



- Expert Verified, Online, **Free**.



CERTIFICATION TEST


- CertificationTest.net - Cheap & Quality Resources With Best Support

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.

Suggested Answer: C

Community vote distribution



Hilab Highly Voted 2 years, 4 months ago

Option B is the best choice.

By using the internal (private) IP address of the Cloud SQL instance, the traffic will stay within the corporate network and will not traverse the public internet. This will help to ensure that the traffic is secure and cannot be intercepted by unauthorized parties. Additionally, using the internal IP address does not require any additional configuration changes to the database instance.

Option A is not recommended as it requires exposing the database instance's external (public) IP address, which can be less secure and may require additional firewall rules.

Option C is a valid option if SSL is enabled on the Cloud SQL instance, but since SSL is disabled in this scenario, this option is not suitable.

Option D is not recommended as it requires exposing the database instance's external (public) IP address, which can be less secure and may require additional firewall rules.

upvoted 8 times

BenMS 2 years, 3 months ago

The Cloud SQL Proxy wraps your connection in an SSL/TLS layer, resolving the concern about compatibility:

<https://cloud.google.com/sql/docs/postgres/connect-auth-proxy>

upvoted 1 times

mennahibi 1 year, 7 months ago

B can be the correct answer, but the most secure and best solution is C, because the Auth Proxy will enable ssl for you without enabling it on the Cloud SQL instance.

upvoted 1 times

dynamic_dba Highly Voted 2 years, 3 months ago

C.

The IP address given is a private IP address and not routable via the internet. Therefore any answer which references a public IP is wrong by definition (A, D). That leaves B and C. B cannot be correct because the app is on a corporate network and thus not on a Google VPC network. Good security practices dictate using Cloud SQL Auth Proxy and a service account which access the Cloud SQL instance via its private IP address.

upvoted 5 times

ccpmad Most Recent 5 months ago

Selected Answer: C

I did the exam on January 2025, and all set of questions were here.

upvoted 1 times

Zakky_09 6 months, 1 week ago

Selected Answer: B

The Cloud SQL instance is configured with a private IP address (192.168.3.48), and your application is on the same corporate network. Therefore, the private IP address is accessible directly, and there is no need to use the Cloud SQL Auth proxy or external (public) IP address.

upvoted 1 times

🗲️ 👤 **CloudKida** 9 months ago

Selected Answer: C

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.

upvoted 4 times

🗲️ 👤 **hanayome** 1 year, 1 month ago

Selected Answer: C

because the most secure way is using Cloud SQL Proxy

upvoted 1 times

🗲️ 👤 **Witcher_PL** 1 year, 8 months ago

Selected Answer: C

C is the valid

upvoted 3 times

🗲️ 👤 **goodsport** 1 year, 9 months ago

Vote for C

upvoted 2 times

🗲️ 👤 **Nirca** 2 years, 3 months ago

Selected Answer: C

Service account is a must.

upvoted 1 times

🗲️ 👤 **__pacman__** 2 years, 4 months ago

Vote for C

upvoted 2 times

🗲️ 👤 **ssaporylo** 2 years, 5 months ago

Vote for C

upvoted 4 times

🗲️ 👤 **chelbsik** 2 years, 6 months ago

Selected Answer: C

Vote for C

upvoted 3 times

🗲️ 👤 **pk349** 2 years, 6 months ago

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.

upvoted 3 times

🗲️ 👤 **pk349** 2 years, 6 months ago

Selected Answer: D

Database Migration Service

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

upvoted 2 times

🗲️ 👤 **GCP72** 2 years, 6 months ago

Selected Answer: C

C is the correct answer

upvoted 2 times

🗨️ 👤 **Kloudgeek** 2 years, 6 months ago

C is correct answer. First of all SSL is disabled and it is not secure to get it exposed to Internet.

https://cloud.google.com/sql/docs/postgres/connect-overview#authentication_options

upvoted 4 times

🗨️ 👤 **juancambb** 2 years, 6 months ago

Selected Answer: C

C is correct, must be private ip because the ip starts with 192... and cloud sql require a proxy to connect because exist on a tenant project

upvoted 2 times

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.

Suggested Answer: BD

Community vote distribution

DE (100%)

  **dynamic_dba** Highly Voted 1 year, 10 months ago

D,E.

A is wrong because why bother configuring manual backups when Cloud SQL will automate that for you.

B seems attractive, but why bother replicating back to on-prem when you can configure a Cloud SQL for HA.

C is wrong because a single zone failure would not give you HA.

That leaves D & E.

upvoted 7 times

  **Zakky_09** Most Recent 6 months, 1 week ago

Selected Answer: DE

To ensure high availability (HA) for your Cloud SQL instance and adhere to Google-recommended practices, you need to prioritize reliable backups and recovery mechanisms.


D. Enable point-in-time recovery:

This allows you to recover your database to a specific point in time in case of accidental data deletion or corruption. It ensures minimal data loss and quick recovery.

E. Schedule automated backups:

Automated backups ensure regular snapshots of your database are taken without manual intervention, providing a reliable recovery point in case of failures.


upvoted 1 times

  **dija123** 7 months, 2 weeks ago

Selected Answer: DE

Agree with D, E

upvoted 2 times

  **hanayome** 7 months, 3 weeks ago

Selected Answer: DE



D and E because point in time recovery and schedule backup are HA efforts and CloudSQL already have these features

upvoted 2 times

  **goodsport** 1 year, 3 months ago

Vote for D, E.

upvoted 2 times

  **Nirca** 1 year, 10 months ago

Selected Answer: DE

Only valid options.

upvoted 3 times

  **chelbsik** 2 years ago

Selected Answer: DE

Vote for DE, seems only reasonable options to me

upvoted 4 times

  **omermahgoub** 2 years ago


To prepare your Cloud SQL instance for high availability, you should do the following:

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.

Note that options A and C are not recommended practices for high availability. Option B is not related to Cloud SQL.

upvoted 3 times

  **GCP72** 2 years ago

Selected Answer: DE

D,E is the correct answer

upvoted 3 times

  **range9005** 2 years ago

Selected Answer: DE

Enable point-in-time recovery.

Schedule automated backups.

upvoted 3 times

  **H_S** 2 years ago

Selected Answer: DE

D. Enable point-in-time recovery.

E. Schedule automated backups

upvoted 1 times

  **omermahgoub** 2 years ago

D. Enable point-in-time recovery.

E. Schedule automated backups.

To prepare your Cloud SQL instance for high availability, Google recommends enabling point-in-time recovery and scheduling automated backups.

Point-in-time recovery allows you to restore your database to a specific point in time, helping you to recover from data loss or corruption.

Scheduling automated backups ensures that you have a recent copy of your database available for recovery in case of an outage or other issue.

upvoted 2 times

  **omermahgoub** 2 years ago

Option A, setting up manual backups, would not be a recommended practice because manual backups are prone to errors and can be time-consuming to create and maintain. Automated backups are a more reliable and efficient way to ensure that you have a recent copy of your database available for recovery.

Option B, creating a PostgreSQL database on-premises as the HA option, would not be a recommended practice because it would not take advantage of the high availability features provided by Cloud SQL.

Option C, configuring single zone availability for automated backups, would not be a recommended practice because it would not provide sufficient protection against outages or other issues. To ensure high availability, it is recommended to use a multi-zone configuration for your Cloud SQL instance.

upvoted 1 times

  **Popa** 2 years ago

Selected Answer: DE

<https://cloud.google.com/sql/docs/mysql/high-availability#backups-and-restores>

upvoted 2 times

  **Popa** 2 years ago

Answer is D and E. <https://cloud.google.com/sql/docs/mysql/high-availability#backups-and-restores>

upvoted 1 times

 **range9005** 2 years ago

Selected Answer: DE

Automated backups and point-in-time recovery must be enabled for high availability (point-in-time recovery uses binary logs).
For more information check here -> <https://cloud.google.com/sql/docs/mysql/high-availability#backups-and-restores>

upvoted 2 times

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).

Suggested Answer: A

Community vote distribution

A (100%)

🗳️ 👤 **Zakky_09** 6 months, 1 week ago

Selected Answer: A

Google Cloud VMware Engine allows you to run your VMware workloads natively on Google Cloud without requiring changes to your applications or infrastructure. This is ideal for minimizing disruption and maintaining compatibility with your current Oracle application footprint.

upvoted 1 times

🗳️ 👤 **sky09** 6 months, 2 weeks ago

Selected Answer: A

Google Cloud VMware Engine (GCVE) provides a fully compatible VMware environment in the cloud.

upvoted 1 times

🗳️ 👤 **hanayome** 7 months, 3 weeks ago

Selected Answer: A

obviously A

upvoted 2 times

🗳️ 👤 **goodsport** 1 year, 3 months ago

Selected Answer: A

Lift , shift. Correct andswer is A.

upvoted 3 times

🗳️ 👤 **Nirca** 1 year, 10 months ago

Selected Answer: A

Classic lift and lift. Everything keeps the same structure. Therefore minimizing impact to zero.

upvoted 3 times

🗳️ 👤 **dynamic_dba** 1 year, 10 months ago

A.

The key here is the current architecture and minimal disruption to it. The simplest way to keep the current architecture is a live migrate using VMware. That can only mean one thing, use Oracle running in GCVE.

You could do B. There's nothing stopping you creating a VM in GCE, copying the Oracle binaries to it and spinning up an Oracle database or several. However, the licensing costs would not be attractive (if even supported), plus the migration would likely be disruptive. C is wrong because Cloud SQL doesn't support Oracle. D is wrong because that represents an architecture change.

upvoted 4 times

🗳️ 👤 **chelbsik** 2 years ago

Selected Answer: A

Since there is no Bare Metal for Oracle option and VMware mentioned -> Choose VMware

upvoted 4 times

🗳️ 👤 **Popa** 2 years ago

Selected Answer: A

Here is the explanation: <https://cloud.google.com/blog/products/databases/migrate-databases-to-google-cloud-vmware-engine-gcve>

upvoted 2 times

🗨️ 👤 **range9005** 2 years ago

Selected Answer: A

A GCVE VMware environment runs natively on Google Cloud bare metal infrastructure in some Google Cloud locations, and the GCVE service includes all the features required to help consume the VMware platforms efficiently and securely

.

<https://cloud.google.com/blog/products/databases/migrate-databases-to-google-cloud-vmware-engine-gcve>

upvoted 1 times

🗨️ 👤 **fredcaram** 2 years ago

Oracle databases can only be migrated to bare metal solutions

upvoted 1 times

🗨️ 👤 **GCP72** 2 years ago

Yes .GCP recommended to use Bare Metal solution for Oracle but option is missing in answers

upvoted 1 times

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.

Suggested Answer: B

Community vote distribution

C (100%)

🗳️ **jnya_1991** Highly Voted 1 year, 6 months ago

Selected Answer: C

Read statistics provide insight into how an application is using the database, and are useful when investigating performance issues.

<https://cloud.google.com/spanner/docs/introspection/read-statistics>

upvoted 5 times

🗳️ **Ral17** Most Recent 6 months ago

Selected Answer: C

SPANNER_SYS.READ_STATS* contains statistics about reads.

upvoted 2 times

🗳️ **Zakky_09** 6 months, 1 week ago

Selected Answer: D

When experiencing high read latency in a Cloud Spanner instance, analyzing the SPANNER_SYS.QUERY_STATS tables provides valuable insights into the performance of your queries. These tables contain statistics about query execution, such as latency, execution count, and resource usage. This information helps identify slow or inefficient queries that may be causing the performance degradation

upvoted 1 times

🗳️ **sky09** 6 months, 2 weeks ago

Selected Answer: D

SPANNER_SYS.QUERY_STATS tables: the first step in diagnosing latency after a code change is often to look at the queries themselves so ans is D

<https://cloud.google.com/spanner/docs/introspection/query-statistics>

upvoted 1 times

🗳️ **theseawillclaim** 1 year, 9 months ago

C! You should analyze the situation before changing the architecture so drastically.

upvoted 2 times

🗳️ **kpkakadiya** 1 year, 9 months ago

Selected Answer: C

C is the correct answer

upvoted 2 times

🗳️ **goodsport** 1 year, 9 months ago

C is definitely the correct answer here. SPANNER_SYS.READ_STATS* contains statistics about reads.

upvoted 2 times

🗳️ **RahulHanumante** 1 year, 10 months ago

C is the correct answer

upvoted 2 times

🗳️ **jamalkhan** 2 years ago

Selected Answer: C

C. Read stats

upvoted 3 times

🗨️ **dynamic_dba** 2 years, 3 months ago

C.

A Query parameters is vague at best. B would not achieve anything. C and D look interesting, but as others have stated, querying the READ_STATS* tables would give you information about what is causing read issues. So C is the best answer.

upvoted 3 times

🗨️ **H_S** 2 years, 3 months ago

Selected Answer: C

C. Use SQL statements to analyze SPANNER_SYS.READ_STATS* tables

upvoted 4 times

🗨️ **H_S** 2 years, 3 months ago

Selected Answer: C

C. Use SQL statements to analyze SPANNER_SYS.READ_STATS* tables

upvoted 3 times

🗨️ **omermahgoub** 2 years, 6 months ago

C. Use SQL statements to analyze SPANNER_SYS.READ_STATS* tables.

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.

upvoted 4 times

🗨️ **pk349** 2 years, 6 months ago

C: Use SQL statements to analyze ***** SPANNER_SYS.READ_STATS* tables.

upvoted 3 times

🗨️ **GCP72** 2 years, 6 months ago

Selected Answer: C

C is the correct answer

upvoted 4 times

🗨️ **yylbgevkujgphocvyh** 2 years, 6 months ago

B - Is correct

upvoted 1 times

🗨️ **range9005** 2 years, 6 months ago

Selected Answer: C

Read statistics provide insight into how an application is using the database, and are useful when investigating performance issues

https://cloud.google.com/spanner/docs/introspection/read-statistics#when_to_use_read_statistics


upvoted 3 times

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.

Suggested Answer: B

Community vote distribution



🗳️ 👤 **Zakky_09** 6 months, 1 week ago

Selected Answer: A

Google-recommended practices for migrating databases to Cloud SQL emphasize using the Database Migration Service (DMS) because it is a native, managed tool that supports:

Minimal downtime migrations: It supports continuous data replication using change data capture (CDC).

Ease of use: DMS is fully integrated with Google Cloud, making it straightforward to configure and monitor migrations.

Monitoring and reliability: It provides tools for monitoring the migration progress and detecting issues.

By using DMS for both one-time and ongoing replication needs, you align with Google's best practices while reducing the complexity of using multiple tools.

upvoted 1 times

🗳️ 👤 **studymoreoften** 7 months ago

Selected Answer: A

database migration service can be used to migrate from on-prem and other clouds (AWS)

upvoted 2 times

🗳️ 👤 **Killerbee05** 9 months, 2 weeks ago

Selected Answer: A

A is a correct answer

upvoted 2 times

🗳️ 👤 **whoosh** 1 year ago

Selected Answer: A

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

upvoted 3 times

🗳️ 👤 **jnya_1991** 1 year ago

Selected Answer: A

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

upvoted 2 times

🗳️ 👤 **goodsport** 1 year, 3 months ago

Selected Answer: A

A is the correct answer here.

upvoted 3 times

🗳️ 👤 **somnathmaddi** 1 year, 7 months ago

Selected Answer: A

The question says to use Google native data migration tools
upvoted 3 times

  **nmm22** 9 months ago

good point, thank you for pointing it out
upvoted 1 times

  **cloudkoala** 1 year, 9 months ago

Selected Answer: A

DMS will do the CDC too.
upvoted 4 times

  **dynamic_dba** 1 year, 9 months ago



A.

The question says to use Google native data migration tools. That eliminates B and D. C doesn't specify the data replication tool in question so it's a reasonable assumption its referring to database native replication which wouldn't be a Google native solution. That eliminates C. That leave A.
upvoted 3 times


  **Nirca** 1 year, 10 months ago

Selected Answer: A

DMS will do the job. For the init time and the CDC phase
upvoted 1 times

  **Hilab** 1 year, 10 months ago

Option A is the most straightforward and recommended solution for migrating PostgreSQL databases to Cloud SQL while following Google-recommended practices and using native data migration tools.
upvoted 1 times

  **realvarex** 1 year, 11 months ago

Selected Answer: A

A
<https://cloud.google.com/blog/products/databases/tips-for-migrating-across-compatible-database-engines>
upvoted 1 times



  **chelbsik** 2 years ago

Selected Answer: A

A is enough
upvoted 4 times



  **pk349** 2 years ago

A: Use Database Migration Service to migrate all databases to Cloud SQL.
upvoted 1 times

  **GCP72** 2 years ago

Selected Answer: C

C is the correct answer, we can use GCP migration tool for onetime or CDC
upvoted 1 times

  **GCP72** 1 year, 11 months ago

My bad ,A is correct answer
upvoted 2 times

  **H_S** 2 years ago

Selected Answer: A

A for live magrations
upvoted 1 times

  **range9005** 2 years ago

Selected Answer: A


Migrate to Cloud SQL and AlloyDB for PostgreSQL from on-premises, Google Cloud, or other clouds
upvoted 2 times

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.

Suggested Answer: C

Community vote distribution



🗲️ 👤 **JustANick** Highly Voted 1 year, 4 months ago

Selected Answer: C

<https://cloud.google.com/bare-metal/docs/bms-setup?hl=en#bms-access-options>

It can't be D. This doc mentions NAT gateway, not Cloud NAT.

upvoted 6 times

🗲️ 👤 **MikeFR** Most Recent 1 month ago

Selected Answer: C

looks like the most appropriate answer

upvoted 1 times

🗲️ 👤 **hmarine** 3 months, 1 week ago

Selected Answer: D

A Cloud NAT gateway is not installed on a compute engine VM. Cloud NAT is a managed service. Cloud NAT gateways are a component of the cloud NAT service.

Therefore, D is the most appropriate and efficient solution.

upvoted 2 times

🗲️ 👤 **Zakky_09** 6 months, 1 week ago

Selected Answer: D

The Bare Metal Solution environment does not have public IP addresses for security reasons, so it requires a Cloud NAT (Network Address Translation) service to enable secure outbound internet access. Cloud NAT allows private resources, such as the Bare Metal Solution environment, to connect to the internet for software updates without exposing them to the internet.

upvoted 1 times

🗲️ 👤 **sky09** 6 months, 2 weeks ago

Selected Answer: D

According to the official documentation provided at Bare Metal Solution setup, the recommended method for providing internet access to your Bare Metal Solution (BMS) environment—particularly for downloading operating system updates—is to use a Cloud NAT gateway.

upvoted 1 times

🗲️ 👤 **Pime13** 8 months ago

Selected Answer: C

<https://cloud.google.com/bare-metal/docs/bms-setup#bms-access-options>

Note: Cloud NAT feature doesn't support transitive endpoints thus it can not be used standalone to provide the internet access to the BMS server. Compute Engine VM must be used along with Cloud NAT.

upvoted 4 times

🗲️ 👤 **Jason_Cloud_at** 10 months, 2 weeks ago

Selected Answer: D

Knows the difference between Cloud NAT Services and Cloud NAT Gateway,

Cloud NAT services - lets your VMs and container pods create outbound connections to the internet or to other Virtual Private Cloud (VPC) networks.

Cloud NAT Gateway - Cloud NAT uses NAT gateway to manage the connections.

Also, Cloud NAT gateway is region and VPC network specific, we can use cloud NAT mapping to a VM instance , not the gateway itself.

upvoted 1 times

  **Jason_Cloud_at** 10 months, 1 week ago

Looks like we can map NAT gateway on GCE, will change to C.

upvoted 1 times

  **PKookNN** 11 months ago

Selected Answer: C

<https://cloud.google.com/bare-metal/docs/bms-setup?hl=en#bms-access-internet-vm-nat>

offers 3 options all of them involve Compute Engine VM and CloudNAT. so C

upvoted 1 times

  **whoosh** 1 year ago

Selected Answer: C


C - The following instructions set up a NAT gateway on a Compute Engine VM to connect the servers in a Bare Metal Solution environment to the internet for purposes such as receiving software updates

upvoted 1 times

  **Artists** 1 year ago

C - The following instructions set up a NAT gateway on a Compute Engine VM to connect the servers in a Bare Metal Solution environment to the internet for purposes such as receiving software updates

upvoted 1 times

  **julioevk** 1 year, 3 months ago

Selected Answer: C

C - <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.

upvoted 3 times

  **goodsport** 1 year, 3 months ago

Selected Answer: C

Voting for C.

upvoted 1 times

  **learnazureportal** 1 year, 3 months ago

Correct answer is D. Option C (setting up a Cloud NAT gateway on a Compute Engine VM) is not a recommended approach for providing internet access to your Bare Metal Solution environment

upvoted 2 times

  **omermahgoub** 1 year, 4 months ago

Selected Answer: D

D:

Cloud NAT is a network address translation (NAT) service that allows you to connect your Bare Metal Solution environment to the internet without having to assign a public IP address to each machine. This is the best option for you because it is the most secure and easiest way to connect your Bare Metal Solution environment to the internet.

<https://cloud.google.com/bare-metal/docs/bms-setup#bms-access-options>

upvoted 1 times

  **JustANick** 1 year, 4 months ago

In this link, there is no mention to Cloud NAT

upvoted 1 times

  **[Removed]** 1 year, 4 months ago

Selected Answer: C

The BMS environment traffic originates outside the VPC, so Cloud NAT can not work. C is the correct one

upvoted 1 times

  **nescafe7** 1 year, 5 months ago

Selected Answer: D

I agree D is the simplest option.

upvoted 1 times

🗨️ 👤 **BenMS** 1 year, 9 months ago

Selected Answer: D

<https://cloud.google.com/bare-metal/docs/bms-setup#bms-access-internet>

The BMS documentation mentions the Cloud NAT service as an option but the provided example involves manually deploying a NAT gateway on a GCE machine, without explaining why you would need this option as opposed to the managed NAT service. However there are no limitations mentioned, so I take it both options are valid.

In this question, there is no mention of an existing GCE machine, therefore a managed NAT service is the simplest option, which avoids additional infrastructure - hence D is my choice.

upvoted 2 times

Your organization is running a MySQL workload in Cloud SQL. Suddenly you see a degradation in database performance. You need to identify the root cause of the performance degradation. What should you do?

- A. Use Logs Explorer to analyze log data.
- B. Use Cloud Monitoring to monitor CPU, memory, and storage utilization metrics.
- C. Use Error Reporting to count, analyze, and aggregate the data.
- D. Use Cloud Debugger to inspect the state of an application.

Suggested Answer: B

Community vote distribution

B (100%)

  **Zakky_09** 6 months, 1 week ago

Selected Answer: B


Performance degradation in a Cloud SQL MySQL instance is often tied to resource utilization. Cloud Monitoring provides real-time insights into CPU, memory, and storage usage, which are critical factors that affect database performance. Analyzing these metrics helps identify issues such as:

High CPU usage due to inefficient queries or excessive connections.

Low available memory causing performance bottlenecks.

Storage nearing capacity or IOPS limitations impacting performance.

upvoted 1 times

  **dynamic_dba** 9 months, 3 weeks ago

B.

