

Alpha

This product or feature is in a pre-release state and might change or have limited support. For more information, see the [product launch stages](/products/#product-launch-stages) (/products/#product-launch-stages).

This tutorial shows you how to train a Keras ResNet model on Cloud TPU using `tf.distribute.TPUStrategy`.

If you are not familiar with Cloud TPU, it is strongly recommended that you go through the [quickstart](https://cloud.google.com/tpu/docs/quickstart) (https://cloud.google.com/tpu/docs/quickstart) to learn how to create a TPU and Compute Engine VM.

Warning: This tutorial uses a third-party dataset. Google provides no representation, warranty, or other guarantees about the validity, or any other aspects of this dataset.

- Create a Cloud Storage bucket to hold your dataset and model output.
- Prepare a fake/imagenet dataset that is similar to the ImageNet dataset.
- Run the training job.
- Verify the output results.

This tutorial uses billable components of Google Cloud, including:

- Compute Engine
- Cloud TPU
- Cloud Storage

Use the [pricing calculator](/products/calculator/) (/products/calculator/) to generate a cost estimate based on your projected usage. New Google Cloud users might be eligible for a [free trial](/free/) (/free/).

Before starting this tutorial, check that your Google Cloud project is correctly set up.

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).

This walkthrough uses billable components of Google Cloud. Check the [Cloud TPU pricing page](/tpu/docs/pricing) (/tpu/docs/pricing) to estimate your costs. Be sure to [clean up](#) (#clean-up) resources you create when you've finished with them to avoid unnecessary charges.

This section provides information on setting up Cloud Storage storage, VM, and Cloud TPU resources for tutorials.

Important: Set up all resources (Compute Engine VM, Cloud TPU, and Cloud Storage bucket) in the same region/zone to reduce network latency and network costs. VMs and TPU nodes are located in [specific zones](/tpu/docs/types-zones#types) (/tpu/docs/types-zones#types), which are subdivisions within a region.

1. Open a Cloud Shell window.

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

2. Create a variable for your project's name.
3. Configure `gcloud` command-line tool to use the project where you want to create Cloud TPU.
4. Create a Cloud Storage bucket using the following command:

This Cloud Storage bucket stores the data you use to train your model and the training results. The `ctpu up` tool used in this tutorial sets up default permissions for the Cloud TPU service account. If you want finer-grain permissions, review the [access level permissions](/tpu/docs/storage-buckets) (/tpu/docs/storage-buckets).

The bucket location must be in the same region as your Compute Engine (VM) and your Cloud TPU node.

5. Launch the Compute Engine VM resource using the `ctpu up` command.

★ **Note:** If you have more than one project, you must specify the project name.

For more information on the CTPU utility, see the [CTPU Reference](/tpu/docs/ctpu-reference) (/tpu/docs/ctpu-reference).

6. When prompted, press `y` to create your Cloud TPU resources.

★ **Note:** The first time you run `ctpu up` on a project it takes about 5 minutes to perform startup tasks such as SSH key propagation and API turnup.

When the `ctpu up` command has finished executing, verify that your shell prompt has changed from `username@project` to `username@tpuname`. This change shows that you are now logged into your Compute Engine VM.

★ **Note:** If you are not connected to the Compute Engine instance, you can connect by running the following command, replacing **`vm-name`** with the name of your VM:

From this point on, a prefix of `(vm)$` means you should run the command on the Compute Engine VM instance.

Set up the following environment variables, replacing **`bucket-name`** with the name of your Cloud Storage bucket:

The training application expects your training data to be accessible in Cloud Storage. The training application also uses your Cloud Storage bucket to store checkpoints during training.

ImageNet is an image database. The images in the database are organized into a hierarchy, with each node of the hierarchy depicted by hundreds and thousands of images.

This tutorial uses a demonstration version of the full ImageNet dataset, referred to as *fake_imagenet*. This demonstration version allows you to test the tutorial, while reducing the storage and time requirements typically associated with running a model against the full ImageNet database.

The fake_imagenet dataset is at this location on Cloud Storage:

The fake_imagenet dataset is only useful for understanding how to use a Cloud TPU and validating end-to-end performance. The accuracy numbers and saved model will not be meaningful.

For information on how to download and process the full ImageNet dataset, see [Downloading, preprocessing, and uploading the ImageNet dataset](/tpu/docs/imagenet-setup) (/tpu/docs/imagenet-setup).

Caution: For this tutorial, make sure you **don't** set the `STORAGE_BUCKET` environment variable to the path of the fake_imagenet dataset. You can read from `gs://cloud-tpu-test-datasets` but you can't write to it. As a result, you can't use it to write out training logs. Make sure the `STORAGE_BUCKET` environment variable is set to your own Cloud Storage bucket, as shown above.

Note: If you want to monitor the model's output and performance, follow the guide to [setting up TensorBoard](/tpu/docs/tensorboard-setup) (/tpu/docs/tensorboard-setup).

