

This page explains how to monitor a Cloud Spanner instance by viewing current and historical data in the Google Cloud Console.

You can also [use Stackdriver Monitoring to monitor your instance](/spanner/docs/monitoring-stackdriver) (/spanner/docs/monitoring-stackdriver). Use Stackdriver Monitoring to create customized charts and dashboards, and to set alerts for metrics that exceed your specified thresholds.

When you open a Cloud Spanner instance in the Cloud Console, you can see the current status of the most essential metrics for your instance, such as CPU utilization. You can also view the current metrics for a single database within the instance.

To view an instance's current status, follow these steps:

1. In the Cloud Console, open the list of Cloud Spanner instances.

[Go to the instance list](https://console.cloud.google.com/spanner/instances) (https://console.cloud.google.com/spanner/instances)

2. Click the name of the instance you want to learn about.

The Cloud Console displays an overview of the instance, including the current values of several key metrics. [Learn what metrics are available](#) (#charts-metrics).

3. Optional: To view the current metrics for a specific database in the instance, click the name of the database in the left pane.

The Cloud Console displays the current values of several key metrics.

The Cloud Console provides charts that show historical data for your instance. You can view charts for several different time periods, ranging from the past 1 hour to the past 30 days. You can also view charts for a single database within the instance.

To view historical data for an instance, follow these steps:

1. In the Cloud Console, open the list of Cloud Spanner instances.

[Go to the instance list \(https://console.cloud.google.com/spanner/instances\)](https://console.cloud.google.com/spanner/instances)

2. Click the name of the instance you want to learn about, then click the **Monitor** tab.

The Cloud Console displays charts of historical data for the instance. [Learn what charts are available \(#charts-metrics\)](#).

3. Optional: To view historical data for a different time period, find the buttons above the first chart, then click the time period you want to view.
4. Optional: To control what data appears in the chart, click one of the drop-down lists in the chart. For example, if the instance uses a [multi-region configuration \(/spanner/docs/instances#configs-multi-region\)](#), some charts provide a drop-down list to view data for a specific region. Not all charts have drop-down lists.
5. Optional: To view historical data for a specific database in the instance, click the name of the database in the left pane.

The Cloud Console provides the following charts and metrics to show an instance's current and historical status. All charts and metrics are available at the instance level. You can also view many charts and metrics for a single database within an instance.

#### Available charts and metrics

<b>Node count</b>	The number of nodes in the instance.  Available only for instances.
<b>CPU utilization - rolling average 24 hour</b>	A rolling average of total <a href="#">CPU utilization (/spanner/docs/cpu-utilization)</a> , as a percentage of the instance's CPU resources, for each database. Each data point is an average for the previous 24 hours.  Available only for instances.

Available charts and metrics	
<b>CPU utilization - high priority</b>	<p>The total <a href="/spanner/docs/cpu-utilization">CPU utilization</a> (/spanner/docs/cpu-utilization), as a percentage of the instance's CPU resources, for high-priority tasks. These tasks include most requests that you initiate, as well as maintenance tasks that Cloud Spanner must complete promptly. <a href="/spanner/docs/cpu-utilization#task-priority">Learn more about high-priority tasks</a> (/spanner/docs/cpu-utilization#task-priority).</p> <p>Available for instances and databases.</p>
<b>CPU utilization - total</b>	<p>The total <a href="/spanner/docs/cpu-utilization">CPU utilization</a> (/spanner/docs/cpu-utilization), as a percentage of the instance's CPU resources.</p> <p>Available for instances and databases:</p> <ul style="list-style-type: none"> <li>• For <i>instances</i>, you can view charts of total CPU utilization by database, or by <a href="/spanner/docs/cpu-utilization#task-priority">task priority</a> (/spanner/docs/cpu-utilization#task-priority).</li> <li>• For <i>databases</i>, you can view charts of total CPU utilization by task priority.</li> </ul>
<b>Latency</b>	<p>The amount of time, in seconds, that Cloud Spanner took to handle a request. This measurement begins when the Cloud Spanner service receives a request, and it ends when the Cloud Spanner service starts to send a response. <a href="/spanner/docs/latency">Learn more about latency metrics</a> (/spanner/docs/latency).</p> <p>You can view latency metrics for 50th and 99th percentile latency:</p> <ul style="list-style-type: none"> <li>• <b>50th percentile latency:</b> The maximum latency, in seconds, for the fastest 50% of all requests</li> <li>• <b>99th percentile latency:</b> The maximum latency, in seconds, for the fastest 99% of all requests</li> </ul> <p>Available for instances and databases.</p>
<b>Operations per second</b>	<p>The number of operations that Cloud Spanner performed per second, or the number of errors that occurred on the Cloud Spanner server per second.</p> <p>You can choose which operations to view in this chart:</p> <ul style="list-style-type: none"> <li>• Reads and writes</li> <li>• Reads only (also includes DML statements)</li> <li>• Writes only (excludes DML statements)</li> <li>• Errors on the Cloud Spanner server</li> </ul> <p>Available for instances and databases.</p>

## Available charts and metrics

### Throughput (MB/sec)

The amount of uncompressed data that was read from, or written to, the instance or database each second. This value is measured in *binary megabytes* (MB), where 1 MB is  $2^{20}$  bytes. This unit of measurement is also known as a mebibyte (MiB). (<https://en.wikipedia.org/wiki/Mebibyte>).

Read throughput includes requests and responses for methods in the read API (/spanner/docs/reads) and for SQL queries. It also includes requests and responses for DML statements.

Write throughput includes requests and responses to commit data through the mutation API (/spanner/docs/modify-mutation-api). It excludes requests and responses for DML statements.

Available for instances and databases.

### Total storage

The amount of data that is stored in the instance or database. This value is measured in *binary byte units*. For example, 1 binary gigabyte (GB) is  $2^{30}$  bytes; this unit of measurement is also known as a gibibyte (GiB). (<https://en.wikipedia.org/wiki/Gibibyte>).

Available for instances and databases.

- Understand the CPU utilization (/spanner/docs/cpu-utilization) and latency (/spanner/docs/latency) metrics for Cloud Spanner.
- Set up customized charts and alerts (/spanner/docs/monitoring-stackdriver) with Stackdriver Monitoring.
- Get details about types of Cloud Spanner instances (/spanner/docs/instances).