


Tableau Server for Healthcare on Linux in the Google Cloud Platform

Administrator Guide



Contents

- Install Tableau Server for Healthcare on the Google Cloud Platform3**
- Tableau Server on the Google Cloud Platform deployment options.....4**
- What you need before you begin.....5**
- Best practices for installing Tableau Server on the Google Cloud Platform.....5**
 - Keeping costs down.....6
- Tableau Server on the Google Cloud Platform topology.....6**
- Selecting a Google Compute Engine virtual machine type and size7**
- Self-deploy a single Tableau Server on the Google Cloud Platform.....9**
 - Step 1: Set up a Google Compute Engine VM9
 - Step 2: Connect to your Google Compute Engine VM10
 - Step 3: Install Tableau Server on your Google Compute Engine VM10
 - Step 4: Access Tableau Server from a remote browser10
 - Step 5: Publish data to Tableau Server on your Google Compute Engine VM11
 - Step 6: Manage Tableau Server licenses12
- Self-deploy Tableau Server on the Google Cloud Platform
in a distributed environment.....12**
 - Step 1: Create a Google Cloud Platform HIPAA-aligned project13
 - Step 2: Deploy three Google Compute Engine VMs14
 - Step 3: Create a Google Cloud Platform load balance
for the Tableau Server cluster.....15
- Troubleshooting Tableau Server on the Google Cloud Platform16**
- About Tableau & additional resources17**

Install Tableau Server for Healthcare on the Google Cloud Platform

When you install Tableau on a Google Compute Engine (GCE) virtual machine (VM), you still have the advanced capabilities of Tableau in a local environment. Running Tableau on a GCE VM is an excellent choice if you want the flexibility of scaling up and scaling out without having to purchase and maintain an expensive fleet of servers. For instance, you can configure Tableau to be highly available, and integrate with all of the common enterprise applications (for example, Active Directory) that many organizations rely on.

Looking for Tableau on Windows?

See [Install Tableau Server](#) on the Google Cloud Platform in Tableau Help.

Tableau can take advantage of many native Google Cloud Services:

- [Google Persistent Disk](#): Cloud-based block storage for virtual machines.
- [Google Cloud Storage](#): Cloud-based unified object storage for data archiving.
- [Google Cloud Dataflow](#): Cloud-based service for stream and batch data processing.
- [Google Cloud Dataproc](#): Cloud-based service for running Spark and Hadoop clusters.

In addition, you can store the data you use with Tableau Server using any of the following Google Cloud Services. Tableau provides native data connectors that enable you to connect to data in the following Google data sources:

- [Google Analytics](#): Web analytics tool for tracking and reporting website traffic.
- [Google BigQuery](#): Managed enterprise data warehouse for analytics.
- [Google Cloud SQL](#): Managed SQL database service.
- [Google Sheets](#): Spreadsheet included in Google Docs.

Tableau Server on the Google Cloud Platform deployment options

You can deploy Tableau Server on a Google Compute Engine virtual machine. A Tableau deployment on GCE supports the following:

Tableau Server	Google Compute Engine VM self-deployment
Production ready	✓
Upgrade Tableau without replacing the VM	✓
Install on Linux	✓
Install on Windows	✓
Scale-up	✓
Scale-out (add nodes)	✓
Active Directory support	✓
14-day trial license	✓
BYOL license	✓

Self-deployment of Tableau Server on a Google Compute Engine VM provides the most flexibility and options to customize Tableau Server to your environment. This is recommended for development, test, and production environments that start simple but may need to scale later.

For more information about self-deployment, see the sections *Self-deploy a single Tableau Server on the Google Cloud Platform* and *Self-deploy Tableau Server on the Google Cloud Platform in a distributed environment*.

What you need before you begin

To work with Tableau Server on the Google Cloud Platform, you will need:

- A Google Cloud Platform account. If you don't already have one, [sign up for a new account](#).
- Ensure that all Google Cloud Services you're using are covered by the Google Cloud Business Associate Agreement (BAA), and sign a Google Cloud BAA with the Google Cloud Platform. For more information, see the Google page [HIPAA Compliance on Google Cloud Platform](#).
- Make sure that your basic Google Cloud setup is complete:
 - Set up [Cloud Identity](#) for your organization.
 - Add an [organization admin](#) in Google Cloud.
 - Set up a [billing account](#).
- A Tableau license.
 - For a standalone deployment, you can use a 14-day trial version of Tableau, which doesn't require the use of a product key. After 14 days, you'll need to purchase a Tableau license and activate the associated product key.
 - For a cluster-based deployment, you need a user-based license (which covers all authorized users of Tableau), a core-based license (with a minimum of 16 cores), or a Tableau subscription license. To obtain a product key, [contact sales](#).

