

Getting Started with SLURM and Python Environments on an HPC Cluster

Md Shahriar Forhad
shahriar.forhad.eee@gmail.com
<https://shahriar88.github.io/>



March 8, 2026

Contents

1	Introduction to CRADLE	5
1.1	Introduction	5
1.1.1	Logging In to CRADLE	5
1.2	A Minimal SLURM Job Script Example	6
1.3	Understanding Key #SBATCH Directives	6
1.3.1	Number of Nodes: <code>-N</code>	7
1.3.2	Partition (Queue): <code>-p</code>	7
1.3.3	GPU Request: <code>-gres</code>	7
1.3.4	CPU Cores per Task: <code>-cpus-per-task</code>	7
1.3.5	Time Limit: <code>-time</code>	8
1.3.6	Summary Table of Core Directives	8
1.4	Why Use Python Environments on an HPC Cluster?	8
2	CRADLE COMMAND	11
2.1	Transferring Files to and from the CRADLE Cluster	11
2.1.1	Using <code>scp</code> (Secure Copy)	11
2.1.2	Working with Directories in the Terminal	12
2.1.3	Cloning a GitHub Repository	12
2.1.4	Recommended Best Practices	12
2.2	Cloning a Shared TensorFlow Environment	13
2.2.1	Step 1: Install the cloning tool (one time)	13
2.2.2	Step 2: Clone the shared environment	13
2.2.3	Step 3: Activate the cloned environment	14
2.2.4	Step 4: Add new packages (example: <code>pandas</code>)	14
2.2.5	Step 5: Use the cloned environment in your SLURM job	14
2.2.6	Mental Model Summary	15
2.3	Creating Your Own Python Virtual Environment	15
2.3.1	Option A (Recommended for CPU-Only): Create from Scratch	15
2.3.2	Option B (Best for GPU): Clone an Optimized Shared Environment	16
2.4	Verifying Your Environment and Common Pitfalls	16
2.4.1	Verifying Python and GPU Visibility	16
2.4.2	Common Mistakes to Avoid	17
2.4.3	Quick Decision Guide	17
2.5	Using a PyTorch Environment on CRADLE	17
2.5.1	Minimal GPU Verification Checklist	17
2.6	Using Micromamba and Portable Conda Environments on the HPC Cluster	18
2.6.1	Exporting a Portable Conda Environment on Your PC	18

2.6.2	Copying the YAML File to CRADLE	18
2.6.3	Installing Micromamba in Your Home Directory	19
2.6.4	Creating a Micromamba Environment from the YAML File	19
2.6.5	Activating the Environment on CRADLE	19
2.6.6	Using a Micromamba Environment in a SLURM Job	20
2.6.7	Verifying GPU Support Inside the Environment	21
2.6.8	Best Practices and Common Pitfalls for Micromamba on HPC	21
2.7	Worked Examples: Environments and SLURM Job Scripts	21
2.7.1	Creating a New CPU-Only Environment	22
2.7.2	Cloning a Shared GPU-Optimized TensorFlow Environment	22
2.7.3	Selecting an Environment in a Job Script	22
2.7.4	Installing Libraries After Environment Activation	23
2.7.5	Checking and Verifying Resources	23
2.8	PyTorch DataParallel (DP) on SLURM	26
2.8.1	When to Use DataParallel	26
2.8.2	When <i>Not</i> to Use DataParallel	26
2.8.3	Minimal DataParallel Training Skeleton	27
2.8.4	SLURM Script for DataParallel (single node)	29
2.8.5	DP vs DDP: Practical Comparison	29
2.8.6	Common DP Pitfalls	29
2.8.7	Rule of Thumb	30
2.9	PyTorch Distributed Data Parallel (DDP) on SLURM	30
2.9.1	What You Need (Checklist)	30
2.9.2	Minimal DDP Training Skeleton (single file)	30
2.9.3	SLURM Script for DDP (single node, 4 GPUs)	32
2.9.4	SLURM Script for DDP (2 nodes, 4 GPUs per node)	32
2.9.5	Common DDP Pitfalls (Fast Debug List)	33
3	Example Commands	35
3.1	Example Commands for File Transfer	35
3.1.1	Quick Mapping: Which Example to Use?	44
3.2	Summary	44

Chapter 1

Introduction to CRADLE

1.1 Introduction

For more details <https://www.geeksforgeeks.org/linux-unix/scp-command-in-linux-with-examples/>

This short tutorial explains how to:

- Write a basic SLURM batch script for running a TensorFlow job.
- Understand key `#SBATCH` directives, including nodes, partitions, GPUs, and time limits.
- Use and extend a shared TensorFlow environment on the cluster.
- Create your own Python virtual environments for CPU and GPU workloads.

