

ITPassLeader



Pass Your Next Certification Exam Fast!

Select a vendor... Select an test... Your email address [Free Download Demo](#)



Instant Download



365 Days Free Updates



Money Back Guarantee



Security & Privacy

Choose the version that fits your needs

PDF Version

Desktop Test Engine

Online Test Engine

Latest and Up-to-Date exam dumps with real exam questions answers.



Get 12-Months free updates without any extra charges.



Experience same exam environment before appearing in the certification exam.



100% exam passing guarantee in the first attempt.



20% discount on more than one license and 30% discount on 5+ license purchases.



100% secure purchase on SSL.



Completely private purchase without sharing your personal info with anyone.



<http://www.itpassleader.com>

High-praise Exam Dumps Questions grant you success by high pass rate - ITPassLeader

Exam : **DBS-C01**

Title : AWS Certified Database -
Specialty (DBS-C01) Exam

Vendor : Amazon

Version : DEMO

NO.1 A database specialist needs to move a table from a database that is running on an Amazon Aurora PostgreSQL DB cluster into a new and distinct database cluster. The new table in the new database must be updated with any changes to the original table that happen while the migration is in progress.

The original table contains a column to store data as large as 2 GB in the form of large binary objects (LOBs).

A few records are large in size, but most of the LOB data is smaller than 32 KB.

What is the FASTEST way to replicate all the data from the original table?

- A.** Use AWS Database Migration Service (AWS DMS) with ongoing replication in full LOB mode.
- B.** Take a snapshot of the database. Create a new DB instance by using the snapshot.
- C.** Use AWS Database Migration Service (AWS DMS) with ongoing replication in limited LOB mode.
- D.** Use AWS Database Migration Service (AWS DMS) with ongoing replication in inline LOB mode.

Answer: D

Explanation

The fastest way to replicate all the data from the original table is to use AWS Database Migration Service (AWS DMS) with ongoing replication in inline LOB mode. This option allows AWS DMS to transfer LOBs smaller than the specified size inline, which means that they are fetched from the database in bulk as VARCHAR data types. This is significantly faster than other methods that require piece-by-piece migration of LOBs. LOBs larger than the specified size are replicated using full LOB mode, which means that they are migrated one at a time. This option is suitable for the scenario where most of the LOB data is smaller than 32 KB, and only a few records are large in size. Using inline LOB mode can also avoid data truncation that might occur in limited LOB mode, where LOBs that exceed the maximum LOB size are cut off and a warning is issued to the log file. Full LOB mode can be quite slow, as it migrates all LOBs regardless of size, and it does not pre-allocate memory or load the LOB data in bulk. Taking a snapshot of the database and creating a new DB instance by using the snapshot is not a valid option, as it does not meet the requirement of updating the new table with any changes to the original table that happen while the migration is in progress

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

NO.2 A company has an AWS CloudFormation template written in JSON that is used to launch new Amazon RDS for MySQL DB instances. The security team has asked a database specialist to ensure that the master password is automatically rotated every 30 days for all new DB instances that are launched using the template.

What is the MOST operationally efficient solution to meet these requirements?

- A.** Save the password in an Amazon S3 object. Encrypt the S3 object with an AWS KMS key. Set the KMS key to be rotated every 30 days by setting the EnableKeyRotation property to true. Use a CloudFormation custom resource to read the S3 object to extract the password.
- B.** Create an AWS Lambda function to rotate the secret. Modify the CloudFormation template to add an AWS::SecretsManager::RotationSchedule resource. Configure the RotationLambdaARN value and, for the RotationRules property, set the AutomaticallyAfterDays parameter to 30.
- C.** Modify the CloudFormation template to use the AWS KMS key as the database password. Configure an Amazon EventBridge rule to invoke the KMS API to rotate the key every 30 days by setting the ScheduleExpression parameter to `*/30*`.
- D.** Integrate the Amazon RDS for MySQL DB instances with AWS IAM and centrally manage the master database user password.

