# SAP HANA deployment guide

This deployment guide shows you how to deploy a SAP HANA system on Google Cloud by using Cloud Deployment Manager (/deployment-manager/docs/fundamentals) and a configuration file template to define your installation. The guide helps you configure Compute Engine virtual machines (VMs) and persistent disks, as well as the Linux operating system, to achieve the best performance for your SAP HANA system. The Deployment Manager template incorporates best practices from both Compute Engine and SAP.

Use this guide to deploy either a single-host scale-up or a multi-host scale-out SAP HANA system that does not include standby hosts.

If you need to include SAP HANA automatic host failover, use the SAP HANA Scale-Out System with Host Auto-Failover Deployment Guide (/solutions/sap/docs/sap-hana-ha-scaleout-deployment-guide) instead.

If you need to deploy a Linux high-availability cluster for a single-host SAP HANA system, use the SAP HANA High Availability Cluster on SLES Deployment Guide (/solutions/sap/docs/sap-hana-ha-deployment-guide).

## Prerequisite tasks

If you don't already have them set up, you need to create a Google Cloud account and project. You also need to set up Virtual Private Cloud networking, as well as a method for controlling access to your VMs. Finally, you need to load the SAP HANA installation media into a Cloud Storage bucket.

id unintentionally exposing your VM instance to the internet, follow these recommendations:

Use a NAT gateway.

Create firewall rules (/vpc/docs/using-firewalls#creating_firewall_rules) that block all external access that yo don't require.

When you create your VMs:

- Specify a network tag for each VM for use in routing and firewall rules. If you use the Deployment Mana templates that Google Cloud provides, specify a tag with **`networkTag: [TAG]`**.

- Create the VMs without an external IP. If you use the Deployment Manager templates that Google Clou provides, specify **`publicIP: No`**.

## Setting up your Google account

A Google account is required to work with GCP.

1. Sign up for a Google account (https://accounts.google.com/SignUp) if you don't already have one.
2. Log in to the Google Cloud Console, and create a new project (https://console.cloud.google.com/project).
3. Enable your billing account (/resource-manager/docs/quickstart#create_a_billing_account).
4. Configure SSH keys so that you are able to use them to SSH into your Compute Engine instances. Use the **`gcloud`** command-line tool to create a new SSH key (/compute/docs/instances/adding-removing-ssh-keys#createsshkeys), or, if you already have an existing SSH key, use the tool to format your existing SSH keys (/compute/docs/instances/adding-removing-ssh-keys#sshkeyformat).
5. Use the **`gcloud`** command-line tool or Cloud Console to add the SSH keys (/compute/docs/instances/adding-removing-ssh-keys#project-wide) to your project metadata. This allows you to access any Compute Engine instance created within this project, except for instances that explicitly disable project-wide SSH keys.

## Creating a network

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)

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

   ```
   gcloud compute networks create [YOUR_NETWORK_NAME] --subnet-mode custom
   ```

   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:

   ```
   gcloud compute networks subnets create [YOUR_SUBNETWORK_NAME] \
           --network [YOUR_NETWORK_NAME] --region [YOUR_REGION] --range [YOUR_I
   ```

   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.

## Setting up a NAT gateway

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.

**tant**: Do not remove the public IP addresses from your new VMs until after the installation of your SAP softwa ete and validated.

To create a NAT gateway:

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

```
gcloud compute instances create [YOUR_VM_NAME] --can-ip-forward \
        --zone [YOUR_ZONE]  --image-family [YOUR_IMAGE_FAMILY] \
        --image-project [YOUR_IMAGE_PROJECT] \
        --machine-type=[YOUR_MACHINE_TYPE] --subnet [YOUR_SUBNETWORK_NAME]
        --metadata startup-script="sysctl -w net.ipv4.ip_forward=1; iptables
        -t nat -A POSTROUTING -o eth0 -j MASQUERADE" --tags [YOUR_VM_TAG]
```

   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:

```
gcloud compute routes create [YOUR_ROUTE_NAME] \
        --network [YOUR_NETWORK_NAME] --destination-range 0.0.0.0/0 \
        --next-hop-instance [YOUR_VM_NAME] --next-hop-instance-zone \
        [YOUR_ZONE] --tags [YOUR_TAG_NAME] --priority 800
```

   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:

```
gcloud compute firewall-rules create allow-ssh --network [YOUR_NETWORK_NAME
```

   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.

