

Exam Questions Professional-Cloud-Database-Engineer

Google Cloud Certified - Professional Cloud Database Engineer

<https://www.2passeasy.com/dumps/Professional-Cloud-Database-Engineer/>



NEW QUESTION 1

Your organization has an existing app that just went viral. The app uses a Cloud SQL for MySQL backend database that is experiencing slow disk performance while using hard disk drives (HDDs). You need to improve performance and reduce disk I/O wait times. What should you do?

- A. Export the data from the existing instance, and import the data into a new instance with solid-state drives (SSDs).
- B. Edit the instance to change the storage type from HDD to SSD.
- C. Create a high availability (HA) failover instance with SSDs, and perform a failover to the new instance.
- D. Create a read replica of the instance with SSDs, and perform a failover to the new instance

Answer: A

Explanation:

<https://stackoverflow.com/questions/72034607/can-i-change-storage-type-from-hdd-to-ssd-on-cloud-sql-after-creating-an-instanc>

NEW QUESTION 2

Your organization works with sensitive data that requires you to manage your own encryption keys. You are working on a project that stores that data in a Cloud SQL database. You need to ensure that stored data is encrypted with your keys. What should you do?

- A. Export data periodically to a Cloud Storage bucket protected by Customer-Supplied Encryption Keys.
- B. Use Cloud SQL Auth proxy.
- C. Connect to Cloud SQL using a connection that has SSL encryption.
- D. Use customer-managed encryption keys with Cloud SQL.

Answer: D

NEW QUESTION 3

Your application uses Cloud SQL for MySQL. Your users run reports on data that relies on near-real time; however, the additional analytics caused excessive load on the primary database. You created a read replica for the analytics workloads, but now your users are complaining about the lag in data changes and that their reports are still slow. You need to improve the report performance and shorten the lag in data replication without making changes to the current reports. Which two approaches should you implement? (Choose two.)

- A. Create secondary indexes on the replica.
- B. Create additional read replicas, and partition your analytics users to use different read replicas.
- C. Disable replication on the read replica, and set the flag for parallel replication on the read replic
- D. Re-enable replication and optimize performance by setting flags on the primary instance.
- E. Disable replication on the primary instance, and set the flag for parallel replication on the primary instanc
- F. Re-enable replication and optimize performance by setting flags on the read replica.
- G. Move your analytics workloads to BigQuery, and set up a streaming pipeline to move data and update BigQuery.

Answer: BC

Explanation:

Replication lag and slow report performance. E is eliminated because using BigQuery would mean changes to the current reports. Report slowness could be the result of poor indexing or just too much read load (or both!). Since excessive load is mentioned in the question, creating additional read replicas and spreading the analytics workload around makes B correct and eliminates A as a way to speed up reporting. That leaves the replication problem. Cloud SQL enables single threaded replication by default, so it stands to reason enabling parallel replication would help the lag. To do that you disable replication on the replica (not the primary), set flags on the replica and optionally set flags on the primary instance to optimize performance for parallel replication. That makes C correct and D incorrect. <https://cloud.google.com/sql/docs/mysql/replication/manage-replicas#configuring-parallel-replication>

NEW QUESTION 4

Your company is migrating all legacy applications to Google Cloud. All on-premises applications are using legacy Oracle 12c databases with Oracle Real Application Cluster (RAC) for high availability (HA) and Oracle Data Guard for disaster recovery. You need a solution that requires minimal code changes, provides the same high availability you have today on-premises, and supports a low latency network for migrated legacy applications. What should you do?

- A. Migrate the databases to Cloud Spanner.
- B. Migrate the databases to Cloud SQL, and enable a standby database.
- C. Migrate the databases to Compute Engine using regional persistent disks.
- D. Migrate the databases to Bare Metal Solution for Oracle.

Answer: D

Explanation:

BMS is the only Google database service which supports Oracle aside from GCVE. It allows you to use all native Oracle features including RAC. Since GCVE isn't mentioned, it has to be D - Bare Metal Solution.

NEW QUESTION 5

Your customer is running a MySQL database on-premises with read replicas. The nightly incremental backups are expensive and add maintenance overhead. You want to follow Google-recommended practices to migrate the database to Google Cloud, and you need to ensure minimal downtime. What should you do?