Answer: B

Explanation

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-resource-secretsmanager-rotationsche>

NO.3 A database specialist needs to move an Amazon RDS DB instance from one AWS account to another AWS account.

Which solution will meet this requirement with the LEAST operational effort?

- A.** Use AWS Database Migration Service (AWS DMS) to migrate the DB instance from the source AWS account to the destination AWS account.
- B.** Create a DB snapshot of the DB instance. Share the snapshot with the destination AWS account. Create a new DB instance by restoring the snapshot in the destination AWS account.
- C.** Create a Multi-AZ deployment for the DB instance. Create a read replica for the DB instance in the source AWS account. Use the read replica to replicate the data into the DB instance in the destination AWS account.
- D.** Use AWS DataSync to back up the DB instance in the source AWS account. Use AWS Resource Access Manager (AWS RAM) to restore the backup in the destination AWS account.

Answer: B

Explanation

The solution that will meet the requirement with the least operational effort is to create a DB snapshot of the DB instance, share the snapshot with the destination AWS account, and create a new DB instance by restoring the snapshot in the destination AWS account. This solution is simple, secure, and does not require any additional services or tools. It also preserves the configuration and data of the original DB instance. The other solutions are either more complex, require additional services or tools, or do not guarantee data consistency or security. For example, using AWS DMS to migrate the DB instance requires setting up replication endpoints, tasks, and network connectivity between the source and destination accounts. Creating a Multi-AZ deployment and a read replica for the DB instance involves additional costs and complexity, and does not ensure data consistency between the source and destination DB instances. Using AWS DataSync and AWS RAM to back up and restore the DB instance requires additional storage and permissions, and does not support encryption or compression for the backup data

NO.4 A corporation intends to migrate a 500-GB Oracle database to Amazon Aurora PostgreSQL utilizing the AWS Schema Conversion Tool (AWS SCT) and AWS Data Management Service (AWS DMS). The database does not have any stored procedures, but does contain several huge or partitioned tables. Because the program is vital to the company, it is preferable to migrate with little downtime.

Which measures should a database professional perform in combination to expedite the transfer process?

(Select three.)

- A.** Use the AWS SCT data extraction agent to migrate the schema from Oracle to Aurora PostgreSQL.
- B.** For the large tables, change the setting for the maximum number of tables to load in parallel and perform a full load using AWS DMS.
- C.** For the large tables, create a table settings rule with a parallel load option in AWS DMS, then perform a full load using DMS.

- D. Use AWS DMS to set up change data capture (CDC) for continuous replication until the cutover date.
- E. Use AWS SCT to convert the schema from Oracle to Aurora PostgreSQL.
- F. Use AWS DMS to convert the schema from Oracle to Aurora PostgreSQL and for continuous replication.

Answer: C D E

NO.5 A healthcare company is running an application on Amazon EC2 in a public subnet and using Amazon DocumentDB (with MongoDB compatibility) as the storage layer. An audit reveals that the traffic between the application and Amazon DocumentDB is not encrypted and that the DocumentDB cluster is not encrypted at rest. A database specialist must correct these issues and ensure that the data in transit and the data at rest are encrypted.

Which actions should the database specialist take to meet these requirements? (Select TWO.)

- A. Download the SSH RSA public key for Amazon DocumentDB. Update the application configuration to use the instance endpoint instead of the cluster endpoint and run queries over SSH.
- B. Download the SSL .pem public key for Amazon DocumentDB. Add the key to the application package and make sure the application is using the key while connecting to the cluster.
- C. Create a snapshot of the unencrypted cluster. Restore the unencrypted snapshot as a new cluster with the `-storage-encrypted` parameter set to true. Update the application to point to the new cluster.
- D. Create an Amazon DocumentDB VPC endpoint to prevent the traffic from going to the Amazon DocumentDB public endpoint. Set a VPC endpoint policy to allow only the application instance's security group to connect.
- E. Activate encryption at rest using the `modify-db-cluster` command with the `-storage-encrypted` parameter set to true. Set the security group of the cluster to allow only the application instance's security group to connect.