All examples are designed to be copied directly into a terminal or SLURM job script on a typical HPC system (e.g., a university cluster).

1.1.1 Logging In to CRADLE

Once you have received your user account information, you can log in to the CRADLE cluster from a terminal or command-line prompt using your University credentials.

Step 1: Connect to the University Network (VPN)

If you are off campus, you must first connect to the University network using the approved VPN software. Install and start the VPN client, then log in with your University credentials. Once the VPN connection is active, you can reach the CRADLE login node.

Step 2: SSH into the Cluster

To log in, use `ssh` with your assigned username. In the examples below, replace `username` with your actual cluster username.

```
1 ssh username@login.cradle.university.edu
```

You will be prompted for your password:

```
1 username@login.cradle.university.edu's password:
```

After you enter your password correctly, you should see a prompt indicating that you are now on the cluster login node, similar to:

```
1 [username@login001 ~]$
```

At this point you are logged into CRADLE and ready to work with SLURM, Python environments, and the other tools described in this tutorial.

1.2 A Minimal SLURM Job Script Example

Listing 1.1 shows a basic SLURM job script that runs a TensorFlow test script in a shared environment.

Listing 1.1: Basic SLURM job script with TensorFlow environment

```
1 #!/bin/bash
2
3 # Job identification
4 #SBATCH --job-name=myFirstJob
5
6 # Standard output and error files (%j = job ID)
7 #SBATCH --output=myFirstJob.out.%j
8 #SBATCH --error=myFirstJob.err.%j
9
10 # Resource requests
11 #SBATCH -N 1
12 #SBATCH -p kimg
13 #SBATCH --gres=gpu:1
14 #SBATCH --time=1:00:00
15
16 # Display GPU assigned by SLURM
17 echo "CUDA_VISIBLE_DEVICES: ${CUDA_VISIBLE_DEVICES}"
18
19 # Activate TensorFlow virtual environment
20 echo "Activating TensorFlow-2.6.2 environment"
21 source /shared/tensorflow-2.6.2/tf_env/bin/activate
22
23 # Run TensorFlow test script
24 echo "Running testTF.py"
25 python3 ~/testTFForSlurm/testTF.py
26
27 # Deactivate environment
28 echo "Deactivating TensorFlow-2.6.2 environment"
29 deactivate
```

To submit this job to SLURM, save the script as `myFirstJob.slurm` (for example) and run:

```
1 sbatch myFirstJob.slurm
```

1.3 Understanding Key #SBATCH Directives

The #SBATCH lines at the top of the script tell SLURM what resources you want. They are not shell commands; they are scheduler directives.

1.3.1 Number of Nodes: -N

```
1 #SBATCH -N 1
```

This requests **1 compute node**. A node is a physical machine that may contain many CPU cores, memory, and possibly GPUs.

More examples:

```
1 #SBATCH -N 1 # Single-node TensorFlow training
2 #SBATCH -N 2 # Distributed job across 2 nodes
3 #SBATCH -N 1-4 # SLURM may allocate anywhere from 1 to 4 nodes
```

1.3.2 Partition (Queue): -p

```
1 #SBATCH -p kimq
```

The partition selects the *group of nodes* (queue) your job will run on.

On CRADLE, the GPU partition used in the examples below is `gpua30q`. Always verify available partitions with `sinfo`, since partition names are cluster-specific.

Typical partitions might be:

- `compute` → Standard CPU nodes
- `bigmem` → High-memory nodes
- `gpu` or `gpua30q` → GPU-enabled nodes

Examples:

```
1 #SBATCH -p kimq # Run on GPU nodes in the kimq partition
2 #SBATCH -p compute # Run on CPU-only nodes
```

1.3.3 GPU Request: -gres

```
1 #SBATCH --gres=gpu:1
```

This requests **1 GPU** using the generic resource (GRES) mechanism.

