



Google

Exam Questions Professional-Cloud-Database-Engineer

Google Cloud Certified - Professional Cloud Database Engineer

About ExamBible

Your Partner of IT Exam

Found in 1998

ExamBible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, ExamBible has its unique advantages that other companies could not achieve.

Our Advances

* 99.9% Uptime

All examinations will be up to date.

* 24/7 Quality Support

We will provide service round the clock.

* 100% Pass Rate

Our guarantee that you will pass the exam.

* Unique Gurantee

If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

NEW QUESTION 1

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-instanc>

NEW QUESTION 2

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 3

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 4

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 5

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 6

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 7

You plan to use Database Migration Service to migrate data from a PostgreSQL on-premises instance to Cloud SQL. You need to identify the prerequisites for creating and automating the task. What should you do? (Choose two.)

- A. Drop or disable all users except database administration users.
- B. Disable all foreign key constraints on the source PostgreSQL database.
- C. Ensure that all PostgreSQL tables have a primary key.
- D. Shut down the database before the Data Migration Service task is started.
- E. Ensure that pglogical is installed on the source PostgreSQL database.

Answer: CE

Explanation:

<https://cloud.google.com/database-migration/docs/postgres/faq>

NEW QUESTION 8

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 9

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 Cloud¹. Bare Metal Solution for Oracle supports Oracle RAC, which is a cluster database that provides high availability, scalability, and performance for Oracle workloads². 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 interconnectivity¹.

NEW QUESTION 10

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 10

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

- A. Store your data in Bigtable, and use the user identifier as the key.
- B. Use one column family to store user profile data, and use another column family to store user preferences.
- C. Use Cloud SQL, and create different tables for user profile data and user preferences from your recommendations mode.
- D. Use SQL to join the user profile data and preferences.
- E. Use Firestore in Native mode, and store user profile data as a document.
- F. Update the user profile with preferences specific to that user and use the user identifier to query.
- G. Use Firestore in Datastore mode, and store user profile data as a document.
- H. Update the user profile with preferences specific to that user and use the user identifier to query.

Answer: C

Explanation:

Use Firestore in Datastore mode for new server projects. Firestore in Datastore mode allows you to use established Datastore server architectures while removing fundamental Datastore limitations. Datastore mode can automatically scale to millions of writes per second. Use Firestore in Native mode for new mobile and web apps. Firestore offers mobile and web client libraries with real-time and offline features. Native mode can automatically scale to millions of concurrent clients.

NEW QUESTION 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.

Answer: D

Explanation:

<https://cloud.google.com/sql/docs/mysql/best-practices#data-arch>

NEW QUESTION 18

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 1 AM, 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 20

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 23

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 28

You are designing a database architecture for a global application that stores information about public parks worldwide. The application uses the database for read-only purposes, and a centralized batch job updates the database nightly. You want to select an open source, SQL-compliant database. What should you do?

- A. Use Bigtable with multi-region clusters.
- B. Use Memorystore for Redis with multi-zones within a region.
- C. Use Cloud SQL for PostgreSQL with cross-region replicas.
- D. Use Cloud Spanner with multi-region configuration.

Answer: C

NEW QUESTION 29

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 30

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 31

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 33

You are writing an application that will run on Cloud Run and require a database running in the Cloud SQL managed service. You want to secure this instance so that it only receives connections from applications running in your VPC environment in Google Cloud. What should you do?

- A. * 1. Create your instance with a specified external (public) IP address.* 2. Choose the VPC and create firewall rules to allow only connections from Cloud Run into your instance.* 3. Use Cloud SQL Auth proxy to connect to the instance.
- B. * 1. Create your instance with a specified external (public) IP address.* 2. Choose the VPC and create firewall rules to allow only connections from Cloud Run into your instance.* 3. Connect to the instance using a connection pool to best manage connections to the instance.
- C. * 1. Create your instance with a specified internal (private) IP address.* 2. Choose the VPC with private service connection configured.* 3. Configure the Serverless VPC Access connector in the same VPC network as your Cloud SQL instance.* 4. Use Cloud SQL Auth proxy to connect to the instance.
- D. * 1. Create your instance with a specified internal (private) IP address.* 2. Choose the VPC with private service connection configured.* 3. Configure the Serverless VPC Access connector in the same VPC network as your Cloud SQL instance.* 4. Connect to the instance using a connection pool to best manage connections to the instance.