For more information about Tableau licensing options, see [Licensing Overview](#) in Tableau Help.

Best practices for installing Tableau Server on the Google Cloud Platform

The following best practices make installing Tableau Server in the cloud a better experience.

- Read [Plan Your Deployment](#) in Tableau Help.
- Read [Security Hardening Checklist](#) in Tableau Help.
- If you're new to the cloud, use the [Google Cloud Platform](#) free tier to get familiar with working in the cloud before you deploy Tableau Server on the Google Cloud Platform.
- Read Google Cloud Platform [best practices for enterprise organizations](#).
- Read about and understand [network topology](#) on the Google Cloud Platform and how to architect it.

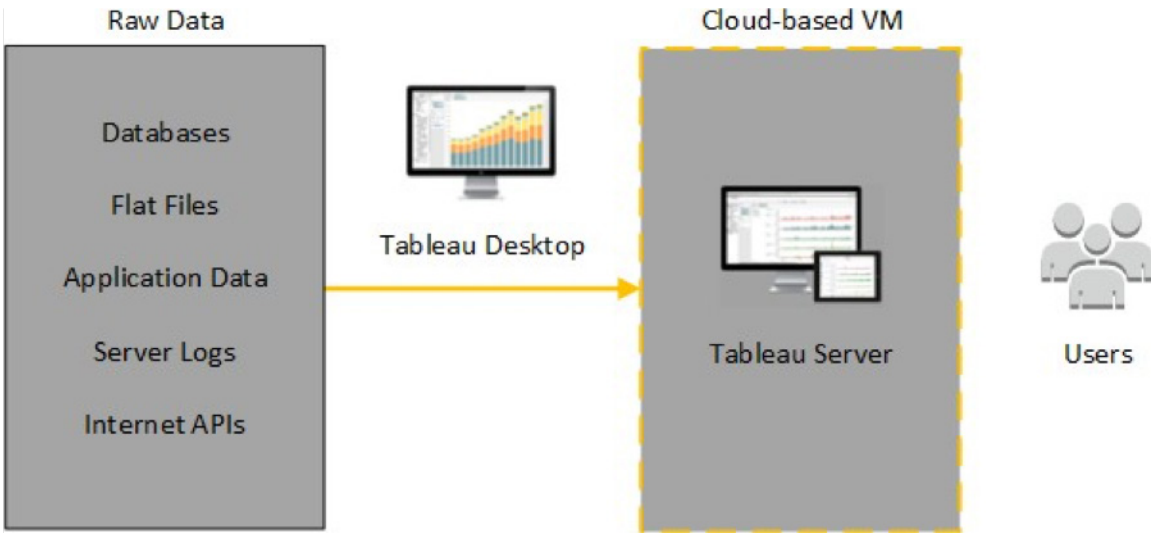
Keeping costs down

The Google Cloud Platform offers cloud-based services on a pay-as-you-go basis. Costs are determined by the services you run and amount of time you use them. Different combinations of instance types and sizes have different costs. For more information about services pricing, or to compare on-premises vs. Google Cloud Platform total cost of ownership (TCO), see [Google Cloud Platform pricing](#). You can estimate your total monthly costs using the [Google Cloud Platform pricing calculator](#).

