

Amazon-Web-Services

Exam Questions DBS-C01

AWS Certified Database - Specialty



NEW QUESTION 1

A database specialist deployed an Amazon RDS DB instance in Dev-VPC1 used by their development team. Dev-VPC1 has a peering connection with Dev-VPC2 that belongs to a different development team in the same department. The networking team confirmed that the routing between VPCs is correct; however, the database engineers in Dev-VPC2 are getting a timeout connections error when trying to connect to the database in Dev- VPC1. What is likely causing the timeouts?

- A. The database is deployed in a VPC that is in a different Region.
- B. The database is deployed in a VPC that is in a different Availability Zone.
- C. The database is deployed with misconfigured security groups.
- D. The database is deployed with the wrong client connect timeout configuration.

Answer: C

Explanation:

"A VPC peering connection is a networking connection between two VPCs that enables you to route traffic between them using private IP addresses. Instances in either VPC can communicate with each other as if they are within the same network. You can create a VPC peering connection between your own VPCs, with a VPC in another AWS account, or with a VPC in a different AWS Region." https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_VPC.Scenarios.html

NEW QUESTION 2

An ecommerce company uses Amazon DynamoDB as the backend for its payments system. A new regulation requires the company to log all data access requests for financial audits. For this purpose, the company plans to use AWS logging and save logs to Amazon S3. How can a database specialist activate logging on the database?

- A. Use AWS CloudTrail to monitor DynamoDB control-plane operation
- B. Create a DynamoDB stream to monitor data-plane operation
- C. Pass the stream to Amazon Kinesis Data Stream
- D. Use that stream as a source for Amazon Kinesis Data Firehose to store the data in an Amazon S3 bucket.
- E. Use AWS CloudTrail to monitor DynamoDB data-plane operation
- F. Create a DynamoDB stream to monitor control-plane operation
- G. Pass the stream to Amazon Kinesis Data Stream
- H. Use that stream as a source for Amazon Kinesis Data Firehose to store the data in an Amazon S3 bucket.
- I. Create two trails in AWS CloudTrai
- J. Use Trail1 to monitor DynamoDB control-plane operation
- K. Use Trail2 to monitor DynamoDB data-plane operations.
- L. Use AWS CloudTrail to monitor DynamoDB data-plane and control-plane operations.

Answer: D

Explanation:

<https://aws.amazon.com/about-aws/whats-new/2021/04/you-now-can-use-aws-cloudtrail-to-log-amazon-dynamo>

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

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 5

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 6

A company has a web-based survey application that uses Amazon DynamoDB. During peak usage, when survey responses are being collected, a Database Specialist sees the ProvisionedThroughputExceededException error.

What can the Database Specialist do to resolve this error? (Choose two.)

- A. Change the table to use Amazon DynamoDB Streams
- B. Purchase DynamoDB reserved capacity in the affected Region
- C. Increase the write capacity units for the specific table
- D. Change the table capacity mode to on-demand
- E. Change the table type to throughput optimized

Answer: CD

Explanation:

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

NEW QUESTION 7

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 version. However, you can do clone with most engine version. Meanwhile, I also consider the performance while restoring snapshot 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 8

A company with branch offices in Portland, New York, and Singapore has a three-tier web application that leverages a shared database. The database runs on Amazon RDS for MySQL and is hosted in the us-west-2 Region. The application has a distributed front end deployed in the us-west-2, ap-southeast-1, and us-east-2 Regions.

This front end is used as a dashboard for Sales Managers in each branch office to see current sales statistics. There are complaints that the dashboard performs more slowly in the Singapore location than it does in Portland or New York. A solution is needed to provide consistent performance for all users in each location.

Which set of actions will meet these requirements?

- A. Take a snapshot of the instance in the us-west-2 Region
- B. Create a new instance from the snapshot in the ap-southeast-1 Region
- C. Reconfigure the ap-southeast-1 front-end dashboard to access this instance.
- D. Create an RDS read replica in the ap-southeast-1 Region from the primary RDS DB instance in the us-west-2 Region
- E. Reconfigure the ap-southeast-1 front-end dashboard to access this instance.
- F. Create a new RDS instance in the ap-southeast-1 Region
- G. Use AWS DMS and change data capture (CDC) to update the new instance in the ap-southeast-1 Region
- H. Reconfigure the ap-southeast-1 front-end dashboard to access this instance.
- I. Create an RDS read replica in the us-west-2 Region where the primary instance reside
- J. Create a read replica in the ap-southeast-1 Region from the read replica located on the us-west-2 Region
- K. Reconfigure the ap-southeast-1 front-end dashboard to access this instance.

Answer: B

Explanation:

<https://aws.amazon.com/rds/features/read-replicas/>

"Amazon RDS Read Replicas provide enhanced performance and durability for RDS database (DB) instances.

They make it easy to elastically scale out beyond the capacity constraints of a single DB instance for read-heavy database workloads. You can create one or more replicas of a given source DB Instance and serve high-volume application read traffic from multiple copies of your data, thereby increasing aggregate read throughput. "

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

NEW QUESTION 9

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 10

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 10

A large company has a variety of Amazon DB clusters. Each of these clusters has various configurations that adhere to various requirements. Depending on the team and use case, these configurations can be organized into broader categories.

A database administrator wants to make the process of storing and modifying these parameters more systematic. The database administrator also wants to ensure that changes to individual categories of configurations are automatically applied to all instances when required.

Which AWS service or feature will help automate and achieve this objective?

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

Answer: B

NEW QUESTION 15

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 key
- B. Use the event time as the sort key
- 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 table
- 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 dash
- I. Add a local secondary index with the player identifier as the sort key.
- J. Create one table for each game
- K. Use the player identifier as the partition key
- L. Use the event time as the sort key.

