replication in the source Region and use the KMS key of the destination Region.

Answer: C

Explanation:

If you want to enable cross-Region snapshot copy for an AWS KMS–encrypted cluster, you must configure a snapshot copy grant for a root key in the destination AWS Region Source-Region : configure a cross-Region snapshot for an AWS KMS–encrypted cluster In Destination AWS Region : choose the AWS Region to which to copy snapshots.

<https://docs.aws.amazon.com/redshift/latest/mgmt/managing-snapshots-console.html#xregioncopy-kms-encrypt>

NEW QUESTION 206

A Database Specialist has migrated an on-premises Oracle database to Amazon Aurora PostgreSQL. The schema and the data have been migrated successfully.

The on-premises database server was also being used to run database maintenance cron jobs written in Python to perform tasks including data purging and generating data exports. The logs for these jobs show that, most of the time, the jobs completed within 5 minutes, but a few jobs took up to 10 minutes to complete. These maintenance jobs need to be set up for Aurora PostgreSQL.

How can the Database Specialist schedule these jobs so the setup requires minimal maintenance and provides high availability?

- A. Create cron jobs on an Amazon EC2 instance to run the maintenance jobs following the required schedule.
- B. Connect to the Aurora host and create cron jobs to run the maintenance jobs following the required schedule.
- C. Create AWS Lambda functions to run the maintenance jobs and schedule them with Amazon CloudWatch Events.
- D. Create the maintenance job using the Amazon CloudWatch job scheduling plugin.

Answer: C

Explanation:

<https://docs.aws.amazon.com/AmazonCloudWatch/latest/events/Create-CloudWatch-Events-Scheduled-Rule.html> <https://docs.aws.amazon.com/prescriptive-guidance/latest/patterns/schedule-jobs-for-amazon-rds-and-aurora-postgresql.html> a job for data extraction or a job for data purging can easily be scheduled using cron. For these jobs, database credentials are typically either hard-coded or stored in a properties file. However, when you migrate to Amazon Relational Database Service (Amazon RDS) or Amazon Aurora PostgreSQL, you lose the ability to log in to the host instance to schedule cron jobs. This pattern describes how to use AWS Lambda and AWS Secrets Manager to schedule jobs for Amazon RDS and Aurora PostgreSQL databases after migration. <https://docs.aws.amazon.com/AmazonCloudWatch/latest/events/RunLambdaSchedule.html>

NEW QUESTION 210

A company just migrated to Amazon Aurora PostgreSQL from an on-premises Oracle database. After the migration, the company discovered there is a period of time every day around 3:00 PM where the response time of the application is noticeably slower. The company has narrowed down the cause of this issue to the database and not the application.

Which set of steps should the Database Specialist take to most efficiently find the problematic PostgreSQL query?

- A. Create an Amazon CloudWatch dashboard to show the number of connections, CPU usage, and disk space consumption.
- B. Watch these dashboards during the next slow period.
- C. Launch an Amazon EC2 instance, and install and configure an open-source PostgreSQL monitoring tool that will run reports based on the output error logs.
- D. Modify the logging database parameter to log all the queries related to locking in the database and then check the logs after the next slow period for this information.
- E. Enable Amazon RDS Performance Insights on the PostgreSQL databases.
- F. Use the metrics to identify any queries that are related to spikes in the graph during the next slow period.

Answer: D

NEW QUESTION 215

A company is using 5 TB Amazon RDS DB instances and needs to maintain 5 years of monthly database backups for compliance purposes. A Database Administrator must provide Auditors with data within 24 hours. Which solution will meet these requirements and is the MOST operationally efficient?

- A. Create an AWS Lambda function to run on the first day of every month to take a manual RDS snapshot. Move the snapshot to the company's Amazon S3 bucket.
- B. Create an AWS Lambda function to run on the first day of every month to take a manual RDS snapshot.
- C. Create an RDS snapshot schedule from the AWS Management Console to take a snapshot every 30 days.
- D. Create an AWS Lambda function to run on the first day of every month to create an automated RDS snapshot.

