



**Google**

## **Exam Questions Professional-Cloud-Database-Engineer**

Google Cloud Certified - Professional Cloud Database Engineer

#### NEW QUESTION 1

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 2

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-instance>

#### NEW QUESTION 3

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 4

You need to redesign the architecture of an application that currently uses Cloud SQL for PostgreSQL. The users of the application complain about slow query response times. You want to enhance your application architecture to offer sub-millisecond query latency. What should you do?

- A. Configure Firestore, and modify your application to offload queries.
- B. Configure Bigtable, and modify your application to offload queries.
- C. Configure Cloud SQL for PostgreSQL read replicas to offload queries.
- D. Configure Memorystore, and modify your application to offload queries.

**Answer:** D

#### Explanation:

"sub-millisecond latency" always involves Memorystore. Furthermore, as we are talking about a relational DB (Cloud SQL), BigTable is not a solution to be considered.

#### NEW QUESTION 5

Your ecommerce application connecting to your Cloud SQL for SQL Server is expected to have additional traffic due to the holiday weekend. You want to follow Google- recommended practices to set up alerts for CPU and memory metrics so you can be notified by text message at the first sign of potential issues. What should you do?

- A. Use a Cloud Function to pull CPU and memory metrics from your Cloud SQL instance and to call a custom service to send alerts.
- B. Use Error Reporting to monitor CPU and memory metrics and to configure SMS notification channels.
- C. Use Cloud Logging to set up a log sink for CPU and memory metrics and to configure a sink destination to send a message to Pub/Sub.
- D. Use Cloud Monitoring to set up an alerting policy for CPU and memory metrics and to configure SMS notification channels.

**Answer:** D

#### Explanation:

Cloud Monitoring collects metrics, events, and metadata from Google Cloud, Amazon Web Services (AWS), hosted uptime probes, and application instrumentation. Using the BindPlane service, you can also collect this data from over 150 common application components, on-premise systems, and hybrid cloud systems.

#### NEW QUESTION 6

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 7

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 8

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 9

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 10

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 10

You are managing a Cloud SQL for PostgreSQL instance in Google Cloud. You need to test the high availability of your Cloud SQL instance by performing a failover. You want to use the cloud command. What should you do?

- A. Use `gcloud sql instances failover <PrimaryInstanceName>`.
- B. Use `gcloud sql instances failover <ReplicaInstanceName>`.
- C. Use `gcloud sql instances promote-replica <PrimaryInstanceName>`.

D. Use gcloud sql instances promote-replica <ReplicaInstanceName>.

**Answer:** A

#### NEW QUESTION 14

Your project is using Bigtable to store data that should not be accessed from the public internet under any circumstances, even if the requestor has a valid service account key. You need to secure access to this data. What should you do?

- A. Use Identity and Access Management (IAM) for Bigtable access control.
- B. Use VPC Service Controls to create a trusted network for the Bigtable service.
- C. Use customer-managed encryption keys (CMEK).
- D. Use Google Cloud Armor to add IP addresses to an allowlist.

**Answer:** B

#### Explanation:

"Users can define a security perimeter around Google Cloud resources such as Cloud Storage buckets, Bigtable instances, and BigQuery datasets to constrain data within a VPC and control the flow of data." <https://cloud.google.com/vpc-service-controls>

#### NEW QUESTION 16

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 20

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 21

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 25

Your company is using Cloud SQL for MySQL with an internal (private) IP address and wants to replicate some tables into BigQuery in near-real time for analytics and machine learning. You need to ensure that replication is fast and reliable and uses Google-managed services. What should you do?

- A. Develop a custom data replication service to send data into BigQuery.
- B. Use Cloud SQL federated queries.
- C. Use Database Migration Service to replicate tables into BigQuery.
- D. Use Datastream to capture changes, and use Dataflow to write those changes to BigQuery.

**Answer:** D

#### Explanation:

"Datastream is a serverless and easy-to-use Change Data Capture (CDC) and replication service that allows you to synchronize data across heterogeneous

databases, storage systems, and applications reliably and with minimal latency. Datastream supports change data streaming from Oracle and MySQL databases to Google Cloud Storage (GCS). The service offers streamlined integration with Dataflow templates to power up to date materialized views in BigQuery for analytics, replicate their databases into Cloud SQL or Cloud Spanner for database synchronization, or leverage the event stream directly from GCS to realize event-driven architectures.”

