

This deployment guide shows you how to deploy and connect to a Compute Engine virtual machine (VM) that is ready for the installation of SAP MaxDB on a Linux operating system.

The instructions in this guide use [Cloud Deployment Manager](/deployment-manager/docs/fundamentals) (/deployment-manager/docs/fundamentals) to deploy and configure the VM, the Linux operating system, and the required disk volumes for SAP MaxDB.

For details on planning your deployment, see the [SAP MaxDB Planning Guide](/solutions/sap/docs/sap-maxdb-planning-guide) (/solutions/sap/docs/sap-maxdb-planning-guide).

If you do not already have a Google Cloud project with billing enabled, you must create one before you can deploy a VM for your SAP MaxDB installation.

To create a project:

1. [Sign in](https://accounts.google.com/Login) (https://accounts.google.com/Login) to your Google Account.

If you don't already have one, [sign up for a new account](https://accounts.google.com/SignUp) (https://accounts.google.com/SignUp).

2. In the Cloud Console, on the project selector page, select or create a Cloud project.

★ **Note:** If you don't plan to keep the resources that you create in this procedure, create a project instead of selecting an existing project. After you finish these steps, you can delete the project, removing all resources associated with the project.

[Go to the project selector page](https://console.cloud.google.com/projectselector2/home/dashboard) (https://console.cloud.google.com/projectselector2/home/dashboard)

3. Make sure that billing is enabled for your Google Cloud project. [Learn how to confirm billing is enabled for your project](/billing/docs/how-to/modify-project) (/billing/docs/how-to/modify-project).

These instructions use [Cloud Shell](/shell/docs/) (/shell/docs/) to enter `gcloud` commands that deploy or configure your Google Cloud resources. Cloud Shell is accessed through the Cloud Console in your browser.

Cloud Shell runs on a VM that Google Cloud provisions each time you start Cloud Shell. The first time you use Cloud Shell, Google Cloud also creates a persistent `$HOME` directory for you, which is restored each time you open Cloud Shell.

The provisioned VM includes the latest [Cloud SDK](#) (`/sdk/docs/`), which provides the `gcloud` command-line interface. Therefore, the `gcloud` commands that you use in Cloud Shell are the same as those you would use in a locally installed instance of the Cloud SDK.

If you have the Cloud SDK installed, you can issue the `gcloud` commands that are used in these instructions from your local machine. However, with a locally installed Cloud SDK you must always make sure that you are using the latest version of the Cloud SDK.

Whether you use Cloud Shell or Cloud SDK, you can set and change the properties of your `gcloud` command environment and save them as a *configuration*. Configurations are collections of key-value pairs that influence the behavior of the `gcloud` commands.

Some basic actions you can take with a configuration in Cloud Shell include:

- Initialize a configuration with `gcloud init`.
- Check the settings of your current `gcloud` configuration with `gcloud config list`.
- Change the Google Cloud project you are working in with `gcloud config set project [PROJECT_ID]` where `[PROJECT_ID]` represents your Google Cloud project.
- Set a default region with `gcloud config set compute/region [REGION]` where `[REGION]` represents a Google Cloud region.
- Set a default zone with `gcloud config set compute/zone [ZONE]` where `[ZONE]` represents a Google Cloud zone.
- Create a new configuration with `gcloud config configurations create [NAME]` where `[NAME]` represents the name for the configuration.

For more information about working with configurations, see [Managing Cloud SDK configurations](#) (`/sdk/docs/configurations`).

For security purposes, create a new network. You can control who has access by adding firewall rules or by using another access control method.

If your project has a default VPC network, don't use it. Instead, create your own VPC network so that the only firewall rules in effect are those that you create explicitly.

During deployment, VM instances typically require access to the internet to download Google's monitoring agent. If you are using one of the SAP-certified Linux images that are available from Google Cloud, the VM instance also requires access to the internet in order to register the license and to access OS vendor repositories. A configuration with a NAT gateway and with VM network tags supports this access, even if the target VMs do not have external IPs.

To set up networking:

1. Go to Cloud Shell.

[Go to Cloud Shell](https://console.cloud.google.com/?cloudshell=true) (<https://console.cloud.google.com/?cloudshell=true>)

2. To create a new network in the custom subnetworks mode, run:

where `[YOUR_NETWORK_NAME]` is the name of the new network. The network name can contain only lowercase characters, digits, and the dash character (-).

Specify `--subnet-mode custom` to avoid using the default auto mode, which automatically creates a subnet in each Compute Engine region. For more information, see [Subnet creation mode](#) (</vpc/docs/vpc#subnet-ranges>).