More examples:

```
1 #SBATCH --gres=gpu:1 # One GPU
2 #SBATCH --gres=gpu:2 # Two GPUs
3 #SBATCH --gres=gpu:a100:1 # One NVIDIA A100 GPU (if supported)
```

Important: Partition and GPU must match. If you request GPUs but choose a CPU-only partition (e.g., `-p compute` with `-gres=gpu:1`), SLURM cannot satisfy the request and the job will stay pending or fail.

1.3.4 CPU Cores per Task: -cpus-per-task

```
1 #SBATCH --cpus-per-task=2
```

This requests CPU cores for the main process. In GPU jobs, these CPU cores are commonly used for data loading, preprocessing, and other host-side tasks. If your Python code uses `DataLoader(num_workers > 0)`, you usually need to request enough CPU cores to support it.

More examples:

```
1 #SBATCH --cpus-per-task=2 # Light data loading
2 #SBATCH --cpus-per-task=4 # Moderate preprocessing
3 #SBATCH --cpus-per-task=8 # Heavier data loading or multi-worker pipeline
```

1.3.5 Time Limit: `-time`

```
1 #SBATCH --time=1:00:00
```

This sets a **wall-clock time limit** of 1 hour for your job. When this time is reached, SLURM terminates the job.

More examples:

```
1 #SBATCH --time=00:30:00 # 30 minutes
2 #SBATCH --time=04:00:00 # 4 hours
3 #SBATCH --time=2-00:00:00 # 2 days
```

1.3.6 Summary Table of Core Directives

1.4 Why Use Python Environments on an HPC Cluster?

On HPC clusters, shared software environments (for example, a shared TensorFlow installation) are usually:

- Optimized for the cluster hardware (correct CUDA, cuDNN, drivers).
- Read-only for users.
- Not modifiable by individual users.

You *cannot* safely run:

```
1 pip install pandas
```

inside a shared environment such as

```
1 /shared/tensorflow-2.6.2/tf_env/
```

because:

- It is shared by everyone.
- You do not have write permission.
- Changes could break other users' jobs.

Solution: Clone the shared environment into your home directory, then customize the copy.

Command	Description	Example
Basic Linux Commands		
<code>pwd</code>	Print current working directory.	<code>pwd</code>
<code>ls</code>	List files in current directory.	<code>ls</code>
<code>ls -l</code>	List files with detailed information.	<code>ls -l</code>
<code>cd <dir></code>	Change directory.	<code>cd projects</code>
<code>mkdir -p <dir></code>	Create directory (including parent directories if needed).	<code>mkdir -p data/images</code>
<code>cp <src> <dst></code>	Copy file or directory.	<code>cp file.txt backup/</code>
<code>mv <src> <dst></code>	Move or rename file.	<code>mv old.txt new.txt</code>
<code>rm <file></code>	Remove file.	<code>rm temp.txt</code>
<code>clear</code>	Clear the terminal screen.	<code>clear</code>
<code>history</code>	Show previously executed commands.	<code>history</code>
<code>whoami</code>	Show current username.	<code>whoami</code>
<code>hostname</code>	Show the current machine/node name.	<code>hostname</code>
Account and Security		
<code>passwd</code>	Change your login password.	<code>passwd</code>
SLURM / HPC Commands		
<code>sinfo</code>	Show available partitions and node status.	<code>sinfo</code>
<code>squeue</code>	Show all running jobs.	<code>squeue</code>
<code>squeue -u <user></code>	Show jobs for specific user.	<code>squeue -u username</code>
<code>sbatch <script.sh></code>	Submit batch job.	<code>sbatch train.slurm</code>
<code>srunch -pty bash</code>	Start interactive session on compute node.	<code>srunch -pty bash</code>
<code>scancel <jobid></code>	Cancel a running job.	<code>scancel 12345</code>
<code>scontrol show job <jobid></code>	Display detailed information about a job.	<code>scontrol show job 12345</code>
GPU and Monitoring		
<code>nvidia-smi</code>	Display GPU usage (compute nodes only).	<code>nvidia-smi</code>
<code>tail -f <file></code>	Monitor job output in real time.	<code>tail -f job.out</code>
Environment and Modules		
<code>module avail</code>	List available software modules.	<code>module avail</code>
<code>module load <name></code>	Load a software module.	<code>module load cuda/12.3</code>
<code>module list</code>	Show currently loaded modules.	<code>module list</code>
<code>which python</code>	Show the path of the active Python interpreter.	<code>which python</code>
Storage and Usage		
<code>du -sh *</code>	Show directory sizes.	<code>du -sh *</code>
<code>df -h</code>	Show disk usage.	<code>df -h</code>
<code>quota -s</code>	Show storage quota.	<code>quota -s</code>
<code>less <file></code>	View large files interactively.	<code>less output.log</code>