Answer: B C

NO.6 A financial services company is developing a shared data service that supports different applications from throughout the company. A Database Specialist designed a solution to leverage Amazon ElastiCache for Redis with cluster mode enabled to enhance performance and scalability. The cluster is configured to listen on port 6379.

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

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

Answer: A C F

Explanation

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

NO.7 A company runs online transaction processing (OLTP) workloads on an Amazon RDS for PostgreSQL Multi-AZ DB instance. Tests were run on the database after work hours, which generated additional database logs.

The free storage of the RDS DB instance is low due to these additional logs.

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

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

Answer: B

Explanation

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

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

NO.8 A company wants to automate the creation of secure test databases with random credentials to be stored safely for later use. The credentials should have sufficient information about each test database to initiate a connection and perform automated credential rotations. The credentials should not be logged or stored anywhere in an unencrypted form.

Which steps should a Database Specialist take to meet these requirements using an AWS CloudFormation template?

- A.** Create the database with the `MasterUserName` and `MasterUserPassword` properties set to the default values. Then, create the secret with the user name and password set to the same default values. Add a Secret Target Attachment resource with the `SecretId` and `TargetId` properties set to the Amazon Resource Names (ARNs) of the secret and the database. Finally, update the secret's password value with a randomly generated string set by the `GenerateSecretString` property.
- B.** Add a Mapping property from the database Amazon Resource Name (ARN) to the secret ARN. Then, create the secret with a chosen user name and a randomly generated password set by the `GenerateSecretString` property. Add the database with the `MasterUserName` and `MasterUserPassword` properties set to the user name of the secret.
- C.** Add a resource of type `AWS::SecretsManager::Secret` and specify the `GenerateSecretString` property. Then, define the database user name in the `SecureStringTemplate` template. Create a resource for the database and reference the secret string for the `MasterUserName` and `MasterUserPassword` properties. Then, add a resource of type `AWS::SecretsManagerSecretTargetAttachment` with the `SecretId` and `TargetId` properties set to the Amazon Resource Names (ARNs) of the secret and the database.

D. Create the secret with a chosen user name and a randomly generated password set by the `GenerateSecretString` property. Add an `SecretTargetAttachment` resource with the `SecretId` property set to the Amazon Resource Name (ARN) of the secret and the `TargetId` property set to a parameter value matching the desired database ARN. Then, create a database with the `MasterUserName` and `MasterUserPassword` properties set to the previously created values in the secret.

Answer: C

NO.9 In North America, a business launched a mobile game that swiftly expanded to 10 million daily active players.

The game's backend is hosted on AWS and makes considerable use of a TTL-configured Amazon DynamoDB table.

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

How should a database administrator resolve this issue?

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

Answer: B

Explanation

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

NO.10 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.

NO.11 A Database Specialist has migrated an on-premises Oracle database to Amazon Aurora PostgreSQL. The schema and the data have been migrated successfully. The on-premises database server was also being used to run database maintenance cron jobs written in Python to perform tasks including data purging and generating data exports. The logs for these jobs show that, most of the time, the jobs completed within 5 minutes, but a few jobs took up to 10 minutes to complete. These maintenance jobs need to be set up for Aurora PostgreSQL.

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

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

Answer: C

Explanation

<https://docs.aws.amazon.com/AmazonCloudWatch/latest/events/Create-CloudWatch-Events-Scheduled-Rule.htm>

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

<https://docs.aws.amazon.com/AmazonCloudWatch/latest/events/RunLambdaSchedule.html>

NO.12 A development team at an international gaming company is experimenting with Amazon DynamoDB to store in-game events for three mobile games. The most popular game hosts a maximum of 500,000 concurrent users, and the least popular game hosts a maximum of 10,000 concurrent users. The average size of an event is

20 KB, and the average user session produces one event each second. Each event is tagged with a time in milliseconds and a globally unique identifier.

