

Product Questions: 141

Version: 6.0

Question: 1

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>

Question: 2

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.

Question: 3

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>

Question: 4

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.

Question: 5

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>

Question: 6

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.

Question: 7

Your organization is running a MySQL workload in Cloud SQL. Suddenly you see a degradation in

database performance. You need to identify the root cause of the performance degradation. What should you do?

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

Answer: B

Explanation:

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

Question: 8

You work for a large retail and ecommerce company that is starting to extend their business globally. Your company plans to migrate to Google Cloud. You want to use platforms that will scale easily, handle transactions with the least amount of latency, and provide a reliable customer experience. You need a storage layer for sales transactions and current inventory levels. You want to retain the same relational schema that your existing platform uses. What should you do?

- A. Store your data in Firestore in a multi-region location, and place your compute resources in one of the constituent regions.
- B. Deploy Cloud Spanner using a multi-region instance, and place your compute resources close to the default leader region.
- C. Build an in-memory cache in Memorystore, and deploy to the specific geographic regions where your application resides.
- D. Deploy a Bigtable instance with a cluster in one region and a replica cluster in another geographic region.

Answer: B

Explanation:

Question: 9

You host an application in Google Cloud. The application is located in a single region and uses Cloud SQL for transactional data.

a. Most of your users are located in the same time zone and expect the application to be available 7 days a week, from 6 AM to 10 PM. You want to ensure regular maintenance updates to your Cloud SQL instance without creating downtime for your users. What should you do?

- A. Configure a maintenance window during a period when no users will be on the system. Control the order of update by setting non-production instances to earlier and production instances to later.
- B. Create your database with one primary node and one read replica in the region.

- C. Enable maintenance notifications for users, and reschedule maintenance activities to a specific time after notifications have been sent.
- D. Configure your Cloud SQL instance with high availability enabled.

Answer: A

Explanation:

Configure a maintenance window during a period when no users will be on the system. Control the order of update by setting non-production instances to earlier and production instances to later.

Question: 10

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

Question: 11

Your organization operates in a highly regulated industry. Separation of concerns (SoC) and security principle of least privilege (PoLP) are critical. The operations team consists of:

Person A is a database administrator.

Person B is an analyst who generates metric reports.

Application C is responsible for automatic backups.

You need to assign roles to team members for Cloud Spanner. Which roles should you assign?

- A.
 - roles/spanner.databaseAdmin for Person A
 - roles/spanner.databaseReader for Person B
 - roles/spanner.backupWriter for Application C
- B.
 - roles/spanner.databaseAdmin for Person A
 - roles/spanner.databaseReader for Person B
 - roles/spanner.backupAdmin for Application C
- C.
 - roles/spanner.databaseAdmin for Person A

roles/spanner.databaseUser for Person B
roles/spanner.databaseReader for Application C
D.
roles/spanner.databaseAdmin for Person A
roles/spanner.databaseUser for Person B
roles/spanner.backupWriter for Application C

Answer: A

Explanation:

<https://cloud.google.com/spanner/docs/iam#spanner.backupWriter>

Question: 12

You are designing an augmented reality game for iOS and Android devices. You plan to use Cloud Spanner as the primary backend database for game state storage and player authentication. You want to track in-game rewards that players unlock at every stage of the game. During the testing phase, you discovered that costs are much higher than anticipated, but the query response times are within the SL

- A. You want to follow Google-recommended practices. You need the database to be performant and highly available while you keep costs low. What should you do?
- A. Manually scale down the number of nodes after the peak period has passed.
- B. Use interleaving to co-locate parent and child rows.
- C. Use the Cloud Spanner query optimizer to determine the most efficient way to execute the SQL query.
- D. Use granular instance sizing in Cloud Spanner and Autoscaler.

Answer: D

Explanation:

Granular instance is available in Public Preview. With this feature, you can run workloads on Spanner at as low as 1/10th the cost of regular instances,

<https://cloud.google.com/blog/products/databases/get-more-out-of-spanner-with-granular-instance-sizing>

Question: 13

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

D.

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>

Question: 14

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. Use Datastream to export data to Cloud Storage and analyze with Dataproc and the Cloud Storage connector.
- C. Use Memorystore to handle your low-latency requirements and for real-time analytics.
- D. 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

Question: 15

Your company uses Cloud Spanner for a mission-critical inventory management system that is globally available. You recently loaded stock keeping unit (SKU) and product catalog data from a company acquisition and observed hot-spots in the Cloud Spanner database. You want to follow Google-recommended schema design practices to avoid performance degradation. What should you do? (Choose two.)