No actual errors are mentioned so using Error reporting would be irrelevant. That eliminates C. Inspecting the state of an application is also irrelevant since so mention of any application changes is made. Eliminate D. That leave A and B and B is the best answer. In Cloud SQL you get monitoring built right in (which you don't by default with GCE VMs). Cloud SQL monitoring metrics include CPU utilization, storage usage, memory usage, r/w operations and egress/ingress bytes. Has to be B.

upvoted 3 times

  **pk349** 1 year ago

B: Use Cloud Monitoring ***** to monitor CPU, memory, and storage utilization metrics.

upvoted 3 times

  **GCP72** 1 year ago

Selected Answer: B

B is the correct answer

upvoted 2 times

  **chelbsik** 1 year ago

Selected Answer: B

If your instance stops responding to connections or performance is degraded, make sure it conforms to the Operational Guidelines

[https://cloud.google.com/sql/docs/mysql/diagnose-](https://cloud.google.com/sql/docs/mysql/diagnose-issues#:~:text=If%20your%20instance%20stops%20responding%20to%20connections%20or%20performance%20is%20degraded%2C%20make%20sure%20it%20conforms%20to%20the%20Operational%20Guidelines)

issues#:~:text=If%20your%20instance%20stops%20responding%20to%20connections%20or%20performance%20is%20degraded%2C%20make%20sure%20it%20conforms%20to%20the%20Operational%20Guidelines

And then checking resource constraints:

Storage full

CPU overloaded

Too many database tables

https://cloud.google.com/sql/docs/mysql/operational-guidelines#resource_constraints

Cloud Monitoring seems like the only way to check 2/3 of those, so for me answer is B

upvoted 3 times

🗨️ 👤 **H_S** 1 year ago

Selected Answer: B

B. Use Cloud Monitoring to monitor CPU, memory, and storage utilization metrics.

upvoted 2 times

🗨️ 👤 **jitu028** 1 year ago

Selected Answer: B

Correct Answer - B

upvoted 2 times

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.

Suggested Answer: B

Community vote distribution

B (100%)

🗳️ 👤 **Zakky_09** 6 months, 1 week ago

Selected Answer: B

Cloud Spanner is the ideal choice for this scenario because:

Global scalability and low latency: Cloud Spanner is designed for globally distributed applications and ensures consistency across multiple regions with minimal latency.

Relational schema: It supports relational schemas and SQL queries, making it suitable for applications requiring traditional relational database structures.

High availability: Multi-region instances provide automatic failover and data replication for reliability and disaster recovery.

Placing compute resources close to the default leader region minimizes latency for write operations, ensuring optimal performance for transactions.
upvoted 1 times

🗳️ 👤 **hanayome** 1 year, 1 month ago

Selected Answer: B

Obviously B

upvoted 2 times

🗳️ 👤 **julioirevk** 1 year, 9 months ago

Selected Answer: B

B and Spanner as soon as it says Global and relational

upvoted 2 times

🗳️ 👤 **goodsport** 1 year, 9 months ago

Selected Answer: B

B, spanner.

upvoted 2 times

🗳️ 👤 **somnathmaddi** 2 years, 1 month ago

Selected Answer: B

It's B. Spanner

upvoted 3 times

🗳️ 👤 **felipeschossler** 2 years, 2 months ago

Selected Answer: B

Global, scale easily and keeping the relation schema. It's B. Spanner. There is no other option.

upvoted 1 times

🗳️ 👤 **dynamic_dba** 2 years, 3 months ago

B.

The clues are "globally" and "relational schema". Relational rules out Firestore (A) and Bigtable (D). Cloud Spanner is both global in scale and relational, so it fits. So B.

upvoted 4 times

🗨️ 👤 **Nirca** 2 years, 3 months ago

Selected Answer: B

Spanner is the right answer

upvoted 2 times

🗨️ 👤 **pk349** 2 years, 6 months ago

B: Deploy Cloud Spanner *** using a multi-region instance, and place your compute resources close to the default leader region.

upvoted 1 times

🗨️ 👤 **GCP72** 2 years, 6 months ago

Selected Answer: B

B is the correct answer

upvoted 1 times

🗨️ 👤 **chelbsik** 2 years, 6 months ago

Selected Answer: B

Spanner seems to be the only option, since it's the only relational DB. Plus "scale easily", which is another clear indication of Spanner.


upvoted 2 times

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. Control the order of update by setting non-production instances to earlier and production instances to later.
- B. Create your database with one primary node and one read replica in the region.
- C. Enable maintenance notifications for users, and reschedule maintenance activities to a specific time after notifications have been sent.
- D. Configure your Cloud SQL instance with high availability enabled.

Suggested Answer: A

Community vote distribution



chelbsik Highly Voted 2 years, 6 months ago

Selected Answer: A

Since we don't really need HA and we have a window that users are not need our app - A is fine, and D looks like an overkill
upvoted 10 times

orvalver Most Recent 6 months, 3 weeks ago

Selected Answer: A

The key term here is the maintenance window, if we see MW we can work to configure it. Letter A.
upvoted 1 times

hanayome 1 year, 1 month ago

Selected Answer: A

A because maintenance should be on the time with minimal impact
upvoted 2 times

goodsport 1 year, 9 months ago

Selected Answer: A

A is the correct answer.
upvoted 3 times

dynamic_dba 2 years, 3 months ago

A.
Google controls maintenance which could cause some downtime. Hence D would be irrelevant. C seems like a lot of work. B is also irrelevant. That leaves A as the best answer since you can choose your maintenance window to be after users will not be using the system. The addition of the earlier and later information is fluff and is not relevant to the question.
upvoted 4 times

Nirca 2 years, 3 months ago

Selected Answer: A

A is right.
upvoted 1 times

ssaporylo 2 years, 5 months ago

Vote for A. Configure time slot for maintainance
HA for fail over but also has maintainance window
upvoted 4 times



pk349 2 years, 6 months ago

A: 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.
upvoted 4 times

GCP72 2 years, 6 months ago



6AM to 10PM

upvoted 1 times

  **GCP72** 2 years, 6 months ago

Selected Answer: A

A is the correct answer because application is used between
upvoted 3 times

  **jitu028** 2 years, 6 months ago

Selected Answer: D


Correct answer - D
<https://cloud.google.com/sql/docs/mysql/high-availability#HA-configuration>
upvoted 2 times

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.

Suggested Answer: C

Community vote distribution



🗳️ 👤 **hanayome** 7 months, 3 weeks ago

Selected Answer: A

obviously A

upvoted 1 times

🗳️ 👤 **Tempingtron** 9 months, 2 weeks ago

Selected Answer: A

Other options don't directly resolve the issue. B is the worst answer since it disrupts the whole read setup.

upvoted 2 times

🗳️ 👤 **whoosh** 1 year ago

Selected Answer: A

Definitely A.

upvoted 3 times

🗳️ 👤 **goodsport** 1 year, 3 months ago

Selected Answer: A

Definitely A.

upvoted 3 times

🗳️ 👤 **felipeschossler** 1 year, 8 months ago

Selected Answer: A

A. Optimize query for resolve replication lag. Docs: https://cloud.google.com/sql/docs/mysql/replication/replication-lag#optimize_queries_and_schema

upvoted 4 times

🗳️ 👤 **BenMS** 1 year, 9 months ago

Selected Answer: B

I would have thought that recreating your replicas should be a standard action as part of a major client software release - especially one that potentially makes structural changes to the DB, as implied by the description here.

Option B seems to me like the most effective solution in this scenario, as well as the simplest.

upvoted 1 times

🗳️ 👤 **dynamic_dba** 1 year, 9 months ago

A.



High replication lag is caused when the write load is too high for the replica to handle. Other causes include slow running queries on the replica, tables not having PKs thus forcing FTS, queries like DELETE...WHERE. Possible solutions are configure parallel replications, increase the size of the replica, send read traffic to the read replica, index the tables, identify and fix slow write queries, recreate the replica. To increase the throughput of replication increase the flag `slave_parallel_workers`. B is possible but should not be the first option. C and D add resource but don't fix the issue. As others have said, the issue could be network related and additional traffic is mentioned in the question. A is still the best answer.

upvoted 4 times

🗳️ 👤 **Nirca** 1 year, 10 months ago

Selected Answer: A

A. But in reality, none. You need to analyze the root cause. Network connection latency or bandwidth might be relevant too.
upvoted 2 times

  **felipeschossler** 1 year, 8 months ago

It's true, none options seems to be right here because you need to analyze everything first.
upvoted 1 times

  **ssaporylo** 1 year, 12 months ago

Vote for A
upvoted 1 times



  **chelbsik** 2 years ago

Selected Answer: A

Vote for A
upvoted 2 times

  **pk349** 2 years ago

A: Identify and optimize slow running queries, or set parallel ***** replication flags.
upvoted 2 times

  **GCP72** 2 years ago

Selected Answer: A

A & C is correct but A is the best answer
upvoted 2 times

  **jitu028** 2 years ago

Selected Answer: A

correct answer - A
<https://cloud.google.com/sql/docs/mysql/replication/manage-replicas#:~:text=Replication%20lag%20is%20consistently,Find%20and%20fix%20them>.
upvoted 3 times

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

Suggested Answer: B

Community vote distribution

A (100%)

 **dynamic_dba** Highly Voted 1 year, 9 months ago


A.

C is wrong because databaseUser (Person B) would allow database writes and the question says generate metric reports, which would be read access only. databaseReader (Application C) doesn't allow backups.

D is wrong because databaseUser (Person B) would allow database writes. That leaves A and B. Based upon Google's own documentation, it must be A. B would work, but backupAdmin for Application C would allow backup deletion as well as creation. backupWriter is described in the docs as "is intended to be used by scripts that automate backup creation".

<https://cloud.google.com/spanner/docs/iam>

upvoted 9 times

 **Tempingtron** Most Recent 9 months, 2 weeks ago

Selected Answer: A

We need an Admin for A, A reader for B and a Writer for C. Therefore A is the correct answer.

upvoted 2 times

 **theseawillclaim** 1 year, 3 months ago

A is the one.

You don't need the backupAdmin.

upvoted 2 times

 **goodsport** 1 year, 3 months ago

Selected Answer: A

Answer is A.

upvoted 2 times


 **cloudkoala** 1 year, 8 months ago

Selected Answer: A

It should be A as per the documentation.



<https://cloud.google.com/spanner/docs/iam#spanner.backupWriter>

upvoted 2 times

 **Nirca** 1 year, 10 months ago

Selected Answer: A

A is the best answer
upvoted 2 times



  **pk349** 2 years ago

A: roles/spanner.databaseAdmin for Person A
roles/spanner.databaseReader for Person B
roles/spanner.backupWriter for Application C
upvoted 3 times

  **chelbsik** 2 years ago


Selected Answer: A

B and C are obviously wrong because application only needs backupWriter permissions.
D is wrong because roles/spanner.databaseUser contains write permissions, and we don't need that.
upvoted 4 times

  **GCP72** 2 years ago

Selected Answer: A

A is the correct answer.
Cloud Spanner Backup Writer
This role is intended to be used by scripts that automate backup creation. A principal with this role can create backups, but cannot update or delete them. Lowest-level resource
upvoted 3 times

  **jitu028** 2 years ago

Selected Answer: A


Correct answer - A
upvoted 2 times

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.

Suggested Answer: C

Community vote distribution



dynamic_dba Highly Voted 1 year, 3 months ago

D.

A is nonsense. Using interleaved tables can help speed up queries, but the question says query response times are OK. So B is wrong. C is wrong for the same reason. That leaves D. The question is about which factors determine the cost of running Spanner. They include region vs. multi-region, compute unit (nodes or processing units), how much storage and how much backup space. From the Google docs, it says "When you create a Cloud Spanner instance, you choose the number of compute capacity nodes or processing units to serve your data. However, if the workload of an instance changes, Cloud Spanner doesn't automatically adjust the size of the instance. This document introduces the Autoscaler tool for Cloud Spanner (Autoscaler), an open source tool that you can use as a companion tool to Cloud Spanner. This tool lets you automatically increase or reduce the number of nodes or processing units in one or more Spanner instances based on how their capacity is being used."

<https://cloud.google.com/spanner/docs/autoscaling-overview>

upvoted 8 times

ewelaz Most Recent 9 months ago

Selected Answer: B

It's D

upvoted 1 times

goodsport 9 months, 2 weeks ago

Selected Answer: D

Autoscaling is the way to go here. D.

upvoted 3 times

CloudKida 1 year ago

Selected Answer: D

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>

upvoted 3 times

BenMS 1 year, 3 months ago

Selected Answer: D

As the others say - use autoscaling to rightsize the cluster

upvoted 3 times

pk349 1 year, 6 months ago

D: Use granular instance sizing in Cloud Spanner and Autoscaler.



upvoted 4 times

GCP72 1 year, 6 months ago

Selected Answer: D

D is the correct answer, <https://cloud.google.com/architecture/autoscaling-cloud-spanner>

upvoted 4 times

  **jitu028** 1 year, 6 months ago

Selected Answer: D

Correct answer - D

upvoted 2 times



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

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

Suggested Answer: D

Community vote distribution

D (100%)

  **dynamic_dba** Highly Voted 1 year, 9 months ago

D.

The question HA for US and APAC. Any answer which mentions Europe must be wrong. That eliminates A and C. HA requires > 1 cluster, so B must be wrong, leaving D. D shows 2 clusters in US and 2 in APAC.

upvoted 12 times

  **hanayome** Most Recent 7 months, 3 weeks ago

Selected Answer: D

obviously D

upvoted 2 times

  **goodsport** 1 year, 3 months ago

Selected Answer: D

D seems right.

upvoted 2 times

  **jteru** 1 year, 5 months ago

Selected Answer: D

D is correct.

upvoted 2 times

  **pk349** 2 years ago

D: Maintain a target of 35% CPU utilization by locating:

cluster-a in zone us-central1-a

cluster-b in zone us-central2-a

cluster-c in zone asia-northeast1-b

cluster-d in zone asia-east1-b

upvoted 2 times

  **chelbsik** 2 years ago

Selected Answer: D

Forgot to vote

upvoted 4 times

  **chelbsik** 2 years ago

Looks like D to me - it's the only one with 2 US and 2 Asia regions. Also <https://cloud.google.com/bigtable/docs/replication-settings#regional-failover>

upvoted 4 times

  **GCP72** 2 years ago

D is the correct answer

upvoted 2 times

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

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

Suggested Answer: B

Community vote distribution

A (79%)

D (21%)

🗳️ 👤 **dynamic_dba** Highly Voted 🍌 1 year, 9 months ago

A.

Cloud SQL could not handle the load, so B is wrong. Memorystore can scale up to 300 GB. The question mentions needing 8 TB, so C must be wrong. BigQuery could not handle the latency requirements of the question, which leaves A. Bigtable could handle the volume of writes at the speeds required.

upvoted 7 times

🗳️ 👤 **pk349** Highly Voted 🍌 2 years ago

A: Write your data into Bigtable ***** and use Dataproc and the Apache Hbase libraries for analysis.

upvoted 6 times

🗳️ 👤 **Jason_Cloud_at** Most Recent 🕒 10 months, 1 week ago

Selected Answer: A

Bigtable is ideal for clickstream and IOT use cases, also it can process high performance read and writes globally.

upvoted 4 times

🗳️ 👤 **ToniTovar** 10 months, 2 weeks ago

Selected Answer: D

This option uses BigQuery, that has a low latency and is a big data

upvoted 2 times

🗳️ 👤 **VG1900** 1 year, 1 month ago

Selected Answer: D

A is not correct because Bigtable is not designed for real-time analytics. It is a good choice for storing and retrieving small amounts of data quickly, but it is not as efficient for analyzing large volumes of data.

B is not correct because it cannot support Million of Read and Write

C is not correct because of storage limitation

D is correct

upvoted 4 times

🗳️ 👤 **goodsport** 1 year, 3 months ago

Selected Answer: A

I would opt for A.

upvoted 2 times

🗳️ 👤 **ArtistS** 1 year ago

Why opt A? It is not real-time, and the question mentions that they want to analysis, why not use bigquery, only 8 TB

upvoted 1 times

🗳️ 👤 **learnazureportal** 1 year, 3 months ago

The correct answer is D. Stream your data into BigQuery and use Dataproc and the BigQuery Storage API to analyze large volumes of data.. A is used for NOSQL

upvoted 2 times

🗨️ 👤 **CloudKida** 1 year, 6 months ago

Selected Answer: A

At a high level, Bigtable is a NoSQL wide-column database. It's optimized for low latency, large numbers of reads and writes, and maintaining performance at scale. Bigtable use cases are of a certain scale or throughput with strict latency requirements, such as IoT, AdTech, FinTech, and so on. If high throughput and low latency at scale are not priorities for you, then another NoSQL database like Firestore might be a better fit.

upvoted 2 times

🗨️ 👤 **Hilab** 1 year, 9 months ago

B. Normalize the data model.

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

When designing a schema for Cloud Spanner, it is important to follow best practices to avoid hotspots and ensure optimal performance. Hotspots occur when too many requests are targeted at a single node or group of nodes, causing them to become overloaded and potentially impacting performance.

upvoted 1 times

🗨️ 👤 **Hilab** 1 year, 9 months ago

Normalization is a recommended best practice in database schema design, including in Cloud Spanner. It involves breaking down large tables into smaller, more manageable tables that are linked together by relationships. This can help reduce duplication of data and improve performance by reducing the amount of data that needs to be read or written to the database.

Promoting high-cardinality attributes in multi-attribute primary keys is also recommended in Cloud Spanner schema design. High-cardinality attributes are those that have a large number of distinct values, such as product IDs or customer IDs. Including these attributes in the primary key can help distribute data more evenly across nodes, reducing the likelihood of hotspots.

Using an auto-incrementing value as the primary key or a bit-reverse sequential value as the primary key can result in hotspots, particularly if new data is being added at a high rate. These approaches can cause all new data to be inserted into a single node, leading to performance issues.

upvoted 1 times

🗨️ 👤 **Hilab** 1 year, 9 months ago

The above answer is for Question #15, my mistake I put the comments here

upvoted 1 times

🗨️ 👤 **Hilab** 1 year, 9 months ago

D. Stream your data into BigQuery and use Dataproc and the BigQuery Storage API to analyze large volumes of data.

BigQuery is a fully managed, serverless data warehouse that allows you to store and analyze large datasets using SQL-like queries. It is designed to handle petabyte-scale data and is optimized for fast query performance. By streaming your clickstream data into BigQuery, you can store and process large amounts of data in real-time.

Dataproc, on the other hand, is a fully-managed cloud service for running Apache Hadoop and Spark clusters. It provides a managed, easy-to-use environment for data processing, which can be used to analyze the data stored in BigQuery.

The BigQuery Storage API allows you to directly access data stored in BigQuery from external applications, including Dataproc, which enables you to run advanced analytics on large volumes of data with low latency.

This approach provides a scalable, low-latency solution for storing and analyzing large volumes of data, making it a good fit for your requirements.

upvoted 1 times

🗨️ 👤 **H_S** 1 year, 9 months ago

Selected Answer: A

A. Write your data into Bigtable and use Dataproc and the Apache Hbase libraries for analysis. Most Voted

upvoted 2 times

🗨️ 👤 **Nirca** 1 year, 9 months ago

Selected Answer: A

A looks like best option.

upvoted 2 times

🗨️ 👤 **GCP72** 2 years ago

Selected Answer: A

A is correct answer, C wouldn't be handled 8TB data

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

upvoted 3 times

  **Kloudgeek** 2 years ago

Answer is A. Click stream and time series data and the size is 8TB. Read low latency with reads and writes. Correct answer is A to use BigTable for storage and use either CBT or Hbase API to interact with data.

upvoted 5 times

  **fredcaram** 2 years ago

Selected Answer: A

B couldn't handle this volume of writes and read, D wouldn't be able to handle the writing and C wouldn't be suited for this.

upvoted 3 times

  **juancambb** 2 years ago

Selected Answer: A

must be A


upvoted 2 times

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 hotspots 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.

Suggested Answer: AD

Community vote distribution



PrtkKA Highly Voted 1 year, 10 months ago

Selected Answer: DE

Spanner needs high cardinality primary key to avoid hotspotting.
upvoted 7 times

PKookNN Most Recent 11 months, 2 weeks ago

Selected Answer: DE

I would go with D and E
upvoted 3 times

julioevk 1 year, 2 months ago

Selected Answer: DE

<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

upvoted 3 times

nescafe7 1 year, 5 months ago

Selected Answer: BD

I agree with Hilab's comment below.
upvoted 2 times

felipeschossler 1 year, 8 months ago

Selected Answer: DE

D and E, the docs are below.

D: <https://cloud.google.com/bigtable/docs/schema-design#row-keys-avoid>

E: https://cloud.google.com/spanner/docs/schema-design#bit_reverse_primary_key

upvoted 4 times

Carpediem78 1 year, 8 months ago

high-cardinality
upvoted 1 times

PATILDXB 1 year, 8 months ago

Correct answers are D,E.

Refer to the link which is self explanatory.

<https://cloud.google.com/spanner/docs/schema-design>

upvoted 3 times

BenMS 1 year, 9 months ago

Selected Answer: DE

A - incrementing values are an explicitly documented antipattern

B - normalising the schema does not specifically address hotspotting

C - low cardinality values in the primary key will also cause hotspotting

D - promoting high cardinality values in the primary key (i.e. moving them nearer the front of the value) is a recommended approach to reduce hotspotting

E - bit-reversed keys are an explicitly recommended best practice

upvoted 3 times



  **dynamic_dba** 1 year, 9 months ago

D, E.

A is wrong because that will promote hotspots. C is wrong because low cardinality attributes being part of the key (particularly at the front multi-attribute keys) will also promote hotspots. That makes D correct by definition. This leaves B or D as the other correct answer. The fact the new data has already been added to the database suggests the data model is already properly normalized. In addition, one of the techniques to reduce or eliminate hotspots is to bit reverse sequential values. It's in Google's docs here:

<https://cloud.google.com/spanner/docs/schema-design>

upvoted 3 times



  **Hilab** 1 year, 9 months ago

B. Normalize the data model.

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

When designing a schema for Cloud Spanner, it is important to follow best practices to avoid hotspots and ensure optimal performance. Hotspots occur when too many requests are targeted at a single node or group of nodes, causing them to become overloaded and potentially impacting performance.

upvoted 2 times

  **Hilab** 1 year, 9 months ago

Normalization is a recommended best practice in database schema design, including in Cloud Spanner. It involves breaking down large tables into smaller, more manageable tables that are linked together by relationships. This can help reduce duplication of data and improve performance by reducing the amount of data that needs to be read or written to the database.

Promoting high-cardinality attributes in multi-attribute primary keys is also recommended in Cloud Spanner schema design. High-cardinality attributes are those that have a large number of distinct values, such as product IDs or customer IDs. Including these attributes in the primary key can help distribute data more evenly across nodes, reducing the likelihood of hotspots.

Using an auto-incrementing value as the primary key or a bit-reverse sequential value as the primary key can result in hotspots, particularly if new data is being added at a high rate. These approaches can cause all new data to be inserted into a single node, leading to performance issues.

upvoted 2 times

  **Nirca** 1 year, 9 months ago

Selected Answer: DE

"hotspots" in a database means that many IOPS (usually writes/updates) are happening on the same data-block; usually due to calling the same DATA.

low cardinality => same value in the column ==> hotspots.

High cardinality => different values in the column ==> avoiding hotspotting.

upvoted 4 times

  **zanhsieh** 1 year, 10 months ago

Selected Answer: CE

CE

A and D: WRONG. Anti-pattern

Since the question specifically stated the hotspots cause by new SKUs and product catalog data added, so the goal would be:

1. The old data keeps distributed without any extra work needed.

2. Resolving the new data hot spots problem.

It seems to me that SKU and product catalog are already normalized, so further normalize might touch the old data. This means B is out. If the new data already normalized, then it must have some high-cardinality attributes, e.g. SKU_id, and some low-cardinality attributes, e.g. category_id. So I picked low-cardinality attributes in multi-attribute primary keys as C. I agreed with E as already Google recommended practice.

