

Google

Exam Questions Professional-Cloud-Database-Engineer

Google Cloud Certified - Professional Cloud Database Engineer



NEW QUESTION 1

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

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

Answer: DE

Explanation:

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

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

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

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

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

NEW QUESTION 2

Your ecommerce website captures user clickstream data to analyze customer traffic patterns in real time and support personalization features on your website. You plan to analyze this data using big data tools. You need a low-latency solution that can store 8 TB of data and can scale to millions of read and write requests per second. What should you do?

- A. Write your data into Bigtable and use Dataproc and the Apache Hbase libraries for analysis.
- B. Deploy a Cloud SQL environment with read replicas for improved performance.
- C. Use Datastream to export data to Cloud Storage and analyze with Dataproc and the Cloud Storage connector.
- D. Use Memorystore to handle your low-latency requirements and for real-time analytics.
- E. Stream your data into BigQuery and use Dataproc and the BigQuery Storage API to analyze large volumes of data.

Answer: A

Explanation:

Start with the lowest tier and smallest size and then grow your instance as needed. Memorystore provides automated scaling using APIs, and optimized node placement across zones for redundancy. Memorystore for Memcached can support clusters as large as 5 TB, enabling millions of QPS at very low latency

NEW QUESTION 3

You need to provision several hundred Cloud SQL for MySQL instances for multiple project teams over a one-week period. You must ensure that all instances adhere to company standards such as instance naming conventions, database flags, and tags. What should you do?

- A. Automate instance creation by writing a Dataflow job.
- B. Automate instance creation by setting up Terraform scripts.
- C. Create the instances using the Google Cloud Console UI.
- D. Create clones from a template Cloud SQL instance.

Answer: B

NEW QUESTION 4

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 5

You are evaluating Cloud SQL for PostgreSQL as a possible destination for your on-premises PostgreSQL instances. Geography is becoming increasingly relevant to customer privacy worldwide. Your solution must support data residency requirements and include a strategy to:

configure where data is stored

control where the encryption keys are stored govern the access to data

What should you do?

- A. Replicate Cloud SQL databases across different zones.
- B. Create a Cloud SQL for PostgreSQL instance on Google Cloud for the data that does not need to adhere to data residency requirement
- C. Keep the data that must adhere to data residency requirements on-premise
- D. Make application changes to support both databases.
- E. Allow application access to data only if the users are in the same region as the Google Cloud region for the Cloud SQL for PostgreSQL database.
- F. Use features like customer-managed encryption keys (CMEK), VPC Service Controls, and Identity and Access Management (IAM) policies.

Answer: D

Explanation:

<https://cloud.google.com/blog/products/identity-security/meet-data-residency-requirements-with-google-cloud>

NEW QUESTION 6

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 7

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 8

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 9

Your organization has strict policies on tracking rollouts to production and periodically shares this information with external auditors to meet compliance requirements. You need to enable auditing on several Cloud Spanner databases. What should you do?

- A. Use replication to roll out changes to higher environments.
- B. Use backup and restore to roll out changes to higher environments.
- C. Use Liquibase to roll out changes to higher environments.
- D. Manually capture detailed DBA audit logs when changes are rolled out to higher environments.

Answer: C

Explanation:

To satisfy audit reporting you would need a way to record what was changed and when. The best answer is one which uses some kind of source code control system (SCCS). That rules out A and B. Any mention of anything manual in a cloud environment should look suspicious, which leave option C. As it happens, Liquibase is an SCCS and can be integrated with Spanner. <https://cloud.google.com/spanner/docs/use-liquibase>

NEW QUESTION 10

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 10

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 15

You use Python scripts to generate weekly SQL reports to assess the state of your databases and determine whether you need to reorganize tables or run statistics. You want to automate this report but need to minimize operational costs and overhead. What should you do?

- A. Create a VM in Compute Engine, and run a cron job.
- B. Create a Cloud Composer instance, and create a directed acyclic graph (DAG).
- C. Create a Cloud Function, and call the Cloud Function using Cloud Scheduler.
- D. Create a Cloud Function, and call the Cloud Function from a Cloud Tasks queue.

Answer: C

Explanation:

Cloud Scheduler triggers actions at regular fixed intervals, whereas Cloud Tasks triggers actions based on how the individual task object is configured. Reference: <https://cloud.google.com/tasks/docs/comp-tasks-sched>

NEW QUESTION 20

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 23

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 26

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 30

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 31

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

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

Answer: A

Explanation:

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

NEW QUESTION 33

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

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

Answer: C

Explanation:

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

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

NEW QUESTION 38

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 43

You are configuring a brand new Cloud SQL for PostgreSQL database instance in Google Cloud. Your application team wants you to deploy one primary instance, one standby instance, and one read replica instance. You need to ensure that you are following Google- recommended practices for high availability. What should you do?

- A. Configure the primary instance in zone A, the standby instance in zone C, and the read replica in zone B, all in the same region.
- B. Configure the primary and standby instances in zone A and the read replica in zone B, all in the same region.
- C. Configure the primary instance in one region, the standby instance in a second region, and the read replica in a third region.
- D. Configure the primary, standby, and read replica instances in zone A, all in the same region.

Answer: A

Explanation:

<https://cloud.google.com/sql/docs/postgres/high-availability#failover-overview>

NEW QUESTION 45

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 50

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 51

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 53

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 57

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 60

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 configuratio
- 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 applicatio
- 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 65

You are the DBA of an online tutoring application that runs on a Cloud SQL for PostgreSQL database. You are testing the implementation of the cross-regional failover configuration. The database in region R1 fails over successfully to region R2, and the database becomes available for the application to process data. During testing, certain scenarios of the application work as expected in region R2, but a few scenarios fail with database errors. The application-related database queries, when executed in isolation from Cloud SQL for PostgreSQL in region R2, work as expected. The application performs completely as expected when the database fails back to region R1. You need to identify the cause of the database errors in region R2. What should you do?