Table 1.1: Common Linux and SLURM commands for CRADLE HPC usage

Directive	Controls	Example
<code>-N</code>	Number of nodes	<code>-N 1</code>
<code>-p</code>	Hardware pool	<code>-p gpua30q</code>
<code>-gres</code>	GPU count	<code>gpu:1</code>
<code>-cpus-per-task</code>	CPU cores per task	<code>2</code>
<code>-time</code>	Runtime limit	<code>1:00:00</code>

Table 1.2: Core SLURM directives and examples.

Chapter 2

CRADLE COMMAND

2.1 Transferring Files to and from the CRADLE Cluster

When working on the CRADLE HPC cluster, you will often need to move files between your local computer and the cluster. Common examples include:

- Uploading scripts, data files, or configuration files to the cluster.
- Downloading results, logs, or trained models back to your local machine.

This section presents two recommended and commonly used methods:

1. Secure Copy (`scp`)
2. Cloning repositories from GitHub

2.1.1 Using `scp` (Secure Copy)

The `scp` command allows you to securely transfer files between your local machine and the cluster over SSH. It works on Linux, macOS, and Windows (with PowerShell, Git Bash, or WSL).

Transferring a Local File to the Cluster

To copy a file from your local directory to your home directory on the cluster:

```
1 scp localFileInMyDirectory.txt \  
2   yourUserName@login.cradle.university.edu:~/destinationForYourFile/
```

Explanation:

- `localFileInMyDirectory.txt` is the file on your local computer.
- `yourUserName@login.cradle.university.edu` is your cluster login.
- `~/destinationForYourFile/` is the target directory on the cluster.

If the destination directory does not exist, create it first after logging in:

```
1 mkdir -p ~/destinationForYourFil
```

Transferring a File from the Cluster to Your Local Machine

To copy a file from the cluster back to your local directory:

```
1 scp yourUserName@login.cradle.university.edu:~/destinationForYourFile/
   localFileInMyDirectory.txt \
2 .
```

Here, the dot (.) represents your current local directory.

2.1.2 Working with Directories in the Terminal

In both `scp` examples above, the file name refers to a file located in your *current working directory*. You can usually identify this directory by looking at your terminal prompt.

For example: `1 [username@mycomputer:~/Desktop/test]$`

This indicates that your current directory is: `1 ~/Desktop/test`

You can explicitly check your current directory using:

On Linux or macOS: `1 pwd`

This prints something like: `1 /home/username/Desktop/test`

On Windows (PowerShell): `1 cd`

Understanding your current directory is essential to avoid copying the wrong files or encountering file not found errors.

2.1.3 Cloning a GitHub Repository

Instead of manually transferring files, it is often easier to store your project code in a GitHub repository and clone it directly on the cluster.

Example: Cloning a Repository

After logging into the cluster, run:

```
1 git clone https://github.com/emenriquez/testTFForSlurm
```

This creates a new directory called `testTFForSlurm` containing the full repository.

Advantages of using GitHub:

- Version control and change tracking.
- Easy synchronization across multiple machines.
- Clean separation between code and large datasets.

2.1.4 Recommended Best Practices

- Use `scp` for small files, configuration files, and quick transfers.
- Use GitHub for project code and scripts.

- Avoid transferring very large datasets via `scp`; instead, place them directly on shared cluster storage if available.
- Keep your home directory organized (e.g., `projects/`, `data/`, `results/`).

These practices help ensure efficient workflows and reduce errors when working between your local system and the CRADLE cluster.