Reference:

<https://cloud.google.com/spanner/docs/schema-design>

upvoted 3 times

  **TFMV** 1 year, 12 months ago

CE. Normalizing the data is not generally recommended if interleaving can suffice.

upvoted 1 times

🗨️ 👤 **pk349** 2 years ago

B: Normalize the data model.

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

upvoted 2 times

🗨️ 👤 **gabrielosluz** 1 year, 10 months ago

Wrong B

upvoted 1 times

🗨️ 👤 **GCP72** 2 years ago

Selected Answer: CE

Looks CE is correct for me

upvoted 2 times

🗨️ 👤 **jitu028** 2 years ago

Selected Answer: CE

Correct answer - CE

upvoted 1 times

🗨️ 👤 **Kloudgeek** 2 years ago

Answer is B & E for schema design. <https://cloud.google.com/spanner/docs/schema-design> . B&E are correct answers


upvoted 1 times

You are managing multiple applications connecting to a database on Cloud SQL for PostgreSQL. You need to be able to monitor database performance to easily identify applications with long-running and resource-intensive queries. What should you do?

- A. Use log messages produced by Cloud SQL.
- B. Use Query Insights for Cloud SQL.
- C. Use the Cloud Monitoring dashboard with available metrics from Cloud SQL.
- D. Use Cloud SQL instance monitoring in the Google Cloud Console.

Suggested Answer: C

Community vote distribution



dynamic_dba Highly Voted 1 year, 9 months ago

B.

Query Insights helps identify performance and load issues at the database application layer. None of the other options do that. So the answer is B.
upvoted 8 times

chelbsik Highly Voted 2 years ago

Selected Answer: B

Vote for B

upvoted 5 times

hanayome Most Recent 7 months, 3 weeks ago

Selected Answer: B

B. Query Insights is built in feature to identify long running query

upvoted 2 times

H_S 1 year, 9 months ago

Selected Answer: B

B. Query insights: for identifying long running SQLs <https://cloud.google.com/sql/docs/mysql/using-query-insights#introduction>

upvoted 3 times

Nirca 1 year, 9 months ago

Selected Answer: B

B. Query insights: for identifying long running SQLs <https://cloud.google.com/sql/docs/mysql/using-query-insights#introduction>

upvoted 3 times

SidsA 1 year, 10 months ago

Answer should be B: As long-running query and resource intensive queries are available in query insight providing details about which queries are taking how much time and resource utilization at what stage.

upvoted 4 times

Blueocean 2 years ago

Option B

<https://cloud.google.com/sql/docs/postgres/using-query-insights>

upvoted 3 times

pk349 2 years ago

B: Use Query Insights for Cloud SQL.

ging multiple applications connecting

upvoted 1 times

GCP72 2 years ago

Selected Answer: B

B is the correct answer, agree with jitu028

upvoted 1 times

🗨️ 👤 **jitu028** 2 years ago

Correct answer - B

<https://www.youtube.com/watch?v=qN7x3ngwz1o>

upvoted 2 times

🗨️ 👤 **range9005** 2 years ago

Selected Answer: D

For Database performance, Cloud SQL System insights dashboard is preferable

upvoted 1 times

🗨️ 👤 **Kloudgeek** 2 years ago

C is correct answer. We can use Cloud SQL Instance monitoring as well, but need to build the custom metrics for the metrics needed. Instead with Cloud Monitoring these metrics are already available. <https://cloud.google.com/sql/docs/mysql/monitor-instance#cloud-monitoring>

<https://cloud.google.com/sql/docs/sqlserver/admin-api/metrics>

upvoted 1 times

🗨️ 👤 **range9005** 2 years ago

Selected Answer: D

The Cloud SQL System insights dashboard helps you detect and analyze system performance problems.

.

<https://cloud.google.com/sql/docs/postgres/monitor-instance#sql-system-insights>


upvoted 1 times

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

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

Suggested Answer: A

Community vote distribution



BenMS Highly Voted 1 year, 9 months ago

Selected Answer: C

Dynamic schema indicates this is a NoSQL solution (ruling out Cloud SQL) and the application use case specifically suits Firestore (the question even refers to storing data in documents) as opposed to BigTable.

Firestore in Native supports realtime client updates, which is needed for the analytics requirement: https://cloud.google.com/firestore/docs/firestore-or-datastore#feature_comparison

upvoted 6 times

dynamic_dba Highly Voted 1 year, 9 months ago

C.

A dynamic schema means the database backend cannot be relational. That eliminates B. No criteria is mentioned that would justify Bigtable (low latency or massive data volume), so eliminate A. That leaves Firestore options which make sense since it's a NoSQL database. Since "website" and "mobile" are both mentioned in the question, Firestore in Native mode must be the correct answer.

upvoted 5 times

abhinav45852 Most Recent 11 months, 2 weeks ago

Selected Answer: A

Firestore in native mode and Bigtable both seems to be correct answer but weightage can be given to Bigtable on the basis of replication to BigQuery via change streams is easier.

upvoted 2 times

Shadab 1 year, 8 months ago

Selected Answer: C

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.

upvoted 3 times

pk349 2 years ago

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

upvoted 2 times

  **GCP72** 2 years ago

Selected Answer: C

C is the correct Answer

upvoted 2 times

  **range9005** 2 years ago

Selected Answer: C

Firestore introduces new features such as:

A new, strongly consistent storage layer

A collection and document data model

Real-time updates

Mobile and Web client libraries

.

https://cloud.google.com/datastore/docs/firestore-or-datastore#in_native_mode

upvoted 3 times

  **fredcaram** 2 years ago

Selected Answer: C

Seems like a better use for firestore and since it needs to reflect changes downstream in real time the native one would be better.

upvoted 2 times

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

- A. Create secondary indexes on the replica.
- B. Create additional read replicas, and partition your analytics users to use different read replicas.
- C. Disable replication on the read replica, and set the flag for parallel replication on the read replica. Re-enable replication and optimize performance by setting flags on the primary instance.
- D. Disable replication on the primary instance, and set the flag for parallel replication on the primary instance. Re-enable replication and optimize performance by setting flags on the read replica.
- E. Move your analytics workloads to BigQuery, and set up a streaming pipeline to move data and update BigQuery.

Suggested Answer: BE

Community vote distribution



dynamic_dba Highly Voted 2 years, 3 months ago

B, C.

You have 2 problems. 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>

upvoted 7 times

cardareel 1 year, 9 months ago

B isn't correct ==> "without making changes to the current reports". If you choose B, reports will need changes to point to the new instances.

upvoted 1 times

hmarine 3 months ago

B is not a problem, as the report content itself is unchanged.

upvoted 1 times

ssaporylo Highly Voted 2 years, 5 months ago

Vote for AC

A <https://cloud.google.com/sql/docs/mysql/replication/read-replica-indexes> increase performance on read operation

C <https://cloud.google.com/sql/docs/mysql/replication/manage-replicas#basic-steps-to-change-parallel-replication-flags>

upvoted 7 times

Lenifia Most Recent 11 months, 3 weeks ago

Selected Answer: AC

The two main issues at hand are replication lag and slow report performance. While option B, creating additional read replicas, could address the load issue, it would require changes to the current reports, which is not desirable. Instead, option A, implementing secondary indexes in Cloud SQL, could enhance report speed without altering the reports. As for the replication lag, option C suggests enabling parallel replication in Cloud SQL, which is a plausible solution. 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. Therefore, considering the constraints and the context, options A and C are the most suitable choices.

upvoted 1 times

studymoreoften 1 year ago

Selected Answer: BC

B - Addresses performance improvements: Reduce the burden on the primary instance by offloading replication work to multiple read replicas.

<https://cloud.google.com/sql/docs/mysql/replication#cascading-replicas>

C- Addresses lag time because Parallel replication reduces replication lag by increasing the number of SQL threads that work to execute these transactions.

<https://cloud.google.com/sql/docs/mysql/replication/manage-replicas#configuring-parallel-replication>

upvoted 1 times

🗳️ 👤 **PKookNN** 1 year, 4 months ago

Selected Answer: C

Just got this question and there is no A, and it is not 'choose two' but one answer only.

upvoted 3 times

🗳️ 👤 **cardareel** 1 year, 9 months ago

There's no discussion about C.

A & B both sounds reasonable. Why I would choose A instead of B? Due to keywords "without making changes to the current reports" and "MySQL".

Option B would require to point to new IP addresses (the new read replicas) and split which group of users which run X reports and which group of user which run Y reports connect to which read replica. Option A (secondary indexes) is only available for Cloud SQL (the question's use case is about MySQL) and explicitly mentions "for reporting purposes".

upvoted 4 times

🗳️ 👤 **KennyHuang** 2 years, 1 month ago

Selected Answer: BC

B. By creating additional read replicas, you can distribute the load of analytics workloads across multiple instances. Partitioning your analytics users to use different read replicas allows you to further distribute the workload and improve performance. This helps to alleviate the excessive load on the primary database and enhances the reporting experience for users.

C. Disabling replication on the read replica can help reduce the data replication lag. By setting the flag for parallel replication on the read replica, you allow parallel execution of replication threads, which can expedite data replication. Additionally, optimizing performance by setting flags on the primary instance can help improve the overall performance of the replication process and reduce the lag experienced by the read replica.

upvoted 3 times

🗳️ 👤 **cardareel** 1 year, 9 months ago

B isn't correct ==> "without making changes to the current reports". If you choose B, reports will need changes to point to the new instances.

upvoted 1 times

🗳️ 👤 **Hilab** 2 years, 3 months ago

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

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

Creating additional read replicas can distribute the analytics workload and reduce the lag in data replication. By partitioning your analytics users to use different read replicas, you can further reduce the load on each replica and improve performance.

upvoted 1 times

🗳️ 👤 **Hilab** 2 years, 3 months ago

Disabling replication on the primary instance and setting the flag for parallel replication can improve the replication speed and reduce the lag in data replication. Once you have optimized performance on the primary instance, you can re-enable replication and optimize performance on the read replica.

Creating secondary indexes on the replica may improve query performance but will not reduce the lag in data replication. Moving your analytics workloads to BigQuery and setting up a streaming pipeline to move data can provide near-real-time data but will require significant changes to your current reports.

upvoted 1 times

🗳️ 👤 **H_S** 2 years, 3 months ago

Selected Answer: BC

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

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

upvoted 1 times

🗳️ 👤 **cardareel** 1 year, 9 months ago

B isn't correct ==> "without making changes to the current reports". If you choose B, reports will need changes to point to the new instances.

upvoted 1 times

🗳️ 👤 **Nirca** 2 years, 3 months ago

Selected Answer: BC

- A. Create secondary indexes on the replica. - No indication that the reports will benefit from indexes.
- B. Create additional read replicas, and partition your analytics users to use different read replicas. --> might rebalance the load.
- C. Disable replication on the read replica, and set the flag for parallel replication on the read replica. Re-enable replication and optimize performance by setting flags on the primary instance. --> might add parallelism to the replication lag.
- D. Disable replication on the primary instance, and set the flag for parallel replication on the primary instance. Re-enable replication and optimize performance by setting flags on the read replica. --> na
- E. Move your analytics workloads to BigQuery, and set up a streaming pipeline to move data and update BigQuery.--> according to question statement , no SQL rewrite is possible.

upvoted 2 times

🗳️ 👤 **Swapnil54** 2 years, 4 months ago

A & C looks fine.

upvoted 2 times

🗳️ 👤 **muky31dec** 2 years, 5 months ago

Selected Answer: AC

Ans is AC

upvoted 3 times

🗳️ 👤 **muky31dec** 2 years, 5 months ago

Creating secondary indexes on the replica can help improve the performance of the reports by allowing the read replica to quickly locate the data it needs without having to scan the entire table. This can help speed up the queries

upvoted 1 times

🗳️ 👤 **muky31dec** 2 years, 5 months ago

Ans is AC

upvoted 1 times

🗳️ 👤 **csrazdan** 2 years, 5 months ago

Selected Answer: BC

The question has 2 issues - replication lag and reports running slow.

B - will address reports running slow since fewer users will be on the replica server

C - will address replication lag.

upvoted 4 times

🗳️ 👤 **muky31dec** 2 years, 5 months ago

AC must correct choice in the situation.

upvoted 1 times

🗳️ 👤 **TFMV** 2 years, 5 months ago

Moving workload to BQ is not an option. That, at a minimum, would require connection changes in the reports and the question specifically states that report changes are unacceptable. Aside from that, we do not know if the reports are being generated by a tool and whether that tool supports BQ.

upvoted 1 times

🗳️ 👤 **sp57** 2 years, 6 months ago

AC - Not understanding votes for E, can't be done without some changes to reports.

upvoted 2 times

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

What should you do?

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

Suggested Answer: C

Community vote distribution

D (100%)

🗳️ **julioevk** 9 months ago

D because
CMEK where the encryption keys are stored
IAM govern the access to data
VPC Service Controls configure where data is stored control
upvoted 4 times

🗳️ **theseawillclaim** 9 months, 1 week ago

Selected Answer: D

D.
C might seem ok, but you'd need some kind of tracking to localize users, and there is no mention of it.
upvoted 1 times

🗳️ **standm** 1 year, 1 month ago

should be CSEK and not CMEK. Then 'D'.
upvoted 1 times

🗳️ **dynamic_dba** 1 year, 3 months ago

D.
Using IAM policies, VPC Service Controls and CMEK is the best answer. A doesn't make sense since Geography would be a factor at the Region level, not zone level. B is a lot of work and GCP is all about making things easier. C address part of the issue, but D addresses more. The link provided by sp57 is spot on.
upvoted 1 times

🗳️ **ralf_cc** 1 year, 5 months ago

C - it is about location of the data
upvoted 1 times

🗳️ **ssaporylo** 1 year, 5 months ago

My vote D
upvoted 2 times

🗳️ **sp57** 1 year, 6 months ago

D, <https://cloud.google.com/blog/products/identity-security/meet-data-residency-requirements-with-google-cloud>
upvoted 4 times

🗳️ **pk349** 1 year, 6 months ago



D. Use features like customer-managed encryption keys (CMEK), VPC Service Controls, and Identity and Access Management (IAM) policies.
upvoted 2 times

🗳️ **GCP72** 1 year, 6 months ago

Selected Answer: D

D is the correct answer

upvoted 1 times

  **range9005** 1 year, 6 months ago

Selected Answer: D

data residency requirements can be achiy with CMEK, VPC and IAM

upvoted 2 times

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.

Suggested Answer: B

Community vote distribution

D (75%)

B (25%)

🗳️ 👤 **dynamic_dba** Highly Voted 1 year, 9 months ago

D.

The question says backups and maintenance are an issue, so moving to a managed service (Cloud SQL) would be the right thing to do. That eliminates C and A. Option B could (depending upon the DB size) require a lot of downtime to export, copy the dump file to Cloud Storage, then import into Cloud SQL. Therefore, the least amount of downtime would be D.

<https://cloud.google.com/sql/docs/mysql/replication/configure-replication-from-external>

upvoted 8 times

🗳️ 👤 **RaphaelG** Most Recent 11 months, 2 weeks ago

The only issue that I have with the D answer is that an external replica is an explicit Cloud SQL concept not a MySQL concept; external replica is when you have a primary instance already in Cloud SQL, however, if you have your MySQL on premises then you are dealing with replication from an external server. I have read documentation for CloudSQL for MySQL before and I am more convinced with B to be frank

upvoted 4 times

🗳️ 👤 **kkjw** 1 year, 3 months ago

External replica is a setup with primary on cloud and the replica external to the cloud .eg.e on prem.. which is the reverse of what the question is looking for

upvoted 2 times

🗳️ 👤 **ArtistS** 1 year ago

Yep, D sounds wired, the question said it should import from the on-premise. So my opinion is B

upvoted 1 times

🗳️ 👤 **KennyHuang** 1 year, 7 months ago

Selected Answer: D

This approach provides a seamless migration process with minimal impact on the application's availability.

upvoted 2 times

🗳️ 👤 **BenMS** 1 year, 9 months ago

Selected Answer: D

Another good explanation from dynamic_dba.

upvoted 2 times

🗳️ 👤 **EueChan** 1 year, 9 months ago

Selected Answer: D

I agree with D

upvoted 2 times

🗳️ 👤 **H_S** 1 year, 9 months ago

Selected Answer: D

D is the better option

upvoted 3 times

🗳️ 👤 **Nirca** 1 year, 9 months ago

Selected Answer: D

minimal downtime!!! And the customer has read replicas + backups are expensive.

D is the better option

upvoted 4 times

🗨️ 👤 **PrtkKA** 1 year, 10 months ago

Selected Answer: D

D : External read replica will help achive minimal down time.

upvoted 2 times

🗨️ 👤 **SidsA** 1 year, 10 months ago

Correct answer is D: Create cloudsql replica of on-prem server and promote with almost-no downtime by pointing app to cloudsql. mysqldump is heavy and time-consuming operation (even though if you run it on on-prem read-replica, it will be identical to migrating to cloudsql using managed database migration service.

upvoted 4 times

🗨️ 👤 **SandyZA** 1 year, 11 months ago

D create a read replica on cloud then promote it

upvoted 3 times

🗨️ 👤 **GCP72** 2 years ago

Selected Answer: B

B is correct answer

upvoted 2 times

🗨️ 👤 **range9005** 2 years ago

Selected Answer: B

<https://cloud.google.com/database-migration/docs/mysql/mysql-dump>


upvoted 3 times

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.

Suggested Answer: C

Community vote distribution



KennyHuang Highly Voted 2 years, 1 month ago

Selected Answer: D

By combining Google Cloud IoT Core for device management and data ingestion, Google Cloud Bigtable for storing and processing IoT data with low latency, and Looker for advanced analytics and visualization, you can build a highly scalable, multi-cloud compatible, and low-latency solution to address your IoT device maintenance requirements effectively.

upvoted 10 times

dynamic_dba Highly Voted 2 years, 3 months ago

D.

The question says a managed solution, so that eliminates C. Firestore and Spanner do not have the scalability or low latency required. This leaves D. Bigtable by itself is a GCP thing, but Looker allows data visualization across multiple cloud environments.

<https://www.looker.com/google-cloud/>

upvoted 6 times

krop Most Recent 2 months ago

Selected Answer: C

Multi-Cloud Strategy Support

MongoDB Atlas is a managed DBaaS available across AWS, Azure, and Google Cloud, enabling hybrid/multi-cloud deployments. This aligns with the need for geographic redundancy and compliance across regions.

upvoted 2 times

bad5fad 3 months, 2 weeks ago

Selected Answer: C

a multicloud solution is needed and MongoDB Atlas fill all the checks. C is correct

upvoted 1 times

887ad17 5 months ago

Selected Answer: D

Scalability: Bigtable is a fully managed, highly scalable NoSQL database that can handle large volumes of data with low latency, making it well-suited for real-time data ingestion from thousands of IoT devices.

Low Latency: Bigtable is designed for high-throughput and low-latency read and write operations, which is essential for real-time data processing and insights.

Multi-Cloud Strategy: While Bigtable itself is a Google Cloud service, it can integrate with other cloud services and tools, and data can be exported or accessed from other cloud environments if needed.

Integration with Looker: Looker is a powerful business intelligence tool that can connect to Bigtable to provide insights and visualizations. This integration allows you to analyze the data and design inspection routines, repair schedules, and replacement schedules based on the insights gathered.

upvoted 1 times

rglearn 8 months, 1 week ago

Selected Answer: D

Managed solution hence Bigtable is right choice

upvoted 2 times

🗨️ 👤 **dija123** 1 year, 1 month ago

Selected Answer: D

Agree with D as You need a managed solution

upvoted 2 times

🗨️ 👤 **Pime13** 1 year, 2 months ago

Selected Answer: C

Bigtable can't be integrated with looker

upvoted 1 times

🗨️ 👤 **Haraprasad** 1 year, 2 months ago

C . As its says to avoid refactoring.

upvoted 1 times

🗨️ 👤 **okkokkoo** 1 year, 2 months ago

Selected Answer: D

Google will only pitch its own products

upvoted 1 times

🗨️ 👤 **james2033** 1 year, 3 months ago

Selected Answer: C

multi-cloud

upvoted 1 times

🗨️ 👤 **ToniTovar** 1 year, 4 months ago

Selected Answer: D

I think Looker is multi-cloud as some of the colleagues said in this post. So I choose D.

<https://services.google.com/fh/files/misc/042420-ppm-multi-cloud-one-sheet-8-5x11-en-web-gc.pdf>

upvoted 1 times

🗨️ 👤 **ToniTovar** 1 year, 4 months ago

I think Looker is multi-cloud as some of the colleagues said in this post. So I choose D.

<https://services.google.com/fh/files/misc/042420-ppm-multi-cloud-one-sheet-8-5x11-en-web-gc.pdf>

upvoted 2 times

🗨️ 👤 **PKookNN** 1 year, 5 months ago

Selected Answer: C

As others said, MongoDB is the only one that support multi-cloud (otherwise D is also a candidate).

upvoted 1 times

🗨️ 👤 **gcp_k** 1 year, 2 months ago

BigTable is nothing but hbase.

upvoted 1 times

🗨️ 👤 **Ravi_Mangalpally** 1 year, 7 months ago

The answer is D. Use Bigtable with Looker.

Bigtable is a fully managed, petabyte-scale NoSQL wide column database service for storing large amounts of data. It is designed for low latency, high throughput, and scalability. Bigtable is a good choice for storing IoT data because it can handle the high volume of data generated by IoT devices and provide low latency for real-time analysis.

Looker is a business intelligence and data analytics platform that can be used to visualize and analyze data stored in Bigtable. Looker provides a variety of features that can be used to design inspection routines, device repair, and replacement schedules based on insights gathered from the data produced by IoT devices. Looker supports hosting on public clouds like AWS and GCP, and in multi-cloud and hybrid environments.

upvoted 4 times

🗨️ 👤 **theseawillclaim** 1 year, 9 months ago

Selected Answer: C

"Multi-cloud" is the key that rules out everything except C.

upvoted 4 times

  **Ramheadhunter** 1 year, 11 months ago

Selected Answer: C


Every thing in this question points to D except for one scenario "multi-cloud". So just becuae multi cloud is specified, I choose 'C'
upvoted 3 times

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.

Suggested Answer: A

Community vote distribution



🗳️ 👤 **Kloudgeek** Highly Voted 👍 2 years, 6 months ago

Correct answer is D. <https://cloud.google.com/sql/docs/mysql/best-practices#data-arch> - Split your large instances into smaller instances, where possible.

upvoted 13 times

🗳️ 👤 **fraloca** Highly Voted 👍 1 year, 5 months ago

Selected Answer: C

The solution D is better if we can execute a massive refactor. So, the best solution is C because if CPU is normal, the issue is the capacity of the I/O. For increase it, it's enough increase the disk storage.

Source: <https://cloud.google.com/sql/docs/sqlserver/best-practices>

upvoted 10 times

🗳️ 👤 **krop** Most Recent 🕒 2 months ago

Selected Answer: D

1. Microservices Best Practices:

- The database-per-service pattern reduces contention by isolating workloads. Each microservice uses a dedicated Cloud SQL instance, preventing cross-service resource competition.
- Smaller instances are easier to scale independently based on individual service needs.

2. Horizontal Scaling:

- Splitting the monolithic database into smaller instances aligns with Google's recommendation for horizontal scaling in microservices (search result).
- Avoids hitting connection limits, I/O bottlenecks, or lock contention that may occur in a single large instance, even if CPU utilization appears normal.

