

Amazon-Web-Services

Exam Questions DBS-C01

AWS Certified Database - Specialty



NEW QUESTION 1

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

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

Answer: A

NEW QUESTION 2

A 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 3

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 CloudTrail
- 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 4

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

What should the database specialist do to meet these requirements?

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

Answer: A

Explanation:

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

NEW QUESTION 5

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

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

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

Answer: BEF

Explanation:

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

NEW QUESTION 6

A large ecommerce company uses Amazon DynamoDB to handle the transactions on its web portal. Traffic patterns throughout the year are usually stable; however, a large event is planned. The company knows that traffic will increase by up to 10 times the normal load over the 3-day event. When sale prices are published during the event, traffic will spike rapidly.

How should a Database Specialist ensure DynamoDB can handle the increased traffic?

- A. Ensure the table is always provisioned to meet peak needs
- B. Allow burst capacity to handle the additional load
- C. Set an AWS Application Auto Scaling policy for the table to handle the increase in traffic
- D. Preprovision additional capacity for the known peaks and then reduce the capacity after the event

Answer: D

Explanation:

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/bp-partition-key-design.html#bp-partition>

"DynamoDB provides some flexibility in your per-partition throughput provisioning by providing burst capacity. Whenever you're not fully using a partition's throughput, DynamoDB reserves a portion of that unused capacity for later bursts of throughput to handle usage spikes. DynamoDB currently retains up to 5 minutes (300 seconds) of unused read and write capacity. During an occasional burst of read or write activity, these extra capacity units can be consumed quickly—even faster than the per-second provisioned throughput capacity that you've defined for your table. DynamoDB can also consume burst capacity for background maintenance and other tasks without prior notice. Note that these burst capacity details might change in the future."

NEW QUESTION 7

A single MySQL database was moved to Amazon Aurora by a business. The production data is stored in a database cluster in VPC PROD, whereas 12 testing environments are hosted in VPC TEST with the same AWS account. Testing has a negligible effect on the test data. The development team requires that each environment be updated nightly to ensure that each test database has daily production data.

Which migration strategy will be the quickest and least expensive to implement?

- A. Run the master in Amazon Aurora MySQL
- B. Create 12 clones in VPC_TEST, and script the clones to be deleted and re-created nightly.
- C. Run the master in Amazon Aurora MySQL
- D. Take a nightly snapshot, and restore it into 12 databases in VPC_TEST using Aurora Serverless.
- E. Run the master in Amazon Aurora MySQL
- F. Create 12 Aurora Replicas in VPC_TEST, and script the replicas to be deleted and re-created nightly.
- G. Run the master in Amazon Aurora MySQL using Aurora Serverless
- H. Create 12 clones in VPC_TEST, and script the clones to be deleted and re-created nightly.

Answer: A

Explanation:

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

NEW QUESTION 8

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 9

A banking company recently launched an Amazon RDS for MySQL DB instance as part of a proof-of-concept project. A database specialist has configured automated database snapshots. As a part of routine testing, the database specialist noticed one day that the automated database snapshot was not created. Which of the following are possible reasons why the snapshot was not created? (Choose two.)

- A. A copy of the RDS automated snapshot for this DB instance is in progress within the same AWS Region.
- B. A copy of the RDS automated snapshot for this DB instance is in progress in a different AWS Region.
- C. The RDS maintenance window is not configured.
- D. The RDS DB instance is in the STORAGE_FULL state.
- E. RDS event notifications have not been enabled.

Answer: AD

Explanation:

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

NEW QUESTION 10

A financial company is running an Amazon Redshift cluster for one of its data warehouse solutions. The company needs to generate connection logs, user logs, and user activity logs. The company also must make these logs available for future analysis.

Which combination of steps should a database specialist take to meet these requirements? (Choose two.)

- A. Edit the database configuration of the cluster by enabling audit logging
- B. Direct the logging to a specified log group in Amazon CloudWatch Logs.
- C. Edit the database configuration of the cluster by enabling audit logging
- D. Direct the logging to a specified Amazon S3 bucket
- E. Modify the cluster by enabling continuous delivery of AWS CloudTrail logs to Amazon S3.
- F. Create a new parameter group with the `enable_user_activity_logging` parameter set to `tru`
- G. Configure the cluster to use the new parameter group.
- H. Modify the system table to enable logging for each user.

Answer: AD

Explanation:

AWS CloudWatch Logs are stored indefinitely and CloudWatch Log Insights is used to analyze the logs and query upon them.

<https://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/AnalyzingLogData.html>

<https://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/WhatIsCloudWatchLogs.html>

"Log retention – By default, logs are kept indefinitely and never expire. You can adjust the retention policy for each log group, keeping the indefinite retention, or choosing a retention period between 10 years and one day."

<https://docs.aws.amazon.com/redshift/latest/mgmt/db-auditing.html>

NEW QUESTION 10

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 13

A Database Specialist is creating Amazon DynamoDB tables, Amazon CloudWatch alarms, and associated infrastructure for an Application team using a development AWS account. The team wants a deployment method that will standardize the core solution components while managing environment-specific settings separately, and wants to minimize rework due to configuration errors.

Which process should the Database Specialist recommend to meet these requirements?

- A. Organize common and environmental-specific parameters hierarchically in the AWS Systems Manager Parameter Store, then reference the parameters dynamically from an AWS CloudFormation template
- B. Deploy the CloudFormation stack using the environment name as a parameter.
- C. Create a parameterized AWS CloudFormation template that builds the required object
- D. Keep separate environment parameter files in separate Amazon S3 bucket
- E. Provide an AWS CLI command that deploys the CloudFormation stack directly referencing the appropriate parameter bucket.
- F. Create a parameterized AWS CloudFormation template that builds the required object
- G. Import the template into the CloudFormation interface in the AWS Management Console
- H. Make the required changes to the parameters and deploy the CloudFormation stack.
- I. Create an AWS Lambda function that builds the required objects using an AWS SD
- J. Set the required parameter values in a test event in the Lambda console for each environment that the Application team can modify, as needed
- K. Deploy the infrastructure by triggering the test event in the console.

Answer: A

Explanation:

<https://aws.amazon.com/blogs/mt/integrating-aws-cloudformation-with-aws-systems-manager-parameter-store/>

NEW QUESTION 17

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

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

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

- D. Allow AWS DMS to handle syncing change data capture (CDC) data from the source to the target database.
- E. Use the AWS Schema Conversion Tool (AWS SCT) to migrate the full load of the source database over a VPN tunnel using the internet for its primary connectio
- F. Allow AWS SCT to handle syncing change data capture (CDC) data from the source to the target database.