- A. Determine whether the versions of Cloud SQL for PostgreSQL in regions R1 and R2 are different.
- B. Determine whether the database patches of Cloud SQL for PostgreSQL in regions R1 and R2 are different.
- C. Determine whether the failover of Cloud SQL for PostgreSQL from region R1 to region R2 is in progress or has completed successfully.
- D. Determine whether Cloud SQL for PostgreSQL in region R2 is a near-real-time copy of region R1 but not an exact copy.

Answer: D

Explanation:

Verify that the replica has processed all the transactions it has received from the primary. This ensures that when promoted, the replica reflects all transactions that were received before the primary became unavailable. https://cloud.google.com/sql/docs/postgres/replication/cross-region-replicas#verify_failover_criteria

NEW QUESTION 68

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 69

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 73

You are running a transactional application on Cloud SQL for PostgreSQL in Google Cloud.

The database is running in a high availability configuration within one region. You have encountered issues with data and want to restore to the last known pristine version of the database. What should you do?

- A. Create a clone database from a read replica database, and restore the clone in the same region.
- B. Create a clone database from a read replica database, and restore the clone into a different zone.
- C. Use the Cloud SQL point-in-time recovery (PITR) feature.
- D. Restore the copy from two hours ago to a new database instance.
- E. Use the Cloud SQL database import feature.
- F. Import last week's dump file from Cloud Storage.

Answer: C

Explanation:

Using import/export from last week is slow for large scale databases and will restore database from last week.

NEW QUESTION 75

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 79

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 Dataproc 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 80

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 82

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 86

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 88

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

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

Answer: C

Explanation:

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

NEW QUESTION 90

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 92

You are starting a large CSV import into a Cloud SQL for MySQL instance that has many open connections. You checked memory and CPU usage, and sufficient resources are available. You want to follow Google-recommended practices to ensure that the import will not time out. What should you do?

- A. Close idle connections or restart the instance before beginning the import operation.
- B. Increase the amount of memory allocated to your instance.
- C. Ensure that the service account has the Storage Admin role.
- D. Increase the number of CPUs for the instance to ensure that it can handle the additional import operation.

Answer: A

Explanation:

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

NEW QUESTION 95

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 98

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 101

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 104

Your organization has a ticketing system that needs an online marketing analytics and reporting application. You need to select a relational database that can manage hundreds of terabytes of data to support this new application. Which database should you use?

- A. Cloud SQL
- B. BigQuery
- C. Cloud Spanner
- D. Bigtable

Answer: B

NEW QUESTION 109

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 113

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 114

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 119

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 120

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 122

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 125

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

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

Answer: A

Explanation:

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

NEW QUESTION 128

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 132

You have deployed a Cloud SQL for SQL Server instance. In addition, you created a cross-region read replica for disaster recovery (DR) purposes. Your company requires you to maintain and monitor a recovery point objective (RPO) of less than 5 minutes. You need to verify that your cross-region read replica meets the allowed RPO. What should you do?

- A. Use Cloud SQL instance monitoring.
- B. Use the Cloud Monitoring dashboard with available metrics from Cloud SQL.
- C. Use Cloud SQL logs.
- D. Use the SQL Server Always On Availability Group dashboard.

Answer: D

Explanation:

Note, you cannot create a read replica in Cloud SQL for SQL Server unless you use an Enterprise Edition. Which is also a requirement for configuring SQL Server AG. That's not a coincidence. That's how Cloud SQL for SQL Server creates SQL Server read replicas. To find out about the replication, use the AG Dashboard in SSMS.

<https://cloud.google.com/sql/docs/sqlserver/replication/manage-replicas#promote-replica>

NEW QUESTION 133

Your company is shutting down their data center and migrating several MySQL and PostgreSQL databases to Google Cloud. Your database operations team is severely constrained by ongoing production releases and the lack of capacity for additional on-premises backups. You want to ensure that the scheduled migrations happen with minimal downtime and that the Google Cloud databases stay in sync with the on-premises data changes until the applications can cut over. What should you do? (Choose two.)

- A. Use an external read replica to migrate the databases to Cloud SQL.
- B. Use a read replica to migrate the databases to Cloud SQL.
- C. Use Database Migration Service to migrate the databases to Cloud SQL.
- D. Use a cross-region read replica to migrate the databases to Cloud SQL.
- E. Use replication from an external server to migrate the databases to Cloud SQL.

Answer: CE

NEW QUESTION 136

Your organization has a busy transactional Cloud SQL for MySQL instance. Your analytics team needs access to the data so they can build monthly sales reports. You need to provide data access to the analytics team without adversely affecting performance. What should you do?

- A. Create a read replica of the database, provide the database IP address, username, and password to the analytics team, and grant read access to required tables to the team.
- B. Create a read replica of the database, enable the cloudsql.iam_authentication flag on the replica, and grant read access to required tables to the analytics team.

- C. Enable the cloudsql.iam_authentication flag on the primary database instance, and grant read access to required tables to the analytics team.
- D. Provide the database IP address, username, and password of the primary database instance to the analytics, team, and grant read access to required tables to the team.

Answer: B

Explanation:

"Read replicas do not have the cloudsql.iam_authentication flag enabled automatically when it is enabled on the primary instance."
https://cloud.google.com/sql/docs/postgres/replication/create-replica#configure_iam_replicas

NEW QUESTION 139

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 142

.....

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](#)