Answer: A

Explanation:

Unlike automated backups, manual snapshots aren't subject to the backup retention period. Snapshots don't expire. For very long-term backups of MariaDB, MySQL, and PostgreSQL data, we recommend exporting snapshot data to Amazon S3. If the major version of your DB engine is no longer supported, you can't restore to that version from a snapshot. https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_CreateSnapshot.html

NEW QUESTION 220

A company runs online transaction processing (OLTP) workloads on an Amazon RDS for PostgreSQL Multi-AZ DB instance. Tests were run on the database after work hours, which generated additional database logs. The free storage of the RDS DB instance is low due to these additional logs.

What should the company do to address this space constraint issue?

- A. Log in to the host and run the `rm $PGDATA/pg_logs/*` command.
- B. Modify the `rds.log_retention_period` parameter to 1440 and wait up to 24 hours for database logs to be deleted.
- C. Create a ticket with AWS Support to have the logs deleted.
- D. Run the `SELECT rds_rotate_error_log()` stored procedure to rotate the logs.

Answer: B

Explanation:

To set the retention period for system logs, use the `rds.log_retention_period` parameter. You can find `rds.log_retention_period` in the DB parameter group associated with your DB instance. The unit for this parameter is minutes. For example, a setting of 1,440 retains logs for one day. The default value is 4,320 (three days). The maximum value is 10,080 (seven days).

https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/USER_LogAccess.Concepts.PostgreSQL.html

NEW QUESTION 225

In one AWS account, a business runs a two-tier ecommerce application. An Amazon RDS for MySQL Multi-AZ database instance serves as the application's backend. A developer removed the database instance in the production environment by accident. Although the organization recovers the database, the incident results in hours of outage and financial loss.

Which combination of adjustments would reduce the likelihood that this error will occur again in the future? (Select three.)

- A. Grant least privilege to groups, IAM users, and roles.
- B. Allow all users to restore a database from a backup.

- C. Enable deletion protection on existing production DB instances.
- D. Use an ACL policy to restrict users from DB instance deletion.
- E. Enable AWS CloudTrail logging and Enhanced Monitoring.

Answer: ACD

NEW QUESTION 227

A company has an ecommerce web application with an Amazon RDS for MySQL DB instance. The marketing team has noticed some unexpected updates to the product and pricing information on the website, which is impacting sales targets. The marketing team wants a database specialist to audit future database activity to help identify how and when the changes are being made.

What should the database specialist do to meet these requirements? (Choose two.)

- A. Create an RDS event subscription to the audit event type.
- B. Enable auditing of CONNECT and QUERY_DML events.
- C. SSH to the DB instance and review the database logs.
- D. Publish the database logs to Amazon CloudWatch Logs.
- E. Enable Enhanced Monitoring on the DB instance.

Answer: BD

Explanation:

<https://aws.amazon.com/blogs/database/configuring-an-audit-log-to-capture-database-activities-for-amazon-rds>

NEW QUESTION 232

A business maintains a SQL Server database on-premises. Active Directory authentication is used to provide users access to the database. The organization transferred their database successfully to Amazon RDS for SQL Server. The organization, however, has reservations regarding user authentication in the AWS Cloud environment.

Which authentication solution should a database professional provide?

- A. Deploy Active Directory Federation Services (AD FS) on premises and configure it with an on-premises Active Director
- B. Set up delegation between the on- premises AD FS and AWS Security Token Service (AWS STS) to map user identities to a role using theAmazonRDSDirectoryServiceAccess managed IAM policy.
- C. Establish a forest trust between the on-premises Active Directory and AWS Directory Service for Microsoft Active Director
- D. Use AWS SSO to configure an Active Directory user delegated to access the databases in RDS for SQL Server.
- E. Use Active Directory Connector to redirect directory requests to the company's on-premises Active Directory without caching any information in the cloud
- F. Use the RDS master user credentials to connect to the DB instance and configure SQL Server logins and users from the Active Directory users and groups.
- G. Establish a forest trust between the on-premises Active Directory and AWS Directory Service for Microsoft Active Director
- H. Ensure RDS for SQL Server is using mixed mode authentication
- I. Use the RDS master user credentials to connect to the DB instance and configure SQL Server logins and users from the Active Directory users and groups.

