

Google Cloud Platform (GCP) provides a custom monitoring agent that collects custom metrics from SAP HANA and sends them to Google Stackdriver Monitoring, GCP's built-in monitoring solution. With Stackdriver Monitoring, you can build out dashboards to visualize your SAP HANA metrics and set up alerts based on metric thresholds. For more information, see the [Stackdriver Monitoring documentation](/monitoring/docs/) (/monitoring/docs/).

This guide assumes that you have deployed SAP HANA on GCP using the method described in the [SAP HANA Deployment Guide](/solutions/sap/docs/sap-hana-deployment-guide) (/solutions/sap/docs/sap-hana-deployment-guide).

To use Google's monitoring agent for SAP HANA, you must set up a new [Workspace](/monitoring/workspaces/) (/monitoring/workspaces/) that is connected to your GCP project. The following steps create a new Workspace:

1. In your browser, go to Stackdriver Monitoring:

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

2. Click **Add Workspace**.
3. On the **Create your free Workspace** page, click the **Google Cloud Platform project** field. A dropdown containing your GCP projects appears.
4. Select your project from the dropdown and click **Create Account**.
5. On the **Add Google Cloud Platform projects to monitor** page, click **Continue**.
6. On the **Monitor AWS accounts** page, click **Skip AWS Setup**.
7. On the **Install the Stackdriver Agents** page, click **Continue**.
8. On the **Get Reports by Email** page, select **No reports** and then click **Continue**.
9. After the message "Finished initial collection!" appears, click **Launch monitoring** to enter your Workspace dashboard.

10. Click **Continue with the trial** to finish setup.

You now have a Workspace monitoring your GCP project, and the monitoring agent can send metrics to Stackdriver.

By default, the monitoring agent uses your Compute Engine virtual machine (VM) instance's default service account (/iam/docs/service-accounts). This service account allows binaries running on the VM instance to write metrics to Stackdriver.

If you choose to use a different service account, you must manually add the IAM role (/iam/docs/understanding-roles) that provides these permissions to your service account.

To add the required IAM role to your service account:

1. Go to the **IAM & Admin** page in the Cloud Console.

[Go to the IAM & Admin page](https://console.cloud.google.com/project/_/iam-admin) (https://console.cloud.google.com/project/_/iam-admin)

2. Select your project and click **Continue**.
3. Identify the service account to which you want to add a role.
 - If the service account isn't already on the members list, it doesn't have any roles assigned to it. Click **Add member** and enter the email address of the service account.
 - If the service account is already on the members list, it has existing roles. Click the current role dropdown list for the service account that you want to edit.
4. Select **Monitoring > Monitoring Metric Writer** from the list of available roles.
5. Click **Add** or **Save** to apply the roles to the service account.

The monitoring agent can be run directly on one or more nodes in your SAP HANA installation, or indirectly as a service on another VM instance. Choose your deployment location based on your use case.

To download and install the agent:

1. Establish an SSH connection with your SAP HANA VM instance.
2. Download and run the installation script:

When you run the above command, the installation script performs the following operations:

1. Creates the following directories:
 - `/usr/local/lib/gcm`, which will contain the agent binary.
 - `/etc/gcm.d`, which will contain the queries you want the agent to run.
 - `/var/spool/gcm`, which contains the metrics that have recently been gathered by the agent. The agent eventually sends these metrics to Stackdriver Monitoring.
2. Downloads the latest release of the agent and saves it to `/usr/local/lib/gcm/`.
3. Creates a symlink to the binary in `/usr/local/bin/gcm`.
4. Creates a basic configuration template, `/etc/default/gcm`.
5. Creates a daemon configuration file, `/etc/systemd/system/gcm.service`.
6. Reloads the `systemd` service definitions.
7. Enables the monitoring agent as a daemon.

After running the install script, configure the VM instances that the agent will monitor and, optionally, the queries that the agent will perform on your database.

When you ran the installation script, it created a configuration template file at the following location:

Using the configuration file, you can configure the agent to query SAP HANA on one or more VM instances. For each VM instance you want to monitor:

1. Add a VM instance definition in the configuration file under `instances`. For each instance, define the name, host, port, user, and password.
2. To enable the agent to collect metrics from that VM instance, add the VM instance name under `enabled_instances`.

Define the attributes in the configuration by using the following format:

The following attributes are required in the configuration file:

- `[PROJECT_ID]`: The ID of the Google Cloud project that contains the VM instances to be monitored. Specify only one project ID.
- `[INSTANCE_NAME]`: The name of the Compute Engine VM instance you want to monitor.

- `[INSTANCE_IP]`: The internal or external IP of the Compute Engine VM instance you want to monitor. Google recommends using internal IPs if possible, as they require less setup than external IPs and are private and secure by default.
- `[DB_USERNAME]`: The SAP HANA database user you want to use.
- `[PORT_NUMBER]`: The port number of your SAP HANA database. Usually, the port number is 39015 for SAP HANA Express or 30015 for SAP HANA Platform Edition.
- `[YOUR_PASSWORD]`: The user password for your database user.
- `[INSTANCE_TYPE]`: Enter `sap_hana` for this field.