2.2 Cloning a Shared TensorFlow Environment

2.2.1 Step 1: Install the cloning tool (one time)

```
1 pip3 install --user virtualenv-clone
```

This:

- Installs `virtualenv-clone` into your user space.
- Does not require admin/root access.
- Only needs to be run once.

2.2.2 Step 2: Clone the shared environment

```
1 virtualenv-clone /shared/tensorflow-2.6.2/tf_env ~/myEnvs/tf_env
```

Here:

- `/shared/tensorflow-2.6.2/tf_env/` is the original, optimized, read-only TensorFlow environment.
- `~/myEnvs/tf_env` is your private, writable copy.

What gets copied:

- Python version.
- TensorFlow build (CUDA, cuDNN matched to cluster GPUs).
- All installed dependencies.
- Environment configuration tuned for the cluster.

Advantages:

- Runs as fast as the shared environment.
- You can safely add or upgrade packages.
- No conflicts with other users.

2.2.3 Step 3: Activate the cloned environment

```
1 source ~/myEnvs/tf_env/bin/activate
```

After activation:

- `python` and `pip` point to your cloned environment.
- Any `pip install` affects only your copy.

The shell prompt typically changes to something like:

```
1 (tf_env) username@node \ $
```

2.2.4 Step 4: Add new packages (example: pandas)

With the environment active:

```
1 pip3 install pandas
```

What happens:

- `pandas` is installed only into `~/myEnvs/tf_env`.
- TensorFlow and CUDA support remain intact.
- The customization is isolated and safe.

You can also add:

```
1 pip install scikit-learn matplotlib seaborn tqdm
```

2.2.5 Step 5: Use the cloned environment in your SLURM job

Update the activation lines in your job script.

Old (shared environment):

```
1 echo "Activating TensorFlow-2.6.2 environment"
2 source /shared/tensorflow-2.6.2/tf_env/bin/activate
```

New (your cloned environment):

```
1 echo "Activating my custom environment"
2 source ~/myEnvs/tf_env/bin/activate
```

This ensures:

- Your job sees TensorFlow and any extra packages (e.g., `pandas`).
- You avoid dependency errors at runtime.
- The behavior of your job is reproducible.

Piece	Purpose
Shared environment	Fast, optimized, read-only
Cloned environment	Same speed, user-modifiable
<code>virtualenv-clone</code>	Copies entire environment
<code>pip install</code>	Safe only in cloned environment
Job script	Must activate your chosen environment

Table 2.1: Conceptual roles of shared and cloned environments.

2.2.6 Mental Model Summary

2.3 Creating Your Own Python Virtual Environment

Sometimes you want a completely independent Python environment. There are two main options; here we focus on standard `venv` and the cloned approach.

2.3.1 Option A (Recommended for CPU-Only): Create from Scratch

Step 1: Choose a location

Create a directory to hold your environments:

```
1 mkdir -p ~/myEnvs
```

Step 2: Create the virtual environment

```
1 python3 -m venv ~/myEnvs/myenv
```

This creates something like:

```
1 ~/myEnvs/myenv/
2 |- bin/
3 |- lib/
4 '- pyvenv.cfg
```

Step 3: Activate the environment

```
1 source ~/myEnvs/myenv/bin/activate
```

Your shell prompt changes to:

```
1 (myenv) username@login $
```

Step 4: Upgrade pip (recommended)

```
1 pip install --upgrade pip
```

Step 5: Install packages

Example:

```
1 pip install tensorflow pandas numpy matplotlib
```

Note: On clusters, installing TensorFlow this way may not include GPU support unless CUDA and drivers are correctly matched. For GPU workloads, the cloned shared environment (Option B) is often safer.

Step 6: Use this environment in a SLURM job

In your job script:

```
1 source ~/myEnvs/myenv/bin/activate
2 python train.py
```

2.3.2 Option B (Best for GPU): Clone an Optimized Shared Environment

This is essentially the workflow described in Section 2.2.

Step 1: Install the clone tool:

```
1 pip install --user virtualenv-clone
```

Step 2: Clone the shared environment:

```
1 virtualenv-clone /shared/tensorflow-2.6.2/tf_env ~/myEnvs/tf_env
```

Step 3: Activate it:

```
1 source ~/myEnvs/tf_env/bin/activate
```

Step 4: Add packages:

```
1 pip install pandas scikit-learn seaborn
```

Step 5: Update your SLURM job to activate:

```
1 source ~/myEnvs/tf_env/bin/activate
```

This preserves CUDA/cuDNN compatibility and performance while allowing you to customize the environment.