3. Create a subnetwork, and specify the region and IP range:

where:

- `[YOUR_SUBNETWORK_NAME]` is the new subnetwork.
- `[YOUR_NETWORK_NAME]` is the name of the network you created in the previous step.
- `[REGION]` is the region where you want the subnetwork.
- `[YOUR_RANGE]` is the IP address range, specified in CIDR format (https://wikipedia.org/wiki/Classless_Inter-Domain_Routing), such as `10.1.0.0/24`. If you plan to add more than one subnetwork, assign non-overlapping CIDR IP ranges for each

subnetwork in the network. Note that each subnetwork and its internal IP ranges are mapped to a single region.

4. Optionally, repeat the previous step and add additional subnetworks.

If you intend to create one or more VMs that will not have public IP addresses, you must create a NAT gateway so that your VMs can access the Internet to download Google's monitoring agent.

If you intend to assign an external public IP address to your VM, you can skip this step.

Important: Do not remove the public IP addresses from your new VMs until after the installation of your SAP software is complete and validated.

To create a NAT gateway:

1. Create a VM to act as the NAT gateway in the subnet you just created:

where:

- [YOUR_VM_NAME] is the name of the VM you are creating that want to use for the NAT gateway.
- [YOUR_ZONE] is the zone where you want the VM.
- [YOUR_IMAGE_FAMILY] and [YOUR_IMAGE_PROJECT] specify the image you want to use (/compute/docs/images#os-compute-support) for the NAT gateway.
- [YOUR_MACHINE_TYPE] is any supported machine type. If you expect high network traffic, choose a machine type with that has at least eight virtual CPUs.
- [YOUR_SUBNETWORK_NAME] is the name of the subnetwork where you want the VM.

- [YOUR_VM_TAG] is a tag that is applied to the VM you are creating. If you use this VM as a bastion host, this tag is used to apply the related firewall rule only to this VM.

2. Create a route that is tagged so that traffic passes through the NAT VM instead of the default Internet gateway:

where:

- [YOUR_ROUTE_NAME] is the name of the route you are creating.
- [YOUR_NETWORK_NAME] is the network you created.
- [YOUR_VM_NAME] is the VM you are using for your NAT gateway.
- [YOUR_ZONE] is the zone where the VM is located.
- [YOUR_TAG_NAME] is the tag on the route that directs traffic through the NAT VM.

3. If you also want to use the NAT gateway VM as a bastion host, run the following command. This command creates a firewall rule that allows inbound SSH access to this instance from the Internet:

where:

- [YOUR_NETWORK_NAME] is the network you created.
- [YOUR_VM_TAG] is the tag you specified when you created the NAT gateway VM. This tag is used so this firewall rule applies only to the VM that hosts the NAT gateway, and not to all VMs in the network.

The following instructions use Deployment Manager to deploy a VM instance with Linux and all of the persistent disks that SAP MaxDB requires.

In these instructions, you define the resource options for your installation in a Deployment Manager configuration file template.

Deployment Manager treats all of the resources that are created for your SAP system as a single entity called a *deployment*. You can view and work with all of the deployments for your project on the [Deployments](https://console.cloud.google.com/dm/deployments) (https://console.cloud.google.com/dm/deployments) page in the Cloud Console.

Be aware of the following behaviors when using Deployment Manager:

- Deleting a deployment deletes all of the resources associated with the deployment, including the VMs, the persistent disks, and any SAP systems that are installed on the VM.
- By default, Deployment Manager uses the **ACQUIRE** resource creation policy. If you specify a VM name that is already in use by another VM in your project, Deployment Manager doesn't create a new VM, but instead adds the existing VM to your new deployment. If your original VM was created by a previous run of Deployment Manager, the VM is associated with two deployments.

If you then delete the new deployment, the acquired VM is deleted from the deployment that originally created it. To avoid such a scenario, either set the Deployment Manager resource policy to **CREATE**, or make sure that you use unique resource names in your new deployment.

For information about the policies you can use while creating resources with Deployment Manager and how to specify them, see the [Deployment Manager documentation](#) (/deployment-manager/docs/deployments/#create-policy).

1. Open the Cloud Shell.

[Go to the Cloud Shell](https://console.cloud.google.com/?cloudshell=true) (https://console.cloud.google.com/?cloudshell=true)

2. Download the template.yaml configuration file template to your working directory by entering the following command in the Cloud Shell:

3. Optionally, rename the template.yaml file to identify the configuration it defines.

4. Open the template.yaml file in the Cloud Shell code editor.

To open the Cloud Shell code editor, click the pencil icon in the upper right corner of the Cloud Shell terminal window.

- In the `template.yaml` file, update the following property values by replacing the brackets and their contents with the values for your installation.