Answer: D

NEW QUESTION 18

A user has a non-relational key-value database. The user is looking for a fully managed AWS service that will offload the administrative burdens of operating and scaling distributed databases. The solution must be cost-effective and able to handle unpredictable application traffic.

What should a Database Specialist recommend for this user?

- A. Create an Amazon DynamoDB table with provisioned capacity mode
- B. Create an Amazon DocumentDB cluster
- C. Create an Amazon DynamoDB table with on-demand capacity mode
- D. Create an Amazon Aurora Serverless DB cluster

Answer: C

NEW QUESTION 20

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 22

An online gaming company is using an Amazon DynamoDB table in on-demand mode to store game scores. After an intensive advertisement campaign in South America, the average number of concurrent users rapidly increases from 100,000 to 500,000 in less than 10 minutes every day around 5 PM.

The on-call software reliability engineer has observed that the application logs contain a high number of DynamoDB throttling exceptions caused by game score insertions around 5 PM. Customer service has also reported that several users are complaining about their scores not being registered.

How should the database administrator remediate this issue at the lowest cost?

- A. Enable auto scaling and set the target usage rate to 90%.
- B. Switch the table to provisioned mode and enable auto scaling.
- C. Switch the table to provisioned mode and set the throughput to the peak value.
- D. Create a DynamoDB Accelerator cluster and use it to access the DynamoDB table.

Answer: B

NEW QUESTION 27

A company developed an AWS CloudFormation template used to create all new Amazon DynamoDB tables in its AWS account. The template configures provisioned throughput capacity using hard-coded values. The company wants to change the template so that the tables it creates in the future have independently configurable read and write capacity units assigned.

Which solution will enable this change?

- A. Add values for the rcuCount and wcuCount parameters to the Mappings section of the template. Configure DynamoDB to provision throughput capacity using the stack's mappings.
- B. Add values for two Number parameters, rcuCount and wcuCount, to the template.
- C. Replace the hard-coded values with calls to the Ref intrinsic function, referencing the new parameters.
- D. Add values for the rcuCount and wcuCount parameters as outputs of the template.
- E. Configure DynamoDB to provision throughput capacity using the stack outputs.
- F. Add values for the rcuCount and wcuCount parameters to the Mappings section of the template.
- G. Replace the hard-coded values with calls to the Ref intrinsic function, referencing the new parameters.

Answer: B

Explanation:

Input parameter and FindInMap You can use an input parameter with the Fn::FindInMap function to refer to a specific value in a map. For example, suppose you have a list of regions and environment types that map to a specific AMI ID. You can select the AMI ID that your stack uses by using an input parameter (EnvironmentType). To determine the region, use the AWS::Region pseudo parameter, which gets the AWS Region in which you create the stack.
<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/parameters-section-structure.html>

NEW QUESTION 31

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 35

A financial organization must ensure that the most current 90 days of MySQL database backups are accessible. Amazon RDS for MySQL DB instances are used to host all MySQL databases. A database expert must create a solution that satisfies the criteria for backup retention with the least amount of development work feasible.

Which strategy should the database administrator take?

- A. Use AWS Backup to build a backup plan for the required retention period.
- B. Assign the DB instances to the backup plan.
- C. Modify the DB instances to enable the automated backup option.
- D. Select the required backup retention period.
- E. Automate a daily cron job on an Amazon EC2 instance to create MySQL dumps, transfer to Amazon S3, and implement an S3 Lifecycle policy to meet the retention requirement.
- F. Use AWS Lambda to schedule a daily manual snapshot of the DB instance.
- G. Delete snapshots that exceed the retention requirement.

Answer: A

Explanation:

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

NEW QUESTION 38

A database professional is developing an application that will respond to single-instance requests. The program will query large amounts of client data and offer end users with results.

These reports may include a variety of fields. The database specialist wants to enable users to query the database using any of the fields offered.

During peak periods, the database's traffic volume will be significant yet changeable. However, the database will see little activity over the rest of the day.

Which approach will be the most cost-effective in meeting these requirements?

- A. Amazon DynamoDB with provisioned capacity mode and auto scaling
- B. Amazon DynamoDB with on-demand capacity mode
- C. Amazon Aurora with auto scaling enabled
- D. Amazon Aurora in a serverless mode

Answer: D

Explanation:

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Limits.html#limits-items>

NEW QUESTION 42

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 47

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 48

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 50

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 52

A company is load testing its three-tier production web application deployed with an AWS CloudFormation template on AWS. The Application team is making changes to deploy additional Amazon EC2 and AWS Lambda resources to expand the load testing capacity. A Database Specialist wants to ensure that the changes made by the Application team will not change the Amazon RDS database resources already deployed. Which combination of steps would allow the Database Specialist to accomplish this? (Choose two.)

- A. Review the stack drift before modifying the template
- B. Create and review a change set before applying it
- C. Export the database resources as stack outputs
- D. Define the database resources in a nested stack
- E. Set a stack policy for the database resources

Answer: BE

Explanation:

https://docs.amazonaws.cn/en_us/AWSCloudFormation/latest/UserGuide/best-practices.html#cfn-best-practices

NEW QUESTION 53

A company wants to migrate its existing on-premises Oracle database to Amazon Aurora PostgreSQL. The migration must be completed with minimal downtime using AWS DMS. A Database Specialist must validate that the data was migrated accurately from the source to the target before the cutover. The migration must have minimal impact on the performance of the source database. Which approach will MOST effectively meet these requirements?