Answer: A

Explanation:

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

NEW QUESTION 19

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

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

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

Answer: C

Explanation:

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

NEW QUESTION 20

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 Regio
- B. Create a new instance from the snapshot in the ap-southeast-1 Regio
- 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 Regio
- E. Reconfigure the ap-southeast-1 front-end dashboard to access this instance.
- F. Create a new RDS instance in the ap-southeast-1 Regio
- G. Use AWS DMS and change data capture (CDC) to update the new instance in the ap-southeast-1 Regio
- 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 Regio
- 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 24

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

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

Answer: D

Explanation:

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

NEW QUESTION 27

An ecommerce company uses a backend application that stores data in an Amazon DynamoDB table. The backend application runs in a private subnet in a VPC and must connect to this table.

The company must minimize any network latency that results from network connectivity issues, even during periods of heavy application usage. A database administrator also needs the ability to use a private connection to connect to the DynamoDB table from the application. Which solution will meet these requirements?

- A. Use network ACLs to ensure that any outgoing or incoming connections to any port except DynamoDB are deactivate
- B. Encrypt API calls by using TLS.
- C. Create a VPC endpoint for DynamoDB in the application's VP
- D. Use the VPC endpoint to access the table.
- E. Create an AWS Lambda function that has access to DynamoD
- F. Restrict outgoing access only to this Lambda function from the application.
- G. Use a VPN to route all communication to DynamoDB through the company's own corporate network infrastructure.

Answer: B

Explanation:

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/vpc-endpoints-dynamodb.html>

NEW QUESTION 32

Recently, an ecommerce business transferred one of its SQL Server databases to an Amazon RDS for SQL Server Enterprise Edition database instance. The corporation anticipates an increase in read traffic as a result of an approaching sale. To accommodate the projected read load, a database professional must establish a read replica of the database instance.

Which procedures should the database professional do prior to establishing the read replica? (Select two.)

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

Answer: BC

Explanation:

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

NEW QUESTION 34

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

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

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

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

Answer: B

Explanation:

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

NEW QUESTION 37

A database professional is tasked with the task of migrating 25 GB of data files from an on-premises storage system to an Amazon Neptune database.

Which method of data loading is the FASTEST?

- A. Upload the data to Amazon S3 and use the Loader command to load the data from Amazon S3 into the Neptune database.
- B. Write a utility to read the data from the on-premises storage and run INSERT statements in a loop to load the data into the Neptune database.
- C. Use the AWS CLI to load the data directly from the on-premises storage into the Neptune database.
- D. Use AWS DataSync to load the data directly from the on-premises storage into the Neptune database.

Answer: A

Explanation:

- * 1. Copy the data files to an Amazon Simple Storage Service (Amazon S3) bucket.
- * 2. Create an IAM role with Read and List access to the bucket.
- * 3. Create an Amazon S3 VPC endpoint.
- * 4. Start the Neptune loader by sending a request via HTTP to the Neptune DB instance.
- * 5. The Neptune DB instance assumes the IAM role to load the data from the bucket.

NEW QUESTION 39

A Database Specialist is designing a disaster recovery strategy for a production Amazon DynamoDB table. The table uses provisioned read/write capacity mode, global secondary indexes, and time to live (TTL). The Database Specialist has restored the latest backup to a new table.

To prepare the new table with identical settings, which steps should be performed? (Choose two.)

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

E. Set the provisioned read and write capacity

Answer: BC

Explanation:

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

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

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

NEW QUESTION 40

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 tables
- E. Use the game name as the partition key in both tables
- 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 41

A database specialist is designing an enterprise application for a large company. The application uses Amazon DynamoDB with DynamoDB Accelerator (DAX). The database specialist observes that most of the queries are not found in the DAX cache and that they still require DynamoDB table reads.

What should the database specialist review first to improve the utility of DAX?

- A. The DynamoDB `ConsumedReadCapacityUnits` metric
- B. The trust relationship to perform the DynamoDB API calls
- C. The DAX cluster's TTL setting
- D. The validity of customer-specified AWS Key Management Service (AWS KMS) keys for DAX encryption at rest

Answer: C

Explanation:

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/DAX.cluster-management.html#DAX.clu>

NEW QUESTION 44

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 49

A company's Security department established new requirements that state internal users must connect to an existing Amazon RDS for SQL Server DB instance using their corporate Active Directory (AD) credentials. A Database Specialist must make the modifications needed to fulfill this requirement.

Which combination of actions should the Database Specialist take? (Choose three.)

- A. Disable Transparent Data Encryption (TDE) on the RDS SQL Server DB instance.

- B. Modify the RDS SQL Server DB instance to use the directory for Windows authentication
- C. Create appropriate new logins.
- D. Use the AWS Management Console to create an AWS Managed Microsoft A
- E. Create a trust relationship with the corporate AD.
- F. Stop the RDS SQL Server DB instance, modify it to use the directory for Windows authentication, and start it again
- G. Create appropriate new logins.
- H. Use the AWS Management Console to create an AD Connecto
- I. Create a trust relationship with the corporate AD.
- J. Configure the AWS Managed Microsoft AD domain controller Security Group.

Answer: BCF

Explanation:

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

NEW QUESTION 53

An online gaming company is planning to launch a new game with Amazon DynamoDB as its data store. The database should be designated to support the following use cases:

Update scores in real time whenever a player is playing the game. Retrieve a player's score details for a specific game session.

A Database Specialist decides to implement a DynamoDB table. Each player has a unique user_id and each game has a unique game_id.

Which choice of keys is recommended for the DynamoDB table?

- A. Create a global secondary index with game_id as the partition key
- B. Create a global secondary index with user_id as the partition key
- C. Create a composite primary key with game_id as the partition key and user_id as the sort key
- D. Create a composite primary key with user_id as the partition key and game_id as the sort key

Answer: D

Explanation:

<https://aws.amazon.com/blogs/database/amazon-dynamodb-gaming-use-cases-and-design-patterns/> "EA uses the user ID as the partition key and primary key (a 1:1 modeling pattern)."

