


Private TensorBoard Quickstart



Step 1: Enable Podman (on LC)

If this is your first time using TensorBoard or Podman, run the following command to enable Podman:

```
enable-podman
```

Step 2: Launch TensorBoard (on LC)

Run the TensorBoard script with the required `-l` option to specify the log directory:

```
/collab/usr/gapps/lcweg/containers/scripts/launch-tensorboard.sh -l </path/to/logs>
```

Step 3: Access TensorBoard (locally)

Once the script runs successfully, it will provide an SSH command and a web address. Run the SSH command from your local machine to establish secure connectivity:

```
ssh <login-node-pub>.llnl.gov -L6006:/var/tmp/<private_socket_directory>/tensorboard.sock
```

Then, open TensorBoard in your browser using the web address:

<http://localhost:6006>

Replace `6006` with the port used in the SSH command, if different.

See: <https://hpc.llnl.gov/software/ai-ml-software-and-tools/using-tensorboard-login-nodes>

Private Tensorboard Architecture

1. Containerize Tensorboard
2. Bind mount your logs
3. Connect the Tensorboard TCP Socket (**Inside the container**) to a UNIX named socket
4. Bind mount the named socket outside the container and use SSH to talk directly to it