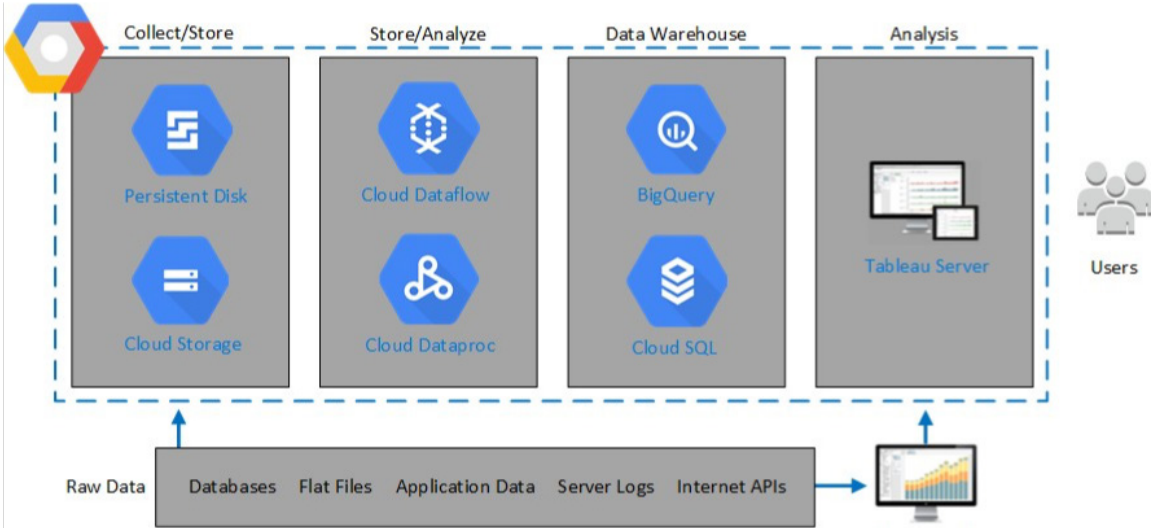
To help monitor and control usage costs on an ongoing basis, you can set up billing alerts for the Google Cloud Platform to alert you when your monthly Google Cloud Platform costs reach your predefined spending threshold. For more information, see the Google page [Set budgets and alerts](#).

Tableau Server on the Google Cloud Platform Topology

When deploying Tableau Server on the Google Cloud Platform, you may choose the level and depth of integration between Tableau Desktop and Tableau Server on your Google Compute Engine VMs. You can use Tableau Desktop and its ability to extract data from your data sources to act as a bridge between your data and Tableau. Based on your needs, or if you already have a lot of data in the Google Cloud Platform, you might instead choose to leverage the full range of Google Cloud Platform services in conjunction with Tableau Server. The following diagram shows data source integration with Tableau Desktop and Tableau.



In the following diagram, all of your data is hosted on the Google Cloud Platform. You can analyze both structured and unstructured data managed in a secure, scalable manner. You may leverage data that resides on the Google Cloud Platform only, unmanaged data that lives outside of the Google Cloud Platform, or a combination of both. This flexibility greatly increases the ability of your organization to support moves to the cloud from on-premise data, since both types of data is supported equally.



Selecting a Google Compute Engine virtual machine type and size

The Google Cloud Platform offers many types of Google Compute Engine VMs in various sizes. Choosing the right VM for your workload is an important factor for a successful Tableau Server deployment. You can choose from a wide range of VMs. For a complete list of all available VM types and sizes, see the [Google page Machine Types](#).

It is important to select a VM that can run Tableau Server. The VM must meet the Tableau Server hardware guidelines (a minimum of 8 cores and 32 GB of RAM).

At minimum, a 64-bit Tableau Server requires a 4-core CPU (the equivalent of 8 Google Compute Engine vCPUs) and 16 GB RAM. However, a total of 8 CPU cores (16 Google Compute Engine vCPUs) and 64GB RAM are strongly recommended for a single production Google Compute Engine VM.

Typical VM types and sizes for development, test, and production environments

· n1-standard-16

Recommended specifications for a single production instance:

Component/Resource	Google Cloud Platform
CPU	16+ vCPU
Operating System	<ul style="list-style-type: none">• Red Hat Enterprise Linux (RHEL) 7.3 and higher (not 8.x)• CentOS 7.3 and higher (not 8.x)• Oracle Linux 7.3 and higher (not 8.x)• The latest versions of Ubuntu 16.04 LTS and 18.04 LTS (not 17.04)• Debian 9
Memory	<ul style="list-style-type: none">• 64+ GB RAM (4GB RAM per vCPU rating system)• 100 GiB or larger volume for Tableau Server
Storage	Two volumes: <ul style="list-style-type: none">• 30-50 GiB volume for the operating system• 100 GiB or larger volume for Tableau Server
Storage type	SSD persistent disk, 200GB++ For more information about SSD persistent disks, see Storage Options at the Google Cloud Platform website.
Disk latency	Less than or equal to 20ms as measured by the <code>iostat</code> command in Linux.

