

This guide provides an overview of how to manage IT ops for a running SAP NetWeaver system on Compute Engine. This guide does not cover the specifics of managing the SAP NetWeaver product. Google Cloud is certified for running SAP NetWeaver application servers ABAP and Java, and SAP products based on these application server stacks.

- To learn how to operate your deployment of SAP NetWeaver, see the [SAP NetWeaver Master Guide](https://help.sap.com/doc/18cb1a50b9924bc3b94c2988cc8c51d9/7.5/en-US/mg_nw_75.pdf) (https://help.sap.com/doc/18cb1a50b9924bc3b94c2988cc8c51d9/7.5/en-US/mg_nw_75.pdf).
- To learn how to plan your deployment of SAP NetWeaver, see the [SAP NetWeaver on Google Cloud Planning Guide](/solutions/sap/docs/netweaver-planning-guide) (/solutions/sap/docs/netweaver-planning-guide) and the [SAP NetWeaver Master Guide](https://help.sap.com/doc/18cb1a50b9924bc3b94c2988cc8c51d9/7.5/en-US/mg_nw_75.pdf) (https://help.sap.com/doc/18cb1a50b9924bc3b94c2988cc8c51d9/7.5/en-US/mg_nw_75.pdf).
- To learn how to deploy and prepare Google Cloud resources for an SAP NetWeaver system, see the deployment guide for your operating system:
 - [SAP NetWeaver on Linux Deployment Guide](/solutions/sap/docs/netweaver-deployment-guide-linux) (/solutions/sap/docs/netweaver-deployment-guide-linux)
 - [SAP NetWeaver on Windows Deployment Guide](/solutions/sap/docs/netweaver-deployment-guide-windows) (/solutions/sap/docs/netweaver-deployment-guide-windows)

This section shows how to perform administrative tasks typically required to operate SAP NetWeaver on Compute Engine, including information about starting, stopping, and cloning systems.

This section provides information about how to manage the running state of your Compute Engine VMs.

Compute Engine offers availability policies that determine how a VM behaves during certain infrastructure-related events. For VM instances in your SAP NetWeaver implementation, it is important that you do not disable the following features:

- [Live migrate](/compute/docs/instances/setting-instance-scheduling-options#live_migrate) (/compute/docs/instances/setting-instance-scheduling-options#live_migrate), which enables Compute Engine to keep your instance running when responding to an infrastructure maintenance event.
- [Automatic restart](/compute/docs/instances/setting-instance-scheduling-options#autorestart) (/compute/docs/instances/setting-instance-scheduling-options#autorestart), which enables Compute Engine to restart your instance in the event of an instance crash. Note that the SAP system does not automatically restart.

You can stop one or multiple SAP NetWeaver hosts at any time; stopping a VM instance shuts down the instance. If the shutdown doesn't complete within 2 minutes, the instance is forced to halt. As a best practice, you should first stop SAP NetWeaver before you stop the instance.

The private IP associated with the instance is not released after you stop the instance, so when you start the instance it will start with the same private IP address, network, and storage configuration as before.

Stopping a VM causes Compute Engine to send the [ACPI power-off signal](https://wikipedia.org/wiki/Advanced_Configuration_and_Power_Interface) (https://wikipedia.org/wiki/Advanced_Configuration_and_Power_Interface) to the VM instance. After it is stopped, you are not billed for the VM instance.

If you have persistent disks attached to the VM, the disks are not deleted and you continue to be charged for them. If the data in the persistent disk is important, you can either:

- Keep the disk.
- Create a snapshot of the persistent disk and then delete the disk. This option can help you save on costs. You can create another disk from the snapshot when you need the data again.

To stop a VM:

1. In the Google Cloud Console, navigate to the:

[VM INSTANCES PAGE](https://console.cloud.google.com/compute/instances) (https://console.cloud.google.com/compute/instances)

2. Select one or more instances that you want to stop.
3. At the top of the **VM instances** page, click **■ STOP**.

For alternatives and more information, see [Stopping an instance](/compute/docs/instances/stopping-or-deleting-an-instance#stop_an_instance) (/compute/docs/instances/stopping-or-deleting-an-instance#stop_an_instance).

To restart a VM:

1. In the Cloud Console, navigate to the:

VM INSTANCES PAGE (<https://console.cloud.google.com/compute/instances>)

2. Select the instances that you want to restart.
3. At the top right-hand of the page, click ► **START**.

For alternatives and more information, see [Restarting an instance](/compute/docs/instances/restarting-an-instance#restarting_a_stopped_instance) (/compute/docs/instances/restarting-an-instance#restarting_a_stopped_instance).