3. Minimal Refactoring:

- Requires only configuration changes (e.g., updating connection strings) rather than rewriting application logic.

upvoted 1 times

🗳️ 👤 **887ad17** 5 months ago

Selected Answer: C

Storage Size Impact: Increasing the storage size for your Cloud SQL instance can improve performance because Cloud SQL automatically increases the IOPS (Input/Output Operations Per Second) as the storage size increases. This can help resolve performance bottlenecks related to disk I/O, which might be the underlying issue given that CPU utilization is normal.

No Major Refactoring: This approach does not require any significant changes to your application architecture or code, making it a straightforward solution.

Scalability: By increasing the storage size, you can ensure that your instance has the necessary IOPS to handle the increasing load as your application grows.

upvoted 1 times

🗳️ 👤 **rglearn** 8 months, 1 week ago

Selected Answer: C

can't do major refactoring. also CPU utilization is under control then mostly issue is with IOPS capacity of CloudSQL we need to increase storage space.

to increase IOPS

upvoted 1 times

🗳️ **TNT87** 8 months, 1 week ago

Selected Answer: D

d

upvoted 1 times

🗳️ **dija123** 1 year, 1 month ago

Selected Answer: C

Agree with C as the CPU Looks normal.

upvoted 1 times

🗳️ **0e75489** 1 year, 6 months ago

Splitting smaller databases require major effort.

Answer should be C

upvoted 2 times

🗳️ **julioevk** 1 year, 9 months ago

Selected Answer: D

It's a microservices architecture and CPU utilization is normal. This means that having multiple Cloud SQL instances will help for each microservice.

upvoted 2 times

🗳️ **theseawillclaim** 1 year, 9 months ago

Wouldn't D be a big refactoring, as well as switching to Spanner, especially if MySQL is considered?

This question is bad.

upvoted 2 times

🗳️ **dynamic_dba** 2 years, 3 months ago

D.

Needing to avoid any major refactoring eliminates A. The question states CPU is not an issue, so that eliminates B. Adding more storage would increase IOPS, but there's no indication network throughput is an issue, so that eliminates C. That leaves D. A microservice architecture is supposed to use a separate database for each microservice, rather than one big database for all the microservices. So D it is. The link provided by Kloudgeek is spot on.

upvoted 6 times

🗳️ **Ramheadhunter** 1 year, 11 months ago

Will splitting single instance to multiple smaller instance amount to re-factoring ?

upvoted 1 times

🗳️ **H_S** 2 years, 3 months ago

Selected Answer: D

D: Use many smaller Cloud SQL instances.

upvoted 1 times

🗳️ **H_S** 2 years, 3 months ago

D: Use many smaller Cloud SQL instances.

upvoted 2 times

🗳️ **pk349** 2 years, 6 months ago

D: Use many smaller ***** Cloud SQL instances.

upvoted 3 times

🗳️ **GCP72** 2 years, 6 months ago

Selected Answer: D

D is the correct answer

upvoted 1 times

🗳️ **range9005** 2 years, 6 months ago

Selected Answer: D

Split CloudSql instance into many small instances to support Microservices

upvoted 2 times

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

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

Suggested Answer: C

Community vote distribution

C (100%)

🗳️ 👤 **PKookNN** 11 months, 2 weeks ago

Selected Answer: C

C is simple and works

upvoted 1 times

🗳️ 👤 **dynamic_dba** 1 year, 9 months ago

C.

The only way to minimize performance impact of running an export on a Cloud SQL instance is to use a serverless export. The fact that no data synchronization is needed since it's a one off eliminates every option apart from C.

upvoted 3 times

🗳️ 👤 **Nirca** 1 year, 9 months ago

Selected Answer: C

C looks simple and ok.

upvoted 1 times

🗳️ 👤 **chelbsik** 2 years ago

Selected Answer: C

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

upvoted 1 times

🗳️ 👤 **pk349** 2 years ago

C: Create a SQL dump file in Cloud Storage using a temporary instance, and then use that file to import into a new instance.

upvoted 1 times

🗳️ 👤 **GCP72** 2 years ago

Selected Answer: C

C is the correct answer


upvoted 1 times

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.

Suggested Answer: B

Community vote distribution



🗳️ 👤 **dynamic_dba** Highly Voted 1 year, 9 months ago

A.

Eliminate D because Cloud SQL Auth Proxy by itself does not provide connection pooling. There's nothing in the question about needing to load balance HTTP traffic specifically, so ignore C. B is eliminated on the basis PgBouncer does not have multi-host configuration, failover, or detection and the question specifically says "mission critical". That leaves A which makes sense since Google Cloud Load Balancer is a regional service and the question specifically mentions a single region.

upvoted 12 times

🗳️ 👤 **bigdawg70** 10 months, 3 weeks ago

This doesn't seem correct:

<https://www.googlecloudcommunity.com/gc/Databases/Load-Balancer-for-Postgres/mp/647563#:~:text=To%20clarify%2C%20Google%20Cloud%20Platform%27s,in%20Cloud%20SQL%20for%20PostgreSQL>.

upvoted 2 times

🗳️ 👤 **Steve8512** Most Recent 2 months, 3 weeks ago

Selected Answer: A

<https://www.pgбouncer.org/faq.html> says pgбouncer cannot handle multi-host configs, so B is out. C is out because there's nothing involving HTTP, D is out because the auth proxy does not provide connection pooling.

upvoted 1 times

🗳️ 👤 **studymoreoften** 6 months, 3 weeks ago

Selected Answer: B

Cloud SQL doesn't provide load balancing between replicas. You can choose to implement load balancing for your Cloud SQL instance. You can also use con performance.

<https://cloud.google.com/sql/docs/postgres/replication#:~:text=Cloud%20SQL%20doesn't%20provide,balancing%20setup%20for%20better%20performance>

upvoted 2 times

🗳️ 👤 **Tempingtron** 9 months, 1 week ago

Should be A. Pgбouncer needs a load balancer or DNS roundrobin in front of it to operate. It can't route the traffic to multiple hosts without it. C, D are wrong.

upvoted 1 times

🗳️ 👤 **Jason_Cloud_at** 10 months, 1 week ago

Selected Answer: B

PgBouncer is especially used to manage connection pools to the PostgreSQL database.

upvoted 1 times

🗳️ 👤 **PKookNN** 11 months ago

Selected Answer: B

Cloud Load Balancing can't be used to LB cloud SQL (it's mostly for VM), so you can choose to use HAProxy or PgBouncer as Google recommend connection pooling (<https://cloud.google.com/sql/docs/postgres/replication#rr-info>)

upvoted 3 times

🗨️ 👤 **whoosh** 1 year ago

Selected Answer: B

By using PgBouncer, you can configure it to distribute the application load between the Cloud SQL primary and read replica instances. PgBouncer will handle connection pooling and load balancing, ensuring efficient utilization of resources and improving performance for your mission-critical application.

upvoted 1 times

🗨️ 👤 **AngieSoccerBall49** 1 year, 1 month ago

It should be (A). Specifically, you'd (most likely) use a TCP Load balancer

<https://www.pgbouncer.org/faq.html#how-to-load-balance-queries-between-several-servers>

"PgBouncer does not have an internal multi-host configuration. It is possible via external tools."

you need to frontend stateless pgBouncer instances with a TCP load balancer.

upvoted 2 times

🗨️ 👤 **Jay_Krish** 1 year, 1 month ago

Selected Answer: B

PgBouncer (Option B): PgBouncer is a lightweight connection pooler for PostgreSQL that can efficiently manage and distribute database connections between the primary and read replica instances. It helps in load balancing the application traffic between the instances.

upvoted 1 times

🗨️ 👤 **KennyHuang** 1 year, 7 months ago

Selected Answer: B

By using PgBouncer, you can configure it to distribute the application load between the Cloud SQL primary and read replica instances. PgBouncer will handle connection pooling and load balancing, ensuring efficient utilization of resources and improving performance for your mission-critical application.

upvoted 2 times

🗨️ 👤 **BenMS** 1 year, 9 months ago

Selected Answer: A

As others have said, A is the only option which could achieve the desired effect, providing TCP load balancing across multiple servers.

upvoted 3 times

🗨️ 👤 **Nirca** 1 year, 9 months ago

Selected Answer: A

A, I think is better.

HAProxy is not same as Cloud Load balancing.

upvoted 2 times

🗨️ 👤 **PrtkKA** 1 year, 10 months ago

Selected Answer: B

Connection pooling !

upvoted 2 times

🗨️ 👤 **Teraflow** 1 year, 10 months ago

Selected Answer: A

<https://cloud.google.com/blog/products/databases/using-haproxy-to-scale-read-only-workloads-on-cloud-sql-for-postgresql>

upvoted 4 times

🗨️ 👤 **PrtkKA** 1 year, 10 months ago

HAProxy is not same as Cloud Load balancing.

upvoted 1 times

🗨️ 👤 **felipeschossler** 1 year, 8 months ago

I think in the same way, it's not the same thing. However pgBouncer is not recommended for Load Balancing, just for connection pooling ☹️

upvoted 1 times

🗨️ 👤 **JayGeotab** 1 year, 11 months ago

A is the best answer, PgBouncer does not have multi-host

upvoted 2 times

🗨️ 👤 **sp57** 2 years ago

Per GPC72's referenced link, you need PgBouncer only does connection pooling, need Load balancing coupled. Since Load Balancing not referenced in B., is not A the best answer?

PgBouncer is a popular connection pooler designed for PostgreSQL, but it is not enough to achieve PostgreSQL High Availability by itself as it doesn't have multi-host configuration, failover, or detection.

Using a Load Balancer is a way to have High Availability in your database topology. It could be useful for redirecting traffic to healthy database nodes, distribute the

upvoted 3 times

🗨️ 👤 **sp57** 2 years ago

And C & D are wrong because they don't pool connections. ref for refuting C...The load balancer doesn't store database credentials (except for the health check user), and it doesn't pool or decrypt/re-encrypt database connections. A single client connection in HAProxy translates to a single client connection in Postgres.

This approach is suitable when the workload is constrained by the database's processing capacity, and not by the number of client connections. You may require an additional connection pooling component (e.g. PgBouncer) if the number of clients becomes an issue, for example, when the database instances exhibit performance or stability issues due to the sheer number of simultaneous database connections.

upvoted 1 times

🗨️ 👤 **sp57** 2 years ago

<https://cloud.google.com/blog/products/databases/using-haproxy-to-scale-read-only-workloads-on-cloud-sql-for-postgresql>

upvoted 1 times

🗨️ 👤 **pk349** 2 years ago

PgBouncer is a light-weight connection pool manager for Greenplum and PostgreSQL. PgBouncer maintains a pool for connections for each database and user combination. PgBouncer either creates a new database connection for a client or reuses an existing connection for the same user and database.

upvoted 1 times

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


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

Suggested Answer: B

Community vote distribution

C (88%)

13%

 **dynamic_dba** Highly Voted 1 year, 9 months ago

C.

The question specifically mentions binary logging and the binary logs are used by point-in-time recovery. D doesn't buy you anything since the corrupt data would also be on the HA replica you fail over to. B looks like a lot of work and if the Cloud SQL instance were instantiated a while ago, option B could take a long time. A would work but the backup could have been taken a while before the corruption began. In which case restoring using that backup would wipe all the good data up to the point of corruption. The question asks for minimal data loss and the only way to ensure that is to restore to a point-in-time just before the corruption began.

upvoted 7 times

 **fff2e69** Most Recent 3 months, 2 weeks ago

Selected Answer: C

Point in time recovery helps us to go back exactly a minute before corruption. This will leave us with most of the data before the corruption

upvoted 1 times

 **PKookNN** 11 months, 2 weeks ago

Selected Answer: C

unrelated but this was also a question from PCA. and the same answer C

upvoted 1 times

 **Hilab** 1 year, 9 months ago

C. Perform a point-in-time recovery of your Cloud SQL for MySQL database, selecting a date and time before the data was corrupted.

Performing a point-in-time recovery is the best option to minimize data loss in case of data corruption. Point-in-time recovery restores the database to a specific point in time before the data was corrupted, by replaying the binary logs that were generated since the selected time. This option is available when binary logging is enabled on Cloud SQL for MySQL with high availability.

Option A, restoring from an automated backup, can lead to data loss because it might not contain all the changes made to the database after the backup was taken. Option B, reloading the database from CSV files, can be time-consuming and may lead to data loss if the files used for initialization are not up to date. Option D, failing over to the Cloud SQL for MySQL HA instance, may not help in this scenario as the data corruption is replicated to the HA instance, and it is intended to be used for high availability and not for disaster recovery.

upvoted 2 times

 **zanhsieh** 1 year, 10 months ago

Selected Answer: A

A. Originally I thought it was C, but after reading mysql best practices as well as Kloudgeek link carefully I changed my answer. In Cloud SQL best-practice:

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

"A point-in-time recovery always creates a new instance; you cannot perform a point-in-time recovery to an existing instance."

Kloudgeek's link (why gcloud command use clone?):

<https://cloud.google.com/sql/docs/mysql/backup-recovery/pitr#perform-pitr-binlog>

"gcloud sql instances clone instance1 \

```
instance1-clone \  
--bin-log-file-name=mysql-bin.0000031 \  
--bin-log-position=107"
```

It seems to me that the question does not expect Cloud SQL instance switched just because of data corruption.

upvoted 1 times

🗉 👤 **pk349** 2 years ago

C: Perform a point-in-time recovery of your Cloud SQL for MySQL database, selecting a date and time ***** before the data was corrupted.

upvoted 1 times

🗉 👤 **GCP72** 2 years ago

Selected Answer: C

C is correct answer

upvoted 1 times

🗉 👤 **range9005** 2 years ago

Selected Answer: C

Binary logging --> Point in Recovery

upvoted 3 times

🗉 👤 **fredcaram** 2 years ago

Selected Answer: C

Since it is retaining transaction log, point in time recovery is enabled and that would be the best option

upvoted 2 times

🗉 👤 **Kloudgeek** 2 years ago

Correct Answer C: 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


upvoted 3 times

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.

Suggested Answer: BE

Community vote distribution



dynamic_dba Highly Voted 1 year, 9 months ago

C, E.

A is wrong because you want user accounts migrated as well as the data. B is nonsense. D is also nonsense since the DMS is an online migration. That leave C and E, both of which are mentioned in the Google doc:

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

upvoted 5 times

DeeData Most Recent 1 year, 3 months ago

You are sure to know the answer is CE if you did the labs

upvoted 4 times

RaphaelG 11 months, 2 weeks ago

so true!! I remember doing the challenge on my own for like 2 days to make certain I've got it all right; so the number of times I actually had to install pglogical and create primary keys on tables were ridiculous! :D

upvoted 3 times

Pilot50 1 year, 8 months ago

Selected Answer: CE

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

upvoted 3 times

Hilab 1 year, 9 months ago

C. Ensure that all PostgreSQL tables have a primary key.

E. Ensure that pglogical is installed on the source PostgreSQL database.

When using Database Migration Service to migrate data from a PostgreSQL on-premises instance to Cloud SQL, it is important to ensure that all PostgreSQL tables have a primary key. This is because Cloud SQL requires tables to have a primary key to enable replication and ensure data consistency.

It is also important to ensure that pglogical is installed on the source PostgreSQL database. This is because pglogical is used by Database Migration Service to replicate data changes from the source database to the target Cloud SQL instance.

Options A and D are not prerequisites for creating and automating the task. Option B is not recommended as it can cause data inconsistencies during the migration. Disabling foreign key constraints may result in data being migrated with foreign key constraint violations.

upvoted 3 times

Nirca 1 year, 9 months ago

Selected Answer: CE

Run the CREATE EXTENSION IF NOT EXISTS pglogical command on every database on your source instance. This installs the pglogical extension into the database.

For tables that don't have primary keys, Database Migration Service supports migration of the initial snapshot and INSERT statements during the change data capture (CDC) phase. You should migrate UPDATE and DELETE statements manually.

upvoted 1 times

🗳️ 👤 **TFMV** 1 year, 12 months ago

CE are correct.

upvoted 2 times

🗳️ 👤 **sp57** 2 years ago

E for sure; presume C is required for "automation" of task, because intervention required for tables without primary key.

For tables that don't have primary keys, Database Migration Service supports migration of the initial snapshot and INSERT statements during the change data capture (CDC) phase. You should migrate UPDATE and DELETE statements manually.

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

upvoted 1 times

🗳️ 👤 **sp57** 1 year, 12 months ago

Found this confirming C - was not option prior to 6/2022 - Tables without primary keys on the source PostgreSQL database are not migrated. For those tables, DMS migrated only the schema. This is no longer a limitation after the June 2022 product update.

Additional good content beyond scope of question in source link...<https://cloud.google.com/blog/products/databases/reduce-downtime-for-postgresql-migration-to-google-cloud-sql>

upvoted 2 times

🗳️ 👤 **chelbsik** 2 years ago

Selected Answer: CE

Vote for CE

upvoted 2 times

🗳️ 👤 **pk349** 2 years ago

C: Ensure that all PostgreSQL ***** tables have a primary key.

E: Ensure that pglogical ***** is installed on the source PostgreSQL database.

upvoted 1 times

🗳️ 👤 **GCP72** 2 years ago

Selected Answer: CE

CE

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

upvoted 1 times

🗳️ 👤 **h3clht** 2 years ago

Selected Answer: CE

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

upvoted 1 times

🗳️ 👤 **h3clht** 2 years ago

CE

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

upvoted 1 times

🗳️ 👤 **range9005** 2 years ago

Selected Answer: BE

Remove all the foreign key constraints

Come up with necessary pglogical

upvoted 2 times

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.

Suggested Answer: B

Community vote distribution



dynamic_dba Highly Voted 2 years, 3 months ago

C.

D is nonsense, so can be eliminated. The question tells us we're already managing BOTH sets of MySQL instances, one on prem and the other in GCE. The question also says an objective is to offload part of the management. That can only mean leverage a managed service. The Data(base) Migration Service is a managed service used to instantiate a new migrated DB in Cloud SQL (or AlloyDB for PostgreSQL but that's not in scope here). The question isn't asking about database migration, so we can eliminate B. A could be used to create replicas, but doesn't help with offloading management operations. That leaves C which does use a managed service which could be leveraged to create replicas.

upvoted 17 times

zanhsieh Highly Voted 2 years, 4 months ago

Selected Answer: A

A

A: CORRECT. External server replication meant to serve the case as the question describe: DC -> GSQL.

B: WRONG. DMS intends for one time migration + CDC but does not mean to serve as primary + replica fashion. The replica is not ready to serve the traffic.

C: WRONG. Cloud SQL for MySQL external replica meant to serve the replica is outside of Google cloud, which means GSQL -> GSQL or GSQL -> DC.

D: WRONG. This option is for one time only.

Reference:

<https://cloud.google.com/sql/docs/mysql/replication/external-server>

<https://cloud.google.com/sql/docs/mysql/replication>

https://cloud.google.com/database-migration/docs/overview#use_cases

upvoted 5 times

zanhsieh 2 years, 4 months ago

Correct the reference links:

<https://cloud.google.com/sql/docs/mysql/replication/configure-replication-from-external#curl>

<https://cloud.google.com/sql/docs/mysql/replication/configure-external-replica>

https://cloud.google.com/database-migration/docs/overview#use_cases

upvoted 1 times

nanay Most Recent 2 months, 1 week ago

Selected Answer: C

Cloud SQL External Replica (✓ C)

"Use this to replicate from an external source (like on-premises) to Cloud SQL. This allows you to offload reads and management to Cloud SQL."

upvoted 1 times

Devops2022 3 months ago

Selected Answer: C

Cloud SQL for MySQL external replica allows you to configure MySQL replicas on Google Cloud that can replicate from your on-premises MySQL databases. This is part of the hybrid configuration, where you can scale read operations and offload some management tasks by using Cloud SQL to manage the replicas. With this approach, Cloud SQL can serve as the replication target, allowing you to scale reads effectively without completely migrating to the cloud.

upvoted 1 times

🗄️ 👤 **887ad17** 5 months, 1 week ago

Selected Answer: C

C to offload part of the management operation.

upvoted 1 times

🗄️ 👤 **TNT87** 8 months, 1 week ago

<https://cloud.google.com/sql/docs/mysql/replication#external-read-replicas>

https://cloud.google.com/sql/docs/mysql/replication#replication_use_cases

upvoted 1 times

🗄️ 👤 **dunhill** 11 months, 1 week ago

Selected Answer: C

External replica is for scaling read

upvoted 1 times

🗄️ 👤 **Pime13** 1 year, 1 month ago

Selected Answer: A

<https://cloud.google.com/sql/docs/mysql/replication/external->

upvoted 1 times

🗄️ 👤 **Jason_Cloud_at** 1 year, 4 months ago

C. MySQL external replica are best for create replicas to scale reads and offload

upvoted 1 times

🗄️ 👤 **BIGQUERY_ALT_ALT** 1 year, 7 months ago

Selected Answer: C

C is the answer as Question mentions offload part of the management operation

upvoted 3 times

🗄️ 👤 **ewelaz** 1 year, 9 months ago

Selected Answer: C

it's c

upvoted 1 times

🗄️ 👤 **Kapello10** 1 year, 9 months ago

Selected Answer: C

The ans is C

upvoted 2 times

🗄️ 👤 **learnazureportal** 1 year, 9 months ago

A is correct - se external server replication ==> This option allows you to set up replication between your on-premises MySQL server (in your data center) and a MySQL server running on GOOGLE CLOUD COMPUTE ENGINE.

upvoted 2 times

🗄️ 👤 **KennyHuang** 2 years, 1 month ago

Selected Answer: A

By using external server replication, you can set up and manage replication between your on-premises MySQL database and a replica instance in Google Cloud. This enables you to scale reads, offload management operations, and distribute the workload between your data center and Google Cloud, providing the desired benefits in a hybrid configuration.

upvoted 4 times

🗄️ 👤 **Ieroygordo** 2 years, 3 months ago

Selected Answer: B

Agree with B.