Self-deploy a single Tableau Server on the Google Cloud Platform

You can install and run Tableau Server on a Google Compute Engine VM. Perform the following tasks to install and configure Tableau Server on a Google Compute Engine VM.

Important: Before you begin, review the [Google Cloud Platform HIPAA overview guide](#) and [Setting up a HIPAA-aligned project Google page](#) to familiarize yourself with the settings you'll need to make so that your solution is HIPAA compliant.

Step 1: Set up a Google Compute Engine VM

Create a Google Compute Engine virtual machine where you'll install Tableau.

1. Sign in to the [Google Cloud Platform Console](#).
2. If you haven't already, create a new project and set up billing. For more information, see [Creating and Managing Projects](#) at the Google website.
3. On the Google Cloud Platform dashboard, in the navigation pane, under **Compute**, click **Compute Engine**.
4. In the **Compute Engine VM instances** dialog box, click **Create**.
5. On the **Create an instance** page, specify the following instance properties:
 - a. **Name:** Enter a name for your VM.
 - b. **Zone:** To help reduce latency, choose the zone closest to your data source and users.
 - c. **Machine type:** Select a VM that meets the system requirements for Tableau Server (see section [Selecting a Google Compute Engine virtual machine type and size](#)). For example, select the n1-standard-16 machine type, click **Customize**, and then increase the memory to **64 GB**.
 - d. **Container:** If selected, clear the **Deploy a container image to this VM instance** check box (the default).
 - e. **Boot disk:** Click **Change**, and then select one of the supported Linux distributions, for example, **Ubuntu 16.04 LTS**, with a 128 GB SSD persistent disk, and then click **Select**.
 - f. **Identity and API access:** Use the default setting.
 - g. **Firewall:** Select **Allow HTTP traffic** and **Allow HTTPS traffic**.
6. When you've finished the settings, review your selections, and then click **Create**.

Step 2: Connect to your Google Compute Engine VM

After you create your Google Compute Engine VM, you must connect to it and prepare it to install Tableau.

1. In the **Compute Engine** console, click **VM instances**.
2. In the list of VM instances, select your instance.
3. In the **Connect** column, click the **SSH** pop-up menu, and then click **Open in browser window**. For more information about how to connect to Google Compute Engine VMs, see the Google page [Connecting to Instances](#).

Step 3: Install Tableau Server on your Google Compute Engine VM

Installing Tableau on a Google Compute Engine virtual machine is similar to installing it on any other server.

1. Copy the Tableau Server installation program to your VM by doing one of the following:
 - **Download the installer directly.** While logged into the VM, open a browser, browse to the [Tableau Server download page](#), and download the OS specific version of Tableau.
 - **Copy the installer to the VM.** If you already have a copy of the Tableau installer on your computer, copy it to the VM. Create a folder on the VM where you want to put the installer. From your computer, upload the Tableau installer file (for example, tableau-server-2020-2-3.x86_64.rpm) to the VM. For more information about how to copy files to a Google Compute Engine VM, see [Transferring Files to Instances](#) at the Google website.
2. Run the Tableau installer on the VM to start the server installation process. The installation process works the same as it does on any other server. For more information, see [Install and Configure Tableau Server](#).

Step 4: Access Tableau Server from a remote browser

Make sure that only authorized users can access the Google Compute Platform VM where you've installed Tableau.

1. In the Google Compute Platform console, under **Compute**, click **Compute Engine**, and then click **VM instances**. Find the external IP address associated with the VM.
2. On another computer (not on the VM), open a browser.
3. In the browser, type the external IP address of the Tableau you installed on the VM.

For example, `http://13.64.47.61`.

Note: The external IP is an ephemeral address. It is released when you stop the instance and a new external IP is assigned on restart. If you stop your instance, you may need to retrieve the new external IP address to access the server using your browser. GCP provides persistent static IP addresses. This time the URL resolves correctly and you reach the sign-in page for your instance of Tableau Server.

4. Sign in using the credentials for the administrator user that you created at the end of the installation process. When you've signed in, you see the Tableau Server Content page.