- A. Create a Google Kubernetes Engine (GKE) cluster, install MySQL on the cluster, and then import the dump file.
- B. Use the mysqldump utility to take a backup of the existing on-premises database, and then import it into Cloud SQL.
- C. Create a Compute Engine VM, install MySQL on the VM, and then import the dump file.
- D. Create an external replica, and use Cloud SQL to synchronize the data to the replica.

Answer: D

Explanation:

<https://cloud.google.com/sql/docs/mysql/replication/configure-replication-from-external>

NEW QUESTION 6

You need to migrate a 1 TB PostgreSQL database from a Compute Engine VM to Cloud SQL for PostgreSQL. You want to ensure that there is minimal downtime during the migration. What should you do?

- A. Export the data from the existing database, and load the data into a new Cloud SQL database.
- B. Use Migrate for Compute Engine to complete the migration.
- C. Use Datastream to complete the migration.
- D. Use Database Migration Service to complete the migration.

Answer: D

Explanation:

<https://www.cloudskillsboost.google/focuses/22792?parent=catalog>

NEW QUESTION 7

Your company is shutting down their on-premises data center and migrating their Oracle databases using Oracle Real Application Clusters (RAC) to Google Cloud. You want minimal to no changes to the applications during the database migration. What should you do?

- A. Migrate the Oracle databases to Cloud Spanner.
- B. Migrate the Oracle databases to Compute Engine.
- C. Migrate the Oracle databases to Cloud SQL.
- D. Migrate the Oracle databases to Bare Metal Solution for Oracle.

Answer: D

Explanation:

This answer is correct because Bare Metal Solution for Oracle is a service that provides dedicated physical servers and networking infrastructure for running Oracle databases on Google Cloud¹. Bare Metal Solution for Oracle supports Oracle RAC, which is a cluster database that provides high availability, scalability, and performance for Oracle workloads². By using Bare Metal Solution for Oracle, you can migrate your Oracle databases with minimal to no changes to the applications, and you can leverage the native Google Cloud services and interconnectivity¹.

NEW QUESTION 8

You are choosing a new database backend for an existing application. The current database is running PostgreSQL on an on-premises VM and is managed by a database administrator and operations team. The application data is relational and has light traffic. You want to minimize costs and the migration effort for this application. What should you do?

- A. Migrate the existing database to Firestore.
- B. Migrate the existing database to Cloud SQL for PostgreSQL.
- C. Migrate the existing database to Cloud Spanner.
- D. Migrate the existing database to PostgreSQL running on Compute Engine.

Answer: B

Explanation:

You could migrate to Spanner leveraging the PostgreSQL dialect, but costs need to be minimized so that wouldn't be the cheapest option. Especially since the load doesn't justify Spanner. Again, you could migrate like-for-like to a GCE VM, but that defeats minimizing the migration effort. The cheapest and easiest way to migrate would be Database Migration Service to Cloud SQL for PostgreSQL.

NEW QUESTION 9

Your company wants you to migrate their Oracle, MySQL, Microsoft SQL Server, and PostgreSQL relational databases to Google Cloud. You need a fully managed, flexible database solution when possible. What should you do?

- A. Migrate all the databases to Cloud SQL.
- B. Migrate the Oracle, MySQL, and Microsoft SQL Server databases to Cloud SQL, and migrate the PostgreSQL databases to Compute Engine.
- C. Migrate the MySQL, Microsoft SQL Server, and PostgreSQL databases to Compute Engine, and migrate the Oracle databases to Bare Metal Solution for Oracle.
- D. Migrate the MySQL, Microsoft SQL Server, and PostgreSQL databases to Cloud SQL, and migrate the Oracle databases to Bare Metal Solution for Oracle.

Answer: D

NEW QUESTION 10

You manage a meeting booking application that uses Cloud SQL. During an important launch, the Cloud SQL instance went through a maintenance event that resulted in a downtime of more than 5 minutes and adversely affected your production application. You need to immediately address the maintenance issue to prevent any unplanned events in the future. What should you do?

- A. Set your production instance's maintenance window to non-business hours.
- B. Migrate the Cloud SQL instance to Cloud Spanner to avoid any future disruptions due to maintenance.
- C. Contact Support to understand why your Cloud SQL instance had a downtime of more than 5 minutes.
- D. Use Cloud Scheduler to schedule a maintenance window of no longer than 5 minutes.