<https://aws.amazon.com/blogs/database/choosing-the-right-dynamodb-partition-key/>

"Partition key and sort key: Referred to as a composite primary key, this type of key is composed of two attributes. The first attribute is the partition key, and the second attribute is the sort key."

NEW QUESTION 58

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

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

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

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

Answer: B

Explanation:

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

NEW QUESTION 59

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

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

- A. Set up a VPN tunnel for encrypting data over the network from the data center to AW
- B. Leverage AWS SCT and apply the converted schema to Amazon Redshift
- C. Once complete, start an AWS DMS task to move the data from the source to Amazon S3. Use AWS Glue to load the data from Amazon S3 to Amazon Redshift.
- D. Leverage AWS SCT and apply the converted schema to Amazon Redshift
- E. Start an AWS DMS task with two AWS Snowball Edge devices to copy data from on-premises to Amazon S3 with AWS KMS encryption
- F. Use AWS DMS to finish copying data to Amazon Redshift.
- G. Leverage AWS SCT and apply the converted schema to Amazon Redshift
- H. Once complete, use a fleet of 10 TB dedicated encrypted drives using the AWS Import/Export feature to copy data from on-premises to Amazon S3 with AWS KMS encryption
- I. Use AWS Glue to load the data to Amazon Redshift.
- J. Set up a VPN tunnel for encrypting data over the network from the data center to AW
- K. Leverage a native database export feature to export the data and compress the file
- L. Use the aws s3 cp multi-part upload command to upload these files to Amazon S3 with AWS KMS encryption
- M. Once complete, load the data to Amazon Redshift using AWS Glue.

Answer: B

Explanation:

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

NEW QUESTION 63

An AWS CloudFormation stack that included an Amazon RDS DB instance was accidentally deleted and recent data was lost. A Database Specialist needs to add RDS settings to the CloudFormation template to reduce the chance of accidental instance data loss in the future. Which settings will meet this requirement? (Choose 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: ACF

NEW QUESTION 67

A corporation wishes to move a 1 TB Oracle database from its current location to an Amazon Aurora PostgreSQL DB cluster. The database specialist at the firm noticed that the Oracle database stores 100 GB of large binary objects (LOBs) across many tables. The Oracle database supports LOBs up to 500 MB in size and an average of 350 MB. AWS DMS was picked by the Database Specialist to transfer the data with the most replication instances. How should the database specialist improve the transfer of the database to AWS DMS?

- A. Create a single task using full LOB mode with a LOB chunk size of 500 MB to migrate the data and LOBs together
- B. Create two tasks: task1 with LOB tables using full LOB mode with a LOB chunk size of 500 MB and task2 without LOBs
- C. Create two tasks: task1 with LOB tables using limited LOB mode with a maximum LOB size of 500 MB and task 2 without LOBs
- D. Create a single task using limited LOB mode with a maximum LOB size of 500 MB to migrate data and LOBs together

Answer: C

Explanation:

https://docs.aws.amazon.com/dms/latest/userguide/CHAP_BestPractices.html#CHAP_BestPractices.LOBS, "AWS DMS migrates LOB data in two phases: 1. AWS DMS creates a new row in the target table and populates the row with all data except the associated LOB value. 2. AWS DMS updates the row in the target table with the LOB data." This means that we would need two tasks, one per phase and use limited LOB mode for best performance.

NEW QUESTION 69

A Database Specialist is creating a new Amazon Neptune DB cluster, and is attempting to load data from Amazon S3 into the Neptune DB cluster using the Neptune bulk loader API. The Database Specialist receives the following error: "Unable to connect to s3 endpoint. Provided source = s3://mybucket/graphdata/ and region = us-east-1. Please verify your S3 configuration." Which combination of actions should the Database Specialist take to troubleshoot the problem? (Choose 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: BD

NEW QUESTION 73

A company is releasing a new mobile game featuring a team play mode. As a group of mobile device users play together, an item containing their statuses is updated in an Amazon DynamoDB table. Periodically, the other users' devices read the latest statuses of their teammates from the table using the BatchGetItem operation. Prior to launch, some testers submitted bug reports claiming that the status data they were seeing in the game was not up-to-date. The developers are unable to replicate this issue and have asked a database specialist for a recommendation. Which recommendation would resolve this issue?

- A. Ensure the DynamoDB table is configured to be always consistent.
- B. Ensure the BatchGetItem operation is called with the ConsistentRead parameter set to false.
- C. Enable a stream on the DynamoDB table and subscribe each device to the stream to ensure all devices receive up-to-date status information.
- D. Ensure the BatchGetItem operation is called with the ConsistentRead parameter set to true.

Answer: D

Explanation:

https://docs.aws.amazon.com/ja_jp/amazondynamodb/latest/developerguide/API_BatchGetItem_v20111205.htm By default, BatchGetItem performs eventually consistent reads on every table in the request. If you want strongly consistent reads instead, you can set ConsistentRead to true for any or all tables.

NEW QUESTION 78

A small startup firm wishes to move a 4 TB MySQL database from on-premises to AWS through an Amazon RDS for MySQL DB instance. Which migration approach would result in 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: D

Explanation:

<https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/MySQL.Procedural.Importing.NonRDSRepl.html>
<https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/MySQL.Procedural.Importing.External.Repl.html>

NEW QUESTION 83

A company is looking to move an on-premises IBM Db2 database running AIX on an IBM POWER7 server. Due to escalating support and maintenance costs, the company is exploring the option of moving the workload to an Amazon Aurora PostgreSQL DB cluster. What is the quickest way for the company to gather data on the migration compatibility?

- A. Perform a logical dump from the Db2 database and restore it to an Aurora DB cluster
- B. Identify the gaps and compatibility of the objects migrated by comparing row counts from source and target tables.
- C. Run AWS DMS from the Db2 database to an Aurora DB cluster
- D. Identify the gaps and compatibility of the objects migrated by comparing the row counts from source and target tables.
- E. Run native PostgreSQL logical replication from the Db2 database to an Aurora DB cluster to evaluate the migration compatibility.
- F. Run the AWS Schema Conversion Tool (AWS SCT) from the Db2 database to an Aurora DB cluster. Create a migration assessment report to evaluate the migration compatibility.

Answer: D

NEW QUESTION 87

A company is going through a security audit. The audit team has identified cleartext master user password in the AWS CloudFormation templates for Amazon RDS for MySQL DB instances. The audit team has flagged this as a security risk to the database team. What should a database specialist do to mitigate this risk?