2.4 Verifying Your Environment and Common Pitfalls

2.4.1 Verifying Python and GPU Visibility

After activating an environment, verify what python and pip you are using:

```
1 which python
2 which pip
```

For TensorFlow GPU detection:

```
1 python -c "import tensorflow as tf; print(tf.config.list_physical_devices('GPU'))"
```

Optional but often required: load CUDA module Some clusters require loading a CUDA module before running GPU-enabled PyTorch or TensorFlow jobs. For example:

```
1 module load cuda/12.3
```

follow the cluster documentation or the instructions provided by your system administrator.

You should see at least one GPU listed when running on a GPU node with correct drivers and environment.

2.4.2 Common Mistakes to Avoid

- Installing packages without activating the environment first.
- Using `pip install -user` inside a virtual environment.
- Mixing Conda environments and `venv` in the same workflow.
- Using the system Python (no virtual environment) for SLURM jobs.
- Requesting GPUs in a CPU-only partition.

2.4.3 Quick Decision Guide

Situation	Best choice
CPU-only work	Option A (new <code>venv</code>)
GPU + TensorFlow	Option B (clone shared env)
Cluster-optimized builds	Option B
Full custom control	Option A

Table 2.2: Choosing between creating and cloning environments.

2.5 Using a PyTorch Environment on CRADLE

If your cluster provides a shared PyTorch environment, activate it in your SLURM script using the path provided by your instructor, system administrator, or cluster documentation.

```
1 source /path/to/shared/pytorch/environment/bin/activate
```

You may also use your own cloned or user-created environment if it includes a CUDA-compatible PyTorch installation.

Then verify that PyTorch can see the GPU:

```
1 python3 - << 'PY'
2 import os, torch
3 print("Torch:", torch.__version__)
4 print("CUDA available:", torch.cuda.is_available())
5 print("Device count:", torch.cuda.device_count())
6 print("CUDA_VISIBLE_DEVICES:", os.environ.get("CUDA_VISIBLE_DEVICES"))
7 if torch.cuda.device_count() > 0:
8     for i in range(torch.cuda.device_count()):
9         print(f"GPU {i}:", torch.cuda.get_device_name(i))
10 PY
```

This check confirms that the active Python environment contains a CUDA-enabled PyTorch build and that the GPU allocated by SLURM is visible inside the job.

2.5.1 Minimal GPU Verification Checklist

For a GPU job, verify the following inside the SLURM script:

1. `echo $CUDA_VISIBLE_DEVICES`

2. `nvidia-smi`
3. Framework-level check such as `torch.cuda.is_available()` or TensorFlow GPU detection

2.6 Using Micromamba and Portable Conda Environments on the HPC Cluster

On many HPC clusters, including CRADLE, traditional Conda/Anaconda installations are discouraged or restricted because they are large, modify shell startup files aggressively, and may conflict with system libraries or CUDA drivers. A lightweight and HPC-friendly alternative is *Micromamba*, a single static binary that implements the Conda ecosystem without requiring a global base environment.

Micromamba allows you to:

- Recreate a Conda environment from a portable YAML file exported on your laptop.
- Install packages into your home directory without administrator rights.
- Activate environments explicitly inside SLURM job scripts.

The overall workflow is:

1. Export a portable Conda environment YAML on your local machine.
2. Copy the YAML file to CRADLE.
3. Install Micromamba in your home directory on CRADLE.
4. Create a Micromamba environment from the YAML file.
5. Activate the environment in your SLURM job scripts.

2.6.1 Exporting a Portable Conda Environment on Your PC

On your local machine where you already use Conda, start by creating a portable YAML description of your environment. Suppose your local environment is called `IR`.

First, activate it:

```
1 conda activate IR
```

Then export a *portable* environment file using `fromhistory`:

```
1 conda env export --from-history > IR_portable.yml
```

This creates a file `IR_portable.yml` containing the channels and high-level dependencies (without build hashes). It is much more likely to recreate successfully on a different system such as CRADLE.

