

# Google

## Exam Questions Professional-Cloud-Database-Engineer

Google Cloud Certified - Professional Cloud Database Engineer



**NEW QUESTION 1**

You released a popular mobile game and are using a 50 TB Cloud Spanner instance to store game data in a PITR-enabled production environment. When you analyzed the game statistics, you realized that some players are exploiting a loophole to gather more points to get on the leaderboard. Another DBA accidentally ran an emergency bugfix script that corrupted some of the data in the production environment. You need to determine the extent of the data corruption and restore the production environment. What should you do? (Choose two.)

- A. If the corruption is significant, use backup and restore, and specify a recovery timestamp.
- B. If the corruption is significant, perform a stale read and specify a recovery timestamp
- C. Write the results back.
- D. If the corruption is significant, use import and export.
- E. If the corruption is insignificant, use backup and restore, and specify a recovery timestamp.
- F. If the corruption is insignificant, perform a stale read and specify a recovery timestamp
- G. Write the results back.

**Answer:** AE

**Explanation:**

<https://cloud.google.com/spanner/docs/pitr#ways-to-recover>

To recover the entire database, backup or export the database specifying a timestamp in the past and then restore or import it to a new database. This is typically used to recover from data corruption issues when you have to revert the entire database to a point-in-time before the corruption occurred.

This part describes significant corruption - A

To recover a portion of the database, perform a stale read specifying a query-condition and timestamp in the past, and then write the results back into the live database. This is typically used for surgical operations on a live database. For example, if you accidentally delete a particular row or incorrectly update a subset of data, you can recover it with this method.

This describes insignificant corruption case – E <https://cloud.google.com/spanner/docs/pitr> <https://cloud.google.com/spanner/docs/backup/restore-backup>

**NEW QUESTION 2**

Your ecommerce website captures user clickstream data to analyze customer traffic patterns in real time and support personalization features on your website. You plan to analyze this data using big data tools. You need a low-latency solution that can store 8 TB of data and can scale to millions of read and write requests per second. What should you do?

- A. Write your data into Bigtable and use Dataproc and the Apache Hbase libraries for analysis.
- B. Deploy a Cloud SQL environment with read replicas for improved performance
- C. Use Datastream to export data to Cloud Storage and analyze with Dataproc and the Cloud Storage connector.
- D. Use Memorystore to handle your low-latency requirements and for real-time analytics.
- E. Stream your data into BigQuery and use Dataproc and the BigQuery Storage API to analyze large volumes of data.

**Answer:** A

**Explanation:**

Start with the lowest tier and smallest size and then grow your instance as needed. Memorystore provides automated scaling using APIs, and optimized node placement across zones for redundancy. Memorystore for Memcached can support clusters as large as 5 TB, enabling millions of QPS at very low latency

**NEW QUESTION 3**

You are responsible for designing a new database for an airline ticketing application in Google Cloud. This application must be able to:

Work with transactions and offer strong consistency. Work with structured and semi-structured (JSON) data.

Scale transparently to multiple regions globally as the operation grows.

You need a Google Cloud database that meets all the requirements of the application. What should you do?

- A. Use Cloud SQL for PostgreSQL with both cross-region read replicas.
- B. Use Cloud Spanner in a multi-region configuration.
- C. Use Firestore in Datastore mode.
- D. Use a Bigtable instance with clusters in multiple regions.

**Answer:** B

**Explanation:**

<https://cloud.google.com/blog/products/databases/manage-semi-structured-data-in-cloud-spanner-with-json>

**NEW QUESTION 4**

You need to provision several hundred Cloud SQL for MySQL instances for multiple project teams over a one-week period. You must ensure that all instances adhere to company standards such as instance naming conventions, database flags, and tags. What should you do?

- A. Automate instance creation by writing a Dataflow job.
- B. Automate instance creation by setting up Terraform scripts.
- C. Create the instances using the Google Cloud Console UI.
- D. Create clones from a template Cloud SQL instance.

**Answer:** B

**NEW QUESTION 5**

Your company wants to migrate its MySQL, PostgreSQL, and Microsoft SQL Server on-premises databases to Google Cloud. You need a solution that provides near-zero downtime, requires no application changes, and supports change data capture (CDC). What should you do?

- A. Use the native export and import functionality of the source database.
- B. Create a database on Google Cloud, and use database links to perform the migration.
- C. Create a database on Google Cloud, and use Dataflow for database migration.

D. Use Database Migration Service.

**Answer:** D

**Explanation:**

Simplify migrations to the cloud. Available now for MySQL and PostgreSQL, with SQL Server and Oracle migrations in preview.

- Migrate to Cloud SQL and AlloyDB for PostgreSQL from on-premises, Google Cloud, or other clouds
- Replicate data continuously for minimal downtime migrations
- Serverless and easy to set up

**NEW QUESTION 6**

You are evaluating Cloud SQL for PostgreSQL as a possible destination for your on- premises PostgreSQL instances. Geography is becoming increasingly relevant to customer privacy worldwide. Your solution must support data residency requirements and include a strategy to:

configure where data is stored

control where the encryption keys are stored govern the access to data

What should you do?

- A. Replicate Cloud SQL databases across different zones.
- B. Create a Cloud SQL for PostgreSQL instance on Google Cloud for the data that does not need to adhere to data residency requirement
- C. Keep the data that must adhere to data residency requirements on-premise
- D. Make application changes to support both databases.
- E. Allow application access to data only if the users are in the same region as the Google Cloud region for the Cloud SQL for PostgreSQL database.
- F. Use features like customer-managed encryption keys (CMEK), VPC Service Controls, and Identity and Access Management (IAM) policies.

