This guide shows you how to set up Google Cloud resources for an IBM Db2 (IBM Db2) high-availability (HA) cluster for SAP on the Linux operating system.

These instructions are complementary to the instructions provided by SAP and IBM in IBM Db2 High Availability Solution: IBM Tivoli System Automation for Multiplatforms (https://help.sap.com/doc/a7ba6b8cc8dc46309c2f93bcfc417133/CURRENT_VERSION/en-US/DB6_SAMP_Guide.pdf). Always refer to the latest documentation provided by SAP and IBM when installing and configuring an IBM Db2 HA cluster on Google Cloud.

These instructions are for an IBM Db2 HA cluster that uses IBM Tivoli System Automation for Multiplatforms (TSAMP) to monitor the system and initiate appropriate action if the system becomes unresponsive. The cluster uses the IBM Db2 high-availability disaster recovery (HADR) function to replicate logged data changes to the standby database.

The cluster uses a floating IP address that is implemented by Google Cloud with either a Google Cloud static route or an alias IP address. In this context, the term "floating IP address" is synonymous with the term "virtual IP address", which is used in the SAP documentation.

These instructions show you how to set up an IBM Db2 HA cluster that consists of a primary IBM Db2 server and one secondary or standby IBM Db2 server, each of which are deployed on a separate Compute Engine virtual machine (VM).

This guide is intended for experienced SAP and IBM Db2 users who are familiar with high-availability clusters.

For more information about planning a Db2 HA cluster, see High-availability IBM Db2 clusters (/solutions/sap/docs/sap-ibm-db2-planning-guide#high-availability_ibm_db2_clusters) in the IBM Db2 for SAP Planning Guide.

The instructions for installing and configuring the SAP and IBM components are provided by SAP in IBM Db2 High Availability Solution: IBM Tivoli System Automation for Multiplatforms (https://help.sap.com/doc/a7ba6b8cc8dc46309c2f93bcfc417133/CURRENT_VERSION/en-US/DB6_SAMP_Guide.pdf).

Read both the SAP and Google Cloud documentation before you start the procedures in these instructions. At various stages of deployment, you might need to refer to both the SAP and Google Cloud documentation.

Before you create the IBM Db2 high-availability cluster, make sure that the following prerequisites are met:

- You or your organization has a Google Cloud account and you have created a project for the IBM Db2 HA cluster deployment. For information about creating Google Cloud projects, see Prerequisites (/solutions/sap/docs/sap-ibm-db2-deployment-guide#prerequisites) in the IBM Db2 for SAP deployment guide.
- You have a Virtual Private Cloud network on Google Cloud. For instructions on configuring a VPC network and firewall rules, as well as instructions for setting up a NAT gateway or bastion host for IBM Db2 for SAP, see IBM Db2 for SAP deployment guide (/solutions/sap/docs/sap-ibm-db2-deployment-guide).

These instructions show you how to deploy two VMs, define a floating IP address, and configure the alias IP address or Google Cloud routes that support the floating IP address. When you need to install the IBM components, you are referred to the SAP documentation.

The main Google Cloud services that you need to set up for an IBM Db2 high-availability cluster are:

- A VPC network and subnetwork
- Firewall rules (if you don't use another form of network access control)
- Compute Engine VMs and persistent disk storage

You also download and use a Google Cloud helper script when you define the custom resource that TSAMP uses to manage switching the floating IP address between hosts. The script enables TSAMP to interact with Google Cloud APIs.

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) 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. In Cloud Shell, download the `template_ha.yaml` configuration file template to your working directory:

2. To open the Cloud Shell code editor, click the pencil ( ) icon in the upper-right corner of the Cloud Shell terminal window.
3. Optionally, rename the `template_ha.yaml` file to identify the configuration it defines. For example, `db2_ha_s123_dh1.yaml`.
4. To open the `template_ha.yaml` in the code editor, double-click it.
5. Define the VMs and persistent disks in the `template_ha.yaml` file. The `template_ha.yaml` file contains two sections: `sap_db2_primary` and `sap_db2_secondary`. Each section contains a set

of property-value pairs followed by comments that include less-frequently used properties.

When you fill out each section, except for the **instanceName**, **zone**, and **otherHost** properties, the definitions for each VM must be identical.