Answer: D

Explanation:

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_SQLServerWinAuth.html

NEW QUESTION 234

A Database Specialist is designing a new database infrastructure for a ride hailing application. The application data includes a ride tracking system that stores GPS coordinates for all rides. Real-time statistics and metadata lookups must be performed with high throughput and microsecond latency. The database should be fault tolerant with minimal operational overhead and development effort.

Which solution meets these requirements in the MOST efficient way?

- A. Use Amazon RDS for MySQL as the database and use Amazon ElastiCache
- B. Use Amazon DynamoDB as the database and use DynamoDB Accelerator
- C. Use Amazon Aurora MySQL as the database and use Aurora's buffer cache
- D. Use Amazon DynamoDB as the database and use Amazon API Gateway

Answer: B

Explanation:

[https://aws.amazon.com/dynamodb/dax/#:~:text=Amazon%20DynamoDB%20Accelerator%20\(DAX\)%20is,mil](https://aws.amazon.com/dynamodb/dax/#:~:text=Amazon%20DynamoDB%20Accelerator%20(DAX)%20is,mil) "Amazon DynamoDB Accelerator (DAX) is a fully managed, highly available, in-memory cache for DynamoDB that delivers up to a 10x performance improvement – from milliseconds to microseconds – even at millions of requests per second. "

NEW QUESTION 237

The Security team for a finance company was notified of an internal security breach that happened 3 weeks ago. A Database Specialist must start producing audit logs out of the production Amazon Aurora PostgreSQL cluster for the Security team to use for monitoring and alerting. The Security team is required to perform real- time alerting and monitoring outside the Aurora DB cluster and wants to have the cluster push encrypted files to the chosen solution.

Which approach will meet these requirements?

- A. Use pg_audit to generate audit logs and send the logs to the Security team.
- B. Use AWS CloudTrail to audit the DB cluster and the Security team will get data from Amazon S3.
- C. Set up database activity streams and connect the data stream from Amazon Kinesis to consumer applications.
- D. Turn on verbose logging and set up a schedule for the logs to be dumped out for the Security team.

Answer: C

Explanation:

<https://aws.amazon.com/about-aws/whats-new/2019/05/amazon-aurora-with-postgresql-compatibility-supports-> "Database Activity Streams for Amazon Aurora with PostgreSQL compatibility provides a near real-time data stream of the database activity in your relational database to help you monitor activity. When

integrated with third party database activity monitoring tools, Database Activity Streams can monitor and audit database activity to provide safeguards for your database and help meet compliance and regulatory requirements."

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Overview.LoggingAndMonitoring.html>

NEW QUESTION 238

A marketing company is using Amazon DocumentDB and requires that database audit logs be enabled. A Database Specialist needs to configure monitoring so that all data definition language (DDL) statements performed are visible to the Administrator. The Database Specialist has set the audit_logs parameter to enabled in the cluster parameter group.

What should the Database Specialist do to automatically collect the database logs for the Administrator?

- A. Enable DocumentDB to export the logs to Amazon CloudWatch Logs
- B. Enable DocumentDB to export the logs to AWS CloudTrail
- C. Enable DocumentDB Events to export the logs to Amazon CloudWatch Logs
- D. Configure an AWS Lambda function to download the logs using the download-db-log-file-portion operation and store the logs in Amazon S3

Answer: C

Explanation:

<https://docs.aws.amazon.com/documentdb/latest/developerguide/event-auditing.html> Auditing Amazon DocumentDB Events

PDF

Kindle RSS

With Amazon DocumentDB (with MongoDB compatibility), you can audit events that were performed in your cluster. Examples of logged events include successful and failed authentication attempts, dropping a collection in a database, or creating an index. By default, auditing is disabled on Amazon DocumentDB and requires that you opt in to use this feature.

