Assignment and Exam Content

Cloud Spanner

Always Delete your Cloud Resources to Avoid $$ Charges.

Cloud Spanner Lab

Cloud Spanner Lab Contains – Three major areas below to say complete Lab ☺

A Launch Cloud Spanner Instance

Understand basic concepts – Locations, Performance Read/ Write IOPS etc

B Exam Tips

Always Delete your Cloud Resources to Avoid $$ Charges.

A

Create Cloud SQL Instance

1 Go To -> STORAGE -> Spanner -> Click on Create Instance

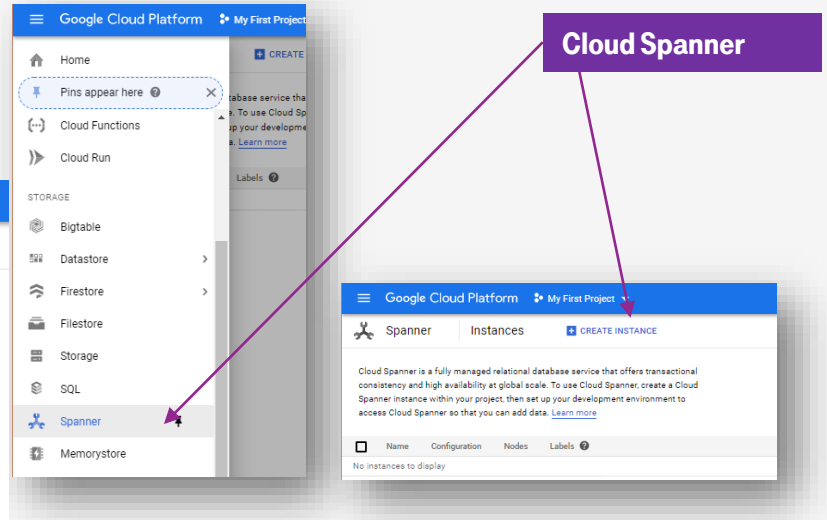

You will need to provide different parameter

2

1. Name : Obviously your instance name
2. Instance ID : if you want to change
3. Choose Config-> Regional vs Multi regional

