



DBS-C01

Certified Database Specialty



EXAMKILLER

Help Pass Your Exam At First Try

Amazon Web Services

Exam DBS-C01

AWS Certified Database - Specialty

Version: 6.0

[Total Questions: 270]

Question No : 1

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-start-formatting>

Question No : 2

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

Which strategy should the database administrator take?

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

Answer: A

Explanation:

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

Question No : 3

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

Question No : 4

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

Which solution satisfies these criteria?

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

Answer: C

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

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

Question No : 5

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.html

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.

Question No : 6

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-accounts-and-regions/

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/stacksets-concepts.html>

Question No : 7

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

Question No : 8

A company is using Amazon RDS for MySQL to redesign its business application. A Database Specialist has noticed that the Development team is restoring their MySQL

database multiple times a day when Developers make mistakes in their schema updates. The Developers sometimes need to wait hours to the restores to complete.

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

Which approach should the Database Specialist take to reduce downtime?

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

Answer: C

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

Question No : 9

A database specialist manages a critical Amazon RDS for MySQL DB instance for a company. The data stored daily could vary from .01% to 10% of the current database size. The database specialist needs to ensure that the DB instance storage grows as needed.

What is the MOST operationally efficient and cost-effective solution?

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

Answer: A

Explanation: If your workload is unpredictable, you can enable storage autoscaling for an Amazon RDS DB instance. With storage autoscaling enabled, when Amazon RDS detects that you are running out of free database space it automatically scales up your storage.
<https://aws.amazon.com/about-aws/whats-new/2019/06/rds-storage-auto-scaling/>
https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_PIOPS.StorageTypes.html#USER_PIOPS.Autoscaling

Question No : 10

A vehicle insurance company needs to choose a highly available database to track vehicle owners and their insurance details. The persisted data should be immutable in the database, including the complete and sequenced history of changes over time with all the owners and insurance transfer details for a vehicle.

The data should be easily verifiable for the data lineage of an insurance claim.

Which approach meets these requirements with MINIMAL effort?

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

Answer: C

Question No : 11

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. Use the mysqldump utility to create an initial snapshot from the on-premises MySQL server, and copy it to an Amazon S3 bucket. Import the snapshot into the DB instance utilizing the MySQL utilities running on an Amazon EC2 instance. Immediately point the application to the DB instance.
- B.** 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. Use the mysqldump utility to create a snapshot of the on-premises MySQL server. Copy the snapshot into the EC2 instance and restore it into the EC2 MySQL instance. Use AWS DMS to migrate data into a new RDS for MySQL DB instance. Point the application to the DB instance.
- C.** 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. Use the mysqldump utility to create a snapshot of the on-premises MySQL server. 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. Point the application to the DB instance.
- D.** Deploy a new RDS for MySQL DB instance and configure it for access from the on-premises data center. Use the mysqldump utility to create an initial snapshot from the on-premises MySQL server, and copy it to an Amazon S3 bucket. Import the snapshot into the DB instance using the MySQL utilities running on an Amazon EC2 instance. Establish replication into the new DB instance using MySQL replication. Stop application access to the on-premises MySQL server and let the remaining transactions replicate over. Point the application to the DB instance.

Answer: B

Question No : 12

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

Question No : 13

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

What should the team do to meet this requirement?

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

Answer: C

Question No : 14

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.

Question No : 15

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. Share the snapshot with the destination account. In the target account, create a DB instance from the snapshot.
- B. 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.
- C. Create a read replica of the DB instance. Give the destination account access to the read replica. In the destination account, create a snapshot of the shared read replica and provision a new RDS for MySQL DB instance.
- D. Use mysqldump to back up the source database. Create an RDS for MySQL DB instance in the destination account. 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>

Question No : 16

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-snapshot-copies-with-the-snapshot-tool-for-amazon-aurora/>

Question No : 17

A Database Specialist is performing a proof of concept with Amazon Aurora using a small instance to confirm a simple database behavior. When loading a large dataset and creating the index, the Database Specialist encounters the following error message from Aurora:

ERROR: cloud not write block 7507718 of temporary file: No space left on device

What is the cause of this error and what should the Database Specialist do to resolve this issue?

- A.** The scaling of Aurora storage cannot catch up with the data loading. The Database Specialist needs to modify the workload to load the data slowly.
- B.** The scaling of Aurora storage cannot catch up with the data loading. The Database Specialist needs to enable Aurora storage scaling.
- C.** The local storage used to store temporary tables is full. The Database Specialist needs to scale up the instance.
- D.** The local storage used to store temporary tables is full. The Database Specialist needs to enable local storage scaling.