Answer: D

Explanation:

<https://cloud.google.com/sql/docs/mysql/connect-run#configure> <https://cloud.google.com/sql/docs/mysql/connect-run#connection-pools>

NEW QUESTION 37

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 38

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 39

You are building a data warehouse on BigQuery. Sources of data include several MySQL databases located on-premises.

You need to transfer data from these databases into BigQuery for analytics. You want to use a managed solution that has low latency and is easy to set up. What should you do?

- A. Create extracts from your on-premises databases periodically, and push these extracts to Cloud Storage. Upload the changes into BigQuery, and merge them with existing tables.
- B. Use Cloud Data Fusion and scheduled workflows to extract data from MySQL.
- C. Transform this data into the appropriate schema, and load this data into your BigQuery database.
- D. Use Datastream to connect to your on-premises database and create a stream.
- E. Have Datastream write to Cloud Storage.
- F. Then use Dataflow to process the data into BigQuery.
- G. Use Database Migration Service to replicate data to a Cloud SQL for MySQL instance.
- H. Create federated tables in BigQuery on top of the replicated instances to transform and load the data into your BigQuery database.

Answer: C

NEW QUESTION 40

Your company is migrating their MySQL database to Cloud SQL and cannot afford any planned downtime during the month of December. The company is also concerned with cost, so you need the most cost-effective solution. What should you do?

- A. Open a support ticket in Google Cloud to prevent any maintenance in that MySQL instance during the month of December.
- B. Use Cloud SQL maintenance settings to prevent any maintenance during the month of December.
- C. Create MySQL read replicas in different zones so that, if any downtime occurs, the read replicas will act as the primary instance during the month of December.
- D. Create a MySQL regional instance so that, if any downtime occurs, the standby instance will act as the primary instance during the month of December.

Answer: B

Explanation:

<https://cloud.google.com/sql/docs/mysql/maintenance?hl=fr>

NEW QUESTION 45

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

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

Answer: C

Explanation:

An external replica is a method that allows you to create a read-only copy of your Cloud SQL instance on an external server, such as a Compute Engine instance or an on-premises database server. An external replica can help you scale reads and offload management operations from your data center to Google Cloud. You can also use an external replica for disaster recovery, migration, or reporting purposes.

To create an external replica, you need to configure a Cloud SQL instance that replicates to one or more replicas external to Cloud SQL, and a source representation instance that represents the source database server in Cloud SQL1. You also need to enable access on the Cloud SQL instance for the IP address of the external replica, create a replication user, and export and import the data from the source database server to the external replica1.

NEW QUESTION 50

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 55

Your organization has hundreds of Cloud SQL for MySQL instances. You want to follow Google-recommended practices to optimize platform costs. What should you do?

- A. Use Query Insights to identify idle instances.
- B. Remove inactive user accounts.
- C. Run the Recommender API to identify overprovisioned instances.
- D. Build indexes on heavily accessed tables.

Answer: C

Explanation:

The Cloud SQL overprovisioned instance recommender helps you detect instances that are unnecessarily large for a given workload. It then provides recommendations on how to resize such instances and reduce cost. This page describes how this recommender works and how to use it. <https://cloud.google.com/sql/docs/mysql/recommender-sql-overprovisioned#:~:text=The%20Cloud%20SQL%20overprovisioned%20instance%20recommender%20helps%20you%20detect%20instances%20that%20are%20unnecessarily%20large%20for%20a%20given%20workload.%20It%20then%20provides%20recommendation%20on%20how%20to%20resize%20such%20instances%20and%20reduce%20cost.%20This%20page%20describes%20how%20this%20recommender%20works%20and%20how%20to%20use%20it.>

NEW QUESTION 57

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 58

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

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

Answer: C

Explanation:

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

NEW QUESTION 61

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 65

Your team uses thousands of connected IoT devices to collect device maintenance data for your oil and gas customers in real time. You want to design inspection routines, device repair, and replacement schedules based on insights gathered from the data produced by these devices. You need a managed solution that is highly scalable, supports a multi-cloud strategy, and offers low latency for these IoT devices. What should you do?

- A. Use Firestore with Looker.
- B. Use Cloud Spanner with Data Studio.
- C. Use MongoDB Atlas with Charts.
- D. Use Bigtable with Looker.

