

# Exam Questions Professional-Cloud-Database-Engineer

Google Cloud Certified - Professional Cloud Database Engineer

<https://www.2passeasy.com/dumps/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](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 timestamps<sup>1</sup>. 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 server<sup>2</sup>.

\* 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 ID<sup>1</sup>. 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 table<sup>2</sup>.

**NEW QUESTION 2**

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 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**

Your organization works with sensitive data that requires you to manage your own encryption keys. You are working on a project that stores that data in a Cloud SQL database. You need to ensure that stored data is encrypted with your keys. What should you do?

- A. Export data periodically to a Cloud Storage bucket protected by Customer-Supplied Encryption Keys.
- B. Use Cloud SQL Auth proxy.
- C. Connect to Cloud SQL using a connection that has SSL encryption.
- D. Use customer-managed encryption keys with Cloud SQL.

**Answer:** D

**NEW QUESTION 7**

You work in the logistics department. Your data analysis team needs daily extracts from Cloud SQL for MySQL to train a machine learning model. The model will be used to optimize next-day routes. You need to export the data in CSV format. You want to follow Google-recommended practices. What should you do?

- A. Use Cloud Scheduler to trigger a Cloud Function that will run a select \* from table(s) query to call the cloudsql.instances.export API.
- B. Use Cloud Scheduler to trigger a Cloud Function through Pub/Sub to call the cloudsql.instances.export API.
- C. Use Cloud Composer to orchestrate an export by calling the cloudsql.instances.export API.
- D. Use Cloud Composer to execute a select \* from table(s) query and export results.

**Answer:** B

**Explanation:**

<https://cloud.google.com/blog/topics/developers-practitioners/scheduling-cloud-sql-exports-using-cloud-functions-and-cloud-scheduler>

**NEW QUESTION 8**

You are designing a database strategy for a new web application. You plan to start with a small pilot in one country and eventually expand to millions of users in a global audience. You need to ensure that the application can run 24/7 with minimal downtime for maintenance. What should you do?

- A. Use Cloud Spanner in a regional configuration.
- B. Use Cloud Spanner in a multi-region configuration.
- C. Use Cloud SQL with cross-region replicas.
- D. Use highly available Cloud SQL with multiple zones.

**Answer:** A

**Explanation:**

<https://docs.google.com/forms/d/e/1FAIpQLSfZ77ZnuUL0NpU-bOtO5QUkC0cnRCe5YKMiuBLXwfV3abBqkg/viewform>

**NEW QUESTION 9**

Your company uses the Cloud SQL out-of-disk recommender to analyze the storage utilization trends of production databases over the last 30 days. Your database operations team uses these recommendations to proactively monitor storage utilization and implement corrective actions. You receive a recommendation that the instance is likely to run out of disk space. What should you do to address this storage alert?

- A. Normalize the database to the third normal form.
- B. Compress the data using a different compression algorithm.
- C. Manually or automatically increase the storage capacity.
- D. Create another schema to load older data.

**Answer:** C

**Explanation:**

<https://cloud.google.com/sql/docs/mysql/instance-settings#storage-capacity-2ndgen>

**NEW QUESTION 10**

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.
- 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 instance.
- F. Re-enable replication and optimize performance by setting flags on the read replica.
- G. Move your analytics workloads to BigQuery, and set up a streaming pipeline to move data and update BigQuery.

**Answer:** BC

**Explanation:**

Replication lag and slow report performance. E is eliminated because using BigQuery would mean changes to the current reports. Report slowness could be the result of poor indexing or just too much read load (or both!). Since excessive load is mentioned in the question, creating additional read replicas and spreading the analytics workload around makes B correct and eliminates A as a way to speed up reporting. That leaves the replication problem. Cloud SQL enables single threaded replication by default, so it stands to reason enabling parallel replication would help the lag. To do that you disable replication on the replica (not the primary), set flags on the replica and optionally set flags on the primary instance to optimize performance for parallel replication. That makes C correct and D incorrect. <https://cloud.google.com/sql/docs/mysql/replication/manage-replicas#configuring-parallel-replication>

**NEW QUESTION 10**

You 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 15**

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 18**

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 22**

You want to migrate an on-premises 100 TB Microsoft SQL Server database to Google Cloud over a 1 Gbps network link. You have 48 hours allowed downtime to migrate this database. What should you do? (Choose two.)

- A. Use a change data capture (CDC) migration strategy.
- B. Move the physical database servers from on-premises to Google Cloud.
- C. Keep the network bandwidth at 1 Gbps, and then perform an offline data migration.
- D. Increase the network bandwidth to 2 Gbps, and then perform an offline data migration.
- E. Increase the network bandwidth to 10 Gbps, and then perform an offline data migration.