The lead developer created a single DynamoDB table for the events with the following schema:

Partition key: game name

Sort key: event identifier

Local secondary index: player identifier

Event time

The tests were successful in a small-scale development environment. However, when deployed to production, new events stopped being added to the table and the logs show DynamoDB failures with the `ItemCollectionSizeLimitExceededException` error code.

Which design change should a database specialist recommend to the development team?

- A.** Use the player identifier as the partition key. Use the event time as the sort key. Add a global secondary index with the game name as the partition key and the event time as the sort key.
- B.** Create two tables. Use the game name as the partition key in both tables. Use the event time as the sort key for the first table. Use the player identifier as the sort key for the second table.
- C.** Replace the sort key with a compound value consisting of the player identifier collated with the event time, separated by a dash. Add a local secondary index with the player identifier as the sort key.
- D.** Create one table for each game. Use the player identifier as the partition key. Use the event time

as the sort key.

Answer: D

Explanation

The correct answer is D. Create one table for each game. Use the player identifier as the partition key. Use the event time as the sort key.

The explanation is as follows:

The `ItemCollectionSizeLimitExceededException` error occurs when an item collection exceeds the 10 GB limit¹. An item collection is a group of items that have the same partition key value but different sort key values². In this case, the item collection is based on the game name, which has only three possible values. This means that all events for each game are stored in the same item collection, which can easily exceed the 10 GB limit given the high volume and size of events.

To avoid this error, a database specialist should recommend a design change that distributes the events across more partitions and reduces the size of each item collection. Option D achieves this by creating one table for each game, and using the player identifier as the partition key. This way, each event is stored in a separate partition based on the player identifier, and sorted by the event time. This design also supports efficient queries by game, player, and time range.

Option A is incorrect because it still uses a single table for all events, which can cause hot partitions and throttling due to uneven access patterns across games. Also, using the player identifier as the partition key can result in many small partitions that are underutilized and waste provisioned capacity. Adding a global secondary index with the game name as the partition key and the event time as the sort key does not solve the problem of item collection size limit, because global secondary indexes have their own item collections that are subject to the same limit³.

Option B is incorrect because it creates two tables with redundant data and increases storage costs. Also, using the game name as the partition key in both tables does not solve the problem of item collection size limit, as explained above.

Option C is incorrect because it still uses a single table for all events, which can cause hot partitions and throttling due to uneven access patterns across games. Also, replacing the sort key with a compound value consisting of the player identifier collated with the event time does not reduce the size of each item collection, because each event still has a unique sort key value. Adding a local secondary index with the player identifier as the sort key does not solve the problem of item collection size limit, because local secondary indexes share the same item collections as their base table⁴.

References: 1: `ItemCollectionSizeLimitExceededException` (AWS SDK for Java - 1.12.547) 2: Best practices for designing and using partition keys effectively - Amazon DynamoDB 3: Global Secondary Indexes - Amazon DynamoDB 4: Local Secondary Indexes - Amazon DynamoDB

NO.13 A marketing company is developing an application to track responses to email message campaigns. The company needs a database storage solution that is optimized to work with highly connected data. The database needs to limit connections and programmatic access to the data by using IAM policies.

Which solution will meet these requirements?

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

Answer: D

Explanation

Amazon Neptune is a fast, reliable, fully managed graph database service that makes it easy to build and run applications that work with highly connected data sets¹. Graph databases are designed to store and query data that has complex relationships and interconnections, such as social networks, recommendation engines, fraud detection, and knowledge graphs². Amazon Neptune supports two popular graph models: Property Graph and Resource Description Framework (RDF), and their respective query languages: Apache TinkerPop Gremlin and SPARQL².