- A. Use an auto-incrementing value as the primary key.

- B. Normalize the data model.
- C. Promote low-cardinality attributes in multi-attribute primary keys.
- D. Promote high-cardinality attributes in multi-attribute primary keys.
- E. Use bit-reverse sequential value as the primary key.

Answer: DE

Explanation:

<https://cloud.google.com/spanner/docs/schema-design> D because high cardinality means you have more unique values in the column. That's a good thing for a hot-spotting issue. E because Spanner specifically has this feature to reduce hot spotting. Basically, it generates unique values https://cloud.google.com/spanner/docs/schema-design#bit_reverse_primary_key

D. Promote high-cardinality attributes in multi-attribute primary keys.

This is a correct answer because promoting high-cardinality attributes in multi-attribute primary keys can help avoid hotspots in Cloud Spanner. [High-cardinality attributes are those that have many distinct values, such as UUIDs, email addresses, or timestamps1. By placing high-cardinality attributes first in the primary key, you can ensure that the rows are distributed more evenly across the key space, and avoid having too many requests sent to the same server2.](#)

E. Use bit-reverse sequential value as the primary key.

This is a correct answer because using bit-reverse sequential value as the primary key can help avoid hotspots in Cloud Spanner. [Bit-reverse sequential value is a technique that reverses the bits of a monotonically increasing value, such as a timestamp or an auto-incrementing ID1. By reversing the bits, you can create a pseudo-random value that spreads the writes across the key space, and avoid having all the inserts occurring at the end of the table2.](#)

Question: 16

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>

Question: 17

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

Question: 18

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

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

Answer: B, C

Explanation:

Replication lag and slow report performance. E is eliminated because using BigQuery would mean changes to the current reports. Report slowness could be the result of poor indexing or just too much read load (or both!). Since excessive load is mentioned in the question, creating additional read replicas and spreading the analytics workload around makes B correct and eliminates A as a way to speed up reporting. That leaves the replication problem. Cloud SQL enables single threaded

replication by default, so it stands to reason enabling parallel replication would help the lag. To do that you disable replication on the replica (not the primary), set flags on the replica and optionally set flags on the primary instance to optimize performance for parallel replication. That makes C correct and D incorrect. <https://cloud.google.com/sql/docs/mysql/replication/manage-replicas#configuring-parallel-replication>

Question: 19

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 requirements. Keep the data that must adhere to data residency requirements on-premises. Make application changes to support both databases.
- C. 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.
- D. 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>

Question: 20

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>

Question: 21

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.

Question: 22

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>

Question: 23

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>

Question: 24

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 instances¹. PgBouncer can also improve performance and scalability by reducing the overhead of creating new connections and reusing existing ones¹. 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 proxy2.](#)

Question: 25

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. 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

Question: 26

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: C, E

Explanation:

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

Question: 27

You are using Compute Engine on Google Cloud and your data center to manage a set of MySQL databases in a hybrid configuration. You need to create replicas to scale reads and to offload part of the management operation. What should you do?

- A. Use external server replication.
- B. Use Data Migration Service.
- C. Use Cloud SQL for MySQL external replica.
- D. Use the mysqldump utility and binary logs.

Answer: C

Explanation:

[An external replica is a method that allows you to create a read-only copy of your Cloud SQL instance on an external server, such as a Compute Engine instance or an on-premises database server¹.](#) An external replica can help you scale reads and offload management operations from your data center to Google Cloud. [You can also use an external replica for disaster recovery, migration, or reporting purposes¹.](#)

[To create an external replica, you need to configure a Cloud SQL instance that replicates to one or more replicas external to Cloud SQL, and a source representation instance that represents the source database server in Cloud SQL¹.](#) [You also need to enable access on the Cloud SQL instance for the IP address of the external replica, create a replication user, and export and import the data from the source database server to the external replica¹.](#)

Question: 28

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: C, E

Explanation:

Question: 29

Your company is migrating the existing infrastructure for a highly transactional application to Google Cloud. You have several databases in a MySQL database instance and need to decide how to transfer the data to Cloud SQL. You need to minimize the downtime for the migration of your 500 GB