



Up-to-date Questions and Answers from authentic resources to improve knowledge and pass the exam at very first attempt. ----- Guaranteed.



DBS-C01 MCQs
DBS-C01 TestPrep
DBS-C01 Study Guide
DBS-C01 Practice Test
DBS-C01 Exam Questions



killexams.com

Amazon

DBS-C01

AWS Certified Database - Specialty

ORDER FULL VERSION

<https://killexams.com/pass4sure/exam-detail/DBS-C01>



Question: 87

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: 88

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

Question: 89

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

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

- A. Contact AWS Support to disable read-only mode on the read replica. Reboot the read replica. Connect to the read replica and create the tables.
- B. Change the read_only parameter to false (read_only=0) in the default parameter group of the read replica. Perform a reboot without failover. Connect to the read replica and create the tables using the local_only MySQL option.
- C. Change the read_only parameter to false (read_only=0) in the default parameter group. Reboot the read replica. Connect to the read replica and create the tables.
- D. Create a new DB parameter group. Change the read_only parameter to false (read_only=0). Associate the read replica with the new group. Reboot the read replica. Connect to the read replica and create the tables.

Answer: D

Explanation:

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

Question: 90

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-custom-endpoints/>

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.

Question: 91

Developers have requested a new Amazon Redshift cluster so they can load new third-party marketing data. The new cluster is ready and the user credentials are given to the developers.

The developers indicate that their copy jobs fail with the following error message:

âAmazon Invalid operation: S3ServiceException:Access Denied,Status 403,Error AccessDenied.â

The developers need to load this data soon, so a database specialist must act quickly to solve this issue.

What is the MOST secure solution?

- A. Create a new IAM role with the same user name as the Amazon Redshift developer user I
- B. Provide the IAM role with read-only access to Amazon S3 with the assume role action.
- C. Create a new IAM role with read-only access to the Amazon S3 bucket and include the assume role action. Modify the Amazon Redshift cluster to add the IAM role.
- D. Create a new IAM role with read-only access to the Amazon S3 bucket with the assume role action. Add this role to the developer IAM user ID used for the copy job that ended with an error message.
- E. Create a new IAM user with access keys and a new role with read-only access to the Amazon S3 bucket. Add this role to the Amazon Redshift cluster. Change the copy job to use the access keys created.

Answer: B

Explanation:

<https://docs.aws.amazon.com/redshift/latest/gsg/rs-gsg-create-an-iam-role.html>

"Now that you have created the new role, your next step is to attach it to your cluster. You can attach

the role when you launch a new cluster or you can attach it to an existing cluster. In the next step, you attach the role to a new cluster."

https://docs.aws.amazon.com/redshift/latest/dg/copy-usage_notes-access-permissions.html

Question: 92

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 Region. Enable point-in-time recovery for both tables.
- C. Use a DynamoDB Accelerator table in another Region. Enable point-in-time recovery for the table.
- D. Create an AWS Backup plan and assign the DynamoDB table as a resource.

Answer: C

Question: 93

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: 94

A software development company is using Amazon Aurora MySQL DB clusters for several use cases, including development and reporting. These use cases place unpredictable and varying demands on the Aurora DB clusters, and can cause momentary spikes in latency. System users run ad-hoc queries sporadically throughout the week. Cost is a primary concern for the company, and a solution that does not require significant rework is needed.

Which solution meets these requirements?

- A. Create new Aurora Serverless DB clusters for development and reporting, then migrate to these new DB clusters.
- B. Upgrade one of the DB clusters to a larger size, and consolidate development and reporting activities on this larger DB cluster.
- C. Use existing DB clusters and stop/start the databases on a routine basis using scheduling tools.
- D. Change the DB clusters to the burstable instance family.

Answer: A

Explanation:

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

Question: 95

A database specialist is building a system that uses a static vendor dataset of postal codes and related territory information that is less than 1 GB in size. The dataset is loaded into the application's cache at start up. The company needs to store this data in a way that provides the lowest cost with a low application startup time.

Which approach will meet these requirements?