- A. Use the AWS Schema Conversion Tool (AWS SCT) to convert source Oracle database schemas to the target Aurora DB cluster
- B. Verify the datatype of the columns.
- C. Use the table metrics of the AWS DMS task created for migrating the data to verify the statistics for the tables being migrated and to verify that the data definition language (DDL) statements are completed.
- D. Enable the AWS Schema Conversion Tool (AWS SCT) premigration validation and review the premigration checklist to make sure there are no issues with the conversion.
- E. Enable AWS DMS data validation on the task so the AWS DMS task compares the source and target records, and reports any mismatches.

Answer: D

Explanation:

"To ensure that your data was migrated accurately from the source to the target, we highly recommend that you use data validation."

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

NEW QUESTION 54

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 58

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 61

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 63

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 65

An ecommerce company is running AWS Database Migration Service (AWS DMS) to replicate an on-premises Microsoft SQL Server database to Amazon RDS for SQL Server. The company has set up an AWS Direct Connect connection from its on-premises data center to AWS. During the migration, the company's security team receives an alarm that is related to the migration. The security team mandates that the DMS replication instance must not be accessible from public IP addresses.

What should a database specialist do to meet this requirement?

- A. Set up a VPN connection to encrypt the traffic over the Direct Connect connection.
- B. Modify the DMS replication instance by disabling the publicly accessible option.
- C. Delete the DMS replication instance
- D. Recreate the DMS replication instance with the publicly accessible option disabled.
- E. Create a new replication VPC subnet group with private subnet
- F. Modify the DMS replication instance by selecting the newly created VPC subnet group.

Answer: C

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/dms-disable-public-access/>

NEW QUESTION 68

A company is running a finance application on an Amazon RDS for MySQL DB instance. The application is governed by multiple financial regulatory agencies. The RDS DB instance is set up with security groups to allow access to certain Amazon EC2 servers only. AWS KMS is used for encryption at rest.

Which step will provide additional security?

- A. Set up NACLs that allow the entire EC2 subnet to access the DB instance
- B. Disable the master user account
- C. Set up a security group that blocks SSH to the DB instance
- D. Set up RDS to use SSL for data in transit

Answer: D

NEW QUESTION 69

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 71

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 Serve
- B. Use a CloudWatch dashboard to identify the root cause.
- C. Enable RDS Performance Insights and determine which query is creating the proble
- 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 76

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 80

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 82

A database specialist manages a critical Amazon RDS for MySQL DB instance for a company. The data stored daily could vary from .01% to 10% of the current database size. The database specialist needs to ensure that the DB instance storage grows as needed.
What is the MOST operationally efficient and cost-effective solution?

- A. Configure RDS Storage Auto Scaling.
- B. Configure RDS instance Auto Scaling.
- C. Modify the DB instance allocated storage to meet the forecasted requirements.
- D. Monitor the Amazon CloudWatch FreeStorageSpace metric daily and add storage as required.

Answer: A

Explanation:

If your workload is unpredictable, you can enable storage autoscaling for an Amazon RDS DB instance. With storage autoscaling enabled, when Amazon RDS detects that you are running out of free database space it automatically scales up your storage.

<https://aws.amazon.com/about-aws/whats-new/2019/06/rds-storage-auto-scaling/>

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_PIOPS.StorageTypes.html#USER_PIOPS.

NEW QUESTION 87

A company is planning to use Amazon RDS for SQL Server for one of its critical applications. The company's security team requires that the users of the RDS for SQL Server DB instance are authenticated with on-premises Microsoft Active Directory credentials. Which combination of steps should a database specialist take to meet this requirement? (Choose three.)

- A. Extend the on-premises Active Directory to AWS by using AD Connector.
- B. Create an IAM user that uses the AmazonRDSDirectoryServiceAccess managed IAM policy.
- C. Create a directory by using AWS Directory Service for Microsoft Active Directory.
- D. Create an Active Directory domain controller on Amazon EC2.
- E. Create an IAM role that uses the AmazonRDSDirectoryServiceAccess managed IAM policy.
- F. Create a one-way forest trust from the AWS Directory Service for Microsoft Active Directory directory to the on-premises Active Directory.

Answer: CEF

NEW QUESTION 88

A database specialist is working on an Amazon RDS for PostgreSQL DB instance that is experiencing application performance issues due to the addition of new workloads. The database has 5 € of storage space with Provisioned IOPS. Amazon CloudWatch metrics show that the average disk queue depth is greater than 200 and that the disk I/O response time is significantly higher than usual.

What should the database specialist do to improve the performance of the application immediately?

- A. Increase the Provisioned IOPS rate on the storage.
- B. Increase the available storage space.
- C. Use General Purpose SSD (gp2) storage with burst credits.
- D. Create a read replica to offload Read IOPS from the DB instance.

Answer: A

NEW QUESTION 89

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 90

A company has an ecommerce website that runs on AWS. The website uses an Amazon RDS for MySQL database. A database specialist wants to enforce the use of temporary credentials to access the database.

Which solution will meet this requirement?

- A. Use MySQL native database authentication.
- B. Use AWS Secrets Manager to rotate the credentials.
- C. Use AWS Identity and Access Management (IAM) database authentication.
- D. Use AWS Systems Manager Parameter Store for authentication.

Answer: C

NEW QUESTION 95

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 97

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 100

A company is developing a multi-tier web application hosted on AWS using Amazon Aurora as the database. The application needs to be deployed to production and other non-production environments. A Database Specialist needs to specify different MasterUsername and MasterUserPassword properties in the AWS CloudFormation templates used for automated deployment. The CloudFormation templates are version controlled in the company's code repository. The company also needs to meet compliance requirement by routinely rotating its database master password for production.

What is most secure solution to store the master password?