When auditing is enabled, Amazon DocumentDB records Data Definition Language (DDL), authentication, authorization, and user management events to Amazon CloudWatch Logs. When auditing is enabled, Amazon DocumentDB exports your cluster's auditing records (JSON documents) to Amazon CloudWatch Logs. You can use Amazon CloudWatch Logs to analyze, monitor, and archive your Amazon DocumentDB auditing events.

NEW QUESTION 241

A company is running its line of business application on AWS, which uses Amazon RDS for MySQL at the persistent data store. The company wants to minimize downtime when it migrates the database to Amazon Aurora.

Which migration method should a Database Specialist use?

- A. Take a snapshot of the RDS for MySQL DB instance and create a new Aurora DB cluster with the option to migrate snapshots.
- B. Make a backup of the RDS for MySQL DB instance using the mysqldump utility, create a new Aurora DB cluster, and restore the backup.
- C. Create an Aurora Replica from the RDS for MySQL DB instance and promote the Aurora DB cluster.
- D. Create a clone of the RDS for MySQL DB instance and promote the Aurora DB cluster.

Answer: C

Explanation:

<https://aws.amazon.com/blogs/database/best-practices-for-migrating-rds-for-mysql-databases-to-amazon-aurora/>

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraPostgreSQL.Migrating.html#Aurora>

NEW QUESTION 245

A business is launching a new Amazon RDS for SQL Server database instance. The organization wishes to allow auditing of the SQL Server database.

Which measures should a database professional perform in combination to achieve this requirement? (Select two.)

- A. Create a service-linked role for Amazon RDS that grants permissions for Amazon RDS to store audit logs on Amazon S3.
- B. Set up a parameter group to configure an IAM role and an Amazon S3 bucket for audit log storage. Associate the parameter group with the DB instance.
- C. Disable Multi-AZ on the DB instance, and then enable auditin
- D. Enable Multi-AZ after auditing is enabled.
- E. Disable automated backup on the DB instance, and then enable auditin
- F. Enable automated backup after auditing is enabled.
- G. Set up an options group to configure an IAM role and an Amazon S3 bucket for audit log storage. Associate the options group with the DB instance.

Answer: AE

Explanation:

<https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Appendix.SQLServer.Options.Audit.html>

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/security_iam_service-with-iam.html

NEW QUESTION 246

A company runs hundreds of Microsoft SQL Server databases on Windows servers in its on-premises data center. A database specialist needs to migrate these databases to Linux on AWS.

Which combination of steps should the database specialist take to meet this requirement? (Choose three.)

- A. Install AWS Systems Manager Agent on the on-premises server
- B. Use Systems Manager Run Command to install the Windows to Linux replatforming assistant for Microsoft SQL Server Databases.
- C. Use AWS Systems Manager Run Command to install and configure the AWS Schema Conversion Tool on the on-premises servers.
- D. On the Amazon EC2 console, launch EC2 instances and select a Linux AMI that includes SQL Server. Install and configure AWS Systems Manager Agent on the EC2 instances.
- E. On the AWS Management Console, set up Amazon RDS for SQL Server DB instances with Linux as the operating system
- F. Install AWS Systems Manager Agent on the DB instances by using an options group.
- G. Open the Windows to Linux replatforming assistant tool
- H. Enter configuration details of the source and destination database
- I. Start migration.
- J. On the AWS Management Console, set up AWS Database Migration Service (AWS DMS) by entering details of the source SQL Server database and the destination SQL Server database on AWS

K. Start migration.

Answer: ACE

Explanation:

<https://docs.aws.amazon.com/AWSEC2/latest/WindowsGuide/replatform-sql-server.html>

https://d1.awsstatic.com/events/reinvent/2019/REPEAT_1_Leverage_automation_to_re-platform_SQL_Server_

NEW QUESTION 248

An online retail company is planning a multi-day flash sale that must support processing of up to 5,000 orders per second. The number of orders and exact schedule for the sale will vary each day. During the sale, approximately 10,000 concurrent users will look at the deals before buying items. Outside of the sale, the traffic volume is very low. The acceptable performance for read/write queries should be under 25 ms. Order items are about 2 KB in size and have a unique identifier. The company requires the most cost-effective solution that will automatically scale and is highly available.

Which solution meets these requirements?

- A. Amazon DynamoDB with on-demand capacity mode
- B. Amazon Aurora with one writer node and an Aurora Replica with the parallel query feature enabled
- C. Amazon DynamoDB with provisioned capacity mode with 5,000 write capacity units (WCUs) and 10,000 read capacity units (RCUs)
- D. Amazon Aurora with one writer node and two cross-Region Aurora Replicas

Answer: A

Explanation:

The number of orders and exact schedule for the sale will vary each day. During the sale, approximately 10,000 concurrent users will look at the deals before buying items. Outside of the sale, the traffic volume is very low ==> Setting provisioning DynamoDB fix read 5000/write 10000 with will waste the resource when the traffic is low. It is not cost-effective.

NEW QUESTION 249

A retail company is about to migrate its online and mobile store to AWS. The company's CEO has strategic plans to grow the brand globally. A Database Specialist has been challenged to provide predictable read and write database performance with minimal operational overhead.

What should the Database Specialist do to meet these requirements?

- A. Use Amazon DynamoDB global tables to synchronize transactions
- B. Use Amazon EMR to copy the orders table data across Regions
- C. Use Amazon Aurora Global Database to synchronize all transactions
- D. Use Amazon DynamoDB Streams to replicate all DynamoDB transactions and sync them

Answer: A

Explanation:

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

With global tables, your globally distributed applications can access data locally in the selected regions to get single-digit millisecond read and write performance. Not Aurora Global Database, as per this link: https://aws.amazon.com/rds/aurora/global-database/?nc1=h_ls . Aurora Global Database lets you easily scale database reads across the world and place your applications close to your users.

NEW QUESTION 251

Amazon Neptune is being used by a corporation as the graph database for one of its products. During an ETL procedure, the company's data science team produced enormous volumes of temporary data by unintentionally. The Neptune DB cluster extended its storage capacity automatically to handle the added data, but the data science team erased the superfluous data.

What should a database professional do to prevent incurring extra expenditures for cluster volume space that is not being used?

- A. Take a snapshot of the cluster volum
- B. Restore the snapshot in another cluster with a smaller volume size.
- C. Use the AWS CLI to turn on automatic resizing of the cluster volume.
- D. Export the cluster data into a new Neptune DB cluster.
- E. Add a Neptune read replica to the cluste
- F. Promote this replica as a new primary DB instanc
- G. Reset the storage space of the cluster.

Answer: C

Explanation:

The only way to shrink the storage space used by your DB cluster when you have a large amount of unused allocated space is to export all the data in your graph and then reload it into a new DB cluster. Creating and restoring a snapshot does not reduce the amount of storage allocated for your DB cluster, because a snapshot retains the original image of the cluster's underlying storage.

NEW QUESTION 252

A company hosts an on-premises Microsoft SQL Server Enterprise edition database with Transparent Data Encryption (TDE) enabled. The database is 20 TB in size and includes sparse tables. The company needs to migrate the database to Amazon RDS for SQL Server during a maintenance window that is scheduled for an upcoming weekend. Data-at-rest encryption must be enabled for the target DB instance.

Which combination of steps should the company take to migrate the database to AWS in the MOST operationally efficient manner? (Choose two.)