Property	Data type	Description
<code>instanceName</code>	String	The name of the VM instance on which SAP MaxDB will be installed. The name must be 13 characters or less and be specified in lowercase letters, numbers, or hyphens.
<code>instanceType</code>	String	The type of Compute Engine virtual machine (/compute/docs/machine-types) on which SAP MaxDB will be installed. Specify a machine type with two or more vCPUs. For example, <code>n1-standard-4</code> . If you are running SAP MaxDB with SAP NetWeaver on the same VM, select a machine type that includes enough CPUs and memory to support both systems. See the SAP NetWeaver Planning Guide (/solutions/sap/docs/netweaver-planning-guide#machine_types).
<code>zone</code>	String	The zone in which you are deploying your SAP MaxDB. It must be in the same region that you selected for your subnetwork.
<code>subnetwork</code>	String	The name of the subnetwork that you created in a previous step. If you are deploying to a shared VPC, specify this value as <code>[SHAREDVPC_PROJECT]/[SUBNETWORK]</code> . For example, <code>myproject/network1</code> .
<code>linuxImage</code>	String	The name of the Linux operating- system image or image family that you are using with SAP MaxDB. To specify an image family, add the prefix <code>family/</code> to the family name. For example, <code>family/rhel-7</code> or <code>family/sles-12-sp3-sap</code> . To specify a specific image, specify only the image name. For the list of available image families, see the Images (https://console.cloud.google.com/compute/images) page in the Cloud console.
<code>linuxImageProject</code>	String	The Google Cloud project that contains the image you are going to use. This project might be your own project or a Google Cloud image project, such as <code>rhel-sap-cloud</code> or <code>suse-sap-cloud</code> . For a list of GCP image projects, see the Images (/compute/docs/images) page in the Compute Engine documentation.
<code>maxdbSID</code>	String	The database instance ID.
<code>maxdbRootSize</code>	Integer	The size in GB of <code>/sapdb</code> , which is the root directory of the database instance. The minimum and default values for <code>maxdbRootSize</code> are both 8 GB.
<code>maxdbDataSize</code>	Integer	The size of <code>/sapdb/[DBSID]/sapdata</code> , which holds the database data files. The minimum and default values for <code>maxdbDataSize</code> are both 30 GB.
<code>maxdbLogSize</code>	Integer	The size of <code>/sapdb/[DBSID]saplog</code> , which holds the database transaction logs. The minimum and default values for <code>maxdbLogSize</code> are both 8 GB.

Property	Data type	Description
maxdbBackupSizeInteger		The size of the <code>/maxdbbackup</code> volume. This property is optional. If set to 0 or omitted, no disk is created.
maxdbDataSSD	boolean	Specifies whether the data drive uses an SSD persistent disk (Yes) or an HDD persistent disk (No). Yes is the default.
maxdbLogSSD	boolean	Specifies whether the log drive uses an SSD persistent disk (Yes) or an HDD persistent disk (No). Yes is the default. SSD is recommended for the log drive.
ursapSize	Integer	Required only if you are installing SAP MaxDB to run with SAP NetWeaver on the same VM instance.
sapmntSize	Integer	Required only if you are installing SAP MaxDB to run with SAP NetWeaver on the same VM instance.
swapSize	Integer	Required only if you are installing SAP MaxDB to run with SAP NetWeaver on the same VM instance.

The following configuration file creates a VM that is configured to run both the SAP MaxDB database server and SAP NetWeaver. The configuration file directs Deployment Manager to deploy an n1-standard-16 VM that is running a SLES 12 SP3 operating system. The VM includes all of the directories that are required to run SAP MaxDB with NetWeaver.

6. Create the VM instance:

The above command invokes the Deployment Manager, which deploys the VM according to the specifications in your template.yaml file. The process might take a few minutes. To check the progress of your deployment, follow the steps in the next section.

The following steps use Stackdriver Logging, which might incur charges. For more information, see [Stackdriver pricing \(/stackdriver/pricing_v2\)](https://cloud.google.com/stackdriver/pricing).

1. Open Stackdriver Logging to check for errors and monitor the progress of the installation.

[Go to Logging \(https://console.cloud.google.com/logs/viewer\)](https://console.cloud.google.com/logs/viewer)

2. On the **Resources** tab, select **Global** as your logging resource. If 'INSTANCE DEPLOYMENT COMPLETE' is displayed for a VM, Deployment Manager processing is complete for the VM.