- A. Change all the databases to use AWS IAM for authentication and remove all the cleartext passwords in CloudFormation templates.
- B. Use an AWS Secrets Manager resource to generate a random password and reference the secret in the CloudFormation template.
- C. Remove the passwords from the CloudFormation templates so Amazon RDS prompts for the password when the database is being created.
- D. Remove the passwords from the CloudFormation template and store them in a separate file
- E. Replace the passwords by running CloudFormation using a sed command.

Answer: B

Explanation:

<https://aws.amazon.com/blogs/infrastructure-and-automation/securing-passwords-in-aws-quick-starts-using-aws>

NEW QUESTION 91

A company has a heterogeneous six-node production Amazon Aurora DB cluster that handles online transaction processing (OLTP) for the core business and OLAP reports for the human resources department. To match compute resources to the use case, the company has decided to have the reporting workload for the human resources department be directed to two small nodes in the Aurora DB cluster, while every other workload goes to four large nodes in the same DB cluster. Which option would ensure that the correct nodes are always available for the appropriate workload while meeting these requirements?

- A. Use the writer endpoint for OLTP and the reader endpoint for the OLAP reporting workload.
- B. Use automatic scaling for the Aurora Replica to have the appropriate number of replicas for the desired workload.
- C. Create additional readers to cater to the different scenarios.
- D. Use custom endpoints to satisfy the different workloads.

Answer: D

Explanation:

<https://aws.amazon.com/about-aws/whats-new/2018/11/amazon-aurora-simplifies-workload-management-with-c> You can now create custom endpoints for Amazon Aurora databases. This allows you to distribute and load balance workloads across different sets of database instances in your Aurora cluster. For example, you may provision a set of Aurora Replicas to use an instance type with higher memory capacity in order to run an analytics workload. A custom endpoint can then help you route the analytics workload to these appropriately-configured instances, while keeping other instances in your cluster isolated from this workload. As you add or remove instances from the custom endpoint to match your workload, the endpoint helps spread the load around.

NEW QUESTION 92

A company is loading sensitive data into an Amazon Aurora MySQL database. To meet compliance requirements, the company needs to enable audit logging on the Aurora MySQL DB cluster to audit database activity. This logging will include events such as

connections, disconnections, queries, and tables queried. The company also needs to publish the DB logs to Amazon CloudWatch to perform real-time data analysis.

Which solution meets these requirements?

- A. Modify the default option group parameters to enable Advanced Auditin
- B. Restart the database for the changes to take effect.
- C. Create a custom DB cluster parameter grou
- D. Modify the parameters for Advanced Auditin
- E. Modify the cluster to associate the new custom DB parameter group with the Aurora MySQL DB cluster.
- F. Take a snapshot of the databas
- G. Create a new DB instance, and enable custom auditing and logging to CloudWate
- H. Deactivate the DB instance that has no logging.
- I. Enable AWS CloudTrail for the DB instanc
- J. Create a filter that provides only connections, disconnections, queries, and tables queried.

Answer: B

Explanation:

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

NEW QUESTION 95

A company has an application that uses an Amazon DynamoDB table as its data store. During normal business days, the throughput requirements from the application are uniform and consist of 5 standard write calls per second to the DynamoDB table. Each write call has 2 KB of data. For 1 hour each day, the company runs an additional automated job on the DynamoDB table that makes 20 write requests per second. No other application writes to the DynamoDB table. The DynamoDB table does not have to meet any additional capacity requirements. How should a database specialist configure the DynamoDB table's capacity to meet these requirements MOST cost-effectively?

- A. Use DynamoDB provisioned capacity with 5 WCUs and auto scaling.
- B. Use DynamoDB provisioned capacity with 5 WCUs and a write-through cache that DynamoDB Accelerator (DAX) provides.
- C. Use DynamoDB provisioned capacity with 10 WCUs and auto scaling.
- D. Use DynamoDB provisioned capacity with 10 WCUs and no auto scaling.

Answer: C

Explanation:

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

NEW QUESTION 97

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

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

Answer: C

Explanation:

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

NEW QUESTION 101

A large company is using an Amazon RDS for Oracle Multi-AZ DB instance with a Java application. As a part of its disaster recovery annual testing, the company would like to simulate an Availability Zone failure and record how the application reacts during the DB instance failover activity. The company does not want to make any code changes for this activity.

What should the company do to achieve this in the shortest amount of time?

- A. Use a blue-green deployment with a complete application-level failover test
- B. Use the RDS console to reboot the DB instance by choosing the option to reboot with failover
- C. Use RDS fault injection queries to simulate the primary node failure
- D. Add a rule to the NACL to deny all traffic on the subnets associated with a single Availability Zone

Answer: B

Explanation:

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_RebootInstance.html <https://exain.wordpress.com/2017/07/12/amazon-rds-multi-az-setup-failover-simulation/>

"Rebooting with failover is beneficial when you want to simulate a failure of a DB instance for testing, or restore operations to the original AZ after a failover occurs."

NEW QUESTION 106

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 109

An ecommerce company has tasked a Database Specialist with creating a reporting dashboard that visualizes critical business metrics that will be pulled from the core production database running on Amazon Aurora. Data that is read by the dashboard should be available within 100 milliseconds of an update.

The Database Specialist needs to review the current configuration of the Aurora DB cluster and develop a cost-effective solution. The solution needs to accommodate the unpredictable read workload from the

reporting dashboard without any impact on the write availability and performance of the DB cluster.

Which solution meets these requirements?

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

Answer: A

NEW QUESTION 113

A company uses the Amazon DynamoDB table contractDB in us-east-1 for its contract system with the following schema:

orderID (primary key) timestamp (sort key) contract (map) createdBy (string) customerEmail (string)

After a problem in production, the operations team has asked a database specialist to provide an IAM policy to read items from the database to debug the application. In addition, the developer is not allowed to access the value of the customerEmail field to stay compliant.

Which IAM policy should the database specialist use to achieve these requirements?