- Multi-cloud continuous replication

Much like the read replicas across regions, if data exists in another cloud provider, a migration job can be set up which continuously replicates the database <<into Google Cloud for multi-cloud read-availability>>. Database Migration Service doesn't support a dual-write scenario, that is writing to and reading from both the source and destination.

https://cloud.google.com/database-migration/docs/overview#use_cases

upvoted 2 times

🗨️ 👤 **H_S** 2 years, 3 months ago

Selected Answer: C

offload part of the management operation => cloud sql => C

upvoted 1 times

🗨️ 👤 **H_S** 2 years, 3 months ago

<https://cloud.google.com/sql/docs/mysql/replication/external-server>

upvoted 2 times

🗨️ 👤 **Nirca** 2 years, 3 months ago

Selected Answer: A

A. Use external server replication.

<https://cloud.google.com/sql/docs/mysql/replication/external-server#config-description>

upvoted 1 times

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

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

Suggested Answer: CE

Community vote distribution

AC (100%)

🗳️ 👤 **Devops2022** 3 months ago

Selected Answer: AD

Can anyone explain why D is not the right answer?
upvoted 1 times

🗳️ 👤 **dunhill** 11 months, 1 week ago

Selected Answer: AC

replication from an external server is for db migration.
upvoted 1 times

🗳️ 👤 **julioevk** 1 year, 9 months ago

Selected Answer: AC

A because Database migration service is the managed offering
C because the external server is used for migration.
The others aren't approaches aren't actual methods for migration.
upvoted 1 times

🗳️ 👤 **Pilot50** 2 years, 2 months ago

Selected Answer: AC

no other choices are correct
upvoted 1 times

🗳️ 👤 **dynamic_dba** 2 years, 3 months ago

A, C.
B doesn't make sense since the DBs aren't in Cloud SQL yet. D and E don't makes sense because no method is attached to either answer and you wouldn't use a read replica as a source anyway. That leaves A and C. A makes sense since it can be scheduled and is online hence little/no downtime. Native DB replication of an external (on prem) server is basically what the Database Migration Service is doing. Which makes C correct as well.
upvoted 1 times

🗳️ 👤 **Ral17** 7 months ago

Why can't you use a read replica as a source for migration if it is being synced with the primary instance?
upvoted 1 times

🗳️ 👤 **Nirca** 2 years, 3 months ago

Selected Answer: AC

A & C is the correct Answer
upvoted 1 times

🗳️ 👤 **pk349** 2 years, 6 months ago

A: Use Database Migration Service ***** to migrate the databases to Cloud SQL.
C: Use replication ***** from an external server to migrate the databases to Cloud SQL.

upvoted 2 times

🗨️ 👤 **GCP72** 2 years, 6 months ago

A & C is the correct Answer

upvoted 2 times

🗨️ 👤 **lapeyus** 2 years, 6 months ago

Selected Answer: AC

<https://cloud.google.com/sql/docs/mysql/replication/manage-replicas#basic-steps-to-change-parallel-replication-flags>

upvoted 1 times

🗨️ 👤 **range9005** 2 years, 6 months ago

Selected Answer: AC

A is Database Migration Service to migrate the databases with CDC

C is Replication from an external server is used to Migration Database to Cloud SQL

upvoted 1 times

🗨️ 👤 **fredcaram** 2 years, 6 months ago

Selected Answer: AC

Using A for CDC makes sense to me and C is an option as well

upvoted 1 times

Your company is migrating the existing infrastructure for a highly transactional application to Google Cloud. You have several databases in a MySQL database instance and need to decide how to transfer the data to Cloud SQL. You need to minimize the downtime for the migration of your 500 GB instance. What should you do?

- A. 1. Create a Cloud SQL for MySQL instance for your databases, and configure Datastream to stream your database changes to Cloud SQL.
- 2. Select the Backfill historical data check box on your stream configuration to initiate Datastream to backfill any data that is out of sync between the source and destination.
- 3. Delete your stream when all changes are moved to Cloud SQL for MySQL, and update your application to use the new instance.
- B. 1. Create migration job using Database Migration Service.
- 2. Set the migration job type to Continuous, and allow the databases to complete the full dump phase and start sending data in change data capture (CDC) mode.
- 3. Wait for the replication delay to minimize, initiate a promotion of the new Cloud SQL instance, and wait for the migration job to complete.
- 4. Update your application connections to the new instance.
- C. 1. Create migration job using Database Migration Service.
- 2. Set the migration job type to One-time, and perform this migration during a maintenance window.
- 3. Stop all write workloads to the source database and initiate the dump. Wait for the dump to be loaded into the Cloud SQL destination database and the destination database to be promoted to the primary database.
- 4. Update your application connections to the new instance.
- D. 1. Use the mysqldump utility to manually initiate a backup of MySQL during the application maintenance window.
- 2. Move the files to Cloud Storage, and import each database into your Cloud SQL instance.
- 3. Continue to dump each database until all the databases are migrated.
- 4. Update your application connections to the new instance.

Suggested Answer: C

Community vote distribution

B (100%)

🗳️ **dynamic_dba** Highly Voted 1 year, 9 months ago

B.

A is wrong because Datastream is a CDC and replication service for data synchronization across heterogeneous databases. It's reasonable to assume you'll be using Cloud SQL for MySQL, so you'll be performing a homogeneous migration. Plus, while the a Datastream source can be MySQL, a Datastream target is either BigQuery or Cloud Storage and not Cloud SQL. See <https://cloud.google.com/datastream/docs/overview>.

C is wrong because a one time migration wouldn't capture all the data changes once the maintenance window ended and the apps were fired back up. Furthermore, stopping all writes during the dump would constitute downtime which the question wants minimized. D would take forever in a rapidly changing source system. B is the cleanest and simplest solution especially since the question puts no time constraint on making the migration happen.

upvoted 8 times

🗳️ **887ad17** Most Recent 5 months, 1 week ago

Selected Answer: B

only B have initiate a promotion of the new Cloud SQL instance

upvoted 1 times

🗳️ **dija123** 7 months, 1 week ago

Selected Answer: B

Agree with B

upvoted 1 times

🗳️ **Pilot50** 1 year, 8 months ago

Selected Answer: B

C isn't correct since it will not minimize the downtime

upvoted 2 times

🗳️ **Nirca** 1 year, 9 months ago

Selected Answer: B

B is the correct answer

upvoted 1 times

  **pk349** 2 years ago



B: Create migration job using Database ***** Migration Service.

Set the migration job type to Continuous, and allow the databases to complete the full dump phase and start sending data in change data capture (CDC) mode.

Wait ***** for the replication delay to minimize, initiate a promotion of the new Cloud SQL instance, and wait for the migration job to complete.

Update your application connections to the new instance.

upvoted 1 times

  **GCP72** 2 years ago

Selected Answer: B

B is the correct answer

upvoted 2 times

  **range9005** 2 years ago

Selected Answer: B

Continuous Migration with CDC

upvoted 2 times

  **fredcaram** 2 years ago

Selected Answer: B

C is not minimizing the downtime

upvoted 2 times

  **Kloudgeek** 2 years ago

Correct option is B. You need to minimize the downtime of the application but option C refers to stop the app while migration to complete.

upvoted 4 times

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.

Suggested Answer: B

Community vote distribution

C (100%)

🗳️ 👤 **Kloudgeek** Highly Voted 👍 1 year ago

Correct answer is C: https://cloud.google.com/sql/docs/mysql/using-ood-recommender#apply_recommendations
upvoted 6 times

🗳️ 👤 **dynamic_dba** Highly Voted 👍 9 months, 3 weeks ago

C.
A is wrong since modifying the schemas to 3NF would use more disk. D is nonsense. B sounds vague at best and probably not supported at the database level. C is the best answer. The link provided by Kloudgeek is spot on.
upvoted 6 times

🗳️ 👤 **Pilot50** Most Recent 🕒 8 months, 4 weeks ago

Selected Answer: C
out of disk means need for more space
upvoted 1 times

🗳️ 👤 **pk349** 1 year ago

C: Manually or automatically increase the storage capacity.
upvoted 1 times

🗳️ 👤 **GCP72** 1 year ago

Selected Answer: C
C is the correct answer
upvoted 2 times

🗳️ 👤 **lapeyus** 1 year ago

Selected Answer: C
<https://cloud.google.com/sql/docs/mysql/instance-settings#storage-capacity-2ndgen>
upvoted 1 times

🗳️ 👤 **range9005** 1 year ago

Selected Answer: C
Manually or automatically increase the storage capacity.
upvoted 2 times

You are managing a mission-critical Cloud SQL for PostgreSQL instance. Your application team is running important transactions on the database when another DBA starts an on-demand backup. You want to verify the status of the backup. What should you do?

- A. Check the `cloudsql.googleapis.com/postgres.log` instance log.
- B. Perform the `gcloud sql operations list` command.
- C. Use Cloud Audit Logs to verify the status.
- D. Use the Google Cloud Console.

Suggested Answer: C

Community vote distribution

B (85%)

Other

🗳️ 👤 **dynamic_dba** Highly Voted 🍌 1 year, 9 months ago

B.

A is wrong. The `cloudsql.googleapis.com/postgres.log` log file could be used to find out who started the backup operation, not the status of the operation. C is wrong for a similar reason. D is partially right. In the console there's an Operations option in the menu on the left. Click that and it shows "Creating backup" together a start time. That's not a million miles different from the `gcloud sql operations list` command which shows similar output except there's a STATUS: line showing the word RUNNING while the backup is in progress. Given the question specifically mentions "status", B would be the better answer. Just.

upvoted 7 times

🗳️ 👤 **honeymania23** Most Recent 🔔 8 months, 3 weeks ago

`gcloud sql backup list --instance= instance-id` will give the status of backup for that particular instance. So B is good.

upvoted 2 times

🗳️ 👤 **julioevk** 1 year, 2 months ago

Selected Answer: B

B to get the status using `gcloud`. D is possible but as mentioned, it doesn't specifically mention where in the console so B is a better answer.

upvoted 2 times

🗳️ 👤 **theseawillclaim** 1 year, 3 months ago

Selected Answer: B

If you want the status of the backup, then `gcloud` is the only viable option, as the Audit Logs will just tell you who started it.

upvoted 1 times

🗳️ 👤 **KennyHuang** 1 year, 7 months ago

Selected Answer: D

Using the Google Cloud Console is the most straightforward and convenient method for verifying the status of an on-demand backup for your mission-critical Cloud SQL for PostgreSQL instance. It provides a graphical interface that displays the backup status and any relevant details, enabling you to quickly assess the situation and ensure the integrity of your important transactions.

upvoted 1 times

🗳️ 👤 **Pilot50** 1 year, 8 months ago

Selected Answer: B

option C is to find out who started the backup, for the status B is correct

upvoted 1 times

🗳️ 👤 **Nirca** 1 year, 9 months ago

Selected Answer: B

B should be the right answer: via

`gcloud sql operations list --instance=<instance_name>`

`gcloud alpha sql operations list --instance=<instance_name>`

upvoted 1 times

🗳️ 👤 **Sekierer** 1 year, 11 months ago

Selected Answer: B

Should be B

Perform the gcloud sql operations list command

<https://cloud.google.com/sql/docs/postgres/backup-recovery/backups#troubleshooting-backups>

Under Troubleshooting:

Issue: "You can't see the current operation's status."

The Google Cloud console reports only success or failure when the operation is done. It isn't designed to show warnings or other updates. Run the gcloud sql operations list command to list all operations for the given Cloud SQL instance.

A and C is wrong because there are used for WHO issued the Backup, but not the current status of the backup

D is wrong only shows success or failure in the Cloud Console but not the current status of the backup

upvoted 4 times

🗲️ 👤 **chelbsik** 2 years ago

Selected Answer: B

Vote for B

upvoted 1 times

🗲️ 👤 **pk349** 2 years ago

B: Perform the gcloud sql operations list ***** command.

upvoted 1 times

🗲️ 👤 **GCP72** 2 years ago

Selected Answer: C

C is the correct answer.log gives more information

upvoted 1 times

🗲️ 👤 **range9005** 2 years ago

Selected Answer: B

Perform the gcloud sql operations list command.

upvoted 1 times

🗲️ 👤 **Kloudgeek** 2 years ago

Correct answer is B. <https://cloud.google.com/sql/docs/postgres/backup-recovery/backups#troubleshooting-backups>


upvoted 2 times

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.

Suggested Answer: A

Community vote distribution



🗳️ 👤 **887ad17** 5 months, 1 week ago

Selected Answer: A

not good answers ;) a think A+B correct
upvoted 1 times

🗳️ 👤 **honeymania23** 8 months, 3 weeks ago

A seems to be the first go to choice, if that does not resolve we can move to other options, But first A.
upvoted 2 times

🗳️ 👤 **njda** 1 year, 3 months ago

Selected Answer: A

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](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, you should manually add more nodes, or set), you should manually add more nodes, or set

Indexes and schema are reviewed post identifying query with slow performance. Refer : <https://cloud.google.com/spanner/docs/troubleshooting-performance>
upvoted 3 times

🗳️ 👤 **learnazureportal** 1 year, 3 months ago

I surprised with the chosen answer. The correct answer is B. When addressing high CPU utilization in a Google Cloud Spanner instance, you should first consider B. Modify the database schema, and add additional indexes. High CPU utilization in a database often occurs due to inefficient queries or lack of appropriate indexes.
upvoted 3 times

🗳️ 👤 **ArtistS** 1 year ago

So, any words mentioned the schema is not right? or it lacks index?
upvoted 1 times

🗳️ 👤 **KennyHuang** 1 year, 7 months ago

Selected Answer: D

By modifying the database schema and adding additional indexes, you can optimize query performance and potentially reduce the CPU utilization in Cloud Spanner. This approach focuses on improving the efficiency of the database and aligning it with the specific requirements of the consumer inventory application. It is important to monitor the impact of these changes and make further optimizations as needed.
upvoted 1 times

🗳️ 👤 **dynamic_dba** 1 year, 9 months ago

A.
B is wrong since that would increase CPU utilization even further and the question does not mention anything being wrong with index design. D is wrong since that would reduce CPU capacity and thus increase the load on the remaining CPUs. Cloud Spanner does not autoscale. It's up to you to allocate the number of nodes or processing units to keep CPU utilization under 65%. So add more processing units.
upvoted 3 times

🗳️ 👤 **pk349** 2 years ago

Compute capacity defines amount of server and storage resources that are available to the databases in an instance. When you create an instance, you specify its compute capacity as a number of processing ***** units or as a number of nodes, with 1000 processing units being equal to 1 node.



range9005

upvoted 2 times

  **pk349** 2 years ago

A: Increase the number of processing units.

upvoted 1 times

  **GCP72** 2 years ago

Selected Answer: A

A is correct answer

B is not correct because modifying schema is not a correct option

upvoted 2 times

  **range9005** 2 years ago

Selected Answer: A

Increase the number of processing units.

upvoted 2 times

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.

Suggested Answer: B

Community vote distribution

A (100%)

🗳️ **julioevk** 9 months ago

Selected Answer: A

SSD because this is the more highly performant PD type
upvoted 1 times

🗳️ **theseawillclaim** 9 months, 1 week ago

Selected Answer: A

A.
Data is split correctly on nodes if the row-key is well designed.
upvoted 1 times

🗳️ **CloudKida** 1 year ago

Selected Answer: A

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.
upvoted 1 times

🗳️ **dynamic_dba** 1 year, 3 months ago

A.
When you create a Bigtable instance you have to choose either SSD or HDD. The SSD options says, "Lower latency and more rows read per second. Typically used for real-time serving use cases, such as ad serving and mobile app recommendations". User facing plus low latency plus read intensive equals SSD.
upvoted 2 times

🗳️ **Nirca** 1 year, 3 months ago

Selected Answer: A

A is correct answer ,Question is about storage type (hardware)
upvoted 2 times

🗳️ **pk349** 1 year, 6 months ago

A: Recommend solid-state drives ***** (SSD).
upvoted 2 times

🗳️ **GCP72** 1 year, 6 months ago

Selected Answer: A

A is correct answer ,Question is about storage type so B is not a correct answer
upvoted 2 times

🗳️ **lapeyus** 1 year, 6 months ago

Selected Answer: A

SSD is significantly faster and has more predictable performance than HDD.
upvoted 1 times

🗳️ **range9005** 1 year, 6 months ago

Selected Answer: A

Recommend solid-state drives (SSD)

upvoted 2 times

  **fredcaram** 1 year, 6 months ago

Selected Answer: A

B is a right answer but it is not a storage type

upvoted 1 times

  **Kludgeek** 1 year, 6 months ago

Correct answer is A. <https://cloud.google.com/bigtable/docs/choosing-ssd-hdd>

upvoted 3 times

Your organization has a critical business app that is running with a Cloud SQL for MySQL backend database. Your company wants to build the most fault-tolerant and highly available solution possible. You need to ensure that the application database can survive a zonal and regional failure with a primary region of us-central1 and the backup region of us-east1. What should you do?

- A. 1. Provision a Cloud SQL for MySQL instance in us-central1-a.
2. Create a multiple-zone instance in us-west1-b.
3. Create a read replica in us-east1-c.
- B. 1. Provision a Cloud SQL for MySQL instance in us-central1-a.
2. Create a multiple-zone instance in us-central1-b.
3. Create a read replica in us-east1-b.
- C. 1. Provision a Cloud SQL for MySQL instance in us-central1-a.
2. Create a multiple-zone instance in us-east-b.
3. Create a read replica in us-east1-c.
- D. 1. Provision a Cloud SQL for MySQL instance in us-central1-a.
2. Create a multiple-zone instance in us-east1-b.
3. Create a read replica in us-central1-b.

Suggested Answer: B

Community vote distribution

B (100%)

🗳️ 👤 **Steve8512** 2 months, 3 weeks ago

Selected Answer: B

The answers are badly worded. B looks right, but reads like there will be two instances, one single-zone and one HA, in the same region. Presumably it's actually describing a single HA instance in central1.

upvoted 1 times

🗳️ 👤 **Ral17** 6 months, 3 weeks ago

Selected Answer: B

Cloud SQL is a regional service with read replicas allowed in other regions. So the answer must reference 2 different zones in the us-central1 region, one for the primary and one for the HA replica. A read replica needs to be in a zone within us-east1.

upvoted 1 times

🗳️ 👤 **dynamic_dba** 1 year, 3 months ago

B.

Cloud SQL is a regional service with read replicas allowed in other regions. So the answer must reference 2 different zones in the us-central1 region, one for the primary and one for the HA replica. A read replica needs to be in a zone within us-east1. The only options which provides that is B.

upvoted 4 times

🗳️ 👤 **sp57** 1 year, 6 months ago

B is correct. DR write-up helps... <https://cloud.google.com/sql/docs/sqlserver/intro-to-cloud-sql-disaster-recovery>

upvoted 2 times

🗳️ 👤 **pk349** 1 year, 6 months ago

D: Provision a Cloud SQL for MySQL instance in us-central1-a.

Create a multiple-zone instance in ***** us-east1-b.

Create a read replica in us-central1-b.

upvoted 1 times

🗳️ 👤 **GCP72** 1 year, 6 months ago

Selected Answer: B

B is the correct answer



upvoted 4 times

🗳️ 👤 **chelbsik** 1 year, 6 months ago

Selected Answer: B

Very confusing description. My only guess is that steps 1 and 2 describe the same action - creating primary instance with multiple zones HA. This eliminates all answers but B, because you can only have HA setup within the same region.

upvoted 2 times

  **range9005** 1 year, 6 months ago

I guess D

Primary Instance Us-Central

For HA, Multi-Region us-east

Replica on the top of primary i.e Us-central for low latency

upvoted 1 times

  **fredcaram** 1 year, 6 months ago

I got a little confused about the text, for B to be the correct answer the multiple-zone instance would be the stand-by instance and the read replica would be the cross-region read replica.

upvoted 1 times

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.

Suggested Answer: B

Community vote distribution

A (100%)

  **dynamic_dba** Highly Voted 9 months, 2 weeks ago

A.

B is wrong since that's not secure and doesn't make sense. C is bizarre and doesn't leverage a GCP serverless database. The key is intermittent internet coverage, meaning real-time syncing is not needed and can be supported. That rules out Spanner, which leaves Firestore. Probably Datastore mode, not that the question mentions that. The link provided by GCP72 is spot on.

upvoted 8 times

  **pk349** Most Recent 1 year ago

A: Use Firestore.

upvoted 1 times

  **GCP72** 1 year ago

Selected Answer: A

A is a correct answer , Cloud Firestone <https://firebase.google.com/docs/firestore>

upvoted 3 times

  **range9005** 1 year ago

Selected Answer: A

Android App -->> Cloud Firestone

upvoted 4 times

  **fredcaram** 1 year ago

Selected Answer: A

A supports offline sync

upvoted 2 times

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

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

Suggested Answer: BE

Community vote distribution

AE (83%)

Other

  **chelbsik** Highly Voted 1 year ago

Selected Answer: AE

<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

upvoted 12 times

  **dynamic_dba** Most Recent 9 months, 2 weeks ago

A, E.

The answers are split between significant and insignificant. For insignificant, the simplest form of recovery would be E. That eliminates D. For significant, let's assume that means a lot of data of the 50 TB total. A stale read and write back would probably be too onerous, so that eliminates B. That leaves A and C. The question doesn't mention anything about logical backups (export) which suggests a restore from a backup would be appropriate of a large amount of data that needed to be recovered.

<https://cloud.google.com/spanner/docs/pitr>

<https://cloud.google.com/spanner/docs/backup/restore-backup>

upvoted 4 times

  **TFMV** 12 months ago

AE are correct.

upvoted 3 times

  **pk349** 1 year ago

B: If the corruption is significant, perform a stale ***** read and specify a recovery timestamp. Write the results back.

D: If the corruption is insignificant, use backup and ***** restore, and specify a recovery timestamp.

upvoted 1 times

  **range9005** 1 year ago

Selected Answer: BD

B. If the corruption is significant, perform a stale read and specify a recovery timestamp. Write the results back.

D. If the corruption is insignificant, use backup and restore, and specify a recovery timestamp.

upvoted 1 times

  **range9005** 1 year ago

Selected Answer: BC

B. If the corruption is significant, perform a stale read and specify a recovery timestamp. Write the results back.

C. If the corruption is significant, use import and export.

upvoted 1 times

  **range9005** 1 year ago

By mistake

upvoted 1 times

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

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

Suggested Answer: C

Community vote distribution

A (100%)

🗳️ **dynamic_dba** 9 months, 2 weeks ago

A.

CPU and memory are OK so that eliminates B and D. C is nonsense which leaves A. This is supported by Google's own documentation (read recommended practices) which says close unused operations and re-start the instance. This is the best way to ensure maximum resources for the import operation.

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

upvoted 2 times

🗳️ **SVGoogle89** 1 year ago

C. for import service account needs storage.buckets.get & storage.objects.get

upvoted 1 times

🗳️ **pk349** 1 year ago

A: Close idle connections or restart the instance before beginning the import operation.

upvoted 2 times

🗳️ **GCP72** 1 year ago

Selected Answer: A

The import operation is taking too long. Too many active connections can interfere with import operations.

Close unused operations. Check the CPU and memory usage of your Cloud SQL instance to make sure there are plenty of resources available. The best way to ensure maximum resources for the import is to restart the instance before beginning the operation.

A restart:

Closes all connections.

Ends any tasks that may be consuming resources

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

upvoted 3 times

🗳️ **chelbsik** 1 year ago

Selected Answer: A

Eliminate B and D because 'You checked memory and CPU usage, and sufficient resources are available.'

Eliminate C because it makes no sense.

upvoted 1 times

🗳️ **chelbsik** 1 year ago

To elaborate on C - it's required for the export into Cloud Storage, which is not the case

https://cloud.google.com/sql/docs/postgres/import-export/import-export-csv#required_roles_and_permissions_for_exporting

upvoted 2 times

🗳️ **range9005** 1 year ago

Selected Answer: A

Close idle connections or restart the instance before beginning the import operation.

upvoted 3 times

  **fredcaram** 1 year ago

Selected Answer: A

It should be A given the amount of opened connections

upvoted 3 times

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.

Suggested Answer: A

Community vote distribution

A (100%)

🗳️ 👤 **dynamic_dba** 9 months, 2 weeks ago

A.

B is wrong since Serverless VPC Access is for connecting to your VPC network from serverless environments (Cloud Run, App Engine, Cloud Functions). C is wrong as this concerns private consumption of services across VPC networks that belong to different groups, teams, projects, or organizations. D is wrong because it concerns application networking for services. Nothing in its documentation mentions BMS. That leaves A. I would prefer to have seen something about VPC network peering, but the clincher is firewall rules which you would use to limit IP traffic sources to the backend Oracle DBs residing in their own Google managed VPC on BMS.

upvoted 2 times

🗳️ 👤 **H_S** 9 months, 3 weeks ago

Selected Answer: A

The answer is for sure A

read the following; <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>

upvoted 2 times

🗳️ 👤 **GCP72** 1 year ago

Selected Answer: A

A is the correct answer

upvoted 3 times

🗳️ 👤 **chelbsik** 1 year ago

Selected Answer: A

B is not applicable here.