## Adding firewall rules

By default, an *implied firewall rule* blocks incoming connections from outside your Virtual Private Cloud (VPC) network. To allow incoming connections, set up a firewall rule for your VM.

After an incoming connection is established with a VM, traffic is permitted in both directions over that connection.

You can also create a firewall rule to allow external access to specified ports, or to restrict access between VMs on the same network. If the `default` VPC network type is used, some additional default rules also apply, such as the `default-allow-internal` rule, which allows connectivity between VMs on the same network on all ports.

Depending on the IT policy that is applicable to your environment, you might need to isolate or otherwise restrict connectivity to your database host, which you can do by creating firewall rules.

Depending on your scenario, you can create firewall rules to allow access for:

- The default SAP ports that are listed in TCP/IP of All SAP Products (https://help.sap.com/viewer/575a9f0e56f34c6e8138439eefc32b16/2.0/en-US/616a3c0b1cc748238de9c0341b15c63c.html) .
- Connections from your computer or your corporate network environment to your Compute Engine VM instance. If you are unsure of what IP address to use, talk to your company's network administrator.
- Communication between VMs when, for example, your database server and application server are running on different VMs. To enable communication between VMs, you must create a firewall rule to allow traffic that originates from the subnetwork.
- SSH connections to your VM instance, including SSH from the browser (https://cloud.google.com/compute/docs/ssh-in-browser).
- Connection to your VM by using a third-party tool in Linux. Create a rule to allow access for the tool through your firewall.

The following procedure is a simplified version of the instructions for creating firewall rules. For more detailed instructions, see the Virtual Private Cloud documentation (/vpc/docs/using-firewalls#creating_firewall_rules).

To create a firewall rule:

Consolegcloud (#gcloud)

1. In the Cloud Console, go to the **Firewall rules** page.

   OPEN FIREWALL RULES (https://console.cloud.google.com/networking/firewalls/list)

2. At the top of the page, click **Create firewall rule**.
   - In the **Network** field, select the network where your VM is located.
   - In the **Targets** field, specify the resources on Google Cloud that this rule applies to. For example, specify **All instances in the network**. Or to to limit the rule to specific instances on Google Cloud, enter tags in **Specified target tags**.
   - In the **Source filter** field, select one of the following:
     - **IP ranges** to allow incoming traffic from specific IP addresses. Specify the range of IP addresses in the **Source IP ranges** field.
     - **Subnets** to allow incoming traffic from a particular subnetwork. Specify the subnetwork name in the following **Subnets** field. You can use this option to allow access between the VMs in a 3-tier or scaleout configuration.
   - In the **Protocols and ports** section, select **Specified protocols and ports** and enter `tcp:[PORT_NUMBER]`.
3. Click **Create** to create your firewall rule.

## Creating a Cloud Storage bucket for the SAP HANA installation files

The installation files that contain the SAP HANA binaries must be stored in a Cloud Storage bucket before you can use Deployment Manager to install SAP HANA. Deployment Manager expects the files in the file formats provided by SAP. Depending on your version of SAP HANA, the file format might be a .zip file or .exe and .rar files.

To download the SAP HANA installation files, create a bucket, and upload the files to the bucket:

1. From SAP Software Downloads (http://support.sap.com/swdc), download all parts of the Linux x86_64 distribution of SAP HANA Platform Edition 1.0 or 2.0, as well as any applicable revision upgrades to your local drive.

   If your SAP Support Portal account does not allow access to the software and you believe that you should be entitled to the software, contact the SAP Global Support Customer Interaction Center (http://support.sap.com/contactus).

2. Use the Cloud Console to create a Cloud Storage bucket (/storage/docs/creating-buckets) for storing the SAP HANA installation files. Note that the bucket name must be unique across GCP.

