

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

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

Answer: D

Explanation:

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

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

NEW QUESTION 3

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 4

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 5

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 6

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 7

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 8

You are a DBA on a Cloud Spanner instance with multiple databases. You need to assign these privileges to all members of the application development team on a specific database:

Can read tables, views, and DDL
Can write rows to the tables
Can add columns and indexes
Cannot drop the database
What should you do?

- A. Assign the Cloud Spanner Database Reader and Cloud Spanner Backup Writer roles.
- B. Assign the Cloud Spanner Database Admin role.
- C. Assign the Cloud Spanner Database User role.
- D. Assign the Cloud Spanner Admin role.

Answer: C

Explanation:

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

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

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

Answer: C

Explanation:

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

NEW QUESTION 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 11

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 12

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 17

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

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

Answer: C

Explanation:

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

NEW QUESTION 20

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 24

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 25

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

Your team is building a new inventory management application that will require read and write database instances in multiple Google Cloud regions around the globe. Your database solution requires 99.99% availability and global transactional consistency. You need a fully managed backend relational database to store inventory changes. What should you do?

- A. Use Bigtable.
- B. Use Firestore.
- C. Use Cloud SQL for MySQL
- D. Use Cloud Spanner.

Answer: D

Explanation:

Spanner covers the SLA

NEW QUESTION 34

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 37

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-](https://cloud.google.com/sql/docs/mysql/high-availability)

availability#failback

NEW QUESTION 42

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 44

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 47

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 1AM, 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 50

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

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

Answer: A

NEW QUESTION 51

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

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

Answer: C

Explanation:

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

NEW QUESTION 53

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

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

Answer: D

Explanation:

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

NEW QUESTION 54

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 59

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 60

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 61

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 63

You are migrating your data center to Google Cloud. You plan to migrate your applications to Compute Engine and your Oracle databases to Bare Metal Solution for Oracle. You must ensure that the applications in different projects can communicate securely and efficiently with the Oracle databases. What should you do?

- A. Set up a Shared VPC, configure multiple service projects, and create firewall rules.
- B. Set up Serverless VPC Access.
- C. Set up Private Service Connect.
- D. Set up Traffic Director.

Answer: A

Explanation:

<https://medium.com/google-cloud/shared-vpc-in-google-cloud-64527e0a409e#:~:text=Unlike%20VPC%20peering%2C%20Shared%20VPC%20connects%20projects%20within%20the%20same%20organization.&text=There%20are%20a%20lot%20of,between%20VPCs%20in%20different%20projects>.

NEW QUESTION 67

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 69

Your organization stores marketing data such as customer preferences and purchase history on Bigtable. The consumers of this database are predominantly data analysts and operations users. You receive a service ticket from the database operations department citing poor database performance between 9 AM-10 AM every day. The application team has confirmed no latency from their logs. A new cohort of pilot users that is testing a dataset loaded from a third-party data provider is experiencing poor database performance. Other users are not affected. You need to troubleshoot the issue. What should you do?

- A. Isolate the data analysts and operations user groups to use different Bigtable instances.
- B. Check the Cloud Monitoring table/bytes_used metric from Bigtable.
- C. Use Key Visualizer for Bigtable.
- D. Add more nodes to the Bigtable cluster.

Answer: C

Explanation:

<https://cloud.google.com/bigtable/docs/performance#troubleshooting>

NEW QUESTION 70

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

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

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

Answer: D

NEW QUESTION 75

You are managing a mission-critical Cloud SQL for PostgreSQL instance. Your application team is running important transactions on the database when another DBA starts an on-demand backup. You want to verify the status of the backup. What should you do?

- A. Check the `cloudsql.googleapis.com/postgres.log` instance log.
- B. Perform the `gcloud sql operations list` command.
- C. Use Cloud Audit Logs to verify the status.
- D. Use the Google Cloud Console.

Answer: B

Explanation:

<https://cloud.google.com/sql/docs/postgres/backup-recovery/backups#troubleshooting-backups> Under Troubleshooting: Issue: "You can't see the current operation's status." The Google Cloud console reports only success or failure when the operation is done. It isn't designed to show warnings or other updates. Run the `gcloud sql operations list` command to list all operations for the given Cloud SQL instance.

NEW QUESTION 79

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 84

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 88

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.

NEW QUESTION 92

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 97

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 99

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%20recommendation%20on%20how%20to%20resize%20such%20instances%20and%20reduce%20cost.%20This%20page%20describes%20how%20this%20recommender%20works%20and%20how%20to%20use%20it.>

NEW QUESTION 101

You need to migrate existing databases from Microsoft SQL Server 2016 Standard Edition on a single Windows Server 2019 Datacenter Edition to a single Cloud SQL for SQL Server instance. During the discovery phase of your project, you notice that your on-premises server peaks at around 25,000 read IOPS. You need to ensure that your Cloud SQL instance is sized appropriately to maximize read performance. What should you do?