**Answer:** D

**Explanation:**

<https://cloud.google.com/blog/products/identity-security/meet-data-residency-requirements-with-google-cloud>

**NEW QUESTION 7**

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 8**

Your organization has strict policies on tracking rollouts to production and periodically shares this information with external auditors to meet compliance requirements. You need to enable auditing on several Cloud Spanner databases. What should you do?

- A. Use replication to roll out changes to higher environments.
- B. Use backup and restore to roll out changes to higher environments.
- C. Use Liquibase to roll out changes to higher environments.
- D. Manually capture detailed DBA audit logs when changes are rolled out to higher environments.

**Answer:** C

**Explanation:**

To satisfy audit reporting you would need a way to record what was changed and when. The best answer is one which uses some kind of source code control system (SCCS). That rules out A and B. Any mention of anything manual in a cloud environment should look suspicious, which leave option C. As it happens, Liquibase is an SCCS and can be integrated with Spanner. <https://cloud.google.com/spanner/docs/use-liquibase>

**NEW QUESTION 9**

You work in the logistics department. Your data analysis team needs daily extracts from Cloud SQL for MySQL to train a machine learning model. The model will be used to optimize next-day routes. You need to export the data in CSV format. You want to follow Google-recommended practices. What should you do?

- A. Use Cloud Scheduler to trigger a Cloud Function that will run a select \* from table(s) query to call the cloudsql.instances.export API.
- B. Use Cloud Scheduler to trigger a Cloud Function through Pub/Sub to call the cloudsql.instances.export API.
- C. Use Cloud Composer to orchestrate an export by calling the cloudsql.instances.export API.
- D. Use Cloud Composer to execute a select \* from table(s) query and export results.

**Answer:** B

**Explanation:**

<https://cloud.google.com/blog/topics/developers-practitioners/scheduling-cloud-sql-exports-using-cloud-functions-and-cloud-scheduler>

**NEW QUESTION 10**

You are designing a database strategy for a new web application. You plan to start with a small pilot in one country and eventually expand to millions of users in a global audience. You need to ensure that the application can run 24/7 with minimal downtime for maintenance. What should you do?

- A. Use Cloud Spanner in a regional configuration.
- B. Use Cloud Spanner in a multi-region configuration.
- C. Use Cloud SQL with cross-region replicas.

D. Use highly available Cloud SQL with multiple zones.

**Answer:** A

**Explanation:**

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

#### NEW QUESTION 10

Your company uses the Cloud SQL out-of-disk recommender to analyze the storage utilization trends of production databases over the last 30 days. Your database operations team uses these recommendations to proactively monitor storage utilization and implement corrective actions. You receive a recommendation that the instance is likely to run out of disk space. What should you do to address this storage alert?

- A. Normalize the database to the third normal form.
- B. Compress the data using a different compression algorithm.
- C. Manually or automatically increase the storage capacity.
- D. Create another schema to load older data.

**Answer:** C

**Explanation:**

<https://cloud.google.com/sql/docs/mysql/instance-settings#storage-capacity-2ndgen>

#### NEW QUESTION 14

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 18

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 21

You want to migrate an existing on-premises application to Google Cloud. Your application supports semi-structured data ingested from 100,000 sensors, and each sensor sends 10 readings per second from manufacturing plants. You need to make this data available for real-time monitoring and analysis. What should you do?

- A. Deploy the database using Cloud SQL.
- B. Use BigQuery, and load data in batches.
- C. Deploy the database using Bigtable.
- D. Deploy the database using Cloud Spanner.

**Answer:** C

**Explanation:**

Bigtable is a scalable, fully managed, and high-performance NoSQL database service that can handle semi-structured data and support real-time monitoring and analysis. Cloud SQL is a relational database service that does not support semi-structured data. BigQuery is a data warehouse service that is optimized for batch processing and analytics, not real-time monitoring. Cloud Spanner is a relational database service that supports semi-structured data with JSON data type, but it is more expensive and complex than Bigtable for this use case.

#### NEW QUESTION 25

You want to migrate an on-premises 100 TB Microsoft SQL Server database to Google Cloud over a 1 Gbps network link. You have 48 hours allowed downtime to migrate this database. What should you do? (Choose two.)

- A. Use a change data capture (CDC) migration strategy.
- B. Move the physical database servers from on-premises to Google Cloud.
- C. Keep the network bandwidth at 1 Gbps, and then perform an offline data migration.

- D. Increase the network bandwidth to 2 Gbps, and then perform an offline data migration.
- E. Increase the network bandwidth to 10 Gbps, and then perform an offline data migration.

**Answer:** AE

**Explanation:**

[https://cloud.google.com/architecture/migration-to-google-cloud-transferring-your-large-datasets#online\\_versus\\_offline\\_transfer](https://cloud.google.com/architecture/migration-to-google-cloud-transferring-your-large-datasets#online_versus_offline_transfer)

#### NEW QUESTION 28

Your company has PostgreSQL databases on-premises and on Amazon Web Services (AWS). You are planning multiple database migrations to Cloud SQL in an effort to reduce costs and downtime. You want to follow Google-recommended practices and use Google native data migration tools. You also want to closely monitor the migrations as part of the cutover strategy. What should you do?

- A. Use Database Migration Service to migrate all databases to Cloud SQL.
- B. Use Database Migration Service for one-time migrations, and use third-party or partner tools for change data capture (CDC) style migrations.
- C. Use data replication tools and CDC tools to enable migration.
- D. Use a combination of Database Migration Service and partner tools to support the data migration strategy.

**Answer:** A

**Explanation:**

<https://cloud.google.com/blog/products/databases/tips-for-migrating-across-compatible-database-engines>

#### NEW QUESTION 29

You are designing a new gaming application that uses a highly transactional relational database to store player authentication and inventory data in Google Cloud. You want to launch the game in multiple regions. What should you do?