Amazon Neptune also supports IAM policies to control access to the database resources and operations. You can use IAM database authentication to authenticate users and applications that connect to a Neptune DB cluster. IAM database authentication works with MySQL and PostgreSQL database clients. You can also use IAM roles to manage access to Neptune from other AWS services, such as Amazon EC2, AWS Lambda, and Amazon SageMaker².

Therefore, Amazon Neptune DB cluster is a suitable solution for the marketing company's requirements, as it can provide a graph database storage solution that is optimized for highly connected data and can limit connections and programmatic access by using IAM policies.

NO.14 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-nativ>

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

NO.15 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

NO.16 A company is running its line of business application on AWS, which uses Amazon RDS for MySQL at the persistent data store. The company wants to minimize downtime when it migrates the database to Amazon Aurora.

Which migration method should a Database Specialist use?

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

Answer: C

Explanation

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

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

NO.17 A Database Specialist is setting up a new Amazon Aurora DB cluster with one primary instance and three Aurora Replicas for a highly intensive, business-critical application. The Aurora DB cluster has one medium-sized primary instance, one large-sized replica, and two medium-sized replicas. The Database Specialist did not assign a promotion tier to the replicas.

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

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

Answer: C

Explanation

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

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

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

in size. If two or more Aurora Replicas share the same priority and size, then Amazon RDS promotes an arbitrary replica in the same promotion tier.

[https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Aurora.Managing.Backups.html](https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Aurora.Managing.Backups.html#Aurora.Ma)

If two or more Aurora Replicas share the same priority, then Amazon RDS promotes the replica that is largest in size. If two or more Aurora Replicas share the same priority and size, then Amazon RDS promotes an arbitrary replica in the same promotion tier.

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Concepts.AuroraHighAvailability.html>

NO.18 Amazon RDS for Oracle with Transparent Data Encryption is used by a financial services organization (TDE).

At all times, the organization is obligated to encrypt its data at rest. The decryption key must be widely distributed, and access to the key must be restricted. The organization must be able to rotate the encryption key on demand to comply with regulatory requirements. If any possible security vulnerabilities are discovered, the organization must be able to disable the key. Additionally, the company's overhead must be kept to a minimal.

What method should the database administrator use to configure the encryption to fulfill these specifications?

- A. AWS CloudHSM
- B. AWS Key Management Service (AWS KMS) with an AWS managed key
- C. AWS Key Management Service (AWS KMS) with server-side encryption
- D. AWS Key Management Service (AWS KMS) CMK with customer-provided material

Answer: D

Explanation

<https://docs.aws.amazon.com/whitepapers/latest/kms-best-practices/aws-managed-and-customer-managed-cmks>.

NO.19 A Database Specialist modified an existing parameter group currently associated with a production Amazon RDS for SQL Server Multi-AZ DB instance. The change is associated with a static parameter type, which controls the number of user connections allowed on the most critical RDS SQL Server DB instance for the company. This change has been approved for a specific maintenance window to help minimize the impact on users.

How should the Database Specialist apply the parameter group change for the DB instance?

- A. Select the option to apply the change immediately
- B. Allow the preconfigured RDS maintenance window for the given DB instance to control when the change is applied
- C. Apply the change manually by rebooting the DB instance during the approved maintenance window
- D. Reboot the secondary Multi-AZ DB instance

Answer: C

Explanation

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_WorkingWithParamGroups.html#USER_W

NO.20 A company needs a data warehouse solution that keeps data in a consistent, highly

structured format. The company requires fast responses for end-user queries when looking at data from the current year, and users must have access to the full 15-year dataset, when needed. This solution also needs to handle a fluctuating number incoming queries. Storage costs for the 100 TB of data must be kept low.

Which solution meets these requirements?

- 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."

NO.21 A company has deployed an application that uses an Amazon RDS for MySQL DB cluster. The DB cluster uses three read replicas. The primary DB instance is an 8XL-sized instance, and the read replicas are each XL-sized instances.