- Select Regional for this demo

4. Select region of your choice (if its multi regional choose area – Watch for configuration details on right side

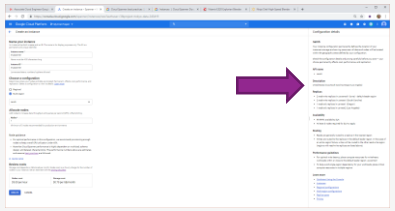

5. Select 1 node.

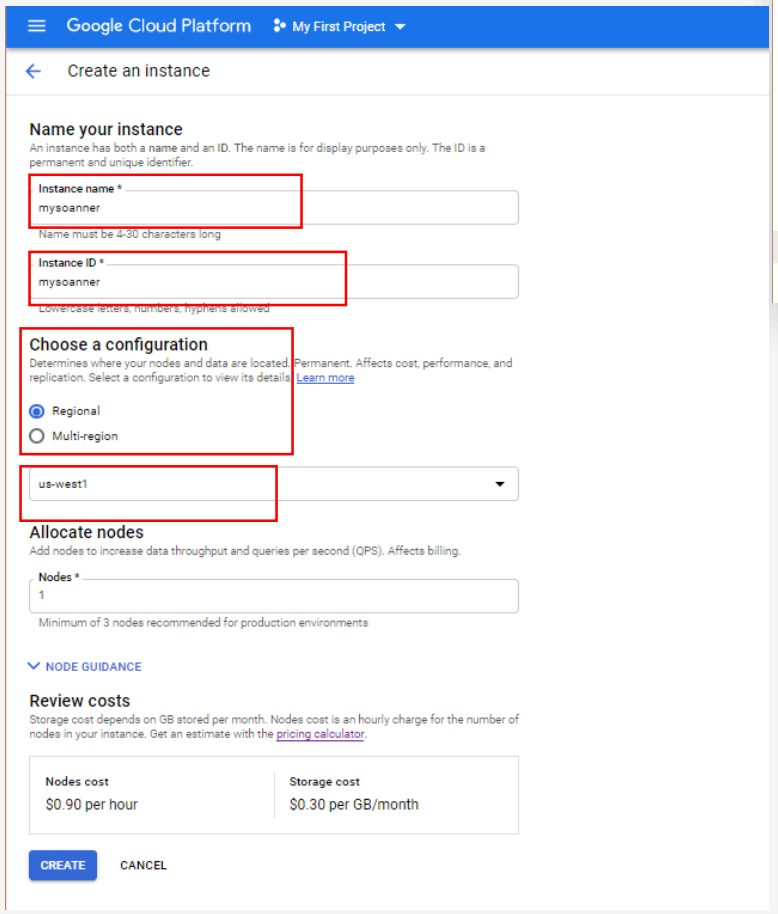

*Cloud Spanner is a fully managed relational database service that offers **transactional consistency and high availability at global scale.***

To use Cloud Spanner, create a Cloud Spanner instance within your project, then set up your development environment to access Cloud Spanner so that you can add data. Learn more

1

Create Cloud SQL Instance

3 Review Cost – (important for exam)

Cloud Spanner cost is dependent on cost per node + Storage cost per GB.

While Choosing Nodes – Please check Node Guidance for

Node Guidance

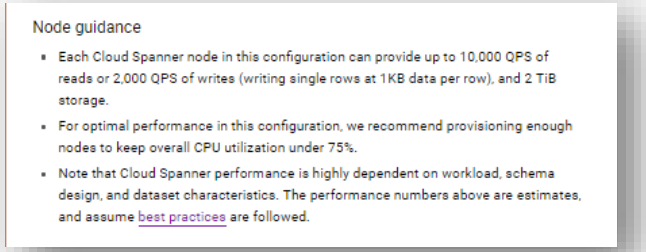

Node guidance

- Each Cloud Spanner node in this configuration can provide up to 10,000 QPS of reads or 2,000 QPS of writes (writing single rows at 1KB data per row), and 2 TiB storage.
- For optimal performance in this configuration, we recommend provisioning enough nodes to keep overall CPU utilization under 75%.
- Note that Cloud Spanner performance is highly dependent on workload, schema design, and dataset characteristics. The performance numbers above are estimates, and assume best practices are followed.

Cloud Spanner Performance

- Performance is dependent on number of Nodes
- Cost also dependent on nodes.
- Please go through Best Practices docs in following link.

https://cloud.google.com/spanner/docs/best-practice-list

Configuration Details

On right side of panel gives you details of configuration you selected.

Go through details and differences when you choose Regional vs multi regional.

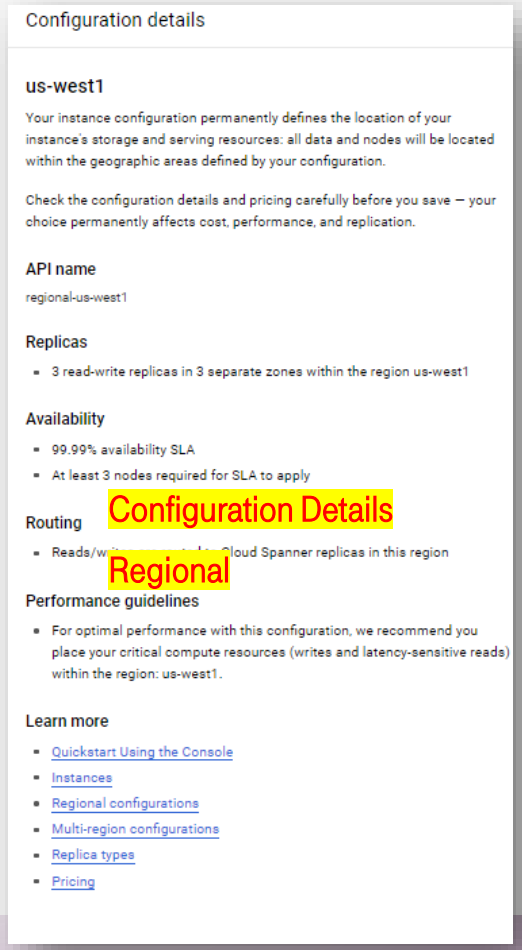

Configuration details

us-west1

Your instance configuration permanently defines the location of your instance's storage and serving resources: all data and nodes will be located within the geographic areas defined by your configuration.