A)

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "IAMPolicy",
      "Effect": "Allow",
      "Action": [
        "dynamodb: Query"
      ],
      "Resource": [
        "arn:aws:dynamodb:us-east-1:123456789012:table/contractDB"
      ],
      "Condition": {
        "ForAllValues:StringLike": {
          "dynamodb:Attributes": [
            "orderID",
            "timestamp",
            "contract",
            "createdBy"
          ]
        },
        "StringEquals": {
          "dynamodb:Select": "SPECIFIC_ATTRIBUTES"
        }
      }
    }
  ]
}
```

B)

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "IAMPolicy",
      "Effect": "Allow",
      "Action": [
        "dynamodb: Query"
      ],
      "Resource": [
        "arn:aws:dynamodb:us-east-1:123456789012:table/contractDB"
      ],
      "Condition": {
        "ForAllValues:StringLike": {
          "dynamodb:Attributes": [
            "customerEmail"
          ]
        },
        "StringEquals": {
          "dynamodb:Select": "SPECIFIC_ATTRIBUTES"
        }
      }
    }
  ]
}
```

C)

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "IAMPolicy",
      "Effect": "Deny",
      "Action": [
        "dynamodb: Query"
      ],
      "Resource": [
        "arn:aws:dynamodb:us-east-1:123456789012:table/contractDB"
      ],
      "Condition": {
        "ForAllValues:StringLike": {
          "dynamodb:Attributes": [
            "customerEmail"
          ]
        },
        "StringEquals": {
          "dynamodb:Select": "SPECIFIC_ATTRIBUTES"
        }
      }
    }
  ]
}
```

D)

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "IAMPolicy",
      "Effect": "Deny",
      "Action": [
        "dynamodb: Query"
      ],
      "Resource": [
        "arn:aws:dynamodb:us-east-1:123456789012:table/contractDB"
      ],
      "Condition": {
        "ForAllValues:StringLike": {
          "dynamodb:Attributes": [
            "orderId",
            "timestamp",
            "contract",
            "createdBy"
          ]
        },
        "StringEquals": {
          "dynamodb:Select": "SPECIFIC_ATTRIBUTES"
        }
      }
    }
  ]
}
```

- A. Option A
- B. Option B
- C. Option C
- D. Option C

Answer: A

NEW QUESTION 116

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

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

Answer: AB

Explanation:

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

NEW QUESTION 120

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 125

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

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

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

Answer: C

NEW QUESTION 130

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

What might account for this? (Choose two.)

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

Answer: AD

Explanation:

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

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

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

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

NEW QUESTION 132

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

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

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

Answer: B

Explanation:

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

NEW QUESTION 134

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

Which solution should the database specialist recommend?

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

Answer: C

Explanation:

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

NEW QUESTION 139

A company has applications running on Amazon EC2 instances in a private subnet with no internet connectivity. The company deployed a new application that uses Amazon DynamoDB, but the application cannot connect to the DynamoDB tables. A developer already checked that all permissions are set correctly. What should a database specialist do to resolve this issue while minimizing access to external resources?

- A. Add a route to an internet gateway in the subnet's route table.
- B. Add a route to a NAT gateway in the subnet's route table.
- C. Assign a new security group to the EC2 instances with an outbound rule to ports 80 and 443.
- D. Create a VPC endpoint for DynamoDB and add a route to the endpoint in the subnet's route table.

Answer: D

Explanation:

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/vpc-endpoints-dynamodb.html>

NEW QUESTION 141

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 144

A company wants to migrate its on-premises MySQL databases to Amazon RDS for MySQL. To comply with the company's security policy, all databases must be encrypted at rest. RDS DB instance snapshots must also be shared across various accounts to provision testing and staging environments.

Which solution meets these requirements?

- A. Create an RDS for MySQL DB instance with an AWS Key Management Service (AWS KMS) customer managed CM
- B. Update the key policy to include the Amazon Resource Name (ARN) of the other AWS accounts as a principal, and then allow the kms:CreateGrant action.
- C. Create an RDS for MySQL DB instance with an AWS managed CM
- D. Create a new key policy to include the Amazon Resource Name (ARN) of the other AWS accounts as a principal, and then allow the kms:CreateGrant action.

- E. Create an RDS for MySQL DB instance with an AWS owned CM
- F. Create a new key policy to include the administrator user name of the other AWS accounts as a principal, and then allow the kms:CreateGrant action.
- G. Create an RDS for MySQL DB instance with an AWS CloudHSM ke
- H. Update the key policy to include the Amazon Resource Name (ARN) of the other AWS accounts as a principal, and then allow the kms:CreateGrant action.

Answer: A

Explanation:

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

NEW QUESTION 145

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

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

Answer: C

Explanation:

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

NEW QUESTION 146

A clothing company uses a custom ecommerce application and a PostgreSQL database to sell clothes to thousands of users from multiple countries. The company is migrating its application and database from its on- premises data center to the AWS Cloud. The company has selected Amazon EC2 for the application and Amazon RDS for PostgreSQL for the database. The company requires database passwords to be changed every 60 days. A Database Specialist needs to ensure that the credentials used by the web application to connect to the database are managed securely. Which approach should the Database Specialist take to securely manage the database credentials?

- A. Store the credentials in a text file in an Amazon S3 bucke
- B. Restrict permissions on the bucket to the IAM role associated with the instance profile onl
- C. Modify the application to download the text file and retrieve the credentials on start u
- D. Update the text file every 60 days.
- E. Configure IAM database authentication for the application to connect to the databas
- F. Create an IAM user and map it to a separate database user for each ecommerce use
- G. Require users to update their passwords every 60 days.
- H. Store the credentials in AWS Secrets Manage
- I. Restrict permissions on the secret to only the IAM role associated with the instance profil
- J. Modify the application to retrieve the credentials from Secrets Manager on start u
- K. Configure the rotation interval to 60 days.
- L. Store the credentials in an encrypted text file in the application AM
- M. Use AWS KMS to store the key for decrypting the text fil
- N. Modify the application to decrypt the text file and retrieve the credentials on start u
- O. Update the text file and publish a new AMI every 60 days.

Answer: C

NEW QUESTION 151

A large IT hardware manufacturing company wants to deploy a MySQL database solution in the AWS Cloud. The solution should quickly create copies of the company's production databases for test purposes. The solution must deploy the test databases in minutes, and the test data should match the latest production data as closely as possible. Developers must also be able to make changes in the test database and delete the instances afterward. Which solution meets these requirements?

- A. Leverage Amazon RDS for MySQL with write-enabled replicas running on Amazon EC2. Create the test copies using a mysqldump backup from the RDS for MySQL DB instances and importing them into the new EC2 instances.
- B. Leverage Amazon Aurora MySQL
- C. Use database cloning to create multiple test copies of the production DB clusters.
- D. Leverage Amazon Aurora MySQL
- E. Restore previous production DB instance snapshots into new test copies of Aurora MySQL DB clusters to allow them to make changes.
- F. Leverage Amazon RDS for MySQL
- G. Use database cloning to create multiple developer copies of the production DB instance.