1. Launch a Cloud TPU resource using the `ctpu` utility.

2. Set the Cloud TPU name variable. This will either be a name you specified with the `--name` parameter to `ctpu up` or the default, your username:

3. The ResNet training script requires a few extra packages. Install them now:

4. Add the top-level `/models` folder to the Python path with the command

The ResNet-50 model is pre-installed on your Compute Engine VM.

5. Navigate to the directory:

6. Run the training script. This uses a `fake_imagenet` dataset and trains ResNet for one epoch.

Parameter	Description
<code>tpu</code>	Uses the name specified in the <code>TPU_NAME</code> variable.
<code>data_dir</code>	Specifies the Cloud Storage path for training input. It is set to the <code>fake_imagenet</code> dataset in this example.
<code>model_dir</code>	Specifies the directory where checkpoints and summaries are stored during model training. If the folder is missing, the program creates one. When using a Cloud TPU, the <code>model_dir</code> must be a Cloud Storage path (<code>gs://...</code>). You can reuse an existing folder to load current checkpoint data and to store additional checkpoints as long as the previous checkpoints were created using TPU of the same size and TensorFlow version.
<code>distribution_strategy</code>	To run ResNet model on a TPU, you must set the <code>distribution_strategy</code> to <code>'tpu'</code> .

This will train ResNet for 1 epoch and will complete on a v3-8 TPU node in under 10 minutes. At the end of the training, output similar to the following appears:

To train the ResNet to convergence, run it for 90 epochs as shown in the following script. Training and evaluation are done together. Each epoch has 1251 steps for a total of 112590 training steps and 48 evaluation steps.

Since the training and evaluation was done on the `fake_imagenet` dataset, the output results do not reflect actual output that would appear if the training and evaluation was performed on a real dataset.

At this point, you can either conclude this tutorial and `clean up` (`#clean-up`) your GCP resources, or you can further explore running the model on Cloud TPU Pods.

You can get results faster by scaling your model with Cloud TPU Pods. The fully supported ResNet-50 model can work with the following Pod slices:

- `v2-32`
- `v3-32`

With Cloud TPU Pods, training and evaluation are done together.

Note: If you have already deleted your Compute Engine instance, create a new one following the steps in [Set up your resources](#) (#set_up_your_resources).

1. Disconnect from the Compute Engine instance, if you have not already done so:

Your prompt should now be `user@projectname`, showing you are in the Cloud Shell.

2. Delete the Cloud TPU resource you created for training the model on a single device.

3. Run the `ctpu up` command, using the `tpu-size` parameter to specify the Pod slice you want to use. For example, the following command uses a v3-32 Pod slice.

★ **Note:** If you are not connected to the Compute Engine instance, you can reconnect by running the following command, replacing `vm-name` with the name of your VM:

4. Update the `MODEL_DIR` directory to store the Cloud TPU Pod training data.

5. Specify your TPU name. Unless you specified the `--name` parameter to `ctpu up`, this will be the same as the name for your VM: your username.

6. Add the top-level `/models` folder to the Python path with the command

7. Navigate to the directory:

8. Train the model.

The procedure trains the model on the fake_imagenet dataset to 1 epoch (312 total training steps and 12 evaluation steps). This training takes approximately 2 minutes on a v3-32 Cloud TPU. When the training and evaluation complete, a message similar to the following appears:

To avoid incurring charges to your Google Cloud Platform account for the resources used in this tutorial:

1. Disconnect from the Compute Engine instance, if you have not already done so:

Your prompt should now be `user@projectname`, showing you are in the Cloud Shell.

2. In your Cloud Shell, run `ctpu delete` with the `--zone` flag you used when you set up the Compute Engine VM and Cloud TPU. This deletes both your VM and your Cloud TPU.

★ **Important:** If you set the TPU resources name when you ran `ctpu up`, you must specify that name with the `--name` flag when you run `ctpu delete` in order to shut down your TPU resources.

3. Run `ctpu status` to make sure you have no instances allocated to avoid unnecessary charges for TPU usage. The deletion might take several minutes. A response like the one below indicates there are no more allocated instances:

4. Run `gsutil` as shown, replacing ***bucket-name*** with the name of the Cloud Storage bucket you created for this tutorial:

Note: For free storage limits and other pricing information, see the [Cloud Storage pricing guide \(/storage/pricing\)](/storage/pricing).

- Explore the other [Cloud TPU tutorials \(/tpu/docs/tutorials\)](/tpu/docs/tutorials).
- Learn to use the [TPU monitoring tools in TensorBoard \(/tpu/docs/cloud-tpu-tools\)](/tpu/docs/cloud-tpu-tools).
- See how to train ResNet with [Cloud TPU and GKE \(/tpu/docs/tutorials/kubernetes-engine-resnet\)](/tpu/docs/tutorials/kubernetes-engine-resnet).
- Speed up your training by [streaming the data from Cloud Bigtable \(/tpu/docs/tutorials/bigtable-resnet\)](/tpu/docs/tutorials/bigtable-resnet).