

Creating a Cloud Bigtable Instance

A Cloud Bigtable instance is a container for up to 4 Cloud Bigtable clusters. Instances that use [replication](/bigtable/docs/replication-overview) (/bigtable/docs/replication-overview) have two or more clusters, and instances without replication have only one cluster.

This page explains how to create an instance with or without replication. Before you read this page, you should be familiar with the [overview of Cloud Bigtable](/bigtable/docs/overview) (/bigtable/docs/overview). You should also read the [overview of instances, clusters, and nodes](/bigtable/docs/instances-clusters-nodes) (/bigtable/docs/instances-clusters-nodes).

Before you begin

Before you begin, you'll need to prepare your environment and do some initial planning:

1. Select or create a GCP project.

[Go to the project selector page](https://console.cloud.google.com/projectselector2/home/dashboard) (https://console.cloud.google.com/projectselector2/home/dashboard)

★ **Important:** A project name must be between 4 and 30 characters. When you type the name, the form suggests a project ID, which you can edit. The project ID must be between 6 and 30 characters, with a lowercase letter as the first character. The last character cannot be a hyphen. Project identifiers might be retained beyond the life of your project, so do not use sensitive information for the project name or ID.

2. Make sure billing is enabled for your Google Cloud project. [Learn how to confirm that billing is enabled for your project](/billing/docs/how-to/modify-project) (/billing/docs/how-to/modify-project).
3. Enable the Cloud Bigtable and Cloud Bigtable Admin APIs.

[Enable the APIs](https://console.cloud.google.com/flows/enableapi?apiid=bigtable,bigtableadmin.google.com) (https://console.cloud.google.com/flows/enableapi?apiid=bigtable,bigtableadmin.google.com)

4. Click **Go to credentials**.
5. On the **Credentials** page, answer the required questions:

- a. In response to "Which API are you using?", choose **Cloud Bigtable API**.

- b. Answer the question about App Engine or Compute Engine. For this quickstart, you are *not* using them.
 - c. Click **What credentials do I need?**
 - d. If you are prompted to create a service account, provide the following information:
 - i. Choose a name for your service account.
 - ii. For **Role**, choose **Cloud Bigtable > Bigtable Administrator**.
 - iii. For **Key type**, select **JSON**.
 - iv. Click **Continue**.
 - v. Save the newly created service account JSON file to your local **HOME** directory or your preferred location.
 - e. If you are not prompted to create a new service account, click **Done**.
6. If you want to use the command-line tools for Cloud Bigtable, [install the Cloud SDK](/bigtable/docs/installing-cloud-sdk) (/bigtable/docs/installing-cloud-sdk) and the [cbt command-line tool](/bigtable/docs/cbt-overview) (/bigtable/docs/cbt-overview) if you haven't already.
7. If you plan to enable replication, do the following:
- Take a few minutes to read the [replication overview](/bigtable/docs/replication-overview) (/bigtable/docs/replication-overview).
 - Identify your [use case for replication](/bigtable/docs/replication-overview#use-cases) (/bigtable/docs/replication-overview#use-cases).
 - Determine the [region or regions](/bigtable/docs/locations) (/bigtable/docs/locations) that your instance should be in, based on your use case and the location of your application and traffic.
 - Decide how you'll use [application profiles](/bigtable/docs/app-profiles) (/bigtable/docs/app-profiles) to route incoming requests.

Creating an instance

To create a Cloud Bigtable instance:

```
Consolegcloud (#gcloud)cbt (#cbt)
```

1. Open the Create Instance page in the Cloud Console.

[Open the Create Instance page \(https://console.cloud.google.com/bigtable/create-instance\)](https://console.cloud.google.com/bigtable/create-instance)

2. Enter a name for the instance.

The Cloud Console displays this name to identify your instance.

3. Enter an instance ID.

The instance ID is a permanent identifier for the instance.

4. Enter a cluster ID for the first cluster.

The cluster ID is a permanent identifier for the cluster.

5. Choose the region and zone (/bigtable/docs/locations) where the first cluster will run.

If you plan to use replication (/bigtable/docs/replication-overview) within a single region, make sure Cloud Bigtable is available in at least one other zone in that region. View the zone list (/bigtable/docs/locations).

6. Choose the number of Cloud Bigtable nodes for the first cluster. If you aren't sure how many nodes you need, use the default. You can add more nodes later. Learn more (/bigtable/docs/instances-clusters-nodes#nodes).

7. To enable replication now, click **Add replicated cluster**, then update the settings for the second cluster. Repeat this step to create up to 4 clusters in the instance. You can also enable replication later by adding a cluster (/bigtable/docs/modifying-instance#clusters-adding).

An instance's clusters must each be in unique zones. You can create an additional cluster in any zone where Cloud Bigtable is available. If the **Add cluster** button is disabled, change the zone for your first cluster.

8. Choose whether to use an SSD or HDD disk for your clusters. In most cases, SSD is best. This choice is permanent. Learn more (/bigtable/docs/choosing-ssd-hdd).

The Cloud Console shows how the disk type and number of nodes affect your throughput and cost.

9. Click **Create** to create the instance.

10. Next, review the replication settings in the default app profile (/bigtable/docs/configuring-app-profiles#updating-app-profile) to see if they make sense for your replication use case (/bigtable/docs/replication-overview#use-cases). You might need to update the default app profile (/bigtable/docs/configuring-app-profiles#updating-app-profile) or create custom app profiles (/bigtable/docs/configuring-app-profiles#creating-app-profile).

What's next

- Find out how Cloud Bigtable uses [instances, clusters, and nodes](/bigtable/docs/instances-clusters-nodes) (/bigtable/docs/instances-clusters-nodes).
- Learn about [Cloud Bigtable replication](/bigtable/docs/replication-overview) (/bigtable/docs/replication-overview).
- [Review and update the default app profile](/bigtable/docs/configuring-app-profiles#updating-app-profile) (/bigtable/docs/configuring-app-profiles#updating-app-profile) for replication, and [create custom app profiles](/bigtable/docs/configuring-app-profiles#creating-app-profile) (/bigtable/docs/configuring-app-profiles#creating-app-profile) as needed.
- Find out how to [modify an existing instance](/bigtable/docs/modifying-instance) (/bigtable/docs/modifying-instance).

Except as otherwise noted, the content of this page is licensed under the [Creative Commons Attribution 4.0 License](https://creativecommons.org/licenses/by/4.0/) (https://creativecommons.org/licenses/by/4.0/), and code samples are licensed under the [Apache 2.0 License](https://www.apache.org/licenses/LICENSE-2.0) (https://www.apache.org/licenses/LICENSE-2.0). For details, see the [Google Developers Site Policies](https://developers.google.com/site-policies) (https://developers.google.com/site-policies). Java is a registered trademark of Oracle and/or its affiliates.

Last updated 2020-06-25 UTC.