You can change various attributes of a VM, including the VM type, after the VM is deployed. Some changes might require you to restore your SAP system from backups, while others only require you to restart the VM.

For more information, see [Modifying VM configurations for SAP systems](/solutions/sap/docs/modifying_vm_configurations) (/solutions/sap/docs/modifying_vm_configurations).

This section describes scenarios that require saving the state of your system, and the Compute Engine features that you can use for these purposes.

You can create a snapshot of a persistent disk at any time to generate a point-in-time copy of the disk state. Compute Engine stores multiple copies of each snapshot, across multiple locations, with automatic checksums to ensure the integrity of your data.

Snapshots are useful for the following use cases:

Use case	Details
----------	---------

Use case	Details
Migrate to a different type of storage.	You can use snapshots to move a persistent disk from one disk type, <i>standard</i> or <i>SSD</i> , to the other type. See Restoring a snapshot to a different disktype (/compute/docs/disks/restore-and-delete-snapshots#restoresnapshotlargerize) in the Compute Engine documentation.
Migrate SAP NetWeaver to another zone.	You can use snapshots to move your SAP NetWeaver system from one zone to another zone in the same region, or even from one region to another region. See Moving VMs between regions and zones (#moving_vms).
Provide an easy, software-independent, and cost-effective backup solution.	Back up your attached persistent disks by using snapshots. You can back up your root disk and SAP NetWeaver installed binaries. Snapshots can be useful for taking backups of your database systems. However, depending on your implementation, you might want to use a different approach. For guidance about how to back up and restore databases, see the guides listed in Database operations (#database_operations).

To obtain a consistent snapshot, you must either stop SAP NetWeaver or stop the database from writing to the file system.

To create a snapshot, follow the Compute Engine instructions for [creating snapshots](#) (/compute/docs/disks/create-snapshots#creating_snapshots). Pay careful attention to the preparation steps, such as flushing the disk buffers to disk, to make sure that the snapshot is consistent.

To clone your SAP NetWeaver system on Google Cloud, follow the standard SAP export-import procedure:

1. Use the Software Provisioning Manager (SWPM) to export the source system.
2. [Copy the data](#) (/compute/docs/disks/gcs-buckets#writetobucket) from the system and database export to your Cloud Storage bucket.
3. Use SWPM to create a new, target system and to import the artifacts that you exported from the source system. You can [mount the Cloud Storage bucket](#) (/compute/docs/disks/gcs-buckets) as a file system for use by the target system.

To capture the state of a boot disk, you can create a *custom image*. An image is different from a backup because you use an image to create new VM instances that are based on a single, source

VM.

Important: Creating an image is a destructive operation with regards to the source VM, so this is not something you would typically do in your live, production system.

When you followed the SAP NetWeaver on Google Cloud Deployment Guide, you should have created one or more images at the end of the deployment steps. However, you might want to create new images after you make important changes to the system, such as installing an update of SAP NetWeaver binaries or upgrading the SAP NetWeaver version.

For instructions, see:

- [Creating, Deleting, and Deprecating Custom Images](/compute/docs/images/create-delete-deprecate-private-images) (/compute/docs/images/create-delete-deprecate-private-images), for Linux-based systems.
- [Creating a Windows image](/compute/docs/instances/windows/creating-windows-os-image) (/compute/docs/instances/windows/creating-windows-os-image), for Windows-based systems.

Compute Engine enables you to move VMs between zones in the same region and zones in different regions. You might want to move a VM if, for example, a new region or zone becomes available that would give you better performance, or if a zone becomes deprecated.

The Compute Engine documentation contains [detailed instructions about how to move your VM to another zone](/compute/docs/instances/moving-instance-across-zones) (/compute/docs/instances/moving-instance-across-zones).

Here are considerations for SAP NetWeaver:

- SAP can run only in certain zones because of machine-type restrictions. See the [SAP NetWeaver on Google Cloud Planning Guide](/solutions/sap/docs/netweaver-planning-guide) (/solutions/sap/docs/netweaver-planning-guide) for details.
- Migrating the VM causes the VM's ID to change. This change triggers an SAP HW Key change, which requires you to import a new SAP license.
- You can use the same hostname in the new zone, if it isn't already in use. If the hostname changes, you need to use the generic operations feature of SAP's SWPM to run a rename operation to change the SAP NetWeaver hostname.