★ **Important:** For now, do not extract the downloaded SAP HANA software. Just stage the files in a Cloud Storage bucket as-is. Extraction and installation is covered later.

   - During bucket creation, choose **Standard** for your storage class.

3. Configure bucket permissions. By default, as owner of the bucket, you have read-write access to the bucket. To give access to other members in your group or individual users, see Using IAM permissions (/storage/docs/access-control/using-iam-permissions).

4. In the Cloud Console, in the Cloud Storage bucket page, choose **Upload Files** to upload the SAP HANA software and any upgrade revision files to your bucket from your local media:

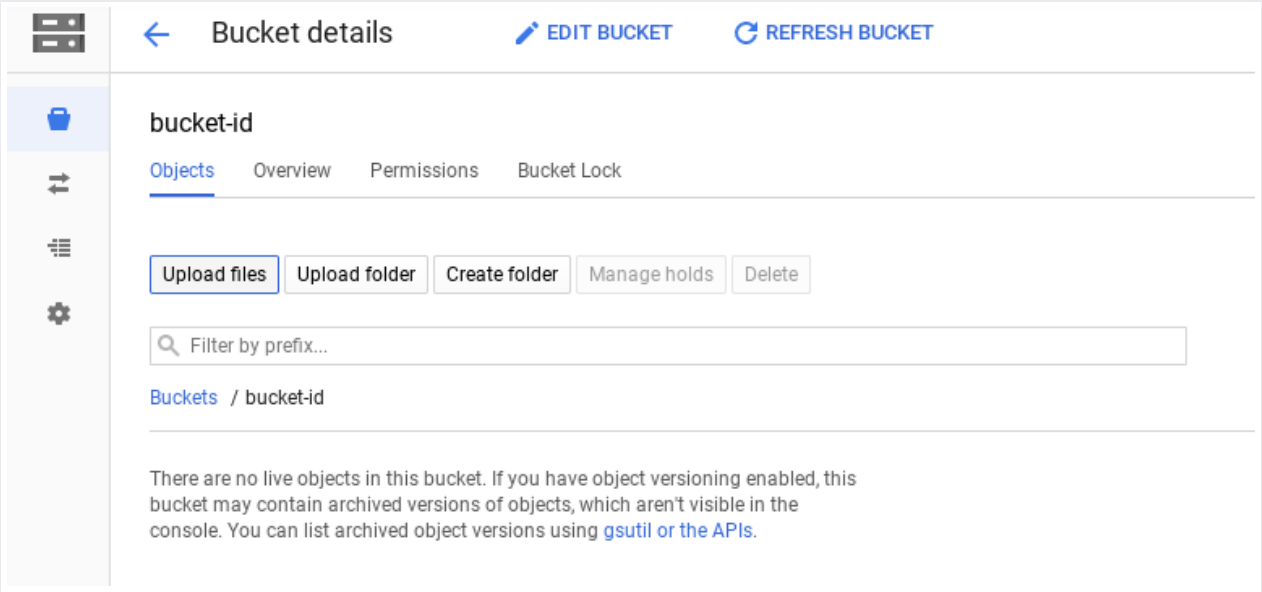


5. Note the name of the bucket that you uploaded the binaries to. You need to use it later when you install SAP HANA.

## Creating a VM with SAP HANA installed

The following instructions use the Deployment Manager to install SAP HANA on one or more VM instances with all of the persistent disks that SAP HANA requires. You define the values for

the installation in a Deployment Manager configuration file template.

Deployment Manager treats your SAP HANA system and all of the VMs, disks, and other resources that are created for the SAP HANA system as a single entity called a *deployment*. You can view all of the deployments for your GCP project on the Deployment Manager Deployments (https://console.cloud.google.com/dm/deployments) page.

Do not delete a deployment unless you need to delete everything associated with the deployment, including th ANA system, the VMs, and the persistent disks that were deployed with it. Deployments are a permanent entation of your SAP HANA installation on GCP.