- A. Store the master password in a parameter file in each environmen
- B. Reference the environment-specific parameter file in the CloudFormation template.
- C. Encrypt the master password using an AWS KMS ke
- D. Store the encrypted master password in the CloudFormation template.
- E. Use the secretsmanager dynamic reference to retrieve the master password stored in AWS Secrets Manager and enable automatic rotation.
- F. Use the ssm dynamic reference to retrieve the master password stored in the AWS Systems Manager Parameter Store and enable automatic rotation.

Answer: C

Explanation:

"By using the secure string support in CloudFormation with dynamic references you can better maintain your infrastructure as code. You'll be able to avoid hard coding passwords into your templates and you can keep these runtime configuration parameters separated from your code. Moreover, when properly used, secure strings will help keep your development and production code as similar as possible, while continuing to make your infrastructure code suitable for continuous deployment pipelines."

<https://aws.amazon.com/blogs/mt/using-aws-systems-manager-parameter-store-secure-string-parameters-in-aws> <https://aws.amazon.com/blogs/security/how-to-use-aws-secrets-manager-rotate-credentials-amazon-rds-database>

NEW QUESTION 103

A development team asks a database specialist to create a copy of a production Amazon RDS for MySQL DB instance every morning. The development team will use the copied DB instance as a testing environment for development. The original DB instance and the copy will be hosted in different VPCs of the same AWS account. The development team wants the copy to be available by 6 AM each day and wants to use the same endpoint address each day.

Which combination of steps should the database specialist take to meet these requirements MOST cost-effectively? (Choose three.)

- A. Create a snapshot of the production database each day before the 6 AM deadline.
- B. Create an RDS for MySQL DB instance from the snapsho
- C. Select the desired DB instance size.
- D. Update a defined Amazon Route 53 CNAME record to point to the copied DB instance.
- E. Set up an AWS Database Migration Service (AWS DMS) migration task to copy the snapshot to the copied DB instance.
- F. Use the CopySnapshot action on the production DB instance to create a snapshot before 6 AM.
- G. Update a defined Amazon Route 53 alias record to point to the copied DB instance.

Answer: ABC

NEW QUESTION 108

A Database Specialist needs to speed up any failover that might occur on an Amazon Aurora PostgreSQL DB cluster. The Aurora DB cluster currently includes the primary instance and three Aurora Replicas.

How can the Database Specialist ensure that failovers occur with the least amount of downtime for the application?

- A. Set the TCP keepalive parameters low
- B. Call the AWS CLI failover-db-cluster command
- C. Enable Enhanced Monitoring on the DB cluster
- D. Start a database activity stream on the DB cluster

Answer: A

Explanation:

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

NEW QUESTION 109

A Database Specialist is planning to create a read replica of an existing Amazon RDS for MySQL Multi-AZ DB instance. When using the AWS Management Console to conduct this task, the Database Specialist discovers that the source RDS DB instance does not appear in the read replica source selection box, so the read replica cannot be created.

What is the most likely reason for this?

- A. The source DB instance has to be converted to Single-AZ first to create a read replica from it.
- B. Enhanced Monitoring is not enabled on the source DB instance.

- C. The minor MySQL version in the source DB instance does not support read replicas.
- D. Automated backups are not enabled on the source DB instance.

Answer: D

Explanation:

>Your source DB instance must have backup retention enabled.

https://docs.aws.amazon.com/AmazonRDS/latest/APIReference/API_CreateDBInstanceReadReplica.html

NEW QUESTION 112

A business just transitioned from an on-premises Oracle database to Amazon Aurora PostgreSQL. Following the move, the organization observed that every day around 3:00 PM, the application's response time is substantially slower. The firm has determined that the problem is with the database, not the application. Which set of procedures should the Database Specialist do to locate the erroneous PostgreSQL query most efficiently?

- 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

Explanation:

<https://aws.amazon.com/blogs/database/optimizing-and-tuning-queries-in-amazon-rds-postgresql-based-on-nativ> "AWS recently released a feature called Amazon RDS Performance Insights, which provides an easy-to-understand dashboard for detecting performance problems in terms of load." "AWS recently released a feature called Amazon RDS Performance Insights, which provides an easy-to-understand dashboard for detecting performance problems in terms of load."

NEW QUESTION 114

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 115

A company is running an Amazon RDS for PostgreSQL DB instance and wants to migrate it to an Amazon Aurora PostgreSQL DB cluster. The current database is 1 TB in size. The migration needs to have minimal downtime.

What is the FASTEST way to accomplish this?

- A. Create an Aurora PostgreSQL DB cluster
- B. Set up replication from the source RDS for PostgreSQL DB instance using AWS DMS to the target DB cluster.
- C. Use the pg_dump and pg_restore utilities to extract and restore the RDS for PostgreSQL DB instance to the Aurora PostgreSQL DB cluster.
- D. Create a database snapshot of the RDS for PostgreSQL DB instance and use this snapshot to create the Aurora PostgreSQL DB cluster.
- E. Migrate data from the RDS for PostgreSQL DB instance to an Aurora PostgreSQL DB cluster using an Aurora Replic
- F. Promote the replica during the cutover.

Answer: D

Explanation:

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraPostgreSQL.Migrating.html> Migrating data from an RDS PostgreSQL DB instance to an Aurora PostgreSQL DB cluster by using an

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

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

NEW QUESTION 117

A company plans to migrate a MySQL-based application from an on-premises environment to AWS. The application performs database joins across several tables and uses indexes for faster query response times. The company needs the database to be highly available with automatic failover.

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

- A. Deploy an Amazon RDS DB instance with a read replica.
- B. Deploy an Amazon RDS Multi-AZ DB instance.
- C. Deploy Amazon DynamoDB global tables.
- D. Deploy multiple Amazon RDS DB instances
- E. Use Amazon Route 53 DNS with failover health checks configured.