- A. Create a SQL Server 2019 Standard on Standard machine type with 4 vCPUs, 15 GB of RAM, and 800 GB of solid-state drive (SSD).
- B. Create a SQL Server 2019 Standard on High Memory machine type with at least 16 vCPUs, 104 GB of RAM, and 200 GB of SSD.
- C. Create a SQL Server 2019 Standard on High Memory machine type with 16 vCPUs, 104 GB of RAM, and 4 TB of SSD.
- D. Create a SQL Server 2019 Enterprise on High Memory machine type with 16 vCPUs, 104 GB of RAM, and 500 GB of SSD.

Answer: C

Explanation:

Given that Google SSD performance is related to the size of the disk in an order of 30 IOPS for each GB, it would require at least 833 GB to handle 25000 IOPS, the only answer that exceeds this value is C. <https://cloud.google.com/compute/docs/disks/performance>

NEW QUESTION 104

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 108

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 111

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

- A. Set the automated backups to occur every other day to lower the frequency of backups.
- B. Change the storage tier of the automated backups from solid-state drive (SSD) to hard disk drive (HDD).
- C. Select a different region to store your backups.
- D. Reduce the number of automated backups that are retained to two (2).

Answer: D

Explanation:

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

NEW QUESTION 115

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 116

Your company is developing a new global transactional application that must be ACID-compliant and have 99.999% availability. You are responsible for selecting the appropriate Google Cloud database to serve as a datastore for this new application. What should you do?

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

Answer: B

NEW QUESTION 117

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 120

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 123

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 128

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 131

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 135

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 136

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>

NEW QUESTION 139

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

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

Answer: D

Explanation:

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

NEW QUESTION 144

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

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

Answer: B

Explanation:

https://groups.google.com/g/google-cloud-sql-discuss/c/M5G5_HPXytY?pli=1

NEW QUESTION 147

Your organization is running a Firestore-backed Firebase app that serves the same top ten news stories on a daily basis to a large global audience. You want to optimize content delivery while decreasing cost and latency. What should you do?

- A. Enable serializable isolation in the Firebase app.
- B. Deploy a US multi-region Firestore location.
- C. Build a Firestore bundle, and deploy bundles to Cloud CDN.
- D. Create a Firestore index on the news story date.

Answer: C

Explanation:

A global audience strongly suggests serving content via Google's Content Delivery Network. Changing the isolation level won't decrease cost or latency

NEW QUESTION 151

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

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

Answer: C

Explanation:

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

NEW QUESTION 156

During an internal audit, you realized that one of your Cloud SQL for MySQL instances does not have high availability (HA) enabled. You want to follow Google-recommended practices to enable HA on your existing instance. What should you do?

- A. Create a new Cloud SQL for MySQL instance, enable HA, and use the export and import option to migrate your data.
- B. Create a new Cloud SQL for MySQL instance, enable HA, and use Cloud Data Fusion to migrate your data.
- C. Use the gcloud instances patch command to update your existing Cloud SQL for MySQL instance.

D. Shut down your existing Cloud SQL for MySQL instance, and enable HA.

Answer: C

Explanation:

Creating a new instance and migrating data can be time-consuming and disruptive to your application's availability. Shutting down the existing instance is not a recommended approach, as it will cause downtime for your application.

The recommended approach is to use the `gcloud instances patch` command to enable high availability on your existing Cloud SQL for MySQL instance. This command updates the instance's configuration to enable the failover replica, configure it, and enable automatic failover.

By following this approach, you can ensure minimal downtime, and your application can continue to operate during the process.

NEW QUESTION 157

You are designing for a write-heavy application. During testing, you discover that the write workloads are performant in a regional Cloud Spanner instance but slow down by an order of magnitude in a multi-regional instance. You want to make the write workloads faster in a multi-regional instance. What should you do?

- A. Place the bulk of the read and write workloads closer to the default leader region.
- B. Use staleness of at least 15 seconds.
- C. Add more read-write replicas.
- D. Keep the total CPU utilization under 45% in each region.

Answer: A

Explanation:

<https://cloud.google.com/spanner/docs/instance-configurations#multi-region-best-practices> Best practices For optimal performance, follow these best practices: Design a schema that prevents hotspots and other performance issues. For optimal write latency, place compute resources for write-heavy workloads within or close to the default leader region. For optimal read performance outside of the default leader region, use staleness of at least 15 seconds. To avoid single-region dependency for your workloads, place critical compute resources in at least two regions. A good option is to place them next to the two different read-write regions so that any single region outage will not impact all of your application. Provision enough compute capacity to keep high priority total CPU utilization under 45% in each region.

NEW QUESTION 159

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 161

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

- A. Use query parameters to speed up frequently executed queries.
- B. Change the Cloud Spanner configuration from multi-region to single region.
- C. Use SQL statements to analyze `SPANNER_SYS.READ_STATS*` tables.
- D. Use SQL statements to analyze `SPANNER_SYS.QUERY_STATS*` tables.

Answer: C

Explanation:

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

NEW QUESTION 162

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 163

.....

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](#)