Check the configuration details and pricing carefully before you save — your choice permanently affects cost, performance, and replication.

API name

regional-us-west1

Replicas

- 3 read-write replicas in 3 separate zones within the region us-west1

Availability

- 99.99% availability SLA
- At least 3 nodes required for SLA to apply

Routing Configuration Details

- Reads/w… Cloud Spanner replicas in this region

Regional

Performance guidelines

- For optimal performance with this configuration, we recommend you place your critical compute resources (writes and latency-sensitive reads) within the region: us-west1.

Learn more

- Quickstart Using the Console
- Instances
- Regional configurations
- Multi-region configurations
- Replica types
- Pricing

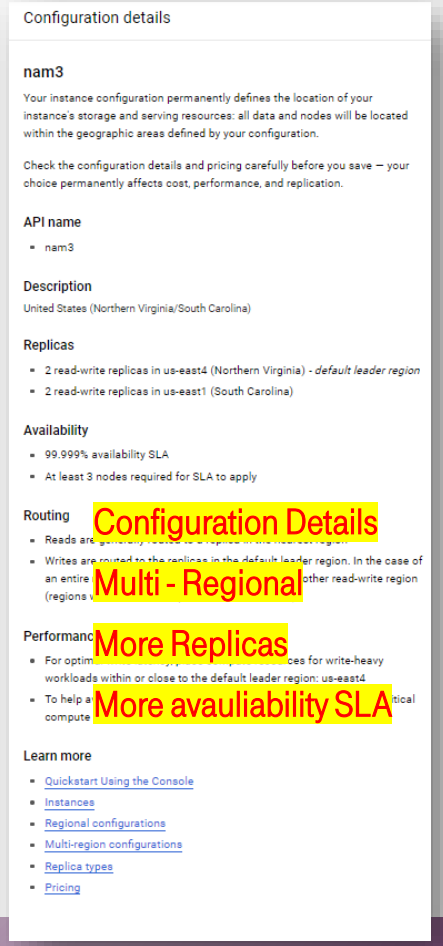

Configuration details

nam3

Your instance configuration permanently defines the location of your instance's storage and serving resources: all data and nodes will be located within the geographic areas defined by your configuration.

Check the configuration details and pricing carefully before you save — your choice permanently affects cost, performance, and replication.

API name

- nam3

Description

United States (Northern Virginia/South Carolina)

Replicas

- 2 read-write replicas in us-east4 (Northern Virginia) - *default leader region*
- 2 read-write replicas in us-east1 (South Carolina)

Availability

- 99.999% availability SLA
- At least 3 nodes required for SLA to apply

Routing Configuration Details

- Reads are …

Multi - Regional

- Writes are … other read-write region (regions …

Performanc More Replicas

- For optim… …ces for write-heavy workloads within or close to the default leader region: us-east4
- To help a… …tical compute

More avauliability SLA

Learn more

- Quickstart Using the Console
- Instances
- Regional configurations
- Multi-region configurations
- Replica types
- Pricing

1

Create Cloud Spanner

4

Click Create at the bottom to create instance

- Your instance is now created, and you can database , add tables , Interleaved tables etc

You can modify instance – using EDIT INSTANCE

You can import and export data using link given on spanner page.