Answer: B

NEW QUESTION 121

A gaming company has recently acquired a successful iOS game, which is particularly popular during the holiday season. The company has decided to add a leaderboard to the game that uses Amazon DynamoDB. The application load is expected to ramp up over the holiday season. Which solution will meet these requirements at the lowest cost?

- A. DynamoDB Streams
- B. DynamoDB with DynamoDB Accelerator
- C. DynamoDB with on-demand capacity mode
- D. DynamoDB with provisioned capacity mode with Auto Scaling

Answer: C

NEW QUESTION 126

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 131

A company has deployed an e-commerce web application in a new AWS account. An Amazon RDS for MySQL Multi-AZ DB instance is part of this deployment with a database-1.xxxxxxxxxx.us-east-1.rds.amazonaws.com endpoint listening on port 3306. The company's Database Specialist is able to log in to MySQL and run queries from the bastion host using these details.

When users try to utilize the application hosted in the AWS account, they are presented with a generic error message. The application servers are logging a "could not connect to server: Connection times out" error message to Amazon CloudWatch Logs.

What is the cause of this error?

- A. The user name and password the application is using are incorrect.
- B. The security group assigned to the application servers does not have the necessary rules to allow inbound connections from the DB instance.
- C. The security group assigned to the DB instance does not have the necessary rules to allow inbound connections from the application servers.
- D. The user name and password are correct, but the user is not authorized to use the DB instance.

Answer: C

NEW QUESTION 134

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 volume
- 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 cluster
- F. Promote this replica as a new primary DB instance
- 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 139

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 141

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 143

A financial services company is developing a shared data service that supports different applications from throughout the company. A Database Specialist designed a solution to leverage Amazon ElastiCache for Redis with cluster mode enabled to enhance performance and scalability. The cluster is configured to listen on port 6379.

Which combination of steps should the Database Specialist take to secure the cache data and protect it from unauthorized access? (Choose three.)

- A. Enable in-transit and at-rest encryption on the ElastiCache cluster.
- B. Ensure that Amazon CloudWatch metrics are configured in the ElastiCache cluster.
- C. Ensure the security group for the ElastiCache cluster allows all inbound traffic from itself and inbound traffic on TCP port 6379 from trusted clients only.
- D. Create an IAM policy to allow the application service roles to access all ElastiCache API actions.
- E. Ensure the security group for the ElastiCache clients authorize inbound TCP port 6379 and port 22 traffic from the trusted ElastiCache cluster's security group.
- F. Ensure the cluster is created with the auth-token parameter and that the parameter is used in all subsequent commands.

Answer: ACF

Explanation:

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

NEW QUESTION 148

A company wants to build a new invoicing service for its cloud-native application on AWS. The company has a small development team and wants to focus on service feature development and minimize operations and maintenance as much as possible. The company expects the service to handle billions of requests and millions of new records every day. The service feature requirements, including data access patterns are well-defined. The service has an availability target of 99.99% with a milliseconds latency requirement. The database for the service will be the system of record for invoicing data.

Which database solution meets these requirements at the LOWEST cost?

- A. Amazon Neptune
- B. Amazon Aurora PostgreSQL Serverless
- C. Amazon RDS for PostgreSQL
- D. Amazon DynamoDB

Answer: D

Explanation:

Known patterns, minimum maintenance, milliseconds latency

NEW QUESTION 149

A company is building a new web platform where user requests trigger an AWS Lambda function that performs an insert into an Amazon Aurora MySQL DB cluster. Initial tests with less than 10 users on the new platform yielded successful execution and fast response times. However, upon more extensive tests with the actual target of 3,000 concurrent users, Lambda functions are unable to connect to the DB cluster and receive too many connections errors.

Which of the following will resolve this issue?

- A. Edit the my.cnf file for the DB cluster to increase max_connections
- B. Increase the instance size of the DB cluster
- C. Change the DB cluster to Multi-AZ
- D. Increase the number of Aurora Replicas

Answer: B

Explanation:

Max_connection is a formula in RDS parameter group:

$\text{GREATEST}(\{\log(\text{DBInstanceClassMemory}/805306368)*45\}, \{\log(\text{DBInstanceClassMemory}/8187281408)*100\})$

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraMySQL.Managing.Performance.htm> You can increase the maximum number of connections to your Aurora MySQL DB instance by scaling the instance up to a DB instance class with more memory, or by setting a larger value for the max_connections parameter in the DB parameter group for your instance, up to 16,000. You must change a larger value for the max_connections parameter in the DB parameter group, not edit my.cnf, it is not physical server hosting MySQL.

NEW QUESTION 151

A company is moving its fraud detection application from on premises to the AWS Cloud and is using Amazon Neptune for data storage. The company has set up a 1 Gbps AWS Direct Connect connection to migrate 25 TB of fraud detection data from the on-premises data center to a Neptune DB instance. The company already has an Amazon S3 bucket and an S3 VPC endpoint, and 80% of the company's network bandwidth is available.

How should the company perform this data load?

- A. Use an AWS SDK with a multipart upload to transfer the data from on premises to the S3 bucket
- B. Use the Copy command for Neptune to move the data in bulk from the S3 bucket to the Neptune DB instance.
- C. Use AWS Database Migration Service (AWS DMS) to transfer the data from on premises to the S3 bucket
- D. Use the Loader command for Neptune to move the data in bulk from the S3 bucket to the Neptune DB instance.
- E. Use AWS DataSync to transfer the data from on premises to the S3 bucket
- F. Use the Loader command for Neptune to move the data in bulk from the S3 bucket to the Neptune DB instance.
- G. Use the AWS CLI to transfer the data from on premises to the S3 bucket
- H. Use the Copy command for Neptune to move the data in bulk from the S3 bucket to the Neptune DB instance.