C is also not the case - don't confuse it with Private Google Access <https://cloud.google.com/bare-metal/docs/bms-security#enforce-a-secure-perimeter-with> We don't have a Service Mesh here, so D is also not an option.

I go for A - <https://cloud.google.com/bare-metal/docs/bms-security#:~:text=As%20shown%20in%20Figure%206%2C%20use%20a%20shared%20VPC%20architecture%20to%20allow%20resources%20from%20different>

upvoted 3 times

🗳️ 👤 **range9005** 1 year ago

Selected Answer: A

When you use Shared VPC, you designate a project as a host project and attach one or more other service projects to it.

upvoted 3 times

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.

Suggested Answer: C

Community vote distribution

C (100%)

🗲️ 👤 **dynamic_dba** Highly Voted 👍 9 months, 2 weeks ago

C.

A and B are wrong because Spanner exports are run as Dataflow jobs. The question says you need a copy of your entire database which means all the tables. You cannot export an entire database using the CSV (text) format, but you can using the Avro format. So that would make it the better option.

<https://cloud.google.com/spanner/docs/import-export-overview#file-format>

upvoted 5 times

🗲️ 👤 **H_S** Most Recent 🔍 9 months, 3 weeks ago

Although I agree with C, I don't know why not D

upvoted 1 times

🗲️ 👤 **pk349** 1 year ago

C: Use Dataflow to run the export in Avro format.

upvoted 2 times

🗲️ 👤 **GCP72** 1 year ago

Selected Answer: C

Answer is C, Dataflow and Avro format.

Cloud functions has timeout Gen-1 6mins Gen-2 1hr

upvoted 3 times

🗲️ 👤 **range9005** 1 year ago

Selected Answer: C

Use Dataflow to run the export in Avro format.

.

<https://cloud.google.com/spanner/docs/export>


upvoted 4 times

You need to redesign the architecture of an application that currently uses Cloud SQL for PostgreSQL. The users of the application complain about slow query response times. You want to enhance your application architecture to offer sub-millisecond query latency. What should you do?

- A. Configure Firestore, and modify your application to offload queries.
- B. Configure Bigtable, and modify your application to offload queries.
- C. Configure Cloud SQL for PostgreSQL read replicas to offload queries.
- D. Configure Memorystore, and modify your application to offload queries.

Suggested Answer: D

Community vote distribution



pk349 Highly Voted 2 years ago

D: Configure Memorystore and modify your application to offload queries.
upvoted 8 times

csrazdan Highly Voted 1 year, 11 months ago

Selected Answer: D

The question says "redesign the architecture of an application" - Caching data using MemoryStore will trigger this. Read replica will require only updating connection details which would not be considered an application redesign.
upvoted 6 times

887ad17 Most Recent 5 months, 1 week ago

Selected Answer: D

Memorystore for caching
upvoted 1 times

theseawillclaim 1 year, 3 months ago

Selected Answer: D

"sub-millisecond latency" always involves Memorystore.
Furthermore, as we are talking about a relational DB (Cloud SQL), BigTable is not a solution to be considered.
upvoted 5 times

KIDDO24 1 year, 3 months ago

MEMORY STORE IS ALSO A NO RELATIONAL DB LIKE BIGTABLE
upvoted 1 times

RaphaelG 11 months, 2 weeks ago

"sub-milisecond latency" always involves Bigtable if anything
upvoted 2 times

learnazureportal 1 year, 3 months ago

The correct answer is ==> B. Configure Bigtable, and modify your application to offload queries.
upvoted 1 times

learnazureportal 1 year, 3 months ago

Memorystore is an in-memory data store that can provide low-latency access to cached data, but it may not be suitable for all types of queries, and achieving sub-millisecond latency depends on factors such as data size and query complexity.
upvoted 1 times

KennyHuang 1 year, 7 months ago

Selected Answer: D

The recommended approach is to configure Memorystore (Redis) and modify your application to offload queries. This will allow you to take advantage of the sub-millisecond query latency provided by in-memory caching and significantly improve the performance of your application.
upvoted 2 times

felipeschossler 1 year, 8 months ago

Selected Answer: D

D. The only thing that achieves submilli seconds of response is Redis: <https://blog.bytebytego.com/p/ep22-latency-numbers-you-should-know>
upvoted 2 times

🗨️ 👤 **PATILDXB** 1 year, 8 months ago

D.

To meet demands of low latency at increased scale and reduced cost you need an in-memory datastore. Redis and Memcached are among the most popular. Memorystore is a fully managed in-memory data store service for Redis and Memcached at Google Cloud.

upvoted 1 times

🗨️ 👤 **dynamic_dba** 1 year, 9 months ago

D.

A is wrong since Firestore would not offer sub-millisecond response. C would help, but sub-millisecond would still be hard to achieve. The question gives no justification for Bigtable and sub-millisecond response does strongly suggest reads from memory rather than disk. That leaves D as the best answer.

upvoted 4 times

🗨️ 👤 **muky31dec** 1 year, 11 months ago

Selected Answer: D

Ans is D

upvoted 1 times

🗨️ 👤 **ssaporylo** 1 year, 12 months ago

C: read replica

upvoted 2 times

🗨️ 👤 **SVGoogle89** 1 year, 12 months ago

sub-millisecs -> BigTable

upvoted 1 times

🗨️ 👤 **ssaporylo** 1 year, 12 months ago

it requires additional efforts to integrate BT

upvoted 1 times

🗨️ 👤 **GCP72** 2 years ago

Selected Answer: C

C is the correct answer

upvoted 2 times

🗨️ 👤 **GCP72** 2 years ago

Sorry D is the correct answer

<https://cloud.google.com/blog/topics/developers-practitioners/what-memorystore/>

upvoted 7 times

🗨️ 👤 **chelbsik** 2 years ago

Selected Answer: B

Can't find any proper docs for this case, so let's try elimination. I'm not so sure though.

A and C - don't give us sub-millisecond query latency.

I don't see how Memorystore can help with queries, since it's cache, so eliminate D as well.

upvoted 2 times

🗨️ 👤 **ssaporylo** 1 year, 12 months ago

C read replica with additional indexes give benefits. PG => BT dev efforts to migrate

upvoted 1 times


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.

Suggested Answer: B

Community vote distribution

C (100%)

 **fredcaram** Highly Voted 1 year, 6 months ago

Selected Answer: C

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>

upvoted 11 times

 **learnazureportal** Most Recent 9 months, 3 weeks ago

The correct answer is ==> 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

focus on "at least" keyword


upvoted 2 times

 **dynamic_dba** 1 year, 3 months ago

C.

D is wrong since the IOPS would not improve based upon the edition of SQL Server. IOPS increases with the amount of storage, so the most amount of storage is C with 4 TB. I checked this using the GCP console and C is correct.


upvoted 2 times

 **GCP72** 1 year, 6 months ago

Selected Answer: C

Agree C is the correct answer

upvoted 2 times

 **pk349** 1 year, 6 months ago

C: Create a SQL Server 2019 Standard on High Memory machine type with 16 vCPUs, ***** 104 GB of RAM, and 4 TB of SSD.

upvoted 1 times

 **Popa** 1 year, 6 months ago

Selected Answer: C

A disk size of 4TB or greater provides more throughput and IOPS. Storage: >= 4TB for the best IOPS

<https://cloud.google.com/sql/docs/sqlserver/best-practices#admin>

upvoted 3 times

You are managing a small Cloud SQL instance for developers to do testing. The instance is not critical and has a recovery point objective (RPO) of several days. You want to minimize ongoing costs for this instance. What should you do?

- A. Take no backups, and turn off transaction log retention.
- B. Take one manual backup per day, and turn off transaction log retention.
- C. Turn on automated backup, and turn off transaction log retention.
- D. Turn on automated backup, and turn on transaction log retention.

Suggested Answer: B

Community vote distribution

C (100%)

🗳️ 👤 **MikeFR** 1 month ago

Selected Answer: C

automated backup to manage deletion. not need to keep transaction logs as no future analysis
upvoted 1 times

🗳️ 👤 **hussain.sain** 1 year ago

Why not A. its a testing environment. why we need to have backup when we are low on budget.
upvoted 2 times

🗳️ 👤 **julioirevk** 1 year, 9 months ago

Selected Answer: C

C and not B because as per: <https://cloud.google.com/sql/docs/mysql/backup-recovery/backups>
On-demand backups are not automatically deleted the way automated backups are. They persist until you delete them or until their instance is deleted. Because they are not automatically deleted, on-demand backups can have a long-term effect on your billing charges.
upvoted 1 times

🗳️ 👤 **theseawillclaim** 1 year, 9 months ago

Selected Answer: C

C is the one.
Manual backups are always discouraged, and transaction log can be removed for a cheap, dev DB.
upvoted 1 times

🗳️ 👤 **felipeschossler** 2 years, 2 months ago

Selected Answer: C

I think that is C but I didn't find any link that corroborates with my opinion 😊
upvoted 1 times

🗳️ 👤 **Pilot50** 2 years, 2 months ago

Selected Answer: C

B is manual process and can't be the right approach for automation
upvoted 1 times

🗳️ 👤 **dynamic_dba** 2 years, 3 months ago

C.
A is wrong since there is an RPO. B requires manual intervention which partly defeats the object of using a managed service like Cloud SQL. D is wrong since retaining transaction logs would permit point-in-time recovery which is not required. That leaves C.
upvoted 3 times

🗳️ 👤 **csrazdan** 2 years, 5 months ago

Selected Answer: C



Automatic backups are incremental where as manual backups are full. Other than compute time for manual backup, storage costs will also increase.
upvoted 1 times

🗳️ 👤 **GCP72** 2 years, 6 months ago

Selected Answer: C

C. Turn on automated backup, and turn off transaction log retention.

upvoted 3 times

  **pk349** 2 years, 6 months ago

C: Turn on automated backup, and turn off transaction log retention.

upvoted 1 times

  **fredcaram** 2 years, 6 months ago

Selected Answer: C

There is no need to have the overload of using a manual backup, you could schedule an automatic one once a day

upvoted 2 times

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.

Suggested Answer: B

Community vote distribution

A (100%)

🗳️ **fredcaram** Highly Voted 2 years ago

Selected Answer: A

A makes more sense than changing products
upvoted 5 times

🗳️ **LaxmanTiwari** Most Recent 11 months, 2 weeks ago

Selected Answer: A
upvoted 1 times

🗳️ **vodkaHN** 1 year ago

Selected Answer: A

Changing service is too risky. ==> A
upvoted 1 times

🗳️ **dynamic_dba** 1 year, 9 months ago

A.
Migrating to a different platform is a bit extreme, so eliminate B. C would be a good idea in any event, but is not the best answer here. D is not how you manage maintenance windows on Cloud SQL. A is the best answer.
upvoted 1 times

🗳️ **GCP72** 2 years ago

Selected Answer: A

A is the correct answer
upvoted 4 times

🗳️ **pk349** 2 years ago

A: Set your production instance's maintenance window to ***** non-business hours.
upvoted 2 times

🗳️ **lapeyus** 2 years ago

Selected Answer: A

maintenance window
upvoted 3 times

🗳️ **Kloudgeek** 2 years ago

Correct Answer is A
upvoted 2 times

You are designing a highly available (HA) Cloud SQL for PostgreSQL instance that will be used by 100 databases. Each database contains 80 tables that were migrated from your on-premises environment to Google Cloud. The applications that use these databases are located in multiple regions in the US, and you need to ensure that read and write operations have low latency. What should you do?

- A. Deploy 2 Cloud SQL instances in the us-central1 region with HA enabled, and create read replicas in us-east1 and us-west1.
- B. Deploy 2 Cloud SQL instances in the us-central1 region, and create read replicas in us-east1 and us-west1.
- C. Deploy 4 Cloud SQL instances in the us-central1 region with HA enabled, and create read replicas in us-central1, us-east1, and us-west1.
- D. Deploy 4 Cloud SQL instances in the us-central1 region, and create read replicas in us-central1, us-east1 and us-west1.

Suggested Answer: B

Community vote distribution

A (67%)

C (33%)

🗳️ 👤 **SandyZA** Highly Voted 1 year, 11 months ago

A is correct. We only have 8000 tables. More than 1 HA env is only required beyond 50000 tables

If you have 50,000 or more database tables on a single instance, it could result in the instance becoming unresponsive or unable to perform maintenance operations, and the instance is not covered by the SLA.

upvoted 8 times

🗳️ 👤 **H_S** 1 year, 9 months ago

although I agree with your answe, table limit is for mysql and not postgresql

https://cloud.google.com/sql/docs/mysql/quotas#table_limit

upvoted 1 times

🗳️ 👤 **dynamic_dba** Highly Voted 1 year, 9 months ago

A.

B and D do not mention HA, so eliminate those. That leaves A and C. C talks about 4 instances with HA which presumably means 2 primaries each with an HA standby. Oddly, there are 4 zones in us-central1. The killer is having a read replica also in us-central1 which would mean the same zone would have a read replica and either a primary or HA standby. Not a good idea. Option A is the best choice. A primary and an HA standby in us-central1 (different zones) and then read replicas in us-east1 and us-west1.

upvoted 7 times

🗳️ 👤 **LunarLander** Most Recent 1 month, 4 weeks ago

Selected Answer: C

The question said to minimize read AND write latency, and that doesn't occur obviously with Read Replicas. The only solution that minimizes both R&W latency AND covers the regions is C. Answer A is invalid for that.

upvoted 1 times

🗳️ 👤 **PKookNN** 11 months, 2 weeks ago

Selected Answer: A

as I focus on low latency

upvoted 2 times

🗳️ 👤 **BIGQUERY_ALT_ALT** 1 year, 1 month ago

The answer is C. you need a instance for a read replica not 2 per read replica.

upvoted 1 times

🗳️ 👤 **Blackstile** 1 year, 2 months ago

Selected Answer: A

Answer C is wrong because the question suggests low-latency and replication for 4 region INCREASE latency.

Answer A is correct because attend the two requirements high availability and low latency.

The Key for this question is low latency.

upvoted 5 times

🗳️ 👤 **dija123** 7 months, 1 week ago

Thanks for the clarification.

upvoted 2 times

🗨️ 👤 **learnazureportal** 1 year, 3 months ago

The correct answer is == C. Deploy 4 Cloud SQL instances in the us-central1 region with HA enabled, and create read replicas in us-central1, us-east1, and us-west1.

upvoted 1 times

🗨️ 👤 **Jason_Cloud_at** 10 months, 1 week ago

we can't have both primary and read replica in same region, it will not be considered as HA
i will stick with A.

upvoted 1 times

🗨️ 👤 **KennyHuang** 1 year, 7 months ago

Selected Answer: C

This option provides high availability with four instances in the primary region, ensuring redundancy and fault tolerance. By creating read replicas in us-central1, us-east1, and us-west1, read operations can be distributed across multiple regions, reducing latency for applications in those regions.

This design allows for efficient and low-latency read operations while maintaining high availability.

Option C is the recommended choice as it combines HA with multiple read replicas in different regions, providing both high availability and low-latency read operations for your multi-region application setup.

upvoted 3 times

🗨️ 👤 **GCP72** 2 years ago

Selected Answer: C

Why not " C" is an answer

upvoted 3 times

🗨️ 👤 **pk349** 2 years ago

A: Deploy 2 Cloud SQL instances in the us-central1 region ***** with HA enabled, and create read replicas in us-east1 and us-west1.

upvoted 4 times

🗨️ 👤 **chelbsik** 2 years ago

Selected Answer: A

We only need options with HA enabled and 4 databases with 3 read replicas (each?) seems overkill to me.


upvoted 6 times

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.

Suggested Answer: A

Community vote distribution



chelbsik Highly Voted 2 years ago

Selected Answer: B

<https://cloud.google.com/blog/topics/developers-practitioners/scheduling-cloud-sql-exports-using-cloud-functions-and-cloud-scheduler>
upvoted 8 times

Pime13 Most Recent 8 months ago

Selected Answer: B

https://cloud.google.com/sql/docs/mysql/backup-recovery/scheduling-backups#creating_a_pub_sub_topic_a_cloud_function_and_a_cloud_scheduler_job
upvoted 2 times

nhiguchi 1 year, 1 month ago

Selected Answer: C

C is correct.
upvoted 2 times

learnazureportal 1 year, 3 months ago

The most accurate answer based on Google-recommended practices would be C
upvoted 1 times

KennyHuang 1 year, 7 months ago

Selected Answer: C

C is the recommended choice as it leverages Cloud Composer's workflow orchestration capabilities to schedule and automate the export process. By calling the `cloudsql.instances.export` API within the workflow, you can ensure that the data is exported from Cloud SQL for MySQL in CSV format as needed by your data analysis team.

upvoted 2 times

KennyHuang 1 year, 7 months ago

Selected Answer: C

Cloud Composer is a fully managed workflow orchestration service. It allows you to define and manage complex workflows using Apache Airflow. By using Cloud Composer, you can create a workflow that includes a task to export data from Cloud SQL for MySQL using the `cloudsql.instances.export` API. You can specify the export format as CSV to meet the requirement of your data analysis team. This approach provides a scalable and manageable solution for regular data exports.

Therefore, option C is the recommended choice as it leverages Cloud Composer's workflow orchestration capabilities to schedule and automate the export process. By calling the `cloudsql.instances.export` API within the workflow, you can ensure that the data is exported from Cloud SQL for MySQL in CSV format as needed by your data analysis team.

upvoted 1 times

dynamic_dba 1 year, 9 months ago

B.

Performing a "SELECT * FROM TABLE" wouldn't give you CSV output and that alone wouldn't call an API. Eliminate A. There's no need for Cloud Composer in this scenario, especially when the solution is a known combination of Cloud Scheduler, Cloud Function and a Pub/Sub Topic, which is B.
upvoted 4 times

🗨️ 👤 **njda** 1 year, 3 months ago

why will someone use 3 services when it can be done using 1 service only, in this case composer. Secondly export can be for multiple tables, and ,cloud functions has a execution time limit.?

C still looks most appropriate solution.

upvoted 3 times

🗨️ 👤 **Nirca** 1 year, 9 months ago

Selected Answer: B

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

upvoted 2 times

🗨️ 👤 **[Removed]** 1 year, 11 months ago

what if the query is large, Cloud Function has limit? shouldn't it be C?

upvoted 2 times

🗨️ 👤 **sp57** 2 years ago

As noted in text from this link, Cloud Function can be triggered with/without pub/sub. I think distinction between A & B is suggestion in A the API s called from query, which is not correct - the 'select query' is passed as parameter in call, so I think B is correct.

"Note that we could also have a Pub/Sub Trigger configured Cloud Function get triggered by the Scheduler by using the same steps given above to create a job. Except that in the target, you will select Pub/Sub and provide the Pub/Sub Topic name. At the scheduler trigger time, the Cloud Scheduler will publish a message to the topic with the message body that you specify."

<https://rominirani.com/google-cloud-functions-tutorial-using-the-cloud-scheduler-to-trigger-your-functions-756160a95c43>

upvoted 3 times

🗨️ 👤 **GCP72** 2 years ago

Selected Answer: A

A is the correct answer, PUBSUB is not required for this task

upvoted 3 times

🗨️ 👤 **cardareel** 1 year, 3 months ago

Doing a SELECT * won't call the API. Pub/Sub is not needed, you are correct, the problem is the SELECT * (nothing to do there).

upvoted 1 times

🗨️ 👤 **pk349** 2 years ago

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.

upvoted 2 times

🗨️ 👤 **cardareel** 1 year, 3 months ago

Doing a SELECT * won't call the API. Pub/Sub is not needed, you are correct, the problem is the SELECT * (nothing to do there).

upvoted 1 times

🗨️ 👤 **gabrielosluz** 1 year, 10 months ago

The answer is letter B. Google recommends this pattern to perform this task: <https://cloud.google.com/sql/docs/mysql/backup-recovery/scheduling-backups>


upvoted 2 times

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.

Suggested Answer: C

Community vote distribution



fredcaram Highly Voted 2 years, 6 months ago

Selected Answer: D

This is a bigtable use case
upvoted 7 times

hussain.sain Most Recent 1 year ago

point to consider is :
IOT data,
scaling for accepting millions of request.
storing historic data of 100TB.
Big Table is used for low latency but its for operational workload and
we need Fully managed Redis and Memcached for sub-millisecond data access.

that's the only argument here.

i think option c is correct.
upvoted 1 times

hussain.sain 1 year ago

<https://medium.com/google-cloud/which-database-should-i-choose-44be039179ea>
upvoted 1 times

theseawillclaim 1 year, 9 months ago

Selected Answer: D

Key is 100TB of historical data.
BigTable is the only service (among the mentioned) that is able to handle that load within the required latency.
upvoted 1 times

fsmathias 2 years ago

Selected Answer: C

Correct answer is C
upvoted 1 times

dynamic_dba 2 years, 3 months ago

D.
Simply a classic use case for Bigtable. Neither Cloud SQL, Firestore nor Memorystore have either the capacity or latency to provide the solution.
upvoted 2 times

lepach 2 years, 5 months ago


https://cloud.google.com/memorystore/docs/redis/redis-overview#what_its_good_for
What it's good for:

Stream Processing: Whether processing a Twitter feed or stream of data from IoT devices, Memorystore for Redis is a perfect fit for streaming

solutions.



Correct answer is C

upvoted 1 times

  **felipeschossler** 2 years, 2 months ago

You just forgot the 100TB of data that needs to be stored ☹

upvoted 2 times

  **GCP72** 2 years, 6 months ago

Selected Answer: D

C is not a correct answer, D- Bigtable is correct choice for this user case.

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

upvoted 2 times

  **range9005** 2 years, 6 months ago

Selected Answer: D

IOT -->> Bigtable

upvoted 2 times

  **Kludgegeek** 2 years, 6 months ago

Answer is D


upvoted 3 times

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.

Suggested Answer: C

Community vote distribution



🗳️ **zeta_xs** 11 months ago

Selected Answer: A

A, fault region

upvoted 1 times

🗳️ **dynamic_dba** 2 years, 3 months ago

A.

Avoiding user disruption if a regional failure occurs means you need to pick a multi-region service. That rules out C and D. Having more control over the EKs means CMEK. That eliminates B.

upvoted 4 times

🗳️ **GCP72** 2 years, 6 months ago

Selected Answer: A

A is the correct answer because " you want to control where you store the encryption key"

upvoted 4 times

🗳️ **pk349** 2 years, 6 months ago

C: Use Cloud SQL with a customer-managed encryption key (CMEK).

upvoted 1 times

🗳️ **chelbsik** 2 years, 6 months ago

Selected Answer: A

A for me: it's A or C because we want to control keys, but C would cause downtime, since we would need to manually failover to another region if a regional failure occurs, and we don't want that.

upvoted 2 times

🗳️ **jit028** 2 years, 6 months ago

Selected Answer: A

Correct Answer - A

upvoted 1 times

🗳️ **range9005** 2 years, 6 months ago

Selected Answer: B

I guess B

Since Google cloud default encryption comes with AES-256 encryption

upvoted 1 times

🗳️ **Tharun1125438** 2 years, 5 months ago


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.

upvoted 1 times

🗳️ **fredcaram** 2 years, 6 months ago

A and C would work for this scenario

upvoted 2 times

  **chelbsik** 2 years, 6 months ago

Right, but C would cause downtime, since you would need to manually failover to another region
upvoted 2 times

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.

