Google

Exam Questions Professional-Cloud-Database-Engineer

Google Cloud Certified - Professional Cloud Database Engineer

NEW QUESTION 1
You released a popular mobile game and are using a 50 TB Cloud Spanner instance to store game data in a PITR-enabled production environment. When you analyzed the game statistics, you realized that some players are exploiting a loophole to gather more points to get on the leaderboard. Another DBA accidentally ran an emergency bugfix script that corrupted some of the data in the production environment. You need to determine the extent of the data corruption and restore the production environment. What should you do? (Choose two.)

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

Answer: AE

Explanation:
https://cloud.google.com/spanner/docs/pitr#ways-to-recover
To recover the entire database, backup or export the database specifying a timestamp in the past and then restore or import it to a new database. This is typically used to recover from data corruption issues when you have to revert the entire database to a point-in-time before the corruption occurred.
This part describes significant corruption - A
To recover a portion of the database, perform a stale read specifying a query-condition and timestamp in the past, and then write the results back into the live database. This is typically used for surgical operations on a live database. For example, if you accidentally delete a particular row or incorrectly update a subset of data, you can recover it with this method.
This describes insignificant corruption case – E https://cloud.google.com/spanner/docs/pitr https://cloud.google.com/spanner/docs/backup/restore-backup

NEW QUESTION 2
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 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 4
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 5
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 6
You are a DBA on a Cloud Spanner instance with multiple databases. You need to assign these privileges to all members of the application development team on a specific database:
Can read tables, views, and DDL Can write rows to the tables
Can add columns and indexes Cannot drop the database What should you do?

A. Assign the Cloud Spanner Database Reader and Cloud Spanner Backup Writer roles.
B. Assign the Cloud Spanner Database Admin role.
C. Assign the Cloud Spanner Database User role.
D. Assign the Cloud Spanner Admin role.

Answer: C

Explanation:
https://cloud.google.com/spanner/docs/iam#spanner.databaseUser

NEW QUESTION 7
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, andthen 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 8
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

NEW QUESTION 9
You are managing a Cloud SQL for PostgreSQL instance in Google Cloud. You need to test the high availability of your Cloud SQL instance by performing a failover. You want to use the cloud command.
What should you do?

A. Use gcloud sql instances failover <PrimaryInstanceName>.
B. Use gcloud sql instances failover <ReplicaInstanceName>.
C. Use gcloud sql instances promote-replica <PrimaryInstanceName>.
D. Use gcloud sql instances promote-replica <ReplicaInstanceName>.

Answer: A

NEW QUESTION 10
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 10
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 15
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 18
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 ke
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 documen
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 documen
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 20
Your company wants you to migrate their Oracle, MySQL, Microsoft SQL Server, and PostgreSQL relational databases to Google Cloud. You need a fully managed, flexible database solution when possible. What should you do?

A. Migrate all the databases to Cloud SQL.
B. Migrate the Oracle, MySQL, and Microsoft SQL Server databases to Cloud SQL, and migrate the PostgreSQL databases to Compute Engine.
C. Migrate the MySQL, Microsoft SQL Server, and PostgreSQL databases to Compute Engine, and migrate the Oracle databases to Bare Metal Solution for Oracle.
D. Migrate the MySQL, Microsoft SQL Server, and PostgreSQL databases to Cloud SQL, and migrate the Oracle databases to Bare Metal Solution for Oracle.

Answer: D

NEW QUESTION 23
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 27

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 32
You are migrating your data center to Google Cloud. You plan to migrate your applications to Compute Engine and your Oracle databases to Bare Metal Solution for Oracle. You must ensure that the applications in different projects can communicate securely and efficiently with the Oracle databases. What should you do?

A. Set up a Shared VPC, configure multiple service projects, and create firewall rules.
B. Set up Serverless VPC Access.
C. Set up Private Service Connect.
D. Set up Traffic Director.

Answer: A

Explanation:
https://medium.com/google-cloud/shared-vpc-in-google-cloud-64527e0a409e#:~:text=Unlike%20VPC%20peering%2C%20Shared%20VPC%20connects%20projects%20within%20the%20same%20organization.&text=There%20are%20a%20lot%20of,between%20VPCs%20in%20different%20projects.