2.6.2 Copying the YAML File to CRADLE

Use `scp` or an equivalent file transfer method to copy `IR_portable.yml` to your home directory on CRADLE. For example, from your local terminal:

```
1 scp IR_portable.yml username@login.cradle.university.edu:~
```

After logging into CRADLE, you should see the file in your home directory:

```
1 ls ~/IR_portable.yml
```

2.6.3 Installing Micromamba in Your Home Directory

Micromamba is distributed as a single binary, which makes it ideal for installation in a user home directory (no root access required).

Create a directory for the binary:

```
1 mkdir -p ~/bin
2 cd ~/bin
```

Download and unpack the latest Micromamba binary (Linux 64-bit):

```
1 curl -L https://micro.mamba.pm/api/micromamba/linux-64/latest \
2 | tar -xvj bin/micromamba
```

Add `$HOME/bin` to your `PATH` in `~/.bashrc`:

```
1 echo 'export PATH=$HOME/bin:$PATH' >> ~/.bashrc
2 source ~/.bashrc
```

Verify that `micromamba` is available:

```
1 micromamba --version
```

2.6.4 Creating a Micromamba Environment from the YAML File

Choose a directory where Micromamba will store its environments, for example:

```
1 mkdir -p ~/micromamba/envs
```

Initialize Micromamba for Bash (one time):

```
1 micromamba shell init -s bash -p ~/micromamba
2 source ~/.bashrc
```

Now create a new environment (for example, `IR_NEW`) from your portable YAML file:

```
1 micromamba create \
2 -f ~/IR_portable.yml \
3 -n IR_NEW \
4 -p ~/micromamba/envs/IR_NEW
```

Here:

- `~/IR_portable.yml` is the YAML exported on your PC.
- `IR_NEW` is the environment name.
- `~/micromamba/envs/IR_NEW` is the full path Micromamba uses.

2.6.5 Activating the Environment on CRADLE

Once created, you can activate the environment on the CRADLE login node:

```
1 micromamba activate ~/micromamba/envs/IR_NEW
```

After activation:

- `python` and `pip` refer to the Micromamba environment.
- Any installs via `pip` or `micromamba install` affect only that environment.

You can verify:

```
1 which python
2 python --version
3 pip list
```

2.6.6 Using a Micromamba Environment in a SLURM Job

In a SLURM job script, you should explicitly set the Micromamba root and activate the environment before running Python code. For example:

```
1 #!/bin/bash
2 #SBATCH --job-name=ir_gpu_job
3 #SBATCH --output=ir_gpu_job.out.%j
4 #SBATCH --error=ir_gpu_job.err.%j
5 #SBATCH -N 1
6 #SBATCH -p kimg # GPU partition
7 #SBATCH --gres=gpu:1 # 1 GPU
8 #SBATCH --time=02:00:00 # 2 hours
9
10 # Ensure Micromamba root is defined
11 export MAMBA_ROOT_PREFIX=$HOME/micromamba
12
13 # Load shell initialization (if micromamba shell init was used)
14 source $HOME/.bashrc
15
16 echo "Host: $(hostname)"
17 echo "CUDA_VISIBLE_DEVICES: $CUDA_VISIBLE_DEVICES"
18
19 echo "Activating Micromamba environment IR_NEW"
20 micromamba activate $HOME/micromamba/envs/IR_NEW
21
22 echo "Checking TensorFlow GPU visibility"
23 python -c "import tensorflow as tf; print(tf.config.list_physical_devices('GPU'))"
24
25 echo "Running training script"
26 python ~/projects/train_ir_model.py
```

This pattern ensures that:

- The correct Micromamba environment is active inside the job.
- GPU visibility can be checked via TensorFlow (or PyTorch).
- The job is reproducible and does not depend on the login-node Conda setup.

2.6.7 Verifying GPU Support Inside the Environment

To confirm that your Micromamba environment has working GPU support for TensorFlow (or another framework), you can run:

```
1 python -c "import tensorflow as tf; print(tf.config.list_physical_devices('GPU'))"
```

On a GPU node with correct drivers and CUDA libraries, this should list one or more GPU devices. You can also check the low-level GPU visibility with:

```
1 nvidia-smi
```

2.6.8 Best Practices and Common Pitfalls for Micromamba on HPC

Best practices:

- Export environments from Conda using `-from-history` or `-no-builds` to avoid hardware-specific build strings.
- Use Micromamba to recreate the environment on the cluster, not full Anaconda.
- Keep all Micromamba files under `~/micromamba` or a dedicated directory in your home space.
- Activate environments explicitly inside SLURM job scripts.