Answer: A

NEW QUESTION 10

You are configuring a brand new Cloud SQL for PostgreSQL database instance in Google Cloud. Your application team wants you to deploy one primary instance, one standby instance, and one read replica instance. You need to ensure that you are following Google- recommended practices for high availability. What should you do?

- A. Configure the primary instance in zone A, the standby instance in zone C, and the read replica in zone B, all in the same region.
- B. Configure the primary and standby instances in zone A and the read replica in zone B, all in the same region.
- C. Configure the primary instance in one region, the standby instance in a second region, and the read replica in a third region.
- D. Configure the primary, standby, and read replica instances in zone A, all in the same region.

Answer: A

Explanation:

<https://cloud.google.com/sql/docs/postgres/high-availability#failover-overview>

NEW QUESTION 11

Your organization is running a MySQL workload in Cloud SQL. Suddenly you see a degradation in database performance. You need to identify the root cause of the performance degradation. What should you do?

- A. Use Logs Explorer to analyze log data.
- B. Use Cloud Monitoring to monitor CPU, memory, and storage utilization metrics.
- C. Use Error Reporting to count, analyze, and aggregate the data.
- D. Use Cloud Debugger to inspect the state of an application.

Answer: B

Explanation:

<https://cloud.google.com/sql/docs/mysql/diagnose-issues#:~:text=If%20your%20instance%20stops%20responding%20to%20connections%20or%20performance%20is%20degraded%2C%20make%20sure%20it%20conforms%20to%20the%20Operational%20Guidelines>

NEW QUESTION 12

You are migrating an on-premises application to Google Cloud. The application requires a high availability (HA) PostgreSQL database to support business-critical functions. Your company's disaster recovery strategy requires a recovery time objective (RTO) and recovery point objective (RPO) within 30 minutes of failure. You plan to use a Google Cloud managed service. What should you do to maximize uptime for your application?

- A. Deploy Cloud SQL for PostgreSQL in a regional configuration
- B. Create a read replica in a different zone in the same region and a read replica in another region for disaster recovery.
- C. Deploy Cloud SQL for PostgreSQL in a regional configuration with HA enable
- D. Take periodic backups, and use this backup to restore to a new Cloud SQL for PostgreSQL instance in another region during a disaster recovery event.
- E. Deploy Cloud SQL for PostgreSQL in a regional configuration with HA enable
- F. Create a cross-region read replica, and promote the read replica as the primary node for disaster recovery.
- G. Migrate the PostgreSQL database to multi-regional Cloud Spanner so that a single region outage will not affect your application
- H. Update the schema to support Cloud Spanner data types, and refactor the application.

Answer: C

Explanation:

The best answer is deploy an HA configuration and have a read replica you could promote to the primary in a different region

NEW QUESTION 14

You are writing an application that will run on Cloud Run and require a database running in the Cloud SQL managed service. You want to secure this instance so that it only receives connections from applications running in your VPC environment in Google Cloud. What should you do?

- A. * 1. Create your instance with a specified external (public) IP address.* 2. Choose the VPC and create firewall rules to allow only connections from Cloud Run into your instance.* 3. Use Cloud SQL Auth proxy to connect to the instance.
- B. * 1. Create your instance with a specified external (public) IP address.* 2. Choose the VPC and create firewall rules to allow only connections from Cloud Run into your instance.* 3. Connect to the instance using a connection pool to best manage connections to the instance.
- C. * 1. Create your instance with a specified internal (private) IP address.* 2. Choose the VPC with private service connection configured.* 3. Configure the Serverless VPC Access connector in the same VPC network as your Cloud SQL instance.* 4. Use Cloud SQL Auth proxy to connect to the instance.
- D. * 1. Create your instance with a specified internal (private) IP address.* 2. Choose the VPC with private service connection configured.* 3. Configure the Serverless VPC Access connector in the same VPC network as your Cloud SQL instance.* 4. Connect to the instance using a connection pool to best manage connections to the instance.

Answer: D

Explanation:

<https://cloud.google.com/sql/docs/mysql/connect-run#configure> <https://cloud.google.com/sql/docs/mysql/connect-run#connection-pools>

NEW QUESTION 15

Your organization has a security policy to ensure that all Cloud SQL for PostgreSQL databases are secure. You want to protect sensitive data by using a key that meets specific locality or residency requirements. Your organization needs to control the key's lifecycle activities. You need to ensure that data is encrypted at rest and in transit. What should you do?