- A. Use Cloud Spanner to deploy the database.
- B. Use Bigtable with clusters in multiple regions to deploy the database
- C. Use BigQuery to deploy the database
- D. Use Cloud SQL with a regional read replica to deploy the database.

**Answer:** A

**Explanation:**

Cloud Spanner is a fully managed, mission-critical, relational database service that offers transactional consistency at global scale, automatic, synchronous replication for high availability, and support for two SQL dialects: Google Standard SQL (ANSI 2011 with extensions) and PostgreSQL.

#### NEW QUESTION 33

You plan to use Database Migration Service to migrate data from a PostgreSQL on- premises instance to Cloud SQL. You need to identify the prerequisites for creating and automating the task. What should you do? (Choose two.)

- A. Drop or disable all users except database administration users.
- B. Disable all foreign key constraints on the source PostgreSQL database.
- C. Ensure that all PostgreSQL tables have a primary key.
- D. Shut down the database before the Data Migration Service task is started.
- E. Ensure that pglogical is installed on the source PostgreSQL database.

**Answer:** CE

**Explanation:**

<https://cloud.google.com/database-migration/docs/postgres/faq>

#### NEW QUESTION 36

You are setting up a Bare Metal Solution environment. You need to update the operating system to the latest version. You need to connect the Bare Metal Solution environment to the internet so you can receive software updates. What should you do?

- A. Setup a static external IP address in your VPC network.
- B. Set up bring your own IP (BYOIP) in your VPC.
- C. Set up a Cloud NAT gateway on the Compute Engine VM.
- D. Set up Cloud NAT service.

**Answer:** C

**Explanation:**

<https://cloud.google.com/bare-metal/docs/bms-setup?hl=en#bms-access-internet-vm-nat> The docs specifically says "Setting up a NAT gateway on a Compute Engine VM" is the way to give BMS internet access.

#### NEW QUESTION 40

Your company wants to move to Google Cloud. Your current data center is closing in six months. You are running a large, highly transactional Oracle application footprint on VMWare. You need to design a solution with minimal disruption to the current architecture and provide ease of migration to Google Cloud. What should you do?

- A. Migrate applications and Oracle databases to Google Cloud VMware Engine (VMware Engine).
- B. Migrate applications and Oracle databases to Compute Engine.
- C. Migrate applications to Cloud SQL.



D. Migrate applications and Oracle databases to Google Kubernetes Engine (GKE).

**Answer:** A

**Explanation:**

<https://cloud.google.com/blog/products/databases/migrate-databases-to-google-cloud-vmware-engine-gcve>

#### NEW QUESTION 44

Your company is evaluating Google Cloud database options for a mission-critical global payments gateway application. The application must be available 24/7 to users worldwide, horizontally scalable, and support open source databases. You need to select an automatically shardable, fully managed database with 99.999% availability and strong transactional consistency. What should you do?

- A. Select Bare Metal Solution for Oracle.
- B. Select Cloud SQL.
- C. Select Bigtable.
- D. Select Cloud Spanner.

**Answer:** D

**Explanation:**

The application must be available 24/7 to users worldwide, horizontally scalable, and support open source databases.

#### NEW QUESTION 47

You have a Cloud SQL instance (DB-1) with two cross-region read replicas (DB-2 and DB-3). During a business continuity test, the primary instance (DB-1) was taken offline and a replica (DB-2) was promoted. The test has concluded and you want to return to the pre-test configuration. What should you do?

- A. Bring DB-1 back online.
- B. Delete DB-1, and re-create DB-1 as a read replica in the same region as DB-1.
- C. Delete DB-2 so that DB-1 automatically reverts to the primary instance.
- D. Create DB-4 as a read replica in the same region as DB-1, and promote DB-4 to primary.

**Answer:** D

**Explanation:**

If you need to have the primary instance in the zone that had the outage, you can do a failback. A failback performs the same steps as the failover, only in the opposite direction, to reroute traffic back to the original instance. To perform a failback, use the procedure in Initiating failover. <https://cloud.google.com/sql/docs/mysql/high-availability#failback>

#### NEW QUESTION 51

Your application follows a microservices architecture and uses a single large Cloud SQL instance, which is starting to have performance issues as your application grows. In the Cloud Monitoring dashboard, the CPU utilization looks normal. You want to follow Google- recommended practices to resolve and prevent these performance issues while avoiding any major refactoring. What should you do?

- A. Use Cloud Spanner instead of Cloud SQL.
- B. Increase the number of CPUs for your instance.
- C. Increase the storage size for the instance.
- D. Use many smaller Cloud SQL instances.

**Answer:** D

**Explanation:**

<https://cloud.google.com/sql/docs/mysql/best-practices#data-arch>

#### NEW QUESTION 56

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 59

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 63**

You are running an instance of Cloud Spanner as the backend of your ecommerce website. You learn that the quality assurance (QA) team has doubled the number of their test cases. You need to create a copy of your Cloud Spanner database in a new test environment to accommodate the additional test cases. You want to follow Google-recommended practices. What should you do?

- A. Use Cloud Functions to run the export in Avro format.
- B. Use Cloud Functions to run the export in text format.
- C. Use Dataflow to run the export in Avro format.
- D. Use Dataflow to run the export in text format.

**Answer:** C

**Explanation:**

<https://cloud.google.com/spanner/docs/import-export-overview#file-format>

**NEW QUESTION 65**

Your organization is running a low-latency reporting application on Microsoft SQL Server. In addition to the database engine, you are using SQL Server Analysis Services (SSAS), SQL Server Reporting Services (SSRS), and SQL Server Integration Services (SSIS) in your on-premises environment. You want to migrate your Microsoft SQL Server database instances to Google Cloud. You need to ensure minimal disruption to the existing architecture during migration. What should you do?

