

This guide shows you how to use Cloud Deployment Manager to deploy a SAP HANA scale-out system that includes the SAP HANA host auto-failover fault-recovery solution on SUSE Linux Enterprise Server (SLES). By using Deployment Manager, you can deploy a system that meets SAP support requirements and adheres to both SAP and Compute Engine best practices. The resulting SAP HANA system includes a master host, up to 15 worker hosts, and up to 3 standby hosts all within a single Compute Engine zone.

Compute Engine [automatic restart](/compute/docs/instances/setting-instance-scheduling-options#autorestart) (/compute/docs/instances/setting-instance-scheduling-options#autorestart), which is faster than host auto-failover, Compute Engine [live migration](/compute/docs/instances/live-migration) (/compute/docs/instances/live-migration), and the [high uptime percentage](/compute/sla) (/compute/sla) of Compute Engine VMs together might meet your availability requirements, making the increased cost and overhead of implementing SAP HANA host auto-failover unnecessary.

Do not use this guide if you do not need the host auto failover feature. Instead, use the [SAP HANA Deployment Guide](/solutions/sap/docs/sap-hana-deployment-guide) (/solutions/sap/docs/sap-hana-deployment-guide). If you need to deploy a Linux high-availability cluster for a single-host scale-up SAP HANA system, use the [SAP HANA High Availability Cluster on SLES Deployment Guide](/solutions/sap/docs/sap-hana-ha-deployment-guide) (/solutions/sap/docs/sap-hana-ha-deployment-guide).

This guide is intended for advanced SAP HANA users who are familiar with SAP scale-out configurations that include standby hosts for high-availability, as well as network file systems.

Before you create the SAP HANA high availability scale-out system, 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 SAP HANA deployment. For information about creating Google Cloud accounts and projects, see [Setting up your Google account](/solutions/sap/docs/sap-hana-deployment-guide#setting-up-your-google-account) (/solutions/sap/docs/sap-hana-deployment-guide#setting-up-your-google-account) in the SAP HANA Deployment Guide.
- The SAP HANA installation media is stored in a Cloud Storage bucket that is available in your deployment project and region. For information about how to upload SAP HANA installation media to a Cloud Storage bucket, see [Creating a Cloud Storage bucket](#)

(/solutions/sap/docs/sap-hana-deployment-guide#creating_a_cloud_storage_bucket_for_the_sap_hana_installation_files) in the SAP HANA Deployment Guide.

- You have an NFS solution, such as the managed [Filestore](#) (/filestore) solution, for sharing the SAP HANA /hana/shared and /hanabackup volumes among the hosts in the scale-out SAP HANA system. You specify the mount points for the NFS servers in the Deployment Manager configuration file before you can deploy the system. To deploy Filestore NFS servers, see [Creating instances](#) (/filestore/docs/creating-instances).

to avoid 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 you 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 Manager 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 Cloud provides, specify **publicIP**: No.

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), 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, 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 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) (<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) (<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 [Virtual Private Cloud documentation \(/vpc/docs/using-firewalls#creating_firewall_rules\)](https://cloud.google.com/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 (<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 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.

In the following instructions, you complete the following actions:

- Create the SAP HANA system by invoking Deployment Manager with a configuration file template that you complete.
- Verify deployment.
- Test the standby host(s) by simulating a host failure.

Some of the steps in the following instructions use Cloud Shell to enter the `gcloud` commands. If you have the latest version of Cloud SDK installed, you can enter the `gcloud` commands from a local terminal instead.

In the following steps, you download and complete a Deployment Manager configuration file template and invoke Deployment Manager, which deploys the VMs, persistent disks, and SAP HANA instances.

1. Confirm that your current quotas for project 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) (https://console.cloud.google.com/iam-admin/quotas)

2. Open Cloud Shell.

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

3. Download the `template.yaml` configuration file template for the SAP HANA high-availability scale-out system to your working directory:

4. Optionally, rename the `template.yaml` file to identify the configuration it defines. For example, you could use a file name like `hana2sp3rev30-scaleout.yaml`.

5. 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.

6. In the `template.yaml` file, update the following property values by replacing the brackets and their contents with the values for your installation. For example, you might replace "[ZONE]" with "us-central1-f".

Property	Data type	Description
<code>instanceName</code>	String	The name of the VM instance for the SAP HANA master host. The name must be specified in lowercase letters, numbers, or hyphens. The VM instances for the worker and standby hosts use the same name with a "w" and the host number appended to the name.
<code>instanceType</code>	String	The type of Compute Engine virtual machine (/solutions/sap/docs/sap-hana-planning-guide#vm_types) 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).
<code>zone</code>	String	The zone in which you are deploying your SAP HANA systems to run. It must be in the region that you selected for your subnet.
<code>subnetwork</code>	String	The name of the subnetwork 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> .

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 <code>family/</code> to the family name. For example, <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.
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. For SLES, specify <code>suse-sap-cloud</code> . For a list of Google Cloud image projects, see the Images (/compute/docs/images) page in the Compute Engine documentation.
sap_hana_deployment_bucket	String	The name of the Cloud Storage bucket in your project that contains the SAP HANA installation files that you uploaded in a previous step.
sap_hana_sid	String	The SAP HANA system ID. The ID must consist of 3 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 8 characters and include at least 1 uppercase letter, 1 lowercase letter, and 1 number.
sap_hana_system_password	String	The password for the database superuser. Passwords must be at least 8 characters and include at least 1 uppercase letter, 1 lowercase letter, and 1 number.
sap_hana_worker_nodes	Integer	The number of additional SAP HANA worker hosts that you need. You can specify 1 to 15 worker hosts. The default is value is 1.
sap_hana_standby_nodes	Integer	The number of additional SAP HANA standby hosts that you need. You can specify 1 to 3 standby hosts. The default value is 1.

Property	Data type	Description
<code>sap_hana_shared_nfs</code>	String	The NFS mount point for the <code>/hana/shared</code> volume. For example, <code>10.151.91.122:/hana_shared_nfs</code> . ★ Note: The NFS mount point for the <code>/hana/shared</code> volume must be empty before you run Deployment Manager.
<code>sap_hana_backup_nfs</code>	String	The NFS mount point for the <code>/hanabackup</code> volume. For example, <code>10.216.41.122:/hana_backup_nfs</code> .
<code>networkTag</code>	String	Optional. One or more comma-separated network tags 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.
<code>publicIP</code>	Boolean	Optional. Determines whether a public IP address is added to your VM instance. The default is Yes . ★ 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.
<code>sap_hana_double_volume_size</code>	Integer	Optional. Doubles the HANA volume size. Useful if you wish to deploy multiple SAP HANA instances or a disaster recovery SAP HANA instance on the same VM. By default, the volume size is automatically calculated to be the minimum size required for your memory footprint, while still meeting the SAP certification and support requirements.
<code>sap_hana_sidadm_uid</code>	Integer	Optional. Overrides the default value of the <code>[SID]adm</code> user ID. The default value is 900. You can change this to a different value for consistency within your SAP landscape.
<code>sap_hana_sapsys_gid</code>	Integer	Optional. Overrides the default group ID for sapsys. The default is 79.
<code>sap_deployment_debug</code>	Boolean	Optional. If this value is set to Yes , the deployment generates verbose deployment logs. Do not turn this setting on unless a Google support engineer asks you to enable debugging.

Property	Data type	Description
<code>post_deployment_script</code>	BooleanOptional	Optional. The URL or storage 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. Begin the value with <code>http://</code> , <code>https://</code> or <code>gs://</code> . Note that this script will be executed on all VMs that the template creates. If you only want to run it on the master instance, you will need to add a check at the top of your script.

The following example shows a completed configuration file that deploys an SAP HANA scale-out system with three worker hosts and one standby host in the us-central1-f zone. Each host is installed on an n1-highmem-32 VM that is running the SLES 12 SP2 operating system. The worker and standby hosts access the `/hana/shared` and `/hanabackup` volumes through the two NFS instances.

7. Create the instances:

The above command invokes 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 can take many minutes to complete.

To verify deployment, you check the deployment logs in Stackdriver, check the disks and services on the VMs of the primary and worker hosts, display the system in SAP HANA Studio, and test the takeover by a standby host.

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

★ **Note:** You might incur costs when completing this step in Stackdriver. For more information, see [Stackdriver pricing](https://cloud.google.com/stackdriver/pricing) (/stackdriver/pricing_v2).

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

2. Select **Global** from the resources list box. If "INSTANCE DEPLOYMENT COMPLETE" is displayed for all VMs, Deployment Manager processing is complete.

The screenshot shows the Stackdriver Logging console. The left sidebar has 'Stackdriver Logging' selected. The main area shows a filter for 'Global' and a list of logs from the last hour. The logs show the following messages:

Time	Source	Message
2018-09-27 11:53:58.457 PDT	hana-scaleout-w-failover	Deployment "Installing 1 standby nodes"
2018-09-27 11:53:58.791 PDT	hana-scaleout-w-failover	Deployment "--- hana-scaleout-w-failoverw4"
2018-09-27 11:54:51.171 PDT	hana-scaleout-w-failover	Deployment "Updating SAP HANA configured roles"
2018-09-27 11:54:52.864 PDT	hana-scaleout-w-failover	Deployment "INSTANCE DEPLOYMENT COMPLETE"

After deployment is complete, confirm that the disks and SAP HANA services have deployed properly by checking the disks and services of the master host and one worker host.

1. On the Compute Engine VM instances page, connect to the VM of the master host and the VM of one worker host by clicking the SSH button on the row of each of the two VM instances.

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

When connecting to the worker host, make sure that you aren't connecting to a standby host. The standby hosts use the same naming convention as the worker hosts, but have the highest numbered worker-host suffix before the first takeover. For example, if you have three worker hosts and one standby host, before the first takeover the standby host has a suffix of "w4".

2. In each terminal window, change to the root user.

3. In each terminal window, display the disk file system.

On the master host, you should see output similar to the following.

On the worker host, notice that the `/hana/data` and `/hana/log` directories have different mounts. On a standby host, the data and log directories are not mounted until the standby host takes over for a failed host.

4. In each terminal window, change to the SAP HANA operating system user. Replace **[SID]** with the [SID] value that you specified in the configuration file template.

5. In each terminal window, ensure that SAP HANA services, such as `hdbnameserver`, `hdbindexserver`, and others, are running on the instance.

On the master host, you should see output similar to the following:

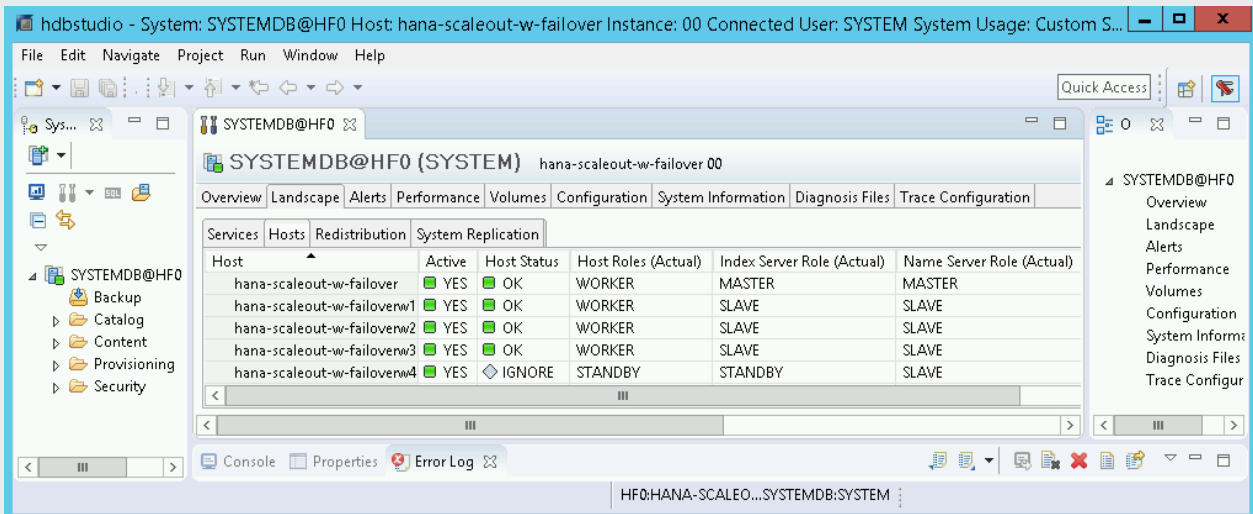
On a worker host, you should see output similar to the following:

1. Connect to the master SAP HANA host from SAP HANA Studio.

You can connect from an instance of SAP HANA Studio that is outside of Google Cloud or from an instance on Google Cloud. You might need to enable network access between the target VMs and SAP HANA Studio.

To use SAP HANA Studio on Google Cloud and enable access to the SAP HANA system, see [Installing SAP HANA Studio on a Compute Engine Windows VM](#) (/solutions/sap/docs/sap-hana-deployment-guide#install_hana_studio).

2. In SAP HANA Studio, click the **Landscape** tab on the default system administration panel. You should see a display similar to the following example.



Host	Active	Host Status	Host Roles (Actual)	Index Server Role (Actual)	Name Server Role (Actual)
hana-scaleout-w-failover	YES	OK	WORKER	MASTER	MASTER
hana-scaleout-w-failoverw1	YES	OK	WORKER	SLAVE	SLAVE
hana-scaleout-w-failoverw2	YES	OK	WORKER	SLAVE	SLAVE
hana-scaleout-w-failoverw3	YES	OK	WORKER	SLAVE	SLAVE
hana-scaleout-w-failoverw4	YES	IGNORE	STANDBY	STANDBY	SLAVE

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.
3. Rerun your deployment.

After you have confirmed that the SAP HANA system deployed successfully, test the failover function.

The following instructions trigger a failover by switching to the SAP HANA operating system user and entering the `HDB stop` command. The `HDB stop` command initiates a graceful shutdown of SAP HANA and detaches the disks from the host, which enables a relatively quick failover.

To simulate a failure of the entire VM, you can switch to the super user with `sudo su -` and enter the `ifconfig eth0 down` command. The `ifconfig eth0 down` command shuts down the network interface to the VM. Failover in this case takes longer because the system performs additional processing to confirm that the host is down and to move the disks.

To perform a failover test:

1. Connect to the VM of a worker host by using SSH. You can connect from the Compute Engine VM instances page by clicking the SSH button for each VM instance, or you can use your preferred SSH method.

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

2. Change to the SAP HANA operating system user. In the following example, replace "[SID]" with the SID that you defined for your system.

3. Simulate a failure by stopping SAP HANA:

The `HDB stop` command initiates a shutdown of SAP HANA, which triggers a failover. During the failover, the disks are detached from the failed host and reattached to the standby host. The failed host is restarted and becomes a standby host.

4. After allowing time for the takeover to complete, reconnect to the host that took over for the failed host by using SSH.

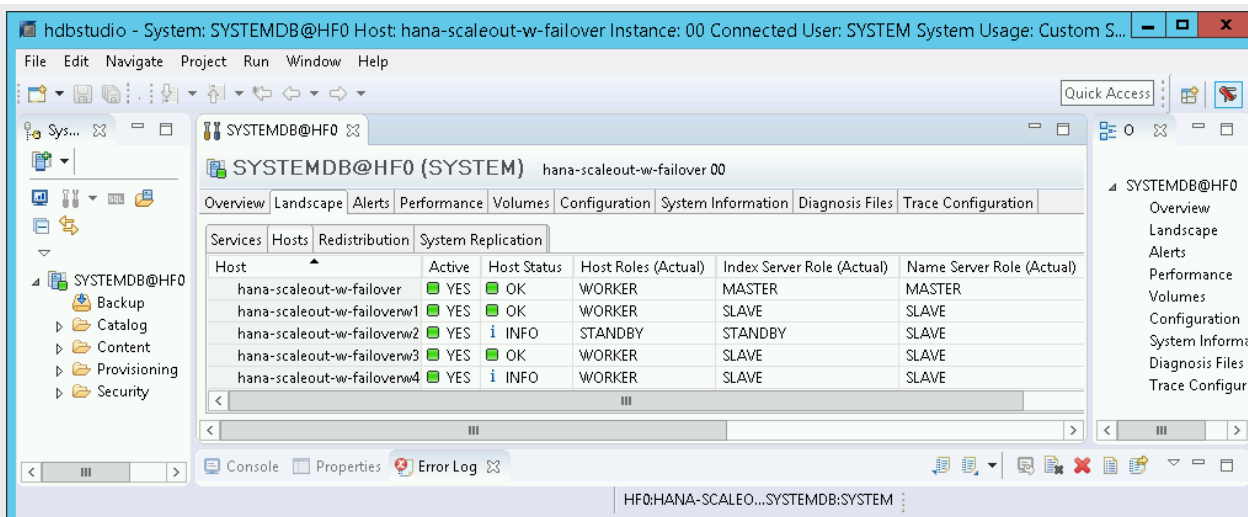
5. Change to the root user.

6. Display the disk file system of the VMs for the master and worker hosts.

You should see output similar to the following. The `/hana/data` and `/hana/log` directories from the failed host are now mounted to the host that took over.

7. In SAP HANA Studio, open the **Landscape** view of the SAP HANA system to confirm that the failover over was successful:

- The status of the hosts involved in the failover should be **INFO**.
- The **Index Server Role (Actual)** column should show the failed host as the new standby host.



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

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

1. In Cloud Shell, add tags to all instances, including the worker hosts:

2. In Cloud Shell, delete external IPs:

Optionally, you can set up Google's monitoring agent for SAP HANA, which collects metrics from SAP HANA and sends them to [Stackdriver Monitoring](/monitoring/docs/) (/monitoring/docs/). Stackdriver 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 Google's monitoring agent for SAP HANA, see the [SAP HANA Monitoring Agent User Guide](/solutions/sap/docs/sap-hana-monitoring-agent-user-guide) (/solutions/sap/docs/sap-hana-monitoring-agent-user-guide).

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 using 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) (https://help.sap.com/hana/SAP_HANA_Studio_Installation_Update_Guide_en.pdf) and access your SAP HANA database.

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) (http://help.sap.com/hana/SAP_HANA_Server_Installation_Guide_en.pdf).

1. Update the SAP HANA software with the latest patches.
2. Install any additional components such as Application Function Libraries (AFL) or Smart Data Access (SDA).
3. 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) (/solutions/sap/docs/sap-hana-operations-guide#backup_and_recovery).

- For more information about VM administration of and monitoring, see the [SAP HANA Operations Guide](/solutions/sap/docs/sap-hana-operations-guide/) (/solutions/sap/docs/sap-hana-operations-guide).