For more information on controlling access to your Tableau Server, see how to configure an Identity Store.

Step 5: Publish data to Tableau Server on your Google Compute Engine VM

After you set up Tableau Server on a Google Compute Engine virtual machine, authorized users in your organization can publish workbooks and data sources to the server. The process works the same way that it does if Tableau is installed on premises, except that users must sign in to the server that is running on the VM.

In a healthcare environment, maintaining data security and integrity is crucial. Through governance, you define the standards, processes, and policies to securely manage data and content. For more information, see [Governance in Tableau](#).

Before you allow users to publish, make sure that each user has been created with the Publisher site role in Tableau Server and that the users have Publisher permissions. For more information, see [Add Users to Tableau Server](#).

To publish workbooks and data sources to Tableau, users must sign in to the server by using the cloud server IP address using credentials they've been assigned. For more information, see [How to Publish Workbooks to a Tableau Server](#).

Step 6: Manage Tableau Server licenses

You don't need to manage your Tableau Server license if you shut down your Google Compute Engine virtual machine using the Google Cloud Platform console or from inside the VM itself. When you [stop a Google Compute Engine VM](#) with persistent SSD disk using the GCP Compute Engine console, the VM does not incur charges, but resources attached to the VM will still incur charges.

Deleting a VM deallocates all of its resources, including external storage. You might delete your VM when you no longer need it. Always deactivate your Tableau license before you shut down and delete your VM.

Before you delete your VM, do the following:

1. Sign in to your VM.
2. Create a backup of Tableau, and then copy the .tsbak file to another location, such as [Google Cloud Storage](#), that will be accessible after the VM is no longer available. For more information, see [Perform a Full Backup and Restore of Tableau](#).
3. Use the `tsm licenses deactivate` command to deactivate your Tableau product key. You can now shut down and delete your VM.

Self-deploy Tableau Server on the Google Cloud Platform in a distributed environment

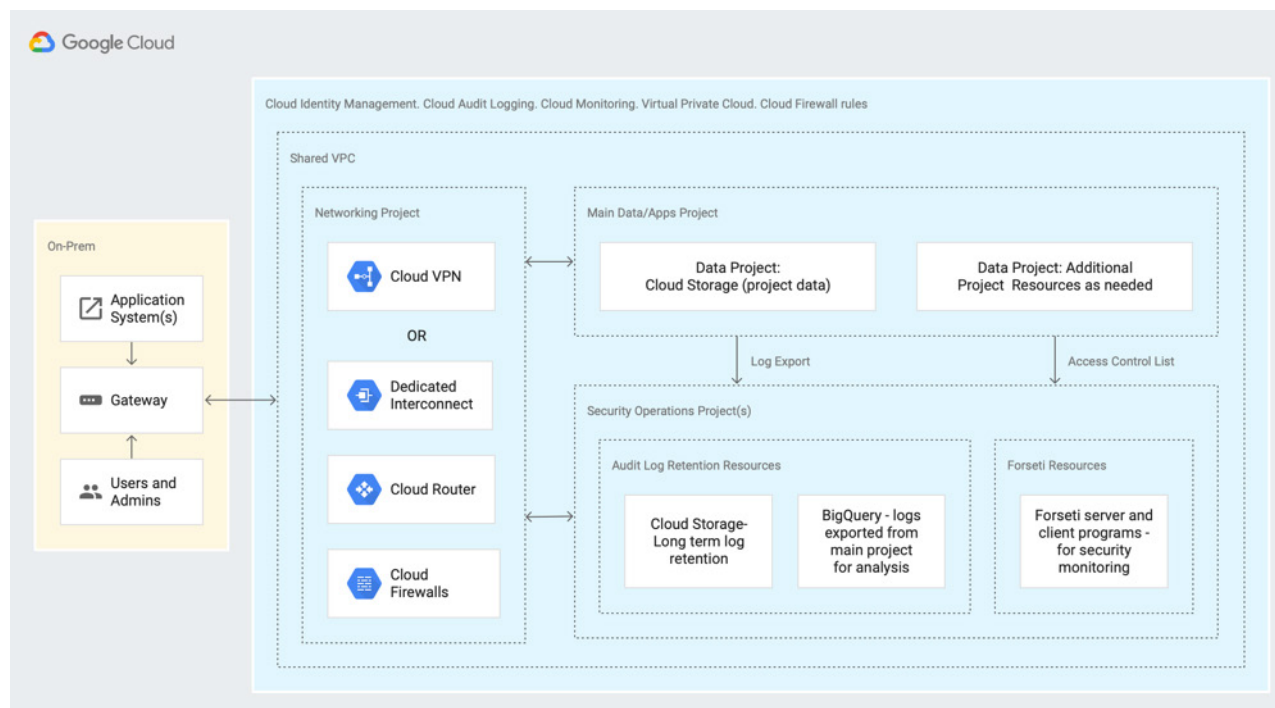
You can run Tableau Server on two VMs in a distributed environment, also known as a cluster. However, if you want to run Tableau Server in a highly available (HA) distributed environment, you need to launch three or more Google Compute Engine VMs of the same type and capacity and configure them as additional nodes. This scenario assumes that you have three Google Compute Engine VMs with Tableau Server installed on each instance. One instance is configured as the initial node, and the other two instances are configured as additional nodes.