Answer: C

Explanation:

"AWS DataSync is an online data transfer service that simplifies, automates, and accelerates moving data between on-premises storage systems and AWS storage services, and also between AWS storage services."

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

NEW QUESTION 154

An online advertising website uses an Amazon DynamoDB table with on-demand capacity mode as its data store. The website also has a DynamoDB Accelerator (DAX) cluster in the same VPC as its web application server. The application needs to perform infrequent writes and many strongly consistent reads from the data store by querying the DAX cluster.

During a performance audit, a systems administrator notices that the application can look up items by using the DAX cluster. However, the QueryCacheHits metric for the DAX cluster consistently shows 0 while the QueryCacheMisses metric continuously keeps growing in Amazon CloudWatch.

What is the MOST likely reason for this occurrence?

- A. A VPC endpoint was not added to access DynamoDB.
- B. Strongly consistent reads are always passed through DAX to DynamoDB.
- C. DynamoDB is scaling due to a burst in traffic, resulting in degraded performance.
- D. A VPC endpoint was not added to access CloudWatch.

Answer: B

Explanation:

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

"If the request specifies strongly consistent reads, DAX passes the request through to DynamoDB. The results from DynamoDB are not cached in DAX. Instead, they are simply returned to the application."

NEW QUESTION 158

A company is developing a new web application. An AWS CloudFormation template was created as a part of the build process.

Recently, a change was made to an AWS::RDS::DBInstance resource in the template. The CharacterSetName property was changed to allow the application to process international text. A change set was generated using the new template, which indicated that the existing DB instance should be replaced during an upgrade.

What should a database specialist do to prevent data loss during the stack upgrade?

- A. Create a snapshot of the DB instance
- B. Modify the template to add the DBSnapshotIdentifier property with the ID of the DB snapshot
- C. Update the stack.
- D. Modify the stack policy using the aws cloudformation update-stack command and the set-stack-policy command, then make the DB resource protected.
- E. Create a snapshot of the DB instance
- F. Update the stack
- G. Restore the database to a new instance.
- H. Deactivate any applications that are using the DB instance
- I. Create a snapshot of the DB instance. Modify the template to add the DBSnapshotIdentifier property with the ID of the DB snapshot
- J. Update the stack and reactivate the applications.

Answer: D

Explanation:

To preserve your data, perform the following procedure:

* 1. Deactivate any applications that are using the DB instance so that there's no activity on the DB instance. * 2. Create a snapshot of the DB instance. For more information about creating DB snapshots

* 3. If you want to restore your instance using a DB snapshot, modify the updated template with your DB instance changes and add the DBSnapshotIdentifier property with the ID of the DB snapshot that you want to use

* 4. Update the stack.

NEW QUESTION 161

A Database Specialist is migrating an on-premises Microsoft SQL Server application database to Amazon RDS for PostgreSQL using AWS DMS. The application

requires minimal downtime when the RDS DB instance goes live.
What change should the Database Specialist make to enable the migration?

- A. Configure the on-premises application database to act as a source for an AWS DMS full load with ongoing change data capture (CDC)
- B. Configure the AWS DMS replication instance to allow both full load and ongoing change data capture (CDC)
- C. Configure the AWS DMS task to generate full logs to allow for ongoing change data capture (CDC)
- D. Configure the AWS DMS connections to allow two-way communication to allow for ongoing change data capture (CDC)

Answer: A

Explanation:

"requires minimal downtime when the RDS DB instance goes live" in order to do CDC: "you must first ensure that ARCHIVELOG MODE is on to provide information to LogMiner. AWS DMS uses LogMiner to read information from the archive logs so that AWS DMS can capture changes"
<https://docs.aws.amazon.com/dms/latest/sbs/chap-oracle2postgresql.steps.configureoracle.html> "If you want to capture and apply changes (CDC), then you also need the following privileges."

NEW QUESTION 164

A company needs to migrate Oracle Database Standard Edition running on an Amazon EC2 instance to an Amazon RDS for Oracle DB instance with Multi-AZ. The database supports an ecommerce website that runs continuously. The company can only provide a maintenance window of up to 5 minutes. Which solution will meet these requirements?

- A. Configure Oracle Real Application Clusters (RAC) on the EC2 instance and the RDS DB instance. Update the connection string to point to the RAC cluster
- B. Once the EC2 instance and RDS DB instance are in sync, fail over from Amazon EC2 to Amazon RDS.
- C. Export the Oracle database from the EC2 instance using Oracle Data Pump and perform an import into Amazon RDS
- D. Stop the application for the entire process
- E. When the import is complete, change the database connection string and then restart the application.
- F. Configure AWS DMS with the EC2 instance as the source and the RDS DB instance as the destination. Stop the application when the replication is in sync, change the database connection string, and then restart the application.
- G. Configure AWS DataSync with the EC2 instance as the source and the RDS DB instance as the destination
- H. Stop the application when the replication is in sync, change the database connection string, and then restart the application.

Answer: C

NEW QUESTION 168

A company uses Amazon Aurora MySQL as the primary database engine for many of its applications. A database specialist must create a dashboard to provide the company with information about user connections to databases. According to compliance requirements, the company must retain all connection logs for at least 7 years.
Which solution will meet these requirements MOST cost-effectively?

- A. Enable advanced auditing on the Aurora cluster to log CONNECT event
- B. Export audit logs from Amazon CloudWatch to Amazon S3 by using an AWS Lambda function that is invoked by an Amazon EventBridge (Amazon CloudWatch Events) scheduled event
- C. Build a dashboard by using Amazon QuickSight.
- D. Capture connection attempts to the Aurora cluster with AWS Cloud Trail by using the DescribeEvents API operation
- E. Create a CloudTrail trail to export connection logs to Amazon S3. Build a dashboard by using Amazon QuickSight.
- F. Start a database activity stream for the Aurora cluster
- G. Push the activity records to an Amazon Kinesis data stream
- H. Build a dynamic dashboard by using AWS Lambda.
- I. Publish the DatabaseConnections metric for the Aurora DB instances to Amazon CloudWatch
- J. Build a dashboard by using CloudWatch dashboards.

