

Google

Exam Questions Professional-Cloud-Database-Engineer

Google Cloud Certified - Professional Cloud Database Engineer



NEW QUESTION 1

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

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

Answer: AE

Explanation:

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

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

This part describes significant corruption - A

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

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

NEW QUESTION 2

Your 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 3

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 4

Your organization has an existing app that just went viral. The app uses a Cloud SQL for MySQL backend database that is experiencing slow disk performance while using hard disk drives (HDDs). You need to improve performance and reduce disk I/O wait times. What should you do?

- A. Export the data from the existing instance, and import the data into a new instance with solid-state drives (SSDs).
- B. Edit the instance to change the storage type from HDD to SSD.
- C. Create a high availability (HA) failover instance with SSDs, and perform a failover to the new instance.
- D. Create a read replica of the instance with SSDs, and perform a failover to the new instance

Answer: A

Explanation:

<https://stackoverflow.com/questions/72034607/can-i-change-storage-type-from-hdd-to-ssd-on-cloud-sql-after-creating-an-instan>

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 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](https://cloud.google.com/sql/docs/mysql/instance-settings#storage-capacity-2ndgen)

NEW QUESTION 8

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](https://cloud.google.com/sql/docs/mysql/replication/manage-replicas#configuring-parallel-replication)

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 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 12

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 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 19

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 22

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 26

Your company is evaluating Google Cloud database options for a mission-critical global payments gateway application. The application must be available 24/7 to users worldwide, horizontally scalable, and support open source databases. You need to select an automatically sharded, 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 27

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 29

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 30

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 31

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 32

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 33

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 37

You are managing a small Cloud SQL instance for developers to do testing. The instance is not critical and has a recovery point objective (RPO) of several days. You want to minimize ongoing costs for this instance. What should you do?

- A. Take no backups, and turn off transaction log retention.
- B. Take one manual backup per day, and turn off transaction log retention.
- C. Turn on automated backup, and turn off transaction log retention.
- D. Turn on automated backup, and turn on transaction log retention.

Answer: C

Explanation:

<https://cloud.google.com/sql/docs/mysql/backup-recovery/backups>

NEW QUESTION 42

Your organization has a security policy to ensure that all Cloud SQL for PostgreSQL databases are secure. You want to protect sensitive data by using a key that meets specific locality or residency requirements. Your organization needs to control the key's lifecycle activities. You need to ensure that data is encrypted at rest and in transit. What should you do?

- A. Create the database with Google-managed encryption keys.
- B. Create the database with customer-managed encryption keys.
- C. Create the database persistent disk with Google-managed encryption keys.

D. Create the database persistent disk with customer-managed encryption keys.

Answer: B

Explanation:

<https://cloud.google.com/sql/docs/postgres/configure-cmek#createcmekinstance>

NEW QUESTION 46

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 50

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 53

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 55

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 60

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 64

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 68

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 72

You are designing a payments processing application on Google Cloud. The application must continue to serve requests and avoid any user disruption if a regional failure occurs. You need to use AES-256 to encrypt data in the database, and you want to control where you store the encryption key. What should you do?

- A. Use Cloud Spanner with a customer-managed encryption key (CMEK).
- B. Use Cloud Spanner with default encryption.
- C. Use Cloud SQL with a customer-managed encryption key (CMEK).
- D. Use Bigtable with default encryption.

Answer: A

Explanation:

Yes default encryption comes with AES-256 but the question states that you need to control where you store the encryption keys. that can be achieved by CMEK.

NEW QUESTION 75

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 76

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 78

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

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

Answer: B

Explanation:

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

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

NEW QUESTION 83

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 84

You need to issue a new server certificate because your old one is expiring. You need to avoid a restart of your Cloud SQL for MySQL instance. What should you do in your Cloud SQL instance?

- A. Issue a rollback, and download your server certificate.
- B. Create a new client certificate, and download it.
- C. Create a new server certificate, and download it.
- D. Reset your SSL configuration, and download your server certificate.

Answer: C

Explanation:

<https://cloud.google.com/sql/docs/sqlserver/configure-ssl-instance#server-certs>

NEW QUESTION 85

Your team is running a Cloud SQL for MySQL instance with a 5 TB database that must be available 24/7. You need to save database backups on object storage with minimal operational overhead or risk to your production workloads. What should you do?

- A. Use Cloud SQL serverless exports.
- B. Create a read replica, and then use the mysqldump utility to export each table.
- C. Clone the Cloud SQL instance, and then use the mysqldump utility to export the data.
- D. Use the mysqldump utility on the primary database instance to export the backup.