The following table describes the properties included in each section. To use a property, replace the placeholder text and brackets with the values for your installation.

Property	Data type	Description
instanceName	String	The name of the VM instance on which you install IBM Db2. The name must be 13 characters or less, specified in lowercase letters, numbers, or hyphens.
instanceType	String	The type of Compute Engine virtual machine (/compute/docs/machine-types) on which you install IBM Db2. Specify a machine type with two or more vCPUs. For example, **n1-standard-4**.
zone	String	The zone in which you are deploying the IBM Db2 instance. It must be in the same region that you selected for your subnetwork.
subnetwork	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 ***shared-vpc-project*/*SUBNETWORK***. For example, `myproject/network1`.
linuxImage	String	The name of the Linux operating- system image or image family that you are using with IBM Db2. To specify an image family, add the prefix `family/` to the family name. For example, `family/rhel-7-sap-apps` or `family/sles-12-sp3-sap`. To specify an image, enter 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 Google Cloud image projects, see the Images (/compute/docs/images) page in the Compute Engine documentation.
db2SID	String	The database instance ID.
db2sidSize	Integer	The size in GB of **/db2/*DBSID***, which is the root directory of the database instance. The minimum and default values for **db2sidSize** are both 8 GB.
db2homeSize	Integer	The size in GB of **/db2/db2*db2sid***, which is the home directory of the database instance. The minimum and default values for **db2homeSize** are both 8 GB.

Property	Data type	Description
db2dumpSize	Integer	The size in GB of **/db2/*DBSID*/db2dump**, which holds the dump files from DB2 that are used for diagnosing problems. The minimum and default values for **db2dumpSize** are both 8 GB.
db2saptmpSize	Integer	The size in GB of **/db2/*DBSID*/saptmp**, which holds the database temporary table space. The minimum and default values for **db2saptmpSize** are both 8 GB.
db2sapdataSize	Integer	The size of **/sapdb/*DBSID*/sapdata**, which holds the database data files. The minimum and default values for **db2sapdataSize** are both 30 GB.
db2logSize	Integer	The size of **/db2/*DBSID*/logdir**, which holds the database transaction logs. The minimum and default values for **db2logSize** are both 8 GB.
db2backupSize	Integer	The size of the **/db2backup** volume. This property is optional. If you set to **0** or omit, no disk is created.
db2sapdataSSD	boolean	Specifies whether the data drive uses an SSD persistent disk (**Yes**) or an HDD persistent disk (**No**). **Yes** is the default.
db2logSSD	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.
usrsapSize	Integer	Required only if you are installing IBM Db2 to run with SAP NetWeaver on the same VM instance.
sapmntSize	Integer	Required only if you are installing IBM Db2 to run with SAP NetWeaver on the same VM instance.
swapSize	Integer	Required only if you are installing IBM Db2 to run with SAP NetWeaver on the same VM instance.
otherHost	String	The name of the other VM instance in the IBM Db2 HA cluster. The VM instance must be defined in the other set of properties in the same **template_ha.yaml** file.
networkTag	String	Optional. A network tag that represents your VM instance for firewall or routing purposes. If you specify **publicIP: No** and don't use a network tag, be sure to provide another means of access to the internet.

Property	Data type	Description
`publicIP`	boolean	Optional. Determines whether a public IP address is added to your VM instance. The default is **`Yes`**. ★ **Note**: Don't enter **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.
`serviceAccount`	String	Optional. If you create your own service account with locked down permissions, enter the name of the account here. By default, VMs are deployed using the default project service account. Note that an incorrectly defined service account prevents a successful deployment. The following is an example of a custom service account: **`myserviceuser@myproject.iam.gserviceaccount.com`**
`sap_deployment_debug`	boolean	Optional. If this value is set to **`Yes`**, the deployment generates verbose deployment logs. Don't turn this setting on unless a Google support engineer asks you to enable debugging.
`post_deployment_script`	String	Optional. Specifies the location of a script to run after the deployment is complete. The script should be hosted on a web server or in a Cloud Storage bucket. The URL should begin with **`http://`**, **`https://`**, or **`gs://`**. Note that this script is executed on all VMs that the template creates. To run it on the master instance only, add a check at the top of your script.