This section provides resources for managing the following database servers on Google Cloud:

- SAP HANA
- SAP ASE
- IBM Db2 for Linux, UNIX and Windows (IBM Db2)
- Microsoft SQL Server

For more information about running SAP HANA on Google Cloud, see the [SAP HANA on Google Cloud Operations Guide](/solutions/sap/docs/sap-hana-operations-guide) (/solutions/sap/docs/sap-hana-operations-guide). That guide provides you with lots of details about administration, backup and recovery, security, networking, and other topics.

For more information about using SAP ASE, see [SAP Adaptive Server Enterprise](https://help.sap.com/viewer/product/SAP_ASE/) (https://help.sap.com/viewer/product/SAP_ASE/).

For more information about using IBM Db2 with SAP, see [SAP on IBM Db2 for Linux, UNIX and Windows](https://help.sap.com/viewer/p/DB6) (https://help.sap.com/viewer/p/DB6).

The following resources provide details about how to run Microsoft SQL Server on Google Cloud:

Resource	Description
----------	-------------

Resource	Description
Best Practices for Microsoft SQL Server (/compute/docs/instances/sql-server/best-practices)	<p>Learn how to configure Microsoft SQL Server for stability and performance on Compute Engine.</p> <p>Note the following important differences in best practices for SAP systems:</p> <ul style="list-style-type: none"> • Do not use local SSD drives. Use persistent disk SSDs, instead. • For parallel query processing, set the max degree of parallelism to 1, not 8. • The transaction log setting must be "FULL" for Microsoft SQL Server databases in production systems. • Do not use the buffer pool extension feature.
Microsoft SQL Server Performance Tuning on Compute Engine (/solutions/sql-server-performance-tuning-compute-engine)	<p>Discusses how to create a Microsoft SQL server on Compute Engine and then use metrics to optimize its performance.</p>
Load Testing Microsoft SQL Server Using HammerDB (/compute/docs/tutorials/load-testing-sql-server-hammerdb)	<p>This tutorial shows how to use HammerDB to perform load testing on a Compute Engine Microsoft SQL Server instance.</p>
Building a Microsoft SQL Server Disaster Recovery Plan on Compute Engine (/solutions/sql-server-disaster-recovery-plan-compute-engine)	<p>Demonstrates how to implement a less-expensive disaster recovery solution for Microsoft SQL Server on Google Cloud.</p>

Controlling access to Google Cloud resources is a critical part of securing and operating your deployment. While SAP provides its own user-management system, GCP's [Cloud Identity and Access Management \(Cloud IAM\)](/iam/docs/overview) (/iam/docs/overview) provides unified control over permissions for Google Cloud resources. You can manage access control by defining who has what access for resources. For example, you can control who can perform control-plane operations on your SAP instances such as creating and modifying VMs, persistent disks, and networking.

For an overview of Cloud IAM in Compute Engine, see [Access Control Options](/compute/docs/access/) (/compute/docs/access/).

From time to time, you will want to add or remove team members from your project or change their permission levels. For details about how to manage team members, see [Add, Remove Team Members, and Change Permissions](#)

(/compute/docs/access/add-remove-change-permissions-for-team-members).

IAM *roles* are key to granting permissions to users. For a reference about roles and which permissions they provide, see [Identity and Access Management Roles](#) (/iam/docs/overview).

By default, Compute Engine automatically manages SSH keys. If you have decided to manage your own SSH keys, you need to add and remove keys from time to time during your normal operations. For detailed steps, see [Adding and Removing SSH Keys](#)

(/compute/docs/instances/adding-removing-ssh-keys).

Cloud IAM's *service accounts* provide a way for you to give permissions to applications and services. It's important to understand how service accounts work in Compute Engine.

If a service account is assigned to a Compute Engine VM, that service account is the default service account for the applications that run on that VM. Any application that uses the VM service account inherits the Cloud IAM roles and permissions that are granted to the VM service account.

For more information, see [Identity and access management for SAP programs on Google Cloud](#) (/solutions/sap/docs/security-for-sap/iam-for-sap-programs).

[Stackdriver Logging](#) (/logging/docs/) is Google Cloud's solution for system-wide logging. Stackdriver Logging allows you to store, search, analyze, monitor, and alert on log data and events. Using Stackdriver Logging requires that you have installed the Stackdriver Logging agent on each VM.

If you didn't install the agent, you can install it now. See [Installing the Agent](#) (/logging/docs/agent/installation).