- A. Create the database with Google-managed encryption keys.
- B. Create the database with customer-managed encryption keys.
- C. Create the database persistent disk with Google-managed encryption keys.
- D. Create the database persistent disk with customer-managed encryption keys.

Answer: B

Explanation:

<https://cloud.google.com/sql/docs/postgres/configure-cmek#createcmekinstance>

NEW QUESTION 17

You want to migrate an on-premises mission-critical PostgreSQL database to Cloud SQL.

The database must be able to withstand a zonal failure with less than five minutes of downtime and still not lose any transactions. You want to follow Google-recommended practices for the migration. What should you do?

- A. Take nightly snapshots of the primary database instance, and restore them in a secondary zone.
- B. Build a change data capture (CDC) pipeline to read transactions from the primary instance, and replicate them to a secondary instance.
- C. Create a read replica in another region, and promote the read replica if a failure occurs.
- D. Enable high availability (HA) for the database to make it regional.

Answer: D

NEW QUESTION 20

Your organization needs to migrate a critical, on-premises MySQL database to Cloud SQL for MySQL. The on-premises database is on a version of MySQL that is supported by Cloud SQL and uses the InnoDB storage engine. You need to migrate the database while preserving transactions and minimizing downtime. What should you do?

- A. Use Database Migration Service to connect to your on-premises database, and choose continuous replication. After the on-premises database is migrated, promote the Cloud SQL for MySQL instance, and connect applications to your Cloud SQL instance.
- B. Build a Cloud Data Fusion pipeline for each table to migrate data from the on-premises MySQL database to Cloud SQL for MySQL. Schedule downtime to run each Cloud Data Fusion pipeline.
- C. Verify that the migration was successful. Re-point the applications to the Cloud SQL for MySQL instance.
- D. Pause the on-premises applications. Use the mysqldump utility to dump the database content in compressed format.
- E. Run gsutil -m to move the dump file to Cloud Storage. Use the Cloud SQL for MySQL import option. After the import operation is complete, re-point the applications to the Cloud SQL for MySQL instance.
- F. Pause the on-premises applications. Use the mysqldump utility to dump the database content in CSV format.
- G. Run gsutil -m to move the dump file to Cloud Storage. Use the Cloud SQL for MySQL import option. After the import operation is complete, re-point the applications to the Cloud SQL for MySQL instance.

Answer: A

Explanation:

<https://cloud.google.com/database-migration/docs/mysql/configure-source-database>

To migrate the database while preserving transactions and minimizing downtime, you should use Database Migration Service. This service will allow you to migrate the database in a way that is transparent to your users and applications. It will also allow you to test the migration before you make it live, so that you can be sure that everything will work as expected.

NEW QUESTION 23

Your company is migrating their MySQL database to Cloud SQL and cannot afford any planned downtime during the month of December. The company is also concerned with cost, so you need the most cost-effective solution. What should you do?

- A. Open a support ticket in Google Cloud to prevent any maintenance in that MySQL instance during the month of December.
- B. Use Cloud SQL maintenance settings to prevent any maintenance during the month of December.
- C. Create MySQL read replicas in different zones so that, if any downtime occurs, the read replicas will act as the primary instance during the month of December.
- D. Create a MySQL regional instance so that, if any downtime occurs, the standby instance will act as the primary instance during the month of December.

Answer: B

Explanation:

<https://cloud.google.com/sql/docs/mysql/maintenance?hl=fr>

NEW QUESTION 26

You are the primary DBA of a Cloud SQL for PostgreSQL database that supports 6 enterprise applications in production. You used Cloud SQL Insights to identify inefficient queries and now need to identify the application that is originating the inefficient queries. You want to follow Google-recommended practices. What should you do?

- A. Shut down and restart each application.
- B. Write a utility to scan database query logs.
- C. Write a utility to scan application logs.
- D. Use query tags to add application-centric database monitoring.