The following instructions use Cloud Shell, but are generally applicable to the Cloud SDK.

1. Confirm that your current quotas for resources such as persistent disks and CPUs are sufficient for the SAP HANA system you are about to install. If your quotas are insufficient, deployment fails. For the SAP HANA quota requirements, see Pricing and quota considerations for SAP HANA (/solutions/sap/docs/sap-hana-planning-guide#costs).

   Go to the quotas page (https://console.cloud.google.com/iam-admin/quotas)

2. Open the Cloud Shell or, if you installed the Cloud SDK on your local workstation, open a terminal.

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

3. Download the `template.yaml` configuration file template to your working directory by entering the following command in the Cloud Shell or Cloud SDK:

   ```
   wget https://storage.googleapis.com/sapdeploy/dm-templates/sap_hana/templat
   ```

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

5. Open the `template.yaml` file in Cloud Shell code editor or, if you are using the Cloud SDK, the text editor of your choice.

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

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

   Some of the property values that you specify for the SAP HANA system, such as [SID] or [PASSWORD], are subject to rules that are defined by SAP. For more information, see the Parameter Reference in the SAP HANA Server Installation and Update Guide (https://help.sap.com/viewer/product/SAP_HANA_PLATFORM/2.0.02/en-US).

   If you want to create a VM instance without installing SAP HANA, delete all of the lines that begin with `sap_hana_`.

| Property | Data type | Description |
|---|---|---|
| instanceName | String | The name of the VM instance for the SAP HANA master host. The name must be specified in lowercase letters, numbers, or hyphens. If any other characters are used, such as "_" (underscore) or a capital letter, deployment fails. The VM instances for any worker hosts use the same name with a "w" and the host number appended to the name. |
| instanceType | String | The type of Compute Engine virtual machine that you need to run SAP HANA on. If you need a custom VM type, specify a predefined VM type (/solutions/sap/docs/sap-hana-planning-guide#vm_types) with a number of vCPUs that is closest to the number you need while still being larger. After deployment is complete, modify the number of vCPUs and the amount of memory. (/solutions/sap/docs/modifying_vm_configurations). |
| zone | String | The zone in which you are deploying your SAP HANA system to run. It must be in the region that you selected for your subnet. |
| subnetwork | String | The name of the subnetwork you created in a previous step. If you are deploying to a shared VPC, specify this value as ***[SHAREDVPC_PROJECT]***/***[SUBNETWORK]***. For example, `myproject/network1`. |

| Property | Data type | Description |
|---|---|---|
| linuxImage | String | The name of the Linux operating-system image or image family that you are using with SAP HANA. To specify an image family, add the prefix `family/` to the family name. For example, `family/rhel-7-4-sap` or `family/sles-12-sp2-sap`. 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. |
| linuxImageProject | 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 `rhel-sap-cloud` or `suse-sap-cloud`. For a list of GCP image projects, see the Images (/compute/docs/images) page in the Compute Engine documentation. |
| sap_hana_deployment_bucket | String | The name of the GCP storage bucket in your project that contains the SAP HANA installation and revision files that you uploaded in a previous step. Any upgrade revision files in the bucket are applied to SAP HANA during the deployment process. |
| sap_hana_sid | String | The SAP HANA system ID. The ID must consist of three alphanumeric characters and begin with a letter. All letters must be uppercase. |
| sap_hana_instance_number | Integer | The instance number, 0 to 99, of the SAP HANA system. The default is 0. |
| sap_hana_sidadm_password | String | The password for the operating system administrator. Passwords must be at least eight characters and include at least one uppercase letter, one lowercase letter, and one number. |
| sap_hana_system_password | String | The password for the database superuser. Passwords must be at least 8 characters and include at least one uppercase letter, one lowercase letter, and one number. |
| sap_hana_scaleout_nodes | Integer | The number of additional SAP HANA worker hosts that you need. The worker hosts are in addition to the primary SAP HANA instance. For example, if you specify 3, four SAP HANA instances are deployed in a scale-out cluster. |

| Property | Data type | Description |
|---|---|---|
| `networkTag` | String | Optional. A network tag that represents your VM instance for firewall or routing purposes. If you specify `publicIP: No` and do not specify a network tag, be sure to provide another means of access to the internet. |
| `publicIP` | Boolean | Optional. Determines whether a public IP address is added to your VM instance. The default is `Yes`.<br><br>★ **Note**: Do not specify **`No`** unless you have a NAT gateway configured with a network tag defined for the VM or you have provided the VM with another route to the internet. If there is no route to the internet, the installation fails. |