▶		2018-08-02 17:21:11.475 PDT	xmp-maxdb-lin	Deployment	"--- /maxdbbackup"	⋮
▶		2018-08-02 17:21:12.832 PDT	xmp-maxdb-lin	Deployment	"Creating file systems for SAP NetWeaver"	⋮
▶		2018-08-02 17:21:13.284 PDT	xmp-maxdb-lin	Deployment	"--- /usr/sap"	⋮
▶		2018-08-02 17:21:17.139 PDT	xmp-maxdb-lin	Deployment	"--- /sapmnt"	⋮
▶		2018-08-02 17:21:21.216 PDT	xmp-maxdb-lin	Deployment	"--- swap volume"	⋮
▶		2018-08-02 17:21:21.838 PDT	xmp-maxdb-lin	Deployment	"INSTANCE DEPLOYMENT COMPLETE"	⋮

3. After the SAP MaxDB system deploys, connect to your VM by using SSH. From the Compute Engine [VM instances page \(https://console.cloud.google.com/compute/instances\)](https://console.cloud.google.com/compute/instances), you can click the SSH button for your VM instance, or you can use your preferred SSH method.

<input type="checkbox"/>	Name ^	Zone	Internal IP	External IP	Connect
<input type="checkbox"/>	xmp-maxdb-lin	us-west1-b	xmp-maxdb-lin (10.2.0.3) (nic0)	35.203.162.107	SSH ⋮

4. Change to the root user.

5. At the command prompt, enter `df -h`. Ensure that you see output similar to the following, such as the `//sapdb/[DBSID]/sapdata` directory.

```
xmp-maxdb-lin:~ # df -h
Filesystem                Size      Used Avail  Use% Mounted on
devtmpfs                  30G         0    30G   0% /dev
tmpfs                     45G         0    45G   0% /dev/shm
tmpfs                     30G        26M    30G   1% /run
tmpfs                    30G         0    30G   0% /sys/fs/cgroup
/dev/sda3                 30G        2.2G    26G   8% /
/dev/sda2                 200M        660K   200M   1% /boot/efi
/dev/mapper/vg_maxdbsid-vol 10G         33M    10G   1% /sapdb/MD1
/dev/mapper/vg_maxdbdata-vol 50G         33M    50G   1% /sapdb/MD1/sapdata
/dev/mapper/vg_maxdblog-vol 30G         33M    30G   1% /sapdb/MD1/saplog
/dev/mapper/vg_maxdbbackup-vol 100G        33M   100G   1% /maxdbbackup
/dev/mapper/vg_usrsap-vol 15G         38M    14G   1% /usr/sap
/dev/mapper/vg_sapmnt-vol 15G         38M    14G   1% /sapmnt
```

6. If you are installing SAP MaxDB on the same VM as SAP NetWeaver, confirm that the swap directory was created by entering the following command:

You should see results similar to the following example:

```
xmp-maxdb-lin:~ # cat /proc/meminfo | grep Swap
SwapCached:          0 kB
SwapTotal:           25161724 kB
SwapFree:            25161724 kB
```

If any of the validation steps show that the installation failed:

1. Correct the error.
2. On the [Deployments](https://console.cloud.google.com/dm/deployments) (https://console.cloud.google.com/dm/deployments) page, delete the deployment to clean up the VMs and persistent disks from the failed installation.
3. Rerun your deployment.

Now that your operating system is configured, you can install your SAP MaxDB database. MaxDB is typically installed with the SAP product that it is integrated with.

To install SAP MaxDB on your VM:

1. Establish an SSH connection to your Linux-based VM.
2. Download the SAP Software Provisioning Manager (SWPM), the SAP product installation media, and MaxDB installation Media according to SAP installation guides.
3. Install your SAP product and SAP MaxDB database according to the SAP installation guides for your SAP product. For additional guidance, see the [SAP MaxDB documentation](https://help.sap.com/viewer/9d62a6d2a95149718db6b65b7ffa2052/113/en-US) (<https://help.sap.com/viewer/9d62a6d2a95149718db6b65b7ffa2052/113/en-US>).

SAP provides additional installation information in [SAP Note 1020175 - FAQ: SAP MaxDB installation, upgrade or applying a patch](https://launchpad.support.sap.com/#/notes/0001020175) (<https://launchpad.support.sap.com/#/notes/0001020175>).

Before using your SAP MaxDB instance, we recommend that you perform the following post-deployment steps:

1. Update your SAP MaxDB software with the latest patches, if available.
2. Install any additional components.
3. Configure and back up your new SAP MaxDB database.

For more information, see [SAP MaxDB Database Administration](https://help.sap.com/viewer/2c2effc99b6746019aeb1af52ad59f5d/113/en-US/1a32c42a0fcf40b5b95ff7040a60aa41.html) (<https://help.sap.com/viewer/2c2effc99b6746019aeb1af52ad59f5d/113/en-US/1a32c42a0fcf40b5b95ff7040a60aa41.html>).

This section contains information about how to correct common issues.

If you are having issues connecting to your VM through SSH, ensure that you have created a firewall rule to open port 22 on the GCP network you are using.

For other possible issues, see [Known issues](/compute/docs/ssh-in-browser#known_issues) (/compute/docs/ssh-in-browser#known_issues) for SSH from the browser.

