This document describes how to set up and run TensorBoard (https://github.com/tensorflow/tensorboard) for visualizing and analyzing program performance on Cloud TPU.

TensorBoard offers a suite of tools designed to present TensorFlow data visually. When used for monitoring, TensorBoard can help identify bottlenecks in processing and suggest ways to improve performance.

The first time you launch TensorBoard, it can take several minutes to initialize. This can cause the web UI to time out ction error. If that happens, re-launch TensorBoard.

The following instructions assume you have already set up your Cloud TPU in Cloud Shell and are ready to run your training application.

If you don't have a model ready to train, you can get started with the MNIST tutorial (/tpu/docs/tutorials/mnist).

Install the current version of `cloud-tpu-profiler 1.15` (/tpu/docs/cloud-tpu-tools#install_cloud_tpu_profiler) to create the `capture-tpu-profile` script.

When you ran `ctpu up` to create your Compute Engine VM and Cloud TPU, the tool automatically set up port forwarding for the Cloud Shell environment to make TensorBoard available. You need to run Tensorboard in a new Cloud Shell, not the shell that's running your training application.

Follow these steps to run Tensorboard in a separate Cloud Shell:

1. Open a second Cloud Shell to capture profiling data and to start TensorBoard.
2. In the second Cloud Shell, run `ctpu up` to set some needed environment variables on the new shell:

   Note the argument `--zone` is necessary in order for `ctpu up` to correctly find your Compute Engine VM.

   This should return output similar to the following:

3. In the second Cloud Shell, create environment variables for your Cloud Storage bucket and model directory. The model directory variable (`MODEL_DIR`) contains the name of the GCP directory where checkpoints, summaries, and TensorBoard output are stored during model training. For example, `MODEL_DIR=${STORAGE_BUCKET}/model`.

There are two ways you can see TensorBoard trace information, static trace viewer or streaming trace viewer. Static trace viewer is limited to 1 million events per Cloud TPU. If you need to access more events, use streaming trace viewer. Both setups are shown below.

1. In the first Cloud Shell, run your TensorFlow model training application. For example, if you're using the MNIST model, run `mnist_tpu.py` as described in the MNIST tutorial (/tpu/docs/tutorials/mnist).

2. Select the type of trace viewer you want to use: static trace viewer (/tpu/docs/cloud-tpu-tools#trace_viewer), or streaming trace viewer. (/tpu/docs/cloud-tpu-tools#stream_tr_viewer)

3. Perform one of the following procedures:

- Explore the Cloud TPU profiling tools in TensorBoard (https://cloud.google.com/tpu/docs/cloud-tpu-tools#profile_tab).