Answer: B

NEW QUESTION 155

A huge gaming firm is developing a centralized method for storing the status of various online games' user sessions. The workload requires low-latency key-value storage and will consist of an equal number of reads and writes. Across the games' geographically dispersed user base, data should be written to the AWS Region nearest to the user. The design should reduce the burden associated with managing data replication across Regions. Which solution satisfies these criteria?

- A. Amazon RDS for MySQL with multi-Region read replicas
- B. Amazon Aurora global database
- C. Amazon RDS for Oracle with GoldenGate
- D. Amazon DynamoDB global tables

Answer: D

Explanation:

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

NEW QUESTION 159

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 metri
- 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 163

A business that specializes in internet advertising is developing an application that will show adverts to its customers. The program stores data in an Amazon DynamoDB database. Additionally, the application caches its reads using a DynamoDB Accelerator (DAX) cluster. The majority of reads come via the GetItem and BatchGetItem queries. The application does not need consistency of readings.

The application cache does not behave as intended after deployment. Specific extremely consistent queries to the DAX cluster are responding in several milliseconds rather than microseconds.

How can the business optimize cache behavior in order to boost application performance?

- A. Increase the size of the DAX cluster.
- B. Configure DAX to be an item cache with no query cache
- C. Use eventually consistent reads instead of strongly consistent reads.
- D. Create a new DAX cluster with a higher TTL for the item cache.

Answer: C

NEW QUESTION 168

A database specialist at a large multi-national financial company is in charge of designing the disaster recovery strategy for a highly available application that is in development. The application uses an Amazon DynamoDB table as its data store. The application requires a recovery time objective (RTO) of 1 minute and a recovery point objective (RPO) of 2 minutes.

Which operationally efficient disaster recovery strategy should the database specialist recommend for the DynamoDB table?

- A. Create a DynamoDB stream that is processed by an AWS Lambda function that copies the data to a DynamoDB table in another Region.
- B. Use a DynamoDB global table replica in another Regio
- C. Enable point-in-time recovery for both tables.
- D. Use a DynamoDB Accelerator table in another Regio
- E. Enable point-in-time recovery for the table.
- F. Create an AWS Backup plan and assign the DynamoDB table as a resource.

Answer: C

NEW QUESTION 172

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

This application has two parts:

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

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

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

Answer: D

NEW QUESTION 174

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 175

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 177

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

How should the company identify the source of the problem?

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

Answer: B

Explanation:

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

NEW QUESTION 180

A company has an on-premises system that tracks various database operations that occur over the lifetime of a database, including database shutdown, deletion, creation, and backup.

The company recently moved two databases to Amazon RDS and is looking at a solution that would satisfy these requirements. The data could be used by other systems within the company.

Which solution will meet these requirements with minimal effort?

- A. Create an Amazon Cloudwatch Events rule with the operations that need to be tracked on Amazon RDS
- B. Create an AWS Lambda function to act on these rules and write the output to the tracking systems.
- C. Create an AWS Lambda function to trigger on AWS CloudTrail API call
- D. Filter on specific RDS API calls and write the output to the tracking systems.
- E. Create RDS event subscription
- F. Have the tracking systems subscribe to specific RDS event system notifications.
- G. Write RDS logs to Amazon Kinesis Data Firehose
- H. Create an AWS Lambda function to act on these rules and write the output to the tracking systems.

Answer: C

NEW QUESTION 181

A company is using Amazon Aurora MySQL as the database for its retail application on AWS. The company receives a notification of a pending database upgrade and wants to ensure upgrades do not occur before or during the most critical time of year. Company leadership is concerned that an Amazon RDS maintenance window will cause an outage during data ingestion.

Which step can be taken to ensure that the application is not interrupted?

- A. Disable weekly maintenance on the DB cluster.

- B. Clone the DB cluster and migrate it to a new copy of the database.
- C. Choose to defer the upgrade and then find an appropriate down time for patching.
- D. Set up an Aurora Replica and promote it to primary at the time of patching.

Answer: C

NEW QUESTION 186

An electric utility company wants to store power plant sensor data in an Amazon DynamoDB table. The utility company has over 100 power plants and each power plant has over 200 sensors that send data every 2 seconds. The sensor data includes time with milliseconds precision, a value, and a fault attribute if the sensor is malfunctioning. Power plants are identified by a globally unique identifier. Sensors are identified by a unique identifier within each power plant. A database specialist needs to design the table to support an efficient method of finding all faulty sensors within a given power plant.

Which schema should the database specialist use when creating the DynamoDB table to achieve the fastest query time when looking for faulty sensors?

- A. Use the plant identifier as the partition key and the measurement time as the sort key
- B. Create a global secondary index (GSI) with the plant identifier as the partition key and the fault attribute as the sort key.
- C. Create a composite of the plant identifier and sensor identifier as the partition key
- D. Use the measurement time as the sort key
- E. Create a local secondary index (LSI) on the fault attribute.
- F. Create a composite of the plant identifier and sensor identifier as the partition key
- G. Use the measurement time as the sort key
- H. Create a global secondary index (GSI) with the plant identifier as the partition key and the fault attribute as the sort key.
- I. Use the plant identifier as the partition key and the sensor identifier as the sort key
- J. Create a local secondary index (LSI) on the fault attribute.

Answer: D

Explanation:

Plant id as partition key and Sensor id as a sort key. Fault can be identified quickly using the local secondary index and associated plant and sensor can be identified easily.

NEW QUESTION 188

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

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

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

Answer: AC

NEW QUESTION 190

A company developed a new application that is deployed on Amazon EC2 instances behind an Application Load Balancer. The EC2 instances use the security group named sg-application-servers. The company needs a database to store the data from the application and decides to use an Amazon RDS for MySQL DB instance. The DB instance is deployed in private DB subnet.

What is the MOST restrictive configuration for the DB instance security group?

- A. Only allow incoming traffic from the sg-application-servers security group on port 3306.
- B. Only allow incoming traffic from the sg-application-servers security group on port 443.
- C. Only allow incoming traffic from the subnet of the application servers on port 3306.
- D. Only allow incoming traffic from the subnet of the application servers on port 443.