The following example shows a completed configuration file, which directs the Deployment Manager to deploy an n1-highmem-96 virtual machine with a scale-out HANA system that includes a master SAP HANA instance with three worker hosts. SAP HANA is running on a SLES 12 SP2 operating system.

```
imports:
- path: https://storage.googleapis.com/sapdeploy/dm-templates/sap_hana/sap_h

resources:
- name: sap_hana
  type: https://storage.googleapis.com/sapdeploy/dm-templates/sap_hana/sap_h
  properties:
    instanceName: example-vm
    instanceType: n1-highmem-96
    zone: us-central1-f
    subnetwork: default
    linuxImage: family/sles-12-sp2-sap
    linuxImageProject: suse-sap-cloud
    sap_hana_deployment_bucket: mybucketname
    sap_hana_sid: ABC
    sap_hana_instance_number: 00
    sap_hana_sidadm_password: Google123
    sap_hana_system_password: Google123
    sap_hana_scaleout_nodes: 3
```

7. Create the instances:

```
gcloud deployment-manager deployments create [DEPLOYMENT-NAME] --config [TEI
```

The above command invokes the Deployment Manager, which deploys the VMs, downloads the SAP HANA software from your storage bucket, and installs SAP HANA, all according to the specifications in your template.yaml file. The process takes approximately 10 to 15 minutes to complete. To check the progress of your deployment, follow the steps in the next section.

## Verifying deployment

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

★ **Note:** You might incur costs when completing this step in Cloud Logging. For more information, see Cloud Logging pricing (/stackdriver/pricing_v2).

Go to Cloud Logging (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 all VMs, Deployment Manager processing is complete and you can proceed to the next step.
    - If you see a quota error:
        a. On the IAM & admin Quotas (https://console.cloud.google.com/iam-admin/quotas) page, increase any of your quotas that do not meet the SAP HANA requirements that are listed in the SAP HANA Planning Guide (/solutions/sap/docs/sap-hana-planning-guide#quotas).
        b. On the Deployment Manager Deployments (https://console.cloud.google.com/dm/deployments) page, delete the deployment to clean up the VMs and persistent disks from the failed installation.
        c. Rerun the Deployment Manager.

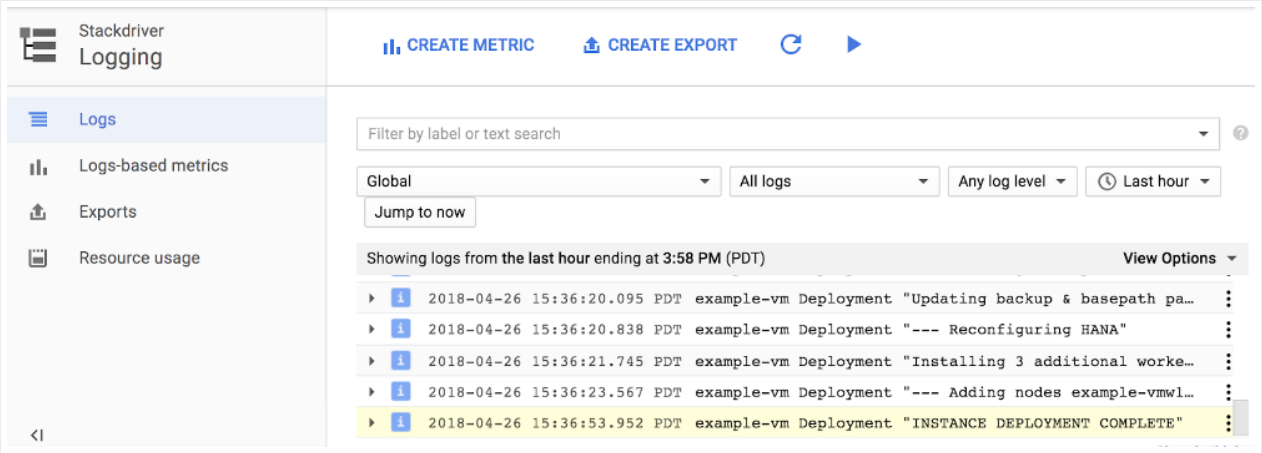


3. After the SAP HANA system deploys without errors, connect to your VM by using SSH. From the Compute Engine VM instances page (https://console.cloud.google.com/compute/instances), you can click the SSH button for your VM instance, or you can use your preferred SSH method.

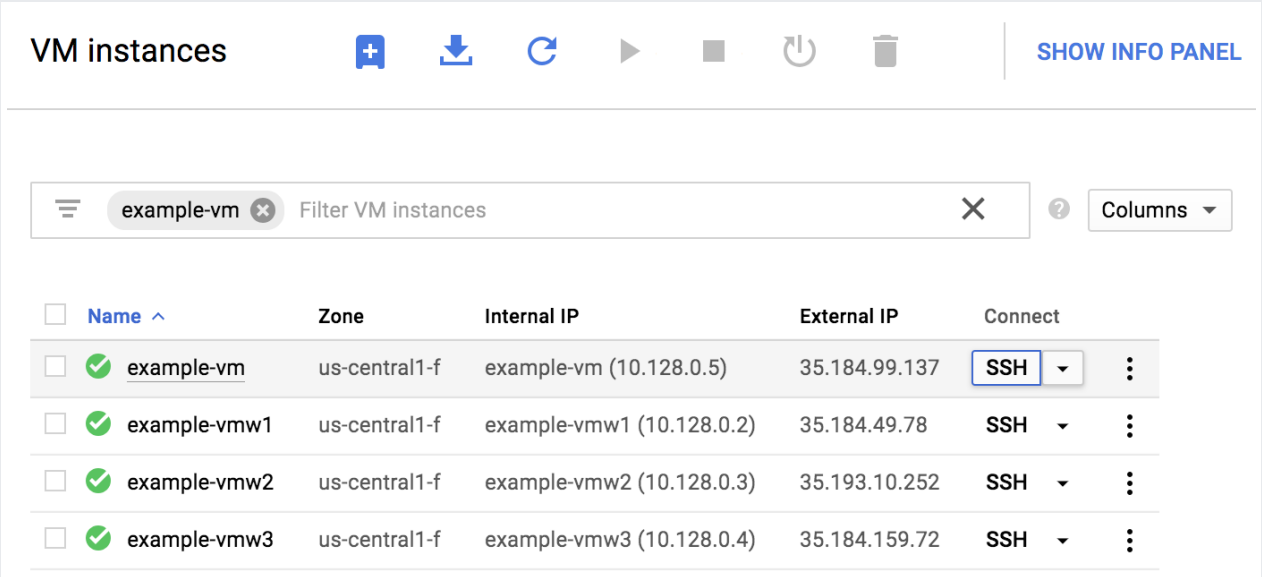


4. Change to the root user.

```
sudo su -
```

5. At the command prompt, enter `df -h`. Ensure that you see output similar to the following, such as the `/hana/data` directory.

```
saphana01:~ # df -h
Filesystem                          Size  Used Avail Use% Mounted on
/dev/sda1                            32G  1.7G   29G   6% /
devtmpfs                            206G     0  206G   0% /dev
tmpfs                               206G     0  206G   0% /dev/shm
tmpfs                               206G   18M  206G   1% /run
tmpfs                               206G     0  206G   0% /sys/fs/cgroup
/dev/mapper/vg_hana-data            1.0T  2.9G 1021G   1% /hana/data
/dev/mapper/vg_hana-log             256G  3.2G  253G   2% /hana/log
/dev/mapper/vg_hana-sap              32G  158M   32G   1% /usr/sap
/dev/mapper/vg_hanabackup-backup    3.0T  3.8G  3.0T   1% /hanabackup
/dev/mapper/vg_hana-shared          512G   35G  478G   7% /hana/shared
```

6. Change to the SAP admin user. Replace ***[SID]*** with the [SID] value that you specified in the configuration file template.

   ```
   su - [SID]adm
   ```

7. Ensure that SAP HANA services, such as `hdbnameserver`, `hdbindexserver`, and others, are running on the instance by entering the following command:

   ```
   HDB info
   ```

If any of the validation steps show that the installation failed, resolve any errors, delete the deployment from the Deployments (https://console.cloud.google.com/dm/deployments) page, and then recreate the instances, as described in the last step of the previous section (#creating_a_vm_with_sap_hana_installed).

## Completing the NAT gateway installation

If you created a NAT gateway, complete the following steps.

do not delete the external IP address from the VM instances that are running SAP HANA until you have compl
ıg both a network and a NAT gateway. After the external IP addresses are deleted, you can access the VM
ces only through the NAT gateway.

1. Add tags to all instances, including the worker hosts:

```
export INSTANCE_NAME="[YOUR_VM_NAME]"
export NETWORK_NAME="[YOUR_NETWORK_NAME]"
export ZONE="[YOUR_ZONE]"
export TAG="[YOUR_TAG_TEXT]"
```

```
gcloud compute instances add-tags "$INSTANCE_NAME" --tags="$TAG" --zone=$ZOI
```

```
gcloud compute instances add-tags "$INSTANCE_NAME"w1 --tags="$TAG" --zone=$Z
```

```
gcloud compute instances add-tags "$INSTANCE_NAME"w2 --tags="$TAG" --zone=$Z
```

```
gcloud compute instances add-tags "$INSTANCE_NAME"w3 --tags="$TAG" --zone=$Z
```

2. Delete external IPs:

```
gcloud compute instances delete-access-config "$INSTANCE_NAME" --access-con
```

```
gcloud compute instances delete-access-config "$INSTANCE_NAME"w1 --access-c
```

```
gcloud compute instances delete-access-config "$INSTANCE_NAME"w2 --access-c
```

```
gcloud compute instances delete-access-config "$INSTANCE_NAME"w3 --access-c
```

# Installing SAP HANA Studio on a Compute Engine Windows VM

You can connect from a SAP HANA instance outside of Google Cloud or from an instance on Google Cloud. To do so, you might need to enable network access to the target VMs from within SAP HANA Studio.

To install SAP HANA Studio on a Windows VM on Google Cloud, use the following procedure.

1. Use the Cloud Shell to invoke the following commands.

   OPEN THE CLOUD SHELL (https://console.cloud.google.com/?cloudshell=true)

   ```
   export NETWORK_NAME="[YOUR_NETWORK_NAME]"
   export REGION="[YOUR_REGION]"
   export ZONE="[YOUR_ZONE]"
   export SUBNET="[YOUR_SUBNETWORK_NAME]"
   export SOURCE_IP_RANGE="[YOUR_WORKSTATION_IP]"
   ```

   ```
   gcloud compute instances create saphanastudio --zone=$ZONE \
   --machine-type=n1-standard-2  --subnet=$SUBNET --tags=hanastudio \
   --image-family=windows-2016  --image-project=windows-cloud \
   --boot-disk-size=100 --boot-disk-type=pd-standard \
   --boot-disk-device-name=saphanastudio
   ```

   ```
   gcloud compute firewall-rules create ${NETWORK_NAME}-allow-rdp \
   --network=$NETWORK_NAME --allow=tcp:3389 --source-ranges=$SOURCE_IP_RANGE \
   --target-tags=hanastudio
   ```

   The above commands set variables for the current Cloud Shell session, create a Windows server in the subnetwork that you created earlier, and create a firewall rule that allows access from your local workstation to the instance through the Remote Desktop Protocol (https://wikipedia.org/wiki/Remote_Desktop_Protocol) (RDP).

