Number: Google
Passing Score: 800
Time Limit: 120 min

Exam Code: Professional Cloud Database Engineer
Exam Name: Professional Cloud Database Engineer
Website: https://VCEup.com/

QUESTION 1
Your team recently released a new version of a highly consumed application to accommodate additional user traffic. Shortly after the release, you received an alert from your production monitoring team that there is consistently high replication lag between your primary instance and the read replicas of your Cloud SQL for MySQL instances. You need to resolve the replication lag.
What should you do?

A. Identify and optimize slow running queries, or set parallel replication flags.

B. Stop all running queries, and re-create the replicas.

C. Edit the primary instance to upgrade to a larger disk, and increase vCPU count.

D. Edit the primary instance to add additional memory.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 2
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

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 3
You are designing an augmented reality game for iOS and Android devices. You plan to use Cloud Spanner as the primary backend database for game state storage and player authentication. You want to track in-game rewards that players unlock at every stage of the game. During the testing phase, you discovered that costs are much higher than anticipated, but the query response times are within the SL

A. You want to follow Google-recommended practices. You need the database to be performant and highly available while you keep costs low. What should you do?

B. Manually scale down the number of nodes after the peak period has passed.

C. Use interleaving to co-locate parent and child rows.

D. Use the Cloud Spanner query optimizer to determine the most efficient way to execute the SQL query.

E. Use granular instance sizing in Cloud Spanner and Autoscaler.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 4
You recently launched a new product to the US market. You currently have two Bigtable clusters in one US region to serve all the traffic. Your marketing team is planning an immediate expansion to APAC. You need to roll out the regional expansion while implementing high availability according to Google-recommended practices. What should you do?

A. Maintain a target of 23% CPU utilization by locating: cluster-a in zone us-central1-a cluster-b in zone europe-west1-d cluster-c in zone asia-east1-b
B. Maintain a target of 23% CPU utilization by locating: cluster-a in zone us-central1-a cluster-b in zone us-central1-b cluster-c in zone us-east1-a
C. Maintain a target of 35% CPU utilization by locating: cluster-a in zone us-central1-a cluster-b in zone australia-southeast1-a cluster-c in zone europe-west1-d cluster-d in zone asia-east1-b
D. Maintain a target of 35% CPU utilization by locating: cluster-a in zone us-central1-a cluster-b in zone us-central2-a cluster-c in zone asia-northeast1-b cluster-d in zone asia-east1-b

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 5
Your ecommerce website captures user clickstream data to analyze customer traffic patterns in real time and support personalization features on your website. You plan to analyze this data using big data tools. You need a low-latency solution that can store 8 TB of data and can scale to millions of read and write requests per second. What should you do?

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

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 6
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.

Correct Answer: AD
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 7
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.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

Explanation:

QUESTION 8
You are building an application that allows users to customize their website and mobile experiences.
The application will capture user information and preferences. User profiles have a dynamic schema, and users can add or delete information from their profile. You need to ensure that user changes automatically trigger updates to your downstream BigQuery data warehouse. What should you do?

A. Store your data in Bigtable, and use the user identifier as the key. Use one column family to store user profile data, and use another column family to store user preferences.

B. Use Cloud SQL, and create different tables for user profile data and user preferences from your recommendations model. Use SQL to join the user profile data and preferences

C. Use Firestore in Native mode, and store user profile data as a document. Update the user profile with preferences specific to that user and use the user identifier to query.

D. Use Firestore in Datastore mode, and store user profile data as a document. Update the user profile with preferences specific to that user and use the user identifier to query.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 9
Your application uses Cloud SQL for MySQL. Your users run reports on data that relies on near-real time; however, the additional analytics caused excessive load on the primary database. You created a read replica for the analytics workloads, but now your users are complaining about the lag in data changes and that their reports are still slow. You need to improve the report performance and shorten the lag in data replication without making changes to the current reports. Which two approaches should you implement? (Choose two.)

A. Create secondary indexes on the replica.

B. Create additional read replicas, and partition your analytics users to use different read replicas.

C. Disable replication on the read replica, and set the flag for parallel replication on the read replica.
Re-enable replication and optimize performance by setting flags on the primary instance.

D. Disable replication on the primary instance, and set the flag for parallel replication on the primary instance. Re-enable replication and optimize performance by setting flags on the read replica.

E. Move your analytics workloads to BigQuery, and set up a streaming pipeline to move data and update BigQuery.

Correct Answer: BE
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 10
You are evaluating Cloud SQL for PostgreSQL as a possible destination for your on-premises PostgreSQL instances. Geography is becoming increasingly relevant to customer privacy worldwide.
Your solution must support data residency requirements and include a strategy to: configure where data is stored control where the encryption keys are stored govern the access to data What should you do?

A. Replicate Cloud SQL databases across different zones.

B. Create a Cloud SQL for PostgreSQL instance on Google Cloud for the data that does not need to adhere to data residency requirements. Keep the data that must adhere to data residency requirements on-premises. Make application changes to support both databases.

C. Allow application access to data only if the users are in the same region as the Google Cloud region for the Cloud SQL for PostgreSQL database.