Use the following steps to install and deploy Tableau Server on a cluster of three Google Compute Engine VMs in a highly available and scalable configuration.

Important: Before you begin, review the [Google Cloud Platform HIPAA overview guide](#) and [Setting up a HIPAA-aligned project](#) [Google page](#) to familiarize yourself with the settings you'll need to make so that your solution is HIPAA compliant.

Step 1: Create a Google Cloud Platform HIPAA-aligned project

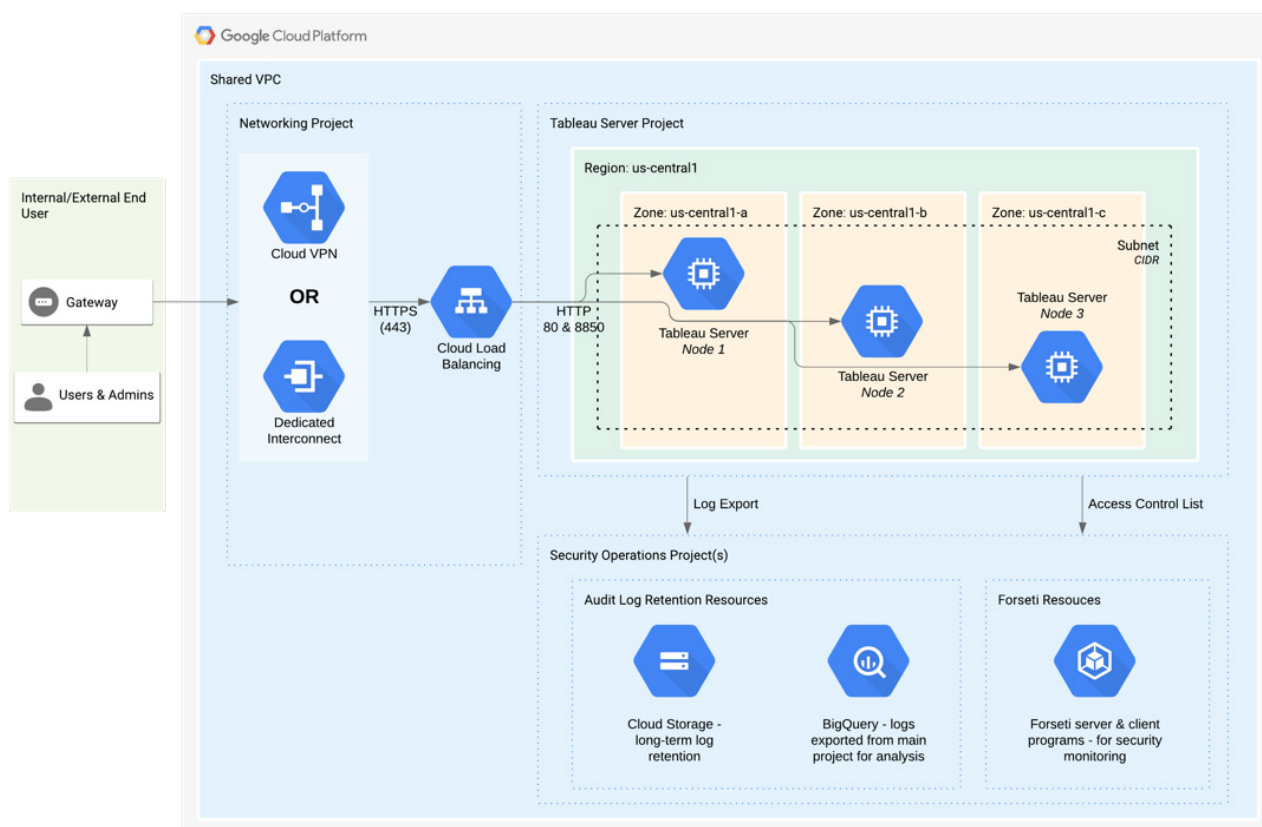
Referring to the [Setting up a HIPAA-aligned Google Cloud project](#) tutorial, and relying on the [Cloud Healthcare Data Protection Toolkit](#), this architecture helps you build a Google Cloud-based infrastructure with few steps by treating the configuration as code. The following diagram illustrates how the architecture helps you meet security and compliance best practices by using reusable building blocks: a Cloud Deployment Manager configuration script and parameterized configuration templates.