- A. Migrate to Cloud SQL for SQL Server.
- B. Migrate to Cloud SQL for PostgreSQL.
- C. Migrate to Compute Engine.
- D. Migrate to Google Kubernetes Engine (GKE).

**Answer:** C

**Explanation:**

<https://cloud.google.com/sql/docs/sqlserver/features>

**NEW QUESTION 67**

You want to migrate your on-premises PostgreSQL database to Compute Engine. You need to migrate this database with the minimum downtime possible. What should you do?

- A. Perform a full backup of your on-premises PostgreSQL, and then, in the migration window, perform an incremental backup.
- B. Create a read replica on Cloud SQL, and then promote it to a read/write standalone instance.
- C. Use Database Migration Service to migrate your database.
- D. Create a hot standby on Compute Engine, and use PgBouncer to switch over the connections.

**Answer:** D

**Explanation:**

PgBouncer maintains a pool for connections for each database and user combination. PgBouncer either creates a new database connection for a client or reuses an existing connection for the same user and database. + PgBouncer is a simple PostgreSQL connection pool that allows for several thousand connections at a time. Using Kubernetes Engine to run a Helm Chart w/ PgBouncer based on the great article from futuretech- industries, we were able to set up an easily deployable system to get the most out of our CloudSQL DBs without breaking the bank. <https://medium.com/google-cloud/increasing-cloud-sql-postgresql-max-connections-w-pgbouncer-kubernetes-engine-49b0b2894820#:~:text=That%20is%20where,breaking%20the%20bank>.

**NEW QUESTION 70**

You are building an Android game that needs to store data on a Google Cloud serverless database. The database will log user activity, store user preferences, and receive in-game updates. The target audience resides in developing countries that have intermittent internet connectivity. You need to ensure that the game can synchronize game data to the backend database whenever an internet network is available. What should you do?

- A. Use Firestore.
- B. Use Cloud SQL with an external (public) IP address.
- C. Use an in-app embedded database.
- D. Use Cloud Spanner.

**Answer:** A

**Explanation:**

<https://firebase.google.com/docs/firestore>

**NEW QUESTION 71**

Your organization deployed a new version of a critical application that uses Cloud SQL for MySQL with high availability (HA) and binary logging enabled to store transactional information. The latest release of the application had an error that caused massive data corruption in your Cloud SQL for MySQL database. You need to minimize data loss. What should you do?

- A. Open the Google Cloud Console, navigate to SQL > Backups, and select the last version of the automated backup before the corruption.
- B. Reload the Cloud SQL for MySQL database using the LOAD DATA command to load data from CSV files that were used to initialize the instance.
- C. Perform a point-in-time recovery of your Cloud SQL for MySQL database, selecting a date and time before the data was corrupted.
- D. Fail over to the Cloud SQL for MySQL HA instance.
- E. Use that instance to recover the transactions that occurred before the corruption.

**Answer:** C

**Explanation:**

Binary Logging enabled, with that you can identify the point of time the data was good and recover from that point time.  
[https://cloud.google.com/sql/docs/mysql/backup-recovery/pitr#perform\\_the\\_point-in-time\\_recovery\\_using\\_binary\\_log\\_positions](https://cloud.google.com/sql/docs/mysql/backup-recovery/pitr#perform_the_point-in-time_recovery_using_binary_log_positions)

**NEW QUESTION 74**

You are managing a set of Cloud SQL databases in Google Cloud. Regulations require that database backups reside in the region where the database is created. You want to minimize operational costs and administrative effort. What should you do?

- A. Configure the automated backups to use a regional Cloud Storage bucket as a custom location.
- B. Use the default configuration for the automated backups location.
- C. Disable automated backups, and create an on-demand backup routine to a regional Cloud Storage bucket.
- D. Disable automated backups, and configure serverless exports to a regional Cloud Storage bucket.

**Answer:** A

**Explanation:**

<https://cloud.google.com/sql/docs/mysql/backup-recovery/backing-up#locationbackups> You can use a custom location for on-demand and automatic backups. For a complete list of valid location values, see the Instance locations.

**NEW QUESTION 79**

You are the DBA of an online tutoring application that runs on a Cloud SQL for PostgreSQL database. You are testing the implementation of the cross-regional failover configuration. The database in region R1 fails over successfully to region R2, and the database becomes available for the application to process data. During testing, certain scenarios of the application work as expected in region R2, but a few scenarios fail with database errors. The application-related database queries, when executed in isolation from Cloud SQL for PostgreSQL in region R2, work as expected. The application performs completely as expected when the database fails back to region R1. You need to identify the cause of the database errors in region R2. What should you do?

- A. Determine whether the versions of Cloud SQL for PostgreSQL in regions R1 and R2 are different.
- B. Determine whether the database patches of Cloud SQL for PostgreSQL in regions R1 and R2 are different.
- C. Determine whether the failover of Cloud SQL for PostgreSQL from region R1 to region R2 is in progress or has completed successfully.
- D. Determine whether Cloud SQL for PostgreSQL in region R2 is a near-real-time copy of region R1 but not an exact copy.

**Answer:** D

**Explanation:**

Verify that the replica has processed all the transactions it has received from the primary. This ensures that when promoted, the replica reflects all transactions that were received before the primary became unavailable. [https://cloud.google.com/sql/docs/postgres/replication/cross-region-replicas#verify\\_failover\\_criteria](https://cloud.google.com/sql/docs/postgres/replication/cross-region-replicas#verify_failover_criteria)

**NEW QUESTION 80**

You finished migrating an on-premises MySQL database to Cloud SQL. You want to ensure that the daily export of a table, which was previously a cron job running on the database server, continues. You want the solution to minimize cost and operations overhead. What should you do?

- A. Use Cloud Scheduler and Cloud Functions to run the daily export.
- B. Create a streaming Datatlow job to export the table.
- C. Set up Cloud Composer, and create a task to export the table daily.
- D. Run the cron job on a Compute Engine instance to continue the export.

