

Google

Exam Questions Professional-Cloud-Database-Engineer

Google Cloud Certified - Professional Cloud Database Engineer



NEW QUESTION 1

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.

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 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 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 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 8

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 9

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

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

Answer: BC

Explanation:

Replication lag and slow report performance. E is eliminated because using BigQuery would mean changes to the current reports. Report slowness could be the result of poor indexing or just too much read load (or both!). Since excessive load is mentioned in the question, creating additional read replicas and spreading the analytics workload around makes B correct and eliminates A as a way to speed up reporting. That leaves the replication problem. Cloud SQL enables single threaded replication by default, so it stands to reason enabling parallel replication would help the lag. To do that you disable replication on the replica (not the

primary), set flags on the replica and optionally set flags on the primary instance to optimize performance for parallel replication. That makes C correct and D incorrect. <https://cloud.google.com/sql/docs/mysql/replication/manage-replicas#configuring-parallel-replication>

NEW QUESTION 10

You are deploying a new Cloud SQL instance on Google Cloud using the Cloud SQL Auth proxy. You have identified snippets of application code that need to access the new Cloud SQL instance. The snippets reside and execute on an application server running on a Compute Engine machine. You want to follow Google-recommended practices to set up Identity and Access Management (IAM) as quickly and securely as possible. What should you do?

- A. For each application code, set up a common shared user account.
- B. For each application code, set up a dedicated user account.
- C. For the application server, set up a service account.
- D. For the application server, set up a common shared user account.

Answer: C

Explanation:

<https://cloud.google.com/sql/docs/mysql/sql-proxy#using-a-service-account>

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

NEW QUESTION 14

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 19

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 21

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 25

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 27

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 32

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

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

Answer: D

Explanation:

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

NEW QUESTION 34

You are working on a new centralized inventory management system to track items available in 200 stores, which each have 500 GB of data. You are planning a gradual rollout of the system to a few stores each week. You need to design an SQL database architecture that minimizes costs and user disruption during each regional rollout and can scale up or down on nights and holidays. What should you do?

- A. Use Oracle Real Application Cluster (RAC) databases on Bare Metal Solution for Oracle.
- B. Use sharded Cloud SQL instances with one or more stores per database instance.
- C. Use a Bigtable cluster with autoscaling.
- D. Use Cloud Spanner with a custom autoscaling solution.

Answer: D

Explanation:

<https://cloud.google.com/spanner/docs/autoscaling-overview>

* 1. CloudSQL max out at 64TB, so unable to hold 100TB of data. https://cloud.google.com/sql/docs/quotas#metrics_collection_limit 2. Scale is done manually on SQL Cloud

NEW QUESTION 37

You are a DBA of Cloud SQL for PostgreSQL. You want the applications to have password-less authentication for read and write access to the database. Which authentication mechanism should you use?

- A. Use Identity and Access Management (IAM) authentication.
- B. Use Managed Active Directory authentication.
- C. Use Cloud SQL federated queries.
- D. Use PostgreSQL database's built-in authentication.

Answer: A

Explanation:

<https://cloud.google.com/sql/docs/postgres/authentication>

NEW QUESTION 41

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 45

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 47

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 52

You manage a production MySQL database running on Cloud SQL at a retail company. You perform routine maintenance on Sunday at midnight when traffic is slow, but you want to skip routine maintenance during the year-end holiday shopping season. You need to ensure that your production system is available 24/7 during the holidays. What should you do?

- A. Define a maintenance window on Sundays between 12 AM and 1 AM, and deny maintenance periods between November 1 and January 15.
- B. Define a maintenance window on Sundays between 12 AM and 5 AM, and deny maintenance periods between November 1 and February 15.
- C. Build a Cloud Composer job to start a maintenance window on Sundays between 12 AM and 1 AM, and deny maintenance periods between November 1 and January 15.
- D. Create a Cloud Scheduler job to start maintenance at 12 AM on Sunday
- E. Pause the Cloud Scheduler job between November 1 and January 15.

Answer: A

Explanation:

"Deny maintenance period. A block of days in which Cloud SQL does not schedule maintenance. Deny maintenance periods can be up to 90 days long." <https://cloud.google.com/sql/docs/mysql/maintenance>

NEW QUESTION 57

Your online delivery business that primarily serves retail customers uses Cloud SQL for MySQL for its inventory and scheduling application. The required recovery time objective (RTO) and recovery point objective (RPO) must be in minutes rather than hours as a part of your high availability and disaster recovery design. You need a high availability configuration that can recover without data loss during a zonal or a regional failure. What should you do?

- A. Set up all read replicas in a different region using asynchronous replication.
- B. Set up all read replicas in the same region as the primary instance with synchronous replication.

- C. Set up read replicas in different zones of the same region as the primary instance with synchronous replication, and set up read replicas in different regions with asynchronous replication.
- D. Set up read replicas in different zones of the same region as the primary instance with asynchronous replication, and set up read replicas in different regions with synchronous replication.