The following example VM definitions from a `template_ha.yaml` configuration file create two VMs for an IBM Db2 HA cluster. For each VM, the configuration file directs Deployment Manager to deploy an `n1-standard-4` VM that is running an operating system from the SLES 12 SP3 image family. The VM includes all of the directories that are required to run an IBM Db2 HA cluster. Deployment Manager doesn't create the SAP NetWeaver directories, because the directory sizes are set to `0`.

6. Deploy the VM instance with Deployment Manager.

Where:

- ***DEPLOYMENT-NAME*** represents a name of your choosing for the current deployment. This name is used to identify this deployment on the Deployments (https://console.cloud.google.com/dm/deployments) page of the Cloud Console.
- ***TEMPLATE-NAME*** represents the name that you gave to the configuration file or, if you didn't change the name of the default file, **template_ha.yaml**.

Deployment Manager reads the specifications in the **template_ha.yaml** file and configures the VM and persistent disks accordingly. 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).

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)

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.

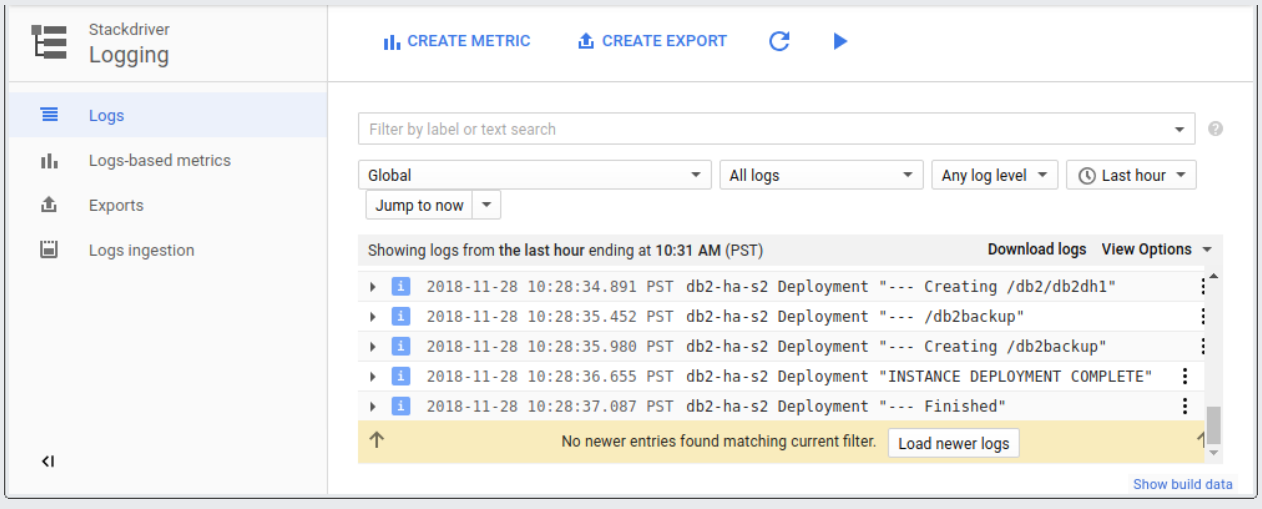

3. After VM deployment is complete, connect to each VM by using `ssh`. From the Compute Engine VM instances page (https://console.cloud.google.com/compute/instances), you can click the SSH

button for each VM instance, or you can use your preferred SSH method.