**Answer:** A

**Explanation:**

<https://cloud.google.com/blog/topics/developers-practitioners/scheduling-cloud-sql-exports-using-cloud-functions-and-cloud-scheduler>

**NEW QUESTION 83**

You are building a data warehouse on BigQuery. Sources of data include several MySQL databases located on-premises. You need to transfer data from these databases into BigQuery for analytics. You want to use a managed solution that has low latency and is easy to set up. What should you do?

- A. Create extracts from your on-premises databases periodically, and push these extracts to Cloud Storage. Upload the changes into BigQuery, and merge them with existing tables.
- B. Use Cloud Data Fusion and scheduled workflows to extract data from MySQL.
- C. Transform this data into the appropriate schema, and load this data into your BigQuery database.
- D. Use Datastream to connect to your on-premises database and create a stream.
- E. Have Datastream write to Cloud Storage.
- F. Then use Dataflow to process the data into BigQuery.
- G. Use Database Migration Service to replicate data to a Cloud SQL for MySQL instance.
- H. Create federated tables in BigQuery on top of the replicated instances to transform and load the data into your BigQuery database.

**Answer:** C

**NEW QUESTION 87**



Your retail organization is preparing for the holiday season. Use of catalog services is increasing, and your DevOps team is supporting the Cloud SQL databases that power a microservices-based application. The DevOps team has added instrumentation through Sqlcommenter. You need to identify the root cause of why certain microservice calls are failing. What should you do?

- A. Watch Query Insights for long running queries.
- B. Watch the Cloud SQL instance monitor for CPU utilization metrics.
- C. Watch the Cloud SQL recommenders for overprovisioned instances.
- D. Watch Cloud Trace for application requests that are failing.

**Answer:** A

**Explanation:**

Cloud Trace doesn't support Cloud SQL. Eliminate D. Cloud SQL recommenders for overprovisioned instances would tell you about Cloud SQL instances which are too large for their workload. Eliminate C. Monitoring CPU utilization wouldn't tell you why microservice calls are failing. Eliminate B. SQLcommenter integrates with Query Insights. So A is the best answer. <https://cloud.google.com/blog/topics/developers-practitioners/introducing-sqlcommenter-open-source-orm-auto-instrumentation-library>

**NEW QUESTION 92**

You are configuring the networking of a Cloud SQL instance. The only application that connects to this database resides on a Compute Engine VM in the same project as the Cloud SQL instance. The VM and the Cloud SQL instance both use the same VPC network, and both have an external (public) IP address and an internal (private) IP address. You want to improve network security. What should you do?

- A. Disable and remove the internal IP address assignment.
- B. Disable both the external IP address and the internal IP address, and instead rely on Private Google Access.
- C. Specify an authorized network with the CIDR range of the VM.
- D. Disable and remove the external IP address assignment.

**Answer:** D

**Explanation:**

It is always more secure to use an internal IP, so removing them doesn't make sense. Eliminate A. You can use Private Google Access when VM instances only have internal IP addresses, so disabling the internal IPs and use Private Google Access doesn't make sense. Eliminate B. Specifying an authorized network when they're on the same subnet doesn't make sense. Eliminate C. A way to improve network security would be to disable external IPs since they're not needed.

**NEW QUESTION 94**

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 97**

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 102**

Your organization has hundreds of Cloud SQL for MySQL instances. You want to follow Google-recommended practices to optimize platform costs. What should you do?

- A. Use Query Insights to identify idle instances.
- B. Remove inactive user accounts.

- C. Run the Recommender API to identify overprovisioned instances.
- D. Build indexes on heavily accessed tables.

**Answer:** C

**Explanation:**

The Cloud SQL overprovisioned instance recommender helps you detect instances that are unnecessarily large for a given workload. It then provides recommendations on how to resize such instances and reduce cost. This page describes how this recommender works and how to use it. <https://cloud.google.com/sql/docs/mysql/recommender-sql-overprovisioned#:~:text=The%20Cloud%20SQL%20overprovisioned%20instance%20recommender%20helps%20you%20detect%20instances%20that%20are%20unnecessarily%20large%20for%20a%20given%20workload.%20It%20then%20provides%20recommendations%20on%20how%20to%20resize%20such%20instances%20and%20reduce%20cost.%20This%20page%20describes%20how%20this%20recommender%20works%20and%20how%20to%20use%20it.>

**NEW QUESTION 107**

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 111**

Your organization has a critical business app that is running with a Cloud SQL for MySQL backend database. Your company wants to build the most fault-tolerant and highly available solution possible. You need to ensure that the application database can survive a zonal and regional failure with a primary region of us-central1 and the backup region of us-east1. What should you do?

- A. Provision a Cloud SQL for MySQL instance in us-central1-
- B. Create a multiple-zone instance in us-west1-b.Create a read replica in us-east1-c.
- C. Provision a Cloud SQL for MySQL instance in us-central1-
- D. Create a multiple-zone instance in us-central1-b.Create a read replica in us-east1-b.
- E. Provision a Cloud SQL for MySQL instance in us-central1-
- F. Create a multiple-zone instance in us-east-b.Create a read replica in us-east1-c.
- G. Provision a Cloud SQL for MySQL instance in us-central1-
- H. Create a multiple-zone instance in us-east1-b.Create a read replica in us-central1-b.

**Answer:** B

**Explanation:**

<https://cloud.google.com/sql/docs/sqlserver/intro-to-cloud-sql-disaster-recovery>

**NEW QUESTION 113**

Your organization is migrating 50 TB Oracle databases to Bare Metal Solution for Oracle. Database backups must be available for quick restore. You also need to have backups available for 5 years. You need to design a cost-effective architecture that meets a recovery time objective (RTO) of 2 hours and recovery point objective (RPO) of 15 minutes. What should you do?