- A. Use AWS Database Migration Service (AWS DMS) to migrate from the on-premises source database to the RDS for SQL Server target database.
- B. Disable TD
- C. Create a database backup without encryptio
- D. Copy the backup to Amazon S3.
- E. Restore the backup to the RDS for SQL Server DB instanc
- F. Enable TDE for the RDS for SQL Server DB instance.
- G. Set up an AWS Snowball Edge devic

- H. Copy the database backup to the device
- I. Send the device to AWS
- J. Restore the database from Amazon S3.
- K. Encrypt the data with client-side encryption before transferring the data to Amazon RDS.

Answer: BC

Explanation:

<https://aws.amazon.com/blogs/database/migrate-tde-enabled-sql-server-databases-to-amazon-rds-for-sql-server/>

NEW QUESTION 257

A company has an Amazon RDS Multi-AZ DB instances that is 200 GB in size with an RPO of 6 hours. To meet the company's disaster recovery policies, the database backup needs to be copied into another Region. The company requires the solution to be cost-effective and operationally efficient. What should a Database Specialist do to copy the database backup into a different Region?

- A. Use Amazon RDS automated snapshots and use AWS Lambda to copy the snapshot into another Region
- B. Use Amazon RDS automated snapshots every 6 hours and use Amazon S3 cross-Region replication to copy the snapshot into another Region
- C. Create an AWS Lambda function to take an Amazon RDS snapshot every 6 hours and use a second Lambda function to copy the snapshot into another Region
- D. Create a cross-Region read replica for Amazon RDS in another Region and take an automated snapshot of the read replica

Answer: C

Explanation:

System snapshot can't fulfill 6 hours requirement. You need to control it by script

<https://aws.amazon.com/blogs/database/%C2%AD%C2%AD%C2%ADautomating-cross-region-cross-account>

NEW QUESTION 258

A small startup company is looking to migrate a 4 TB on-premises MySQL database to AWS using an Amazon RDS for MySQL DB instance. Which strategy would allow for a successful migration with the LEAST amount of downtime?

- A. Deploy a new RDS for MySQL DB instance and configure it for access from the on-premises data center
- B. Use the mysqldump utility to create an initial snapshot from the on-premises MySQL server, and copy it to an Amazon S3 bucket
- C. Import the snapshot into the DB instance utilizing the MySQL utilities running on an Amazon EC2 instance
- D. Immediately point the application to the DB instance.
- E. Deploy a new Amazon EC2 instance, install the MySQL software on the EC2 instance, and configure networking for access from the on-premises data center
- F. Use the mysqldump utility to create a snapshot of the on-premises MySQL server
- G. Copy the snapshot into the EC2 instance and restore it into the EC2 MySQL instance
- H. Use AWS DMS to migrate data into a new RDS for MySQL DB instance
- I. Point the application to the DB instance.
- J. Deploy a new Amazon EC2 instance, install the MySQL software on the EC2 instance, and configure networking for access from the on-premises data center
- K. Use the mysqldump utility to create a snapshot of the on-premises MySQL server
- L. Copy the snapshot into an Amazon S3 bucket and import the snapshot into a new RDS for MySQL DB instance using the MySQL utilities running on an EC2 instance
- M. Point the application to the DB instance.
- N. Deploy a new RDS for MySQL DB instance and configure it for access from the on-premises data center
- O. Use the mysqldump utility to create an initial snapshot from the on-premises MySQL server, and copy it to an Amazon S3 bucket
- P. Import the snapshot into the DB instance using the MySQL utilities running on an Amazon EC2 instance
- Q. Establish replication into the new DB instance using MySQL replication
- R. Stop application access to the on-premises MySQL server and let the remaining transactions replicate over
- S. Point the application to the DB instance.

Answer: B

NEW QUESTION 263

.....

Relate Links

100% Pass Your DBS-C01 Exam with Exam Bible Prep Materials

<https://www.exambible.com/DBS-C01-exam/>

Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - <https://www.exambible.com/>