See [Compute Engine Logs](#) (/logging/docs/view/service/compute-logs) for details about supported logs.

Stackdriver logging provides granular access control to logs and logging operations. For details, see the [Access Control Guide](/logging/docs/access-control/) (/logging/docs/access-control).

[Cloud Audit Logs](/logging/docs/audit/) (/logging/docs/audit/) provides key information about activities happening in GCP through two log types: [Admin Activity](/logging/docs/audit/#admin-activity) (/logging/docs/audit/#admin-activity) and [Data Access](/logging/docs/audit/#data-access) (/logging/docs/audit/#data-access). You can view the Activity Feed and the Logs Viewer in the Cloud Console.

To help you perform tasks such as analyzing system performance and detecting and diagnosing problems early, SAP NetWeaver provides a central monitoring system that collects data about the components and activities system-wide. Google Cloud provides its own monitoring system, [Stackdriver Monitoring](/monitoring/docs/) (/monitoring/docs/), to collect metrics, events, and metadata. As you implement and operate SAP NetWeaver on Google Cloud, dealing with two disconnected systems and trying to figure out where the real issues exist can become challenging for support personnel. To make things easier, Google and SAP have worked together to create a *monitoring agent* for SAP NetWeaver running on Google Cloud.

You must deploy Google's monitoring agent to get support from SAP to and to enable SAP to meet its service-level nents (SLAs).

When managing monitoring operations, it's helpful to understand what Google's monitoring agent is doing. In general, here's how it works:

- [Monitoring](/monitoring/docs/) (/monitoring/docs/) has a local agent that collects metrics, events, and metadata from Google Cloud. Compute Engine also provides APIs that provide monitoring functionality.
- Each VM in your deployment must host an instance of Google's monitoring agent for SAP NetWeaver. The monitoring agent runs as a Windows service or a Linux process.
- The monitoring agent combines monitoring data from Monitoring and the Compute Engine APIs.

- The SAP Host Agent polls Google's monitoring agent for its cached data, over HTTP. It aggregates the metrics, reports them, and stores them in the SAP NetWeaver database.
- SAP's transaction ST06 or the SAPOSCOL command line interface displays the aggregated metrics.
- You can view the data from Google's monitoring agent by running a command in a terminal window (`#viewing_monitoring_data`).

When you install the monitoring agent, the provided start-up script completes the following tasks:

- Downloads the latest version of the monitoring agent.
- Downloads the dependencies, such as the OpenJDK (<http://openjdk.java.net/>) and SIGAR (<https://github.com/hyperic/sigar>) libraries.
- In Linux, creates two `cron` jobs as root that perform the following tasks:
 - Monitors whether the metrics agent is running and restarts the agent if necessary.
 - Refreshes the agent when Google updates the metrics agent provider.

ng: Do not delete these `cron` entries. They are required for the continuous operation of Google's monitoring agent.

- To view these `cron` entries:

- In Windows, the monitoring system runs as a Windows service, named "GCP Metrics Provider". The service restarts automatically, when needed. The monitoring system also depends on a Windows Task Scheduler task that runs once a day.

ng: Do not manually stop the monitoring system's Windows service, or delete the Task Scheduler task.

You can install the monitoring agent by downloading and running the script from the following public Cloud Storage bucket:

- For Linux (https://storage.googleapis.com/sap-netweaver-on-gcp/setupagent_linux.sh)

- For Windows (<https://storage.googleapis.com/sap-netweaver-on-gcp/installer.ps1>)

If you didn't install Google's monitoring agent while creating your VM, you can install the agent in the following ways:

The remainder of this section provides details about managing Google's monitoring agent.

This section describes actions you can take to keep the monitoring agent working properly.

Google's monitoring agent is a local HTTP server. You can check whether the monitoring agent is running by polling for a *health check* from the server. Follow these steps:

If the monitoring agent is functioning properly, the status is `UP`. For example:

If the monitoring agent isn't running, see [Restarting Google's monitoring agent \(#restarting_monitoring\)](#).

To check whether the connection between Google's monitoring agent and SAP NetWeaver works, enter transaction `ST06` in your SAP NetWeaver ABAP System. In the overview pane, check the availability and content of the following fields for the correct, end-to-end setup of the SAP and Google monitoring infrastructure:

- **Cloud Provider:** "Google Cloud Platform"
- **Enhanced Monitoring Access:** "TRUE"

- **Enhanced Monitoring Details: "ACTIVE"**

You can view the monitoring agent's activity by polling the server. Follow these steps:

You should see output from the monitoring agent logged in the terminal window. Here's an example of a line of output:

If the monitoring agent stops working, follow these steps to restart:

This section describes issues you can investigate if the monitoring agent isn't working.