Suggested Answer: B

Community vote distribution



🗳️ 👤 **dija123** 7 months, 1 week ago

Selected Answer: C

Agree with C

upvoted 1 times

🗳️ 👤 **DevBugger** 9 months ago

Selected Answer: C

For MySQL, this should be C

upvoted 1 times

🗳️ 👤 **PKookNN** 11 months ago

Selected Answer: C

For MySQL, this should be C

upvoted 1 times

🗳️ 👤 **kyou1313** 11 months ago

Selected Answer: C

When a Cloud SQL read replica becomes unavailable, Cloud SQL automatically tries to create a new replica to replace the failed one. This process is part of the automated failover mechanism in Cloud SQL.

B. Manually creating a new failover replica from backup might involve additional steps and may not be necessary for a temporary unavailability of the read replica.

upvoted 2 times

🗳️ 👤 **whoosh** 1 year ago

Selected Answer: C

"During a zonal outage, traffic stops to read replicas in that zone. After the zone becomes available again, any read replicas in the zone resume replication from the primary instance."

https://cloud.google.com/sql/docs/mysql/high-availability#read_replicas

upvoted 1 times

🗳️ 👤 **kfiry0** 1 year, 1 month ago

C is the correct answer.

"During a zonal outage, traffic stops to read replicas in that zone. After the zone becomes available again, any read replicas in the zone resume replication from the primary instance."

https://cloud.google.com/sql/docs/mysql/high-availability#read_replicas

upvoted 1 times

🗳️ 👤 **ArtistS** 1 year ago

I think is D, if sth wrong it didn't creat automatically. So B is correct

upvoted 1 times

🗳️ 👤 **ArtistS** 1 year ago

srr it is B, I type it by mistake

upvoted 1 times

🗳️ 👤 **learnazureportal** 1 year, 3 months ago

B is correct answer. Creating a new failover replica from a backup is a reliable way to restore replication and ensure that the read replica is up-to-date.
upvoted 1 times

🗳️ 👤 **[Removed]** 1 year, 5 months ago

Selected Answer: B

B is correct
upvoted 1 times

🗳️ 👤 **Sandipcst** 1 year, 6 months ago

C

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.

upvoted 2 times

🗳️ 👤 **KennyHuang** 1 year, 7 months ago

Selected Answer: C

By verifying that the new replica is created automatically, you can ensure that the read replica instance is functioning and replication is maintained even after the temporary unavailability of the replica in Zone B.

upvoted 2 times

🗳️ 👤 **wolfie09** 1 year, 6 months ago

A read replica is not recreated automatically.

upvoted 1 times

🗳️ 👤 **Pilot50** 1 year, 8 months ago

Selected Answer: D

<https://cloud.google.com/sql/docs/mysql/replication#:~:text=If%20replication%20is,a%20new%20one>.

upvoted 2 times

🗳️ 👤 **Pilot50** 1 year, 8 months ago

Selected Answer: D

Resume replication

upvoted 1 times

🗳️ 👤 **marpayer** 1 year, 11 months ago

B is wrong, failover replica is NOT a read replica

A and C makes no sense

upvoted 2 times

🗳️ 👤 **sp57** 2 years ago

How is B correct? You don't create a replica from backup, and there's no mention of HA that would point to "failover" distinction. But don't like other answers either. Will go with B unless enlightened by subsequent contributor,

upvoted 1 times

🗳️ 👤 **GCP72** 2 years ago

Selected Answer: B

B is the correct answer

upvoted 3 times

🗳️ 👤 **chelbsik** 2 years ago

Selected Answer: B

Vote for B

upvoted 2 times

  **chelbsik** 2 years ago

C - makes no sense

D - nobody said primary instance was offline, plus you can't stop/resture replication on the primary instance, only on read replica

A makes no sense: you can't create clone database from read replica, and if it means to create one from the primary instance - how would that help to insure that replica is still working?

I'll go for B

upvoted 4 times

You are migrating an on-premises application to Google Cloud. The application requires a high availability (HA) PostgreSQL database to support business-critical functions. Your company's disaster recovery strategy requires a recovery time objective (RTO) and recovery point objective (RPO) within 30 minutes of failure. You plan to use a Google Cloud managed service. What should you do to maximize uptime for your application?

- A. Deploy Cloud SQL for PostgreSQL in a regional configuration. Create a read replica in a different zone in the same region and a read replica in another region for disaster recovery.
- B. Deploy Cloud SQL for PostgreSQL in a regional configuration with HA enabled. Take periodic backups, and use this backup to restore to a new Cloud SQL for PostgreSQL instance in another region during a disaster recovery event.
- C. Deploy Cloud SQL for PostgreSQL in a regional configuration with HA enabled. Create a cross-region read replica, and promote the read replica as the primary node for disaster recovery.
- D. Migrate the PostgreSQL database to multi-regional Cloud Spanner so that a single region outage will not affect your application. Update the schema to support Cloud Spanner data types, and refactor the application.

Suggested Answer: C

Community vote distribution

C (100%)

🗳️ 👤 **dynamic_dba** 9 months, 2 weeks ago

C.

D might be possible but it's a lot of effort to migrate to a different platform. Eliminate D. A does not mention HA. Eliminate A. B says to take periodic backups which doesn't support an RTO/RPO of 30 minutes. The best answer is deploy an HA configuration and have a read replica you could promote to the primary in a different region. C is the best answer.

upvoted 4 times

🗳️ 👤 **GCP72** 1 year ago

Selected Answer: C

C is the correct answer

upvoted 3 times

🗳️ 👤 **pk349** 1 year ago

C: Deploy Cloud SQL for PostgreSQL in a regional configuration with HA enabled. Create a cross-region read replica, and promote ***** the read replica as the primary node for disaster recovery.

upvoted 1 times

🗳️ 👤 **chelbsik** 1 year ago

Selected Answer: C

I'll go with C

upvoted 1 times

Your team is running a Cloud SQL for MySQL instance with a 5 TB database that must be available 24/7. You need to save database backups on object storage with minimal operational overhead or risk to your production workloads. What should you do?

- A. Use Cloud SQL serverless exports.
- B. Create a read replica, and then use the mysqldump utility to export each table.
- C. Clone the Cloud SQL instance, and then use the mysqldump utility to export the data.
- D. Use the mysqldump utility on the primary database instance to export the backup.

Suggested Answer: C

Community vote distribution

A (90%)

10%

🗳️ 👤 **GCP72** Highly Voted 👍 2 years ago

Selected Answer: A

A is the correct answer

<https://cloud.google.com/blog/products/databases/introducing-cloud-sql-serverless-exports>

upvoted 5 times

🗳️ 👤 **Pime13** Most Recent 🕒 8 months ago

Selected Answer: A

A: <https://cloud.google.com/blog/products/databases/introducing-cloud-sql-serverless-exports>

Serverless exports enables you to export data from your MySQL and PostgreSQL database instances without any impact on performance or risk to your production workloads.

upvoted 1 times

🗳️ 👤 **dynamic_dba** 1 year, 9 months ago

A.

Minimal operational overhead eliminates B and C. Minimal risk to production workloads eliminates D. That leaves A. Least amount of work and doesn't impact the primary instance.

upvoted 4 times

🗳️ 👤 **Nirca** 1 year, 9 months ago

Selected Answer: C

Cloud SQL backups are incremental. They contain only data that changed after the previous backup was taken. Your oldest backup is a similar size to your database, but the sizes of subsequent backups depend on the rate of change of your data. When the oldest backup is deleted, the size of the next oldest backup increases so that a full backup still exists.

upvoted 1 times

🗳️ 👤 **Nirca** 1 year, 9 months ago

With serverless export, Cloud SQL creates a separate, temporary instance to offload the export operation. Offloading the export operation allows databases on the primary instance to continue to serve queries and perform operations at the usual performance rate. BUT is is export (logical backup) and will never be incremental. and the recovery is slow. for 5TB server is it not an option. (only for mini databases). I believe The better option is C

upvoted 1 times

🗳️ 👤 **pk349** 2 years ago

A: Use Cloud SQL serverless exports.

Serverless exports enables you to export data from your MySQL and PostgreSQL database instances without any impact on performance or risk to your production workloads. Cloud SQL exports, which offer portable data formats (SQL, CSV), can be triggered anytime and are written to Cloud Storage buckets that you control.



upvoted 1 times

🗳️ 👤 **chelbsik** 2 years ago

Selected Answer: A

<https://cloud.google.com/sql/docs/mysql/import-export#:~:text=Use%20serverless%20export,is%20deleted%20automatically.>

upvoted 1 times

  **jitu028** 2 years ago

Selected Answer: A

Correct Answer - A

upvoted 1 times

  **range9005** 2 years ago

Selected Answer: A

Use Cloud SQL serverless exports.

upvoted 1 times

You are deploying a new Cloud SQL instance on Google Cloud using the Cloud SQL Auth proxy. You have identified snippets of application code that need to access the new Cloud SQL instance. The snippets reside and execute on an application server running on a Compute Engine machine. You want to follow Google-recommended practices to set up Identity and Access Management (IAM) as quickly and securely as possible. What should you do?

- A. For each application code, set up a common shared user account.
- B. For each application code, set up a dedicated user account.
- C. For the application server, set up a service account.
- D. For the application server, set up a common shared user account.

Suggested Answer: B

Community vote distribution

C (100%)

🗳️ 👤 **Lucker** 10 months, 2 weeks ago

the best answer would be B . Need granular control for each of the code snippets . The advantage of using a service account for this purpose is that you can create a credential file specifically for the Cloud SQL Auth Proxy, and it is explicitly and permanently linked to the Cloud SQL Auth Proxy as long as it is running. For this reason, using a service account is the recommended method for production instances not running on a Compute Engine instance.

upvoted 1 times

🗳️ 👤 **felipeschossler** 2 years, 2 months ago

C. The docs proving this are here: Create and configure a Google Cloud service account that has the Cloud SQL Client role with permissions to connect to Cloud SQL.

upvoted 1 times

🗳️ 👤 **dynamic_dba** 2 years, 3 months ago

C.

The Google recommendation would be for an application to use a service account. None of the other options make sense in light of this.

upvoted 2 times

🗳️ 👤 **GCP72** 2 years, 6 months ago

Selected Answer: C

Correct answer is C, service account is the GCP recommended option

upvoted 3 times

🗳️ 👤 **pk349** 2 years, 6 months ago

B: For each application code, set up a dedicated ***** user account.

upvoted 1 times

🗳️ 👤 **chelbsik** 2 years, 6 months ago

Selected Answer: C

<https://cloud.google.com/sql/docs/mysql/sql-proxy#using-a-service-account>

upvoted 3 times

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).

Suggested Answer: B

Community vote distribution

C (100%)

  **dynamic_dba** Highly Voted 2 years, 3 months ago

C.

Cloud SQL doesn't support SSAS, SSRS or SSIS. Eliminate A. Migrating to a different database platform doesn't make sense. Eliminate B. Minimal disruption to the existing architecture rules out GKE. Eliminate D. Leaves C. Migrate to a GCE VM.

upvoted 5 times


  **cert_pz** Most Recent 8 months, 1 week ago

A)

Migrating to Cloud SQL for SQL Server (option A) is the optimal solution for ensuring minimal disruption while leveraging a managed service that maintains compatibility with your existing Microsoft SQL Server architecture. This approach streamlines the migration process and minimizes operational overhead.

C. Migrate to Compute Engine: While deploying SQL Server on Compute Engine gives you more control over the environment, it requires more management effort (e.g., handling backups, patches, scaling) compared to a managed service like Cloud SQL. This could lead to higher operational overhead and potential disruptions during migration.

upvoted 1 times

  **Pime13** 1 year, 1 month ago

Selected Answer: C

Still C: <https://cloud.google.com/sql/docs/sqlserver/features>

even though SSRS and SSIS are available (running runs a separate host and connects to Cloud SQL) SSAS is not supported.

upvoted 2 times



  **theseawillclaim** 1 year, 9 months ago

Selected Answer: C

You'd need GCE, as Cloud SQL does not support SSAS, SRS and SSIS, as of September 26th, 2023.

<https://cloud.google.com/sql/docs/sqlserver/features>

upvoted 1 times

  **Pilot50** 2 years, 2 months ago

Selected Answer: C

I always wonder why the answers provided are wrong! so wrong and not even close to the reality! why examtopics.com are doing this?

upvoted 3 times

  **julioirevk** 1 year, 9 months ago

They really should just get rid of the answers and just have the discussion unless they plan on making the answers or at least more accurate.

upvoted 2 times

  **ArtistS** 1 year, 6 months ago

If they provide the answer as well, you will never see the exam topic any more

upvoted 2 times

🗨️ 👤 **absero1609** 2 years, 2 months ago

I think they do this on purpose.

They simply give a random answer on each question and let us fight in the comment section to find the truth answer. The good thing is that we have to work out ass to find the correct answer and we might learn something new along the way (or at least in my case) :D

upvoted 6 times

🗨️ 👤 **felipeschossler** 2 years, 2 months ago

I have to agree 100% with absero1609 kkkkkkkkkkkk, researching in the comments and reading the docs that actually let me learn fast for the exam.

upvoted 3 times

🗨️ 👤 **Nirca** 2 years, 3 months ago

Selected Answer: C

C. Migrate to Compute Engine is the minimal disruption method. customer will keep his SA logins and other system/dba/admin capabilities. and the SSAS SRS SSIS solutions will be migrated as well as "lift & shift" method.

upvoted 1 times

🗨️ 👤 **GCP72** 2 years, 6 months ago

Selected Answer: C

C is the correct answer

https://cloud.google.com/sql/docs/sqlserver/features#general_features_unavailable_for

upvoted 3 times

🗨️ 👤 **pk349** 2 years, 6 months ago

The following features are unavailable in Cloud SQL and not supported by Google Cloud support.

General features unavailable for Cloud SQL

- SQL Server Reporting Services (SSRS)
- SQL Server Analysis Services (SSAS)
- SQL Server Integration Services (SSIS)

upvoted 2 times

🗨️ 👤 **pk349** 2 years, 6 months ago

C: Migrate to Compute Engine

upvoted 1 times

🗨️ 👤 **chelbsik** 2 years, 6 months ago

Selected Answer: C

SSAS, SSRS, SSIS are unavailable for CloudSQL <https://cloud.google.com/sql/docs/sqlserver/features>

upvoted 1 times

🗨️ 👤 **range9005** 2 years, 6 months ago

I guess C

Migrate to Compute Engine

Because Cloud SQL doesn't support SSAS....


upvoted 1 times

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

Suggested Answer: C

Community vote distribution



🗳️ 👤 **887ad17** 5 months ago

Selected Answer: D

roles/cloudsql.instanceUser: This role allows the service account to connect to Cloud SQL instances. It provides the necessary permissions to access and read data from the Cloud SQL database without granting excessive permissions like editing or administrative capabilities.

roles/spanner.databaseUser: This role grants the service account permission to read and write data within Cloud Spanner databases. It allows the service account to update tables in Cloud Spanner, which is necessary for your use case.

upvoted 1 times

🗳️ 👤 **CosmoKramer** 5 months, 2 weeks ago

Selected Answer: D

Y'all should read the documentation. Seems like none of you actually know anything about IAM?

the cloudsql.viewer roles has:

```
cloudsql.*.export
cloudsql.*.get
cloudsql.*.list
cloudsql.instances.listServerCas
cloudsql.instances.listServerCertificates
```

in addition to a bunch of recommender permissions.

roles/cloudsql.instanceUser only has two permissions:

```
cloudsql.instances.get
cloudsql.instances.login
```

Some of you may think that this includes "Cloud SQL Admin", but that's wrong. As mentioned, read the docs more thoroughly, it says: "The following table lists each permission that Cloud SQL supports, the Cloud SQL roles that include it, and its basic role".

A role has a permission, not the other way around. So, cloudsql.instances.get is part of Cloud SQL Admin, Client, Editor etc.

upvoted 2 times

🗳️ 👤 **omermahgoub** 8 months, 1 week ago

Should be A. Because roles/cloudsql.instanceUser has the cloudsql.instances.get role, which has the following roles included:

```
Cloud SQL Admin
Cloud SQL Client
Cloud SQL Editor
Cloud SQL Viewer
```

This compromise "least privilege" requirements

upvoted 1 times

🗳️ 👤 **bigdawg70** 11 months ago

I think it should be A because the roles/cloudsql.instanceUser role only has:
cloudsql.instances.get
cloudsql.instances.login

You won't even be able to view anything with that role.

<https://cloud.google.com/sql/docs/mysql/iam-roles>
upvoted 1 times

🗳️ 👤 **RaphaelG** 11 months, 2 weeks ago

Selected Answer: D

To me, it is also "D", InstanceUser only has 2 permissions and Viewer has like 50 of them
upvoted 1 times

🗳️ 👤 **learnazureportal** 1 year, 3 months ago

The current answer is => D. roles/cloudsql.instanceUser and roles/spanner.databaseUser. roles/cloudsql.instanceUser: This role allows the service account to connect to Cloud SQL instances
upvoted 2 times

🗳️ 👤 **DeeData** 1 year, 3 months ago

I think you are missing out the "least privilege" part
upvoted 1 times

🗳️ 👤 **learnazureportal** 1 year, 3 months ago

No I am not. pay attention to "read data out of Cloud SQL".
upvoted 1 times

🗳️ 👤 **DPonly** 1 year, 3 months ago

Selected Answer: A

Will go by A
upvoted 3 times

🗳️ 👤 **Mithi_Kumar011** 1 year, 5 months ago

Correct Ans = 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.

To update a table in Cloud Spanner, you need to use a service account with the roles/spanner.databaseUser role on the Cloud Spanner instance. This role grants the service account permission to read and write data in the Spanner database.
upvoted 4 times

🗳️ 👤 **ziiMexx** 1 year, 7 months ago

Selected Answer: D

between A or D. But roles/cloudsql.viewer is too broad, so i choose D
upvoted 1 times

🗳️ 👤 **felipeschossler** 1 year, 8 months ago

Selected Answer: D

D. I think that instanceUser had the necessary permissions to read data: <https://cloud.google.com/sql/docs/sqlserver/iam-roles#roles~:text=roles/cloudsql.instanceUser>
upvoted 1 times

🗳️ 👤 **absero1609** 1 year, 8 months ago

Selected Answer: A

We want to apply least privilege and need to read data out of Cloud SQL for SQL Server only, `roles/cloudsql.viewer` is good enough to satisfy the those requirement, that filters out B, C, and D already
<https://cloud.google.com/sql/docs/sqlserver/iam-roles#roles>
upvoted 3 times

🗳️ 👤 **filotti** 1 year, 9 months ago

Selected Answer: D

To read data out of Cloud SQL for SQL Server, you need to use a service account with the roles/cloudsql.instanceUser role on the Cloud SQL instance. This role grants the service account permission to read data from the instance.

To update a table in Cloud Spanner, you need to use a service account with the roles/spanner.databaseUser role on the Cloud Spanner instance. This role grants the service account permission to read and write data in the Spanner database.

Therefore, to grant least privilege access, you should assign the service account only the required roles, which are roles/cloudsql.instanceUser and roles/spanner.databaseUser.

upvoted 2 times

🗲️ 👤 **dynamic_dba** 1 year, 9 months ago

A.

You need read access in Cloud SQL for SQL Server and read/write access in Cloud Spanner. Admin permissions are not required so eliminate B. roles/spanner.database Reader would not provide write access, so eliminate C. roles/cloudsql.viewer provides read only access to Cloud SQL resources. That eliminates D, leaving A.

upvoted 1 times

🗲️ 👤 **SidsA** 1 year, 10 months ago

Ans should be D for minimum access

roles/cloudsql.instanceUser : Allowing to login and get

roles/spanner.databaseUser: read and write

upvoted 1 times

🗲️ 👤 **GCP72** 2 years ago

Selected Answer: A

A is the correct answer

upvoted 3 times

🗲️ 👤 **pk349** 2 years ago

A: roles/cloudsql.viewer ***** and roles/spanner.databaseUser *****

upvoted 2 times

🗲️ 👤 **chelbsik** 2 years ago

Selected Answer: A

[https://cloud.google.com/spanner/docs/iam#:~:text=roles/spanner.databaseUser%20contains%20the%20permissions%20spanner.databases.read%20and%](https://cloud.google.com/spanner/docs/iam#:~:text=roles/spanner.databaseUser%20contains%20the%20permissions%20spanner.databases.read%20and%20)

<https://cloud.google.com/sql/docs/mysql/iam-roles#:~:text=roles/cloudsql.viewer,to%20all%20Cloud%20SQL%20resources.>

upvoted 4 times

You are responsible for designing a new database for an airline ticketing application in Google Cloud. This application must be able to:


- Work with transactions and offer strong consistency.
- Work with structured and semi-structured (JSON) data.
- Scale transparently to multiple regions globally as the operation grows.

You need a Google Cloud database that meets all the requirements of the application. What should you do?

- A. Use Cloud SQL for PostgreSQL with both cross-region read replicas.
- B. Use Cloud Spanner in a multi-region configuration.
- C. Use Firestore in Datastore mode.
- D. Use a Bigtable instance with clusters in multiple regions.

Suggested Answer: A

Community vote distribution



absero1609 Highly Voted 2 years, 2 months ago

Selected Answer: B

I'd choose B.

To me, there is no totally true option.

I agree with dynamic_dba that A and D are wrong.

B. Spanner supports JSON, but not that good (as Pilot50 pointed out)

C. Firestore in Datastore mode does provide some degree of consistency guarantees, specifically "eventual consistency" for global queries and strong consistency for ancestor queries. However, it does not provide strong consistency for non-ancestor queries.

In contrast, Cloud Spanner provides strong consistency for all transactions, which is especially important for applications that require ACID transactions across multiple regions.

So between B and C, I'll pick Spanner for strong consistency

upvoted 7 times

Alfred_Magno Most Recent 10 months ago

I choose option C. Spanner could work, but it would be very tight. Firestore is the best fit for the use case since it auto-scales transparently and works perfectly with semi-structured data

upvoted 1 times

Pime13 1 year, 2 months ago

Selected Answer: B

b: https://cloud.google.com/spanner/docs/working-with-json#spanner_query_with_json_parameter-java

upvoted 1 times

AngieSoccerBall49 1 year, 7 months ago

Selected Answer: B

A. Use Cloud SQL for PostgreSQL with both cross-region read replicas: This won't scale transparently or globally

B. Use Cloud Spanner in a multi-region configuration: supports JSON, designed for mission critical workloads, can be scaled transparently with an autoscaler across the globe.

C. Use Firestore in Datastore mode: Firestore can scale transparently in a multi-region; but it can't be scaled GLOBALLY.

D. Use a Bigtable instance with clusters in multiple regions: This doesn't necessarily fit the supplied workload.

upvoted 2 times

Pilot50 2 years, 2 months ago

Selected Answer: C

Spanner can't support JSON and not scaling automatically



upvoted 2 times

Pilot50 2 years, 2 months ago

Correction: Spanner supports JSON but not scaling automatically

<https://cloud.google.com/blog/products/databases/manage-semi-structured-data-in-cloud-spanner-with-json>

upvoted 2 times

  **Pilot50** 2 years, 2 months ago

Correction: the correct answer is B,

Firestore in Datastore mode does not provide strong consistency

upvoted 3 times

  **dynamic_dba** 2 years, 3 months ago

C.

