

Stackdriver Monitoring in the Cloud Console is now Generally Available and the default experience. For a limited period you also have the option to use the classic Stackdriver Monitoring console. For more information, see [Stackdriver Monitoring in the Cloud Console \(/monitoring/docs/monitoring_in_console\)](/monitoring/docs/monitoring_in_console).

This document describes how to use the Stackdriver Monitoring console to monitor your Cloud Spanner instances.

The Stackdriver Monitoring console provides several monitoring tools for Cloud Spanner:

- *A curated dashboard*, which shows pre-made charts for your Cloud Spanner resources
- *Custom charts*, including [ad-hoc charts in the Metrics Explorer \(/monitoring/charts/metrics-explorer\)](/monitoring/charts/metrics-explorer) as well as [charts in custom dashboards \(/monitoring/charts/\)](/monitoring/charts/)
- *Alerts*, which notify you if a metric exceeds a threshold that you specify

If you prefer to monitor Cloud Spanner programmatically, use the [Cloud Client Libraries for Stackdriver Monitoring \(/monitoring/docs/reference/libraries\)](/monitoring/docs/reference/libraries) to retrieve metrics.

You can also monitor your instances by [viewing charts in the Google Cloud Console \(/spanner/docs/monitoring-console\)](/spanner/docs/monitoring-console) in the Cloud Console to get a quick view of the most important metrics for your instance.

To use Stackdriver Monitoring, you must have a Stackdriver account. If you do not have a Stackdriver account, you can [create a new account \(/monitoring/accounts/guide#setup-account\)](/monitoring/accounts/guide#setup-account).

Stackdriver Monitoring provides you with a curated dashboard that summarizes key information about your Cloud Spanner instances, including:

- *Incidents*: User-created monitoring alerts that are open, active, or resolved
- *Events*: A list of Cloud Spanner [audit logs \(/spanner/docs/logs\)](/spanner/docs/logs) (if enabled and available)

- *Instances*: A high-level summary of your Cloud Spanner instances, including node count, database count, and instance health
- *Aggregated charts* of throughput and storage use

To view the Cloud Spanner dashboard, do the following:

1. In the Google Cloud Console, select **Monitoring**, or use the following button:

[Go to Monitoring \(https://console.cloud.google.com/monitoring\)](https://console.cloud.google.com/monitoring)

2. If **Resources** is shown in the navigation pane, then select **Resources** and then select **Cloud Spanner**. Otherwise, select **Dashboards** and then select the dashboard named **Cloud Spanner**.

When you open the curated dashboard for Cloud Spanner, it shows aggregated data for all of your instances. You can view more details about a specific instance by clicking the instance's name under **Instances**.

The dashboard displays information such as instance metadata, databases in the instance, and charts of various metrics broken down by region.

From the instance dashboard page, you can also see charts for a specific database in the instance:

1. On the right-hand side, above the instance metrics charts, click **Database metrics**.
2. In the **Select a breakdown** drop-down list, select the database that you want to examine.

The Stackdriver Monitoring console displays charts for the database.

You can use Stackdriver Monitoring to create custom charts for Cloud Spanner metrics. You can use the Metrics Explorer to create temporary, ad-hoc charts, or you can create charts that appear on custom dashboards.

In particular, Stackdriver Monitoring enables you to create a custom chart that shows whether two or more metrics are correlated with each other. For example, you can check for a correlation between [CPU utilization \(/spanner/docs/cpu-utilization\)](/spanner/docs/cpu-utilization) and [latency \(/spanner/docs/latency\)](/spanner/docs/latency) in a Cloud Spanner instance, which might indicate that your instance needs more nodes or that some of your queries are causing high CPU utilization.

To get started with this example, follow these steps:

1. In the Google Cloud Console, select **Monitoring**, or use the following button:

[Go to Monitoring \(https://console.cloud.google.com/monitoring\)](https://console.cloud.google.com/monitoring)

2. If **Metrics Explorer** is shown in the navigation pane, select it. Otherwise, select **Resources** and then select **Metrics Explorer**.
3. Click the **View options** tab, then select the **Log scale on Y-axis** checkbox. This option helps you compare multiple metrics when one metric has much larger values than the others.
4. In the drop-down list above the right pane, select **Line**.
5. Click the **Metrics** tab. You can now add metrics to the chart.

To add latency metrics to the chart, follow these steps:

1. In the **Find resource type and metric** box, enter the value `spanner.googleapis.com/api/request_latencies`, then click the row that appears below the box.
2. In the **Filter** box, enter the value `instance_id`, then enter the instance ID you want to examine and click **Apply**.
3. In the **Aggregator** drop-down list, click **max**.
4. Optional: Change the latency percentile:
 - a. Click **Show advanced options**.
 - b. Click the **Aligner** drop-down list, then click the latency percentile that you want to view.

In most cases, you should look at either the 50th percentile latency, to understand the typical amount of latency, or the 99th percentile latency, to understand the latency for the slowest 1% of requests.

To add CPU utilization metrics to the chart, follow these steps:

1. Click **+ Add metric**.
2. In the **Find resource type and metric** box, enter the value `spanner.googleapis.com/instance/cpu/utilization`, then click the row that appears below the box.
3. In the **Filter** box, enter the value `instance_id`, then enter the instance ID you want to examine and click **Apply**.

4. In the **Aggregator** drop-down list, click **max**.


You now have a chart that shows the CPU utilization and latency metrics for a Cloud Spanner instance. If both metrics are higher than expected at the same time, you can [take additional steps to correct the issue](#) (/spanner/docs/latency#diagnose-high-spanner-latency).


For more information about creating custom charts, see the [Stackdriver Monitoring documentation](#) (/monitoring/charts/).


When you create a Cloud Spanner [instance](#) (/spanner/docs/instances), you choose the number of nodes that provide compute resources for the instance. As the instance's workload changes, Cloud Spanner does not automatically adjust the number of nodes in the instance. As a result, you need to set up several alerts to ensure that the instance stays within the [recommended maximums for CPU utilization](#) (/spanner/docs/cpu-utilization#create-stackdriver-alerts) and the [recommended limit for storage per node](#) (/spanner/docs/limits).

You can also create alerts for other metrics. To find out what metrics are available, see the [metrics list for Cloud Spanner](#) (/spanner/docs/api/metrics_gcp#gcp-spanner).

To create the recommended alerting policies, use the instructions and settings displayed in the following table:







Cloud Spanner usage amounts are calculated in *binary terabytes*, where 1 TB is 2^{40} bytes. This unit of measurement is as a tebibyte (TiB). (<https://en.wikipedia.org/wiki/Tebibyte>).

- Understand the CPU utilization (/spanner/docs/cpu-utilization) and latency (/spanner/docs/latency) metrics for Cloud Spanner.

- [Use the Cloud Console \(/spanner/docs/monitoring-console\)](/spanner/docs/monitoring-console) to get a quick view of the most important metrics for your instance.
- Learn more about [Stackdriver Monitoring \(/monitoring/docs/\)](/monitoring/docs/).