Answer: A

Explanation:

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

NEW QUESTION 169

To meet new data compliance requirements, a company needs to keep critical data durably stored and readily accessible for 7 years. Data that is more than 1 year old is considered archival data and must automatically be moved out of the Amazon Aurora MySQL DB cluster every week. On average, around 10 GB of new data is added to the database every month. A database specialist must choose the most operationally efficient solution to migrate the archival data to Amazon S3. Which solution meets these requirements?

- A. Create a custom script that exports archival data from the DB cluster to Amazon S3 using a SQL view, then deletes the archival data from the DB cluster
- B. Launch an Amazon EC2 instance with a weekly cron job to execute the custom script.
- C. Configure an AWS Lambda function that exports archival data from the DB cluster to Amazon S3 using a SELECT INTO OUTFILE S3 statement, then deletes the archival data from the DB cluster
- D. Schedule the Lambda function to run weekly using Amazon EventBridge (Amazon CloudWatch Events).
- E. Configure two AWS Lambda functions: one that exports archival data from the DB cluster to Amazon S3 using the mysqldump utility, and another that deletes the archival data from the DB cluster
- F. Schedule both Lambda functions to run weekly using Amazon EventBridge (Amazon CloudWatch Events).
- G. Use AWS Database Migration Service (AWS DMS) to continually export the archival data from the DB cluster to Amazon S3. Configure an AWS Data Pipeline process to run weekly that executes a custom SQL script to delete the archival data from the DB cluster.

Answer: B

Explanation:

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraMySQL.Integrating.SaveIntoS3.htm>

NEW QUESTION 173

A security team is conducting an audit for a financial company. The security team discovers that the database credentials of an Amazon RDS for MySQL DB instance are hardcoded in the source code. The source code is stored in a shared location for automatic deployment and is exposed to all users who can access the location.

A database specialist must use encryption to ensure that the credentials are not visible in the source code. Which solution will meet these requirements?

- A. Use an AWS Key Management Service (AWS KMS) key to encrypt the most recent database backup. Restore the backup as a new database to activate encryption.
- B. Store the source code to access the credentials in an AWS Systems Manager Parameter Store secure string parameter that is encrypted by AWS Key Management Service (AWS KMS). Access the code with calls to Systems Manager.
- C. Store the credentials in an AWS Systems Manager Parameter Store secure string parameter that is encrypted by AWS Key Management Service (AWS KMS). Access the credentials with calls to Systems Manager.
- D. Use an AWS Key Management Service (AWS KMS) key to encrypt the DB instance at rest.
- E. Activate RDS encryption in transit by using SSL certificates.

Answer: C

Explanation:

only creds in system manager secure parameter.

NEW QUESTION 174

A startup company in the travel industry wants to create an application that includes a personal travel assistant to display information for nearby airports based on user location. The application will use Amazon DynamoDB and must be able to access and display attributes such as airline names, arrival times, and flight numbers. However, the application must not be able to access or display pilot names or passenger counts.

Which solution will meet these requirements MOST cost-effectively?

- A. Use a proxy tier between the application and DynamoDB to regulate access to specific tables, items, and attributes.
- B. Use IAM policies with a combination of IAM conditions and actions to implement fine-grained access control.
- C. Use DynamoDB resource policies to regulate access to specific tables, items, and attributes.
- D. Configure an AWS Lambda function to extract only allowed attributes from tables based on user profiles.

Answer: B

Explanation:

<https://aws.amazon.com/blogs/aws/fine-grained-access-control-for-amazon-dynamodb/>

NEW QUESTION 179

A Database Specialist migrated an existing production MySQL database from on-premises to an Amazon RDS for MySQL DB instance. However, after the migration, the database needed to be encrypted at rest using AWS KMS. Due to the size of the database, reloading, the data into an encrypted database would be too time-consuming, so it is not an option.

How should the Database Specialist satisfy this new requirement?

- A. Create a snapshot of the unencrypted RDS DB instance.
- B. Create an encrypted copy of the unencrypted snapshot.
- C. Restore the encrypted snapshot copy.
- D. Modify the RDS DB instance.
- E. Enable the AWS KMS encryption option that leverages the AWS CLI.
- F. Restore an unencrypted snapshot into a MySQL RDS DB instance that is encrypted.
- G. Create an encrypted read replica of the RDS DB instance.
- H. Promote it to the master.

Answer: A

Explanation:

"However, because you can encrypt a copy of an unencrypted DB snapshot, you can effectively add encryption to an unencrypted DB instance. That is, you can create a snapshot of your DB instance, and then create an encrypted copy of that snapshot. You can then restore a DB instance from the encrypted snapshot, and thus you have an encrypted copy of your original DB instance. For more information, see Copying a Snapshot."

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

NEW QUESTION 183

A business needs a data warehouse system that stores data consistently and in a highly organized fashion. The organization demands rapid response times for end-user inquiries including current-year data, and users must have access to the whole 15-year dataset when necessary. Additionally, this solution must be able to manage a variable volume of incoming inquiries. Costs associated with storing the 100 TB of data must be maintained to a minimum.

Which solution satisfies these criteria?

- A. Leverage an Amazon Redshift data warehouse solution using a dense storage instance type while keeping all the data on local Amazon Redshift storage.
- B. Provision enough instances to support high demand.
- C. Leverage an Amazon Redshift data warehouse solution using a dense storage instance to store the most recent data and keep historical data on Amazon S3 and access it using the Amazon Redshift Spectrum layer.
- D. Keep historical data on Amazon S3 and access it using the Amazon Redshift Spectrum layer.
- E. Provision enough instances to support high demand.
- F. Leverage an Amazon Redshift data warehouse solution using a dense storage instance to store the most recent data and keep historical data on Amazon S3 and access it using the Amazon Redshift Spectrum layer.
- G. Keep historical data on Amazon S3 and access it using the Amazon Redshift Spectrum layer.
- H. Enable Amazon Redshift Concurrency Scaling.
- I. Leverage an Amazon Redshift data warehouse solution using a dense storage instance to store the most recent data and keep historical data on Amazon S3 and access it using the Amazon Redshift Spectrum layer.
- J. Keep historical data on Amazon S3 and access it using the Amazon Redshift Spectrum layer.
- K. Leverage Amazon Redshift elastic resize.

Answer: C

Explanation:

<https://docs.aws.amazon.com/redshift/latest/dg/concurrency-scaling.html>

"With the Concurrency Scaling feature, you can support virtually unlimited concurrent users and concurrent queries, with consistently fast query performance. When concurrency scaling is enabled, Amazon Redshift automatically adds additional cluster capacity when you need it to process an increase in concurrent read queries. Write operations continue as normal on your main cluster. Users always see the most current data, whether the queries run on the main cluster or on a concurrency scaling cluster. You're charged for concurrency scaling clusters only for the time they're in use. For more information about pricing, see Amazon Redshift pricing. You manage which queries are sent to the concurrency scaling cluster by configuring WLM queues. When you enable concurrency scaling for a queue, eligible queries are sent to the concurrency scaling cluster instead of waiting in line."

NEW QUESTION 184

A database specialist needs to review and optimize an Amazon DynamoDB table that is experiencing performance issues. A thorough investigation by the database specialist reveals that the partition key is causing hot partitions, so a new partition key is created. The database specialist must effectively apply this new partition key to all existing and new data. How can this solution be implemented?

- A. Use Amazon EMR to export the data from the current DynamoDB table to Amazon S3. Then use Amazon EMR again to import the data from Amazon S3 into a new DynamoDB table with the new partition key.
- B. Use AWS DMS to copy the data from the current DynamoDB table to Amazon S3. Then import the DynamoDB table to create a new DynamoDB table with the new partition key.
- C. Use the AWS CLI to update the DynamoDB table and modify the partition key.
- D. Use the AWS CLI to back up the DynamoDB table.
- E. Then use the restore-table-from-backup command and modify the partition key.

Answer: A

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/back-up-dynamodb-s3/>

NEW QUESTION 189

A startup company is building a new application to allow users to visualize their on-premises and cloud networking components. The company expects billions of components to be stored and requires responses in milliseconds. The application should be able to identify:

- The networks and routes affected if a particular component fails.
- The networks that have redundant routes between them.
- The networks that do not have redundant routes between them.
- The fastest path between two networks.

Which database engine meets these requirements?

- A. Amazon Aurora MySQL
- B. Amazon Neptune
- C. Amazon ElastiCache for Redis
- D. Amazon DynamoDB

Answer: B

NEW QUESTION 190

A company's database specialist disabled TLS on an Amazon DocumentDB cluster to perform benchmarking tests. A few days after this change was implemented, a database specialist trainee accidentally deleted multiple tables. The database specialist restored the database from available snapshots. An hour after restoring the cluster, the database specialist is still unable to connect to the new cluster endpoint. What should the database specialist do to connect to the new, restored Amazon DocumentDB cluster?

- A. Change the restored cluster's parameter group to the original cluster's custom parameter group.
- B. Change the restored cluster's parameter group to the Amazon DocumentDB default parameter group.
- C. Configure the interface VPC endpoint and associate the new Amazon DocumentDB cluster.
- D. Run the syncInstances command in AWS DataSync.

Answer: A

Explanation:

You can't modify the parameter settings of the default parameter groups. You can use a DB parameter group to act as a container for engine configuration values that are applied to one or more DB instances. If you create a DB instance without specifying a DB parameter group, the DB instance uses a default DB parameter group. Each default DB parameter group contains database engine defaults and Amazon RDS system defaults. You can't modify the parameter settings of a default parameter group. Instead, you create your own parameter group where you choose your own parameter settings. Not all DB engine parameters can be changed in a parameter group that you create.

NEW QUESTION 191

In North America, a business launched a mobile game that swiftly expanded to 10 million daily active players. The game's backend is hosted on AWS and makes considerable use of a TTL-configured Amazon DynamoDB table.

When an item is added or changed, its TTL is set to 600 seconds plus the current epoch time. The game logic is reliant on the purging of outdated data in order to compute rewards points properly. At times, items from the table are read that are many hours beyond their TTL expiration.

How should a database administrator resolve this issue?

- A. Use a client library that supports the TTL functionality for DynamoDB.
- B. Include a query filter expression to ignore items with an expired TTL.
- C. Set the ConsistentRead parameter to true when querying the table.
- D. Create a local secondary index on the TTL attribute.

Answer: B

Explanation:

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/howitworks-ttl.html>

NEW QUESTION 192

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questions and Answers in PDF Format

DBS-C01 Practice Exam Features:

- * DBS-C01 Questions and Answers Updated Frequently
- * DBS-C01 Practice Questions Verified by Expert Senior Certified Staff
- * DBS-C01 Most Realistic Questions that Guarantee you a Pass on Your First Try
- * DBS-C01 Practice Test Questions in Multiple Choice Formats and Updates for 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The DBS-C01 Practice Test Here](#)