- A. Create the database on a Bare Metal Solution server with the database running on flash storage.Keep a local backup copy on all flash storage.Keep backups older than one day stored in Actifio OnVault storage.
- B. Create the database on a Bare Metal Solution server with the database running on flash storage.Keep a local backup copy on standard storage.Keep backups older than one day stored in Actifio OnVault storage.
- C. Create the database on a Bare Metal Solution server with the database running on flash storage.Keep a local backup copy on standard storage.Use the Oracle Recovery Manager (RMAN) backup utility to move backups older than oneday to a Coldline Storage bucket.
- D. Create the database on a Bare Metal Solution server with the database running on flash storage.Keep a local backup copy on all flash storage.Use the Oracle Recovery Manager (RMAN) backup utility to move backups older than one day to an Archive Storage bucket.

**Answer:** B

**Explanation:**

This answer meets the RTO and RPO requirements by using flash storage for the database and standard storage for the local backup copy. It also meets the cost-effectiveness requirement by using Actifio OnVault storage, which is a low-cost, high-performance, and scalable storage solution that integrates with Google Cloud Backup and DR Service1.

References := 1: Solution Guide: Google Cloud Backup and DR for Oracle on Bare Metal Solution1

**NEW QUESTION 116**

You have a large Cloud SQL for PostgreSQL instance. The database instance is not mission-critical, and you want to minimize operational costs. What should you do to lower the cost of backups in this environment?

- A. Set the automated backups to occur every other day to lower the frequency of backups.
- B. Change the storage tier of the automated backups from solid-state drive (SSD) to hard disk drive (HDD).

- C. Select a different region to store your backups.
- D. Reduce the number of automated backups that are retained to two (2).

**Answer:** D

**Explanation:**

By default, for each instance, Cloud SQL retains seven automated backups, in addition to on-demand backups. You can configure how many automated backups to retain (from 1 to 365). We charge a lower rate for backup storage than for other types of instances. <https://cloud.google.com/sql/docs/mysql/backup-recovery/backups>

**NEW QUESTION 118**

Your team recently released a new version of a highly consumed application to accommodate additional user traffic. Shortly after the release, you received an alert from your production monitoring team that there is consistently high replication lag between your primary instance and the read replicas of your Cloud SQL for MySQL instances. You need to resolve the replication lag. What should you do?

- A. Identify and optimize slow running queries, or set parallel replication flags.
- B. Stop all running queries, and re-create the replicas.
- C. Edit the primary instance to upgrade to a larger disk, and increase vCPU count.
- D. Edit the primary instance to add additional memory.

**Answer:** A

**Explanation:**

[https://cloud.google.com/sql/docs/mysql/replication/replication-lag#optimize\\_queries\\_and\\_schema](https://cloud.google.com/sql/docs/mysql/replication/replication-lag#optimize_queries_and_schema)

**NEW QUESTION 123**

Your organization has a ticketing system that needs an online marketing analytics and reporting application. You need to select a relational database that can manage hundreds of terabytes of data to support this new application. Which database should you use?

- A. Cloud SQL
- B. BigQuery
- C. Cloud Spanner
- D. Bigtable

**Answer:** B

**NEW QUESTION 124**

You are managing multiple applications connecting to a database on Cloud SQL for PostgreSQL. You need to be able to monitor database performance to easily identify applications with long-running and resource-intensive queries. What should you do?

- A. Use log messages produced by Cloud SQL.
- B. Use Query Insights for Cloud SQL.
- C. Use the Cloud Monitoring dashboard with available metrics from Cloud SQL.
- D. Use Cloud SQL instance monitoring in the Google Cloud Console.

**Answer:** B

**Explanation:**

<https://cloud.google.com/sql/docs/mysql/using-query-insights#introduction>

**NEW QUESTION 129**

You are managing a Cloud SQL for MySQL environment in Google Cloud. You have deployed a primary instance in Zone A and a read replica instance in Zone B, both in the same region. You are notified that the replica instance in Zone B was unavailable for 10 minutes. You need to ensure that the read replica instance is still working. What should you do?

- A. Use the Google Cloud Console or gcloud CLI to manually create a new clone database.
- B. Use the Google Cloud Console or gcloud CLI to manually create a new failover replica from backup.
- C. Verify that the new replica is created automatically.
- D. Start the original primary instance and resume replication.

**Answer:** C

**Explanation:**

Recovery Process: Once Zone-B becomes available again, Cloud SQL will initiate the recovery process for the impacted read replica. The recovery process involves the following steps: 1. Synchronization: Cloud SQL will compare the data in the recovered read replica with the primary instance in Zone-A. If there is any data divergence due to the unavailability period, Cloud SQL will synchronize the read replica with the primary instance to ensure data consistency. 2. Catch-up Replication: The recovered read replica will start catching up on the changes that occurred on the primary instance during its unavailability. It will apply the necessary updates from the primary instance's binary logs (binlogs) to bring the replica up to date. 3. Resuming Read Traffic: Once the synchronization and catch-up replication processes are complete, the read replica in Zone-B will resume its normal operation. It will be able to serve read traffic and stay updated with subsequent changes from the primary instance.

**NEW QUESTION 130**

Your organization is running a critical production database on a virtual machine (VM) on Compute Engine. The VM has an ext4-formatted persistent disk for data files. The database will soon run out of storage space. You need to implement a solution that avoids downtime. What should you do?

- A. In the Google Cloud Console, increase the size of the persistent disk, and use the `resize2fs` command to extend the disk.
- B. In the Google Cloud Console, increase the size of the persistent disk, and use the `fdisk` command to verify that the new space is ready to use

- C. In the Google Cloud Console, create a snapshot of the persistent disk, restore the snapshot to a new larger disk, unmount the old disk, mount the new disk, and restart the database service.
- D. In the Google Cloud Console, create a new persistent disk attached to the VM, and configure the database service to move the files to the new disk.