#### NEW QUESTION 26

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 30

You are building an application that allows users to customize their website and mobile experiences. The application will capture user information and preferences. User profiles have a dynamic schema, and users can add or delete information from their profile. You need to ensure that user changes automatically trigger updates to your downstream BigQuery data warehouse. What should you do?

- A. Store your data in Bigtable, and use the user identifier as the key.
- B. Use one column family to store user profile data, and use another column family to store user preferences.
- C. Use Cloud SQL, and create different tables for user profile data and user preferences from your recommendations mode.
- D. Use SQL to join the user profile data and preferences.
- E. Use Firestore in Native mode, and store user profile data as a document.
- F. Update the user profile with preferences specific to that user and use the user identifier to query.
- G. Use Firestore in Datastore mode, and store user profile data as a document.
- H. Update the user profile with preferences specific to that user and use the user identifier to query.

**Answer:** C

#### Explanation:

Use Firestore in Datastore mode for new server projects. Firestore in Datastore mode allows you to use established Datastore server architectures while removing fundamental Datastore limitations. Datastore mode can automatically scale to millions of writes per second. Use Firestore in Native mode for new mobile and web apps. Firestore offers mobile and web client libraries with real-time and offline features. Native mode can automatically scale to millions of concurrent clients.

#### NEW QUESTION 32

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 36

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 39

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 44**

Your DevOps team is using Terraform to deploy applications and Cloud SQL databases. After every new application change is rolled out, the environment is torn down and recreated, and the persistent database layer is lost. You need to prevent the database from being dropped. What should you do?

- A. Set Terraform deletion\_protection to true.
- B. Rerun terraform apply.
- C. Create a read replica.
- D. Use point-in-time-recovery (PITR) to recover the database.

**Answer:** A

**Explanation:**

From Google's documentation, "For stateful resources, such as databases, ensure that deletion protection is enabled. The syntax is: lifecycle { prevent\_destroy = true } <https://cloud.google.com/docs/terraform/best-practices-for-terraform#stateful-resources>

**NEW QUESTION 47**

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 48**

Your digital-native business runs its database workloads on Cloud SQL. Your website must be globally accessible 24/7. You need to prepare your Cloud SQL instance for high availability (HA). You want to follow Google-recommended practices. What should you do? (Choose two.)

- A. Set up manual backups.
- B. Create a PostgreSQL database on-premises as the HA option.
- C. Configure single zone availability for automated backups.
- D. Enable point-in-time recovery.
- E. Schedule automated backups.

**Answer:** DE

**Explanation:**

D. Enable point-in-time recovery - This feature allows you to restore your database to a specific point in time. It helps protect against data loss and can be used in the event of data corruption or accidental data deletion. E. Schedule automated backups - Automated backups allow you to take regular backups of your database without manual intervention. You can use these backups to restore your database in the event of data loss or corruption.

**NEW QUESTION 51**

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 53**

You recently launched a new product to the US market. You currently have two Bigtable clusters in one US region to serve all the traffic. Your marketing team is planning an immediate expansion to APAC. You need to roll out the regional expansion while implementing high availability according to Google-recommended practices. What should you do?

- A. Maintain a target of 23% CPU utilization by locating: cluster-a in zone us-central1-a cluster-b in zone europe-west1-d cluster-c in zone asia-east1-b
- B. Maintain a target of 23% CPU utilization by locating: cluster-a in zone us-central1-a cluster-b in zone us-central1-b cluster-c in zone us-east1-a C. Maintain a target of 35% CPU utilization by locating: cluster-a in zone us-central1-a cluster-b in zone australia-southeast1-a cluster-c in zone europe-west1-d cluster-d in zone asia-east1-b
- C. Maintain a target of 35% CPU utilization by locating: cluster-a in zone us-central1-a cluster-b in zone us-central2-a cluster-c in zone asia-northeast1-b cluster-d in zone asia-east1-b

**Answer:** D

**Explanation:**

<https://cloud.google.com/bigtable/docs/replication-settings#regional-failover>

**NEW QUESTION 54**

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

**NEW QUESTION 58**

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 63**

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 64**

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 68

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 71

Your team uses thousands of connected IoT devices to collect device maintenance data for your oil and gas customers in real time. You want to design inspection routines, device repair, and replacement schedules based on insights gathered from the data produced by these devices. You need a managed solution that is highly scalable, supports a multi-cloud strategy, and offers low latency for these IoT devices. What should you do?