Answer: A

Explanation:

<https://cloud.google.com/blog/products/databases/introducing-cloud-sql-serverless-exports>

NEW QUESTION 87

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 88

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

- A. Create a Cloud SQL for MySQL instance for your databases, and configure Datastream to stream your database changes to Cloud SQL. Select the Backfill historical data check box on your stream configuration to initiate Datastream to backfill any data that is out of sync between the source and destination.
- B. Delete your stream when all changes are moved to Cloud SQL for MySQL, and update your application to use the new instance.
- C. Create migration job using Database Migration Service. Set the migration job type to Continuous, and allow the databases to complete the full dump phase and start sending data in change data capture (CDC) mode. Wait for the replication delay to minimize, initiate a promotion of the new Cloud SQL instance, and wait for the migration job to complete. Update your application connections to the new instance.
- D. Create migration job using Database Migration Service. Set the migration job type to One-time, and perform this migration during a maintenance window. Stop all write workloads to the source database and initiate the dump.
- E. Wait for the dump to be loaded into the Cloud SQL destination database and the destination database to be promoted to the primary database. Update your application connections to the new instance.
- F. Use the mysqldump utility to manually initiate a backup of MySQL during the application maintenance window. Move the files to Cloud Storage, and import each database into your Cloud SQL instance.
- G. Continue to dump each database until all the databases are migrated. Update your application connections to the new instance.

Answer: B

Explanation:

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

NEW QUESTION 91

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 95

You are designing an augmented reality game for iOS and Android devices. You plan to use Cloud Spanner as the primary backend database for game state storage and player authentication. You want to track in-game rewards that players unlock at every stage of the game. During the testing phase, you discovered that costs are much higher than anticipated, but the query response times are within the SLA. You want to follow Google- recommended practices. You need the database to be performant and highly available while you keep costs low. What should you do?

- A. Manually scale down the number of nodes after the peak period has passed.
- B. Use interleaving to co-locate parent and child rows.
- C. Use the Cloud Spanner query optimizer to determine the most efficient way to execute the SQL query.
- D. Use granular instance sizing in Cloud Spanner and Autoscaler.

Answer: D

Explanation:

Granular instance is available in Public Preview. With this feature, you can run workloads on Spanner at as low as 1/10th the cost of regular instances, <https://cloud.google.com/blog/products/databases/get-more-out-of-spanner-with-granular-instance-sizing>

NEW QUESTION 100

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

- A. Use log messages produced by Cloud SQL.

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

Answer: B

Explanation:

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

NEW QUESTION 103

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 106

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

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

Answer: B

Explanation:

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

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

This answer is correct because PgBouncer is a lightweight connection pooler for PostgreSQL that can help you distribute read requests between the Cloud SQL primary and read replica instances¹. PgBouncer can also improve performance and scalability by reducing the overhead of creating new connections and reusing existing ones¹. You can install PgBouncer on a Compute Engine instance and configure it to connect to the Cloud SQL instances using private IP addresses or the Cloud SQL Auth proxy².

NEW QUESTION 109

Your company uses Bigtable for a user-facing application that displays a low-latency real-time dashboard. You need to recommend the optimal storage type for this read-intensive database. What should you do?

- A. Recommend solid-state drives (SSD).
- B. Recommend splitting the Bigtable instance into two instances in order to load balance the concurrent reads.
- C. Recommend hard disk drives (HDD).
- D. Recommend mixed storage types.

Answer: A

Explanation:

If you plan to store extensive historical data for a large number of remote-sensing devices and then use the data to generate daily reports, the cost savings for HDD storage might justify the performance tradeoff. On the other hand, if you plan to use the data to display a real-time dashboard, it probably would not make sense to use HDD storage—reads would be much more frequent in this case, and reads that are not scans are much slower with HDD storage.

NEW QUESTION 111

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 114

You are designing a database strategy for a new web application in one region. You need to minimize write latency. What should you do?

- A. Use Cloud SQL with cross-region replicas.
- B. Use high availability (HA) Cloud SQL with multiple zones.
- C. Use zonal Cloud SQL without high availability (HA).
- D. Use Cloud Spanner in a regional configuration.

Answer: D

Explanation:

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

NEW QUESTION 118

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 122

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 125

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 129

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 131

You want to migrate your PostgreSQL database from another cloud provider to Cloud SQL. You plan on using Database Migration Service and need to assess the impact of any known limitations. What should you do? (Choose two.)

- A. Identify whether the database has over 512 tables.
- B. Identify all tables that do not have a primary key.
- C. Identify all tables that do not have at least one foreign key.
- D. Identify whether the source database is encrypted using pgcrypto extension.
- E. Identify whether the source database uses customer-managed encryption keys (CMEK).

Answer: CE

NEW QUESTION 134

.....

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](#)