- A. Use an Amazon RDS DB instance. Shut down the instance once the data has been read.
- B. Use Amazon Aurora Serverless. Allow the service to spin resources up and down, as needed.
- C. Use Amazon DynamoDB in on-demand capacity mode.
- D. Use Amazon S3 and load the data from flat files.

Answer: D

Explanation:

<https://www.sumologic.com/insight/s3-cost-optimization/>

For example, for 1 GB file stored on S3 with 1 TB of storage provisioned, you are billed for 1 GB only. In a lot of other services such as Amazon EC2, Amazon Elastic Block Storage (Amazon EBS) and Amazon DynamoDB you pay for provisioned capacity. For example, in the case of Amazon EBS disk you pay for the size of 1 TB of disk even if you just save 1 GB file. This makes managing S3 cost easier than many other services including Amazon EBS and Amazon EC2. On S3 there is no risk of over-provisioning and no need to manage disk utilization.

Question: 96

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: 97

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.
- 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-secrets-manager/>

Question: 98

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

What should the database specialist do to connect to the new, restored Amazon DocumentDB cluster?

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

Answer: A

Explanation:

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

Question: 99

A company runs a customer relationship management (CRM) system that is hosted on-premises with a MySQL database as the backend. A custom stored procedure is used to send email notifications to another system when data is inserted into a table. The company has noticed that the performance of the CRM system has decreased due to database reporting applications used by various teams. The company requires an AWS solution that would reduce maintenance, improve performance, and

accommodate the email notification feature.

Which AWS solution meets these requirements?

- A. Use MySQL running on an Amazon EC2 instance with Auto Scaling to accommodate the reporting applications. Configure a stored procedure and an AWS Lambda function that uses Amazon SES to send email notifications to the other system.
- B. Use Amazon Aurora MySQL in a multi-master cluster to accommodate the reporting applications. Configure Amazon RDS event subscriptions to publish a message to an Amazon SNS topic and subscribe the other system's email address to the topic.
- C. Use MySQL running on an Amazon EC2 instance with a read replica to accommodate the reporting applications. Configure Amazon SES integration to send email notifications to the other system.
- D. Use Amazon Aurora MySQL with a read replica for the reporting applications. Configure a stored procedure and an AWS Lambda function to publish a message to an Amazon SNS topic. Subscribe the other system's email address to the topic.

Answer: D

Explanation:

RDS event subscriptions do not cover "data is inserted into a table" - see

https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/USER_Events.Messages.html We can use stored procedure to invoke Lambda function -
<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraMySQL.Integrating.Lambda.html>

Question: 100

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

Which solution will meet these requirements?

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

Answer: B

Explanation:

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

Killexams.com is a leading online platform specializing in high-quality certification exam preparation. Offering a robust suite of tools, including MCQs, practice tests, and advanced test engines, Killexams.com empowers candidates to excel in their certification exams. Discover the key features that make Killexams.com the go-to choice for exam success.



Exam Questions:

Killexams.com provides exam questions that are experienced in test centers. These questions are updated regularly to ensure they are up-to-date and relevant to the latest exam syllabus. By studying these questions, candidates can familiarize themselves with the content and format of the real exam.

Exam MCQs:

Killexams.com offers exam MCQs in PDF format. These questions contain a comprehensive collection of questions and answers that cover the exam topics. By using these MCQs, candidate can enhance their knowledge and improve their chances of success in the certification exam.

Practice Test:

Killexams.com provides practice test through their desktop test engine and online test engine. These practice tests simulate the real exam environment and help candidates assess their readiness for the actual exam. The practice test cover a wide range of questions and enable candidates to identify their strengths and weaknesses.

Guaranteed Success:

Killexams.com offers a success guarantee with the exam MCQs. Killexams claim that by using this materials, candidates will pass their exams on the first attempt or they will get refund for the purchase price. This guarantee provides assurance and confidence to individuals preparing for certification exam.

Updated Contents:

Killexams.com regularly updates its question bank of MCQs to ensure that they are current and reflect the latest changes in the exam syllabus. This helps candidates stay up-to-date with the exam content and increases their chances of success.