**Answer:** A

**Explanation:**

[https://cloud.google.com/compute/docs/disks/resize-persistent-disk#resize\\_partitions](https://cloud.google.com/compute/docs/disks/resize-persistent-disk#resize_partitions)

#### NEW QUESTION 131

You are running a mission-critical application on a Cloud SQL for PostgreSQL database with a multi-zonal setup. The primary and read replica instances are in the same region but in different zones. You need to ensure that you split the application load between both instances. What should you do?

- A. Use Cloud Load Balancing for load balancing between the Cloud SQL primary and read replica instances.
- B. Use PgBouncer to set up database connection pooling between the Cloud SQL primary and read replica instances.
- C. Use HTTP(S) Load Balancing for database connection pooling between the Cloud SQL primary and read replica instances.
- D. Use the Cloud SQL Auth proxy for database connection pooling between the Cloud SQL primary and read replica instances.

**Answer:** B

**Explanation:**

<https://severalnines.com/blog/how-achieve-postgresql-high-availability-pgbouncer/>

<https://cloud.google.com/blog/products/databases/using-haproxy-to-scale-read-only-workloads-on-cloud-sql-for-postgresql>

This answer is correct because PgBouncer is a lightweight connection pooler for PostgreSQL that can help you distribute read requests between the Cloud SQL primary and read replica instances<sup>1</sup>. PgBouncer can also improve performance and scalability by reducing the overhead of creating new connections and reusing existing ones<sup>1</sup>. You can install PgBouncer on a Compute Engine instance and configure it to connect to the Cloud SQL instances using private IP addresses or the Cloud SQL Auth proxy<sup>2</sup>.

#### NEW QUESTION 135

You are choosing a database backend for a new application. The application will ingest data points from IoT sensors. You need to ensure that the application can scale up to millions of requests per second with sub-10ms latency and store up to 100 TB of history. What should you do?

- A. Use Cloud SQL with read replicas for throughput.
- B. Use Firestore, and rely on automatic serverless scaling.
- C. Use Memorystore for Memcached, and add nodes as necessary to achieve the required throughput.
- D. Use Bigtable, and add nodes as necessary to achieve the required throughput.

**Answer:** D

**Explanation:**

<https://cloud.google.com/memorystore/docs/redis/redis-overview>

#### NEW QUESTION 139

You are migrating an on-premises application to Compute Engine and Cloud SQL. The application VMs will live in their own project, separate from the Cloud SQL instances which have their own project. What should you do to configure the networks?

- A. Create a new VPC network in each project, and use VPC Network Peering to connect the two together.
- B. Create a Shared VPC that both the application VMs and Cloud SQL instances will use.
- C. Use the default networks, and leverage Cloud VPN to connect the two together.
- D. Place both the application VMs and the Cloud SQL instances in the default network of each project.

**Answer:** B

**Explanation:**

[https://groups.google.com/g/google-cloud-sql-discuss/c/M5G5\\_HPXytY?pli=1](https://groups.google.com/g/google-cloud-sql-discuss/c/M5G5_HPXytY?pli=1)

#### NEW QUESTION 140

Your hotel booking company is expanding into Country A, where personally identifiable information (PII) must comply with regional data residency requirements and audits. You need to isolate customer data in Country A from the rest of the customer data. You want to design a multi-tenancy strategy to efficiently manage costs and operations. What should you do?

- A. Apply a schema data management pattern.
- B. Apply an instance data management pattern.
- C. Apply a table data management pattern.
- D. Apply a database data management pattern.

**Answer:** B

**Explanation:**

<https://cloud.google.com/solutions/implementing-multi-tenancy-cloud-spanner#multi-tenancy-data-management-patterns>

<https://cloud.google.com/solutions/implementing-multi-tenancy-cloud-spanner>

#### NEW QUESTION 143

You are developing a new application on a VM that is on your corporate network. The application will use Java Database Connectivity (JDBC) to connect to Cloud SQL for PostgreSQL. Your Cloud SQL instance is configured with IP address 192.168.3.48, and SSL is disabled. You want to ensure that your application can access your database instance without requiring configuration changes to your database. What should you do?



- A. Define a connection string using your Google username and password to point to the external (public) IP address of your Cloud SQL instance.
- B. Define a connection string using a database username and password to point to the internal (private) IP address of your Cloud SQL instance.
- C. Define a connection string using Cloud SQL Auth proxy configured with a service account to point to the internal (private) IP address of your Cloud SQL instance.
- D. Define a connection string using Cloud SQL Auth proxy configured with a service account to point to the external (public) IP address of your Cloud SQL instance.

**Answer:** C

**Explanation:**

The Cloud SQL connectors are libraries that provide encryption and IAM- based authorization when connecting to a Cloud SQL instance. They can't provide a network path to a Cloud SQL instance if one is not already present. Other ways to connect to a Cloud SQL instance include using a database client or the Cloud SQL Auth proxy. <https://cloud.google.com/sql/docs/postgres/connect-connectors> <https://github.com/GoogleCloudPlatform/cloud-sql-jdbc-socket-factory/blob/main/docs/jdbc-postgres.md>

**NEW QUESTION 147**

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 148**

Your company's mission-critical, globally available application is supported by a Cloud Spanner database. Experienced users of the application have read and write access to the database, but new users are assigned read-only access to the database. You need to assign the appropriate Cloud Spanner Identity and Access Management (IAM) role to new users being onboarded soon. What roles should you set up?