Google's monitoring agent uses HTTP as its communication protocol and it requires access to your network through a specific port. If the agent isn't working, follow these steps:

1. In the Google Cloud Console, go to Firewall rules.

[GO TO THE FIREWALL RULES PAGE \(https://console.cloud.google.com/networking/firewalls\)](https://console.cloud.google.com/networking/firewalls)

2. Create a new rule or edit an existing rule to open port 18181.

On Windows-based systems, also ensure that the Windows Firewall doesn't block the monitoring agent from using the port.

Check the logs available for the agent:

- For Linux: `/var/log/gcp-metric-provider/gcp-metric-provider.log`.
- For Windows: `C:\Program Files\Google\GCP Metrics Provider\Logs\gcp-metric-provider.log`

If the logs contain an `OutOfMemoryError` entry, restart the agent (`#restarting_monitoring`).

Google's monitoring agent uses the service account of its host VM to retrieve Stackdriver Monitoring metrics. Consequently, the monitoring agent requires that the host VM have a service account and that the service account includes the `monitoring.timeSeries.list` permission, which is contained in the predefined Monitoring Viewer role.

If you install Stackdriver Monitoring agent, you might need to grant additional Cloud IAM permissions to your VM service account, such as the predefined Monitoring Metric Writer role. To confirm the permissions that the Monitoring agent requires, see the Stackdriver Monitoring documentation:

- [Authorizing the agent](#) (`/monitoring/agent/authorization`)
- [Access control](#) (`/monitoring/access-control`)

[Access scopes](#) (`/compute/docs/access/service-accounts#accessscopesiam`) are the legacy method of specifying permissions for your instance.

A best practice is to set the full `cloud-platform` access scope on the instance, then securely limit the service account's API access with Cloud IAM roles. For example:

- `'https://www.googleapis.com/auth/cloud-platform'`

If you do limit the access scopes of your VM, you must ensure that the host VM has the following access scopes:

- `https://www.googleapis.com/auth/source.read_write`
- `https://www.googleapis.com/auth/compute`
- `https://www.googleapis.com/auth/servicecontrol`
- `https://www.googleapis.com/auth/service.management.readonly`
- `https://www.googleapis.com/auth/logging.admin`
- `https://www.googleapis.com/auth/monitoring`

- <https://www.googleapis.com/auth/trace.append>
- https://www.googleapis.com/auth/devstorage.full_control

To change the access scopes, you need to stop your VM instance, make the changes, and then restart the instance. For instructions, see the [Compute Engine documentation](#) (</compute/docs/access/create-enable-service-accounts-for-instances#changeserviceaccountandscopes>). You don't need to make any changes to permissions for IAM roles for this issue.

For the monitoring system to work, your SAP NetWeaver system must have the SAP Host Agent installed and the required minimum patch level for the Host Agent. For instructions to install the SAP Host Agent, see [the SAP documentation](#)

(https://help.sap.com/saphelp_nw74/helpdata/en/21/98c443122744efae67c0352033691d/frameset.htm).

For version requirements for the SAP Host Agent, refer to the following SAP Notes:

- Linux: SAP Note [2460297 - SAP on Linux on Google Cloud Platform: Enhanced Monitoring](#) (<https://launchpad.support.sap.com/#/notes/2460297>)
- Windows: SAP Note [1409604 - Virtualization on Windows: Enhanced Monitoring](#) (<https://launchpad.support.sap.com/#/notes/0001409604>)

If the VM running Google's monitoring agent was created without a public IP address, the monitoring agent cannot be downloaded. For a description about how to set up a NAT gateway that gives the VM outbound access to the internet, see the SAP NetWeaver deployment guide for your operating system:

- [SAP NetWeaver on Linux Deployment Guide](#) (</solutions/sap/docs/netweaver-deployment-guide-linux>)
- [SAP NetWeaver on Windows Deployment Guide](#) (</solutions/sap/docs/netweaver-deployment-guide-windows>)

Google's monitoring listens for requests on port 18181. This port must be available or the monitoring agent cannot start up. In this case, the SAP Host Agent logs will show a **Connection Refused** error. Make sure port 18181 is available for the monitoring agent. You might need to restart whatever other service is using port 18181 in order to free the port.