**Answer:** AE

**Explanation:**

[https://cloud.google.com/architecture/migration-to-google-cloud-transferring-your-large-datasets#online\\_versus\\_offline\\_transfer](https://cloud.google.com/architecture/migration-to-google-cloud-transferring-your-large-datasets#online_versus_offline_transfer)

**NEW QUESTION 27**

Your project is using Bigtable to store data that should not be accessed from the public internet under any circumstances, even if the requestor has a valid service account key. You need to secure access to this data. What should you do?

- A. Use Identity and Access Management (IAM) for Bigtable access control.

- B. Use VPC Service Controls to create a trusted network for the Bigtable service.
- C. Use customer-managed encryption keys (CMEK).
- D. Use Google Cloud Armor to add IP addresses to an allowlist.

**Answer:** B

**Explanation:**

"Users can define a security perimeter around Google Cloud resources such as Cloud Storage buckets, Bigtable instances, and BigQuery datasets to constrain data within a VPC and control the flow of data." <https://cloud.google.com/vpc-service-controls>

**NEW QUESTION 29**

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 32**

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 35**

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 40**

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

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

**Answer:** D

**Explanation:**

This answer is correct because Bare Metal Solution for Oracle is a service that provides dedicated physical servers and networking infrastructure for running Oracle databases on Google Cloud<sup>1</sup>. Bare Metal Solution for Oracle supports Oracle RAC, which is a cluster database that provides high availability, scalability, and performance for Oracle workloads<sup>2</sup>. By using Bare Metal Solution for Oracle, you can migrate your Oracle databases with minimal to no changes to the applications, and you can leverage the native Google Cloud services and interconnectivity<sup>1</sup>.

**NEW QUESTION 44**

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 45**

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

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

**Answer:** D

**Explanation:**

Spanner covers the SLA

**NEW QUESTION 48**

You have a Cloud SQL instance (DB-1) with two cross-region read replicas (DB-2 and DB-3). During a business continuity test, the primary instance (DB-1) was taken offline and a replica (DB-2) was promoted. The test has concluded and you want to return to the pre-test configuration. What should you do?

- A. Bring DB-1 back online.
- B. Delete DB-1, and re-create DB-1 as a read replica in the same region as DB-1.
- C. Delete DB-2 so that DB-1 automatically reverts to the primary instance.
- D. Create DB-4 as a read replica in the same region as DB-1, and promote DB-4 to primary.

**Answer:** D

**Explanation:**

If you need to have the primary instance in the zone that had the outage, you can do a failback. A failback performs the same steps as the failover, only in the opposite direction, to reroute traffic back to the original instance. To perform a failback, use the procedure in Initiating failover. <https://cloud.google.com/sql/docs/mysql/high-availability#failback>

**NEW QUESTION 53**

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

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

**Answer:** A

**Explanation:**

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

**NEW QUESTION 58**

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 61**

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 65**

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 70**

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](https://cloud.google.com/sql/docs/mysql/backup-recovery/pitr#perform_the_point-in-time_recovery_using_binary_log_positions)

**NEW QUESTION 74**

Your digital-native business runs its database workloads on Cloud SQL. Your website must be globally accessible 24/7. You need to prepare your Cloud SQL instance for high availability (HA). You want to follow Google-recommended practices. What should you do? (Choose two.)

- A. Set up manual backups.
- B. Create a PostgreSQL database on-premises as the HA option.
- C. Configure single zone availability for automated backups.
- D. Enable point-in-time recovery.
- E. Schedule automated backups.

**Answer:** DE

**Explanation:**

D. Enable point-in-time recovery - This feature allows you to restore your database to a specific point in time. It helps protect against data loss and can be used in the event of data corruption or accidental data deletion. E. Schedule automated backups - Automated backups allow you to take regular backups of your database without manual intervention. You can use these backups to restore your database in the event of data loss or corruption.

**NEW QUESTION 77**

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 82**

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 83**

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 88**

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 92**

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

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

**Answer:** A

**NEW QUESTION 93**

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 96

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

#### NEW QUESTION 97

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 98

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 100

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 102

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

- A. Identify and optimize slow running queries, or set parallel replication flags.
- B. Stop all running queries, and re-create the replicas.
- C. Edit the primary instance to upgrade to a larger disk, and increase vCPU count.
- D. Edit the primary instance to add additional memory.

**Answer:** A

**Explanation:**

[https://cloud.google.com/sql/docs/mysql/replication/replication-lag#optimize\\_queries\\_and\\_schema](https://cloud.google.com/sql/docs/mysql/replication/replication-lag#optimize_queries_and_schema)

**NEW QUESTION 106**

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

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

**Answer:** C

**Explanation:**

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

**NEW QUESTION 109**

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 114**

You work for a large retail and ecommerce company that is starting to extend their business globally. Your company plans to migrate to Google Cloud. You want to use platforms that will scale easily, handle transactions with the least amount of latency, and provide a reliable customer experience. You need a storage layer for sales transactions and current inventory levels. You want to retain the same relational schema that your existing platform uses. What should you do?