Answer: A

Explanation:

most restrictive approach is to allow only incoming connections from SG of EC2 instance on port 3306

NEW QUESTION 192

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

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

Answer: A

Explanation:

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

NEW QUESTION 196

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 198

A business is transferring its on-premises database workloads to the Amazon Web Services (AWS) Cloud. A database professional migrating an Oracle database with a huge table to Amazon RDS has picked AWS DMS. The database professional observes that AWS DMS is consuming considerable time migrating the data.

Which activities would increase the pace of data migration? (Select three.)

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

Answer: ACD

Explanation:

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

NEW QUESTION 202

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 207

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 snapshot
- 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 208

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 213

A gaming company uses Amazon Aurora Serverless for one of its internal applications. The company's developers use Amazon RDS Data API to work with the Aurora Serverless DB cluster. After a recent security review, the company is mandating security enhancements. A database specialist must ensure that access to RDS Data API is private and never passes through the public internet. What should the database specialist do to meet this requirement?

- A. Modify the Aurora Serverless cluster by selecting a VPC with private subnets.
- B. Modify the Aurora Serverless cluster by unchecking the publicly accessible option.
- C. Create an interface VPC endpoint that uses AWS PrivateLink for RDS Data API.
- D. Create a gateway VPC endpoint for RDS Data API.

Answer: C

Explanation:

<https://aws.amazon.com/about-aws/whats-new/2020/02/amazon-rds-data-api-now-supports-aws-privatelink/>

NEW QUESTION 214

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

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

Which solution will meet these requirements?

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

Answer: C

Explanation:

Step 1: Download S3 Files

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

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

NEW QUESTION 217

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

What should a database specialist do so that point-in-time recovery can be successful?

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

Answer: C

Explanation:

You can restore a DB instance to a specific point in time (PITR), creating a new DB instance. To support PITR, your DB instances must have backup retention set to a nonzero value. <https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/custom-backup-sqlserver.html>

<https://aws.amazon.com/blogs/database/setting-up-a-binlog-server-for-amazon-rds-mysql-and-mariadb-using-m> "After you run the command, it's okay to enable backup retention on the RDS instance by using the AWS CLI or the console. Enabling backup retention also enables binary logging."

<https://aws.amazon.com/blogs/storage/point-in-time-recovery-and-continuous-backup-for-amazon-rds-with-aws>

NEW QUESTION 218

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

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

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

- A. Deploy an Amazon RDS DB instance in us-east-1 with a read replica instance in eu-west-2. Create an Amazon ElastiCache cluster in eu-west-2 to cache data from the read replica to generate the dashboards.
- B. Use an Amazon DynamoDB global table in us-east-1 with replication into eu-west-2. Use multi-active replication to ensure that updates are quickly propagated

to eu-west-2.

C. Use an Amazon Aurora global database.

D. Deploy the primary DB cluster in us-east-1. Deploy the secondary DB cluster in eu-west-2. Configure the dashboard application to read from the secondary cluster.

E. Deploy an Amazon RDS for MySQL DB instance in us-east-1 with a read replica instance in eu-west-2. Configure the dashboard application to read from the read replica.

Answer: C

Explanation:

Amazon Aurora global databases span multiple AWS Regions, enabling low latency global reads and providing fast recovery from the rare outage that might affect an entire AWS Region. An Aurora global database has a primary DB cluster in one Region, and up to five secondary DB clusters in different Regions.

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

NEW QUESTION 219

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

A. Create an AWS Lambda function to run on the first day of every month to take a manual RDS snapshot. Move the snapshot to the company's Amazon S3 bucket.

B. Create an AWS Lambda function to run on the first day of every month to take a manual RDS snapshot.

C. Create an RDS snapshot schedule from the AWS Management Console to take a snapshot every 30 days.

D. Create an AWS Lambda function to run on the first day of every month to create an automated RDS snapshot.

Answer: A

Explanation:

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

NEW QUESTION 224

A database specialist is managing an application in the us-west-1 Region and wants to set up disaster recovery in the us-east-1 Region. The Amazon Aurora MySQL DB cluster needs an RPO of 1 minute and an RTO of 2 minutes.

Which approach meets these requirements with no negative performance impact?

A. Enable synchronous replication.

B. Enable asynchronous binlog replication.

C. Create an Aurora Global Database.

D. Copy Aurora incremental snapshots to the us-east-1 Region.

Answer: C

Explanation:

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/aurora-global-database-disaster-recovery.html>

NEW QUESTION 228

In one AWS account, a business runs a two-tier ecommerce application. An Amazon RDS for MySQL

Multi-AZ database instance serves as the application's backend. A developer removed the database instance in the production environment by accident. Although the organization recovers the database, the incident results in hours of outage and financial loss.

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

A. Grant least privilege to groups, IAM users, and roles.

B. Allow all users to restore a database from a backup.

C. Enable deletion protection on existing production DB instances.

D. Use an ACL policy to restrict users from DB instance deletion.

E. Enable AWS CloudTrail logging and Enhanced Monitoring.

Answer: ACD

NEW QUESTION 230

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

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

A. Create an RDS event subscription to the audit event type.

B. Enable auditing of CONNECT and QUERY_DML events.

C. SSH to the DB instance and review the database logs.

D. Publish the database logs to Amazon CloudWatch Logs.

E. Enable Enhanced Monitoring on the DB instance.

Answer: BD

Explanation:

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

NEW QUESTION 231

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 instance
- E. Use Amazon Route 53 DNS with failover health checks configured.

Answer: B

NEW QUESTION 232

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 234

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

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

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

Answer: C

Explanation:

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

PDF

Kindle RSS

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

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

NEW QUESTION 239

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 241

A business is launching a new Amazon RDS for SQL Server database instance. The organization wishes to allow auditing of the SQL Server database. Which measures should a database professional perform in combination to achieve this requirement? (Select two.)

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

G. Set up an options group to configure an IAM role and an Amazon S3 bucket for audit log storage. Associate the options group with the DB instance.

Answer: AE

Explanation:

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

NEW QUESTION 243

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

Which solution meets these requirements?

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

Answer: A

Explanation:

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

NEW QUESTION 245

A company has a database monitoring solution that uses Amazon CloudWatch for its Amazon RDS for SQL Server environment. The cause of a recent spike in CPU utilization was not determined using the standard metrics that were collected. The CPU spike caused the application to perform poorly, impacting users. A Database Specialist needs to determine what caused the CPU spike.