D. Use features like customer-managed encryption keys (CMEK), VPC Service Controls, and Identity and Access Management (IAM) policies.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 11
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.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 12
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 MongoD8 Atlas with Charts.
D. Use Bigtable with Looker.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 13
Your application follows a microservices architecture and uses a single large Cloud SQL instance, which is starting to have performance issues as your application grows. in the Cloud Monitoring dashboard, the CPU utilization looks normal You want to follow Google-recommended practices to resolve and prevent these performance issues while avoiding any major refactoring. What should you do?

A. Use Cloud Spanner instead of Cloud SQL.
B. Increase the number of CPUs for your instance.
C. Increase the storage size for the instance.
D. Use many smaller Cloud SQL instances.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 14
You need to perform a one-time migration of data from a running Cloud SQL for MySQL instance in the us-central1 region to a new Cloud SQL for MySQL instance in the us-east1 region. You want to follow Google-recommended practices to minimize performance impact on the currently running instance. What should you do?

A. Create and run a Dataflow job that uses JdbcIO to copy data from one Cloud SQL instance to another.
B. Create two Datastream connection profiles, and use them to create a stream from one Cloud SQL instance to another.
C. Create a SQL dump file in Cloud Storage using a temporary instance, and then use that file to import into a new instance.
D. Create a CSV file by running the SQL statement SELECT...INTO OUTFILE, copy the file to a Cloud Storage bucket, and import it into a new instance.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 15
You are running a mission-critical application on a Cloud SQL for PostgreSQL database with a multizonal 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 andread replica instances.
D. Use the Cloud SQL Auth proxy for database connection pooling between the Cloud SQL primary and read replica instances.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 16
Your organization deployed a new version of a critical application that uses Cloud SQL for MySQL with high availability (HA) and binary logging enabled to store transactional information. The latest release of the application had an error that caused massive data corruption in your Cloud SQL for MySQL database. You need to minimize data loss. What should you do?

A. Open the Google Cloud Console, navigate to SQL > Backups, and select the last version of the automated backup before the corruption.
B. Reload the Cloud SQL for MySQL database using the LOAD DATA command to load data from CSV files that were used to initialize the instance.
C. Perform a point-in-time recovery of your Cloud SQL for MySQL database, selecting a date and time before the data was corrupted.
D. Fail over to the Cloud SQL for MySQL HA instance. Use that instance to recover the transactions that occurred before the corruption.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 17
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.

Correct Answer: BE
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 18
You are using Compute Engine on Google Cloud and your data center to manage a set of MySQL databases in a hybrid configuration. You need to create replicas to scale reads and to offload part of the management operation. What should you do?

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

Correct Answer: B

Explanation/Reference:
Explanation:

QUESTION 19
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 Database Migration Service to migrate the databases to Cloud SQL.

B. Use a cross-region read replica to migrate the databases to Cloud SQL.

C. Use replication from an external server to migrate the databases to Cloud SQL.

D. Use an external read replica to migrate the databases to Cloud SQL.

E. Use a read replica to migrate the databases to Cloud SQL.

Correct Answer: CE
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 20
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.
   Delete your stream when all changes are moved to Cloud SQL for MySQL, and update your application to use the new instance.
B. 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.
C. 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. 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.
D. 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.
   Continue to dump each database until all the databases are migrated.
   Update your application connections to the new instance.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 21
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.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 22
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.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 23
You support a consumer inventory application that runs on a multi-region instance of Cloud Spanner.
A customer opened a support ticket to complain about slow response times. You notice a Cloud Monitoring alert about high CPU utilization. You want to follow Google-recommended practices to address the CPU performance issue. What should you do first?

A. Increase the number of processing units.
B. Modify the database schema, and add additional indexes.
C. Shard data required by the application into multiple instances.
D. Decrease the number of processing units.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 24
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.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 25
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-a.
   Create a multiple-zone instance in us-west1-b.
   Create a read replica in us-east1-c.
B. Provision a Cloud SQL for MySQL instance in us-central1-a.
   Create a multiple-zone instance in us-central1-b.
   Create a read replica in us-east1-b.
C. Provision a Cloud SQL for MySQL instance in us-central1-a.
   Create a multiple-zone instance in us-east-b.
   Create a read replica in us-east1-c.
D. Provision a Cloud SQL for MySQL instance in us-central1-a.
   Create a multiple-zone instance in us-east1-b.
   Create a read replica in us-central1-b.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 26
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.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 27
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. Write the results back.
C. If the corruption is significant, use import and export.
D. If the corruption is insignificant, use backup and restore, and specify a recovery timestamp.
E. If the corruption is insignificant, perform a stale read and specify a recovery timestamp. Write the results back.

Correct Answer: BE
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 28
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.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 29
You are migrating your data center to Google Cloud. You plan to migrate your applications to Compute Engine and your Oracle databases to Bare Metal Solution for Oracle. You must ensure that the applications in different projects can communicate securely and efficiently with the Oracle databases. What should you do?

A. Set up a Shared VPC, configure multiple service projects, and create firewall rules.

B. Set up Serverless VPC Access.

C. Set up Private Service Connect.

D. Set up Traffic Director.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:

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.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation: