

This deployment guide shows you how to deploy the Google Cloud resources for an IBM Db2 for Linux, Unix, and Windows system that supports SAP applications.

For details on planning your deployment, see the [IBM Db2 for SAP planning guide](/solutions/sap/docs/sap-ibm-db2-planning-guide) (/solutions/sap/docs/sap-ibm-db2-planning-guide).

If you don't already have a Google Cloud project with billing enabled, you must create one before you can deploy a virtual machine (VM) for your IBM Db2 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 \(/sdk/docs/\)](/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 \(/sdk/docs/configurations\)](/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) (/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) (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.

By default, incoming connections from outside your Google Cloud network are blocked. To allow incoming connections, set up a firewall rule for your VM. Firewall rules regulate only new incoming connections to a VM. After a connection is established with a VM, traffic is permitted in both directions over that connection.

You can create a firewall rule to allow access to specified ports, or to allow access between VMs on the same subnetwork.

Create firewall rules to allow access for such things as:

- The default ports used by SAP NetWeaver, as documented in [TCP/IP Ports of All SAP Products](https://help.sap.com/viewer/575a9f0e56f34c6e8138439eefc32b16/2.0/en-US/616a3c0b1cc748238de9c0341b15c63c.html) (<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 admin.
- Communication between VMs in a 3-tier or scaleout configuration. For example, if you are deploying a 3-tier system, you will have at least 2 VMs in your subnetwork: the VM for SAP NetWeaver, and another VM for the database server. To enable communication between the two VMs, you must create a firewall rule to allow traffic that originates from the subnetwork.

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

To create a firewall rule:

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

[Open Firewall Rules page](https://console.cloud.google.com/networking/firewalls/list) (<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, select **All instances in the network**.
- 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 specify **tcp: [PORT\_NUMBER]** ;.

3. Click **Create** to create your firewall rule.

When you create a VM, you can specify several options, including the operating system, region, machine type, and persistent disks. You must also specify a startup script to install Google's monitoring agent on the VM.

To create a VM:

1. In the Cloud Console, go to the **Images** page:

[Go to the Images page](https://console.cloud.google.com/compute/images) (<https://console.cloud.google.com/compute/images>)

2. To use a public image, choose a supported image from the list. SAP has certified the following images for IBM Db2 on Google Cloud:

- Red Hat Enterprise Linux 7.4
- SUSE Linux Enterprise Server 12 SP2 or higher
- Windows Server 2012 R2 and higher

3. Click **Create Instance**.

4. Enter a name for your VM. Limit the name to 13 characters, as this is the maximum supported by SAP. For more information, see SAP Note [611361: Hostnames of SAP servers](https://launchpad.support.sap.com/#/notes/611361) (<https://launchpad.support.sap.com/#/notes/611361>).

5. Select the zone for your VM based on the location of your internal resources and users, and the machine type you want to use.

See the following guides and SAP Notes for more details on the zones supported for SAP Netweaver:

- [SAP Netweaver on Google Cloud planning guide](/solutions/sap/docs/netweaver-planning-guide#planning_regions_zones) ([/solutions/sap/docs/netweaver-planning-guide#planning\\_regions\\_zones](/solutions/sap/docs/netweaver-planning-guide#planning_regions_zones))
- SAP Note [2446441 - Linux on Google Cloud \(IaaS\): Adaptation of your SAP License](https://launchpad.support.sap.com/#/notes/2446441) (<https://launchpad.support.sap.com/#/notes/2446441>)

★ **Note:** The list of zones displayed in the Cloud Console is randomized within each region to encourage use across multiple zones.

6. Change the machine type to one of the supported machine types. To see the supported machine types and their persistent-disk limitations, see the [IBM Db2 for SAP planning guide](#)

(/solutions/sap/docs/sap-ibm-db2-planning-guide).

7. Optionally, in the **Boot disk** section, click **Change**. In the **Boot disk** dialog, you can configure the size and type of your boot disk. After you finish configuring your boot disk, click **Select** to commit your changes.
8. Expand the **Management, security, disks, networking, sole tenancy** section.
9. On the **Management** tab, under **Availability policy**, verify that the following defaults are set:
  - To ensure availability of your SAP systems, keep the **Preemptibility** setting set to **Off (recommended)**.
  - To ensure that your VM can restart if there's a maintenance or failure event, keep the **Automatic restart** setting set to **On (recommended)**.
  - To ensure that your VM is migrated to other hardware during infrastructure maintenance, keep the **On host maintenance** setting set to **Migrate VM instance (recommended)**.
10. On the **Disks** tab, clear the **Delete boot disk when instance is deleted** checkbox.

Next, create a set of disk volumes for your database data, database logs, and stored procedures. For each disk volume, perform the following steps:

1. On the **Disks** tab, under **Additional disks**, click **Add new disk** to add persistent disks for storage.
2. Optionally, enter a name in the **Name** field.
3. Under **Type**, select the disk type for the disk.
4. Under **Source type**, select **None (blank disk)**.
5. Specify the size of your disk. Disk sizes are subject to SAP and IBM Db2 requirements. See:
  - [SAP Note 1707361 - Inst. Systems Based on NW 7.1 and Higher: UNIX Db2 for LUW](https://launchpad.support.sap.com/#/notes/1707361)  
(<https://launchpad.support.sap.com/#/notes/1707361>)
  - The appropriate [installation guide](https://help.sap.com/viewer/30839dda13b2485889466316ce5b39e9/CURRENT_VERSION/en-US/c8ed609927fa4e45988200b153ac63d1.html)  
([https://help.sap.com/viewer/30839dda13b2485889466316ce5b39e9/CURRENT\\_VERSION/en-US/c8ed609927fa4e45988200b153ac63d1.html](https://help.sap.com/viewer/30839dda13b2485889466316ce5b39e9/CURRENT_VERSION/en-US/c8ed609927fa4e45988200b153ac63d1.html))  
for your SAP system with IBM Db2
6. To add the disk, click **Done**.
7. Repeat the preceding steps for each disk you need for your system.

Finally, configure your networking details and create the VM:

1. On the **Networking** tab, under **Network interfaces**, choose the network that you created earlier.



2. If you are using a NAT gateway, under **Network tags**, add the tag that you specified as `[YOUR_TAG_NAME]` when you set up the route directing traffic through the gateway.
3. On the **SSH Keys** tab, select **Block project-wide SSH keys**.
4. To create and start the VM, click **Create**.

At the bottom of the page, you can click **REST** or **command line** to see the equivalent REST and `gcloud` commands for the VM instance you are creating. These can be useful for creating additional VMs.

Next, establish a connection to your VM.

The easiest ways to connect to a Linux VM are to connect from your browser through the Cloud Console, or from the command line through the `gcloud` command-line tool. You can also connect to a Linux VM by using common SSH tools.

Alternatively, you can generate a new key-pair for your Linux VM and apply it to your project, which lets you connect by using third-party tools, such as [PuTTY on Windows workstations](#) (`/compute/docs/instances/connecting-advanced#thirdpartytools`). For more details, see [Creating a new SSH key pair](#) (`/compute/docs/instances/adding-removing-ssh-keys#createsshkeys`).

You can also connect to your VM by using other configurations. For a list of supported environments and known issues, see [SSH from the browser](#) (`/compute/docs/ssh-in-browser`).

To connect to a Windows-based VM, you must first generate a password for the VM. You can then connect to the VM by using RDP or PowerShell.

After you create a Windows VM, you must generate a password for the VM before you can connect to it:

1. In the Cloud Console, go to the **VM instances** page:

[Go to the VM Instance page](https://console.cloud.google.com/compute/instances) (https://console.cloud.google.com/compute/instances)

2. Click the Windows instance where you want to generate a new password.
3. On the **Instance details** page, click **Create or reset Windows Password**. A password is generated for you.

You can change your password later by signing into the system and using the command:

1. To create a new user and password, specify a user name.
2. Click **Set**. Note the username and password so you can sign in to the VM.

You can use RDP to connect to a Windows instance and start a Remote Desktop session. Alternatively, you can [connect to Windows instances using the PowerShell terminal](#) (/solutions/sap/docs/netweaver-deployment-windows-dm#connecting\_to\_your\_vm).

Connect to the remote desktop on Windows instances by using one of the following methods:

- If you use the Chrome browser, you can connect through the [Chrome RDP for Compute Engine](#) (<https://chrome.google.com/webstore/detail/chrome-rdp-for-google-clo/mpbbnannobiobpfnblimoapbephgifkm/>) extension from [Fusion Labs](#) (<http://www.fusionlabs.net/>). This extension lets you connect through the Cloud Console.
- Download the RDP file and manually connect through the Windows Remote Desktop Connection client or a third-party client.

1. If you haven't created a username and password on the Windows VM yet, [create or reset your Windows password](/solutions/sap/docs/netweaver-deployment-windows-dm#connecting_to_your_vm) (/solutions/sap/docs/netweaver-deployment-windows-dm#connecting\_to\_your\_vm).
2. [Add a firewall rule](/solutions/sap/docs/netweaver-deployment-windows-dm#adding_firewall_rules) (/solutions/sap/docs/netweaver-deployment-windows-dm#adding\_firewall\_rules) or edit your existing firewall rule to open port **5986** on the [Google Cloud network](/compute/docs/networking) (/compute/docs/networking) where your Windows Server VM is located.
3. On your local workstation, open the PowerShell terminal.
4. Optionally, you can initialize a variable to hold your user credentials so you don't need to enter them each time you connect to the instance. If you skip this step, you receive a prompt for your username and password later.

## 5. Use the Enter-PSSession command

(<https://msdn.microsoft.com/en-us/powershell/reference/5.1/microsoft.powershell.core/enter-pssession>) to start a remote PowerShell session and include the flags to use SSL and skip credentials checks.

Next, format your disk drives and mount them to your VM.

After you connect to your VM, format and mount your disk drives.

Optionally, you can use Logical Volume Manager (LVM) to format a single disk and split it into the required directories. Both LVM and non-LVM methods are shown in the following procedures.

To format and mount your disk drives without LVM:

1. In the terminal where you are connected to your VM, create directories to which you can mount each device:

2. In the terminal, use the `ls` command to identify the disks that you want to mount.

The disk ID usually includes the name of the disk with a `google-` prefix or a `scsi-0Google_` prefix. This example uses the `google-` ID.

3. Format the disks. If you are using Db2 V11.1, either the `xfs` or `ext4` file system is recommended, because both support the fast preallocation file system feature. Db2 V10.5 doesn't support `xfs` or `ext4`.

The following commands format each disk with a single `ext3` file system and no partition table:

★ **Note:** For more information about the Db2-supported file systems, see [SAP Note 405827 - Linux: Recommended file systems](https://launchpad.support.sap.com/#/notes/405827) (<https://launchpad.support.sap.com/#/notes/405827>) .

4. Mount each disk to the VM:

5. Configure read and write access to the device. For this example, grant write access to the device for all users:

6. Optionally, you can add the persistent disks to the `/etc/fstab` file so that the device automatically mounts again when the VM restarts.

When you specify the entry `/etc/fstab` file, include the `nofail` option so that the instance can continue to boot even if the disk isn't present. For example, if you take a snapshot of the boot disk and create a new instance without any persistent disks attached, the instance can continue through the startup process and not pause indefinitely.

7. Create the `/etc/fstab` entry. Use the `blkid` command to find the UUID for the file system on the device and edit the `/etc/fstab` file to include that UUID with the mount options. You can complete this step with a single command for each disk drive:

where `[NOFAIL]` is the option that specifies what the system should do if it is unable to mount this disk.

8. Use the `cat` command to verify that your `/etc/fstab` entries are correct:

If you detach this persistent disk or create a snapshot from the boot disk for this instance, edit the `/etc/fstab` file and remove the entry for this persistent disk. Even with the `nofail` option in place, keep the `/etc/fstab` file in sync with the devices that are attached to your instance and remove these entries before you create your boot disk snapshot.

To format and mount a single disk with logical volumes for all required drives by using LVM:

1. Make sure the disk you have created for your drives is large enough to provide all volumes with the storage that they require. For example, in the following procedure, the size specified for the disk is 100 GB.
2. Run `pvscan` to scan all disks for physical volumes:
3. Initialize the disk or a partition for use by LVM:
4. Create a volume group:
5. Create logical volumes for each drive, using the sizes you need:
6. Format the volumes. If you are using Db2 V11.1, either the `xfs` or `ext4` file system is recommended, because both support the fast preallocation file system feature. Db2 V10.5 doesn't support `xfs` or `ext4`.

The following commands format each volume with a single `ext3` file system and no partition table:



★ **Note:** For the Db2-supported file systems, see [SAP Note 405827 - Linux: Recommended file systems](https://launchpad.support.sap.com/#/notes/405827) (<https://launchpad.support.sap.com/#/notes/405827>)

7. Create a mount point:

8. Mount the file systems to the mount point:

The resulting drives look similar to this:

9. Optionally, you can add the persistent disk to the `/etc/fstab` file so that the device automatically mounts again when the VM restarts.

When you specify the entry `/etc/fstab` file, include the `nofail` option so that the instance can continue to boot even if the disk isn't present. For example, if you take a snapshot of the boot disk and create a new instance without any persistent disks attached, the instance can continue through the startup process and not pause indefinitely.

10. Create the `/etc/fstab` entry. Use the `blkid` command to find the UUID for the file system on the device and edit the `/etc/fstab` file to include that UUID with the mount options. You can complete this step with a single command for each disk drive:

where `[NOFAIL]` is the option that specifies what the system should do if it is unable to mount this disk.

11. Use the `cat` command to verify that your `/etc/fstab` entries are correct:

If you detach this persistent disk or create a snapshot from the boot disk for this instance, edit the `/etc/fstab` file and remove the entry for this persistent disk. Even with the `nofail` option in place, keep the `/etc/fstab` file in sync with the devices that are attached to your instance and remove these entries before you create your boot disk snapshot.

After you have connected to your Windows VM, format your disks so that you can begin using them. You also configure the Windows pagefile in the following steps:

1. From the **Start** menu, search for and open the Server Manager.
2. Select **File and Storage Services**, and then select **Disks**.

Server Manager > File and Storage Services > Volumes > Disks

DISKS  
All disks | 5 total

Filter

+ Add criteria

Number	Virtual Disk	Status	Capacity	Unallocated	Partition	Read Only	Clustered	Subsystem	Bus Type	Name	Disk ID
win-wsv8000 (5)											
0		Online	50.0 GB	0.00 B	MBR			SAS		Google PersistentDisk SCSI Disk Device	Google win-wsv8000
1		Online	200 GB	0.00 B	GPT			SAS		Google PersistentDisk SCSI Disk Device	Google win-wsv8000-data
2		Online	50.0 GB	0.00 B	GPT			SAS		Google PersistentDisk SCSI Disk Device	Google win-wsv8000-sapmnt
3		Online	75.0 GB	0.00 B	GPT			SAS		Google PersistentDisk SCSI Disk Device	Google win-wsv8000-tlog
4		Online	25.0 GB	0.00 B	GPT			SAS		Google PersistentDisk SCSI Disk Device	Google win-wsv8000-page-tempdb

- In the **Disks** dialog, right-click the first non-MBR disk, and then click **New Volume**.
- Proceed with the defaults and enter a disk label.
- When you get to the file-system-settings step, change the **Allocation unit size** to a value from the following list:
  - Database disks: 32 KB
  - Pagefile: 8 KB
  - Other disks: default of 4 KB.
- Enter a volume label that describes the disk with a meaningful name.

New Volume Wizard

## Select file system settings

Before You Begin

Server and Disk

Size

Drive Letter or Folder

**File System Settings**

Confirmation

Results

File system: NTFS

Allocation unit size: 8192

Volume label: pagefile

Generate short file names (not recommended)

Short file names (8 characters with 3-character extensions) are recommended for applications running on client computers, but make file operation

- Repeat the preceding steps for each additional disk.

After you format and mount your disk drives, prepare your operating system.

After you create your VM, consult the relevant SAP notes on installation and ensure that your system includes the software components specified:

- [1984787 - SUSE LINUX Enterprise Server 12: Installation Notes](https://launchpad.support.sap.com/#/notes/1984787)  
(<https://launchpad.support.sap.com/#/notes/1984787>)

★ **Note:** On SLES 12 SP2, use **zypper update** to update OS packages. Otherwise, some software that runs on Java, such as JDBC, might core dump.

- [2002167 - Red Hat Enterprise Linux 7.x: Installation and Upgrade](https://launchpad.support.sap.com/#/notes/2002167)  
(<https://launchpad.support.sap.com/#/notes/2002167>)
- If planned, join the Windows Server OS to a domain.
- [Configure the High Performance power plan](#)  
([/compute/docs/tutorials/creating-high-performance-sql-server-instance#setting\\_the\\_power\\_profile](/compute/docs/tutorials/creating-high-performance-sql-server-instance#setting_the_power_profile)).
- Calculate the page file size referencing SAP Note [1518419: Page file and virtual memory required by the SAP system](https://launchpad.support.sap.com/#/notes/1518419) (<https://launchpad.support.sap.com/#/notes/1518419>) and [configure the Windows Pagefile](https://help.sap.com/saphelp_nw73ehp1/helpdata/EN/4d/0b7b6d277f5ca8e10000000a42189e/content.htm)  
([https://help.sap.com/saphelp\\_nw73ehp1/helpdata/EN/4d/0b7b6d277f5ca8e10000000a42189e/content.htm](https://help.sap.com/saphelp_nw73ehp1/helpdata/EN/4d/0b7b6d277f5ca8e10000000a42189e/content.htm))
- 

Now that your operating system is configured, you can install your IBM Db2 database server.

For guidance on installing SAP NetWeaver with IBM Db2, see the [installation guide](https://help.sap.com/viewer/30839dda13b2485889466316ce5b39e9/CURRENT_VERSION/en-US/c8ed609927fa4e45988200b153ac63d1.html) ([https://help.sap.com/viewer/30839dda13b2485889466316ce5b39e9/CURRENT\\_VERSION/en-US/c8ed609927fa4e45988200b153ac63d1.html](https://help.sap.com/viewer/30839dda13b2485889466316ce5b39e9/CURRENT_VERSION/en-US/c8ed609927fa4e45988200b153ac63d1.html)) that is specific to your SAP NetWeaver system.

To install IBM Db2 on your VM:

1. Establish an `ssh` connection to your Linux-based VM.

2. Download or copy the complete SAP media set for Db2 to your VM. You can download the SAP media set from the [SAP support portal](http://support.sap.com/swdc) (<http://support.sap.com/swdc>).
3. Install the IBM Db2 database server with the SAP Software Provisioning Manager.
4. Install your IBM Db2 license file. For more information about installing a Db2 license that you obtained from SAP, see [SAP Note 816773 - DB6: Installing an SAP OEM license](https://launchpad.support.sap.com/#/notes/816773) (<https://launchpad.support.sap.com/#/notes/816773>).

To install IBM Db2 on your VM:

1. Use RDP or Windows PowerShell to connect to your Windows-based VM.
2. Download or copy the complete SAP media set for Db2 to your VM. You can download the SAP media set from the [SAP support portal](http://support.sap.com/swdc) (<http://support.sap.com/swdc>).
3. Install the IBM Db2 database server.
4. Install your IBM Db2 license file. For more information about installing a Db2 license that you obtained from SAP, see [SAP Note 816773 - DB6: Installing an SAP OEM license](https://launchpad.support.sap.com/#/notes/816773) (<https://launchpad.support.sap.com/#/notes/816773>).

Before using your IBM Db2 instance, we recommend that you perform the following post-deployment steps:

1. Update your IBM Db2 software with the latest patches, if available.
2. Install any additional components.
3. Configure and back up your new IBM Db2 database.

For additional post-deployment guidance, see the Post-installation Tasks section of the [installation guide](#)

([https://help.sap.com/viewer/30839dda13b2485889466316ce5b39e9/CURRENT\\_VERSION/en-US/c8ed609927fa4e45988200b153ac63d1.html](https://help.sap.com/viewer/30839dda13b2485889466316ce5b39e9/CURRENT_VERSION/en-US/c8ed609927fa4e45988200b153ac63d1.html))

that applies to the SAP system that you are using with IBM Db2.

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

Don't delete the external IP address from the VM instances that are running IBM Db2 until you have completed creating a VPC network and a NAT gateway. After the external IP addresses are deleted, you can access the VM instances only through the NAT gateway.

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

2. Delete external IP addresses:

Because these NAT gateway instructions don't use an external IP address for IBM Db2, you can only connect to the IBM Db2 instances through the bastion instance by using `ssh`.

- To connect to IBM Db2 through the bastion instance, connect to the bastion host, and then to the IBM Db2 instances by using an SSH client of your choice.

---

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 Google Cloud 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 connecting through `ssh` from the browser.