- A. Use Firestore with Looker.
- B. Use Cloud Spanner with Data Studio.
- C. Use MongoDB Atlas with Charts.
- D. Use Bigtable with Looker.

**Answer:** C

#### Explanation:

This scenario has BigTable written all over it - large amounts of data from many devices to be analysed in realtime. I would even argue it could qualify as a multicloud solution, given the links to HBASE. BUT it does not support SQL queries and is not therefore compatible (on its own) with Looker. Firestore + Looker has the same problem. Spanner + Data Studio is at least a compatible pairing, but I agree with others that it doesn't fit this use-case - not least because it's Google-native. By contrast, MongoDB Atlas is a managed solution (just not by Google) which is compatible with the proposed reporting tool (Mongo's own Charts), it's specifically designed for this type of solution and of course it can run on any cloud.

#### NEW QUESTION 76

You need to perform a one-time migration of data from a running Cloud SQL for MySQL instance in the us-central1 region to a new Cloud SQL for MySQL instance in the us-east1 region. You want to follow Google-recommended practices to minimize performance impact on the currently running instance. What should you do?

- A. Create and run a Dataflow job that uses JdbcIO to copy data from one Cloud SQL instance to another.
- B. Create two Datastream connection profiles, and use them to create a stream from one Cloud SQL instance to another.
- C. Create a SQL dump file in Cloud Storage using a temporary instance, and then use that file to import into a new instance.
- D. Create a CSV file by running the SQL statement SELECT...INTO OUTFILE, copy the file to a Cloud Storage bucket, and import it into a new instance.

**Answer:** C

#### Explanation:

<https://cloud.google.com/sql/docs/mysql/import-export#serverless>

#### NEW QUESTION 77

You support a consumer inventory application that runs on a multi-region instance of Cloud Spanner. A customer opened a support ticket to complain about slow response times. You notice a Cloud Monitoring alert about high CPU utilization. You want to follow Google-recommended practices to address the CPU performance issue. What should you do first?

- A. Increase the number of processing units.
- B. Modify the database schema, and add additional indexes.
- C. Shard data required by the application into multiple instances.
- D. Decrease the number of processing units.

**Answer:** A

#### Explanation:

In case of high CPU utilization like, mentioned in question, refer: <https://cloud.google.com/spanner/docs/identify-latency-point#:~:text=Check%20the%20CPU%20utilization%20of%20the%20instance.%20If%20the%20CPU%20utilization%20of%20the%20instance%20is%20above%20the%20recommended%20level%2C%20you%20should%20manually%20add%20more%20nodes%2C%20or%20set%20up%20auto%20scaling.> "Check the CPU utilization of the instance. If the CPU utilization of the instance is above the recommended level, you should manually add more nodes, or set up auto scaling." Indexes and schema are reviewed post identifying query with slow performance. Refer : <https://cloud.google.com/spanner/docs/troubleshooting-performance-regressions#review-schema>

#### NEW QUESTION 81

You are designing a physician portal app in Node.js. This application will be used in hospitals and clinics that might have intermittent internet connectivity. If a connectivity failure occurs, the app should be able to query the cached data. You need to ensure that the application has scalability, strong consistency, and multi-region replication. What should you do?

- A. Use Firestore and ensure that the PersistenceEnabled option is set to true.
- B. Use Memorystore for Memcached.
- C. Use Pub/Sub to synchronize the changes from the application to Cloud Spanner.
- D. Use Table.read with the exactStaleness option to perform a read of rows in Cloud Spanner.

**Answer:** A

**Explanation:**

<https://firebase.google.com/docs/firestore/manage-data/enable-offline>

**NEW QUESTION 83**

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 87**

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 92**

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 95**

.....

## 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](#)