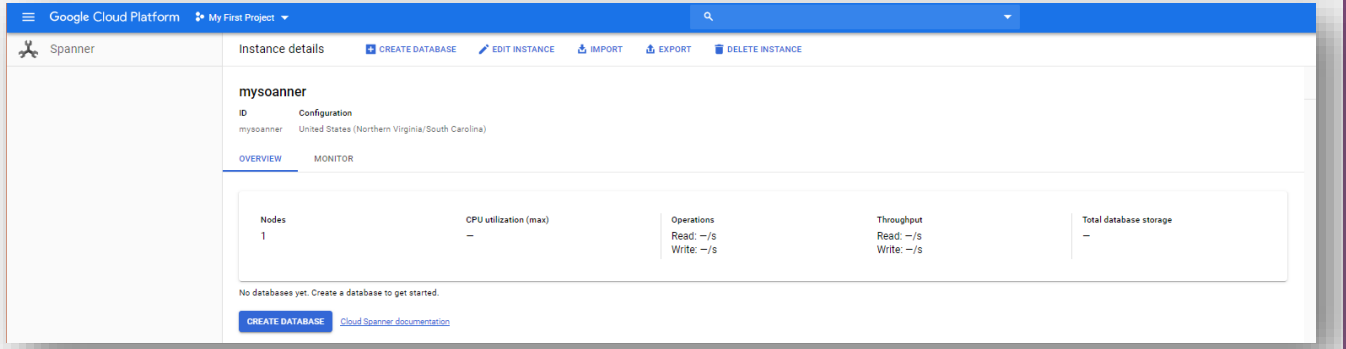

To create Tables and add data – please follow google docs link below

https://cloud.google.com/spanner/docs/quickstart-console

You can manage permissions in right side of panel at Spanner home page.

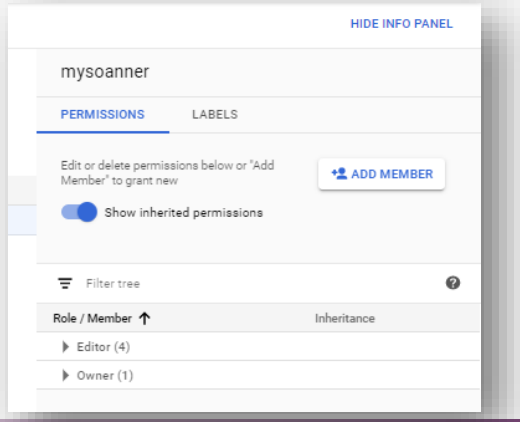

2

Create Cloud Spanner Instance

1

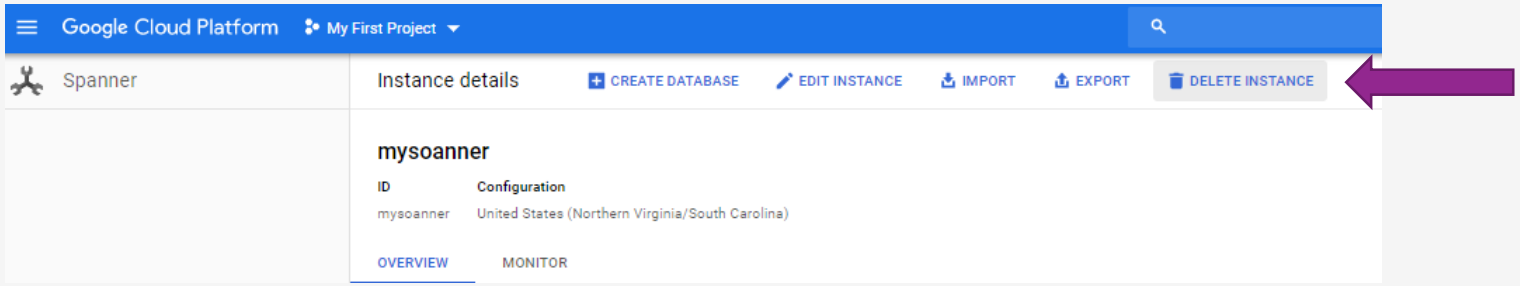

Before Creating next Instance

- Delete Old Instance and Proceed further

Always Delete your Cloud Resources to Avoid $$ Charges.

Cloud Spanner : Try Yourself

1. Edit Instance and see what you can edit and what you can not.

2. **Exam Tips**

Important concepts are

1. **Global Relational database with high consistency**
2. **Performance is dependent on number of nodes**
3. **Cost is dependent on number of nodes.**
4. **Interleaved tables exists in Spanner.**
5. **If you need Global scale database – You can use spanner.**

End of Cloud Spanner Assignment