NEW QUESTION 35
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 configuring the networking of a Cloud SQL instance. The only application that connects to this database resides on a Compute Engine VM in the same project as the Cloud SQL instance. The VM and the Cloud SQL instance both use the same VPC network, and both have an external (public) IP address and an internal (private) IP address. You want to improve network security. What should you do?

A. Disable and remove the internal IP address assignment.
B. Disable both the external IP address and the internal IP address, and instead rely on Private Google Access.
C. Specify an authorized network with the CIDR range of the VM.
D. Disable and remove the external IP address assignment.

Answer: D

Explanation:
It is always more secure to use an internal IP, so removing them doesn't make sense. Eliminate A. You can use Private Google Access when VM instances only have internal IP addresses, so disabling the internal IPs and use Private Google Access doesn't make sense. Eliminate B. Specifying an authorized network when they're on the same subnet doesn't make sense. Eliminate C. A way to improve network security would be to disable external IPs since they're not needed.

NEW QUESTION 43
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 47
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 48
Your organization is migrating 50 TB Oracle databases to Bare Metal Solution for Oracle. Database backups must be available for quick restore. You also need to have backups available for 5 years. You need to design a cost-effective architecture that meets a recovery time objective (RTO) of 2 hours and recovery point objective (RPO) of 15 minutes. What should you do?

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

Answer: B

Explanation:
This answer meets the RTO and RPO requirements by using flash storage for the database and standard storage for the local backup copy. It also meets the cost- effectiveness requirement by using Actifio OnVault storage, which is a low-cost, high- performance, and scalable storage solution that integrates with Google Cloud Backup and DR Service1.
References := 1: Solution Guide: Google Cloud Backup and DR for Oracle on Bare Metal Solution1

NEW QUESTION 52
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 57
Your team uses thousands of connected IoT devices to collect device maintenance data for your oil and gas customers in real time. You want to design inspection routines, device repair, and replacement schedules based on insights gathered from the data produced by these devices. You need a managed solution that is highly scalable, supports a multi-cloud strategy, and offers low latency for these IoT devices. What should you do?

A. Use Firestore with Looker.
B. Use Cloud Spanner with Data Studio.
C. Use MongoD8 Atlas with Charts.
D. Use Bigtable with Looker.

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 59
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 64
You are choosing a database backend for a new application. The application will ingest data points from IoT sensors. You need to ensure that the application can scale up to millions of requests per second with sub-10ms latency and store up to 100 TB of history. What should you do?

A. Use Cloud SQL with read replicas for throughput.
B. Use Firestore, and rely on automatic serverless scaling.
C. Use Memorystore for Memcached, and add nodes as necessary to achieve the required throughput.
D. Use Bigtable, and add nodes as necessary to achieve the required throughput.

Answer: D

Explanation:
https://cloud.google.com/memorystore/docs/redis/redis-overview

NEW QUESTION 65
You are migrating an on-premises application to Compute Engine and Cloud SQL. The application VMs will live in their own project, separate from the Cloud SQL instances which have their own project. What should you do to configure the networks?

A. Create a new VPC network in each project, and use VPC Network Peering to connect the two together.
B. Create a Shared VPC that both the application VMs and Cloud SQL instances will use.
C. Use the default networks, and leverage Cloud VPN to connect the two together.
D. Place both the application VMs and the Cloud SQL instances in the default network of each project.

Answer: B

Explanation:
https://groups.google.com/g/google-cloud-sql-discuss/c/M5G5_HPXytY?pli=1

NEW QUESTION 67
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 68
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 70
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 73
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 77
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 82
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 85
Your organization has a busy transactional Cloud SQL for MySQL instance. Your analytics team needs access to the data so they can build monthly sales reports. You need to provide data access to the analytics team without adversely affecting performance. What should you do?

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

Answer: B

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

NEW QUESTION 88
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. Identity 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 90
.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questons and Answers in PDF Format

Professional-Cloud-Database-Engineer Practice Exam Features:

* Professional-Cloud-Database-Engineer Questions and Answers Updated Frequently

* Professional-Cloud-Database-Engineer Practice Questions Verified by Expert Senior Certified Staff

* Professional-Cloud-Database-Engineer Most Realistic Questions that Guarantee you a Pass on Your FirstTry

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