Answer: C

Explanation:

This scenario has BigTable written all over it - large amounts of data from many devices to be analysed in realtime. I would even argue it could qualify as a multicloud solution, given the links to HBASE. BUT it does not support SQL queries and is not therefore compatible (on its own) with Looker. Firestore + Looker has the same problem. Spanner + Data Studio is at least a compatible pairing, but I agree with others that it doesn't fit this use-case - not least because it's Google-native. By contrast, MongoDB Atlas is a managed solution (just not by Google) which is compatible with the proposed reporting tool (Mongo's own Charts), it's specifically designed for this type of solution and of course it can run on any cloud.

NEW QUESTION 69

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 73

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 77

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 82

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 84

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.

Answer: A

Explanation:

In case of high CPU utilization like, mentioned in question, refer: <https://cloud.google.com/spanner/docs/identify-latency-point#:~:text=Check%20the%20CPU%20utilization%20of%20the%20instance.%20If%20the%20CPU%20utilization%20of%20the%20instance%20is%20above%20the%20recommended%20level%2C%20you%20should%20manually%20add%20more%20nodes%2C%20or%20set%20up%20auto%20scaling.> "Check the CPU utilization of the instance. If the CPU utilization of the instance is above the recommended level, you should manually add more nodes, or set up auto scaling." Indexes and schema are reviewed post identifying query with slow performance. Refer : <https://cloud.google.com/spanner/docs/troubleshooting-performance-regressions#review-schema>

NEW QUESTION 85

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 86

Your hotel booking company is expanding into Country A, where personally identifiable information (PII) must comply with regional data residency requirements and audits. You need to isolate customer data in Country A from the rest of the customer data. You want to design a multi-tenancy strategy to efficiently manage costs and operations. What should you do?

- A. Apply a schema data management pattern.
- B. Apply an instance data management pattern.
- C. Apply a table data management pattern.
- D. Apply a database data management pattern.

Answer: B

Explanation:

<https://cloud.google.com/solutions/implementing-multi-tenancy-cloud-spanner#multi-tenancy-data-management-patterns>
<https://cloud.google.com/solutions/implementing-multi-tenancy-cloud-spanner>

NEW QUESTION 87

You are developing a new application on a VM that is on your corporate network. The application will use Java Database Connectivity (JDBC) to connect to Cloud SQL for PostgreSQL. Your Cloud SQL instance is configured with IP address 192.168.3.48, and SSL is disabled. You want to ensure that your application can

access your database instance without requiring configuration changes to your database. What should you do?

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

Answer: C

Explanation:

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

NEW QUESTION 88

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 89

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 92

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 94

You are running a large, highly transactional application on Oracle Real Application Cluster (RAC) that is multi-tenant and uses shared storage. You need a solution that ensures high- performance throughput and a low-latency connection between applications and databases. The solution must also support existing Oracle features and provide ease of migration to Google Cloud. What should you do?

- A. Migrate to Compute Engine.
- B. Migrate to Bare Metal Solution for Oracle.
- C. Migrate to Google Kubernetes Engine (GKE)
- D. Migrate to Google Cloud VMware Engine

Answer: B

Explanation:

Oracle is neither licensed nor supported in GCE. The only platform which supports RAC and all existing Oracle features is BMS.

NEW QUESTION 96

You are migrating a telehealth care company's on-premises data center to Google Cloud. The migration plan specifies:

PostgreSQL databases must be migrated to a multi-region backup configuration with cross-region replicas to allow restore and failover in multiple scenarios.

MySQL databases handle personally identifiable information (PII) and require data residency compliance at the regional level.

You want to set up the environment with minimal administrative effort. What should you do?

- A. Set up Cloud Logging and Cloud Monitoring with Cloud Functions to send an alert every time a new database instance is created, and manually validate the region.
- B. Set up different organizations for each database type, and apply policy constraints at the organization level.
- C. Set up Pub/Sub to ingest data from Cloud Logging, send an alert every time a new database instance is created, and manually validate the region.
- D. Set up different projects for PostgreSQL and MySQL databases, and apply organizational policy constraints at a project level.

Answer: D

NEW QUESTION 97

.....

Relate Links

100% Pass Your Professional-Cloud-Database-Engineer Exam with Exambible Prep Materials

<https://www.exambible.com/Professional-Cloud-Database-Engineer-exam/>

Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - <https://www.exambible.com/>