Users report that database queries are returning stale data. The replication lag indicates that the replicas are 5 minutes behind the primary DB instance. Status queries on the replicas show that the SQL_THREAD is 10 binlogs behind the IO_THREAD and that the IO_THREAD is 1 binlog behind the primary.

Which changes will reduce the lag? (Choose two.)

- A.** Deploy two additional read replicas matching the existing replica DB instance size.
- B.** Migrate the primary DB instance to an Amazon Aurora MySQL DB cluster and add three Aurora Replicas.
- C.** Move the read replicas to the same Availability Zone as the primary DB instance.
- D.** Increase the instance size of the primary DB instance within the same instance class.
- E.** Increase the instance size of the read replicas to the same size and class as the primary DB instance.

Answer: B E

Explanation

<https://www.quora.com/What-is-the-difference-between-a-RDS-read-replica-and-an-Aurora-Read-repli>

<https://aws.amazon.com/premiumsupport/knowledge-center/rds-mysql-high-replica-lag/>

NO.22 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

NO.23 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.L>

NO.24 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. Select the desired DB instance size.
- C. Update a defined Amazon Route 53 CNAME record to point to the copied DB instance.
- D. Set up an AWS Database Migration Service (AWS DMS) migration task to copy the snapshot to the copied DB instance.
- E. Use the CopySnapshot action on the production DB instance to create a snapshot before 6 AM.
- F. Update a defined Amazon Route 53 alias record to point to the copied DB instance.

Answer: A B C

NO.25 A database specialist has been entrusted by an ecommerce firm with designing a reporting dashboard that visualizes crucial business KPIs derived from the company's primary production database running on Amazon Aurora. The dashboard should be able to read data within 100 milliseconds after an update.

The Database Specialist must conduct an audit of the Aurora DB cluster's present setup and provide a cost-effective alternative. The solution must support the unexpected read demand generated by the reporting dashboard without impairing the DB cluster's write availability and performance.

Which solution satisfies these criteria?

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

Answer: D

NO.26 A company has an application that uses an Amazon DynamoDB table to store user data. Every morning, a single-threaded process calls the DynamoDB API Scan operation to scan the entire table and generate a critical start-of-day report for management. A successful marketing campaign recently doubled the number of items in the table, and now the process takes too long to run and the report is not generated in time.

A database specialist needs to improve the performance of the process. The database specialist notes that, when the process is running, 15% of the table's provisioned read capacity units (RCUs) are being used.

What should the database specialist do?

- A. Enable auto scaling for the DynamoDB table.
- B. Use four threads and parallel DynamoDB API Scan operations.
- C. Double the table's provisioned RCUs.
- D. Set the Limit and Offset parameters before every call to the API.

Answer: B

Explanation

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Scan.html#Scan.ParallelScan>

NO.27 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.

NO.28 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

NO.29 A coffee machine manufacturer is equipping all of its coffee machines with 10T sensors. The 10T core application is writing measurements for each record to Amazon Timestream. The records have multiple dimensions and measures. The measures include multiple measure names and values. An analysis application is running queries against the Timestream database and is focusing on data from the current week. A database specialist needs to optimize the query costs of the analysis application.

Which solution will meet these requirements?

- A.** Ensure that queries contain whole records over the relevant time range.
- B.** Use time range, measure name, and dimensions in the WHERE clause of the query.
- C.** Avoid canceling any query after the query starts running.
- D.** Implement exponential backoff in the application.

Answer: B

Explanation

Use time range, measure name, and dimensions in the WHERE clause of the query.

Explanation from Amazon documents:

Amazon Timestream is a serverless time series database service that allows you to store and analyze time series data at any scale. To optimize the cost of queries, you should use the following best practices¹:

Include only the measure and dimension names essential to query. Adding extraneous columns will increase data scans and therefore will also increase the query cost.

Include a time range in the WHERE clause of your query. For example, if you only need the last one hour of data in your dataset, include a time predicate such as `time > ago (1h)`.