Answer: C

Explanation:

This answer meets the RTO and RPO requirements by using synchronous replication within the same region, which ensures that all writes made to the primary instance are replicated to disks in both zones before a transaction is reported as committed¹. This minimizes data loss and downtime in case of a zonal or an instance failure, and allows for a quick failover to the standby instance¹.

This answer also meets the high availability and disaster recovery requirements by using asynchronous replication across different regions, which ensures that the data changes made to the primary instance are replicated to the read replicas in other regions with minimal delay². This provides additional redundancy and backup in case of a regional failure, and allows for a manual failover to the read replica in another region².

NEW QUESTION 60

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 61

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 65

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

NEW QUESTION 70

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 73

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

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

Answer: B

Explanation:

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

NEW QUESTION 77

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 81

You are configuring a new application that has access to an existing Cloud Spanner database. The new application reads from this database to gather statistics for a dashboard. You want to follow Google-recommended practices when granting Identity and Access Management (IAM) permissions. What should you do?

- A. Reuse the existing service account that populates this database.
- B. Create a new service account, and grant it the Cloud Spanner Database Admin role.
- C. Create a new service account, and grant it the Cloud Spanner Database Reader role.
- D. Create a new service account, and grant it the spanner.databases.select permission.

Answer: C

Explanation:

<https://cloud.google.com/iam/docs/overview>

NEW QUESTION 85

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 88

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 93

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 96

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 101

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 103

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 104

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 109

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

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

Answer: A

NEW QUESTION 113

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

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

Answer: B

Explanation:

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

NEW QUESTION 115

You are managing a Cloud SQL for PostgreSQL instance in Google Cloud. You have a primary instance in region 1 and a read replica in region 2. After a failure of region 1, you need to make the Cloud SQL instance available again. You want to minimize data loss and follow Google-recommended practices. What should you do?

- A. Restore the Cloud SQL instance from the automatic backups in region 3.
- B. Restore the Cloud SQL instance from the automatic backups in another zone in region 1.
- C. Check "Lag Bytes" in the monitoring dashboard for the primary instance in the read replica instance.
- D. Check the replication status using `pg_catalog.pg_last_wal_receive_lsn()`. Then, fail over to region 2 by promoting the read replica instance.
- E. Check your instance operational log for the automatic failover status.
- F. Look for time, type, and status of the operation.
- G. If the failover operation is successful, no action is necessary.
- H. Otherwise, manually perform `gcloud sql instances failover`.

Answer: C

Explanation:

https://cloud.google.com/sql/docs/postgres/replication/cross-region-replicas#disaster_recovery

NEW QUESTION 119

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 123

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 128

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

NEW QUESTION 130

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 133

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 137

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

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

Answer: C

Explanation:

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

routing. Since you need multi-cluster routing, you're going to need a custom app profile. So C is correct. <https://cloud.google.com/bigtable/docs/app-profiles#default-app-profile>

NEW QUESTION 140

You host an application in Google Cloud. The application is located in a single region and uses Cloud SQL for transactional data. 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
- B. Control the order of update by setting non-production instances to earlier and production instances to later.
- C. Create your database with one primary node and one read replica in the region.
- D. Enable maintenance notifications for users, and reschedule maintenance activities to a specific time after notifications have been sent.
- E. 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.

NEW QUESTION 143

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 144

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

NEW QUESTION 145

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 150

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 154

An analytics team needs to read data out of Cloud SQL for SQL Server and update a table in Cloud Spanner. You need to create a service account and grant least privilege access using predefined roles. What roles should you assign to the service account?

- A. roles/cloudsql.viewer and roles/spanner.databaseUser
- B. roles/cloudsql.editor and roles/spanner.admin
- C. roles/cloudsql.client and roles/spanner.databaseReader
- D. roles/cloudsql.instanceUser and roles/spanner.databaseUser

Answer: A

Explanation:

To read data out of Cloud SQL for SQL Server, you need to use a service account with the roles/cloudsql.viewer role on the Cloud SQL instance. This role grants the service account permission to read data from the instance. Whereas roles/cloudsql.instanceUser will only allow to login to cloud SQL instance. No resource will be allowed to view.

NEW QUESTION 156

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 157

You are migrating a telehealth care company's on-premises data center to Google Cloud. The migration plan specifies: PostgreSQL databases must be migrated to a multi-region backup configuration with cross- region replicas to allow restore and failover in multiple scenarios. MySQL databases handle personally identifiable information (PII) and require data residency compliance at the regional level. You want to set up the environment with minimal administrative effort. What should you do?

- A. Set up Cloud Logging and Cloud Monitoring with Cloud Functions to send an alert every time a new database instance is created, and manually validate the region.
- B. Set up different organizations for each database type, and apply policy constraints at the organization level.
- C. Set up Pub/Sub to ingest data from Cloud Logging, send an alert every time a new database instance is created, and manually validate the region.
- D. Set up different projects for PostgreSQL and MySQL databases, and apply organizational policy constraints at a project level.

Answer: D

NEW QUESTION 162

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

NEW QUESTION 166

.....

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](#)