Answer: D

Explanation:

https://cloud.google.com/sql/docs/postgres/using-query-insights#filter_by_query_tags

NEW QUESTION 29

You are managing two different applications: Order Management and Sales Reporting. Both applications interact with the same Cloud SQL for MySQL database. The Order Management application reads and writes to the database 24/7, but the Sales Reporting application is read-only. Both applications need the latest data. You need to ensure that the Performance of the Order Management application is not affected by the Sales Reporting application. What should you do?

- A. Create a read replica for the Sales Reporting application.
- B. Create two separate databases in the instance, and perform dual writes from the Order Management application.
- C. Use a Cloud SQL federated query for the Sales Reporting application.
- D. Queue up all the requested reports in PubSub, and execute the reports at night.

Answer: A

NEW QUESTION 32

Your company wants to migrate an Oracle-based application to Google Cloud. The application team currently uses Oracle Recovery Manager (RMAN) to back up the database to tape for long-term retention (LTR). You need a cost-effective backup and restore solution that meets a 2-hour recovery time objective (RTO) and a 15-minute recovery point objective (RPO). What should you do?

- A. Migrate the Oracle databases to Bare Metal Solution for Oracle, and store backups on tapes on-premises.
- B. Migrate the Oracle databases to Bare Metal Solution for Oracle, and use Actifio to store backup files on Cloud Storage using the Nearline Storage class.
- C. Migrate the Oracle databases to Bare Metal Solution for Oracle, and back up the Oracle databases to Cloud Storage using the Standard Storage class.
- D. Migrate the Oracle databases to Compute Engine, and store backups on tapes on-premises.

Answer: B

Explanation:

<https://www.actifio.com/solutions/cloud/google/>

NEW QUESTION 33

Your team is building an application that stores and analyzes streaming time series financial data. You need a database solution that can perform time series-based scans with sub-second latency. The solution must scale into the hundreds of terabytes and be able to write up to 10k records per second and read up to 200 MB per second. What should you do?

- A. Use Firestore.
- B. Use Bigtable
- C. Use BigQuery.
- D. Use Cloud Spanner.

Answer: B

Explanation:

Financial data, such as transaction histories, stock prices, and currency exchange rates.

<https://cloud.google.com/bigtable/docs/overview#what-its-good-for>

With SSD:

Reads - up to 10,000 rows per second Writes - up to 10,000 rows per second Scans - up to 220 MB/s

<https://cloud.google.com/bigtable/docs/performance#typical-workloads>

NEW QUESTION 34

Your team is running a Cloud SQL for MySQL instance with a 5 TB database that must be available 24/7. You need to save database backups on object storage with minimal operational overhead or risk to your production workloads. What should you do?

- A. Use Cloud SQL serverless exports.
- B. Create a read replica, and then use the mysqldump utility to export each table.
- C. Clone the Cloud SQL instance, and then use the mysqldump utility to export the data.
- D. Use the mysqldump utility on the primary database instance to export the backup.

Answer: A

Explanation:

<https://cloud.google.com/blog/products/databases/introducing-cloud-sql-serverless-exports>

NEW QUESTION 38

You are the database administrator of a Cloud SQL for PostgreSQL instance that has pgaudit disabled. Users are complaining that their queries are taking longer to execute and performance has degraded over the past few months. You need to collect and analyze query performance data to help identify slow-running queries. What should you do?

- A. View Cloud SQL operations to view historical query information.
- B. Write a Logs Explorer query to identify database queries with high execution times.
- C. Review application logs to identify database calls.
- D. Use the Query Insights dashboard to identify high execution times.

Answer: D

Explanation:

A Cloud SQL instance configured for HA is also called a regional instance and has a primary and secondary zone within the configured region. Within a regional instance, the configuration is made up of a primary instance and a standby instance. Through synchronous replication to each zone's persistent disk, all writes made to the primary instance are replicated to disks in both zones before a transaction is reported as committed. In the event of an instance or zone failure, the standby instance becomes the new primary instance. Users are then rerouted to the new primary instance. This process is called a failover.

NEW QUESTION 43

You have an application that sends banking events to Bigtable cluster-a in us-east. You decide to add cluster-b in us-central1. Cluster-a replicates data to cluster-b. You need to ensure that Bigtable continues to accept read and write requests if one of the clusters becomes unavailable and that requests are routed automatically to the other cluster. What deployment strategy should you use?