2. Install SAP HANA Studio (https://help.sap.com/viewer/a2a49126a5c546a9864aae22c05c3d0e/2.0.03/en-US) on this server.

a. Upload (https://cloud.google.com/storage/docs/uploading-objects) the SAP HANA Studio installation files and the SAPCAR extraction tool to a Cloud Storage bucket in your Google Cloud project.

b. Connect to the new Windows VM by using RDP or your preferred method.

c. In Windows, with administrator permissions, open the Google Cloud SDK Shell or other command-line interface.

d. Copy the SAP HANA Studio installation files and the SAPCAR extraction tool from the storage bucket to the VM by entering the `gsutil cp` command (https://cloud.google.com/storage/docs/gsutil/commands/cp) in the command interface. For example:

```
gsutil cp gs://[SOURCE_BUCKET]/IMC_STUDIO2_232_0-80000323.SAR C:\[TARGE
gsutil cp gs://[SOURCE_BUCKET]/SAPCAR_1014-80000938.EXE C:\[TARGET_DIRE
```

e. Change the directory to your target directory.

```
cd C:\[TARGET_DIRECTORY]
```

f. Run the SAPCAR program to extract the SAP HANA Studio installation file.

```
SAPCAR_1014-80000938.EXE -xvf IMC_STUDIO2_232_0-80000323.SAR
```

g. Run the extracted `hdbinst` program to install SAP HANA Studio.

## Setting up the Google monitoring agent for SAP HANA

Optionally, you can set up the Google monitoring agent for SAP HANA, which collects metrics from SAP HANA and sends them to Cloud Monitoring (/monitoring/docs). Cloud Monitoring allows you to create dashboards for your metrics, set up custom alerts based on metric thresholds, and more. For more information on setting up and configuring the Google

monitoring agent for SAP HANA, see the SAP HANA Monitoring Agent User Guide (/solutions/sap/docs/sap-hana-monitoring-agent-user-guide).

## Connecting to SAP HANA

Note that because these instructions don't use an external IP for SAP HANA, you can only connect to the SAP HANA instances through the bastion instance using SSH or through the Windows server through SAP HANA Studio.

- To connect to SAP HANA through the bastion instance, connect to the bastion host, and then to the SAP HANA instance(s) by using an SSH client of your choice.
- To connect to the SAP HANA database through SAP HANA Studio, use a remote desktop client to connect to the Windows Server instance. After connection, manually install SAP HANA Studio (https://help.sap.com/hana/SAP_HANA_Studio_Installation_Update_Guide_en.pdf) and access your SAP HANA database.

## Performing post-deployment tasks

Before using your SAP HANA instance, we recommend that you perform the following post-deployment steps. For more information, see SAP HANA Installation and Update Guide (http://help.sap.com/hana/SAP_HANA_Server_Installation_Guide_en.pdf).

1. Install your permanent SAP HANA license. If you do not, SAP HANA might go into *database lockdown* after the temporary license expires.

   For more information from SAP about managing your SAP HANA licenses, see License Keys for the SAP HANA Database (https://help.sap.com/viewer/6b94445c94ae495c83a19646e7c3fd56/2.0.03/en-US/9fd02a6c45eb485b9b3b0bc845586a6a.html)
   .
2. Update the SAP HANA software with the latest patches.
3. Install any additional components such as Application Function Libraries (AFL) or Smart Data Access (SDA).

4. Configure and backup your new SAP HANA database. For more information, see the SAP HANA operations guide (/solutions/sap/docs/sap-hana-operations-guide#backup_and_recovery).

## What's next

- If you need to use NetApp Cloud Volumes Service for Google Cloud instead of persistent disks for your SAP HANA directories, see the NetApp Cloud Volumes Service deployment information in the SAP HANA planning guide (/solutions/sap/docs/sap-hana-planning-guide#netapp_cvs).
- For more information about VM administration of and monitoring, see the SAP HANA Operations Guide (/solutions/sap/docs/sap-hana-operations-guide).



Last updated 2020-12-11 UTC.