- A. Store your data in Firestore in a multi-region location, and place your compute resources in one of the constituent regions.
- B. Deploy Cloud Spanner using a multi-region instance, and place your compute resources close to the default leader region.
- C. Build an in-memory cache in Memorystore, and deploy to the specific geographic regions where your application resides.
- D. Deploy a Bigtable instance with a cluster in one region and a replica cluster in another geographic region.

**Answer:** B

**NEW QUESTION 115**

Your organization is running a critical production database on a virtual machine (VM) on Compute Engine. The VM has an ext4-formatted persistent disk for data files. The database will soon run out of storage space. You need to implement a solution that avoids downtime. What should you do?

- A. In the Google Cloud Console, increase the size of the persistent disk, and use the `resize2fs` command to extend the disk.
- B. In the Google Cloud Console, increase the size of the persistent disk, and use the `fdisk` command to verify that the new space is ready to use
- C. In the Google Cloud Console, create a snapshot of the persistent disk, restore the snapshot to a new larger disk, unmount the old disk, mount the new disk, and restart the database service.
- D. In the Google Cloud Console, create a new persistent disk attached to the VM, and configure the database service to move the files to the new disk.

**Answer:** A

**Explanation:**

[https://cloud.google.com/compute/docs/disks/resize-persistent-disk#resize\\_partitions](https://cloud.google.com/compute/docs/disks/resize-persistent-disk#resize_partitions)

**NEW QUESTION 116**

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 instances<sup>1</sup>. PgBouncer can also improve performance and scalability by reducing the overhead of creating new connections and reusing existing ones<sup>1</sup>. 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 proxy<sup>2</sup>.

**NEW QUESTION 118**

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**

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 124**

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

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

**Answer:** A

**Explanation:**

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

**NEW QUESTION 126**

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 127

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

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

**Answer:** A

#### Explanation:

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

#### NEW QUESTION 131

Your company's mission-critical, globally available application is supported by a Cloud Spanner database. Experienced users of the application have read and write access to the database, but new users are assigned read-only access to the database. You need to assign the appropriate Cloud Spanner Identity and Access Management (IAM) role to new users being onboarded soon. What roles should you set up?

- A. roles/spanner.databaseReader
- B. roles/spanner.databaseUser
- C. roles/spanner.viewer
- D. roles/spanner.backupWriter

**Answer:** A

#### Explanation:

<https://cloud.google.com/spanner/docs/iam?hl=it>

#### NEW QUESTION 135

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 138

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 141

You work for a financial services company that wants to use fully managed database services. Traffic volume for your consumer services products has increased annually at a constant rate with occasional spikes around holidays. You frequently need to upgrade the capacity of your database. You want to use Cloud Spanner and include an automated method to increase your hardware capacity to support a higher level of concurrency. What should you do?

- A. Use linear scaling to implement the Autoscaler-based architecture
- B. Use direct scaling to implement the Autoscaler-based architecture.
- C. Upgrade the Cloud Spanner instance on a periodic basis during the scheduled maintenance window.
- D. Set up alerts that are triggered when Cloud Spanner utilization metrics breach the threshold, and then schedule an upgrade during the scheduled maintenance window.

**Answer:** A

**Explanation:**

Linear scaling is best used with load patterns that change more gradually or have a few large peaks. The method calculates the minimum number of nodes or processing units required to keep utilization below the scaling threshold. The number of nodes or processing units added or removed in each scaling event is not limited to a fixed step amount. <https://cloud.google.com/spanner/docs/autoscaling-overview#linear>

**NEW QUESTION 143**

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 148**

Your organization is currently updating an existing corporate application that is running in another public cloud to access managed database services in Google Cloud. The application will remain in the other public cloud while the database is migrated to Google Cloud. You want to follow Google-recommended practices for authentication. You need to minimize user disruption during the migration. What should you do?

- A. Use workload identity federation to impersonate a service account.
- B. Ask existing users to set their Google password to match their corporate password.
- C. Migrate the application to Google Cloud, and use Identity and Access Management (IAM).
- D. Use Google Workspace Password Sync to replicate passwords into Google Cloud.

**Answer: A**

**Explanation:**

Updating passwords represents user disruption. Eliminate B. Eliminate C for the same reason. D doesn't make sense, leaves A. From Google's documentation, "Traditionally, applications running outside Google Cloud can use service account keys to access Google Cloud resources. However, service account keys are powerful credentials, and can present a security risk if they are not managed correctly. With identity federation, you can use Identity and Access Management (IAM) to grant external identities IAM roles, including the ability to impersonate service accounts. This approach eliminates the maintenance and security burden associated with service account keys." <https://cloud.google.com/iam/docs/workload-identity-federation>

**NEW QUESTION 149**

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 152**

.....

## THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual Professional-Cloud-Database-Engineer Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the Professional-Cloud-Database-Engineer Product From:

<https://www.2passeasy.com/dumps/Professional-Cloud-Database-Engineer/>

## Money Back Guarantee

### **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 FirstTry
- \* Professional-Cloud-Database-Engineer Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year