Answer: C

Reference: <https://serverfault.com/QUESTION NO:s/109828/how-can-i-tune-postgres-to-avoid-this-error>

Question No : 18

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

Which process should the database specialist recommend?

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

Answer: C

Explanation:

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

Question No : 19

A company stores critical data for a department in Amazon RDS for MySQL DB instances. The department was closed for 3 weeks and notified a database specialist that access to the RDS DB instances should not be granted to anyone during this time. To meet this requirement, the database specialist stopped all the

DB instances used by the department but did not select the option to create a snapshot. Before the 3 weeks expired, the database specialist discovered that users could connect to the database successfully.

What could be the reason for this?

- A.** When stopping the DB instance, the option to create a snapshot should have been selected.
- B.** When stopping the DB instance, the duration for stopping the DB instance should have been selected.
- C.** Stopped DB instances will automatically restart if the number of attempted connections exceeds the threshold set.
- D.** Stopped DB instances will automatically restart if the instance is not manually started after 7 days.

Answer: D

Explanation: <https://aws.amazon.com/premiumsupport/knowledge-center/rds-stop-seven-days/>

Question No : 20

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. Recreate the DMS replication instance with the publicly accessible option disabled.
- D. Create a new replication VPC subnet group with private subnets. 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/>

Question No : 21

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

How can this solution be implemented?

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

Answer: A

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

Question No : 22

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

How could a database professional approach these needs while using DynamoDB?

- A.** Enable Amazon DynamoDB Streams on the table. Create an AWS Lambda function triggered by the stream. Write the log entries to an Amazon S3 object.
- B.** Create a backup plan in AWS Backup to back up the DynamoDB table once a day. Create an AWS Lambda function that restores the backup in another table and compares both tables for changes. Generate the log entries and write them to an Amazon S3 object.
- C.** Enable AWS CloudTrail logs on the table. Create an AWS Lambda function that reads the log files once an hour and filters DynamoDB API actions. Write the filtered log files to Amazon S3.
- D.** Enable Amazon DynamoDB Streams on the table. Create an AWS Lambda function triggered by the stream. Write the log entries to an Amazon Kinesis Data Firehose delivery stream with buffering and Amazon S3 as the destination.

Answer: D

Question No : 23

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

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

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

24 hours.

Answer: C

Explanation:

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/CreateBackup.html#:~:text=If%20you%20don%27t%20want%20to%20create%20scheduling%20scripts%20and%20cleanup%20jobs%2C%20you%20can%20use%20AWS%20Backup%20to%20create%20backup%20plans%20with%20schedules%20and%20retention%20policies%20for%20your%20DynamoDB%20tables.%20AWS%20Backup%20runs%20the%20backups%20and%20deletes%20them%20when%20they%20expire.%20For%20more%20information%2C%20see%20the%20AWS%20Backup%20Developer%20Guide.>

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

Question No : 24

A business just transitioned from an on-premises Oracle database to Amazon Aurora PostgreSQL. Following the move, the organization observed that every day around 3:00 PM, the application's response time is substantially slower. The firm has determined that the problem is with the database, not the application.

Which set of procedures should the Database Specialist do to locate the erroneous PostgreSQL query most efficiently?

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

Answer: D

Explanation:

<https://aws.amazon.com/blogs/database/optimizing-and-tuning-queries-in-amazon-rds-postgresql-based-on-native-and-external-tools/>

"AWS recently released a feature called Amazon RDS Performance Insights, which provides an easy-to-understand dashboard for detecting performance problems in terms of load." "AWS recently released a feature called Amazon RDS Performance Insights, which provides an easy-to-understand dashboard for detecting performance problems in terms of load."

Question No : 25

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 mode. Connect the Amazon Kinesis data stream to Kinesis Data Firehose. Set the Kinesis Data Firehose destination to an Amazon S3 bucket.
- B.** Create an AWS CloudTrail trail in the Region where the database runs. Associate the database activity logs with the trail.
- C.** Enable Aurora Database Activity Streams on the database in asynchronous mode. Connect the Amazon Kinesis data stream to Kinesis Data Firehose. Set the Firehose destination to an Amazon S3 bucket.
- D.** Allow connections to the DB cluster through a bastion host only. Restrict database access to the bastion host and application servers. 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>

Question No : 26

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: B,D

Reference: <https://aws.amazon.com/premiumsupport/knowledge-center/s3-could-not-connect-endpoint-url/>

Question No : 27

The website of a manufacturing firm makes use of an Amazon Aurora PostgreSQL database cluster.

Which settings will result in the LEAST amount of downtime for the application during failover? (Select 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: A,C,E

Explanation:

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraPostgreSQL.Be>