Cloud SQL is a regional service so cannot scale to multiple regions. Eliminate A. Spanner doesn't support semi-structured data and doesn't scale transparently. Eliminate B. Bigtable doesn't work with JSON data. Eliminate D. That leaves C, Firestore in Datastore mode which satisfies all the requirements.

upvoted 2 times

  **bigdawg70** 1 year, 4 months ago

Spanner definitely supports semi-structured data:

<https://console.cloud.google.com/marketplace/product/google-cloud-platform/cloud-spanner>

Cloud Spanner is ideal for relational, structured, and semi-structured data that requires high availability, strong consistency, and transactional reads and writes. It offers all the traditional benefits of a relational database – such as ACID transactions and SQL semantics

upvoted 1 times

  **H_S** 2 years, 3 months ago

Selected Answer: C

Hello Team,

I hope this comment helps

Work with transactions and offer strong consistency

Work with structured and semi-structured (JSON) data => NoSQL (cloud SQL + spanner are not suitable here)

Scale transparently to multiple regions globally as the operation grows => scale transparently => you don't manage anything

Big Table is for time series, millisecond access and for JSON

Reference ; <https://cloud.google.com/firestore#all-features>

upvoted 3 times

  **H_S** 2 years, 3 months ago


more reference ; https://cloud.google.com/firestore/docs/firestore-or-datastore#in_datastore_mode

upvoted 2 times

  **Ayush9596** 2 years, 4 months ago

Why not bigtable ?

upvoted 1 times

  **orbo** 2 years, 5 months ago

Selected Answer: B

Spanner is multiple region and can scale transparently

upvoted 1 times



  **H_S** 2 years, 3 months ago

no you scale spanner, it's not transparently,

Firestore => With automatic multi-region replication and strong consistency, your data is safe and has a 99.999% availability guarantee, even when disasters strike.

<https://cloud.google.com/firestore#all-features>

upvoted 1 times

  **GCP72** 2 years, 6 months ago

Selected Answer: B

A and B are correct answer but B looks best answer because of "Scale transparently to multiple regions globally as the operation grows."
upvoted 3 times

🗨️ 👤 **H_S** 2 years, 3 months ago

what about JSON

upvoted 1 times

🗨️ 👤 **pk349** 2 years, 6 months ago

A: Use Cloud SQL for PostgreSQL ***** with both cross-region read replicas.

upvoted 1 times

🗨️ 👤 **gabrielosluz** 2 years, 4 months ago

Why Cloud SQL? You comment the wrong answer in multiple question around here

upvoted 1 times

🗨️ 👤 **chelbsik** 2 years, 6 months ago

Selected Answer: B

I'll go for Spanner because of 'Scale transparently to multiple regions globally as the operation grows'

upvoted 4 times

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.

Suggested Answer: C

Community vote distribution

D (65%)

C (35%)

 **chelbsik**  2 years, 6 months ago


Selected Answer: D

It's D, CloudSQL Auth proxy is not used when connecting to Private IP

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

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

upvoted 11 times

 **gabrielosluz** 2 years, 4 months ago

The Cloud SQL Auth proxy works with both public and private IP endpoints: <https://cloud.google.com/sql/docs/mysql/connect-auth-proxy>
upvoted 5 times

 **dynamic_dba**  2 years, 3 months ago

D.

Cloud Run to Cloud SQL connectivity can be done using private IPs. Eliminate A and B. C would be right except you wouldn't use Cloud SQL Auth Proxy because Serverless VPC Access would connect directly to the Cloud SQL instance. The connection pool reference in D puts you off, but it is the right answer. The link provided by SVGGoogle89 is spot on.

upvoted 5 times

 **krop**  1 month, 3 weeks ago

Selected Answer: C

Correct answer is C.


upvoted 1 times

 **rglearn** 8 months, 1 week ago

Selected Answer: D

cloud Run has built in SQL auth Proxy hence we dont need to it use it explicitly.


upvoted 1 times

 **Pime13** 1 year, 2 months ago

Selected Answer: D

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

upvoted 2 times

 **PKookNN** 1 year, 5 months ago

Selected Answer: D

https://cloud.google.com/sql/docs/mysql/connect-run#connect_to - said clearly that there is no need for Cloud SQL Auth Proxy when using with Cloud Run.

upvoted 2 times

🗳️ 👤 **LaxmanTiwari** 1 year, 5 months ago

API Quota Limits

Cloud Run provides a mechanism that connects using the Cloud SQL Auth Proxy, which uses the Cloud SQL Admin API. API quota limits apply to the Cloud SQL Auth Proxy. The Cloud SQL Admin API quota used is approximately two times the number of Cloud SQL instances configured by the number of Cloud Run instances of a particular service deployed at any one time. You can cap or increase the number of Cloud Run instances to modify the expected API quota consumed.

upvoted 1 times

🗳️ 👤 **LaxmanTiwari** 1 year, 5 months ago

Selected Answer: C

upvoted 2 times

🗳️ 👤 **examprof** 1 year, 6 months ago

Option D.

This link explicitly indicates that "For private IP paths, your application will connect directly to your instance through Serverless VPC Access. This method uses TCP to connect directly to the Cloud SQL instance without using the Cloud SQL Auth Proxy."

https://cloud.google.com/sql/docs/mysql/connect-run#connect_to

upvoted 3 times

🗳️ 👤 **AngieSoccerBall49** 1 year, 7 months ago

D is the correct answer.

A common misconception.

upvoted 1 times

🗳️ 👤 **AngieSoccerBall49** 1 year, 7 months ago

When using Cloud Run to connect to Cloud SQL Private IP addresses, it is unnecessary to use the SQL Auth Proxy in Private IP mode.

The Serverless VPC Access Connector (which has been superseded by Cloud Run's direct VPC Egress) should connect directly and leverage a connection pooler (potentially in your application via client library, probably better as a separate instance) for more consistent connections to the Cloud SQL DB.

You would create a specific "user" for this purpose in your database.

upvoted 1 times

🗳️ 👤 **nhiguchi** 1 year, 7 months ago

Selected Answer: C

C is correct.

upvoted 1 times

🗳️ 👤 **julioevk** 1 year, 8 months ago

Selected Answer: C

<https://cloud.google.com/sql/docs/mysql/connect-auth-proxy>

"Works with both public and private IP endpoints"

upvoted 1 times

🗳️ 👤 **ewelaz** 1 year, 9 months ago

Selected Answer: D

it's d, Auth proxy is not used when connecting to Private IP

upvoted 2 times

🗳️ 👤 **nqthien041292** 1 year, 9 months ago

Selected Answer: C

Vote C

upvoted 1 times

🗳️ 👤 **learnazureportal** 1 year, 9 months ago

correct answer is C - The Cloud SQL Auth proxy acts as a secure intermediary between your Cloud Run application and the Cloud SQL instance, allowing for secure and authenticated database connections while keeping the database inaccessible from the public internet.

upvoted 1 times

  **CloudKida** 2 years ago

Selected Answer: C

<https://cloud.google.com/sql/docs/mysql/connect-overview>

Configuring your instance with a private IP is preferred when connecting from a client on a resource with access to a VPC. For more information about what resources can use private IP, see Requirements for Private IP.

For private IP paths, the following services and applications connect directly to your instance through Serverless VPC Access:

App Engine standard environment

App Engine flexible environment

Cloud Functions

Cloud Run

upvoted 2 times

  **Sandipcst** 2 years ago

C

Cloud SQL Auth Proxy can connect Cloud SQL instance with private ip by specifying --private-ip argument in same VPC. Cloud Run can run a container that gets the auth proxy installable files and run the auth proxy in cloud Run in same VPC.

upvoted 1 times

  **abdenago** 2 years ago

Selected Answer: D

auth proxy isn't required with private serverless access, the connection pool increases reliability of the connection


upvoted 2 times

You are troubleshooting a connection issue with a newly deployed Cloud SQL instance on Google Cloud. While investigating the Cloud SQL Proxy logs, you see the message Error 403: Access Not Configured. What should you do?

- A. Check the app.yaml value cloud_sql_instances for a misspelled or incorrect instance connection name.
- B. Check whether your service account has cloudsql.instances.connect permission.
- C. Enable the Cloud SQL Admin API.
- D. Ensure that you are using an external (public) IP address interface.

Suggested Answer: A

Community vote distribution



🗳️ 👤 **Pime13** 8 months ago

Selected Answer: C

<https://cloud.google.com/sql/docs/mysql/connect-auth-proxy#troubleshooting>
upvoted 3 times

🗳️ 👤 **AngieSoccerBall49** 1 year, 1 month ago

Selected Answer: C

The answer is C.

If you're missing cloudsql.instances.connect on the attached service account, the error is 403 FORBIDDEN.

If you haven't enabled the Cloud SQL Admin API, the error is 403 ACCESS NOT CONFIGURED.

upvoted 4 times

🗳️ 👤 **juliorevk** 1 year, 2 months ago

Selected Answer: C

<https://cloud.google.com/sql/docs/mysql/connect-auth-proxy#troubleshooting>

C because in docs it says "Make sure to enable the Cloud SQL Admin API. If it is not, you see output like Error 403: Access Not Configured in your Cloud SQL Auth Proxy logs."

B is not correct because this is for a different error: "If you are getting a 403 notAuthorized error, and you are using a service account to authenticate the Cloud SQL Auth Proxy, make sure the service account has the correct permissions."

upvoted 2 times

🗳️ 👤 **ewelaz** 1 year, 3 months ago

Selected Answer: B

it's b

upvoted 1 times

🗳️ 👤 **learnazureportal** 1 year, 3 months ago

The "Error 403: Access Not Configured" message typically indicates a permission issue related to the Cloud SQL instance. The correct answer is B
B. Check whether your service account has cloudsql.instances.connect permission.

upvoted 1 times

🗳️ 👤 **DPonly** 1 year, 5 months ago

Selected Answer: C

<https://cloud.google.com/sql/docs/mysql/connect-auth-proxy#troubleshooting>

upvoted 1 times

🗳️ 👤 **Pilot50** 1 year, 8 months ago

Selected Answer: C

<https://cloud.google.com/sql/docs/mysql/connect-auth-proxy#troubleshooting?text=Make%20sure%20to,Auth%20proxy%20logs>.

upvoted 1 times

🗳️ 👤 **dynamic_dba** 1 year, 9 months ago

C.

D is not a recommended way to access a Cloud SQL instance and you wouldn't get a 403 error connecting via a public IP anyway. A and B could trigger a 403 error, but the text would be "not authorized". The answer is in Google's documentation. The link provided by sp57 is spot on.

upvoted 2 times

SVGoogle89 2 years ago

typo.. its C

Make sure to enable the Cloud SQL Admin API.

If it is not, you see output like Error 403: Access Not Configured in your Cloud SQL Auth proxy logs.

upvoted 1 times

SVGoogle89 2 years ago

B

403 notAuthorized error, and you are using a service account to authenticate the Cloud SQL Auth proxy, make sure the service account has the correct permissions.

upvoted 1 times

sp57 2 years ago

wrong - error is 403 Access not Configured - see 4 bullet pts down at link <https://cloud.google.com/sql/docs/mysql/connect-admin-proxy#troubleshooting>

C is correct

upvoted 1 times

GCP72 2 years ago

Selected Answer: C

Correct answer is c

<https://groups.google.com/g/google-cloud-sql-discuss/c/yoGQzTRXa0k?pli=1>

upvoted 1 times

pk349 2 years ago

C: Enable the Cloud SQL Admin API.

upvoted 1 times

chelbsik 2 years ago

Selected Answer: C

Make sure to enable the Cloud SQL Admin API.

If it is not, you see output like Error 403: Access Not Configured in your Cloud SQL Auth proxy logs.

upvoted 1 times

jitu028 2 years ago

Selected Answer: B

Correct Answer - B

<https://cloud.google.com/sql/docs/mysql/connect-admin-proxy#troubleshooting:~:text=If%20you%20are%20getting,have%20this%20permission.>

upvoted 1 times

range9005 2 years ago

Selected Answer: C

Enable the Cloud SQL Admin API

upvoted 1 times

juancambb 2 years ago

Selected Answer: C

Answer C, <https://cloud.google.com/sql/docs/mysql/connect-admin-proxy#troubleshooting>


upvoted 2 times

You are working on a new centralized inventory management system to track items available in 200 stores, which each have 500 GB of data. You are planning a gradual rollout of the system to a few stores each week. You need to design an SQL database architecture that minimizes costs and user disruption during each regional rollout and can scale up or down on nights and holidays. What should you do?

- A. Use Oracle Real Application Cluster (RAC) databases on Bare Metal Solution for Oracle.
- B. Use sharded Cloud SQL instances with one or more stores per database instance.
- C. Use a Bigtable cluster with autoscaling.
- D. Use Cloud Spanner with a custom autoscaling solution.

Suggested Answer: B

Community vote distribution



dynamic_dba Highly Voted 2 years, 3 months ago

D.

A SQL database architecture rules out Bigtable. Minimizing costs rules out Oracle RAC. That leaves B and D. B would work with each Cloud SQL instance being dedicated to a few stores which would not impact other Cloud SQL instances already running. The downside is the scaling up/down. To change the vCPU on a Cloud SQL instance requires it to be re-started and that's disruption. Spanner also doesn't autoscale by itself, but there's a tool available for Spanner called Autoscaler which can automate scaling up/down. So on balance, D is the best answer.

<https://cloud.google.com/spanner/docs/autoscaling-overview>

upvoted 12 times

Jason_Cloud_at 1 year, 4 months ago

i agree with you on ruling out A and C, however cloud spanner is much costlier than cloud SQL , so B wouldnt be the best answer ?

upvoted 1 times

orbo Most Recent 5 months, 3 weeks ago

Selected Answer: D

D- and Why D?

Max Instance size in Cloud SQL is 64TB, 200 Store each 500GB DB Size and growing it is way over 65 TB

upvoted 1 times

rglearn 8 months, 1 week ago

Selected Answer: D

CloudSQL auto scale up/down causes disruption.

in case of cloudspanner it is automatic. Plus Global solution by spanner is big plus.

but again technical downside of multi regional spanner is, only one region will be a leader, because of which app from non-leader region may observe bit higher write latency.

upvoted 1 times

RaphaelG 1 year, 5 months ago

Selected Answer: B

To me, it is "B" here. A couple of reasons; the statement "minimizes costs and user disruption during each regional rollout" explicitly emphasises that the disruption relates to the rollout not the scale-up or scale-down, and considering it's only a few stores per db, I would presume it would take like a few minutes tops. Lastly, the cost is a bit of a give away as well since Cloud SQL is like twice as cheap (I did some rough estimates recently using europe-west2 as my benchmark).

upvoted 2 times

learnazureportal 1 year, 9 months ago

The correct answer is ==> B. Use sharded Cloud SQL instances with one or more stores per database instance.

upvoted 1 times

Mitra123 1 year, 11 months ago

B

Guys, either B or D. The keyword is :Minimize cost" Although D is the best solution, B is less costly

upvoted 2 times

🗨️ 👤 **Nirca** 2 years, 3 months ago

Selected Answer: D

1. CloudSQL max out at 64TB, so unable to hold 100TB of data. https://cloud.google.com/sql/docs/quotas#metrics_collection_limit

2. Scale is done manually on SQL Cloud.

upvoted 3 times

🗨️ 👤 **zanhseh** 2 years, 4 months ago

Selected Answer: D

CloudSQL max out at 64TB, so unable to hold 200 * 500 GB of data.

A: No. Oracle RAC cannot scale up or down

B: No. Cloud SQL cannot scale up or down manually and the sharded Cloud SQL sound weird, and doesn't meet the "minimize costs and user disruption during each REGIONAL rollout". Also can't break storage limits 64TB.

C: No. BigTable handles rational db poorly

<https://cloud.google.com/sql/docs/quotas#:~:text=Cloud%20SQL%20storage%20limits,core%3A%20Up%20to%203%20TB>.

upvoted 2 times

🗨️ 👤 **JayGeotab** 2 years, 5 months ago

Selected Answer: D

CloudSQL max out at 64TB, so unable to hold 200 * 500 GB of data.

D: choose Spanner

upvoted 2 times

🗨️ 👤 **wolfie09** 2 years ago

it's written 500 GB of data per store and cloud sql instance per 1 or couple of stores, what's so hard to understand?

upvoted 1 times

🗨️ 👤 **marpayer** 2 years, 5 months ago

A - No- oracle RAC cannot scale up or down

C- No, Big Table is for nonSQL

B or D

B - Cloud SQL cannot scale up or down manually and the sharded Cloud SQL sound weird, and doesn't meet the "minimize costs and user disruption during each REGIONAL rollout"

For my, D is the best option as it

upvoted 1 times

🗨️ 👤 **chelbsik** 2 years, 6 months ago

Selected Answer: B

Cloud SQL sharding looks like a good option since we need to minimize costs

and we don't need global scaling <https://cloud.google.com/community/tutorials/horizontally-scale-mysql-database-backend-with-google-cloud-sql-and-proxysql>

upvoted 4 times

🗨️ 👤 **pk349** 2 years, 6 months ago

B: Use sharded ***** Cloud SQL instances with one or more stores per database instance.

Sharding makes horizontal scaling possible by partitioning the database into smaller, more manageable parts (shards), then deploying the parts across a cluster of machines. Data queries are routed to the corresponding server automatically, usually with rules embedded in application logic or a query router.


upvoted 1 times

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

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

Suggested Answer: B

Community vote distribution



dynamic_dba Highly Voted 1 year, 9 months ago

C.

To satisfy audit reporting you would need a way to record what was changed and when. The best answer is one which uses some kind of source code control system (SCCS). That rules out A and B. Any mention of anything manual in a cloud environment should look suspicious, which leave option C.

As it happens, Liquibase is an SCCS and can be integrated with Spanner.

<https://cloud.google.com/spanner/docs/use-liquibase>

upvoted 6 times

Pilot50 Highly Voted 1 year, 8 months ago

Selected Answer: C

<https://cloud.google.com/spanner/docs/use-liquibase>

upvoted 5 times

dija123 Most Recent 7 months, 2 weeks ago

Selected Answer: C

Use Liquibase to roll out changes to higher environments.

upvoted 1 times

learnazureportal 1 year, 3 months ago

The correct answer is => D. Manually capturing audit logs provides the level of granularity and control required for compliance auditing.

upvoted 1 times

Mitra123 1 year, 5 months ago

D.

Auditors wants to see what changed. The only way to showcase that is to capture the audit logs, who changed what... Hence D

upvoted 2 times

Nirca 1 year, 9 months ago

Selected Answer: C

C, liquidbase for db source control

upvoted 2 times

zanhsieh 1 year, 10 months ago

Selected Answer: C

I assume production environment shall be a standalone project. Based on this,

A: No. Replication can't be done cross projects.

B: No. Backup / restore only get a whole set drop / load statements from audit logs. External auditors would focus on "what changes".

C: Yes. Liguibase could export change.yaml to hand over the external auditors.

D: No. No way to capture audit logs and ship to the target project. Spanner is a managed service.

Reference:

<https://cloud.google.com/spanner/docs/backup/gcloud>

<https://cloud.google.com/spanner/docs/use-liquibase>

upvoted 4 times

chelbsik 2 years ago

Selected Answer: B

https://cloud.google.com/spanner/docs/audit-logging#audited_operations

upvoted 3 times

🗨️ 👤 **SVGoogle89** 2 years ago

C, liquidbase for db source control

upvoted 3 times

🗨️ 👤 **pk349** 2 years ago

B: Use backup and restore ***** to roll out changes to higher environments.

[1] Even though restoring a database requires authorization on two resources (the backup ***** and restored database, which might reside in different instances), the RestoreDatabase event is logged only once as a single entry in the instance of the restored database. Within this entry, there will be two authorizationInfo entries: one for the database, checking the spanner.databases.create permission, and one for the backup, checking the spanner.backups.restoreDatabase permission.

upvoted 1 times

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.

Suggested Answer: A

Community vote distribution

A (100%)

🗳️ 👤 **887ad17** 5 months ago

Selected Answer: A

`gcloud sql instances patch [INSTANCE_NAME] --cpu=[NUMBER_OF_VCPUS]`
upvoted 1 times

🗳️ 👤 **Pime13** 8 months ago

Selected Answer: A

<https://cloud.google.com/sdk/gcloud/reference/sql/instances/patch>
upvoted 1 times

🗳️ 👤 **julio revk** 1 year, 2 months ago

Selected Answer: A

<https://cloud.google.com/sdk/gcloud/reference/sql/instances/patch>
upvoted 1 times

🗳️ 👤 **dynamic_dba** 1 year, 9 months ago

A.
D would not represent minimal interruption or effort. B would not change the number of vCPUs available to the instance as a whole. Option C is for GCE VMs, not Cloud SQL.
upvoted 4 times

🗳️ 👤 **Nirca** 1 year, 9 months ago

Selected Answer: A

"`gcloud sql instances patch`" -> updates the settings of a Cloud SQL instance
upvoted 1 times

🗳️ 👤 **chelbsik** 2 years ago

Selected Answer: A

<https://cloud.google.com/sdk/gcloud/reference/sql/instances/patch#:~:text=existing%20authorized%20networks.-%2D%2Dcpu,-%3DCPU>
upvoted 2 times

🗳️ 👤 **pk349** 2 years ago

A: Issue a `gcloud sql instances *** patch` command to increase the number of vCPUs.
Cores
Y 1 to 96 (must be either 1 or an even number)
Partial for shared vCPU
`gcloud sql instances patch INSTANCE [--activation-policy=ACTIVATION_POLICY] [--active-directory-domain=ACTIVE_DIRECTORY_DOMAIN] [--no-assign-ip] [--async] [--audit-bucket-path=AUDIT_BUCKET_PATH] [--audit-retention-interval=AUDIT_RETENTION_INTERVAL] [--audit-upload-interval=AUDIT_UPLOAD_INTERVAL] [--availability-type=AVAILABILITY_TYPE] [--clear-password-policy] [--connector-enforcement=CONNECTOR_ENFORCEMENT] [--cpu=CPU] [--database-version`
upvoted 2 times

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

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

Suggested Answer: B

Community vote distribution

A (100%)

🗳️ 👤 **dynamic_dba** Highly Voted 👍 1 year, 3 months ago

A.

Cloud SQL is a regional service so the primary and standby instances must be in the same region. A recommended practice would have them in different zones. A read replica could be in a different region, but that option isn't offered probably because it would be too far away from the user base. Having it in a different zone from the primary and standby instances is therefore the best and only answer.

upvoted 9 times

🗳️ 👤 **887ad17** Most Recent 🕒 5 months ago

Selected Answer: A

A. Configure the primary instance in zone A, the standby instance in zone C, and the read replica in zone B, all in the same region.

This option follows Google-recommended practices for high availability by distributing the instances across multiple zones within the same region. This setup ensures that:

The primary instance and standby instance are in different zones, providing resilience against zonal failures.

The read replica is also in a different zone, allowing for load balancing of read operations and further protection against zone-specific outages.

upvoted 1 times

🗳️ 👤 **learnazureportal** 9 months, 3 weeks ago

The correct answer is ==> B. Configure the primary and standby instances in zone A and the read replica in zone B, all in the same region.

upvoted 1 times

🗳️ 👤 **chelbsik** 1 year, 6 months ago

Selected Answer: A

HA standby instance can't be located in different region, and Google recommends to use different zones for all three, so answer is A

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

upvoted 3 times

🗳️ 👤 **pk349** 1 year, 6 months ago

The HA configuration provides data redundancy. A Cloud SQL instance configured for HA is also called a regional instance and has a primary and secondary zone within the configured region. Within a regional instance, the configuration is made up of a primary instance and a standby instance. Through synchronous replication to each zone's persistent disk, all writes made to the primary instance are replicated to disks in both zones before a transaction is reported as committed. In the event of an instance or zone failure, the standby instance becomes the new primary instance. Users are then rerouted to the new primary instance. This process is called a failover.

upvoted 3 times

🗳️ 👤 **pk349** 1 year, 6 months ago

A: Configure the primary instance in zone A, the standby instance in zone C, and the read replica in zone B, all in the same region.

upvoted 3 times