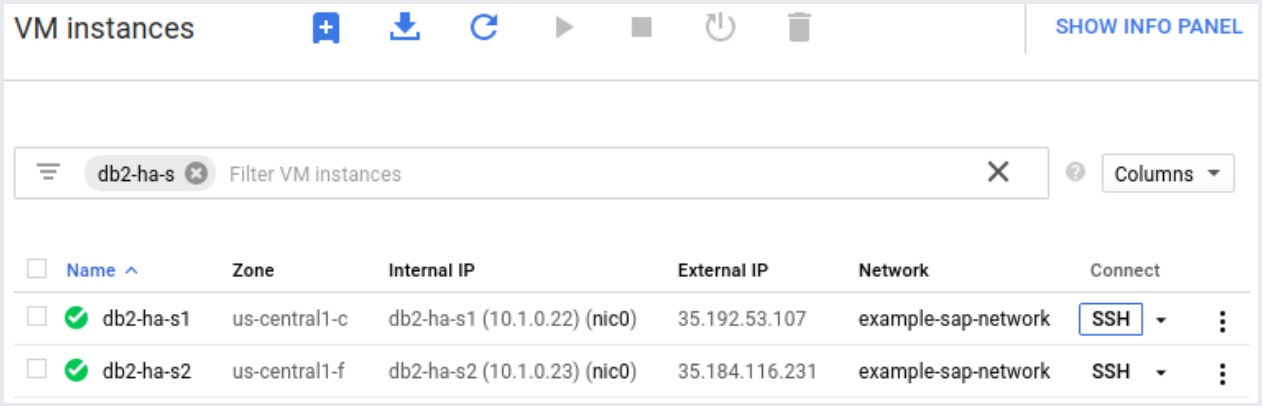

4. Change to the root user.

5. At the command prompt, enter `df -h`. Ensure that you see output similar to the following:

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) page, delete the deployment to clean up the VMs and persistent disks from the failed installation.
3. Rerun your deployment.

You need to select an IP address to use as the floating IP address. You need this IP address later when you set the host VM instance metadata and when you install and configure IBM db2 and the HA cluster.

Depending on whether you choose a route or alias IP implementation type (/solutions/sap/docs/sap-ibm-db2-planning-guide#db2_ha_floating_ip) for your floating IP address, the requirements for the floating IP address are different.

If you are using the static route implementation for your floating IP address, the IP address must be outside of your subnetwork IP address range and cannot be used by anything else in your organization's extended networks. Consult with your network admin to determine an appropriate IP address to use.

If you are using the alias IP address implementation for your floating IP address, you must reserve an IP address from the IP address range of the subnetwork that the hosts are using.

For alias IP address implementations only, reserve an alias IP address:

1. Open a terminal on a host VM or open Cloud Shell.

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

2. Reserve an IP address.

   Specifying the addresses property is optional. If you don't enter an IP address, Compute Engine selects an IP address from your subnetwork for you.

3. Display and make a note the reserved IP address to use when you install the database server and configure the HA cluster.

   For example:

You specify information about your floating IP address, including your chosen route or alias IP implementation type, as custom metadata for each VM instance in the cluster. For more information about choosing an implementation type for your floating IP address, see Floating IP addresses for IBM Db2 HA clusters on Google Cloud (/solutions/sap/docs/sap-ibm-db2-planning-guide#db2_ha_floating_ip).

Depending on your implementation type, the metadata parameters that you set are different. In the following two sections, follow the instructions in the section that applies to your floating IP address implementation

If you are using a route implementation for your floating IP address, use the parameters in the following table and the procedure that follows the table to set the instance metadata.

Parameter	Value	Purpose
`sap_ibm_vip_solution`	`route`	Indicates that this is a multi-zone deployment that uses a Google Cloud static route to support switching the floating IP address between hosts.
`sap_ibm_db2_vip`	***ip-address***	Specifies the floating IP address that you reserved in the previous step.
`sap_ibm_db2_routename`	***route-name***	Specifies an arbitrary name for the static route. For example, you could use `db2-dh1-vip-route`

Parameter	Value	Purpose
`sap_ibm_db2_routenet`	***vpc-network-name***	Specifies the VPC network that contains the IBM Db2 HA cluster.

To set your instance metadata for a static route implementation of your floating IP address:

1. Open a terminal on a host VM or open Cloud Shell.

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

2. For each host VM instance in the cluster, specify the same metadata for the route implementation of the floating IP address.

If you are using an alias IP address implementation for your floating IP address, use the parameters in the following table and the procedure that follows the table to set the instance metadata.

Parameter	Value	Purpose
`sap_ibm_vip_solution`	`alias`	Indicates that this is a single-zone deployment that uses a Google Cloud alias IP address to support switching the floating IP address between hosts.
`sap_ibm_db2_vip`	***ip-address***	Specifies the floating IP address that you reserved in the previous step.
`sap_ibm_db2_vip_range`	***alias-ip-range-name***	Optionally, specifies an arbitrary name for the alias IP range. For example, you could use `db2-dh1-vip-alias` The default value is the subnetwork name.

To set your instance metadata for an alias IP implementation of your floating IP address:

1. Open a terminal on a host VM or open Cloud Shell.

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

2. For each host VM instance in the cluster, specify the same metadata for the alias IP address implementation of the floating IP address.

To review the instance metadata that you set.

If you need to change your custom metadata.

During cluster setup, the SAP cluster setup tool validates the host names and internal IP addresses of each host VM and of the floating IP address. To ensure that the validation is successful, add the IP address, host name, and the VPC internal DNS name (https://cloud.google.com/compute/docs/internal-dns) for each host VM and the floating IP address to the `/etc/hosts` file on each host VM by using your preferred editor.

For example, as root, the following example updates `/etc/hosts`:

In the preceding example, the string between the host name and `>>` on each line is a VPC internal DNS name, which is used by the VPC internal DNS service.

The host VMs use a *zonal* internal DNS name, which includes a field for the zone. The floating IP address uses a *global* internal DNS name, which doesn't include a zone field.

For a VM host, you can retrieve the internal DNS name by entering the following command from a terminal on the host VM:

For a floating IP address, you can enter it yourself by using the following format.

After you update the `/etc/hosts` file, the relevant information in the `/etc/hosts` file should look similar to the following example:

After you have created your VM, prepare the operating system for the IBM Db2 HA cluster.

The requirements are defined by SAP and IBM. The SAP documentation calls for installing software, such as Perl and Korn Shell, that might not be pre-installed on your Compute Engine host VMs.

Check the following documents for the latest requirements:

- IBM Db2 high availability solution: IBM Tivoli System Automation for Multiplatforms (https://help.sap.com/doc/a7ba6b8cc8dc46309c2f93bcfc417133/CURRENT_VERSION/en-US/DB6_SAMP_Guide.pdf)
- 1984787 - SUSE LINUX Enterprise Server 12: Installation notes (https://launchpad.support.sap.com/#/notes/1984787)

★ **Note:** On SLES 12 SP2, customers must 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)

Before you start following the instructions in IBM Db2 High availability solution: IBM Tivoli System Automation for Multiplatforms (https://help.sap.com/doc/a7ba6b8cc8dc46309c2f93bcfc417133/CURRENT_VERSION/en-US/DB6_SAMP_Guide.pdf) to install and configure IBM Db2 and the HA cluster, review the overview of the following procedure, paying particular attention to the notes.

To install SAP NetWeaver and the primary application server, see the Google Cloud SAP NetWeaver documentation (/solutions/sap/docs/netweaver-guides) and the applicable SAP installation guides available from the SAP Help Portal (https://help.sap.com/viewer/index).

The following steps are an overview of the installation procedure. Refer to the SAP documentation for details.

1. Establish key-based SSH connectivity between the primary and secondary instances and between each instance and itself as described in the SAP documentation. SSH is used by the SAP cluster setup tool. Test all connections on each host. For example, on `db2-ha-s1` test both of the following.
   -
   -

   ★ **Note:** Each Compute Engine VM comes with its own **`id_rsa`** and **`id_rsa.pub`** files already generated. For Db2 HA clusters, the rsa key from the **`id_rsa.pub`** file of each host VM is added for you to the **`authorized_keys`** file of every other host VM in the cluster.

2. Download or copy the complete SAP media set for Db2 to your VM from the SAP support portal (http://support.sap.com/swdc).
3. On the primary host VM, use the SAP Software Provisioning Manager (SWPM) to install the IBM Db2 database server.

★ **Note:** On the **Cluster Configuration - Network** panel in SWPM you specify the floating IP address that you reserved and specify the tiebreaker IP address. For information about what you can use as a tiebreaker, see The network tiebreaker (/solutions/sap/docs/sap-ibm-db2-planning-guide#the_network_tiebreaker).

4. On the secondary host VM, set up the standby database by using a method such as SAP homogeneous system copy.
5. On both host VMs, install the license files for IBM Db2 and IBM TSAMP. For more information about installing the IBM licenses that you obtained from SAP, see SAP Note 816773 - DB6: Installing an SAP OEM license (https://launchpad.support.sap.com/#/notes/816773).
6. On both host VMs, install the latest version of TSAMP, as supported by your database version and operating system version.
7. On the primary host VM, use the latest version of the SAP cluster setup tool `sapdb2cluster.sh` to configure and create the IBM Db2 HA cluster.

★ **Note:** Before you create the Db2 HA cluster, you need to configure the SAP cluster setup tool to bypass the creation of the standard **`IBM.ServiceIP`** resource by specifying **`DB2_NO_SERVICE_IP=true`** in the **`sapdb2cluster.sh`** setup dialog. The cluster uses the **`IBM.Application`** resource instead.

For more information, see Best practices for floating IP addresses (https://cloud.google.com/solutions/best-practices-floating-ip-addresses).

8. After the cluster is created, on the primary host, use the DB2 high-availability instance configuration utility (db2haicu) utility to test that the cluster can failover.
    a. Exit the SAP cluster setup tool and the Korn shell.
    b. On the primary instance, confirm that the primary database server is online.

In the following example excerpt from the `lssam` output, the primary database instance is online:

c. Switch to the database instance user.

d. Start the db2haicu utility.

e. In the db2haicu interface, select option **5** and follow the prompts.

f. Exit the db2haicu utility.

g. On the primary host, check that the secondary host is now online.

In the following example excerpt from the lssam output, the secondary database instance is online:

To complete the cluster configuration, follow the instructions in the next section to create a custom TSAMP resource for the floating IP address and associate it in TSAMP with the IBM Db2 instance resource.

To enable TSAMP to manage the floating IP address, you need to create a TSAMP custom resource for it. To enable TSAMP to interact with Google Cloud while managing the floating IP address resource, you need to download and configure a helper script from Google Cloud.

On each host in the cluster, download the Google Cloud helper script and set its permissions.

1. On both the primary and standby hosts, as the root user from the `/root` directory on the primary VM, download the script.

2. On both hosts, set the permissions on the script.

On any one host in the cluster, create and configure a TSAMP custom resources for the floating IP address.

1. On any one host, use your preferred method to create a configuration file called `cluster_res.conf` and paste the following text in it after you update the NodeNameList parameter with your host names.

2. On the primary host as the root user, create the TSAMP custom resource with the following commands.

3. On the primary host as the root user, confirm that the Db2 instance resource that is online is on the same host as the online floating IP resource.

   In the output, the online resources should all be on the same host VMs:

   If the floating IP address resource isn't online on the same host that the database instance is, move the floating IP address resource.

4. As the root user on the primary host, establish a relationship in TSAMP between the database instance resource and the floating IP address resource.

   After you establish a relationship between the database instance resource and the floating IP address resource, you can test failover again, as described in the next section.

To confirm that the IBM Db2 HA cluster is configured correctly, trigger a failover and check that all of the online resources move from one host VM to the other.

To perform a failover test.

1. On the primary host as the root user, note which host VM the online resources are currently on.

2. On the primary host, change to the db2 instance user.

3. Start the db2haicu utility.

4. In the db2haicu utility interface, trigger a failover by selecting option **5** and following the prompts.
5. After the db2haicu utility finishes processing, exit the utility.
6. Switch to the root user.
7. Confirm that the online resources have moved to the other host VM.

Before using your IBM Db2 high-availability system on Google Cloud, we recommend that you complete all of the post-installation activities that are documented in IBM Db2 High Availability Solution: IBM Tivoli System Automation for Multiplatforms (https://help.sap.com/doc/a7ba6b8cc8dc46309c2f93bcfc417133/CURRENT_VERSION/en-US/DB6_SAMP_Guide.pdf), including.

1. Validating the database cluster.
2. Backing up the TSAMP core policy.
3. Updating the database fix packs.
4. Updating Db2 client connections to use the host name and IP address of the floating IP address. For example, update the `db2cli.ini` file on the SAP ABAP application servers.

If you are using a NAT gateway with your DB2 HA cluster, complete the set up of the NAT gateway, as described in Completing the NAT gateway installation (/solutions/sap/docs/sap-ibm-db2-deployment-guide#complete_nat_gateway_db2) in the IBM Db2 for SAP deployment guide.