Include the measure names in the WHERE clause of the query when a query accesses a subset of measures in a table.

Option B follows these best practices, while option A does not. Option C is incorrect because canceling a query can save on cost if the query will not return the desired results¹. Option D is irrelevant because exponential backoff is a technique to handle throttling errors, not to optimize query costs².

NO.30 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. 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.
- D.** 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>

NO.31 An online retailer uses Amazon DynamoDB for its product catalog and order data. Some popular items have led to frequently accessed keys in the data, and the company is using DynamoDB Accelerator (DAX) as the caching solution to cater to the frequently accessed keys. As the number of popular products is growing, the company realizes that more items need to be cached. The company observes a high cache miss rate and needs a solution to address this issue.

What should a database specialist do to accommodate the changing requirements for DAX?

- A.** Increase the number of nodes in the existing DAX cluster.
- B.** Create a new DAX cluster with more nodes. Change the DAX endpoint in the application to point to the new cluster.
- C.** Create a new DAX cluster using a larger node type. Change the DAX endpoint in the application to point to the new cluster.
- D.** Modify the node type in the existing DAX cluster.

Answer: C

Explanation

Create a new DAX cluster using a larger node type. Change the DAX endpoint in the application to point to the new cluster.

Explanation from Amazon documents:

The cache miss rate is the percentage of read requests that are not satisfied by the DAX cache and have to be forwarded to DynamoDB1. A high cache miss rate indicates that the DAX cluster does not have enough memory to store all the frequently accessed items. Increasing the number of nodes in the existing DAX cluster (option A) or creating a new DAX cluster with more nodes (option B) will not increase the total memory available for caching, because DAX uses a partitioned cache model, where each node is responsible for caching a subset of the data2. Modifying the node type in the existing DAX cluster (option D) will cause downtime and data loss, because DAX does not support online resizing of clusters3. Therefore, the best option is to create a new DAX cluster using a larger node type (option C), which will provide more memory per node and allow more items to be cached. The application will need to change the DAX endpoint to point to the new cluster, which can be done with minimal disruption by using DNS aliasing or load balancing3.

NO.32 A company is using an Amazon Aurora PostgreSQL database for a project with a government agency. All database communications must be encrypted in transit. All non-SSL/TLS connection requests must be rejected.

What should a database specialist do to meet these requirements?

- A.** Set the `rds.force_ssl` parameter in the DB cluster parameter group to default.
- B.** Set the `rds.force_ssl` parameter in the DB cluster parameter group to 1.
- C.** Set the `rds.force_ssl` parameter in the DB cluster parameter group to 0.
- D.** Set the `SQLNET.SSL_VERSION` option in the DB cluster option group to 12.

Answer: B

Option A is incorrect because setting the `rds.force_ssl` parameter to default (which is 0) does not enforce SSL/TLS encryption for the database connections. It allows both encrypted and unencrypted connections, which does not meet the requirement of rejecting all non-SSL/TLS connection requests. Option B is correct because setting the `rds.force_ssl` parameter to 1 requires that all database

connections use SSL/TLS encryption. This parameter also modifies the PostgreSQL `pg_hba.conf` file to require SSL/TLS for all connections. This meets the requirement of encrypting all database communications in transit and rejecting all non-SSL/TLS connection requests.

Option C is incorrect because setting the `rds.force_ssl` parameter to 0 is equivalent to setting it to default, which does not enforce SSL/TLS encryption for the database connections. It allows both encrypted and unencrypted connections, which does not meet the requirement of rejecting all non-SSL/TLS connection requests.

Option D is incorrect because `SQLNET.SSL_VERSION` is an option for Oracle databases, not PostgreSQL databases. Amazon Aurora PostgreSQL does not use option groups, but parameter groups to configure database settings. Moreover, `SQLNET.SSL_VERSION` specifies the version of SSL/TLS protocol to use, not whether to enforce SSL/TLS encryption for the database connections.