- A. roles/spanner.databaseReader
- B. roles/spanner.databaseUser
- C. roles/spanner.viewer
- D. roles/spanner.backupWriter

**Answer:** A

**Explanation:**

<https://cloud.google.com/spanner/docs/iam?hl=it>

**NEW QUESTION 150**

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 152**

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 157**

Your company is shutting down their data center and migrating several MySQL and PostgreSQL databases to Google Cloud. Your database operations team is

severely constrained by ongoing production releases and the lack of capacity for additional on- premises backups. You want to ensure that the scheduled migrations happen with minimal downtime and that the Google Cloud databases stay in sync with the on-premises data changes until the applications can cut over. What should you do? (Choose two.)

- A. Use an external read replica to migrate the databases to Cloud SQL.
- B. Use a read replica to migrate the databases to Cloud SQL.
- C. Use Database Migration Service to migrate the databases to Cloud SQL.
- D. Use a cross-region read replica to migrate the databases to Cloud SQL.
- E. Use replication from an external server to migrate the databases to Cloud SQL.

**Answer:** CE

#### NEW QUESTION 161

You work for a financial services company that wants to use fully managed database services. Traffic volume for your consumer services products has increased annually at a constant rate with occasional spikes around holidays. You frequently need to upgrade the capacity of your database. You want to use Cloud Spanner and include an automated method to increase your hardware capacity to support a higher level of concurrency. What should you do?

- A. Use linear scaling to implement the Autoscaler-based architecture
- B. Use direct scaling to implement the Autoscaler-based architecture.
- C. Upgrade the Cloud Spanner instance on a periodic basis during the scheduled maintenance window.
- D. Set up alerts that are triggered when Cloud Spanner utilization metrics breach the threshold, and then schedule an upgrade during the scheduled maintenance window.

**Answer:** A

#### Explanation:

Linear scaling is best used with load patterns that change more gradually or have a few large peaks. The method calculates the minimum number of nodes or processing units required to keep utilization below the scaling threshold. The number of nodes or processing units added or removed in each scaling event is not limited to a fixed step amount. <https://cloud.google.com/spanner/docs/autoscaling-overview#linear>

#### NEW QUESTION 162

Your customer has a global chat application that uses a multi-regional Cloud Spanner instance. The application has recently experienced degraded performance after a new version of the application was launched. Your customer asked you for assistance. During initial troubleshooting, you observed high read latency. What should you do?

- A. Use query parameters to speed up frequently executed queries.
- B. Change the Cloud Spanner configuration from multi-region to single region.
- C. Use SQL statements to analyze SPANNER\_SYS.READ\_STATS\* tables.
- D. Use SQL statements to analyze SPANNER\_SYS.QUERY\_STATS\* tables.

**Answer:** C

#### Explanation:

To troubleshoot high read latency, you can use SQL statements to analyze the SPANNER\_SYS.READ\_STATS\* tables. These tables contain statistics about read operations in Cloud Spanner, including the number of reads, read latency, and the number of read errors. By analyzing these tables, you can identify the cause of the high read latency and take appropriate action to resolve the issue. Other options, such as using query parameters to speed up frequently executed queries or changing the Cloud Spanner configuration from multi-region to single region, may not be directly related to the issue of high read latency. Similarly, analyzing the SPANNER\_SYS.QUERY\_STATS\* tables, which contain statistics about query operations, may not be relevant to the issue of high read latency.

#### NEW QUESTION 164

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 169

Your organization has a busy transactional Cloud SQL for MySQL instance. Your analytics team needs access to the data so they can build monthly sales reports. You need to provide data access to the analytics team without adversely affecting performance. What should you do?

- A. Create a read replica of the database, provide the database IP address, username, and password to the analytics team, and grant read access to required tables to the team.
- B. Create a read replica of the database, enable the `cloudsql.iam_authentication` flag on the replica, and grant read access to required tables to the analytics team.
- C. Enable the `cloudsql.iam_authentication` flag on the primary database instance, and grant read access to required tables to the analytics team.
- D. Provide the database IP address, username, and password of the primary database instance to the analytics, team, and grant read access to required tables to the team.

**Answer:** B

#### Explanation:

"Read replicas do not have the cloudsql.iam\_authentication flag enabled automatically when it is enabled on the primary instance."  
[https://cloud.google.com/sql/docs/postgres/replication/create-replica#configure\\_iam\\_replicas](https://cloud.google.com/sql/docs/postgres/replication/create-replica#configure_iam_replicas)

#### NEW QUESTION 173

Your organization is currently updating an existing corporate application that is running in another public cloud to access managed database services in Google Cloud. The application will remain in the other public cloud while the database is migrated to Google Cloud. You want to follow Google-recommended practices for authentication. You need to minimize user disruption during the migration. What should you do?

- A. Use workload identity federation to impersonate a service account.
- B. Ask existing users to set their Google password to match their corporate password.
- C. Migrate the application to Google Cloud, and use Identity and Access Management (IAM).
- D. Use Google Workspace Password Sync to replicate passwords into Google Cloud.

**Answer:** A

#### Explanation:

Updating passwords represents user disruption. Eliminate B. Eliminate C for the same reason. D doesn't make sense, leaves A. From Google's documentation, "Traditionally, applications running outside Google Cloud can use service account keys to access Google Cloud resources. However, service account keys are powerful credentials, and can present a security risk if they are not managed correctly. With identity federation, you can use Identity and Access Management (IAM) to grant external identities IAM roles, including the ability to impersonate service accounts. This approach eliminates the maintenance and security burden associated with service account keys." <https://cloud.google.com/iam/docs/workload-identity-federation>

#### NEW QUESTION 178

.....

## Thank You for Trying Our Product

### We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

### 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 First Try
- \* Professional-Cloud-Database-Engineer Practice Test Questions in Multiple Choice Formats and Updates for 1 Year

**100% Actual & Verified — Instant Download, Please Click**  
**[Order The Professional-Cloud-Database-Engineer Practice Test Here](#)**