Optionally, you can configure the agent using the following flags. These parameters flags must be set before the `config` item:

- `timestamps_in_output` (*boolean*). If you're running the agent as a daemon under `systemd`, set this to `no`, as `systemd` provides timestamps. Default value is `no`.
- `debug_messages` (*boolean*): If set, show debug messages. Default value is `yes`.
- `skip_default_queries` (*boolean*): Skips the default set of queries. Does not affect user-defined queries. Default value is `no`.
- `queries_directory` (*string*): The directory in which to look for user-defined queries. Default value is `/etc/gcm.d`.
- `spool_directory` (*string*): The directory in which to collect spooled metrics. Default value is `/var/spool/gcm`.

By default, the agent gathers a default set of metrics from your SAP HANA database. For descriptions of these metrics, as well as the queries that generate the metrics, see [SAP HANA Default Metrics \(/solutions/sap/docs/sap-hana-default-metrics\)](/solutions/sap/docs/sap-hana-default-metrics).

You can add additional queries by creating one or more custom YAML files in the following directory:

The following is an example query file:

Each YAML file comprises one or more items that represent the results of a given SQL query into your SAP HANA database's metrics. Each item has the following attributes:

- **root**: A descriptive namespace for a given group of metrics.
- **query**: The SQL query.
- **columns**: The data type, value type, and column name for each field in the query.
 - **type**: Describes how the data is reported to Stackdriver. You can set **type** to **LABEL**, which indicates that the type is a column head, or to **GAUGE**, one of the metric types defined in the Stackdriver API's [MetricKind enum](#) (`/monitoring/api/ref_v3/rest/v3/projects.metricDescriptors#MetricKind`). Currently, **GAUGE** is the only metric type supported by the agent.
 - **value_type**: The value type of the metric. The value of this parameter can be any of the value types defined in the Stackdriver API's [ValueType enum](#) (`/monitoring/api/ref_v3/rest/v3/projects.metricDescriptors#ValueType`).
 - **name**: The name of the column.

For a full list of the system views that SAP HANA makes available for querying, see the [SAP HANA SQL and System Views Reference](#)

(<https://help.sap.com/doc/4fe29514fd584807ac9f2a04f6754767/2.0.01/en-US/index.html>).

This section describes how to perform basic operations with the SAP HANA monitoring agent. For a full list of configurable options, run `gcm --help`.

When you ran the installation script, the installer created a `systemd` unit file for the agent, allowing you to manage the agent using standard `systemctl` commands. The following commands start, stop, and poll the status of the agent, respectively:

To read the logs generated by `systemctl`, query the contents of the `systemd` journal by using the following command:

To see the last few lines that were logged, you can add the `-f` flag, which functions like a pipe to `tail -f`:

By default, the agent runs its default queries in addition to any custom queries you've defined. You can disable this default set of queries by setting the `skip_default_queries` flag to `yes` in your config

file, or by setting the `--no-defaults` flag when running the tool manually. For more information, see [Defining the configuration file](#) (`#defining_the_configuration_file`).

To verify that the agent works as expected, you can manually run commands to run your queries, collect the resulting metrics, and then send them to Stackdriver Monitoring. To run your queries and collect the resulting metrics exactly once, run the following command:

To send the metrics collected by the agent when you ran your queries, run the following command:

The following steps assume that you are using the default queries.

To display the metrics collected by Stackdriver Monitoring in your own charts and dashboards:

1. In the Cloud Console, go to the **Monitoring** page.

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

2. Select **Dashboards > Create Dashboard**.
3. Click **Add Chart**.
4. In the **Resource Type** menu, select **Custom Metrics**.
5. In the **Metric** menu, select **by_component/mem_used_mb**. Leave the other fields with their default values. You see the chart data in the Preview section of the panel.
6. Click **Save**.

You now have a simple dashboard displaying live metrics from your SAP HANA VM instance or instances.

Make sure that your SAP HANA database user has access to the following SAP HANA system views:

- `M_CS_ALL_COLUMNS`
- `M_CS_TABLES`
- `M_EXPENSIVE_STATEMENTS`
- `M_HOST_RESOURCE_UTILIZATION`
- `M_SERVICE_COMPONENT_MEMORY`
- `M_SERVICE_MEMORY`

Make sure that your GCP service account has permission to write metrics to Stackdriver Monitoring. For details, see [Setting the required IAM roles \(#setting_the_required_iam_roles\)](#).

Google Cloud customers either with a Production Support Role or with Enterprise Support can request assistance with the provisioning and configuration of the Google Cloud resources that are required for SAP systems. Google Cloud Production-level support or Enterprise support is required for support of SAP systems in production environments.

For more information about Google Cloud support options, see [Google Cloud Support \(/support/\)](#).

For SAP product-related issues, log your support request with [SAP support](#) (<https://support.sap.com/support-programs-services/about/getting-started.html>). SAP evaluates the support ticket and, if it appears to be a Google Cloud infrastructure issue, transfers the ticket to the Google Cloud queue.

- Review the [default metrics \(/solutions/sap/docs/sap-hana-default-metrics\)](#) collected by the agent.

- [Learn more about Stackdriver Monitoring \(/monitoring/docs/\)](#).