Which combination of steps should be taken to provide more visibility into the processes and queries running during an increase in CPU load? (Choose two.)

- A. Enable Amazon CloudWatch Events and view the incoming T-SQL statements causing the CPU to spike.
- B. Enable Enhanced Monitoring metrics to view CPU utilization at the RDS SQL Server DB instance level.
- C. Implement a caching layer to help with repeated queries on the RDS SQL Server DB instance.
- D. Use Amazon QuickSight to view the SQL statement being run.
- E. Enable Amazon RDS Performance Insights to view the database load and filter the load by waits, SQL statements, hosts, or users.

Answer: BE

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/rds-instance-high-cpu/> "Several factors can cause an increase in CPU utilization. For example, user-initiated heavy workloads, analytic queries, prolonged deadlocks and lock waits, multiple concurrent transactions, long-running transactions, or other processes that utilize CPU resources. First, you can identify the source of the CPU usage by: Using Enhanced Monitoring Using Performance Insights"

NEW QUESTION 246

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 248

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 252

The Development team recently executed a database script containing several data definition language (DDL) and data manipulation language (DML) statements on an Amazon Aurora MySQL DB cluster. The release accidentally deleted thousands of rows from an important table and broke some application functionality. This was discovered 4 hours after the release. Upon investigation, a Database Specialist tracked the issue to a DELETE command in the script with an incorrect WHERE clause filtering the wrong set of rows. The Aurora DB cluster has Backtrack enabled with an 8-hour backtrack window. The Database Administrator also took a manual snapshot of the DB cluster before the release started. The database needs to be returned to the correct state as quickly as possible to resume full application functionality. Data loss must be minimal. How can the Database Specialist accomplish this?

- A. Quickly rewind the DB cluster to a point in time before the release using Backtrack.
- B. Perform a point-in-time recovery (PITR) of the DB cluster to a time before the release and copy the deleted rows from the restored database to the original database.
- C. Restore the DB cluster using the manual backup snapshot created before the release and change the application configuration settings to point to the new DB cluster.
- D. Create a clone of the DB cluster with Backtrack enabled
- E. Rewind the cloned cluster to a point in time before the release
- F. Copy deleted rows from the clone to the original database.

Answer: A

NEW QUESTION 256

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 258

A company recently acquired a new business. A database specialist must migrate an unencrypted 12 TB Amazon RDS for MySQL DB instance to a new AWS account. The database specialist needs to minimize the amount of time required to migrate the database. Which solution meets these requirements?

- A. Create a snapshot of the source DB instance in the source account
- B. Share the snapshot with the destination account
- C. In the target account, create a DB instance from the snapshot.
- D. Use AWS Resource Access Manager to share the source DB instance with the destination account. Create a DB instance in the destination account using the shared resource.
- E. Create a read replica of the DB instance
- F. Give the destination account access to the read replica
- G. In the destination account, create a snapshot of the shared read replica and provision a new RDS for MySQL DB instance.
- H. Use mysqldump to back up the source database
- I. Create an RDS for MySQL DB instance in the destination account
- J. Use the mysql command to restore the backup in the destination database.

Answer: A

Explanation:

Sharing an unencrypted manual DB snapshot enables authorized AWS accounts to directly restore a DB instance from the snapshot instead of taking a copy of it and restoring from that. https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_ShareSnapshot.html However Resource Access Manager could not share non-Aurora cluster. <https://docs.aws.amazon.com/ram/latest/userguide/shareable.html>

NEW QUESTION 259

A gaming company has implemented a leaderboard in AWS using a Sorted Set data structure within Amazon ElastiCache for Redis. The ElastiCache cluster has been deployed with cluster mode disabled and has a replication group deployed with two additional replicas. The company is planning for a worldwide gaming event and is anticipating a higher write load than what the current cluster can handle.

Which method should a Database Specialist use to scale the ElastiCache cluster ahead of the upcoming event?

- A. Enable cluster mode on the existing ElastiCache cluster and configure separate shards for the Sorted Set across all nodes in the cluster.
- B. Increase the size of the ElastiCache cluster nodes to a larger instance size.
- C. Create an additional ElastiCache cluster and load-balance traffic between the two clusters.
- D. Use the EXPIRE command and set a higher time to live (TTL) after each call to increment a given key.

Answer: B

NEW QUESTION 261

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 264

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 269

A company uses an on-premises Microsoft SQL Server database to host relational and JSON data and to run daily ETL and advanced analytics. The company wants to migrate the database to the AWS Cloud. Database specialist must choose one or more AWS services to run the company's workloads.

Which solution will meet these requirements in the MOST operationally efficient manner?

- A. Use Amazon Redshift for relational data
- B. Use Amazon DynamoDB for JSON data
- C. Use Amazon Redshift for relational data and JSON data.
- D. Use Amazon RDS for relational data
- E. Use Amazon Neptune for JSON data
- F. Use Amazon Redshift for relational data
- G. Use Amazon S3 for JSON data.

Answer: B

Explanation:

<https://docs.aws.amazon.com/redshift/latest/dg/super-overview.htm>

NEW QUESTION 273

A company has a production Amazon Aurora DB cluster that serves both online transaction processing (OLTP) transactions and compute-intensive reports. The reports run for 10% of the total cluster uptime while the OLTP transactions run all the time. The company has benchmarked its workload and determined that a six-node Aurora DB cluster is appropriate for the peak workload.

The company is now looking at cutting costs for this DB cluster, but needs to have a sufficient number of nodes in the cluster to support the workload at different times. The workload has not changed since the previous benchmarking exercise.

How can a Database Specialist address these requirements with minimal user involvement?

- A. Split up the DB cluster into two different clusters: one for OLTP and the other for reporting
- B. Monitor and set up replication between the two clusters to keep data consistent.
- C. Review and evaluate the peak combined workload
- D. Ensure that utilization of the DB cluster node is at an acceptable level
- E. Adjust the number of instances, if necessary.
- F. Use the stop cluster functionality to stop all the nodes of the DB cluster during times of minimal workload
- G. The cluster can be restarted again depending on the workload at the time.
- H. Set up automatic scaling on the DB cluster
- I. This will allow the number of reader nodes to adjust automatically to the reporting workload, when needed.

Answer: D

NEW QUESTION 276

.....

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](#)