Step 2: Deploy three Google Compute Engine VMs

Deploy three Google Compute Engine virtual machines of the same type and capacity, in the configured subnets, with each VM in its own zone. Select VMs that meet the system requirements for Tableau Server. You'll install Tableau Server on each of the Google Compute Engine VMs, and use one node as the initial server and the two other nodes as additional servers.

For more information about installing Tableau in a distributed Linux environment, see [Distributed and Highly Available Tableau Server Installations](#) and [Install and Configure Additional Nodes](#) in Tableau Help.



Step 3: Create a Google Cloud Platform load balancer for the Tableau Server cluster

Read [Setting up a load balancer](#) at the Google website, and follow the steps to launch a load balancer within your Google Cloud Platform Virtual Private Cloud (VPC).

1. Create and launch an internally accessible load balancer and follow the steps in [Setting Up Internal Load Balancing](#), and select the two private subnets. Make sure that you configure the load balancer with SSL, as explained by Google in [SSL Certificates](#).

Important: Ensure that you lock down the entry point of your cloud environment to a [Cloud VPN](#) or a [Dedicated Interconnect](#) that is connected to your private network. For more information, see the [Google security whitepaper](#).

2. Ensure that your VPC network firewall rules are configured to allow access on port 80 or 443 only, with the source limited to hosts or ranges of hosts that will access Tableau. For more information about firewall rules, see the Google page [Firewall Rules Overview](#).
3. Specify the ping path as /.
4. Update Tableau to use the load balancer. For more information, see [Add a Load Balancer](#) in Tableau Help.

Troubleshooting Tableau Server on the Google Cloud Platform

If you're having trouble with Tableau on your Google Compute Engine VM, or with getting access to Tableau after installation, review these troubleshooting tips.

Cannot access the server from a remote browser

One or more of the following issues might prevent you from accessing the server from a remote browser:

- Sign in to the instance and make sure that the server is running.
- In the GCP Compute Engine console, make sure that you've configured an endpoint to allow HTTP traffic to the instance.

You receive the following error message: Tableau Server is unlicensed

This error can occur if you've shut down instances without first deactivating the license. Contact [Tableau Support](#).

Tableau Server does not use all CPU cores

One or more of the following symptoms might occur when working with Tableau Server core-based licensing:

- When installing Tableau Server, the installer might show fewer cores than you would expect to see for the computer running Tableau Server.
- On the Licenses page on Tableau Server, the number of cores listed under Licenses in Use is less than you would expect to see.
- When you run `tsm licenses list`, Tableau Server might report fewer cores in use that you would expect to see.

For more information, see [Not All Cores Recognized by Tableau Server](#).

About Tableau

Tableau is a complete, integrated, and enterprise-ready visual analytics platform that helps people and organizations become more data driven. Whether on-premises or in the cloud, on Windows or Linux, Tableau leverages your existing technology investments and scales with you as your data environment shifts and grows. Unleash the power of your most valuable assets: your data and your people.

Additional resources

[Tableau Help: Tableau Server on Linux](#)

[Tableau for IT](#)

[Tableau solutions for healthcare analytics](#)

[Tableau solutions for Google](#)