- A. Use the default app profile with single-cluster routing.
- B. Use the default app profile with multi-cluster routing.
- C. Create a custom app profile with multi-cluster routing.
- D. Create a custom app profile with single-cluster routing.

Answer: C

Explanation:

<https://cloud.google.com/bigtable/docs/app-profiles#default-app-profile> The question states that a single cluster existed first, then a second cluster was added. Google's documentation states, "if you created the instance with one cluster, the default app profile uses single-cluster routing. This ensures that adding additional clusters later does not change the behavior of your existing applications". Simply adding a second cluster does not change the default profile from single-cluster routing to multi-cluster routing. Since you need multi-cluster routing, you're going to need a custom app profile. So C is correct. <https://cloud.google.com/bigtable/docs/app-profiles#default-app-profile>

NEW QUESTION 48

You are designing a database strategy for a new web application in one region. You need to minimize write latency. What should you do?

- A. Use Cloud SQL with cross-region replicas.
- B. Use high availability (HA) Cloud SQL with multiple zones.
- C. Use zonal Cloud SQL without high availability (HA).
- D. Use Cloud Spanner in a regional configuration.

Answer: D

Explanation:

<https://docs.google.com/forms/d/e/1FAIpQLSfZ77ZnuUL0NpU-bOtO5QUkC0cnRCe5YKMiubLXwfV3abBqkg/viewform>

NEW QUESTION 53

Your organization has a production Cloud SQL for MySQL instance. Your instance is configured with 16 vCPUs and 104 GB of RAM that is running between 90% and 100% CPU utilization for most of the day. You need to scale up the database and add vCPUs with minimal interruption and effort. What should you do?

- A. Issue a `gcloud sql instances patch` command to increase the number of vCPUs.
- B. Update a MySQL database flag to increase the number of vCPUs.
- C. Issue a `gcloud compute instances update` command to increase the number of vCPUs.
- D. Back up the database, create an instance with additional vCPUs, and restore the database.

Answer: A

Explanation:

<https://cloud.google.com/sdk/gcloud/reference/sql/instances/patch>

NEW QUESTION 58

You have deployed a Cloud SQL for SQL Server instance. In addition, you created a cross-region read replica for disaster recovery (DR) purposes. Your company requires you to maintain and monitor a recovery point objective (RPO) of less than 5 minutes. You need to verify that your cross-region read replica meets the allowed RPO. What should you do?

- A. Use Cloud SQL instance monitoring.
- B. Use the Cloud Monitoring dashboard with available metrics from Cloud SQL.
- C. Use Cloud SQL logs.
- D. Use the SQL Server Always On Availability Group dashboard.

Answer: D

Explanation:

Note, you cannot create a read replica in Cloud SQL for SQL Server unless you use an Enterprise Edition. Which is also a requirement for configuring SQL Server AG. That's not a coincidence. That's how Cloud SQL for SQL Server creates SQL Server read replicas. To find out about the replication, use the AG Dashboard in SSMS.

<https://cloud.google.com/sql/docs/sqlserver/replication/manage-replicas#promote-replica>

NEW QUESTION 63

You are running a large, highly transactional application on Oracle Real Application Cluster (RAC) that is multi-tenant and uses shared storage. You need a solution that ensures high-performance throughput and a low-latency connection between applications and databases. The solution must also support existing Oracle features and provide ease of migration to Google Cloud. What should you do?

- A. Migrate to Compute Engine.
- B. Migrate to Bare Metal Solution for Oracle.
- C. Migrate to Google Kubernetes Engine (GKE)
- D. Migrate to Google Cloud VMware Engine

Answer: B

Explanation:

Oracle is neither licensed nor supported in GCE. The only platform which supports RAC and all existing Oracle features is BMS.

NEW QUESTION 65

.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual Professional-Cloud-Database-Engineer Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the Professional-Cloud-Database-Engineer Product From:

<https://www.2passeasy.com/dumps/Professional-Cloud-Database-Engineer/>

Money Back Guarantee

Professional-Cloud-Database-Engineer Practice Exam Features:

- * Professional-Cloud-Database-Engineer Questions and Answers Updated Frequently
- * Professional-Cloud-Database-Engineer Practice Questions Verified by Expert Senior Certified Staff
- * Professional-Cloud-Database-Engineer Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- * Professional-Cloud-Database-Engineer Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year