Common pitfalls:

- Installing full Anaconda on an HPC login node (disk usage and conflicts).
- Relying on Conda's `base` environment in batch jobs.
- Mixing Conda and Micromamba environments in the same workflow.
- Installing GPU frameworks with incompatible `cuda-toolkit` versions; always check which CUDA versions the cluster supports.

Using Micromamba with a portable Conda YAML exported from your local machine gives you a reproducible, HPC-safe way to manage Python environments on CRADLE while respecting cluster policies and maximizing compatibility with the installed NVIDIA drivers and CUDA libraries.

2.7 Worked Examples: Environments and SLURM Job Scripts

This section provides concrete, copy-pasteable examples covering:

1. Creating a new environment.
2. Cloning a shared TensorFlow environment.
3. Selecting (activating) an environment.
4. Installing extra libraries.
5. Checking and verifying resources.
6. Writing SLURM scripts for different resource requirements.

2.7.1 Creating a New CPU-Only Environment

```

1 # 1. Create a directory to store your environments
2 mkdir -p ~/myEnvs
3
4 # 2. Create a new CPU-only Python virtual environment
5 python3 -m venv ~/myEnvs/cpu_env
6
7 # 3. Activate the environment
8 source ~/myEnvs/cpu_env/bin/activate
9
10 # 4. Upgrade pip
11 pip install --upgrade pip
12
13 # 5. Install some common libraries
14 pip install numpy scipy pandas matplotlib scikit-learn
15
16 # 6. When finished, deactivate
17 deactivate

```

2.7.2 Cloning a Shared GPU-Optimized TensorFlow Environment

```

1 # 1. Install the clone tool (only once per user)
2 pip3 install --user virtualenv-clone
3
4 # 2. Create a directory to store your environments (if not already created)
5 mkdir -p ~/myEnvs
6
7 # 3. Clone the shared TensorFlow environment into your home directory
8 virtualenv-clone /shared/tensorflow-2.6.2/tf_env ~/myEnvs/tf_env
9
10 # 4. Activate the cloned environment
11 source ~/myEnvs/tf_env/bin/activate
12
13 # 5. Install extra libraries needed for your project
14 pip install pandas scikit-learn matplotlib seaborn tqdm
15
16 # 6. Deactivate when done
17 deactivate

```

2.7.3 Selecting an Environment in a Job Script

Activate the CPU-only environment:

```

1 # Inside your SLURM job script
2 echo "Activating CPU-only environment"
3 source ~/myEnvs/cpu_env/bin/activate

```

Activate the cloned GPU TensorFlow environment:

```

1 # Inside your SLURM job script
2 echo "Activating GPU TensorFlow environment"
3 source ~/myEnvs/tf_env/bin/activate

```

2.7.4 Installing Libraries After Environment Activation

Inside the CPU-only environment:

```
1 source ~/myEnvs/cpu_env/bin/activate
2 pip install jupyter notebook ipykernel
3 deactivate
```

Inside the cloned GPU environment:

```
1 source ~/myEnvs/tf_env/bin/activate
2 pip install opencv-python pillow tensorboard
3 deactivate
```

2.7.5 Checking and Verifying Resources

This subsection shows how to verify what resources are available on the cluster and what resources SLURM has actually allocated to your jobs.

Check Available Partitions and Nodes

To see which partitions and nodes exist and their current status:

```
1 # Show all partitions and their status
2 sinfo
```

For more detailed information about each node:

```
1 # Detailed node information
2 scontrol show nodes
```

Check Your Jobs

To see all of your jobs in the queue:

```
1 # Show jobs for your user
2 squeue -u $USER
```

To inspect a specific job by its job ID:

```
1 # Show a specific job
2 squeue -j <JOBID>
```

Inspect Resources Allocated by SLURM

To see precisely what resources were allocated to a job:

```
1 # Show full job allocation details
2 scontrol show job <JOBID>
```

Look for fields such as NumNodes, NumCPUs, Gres, and NodeList.

Check Resources Inside a Job Script

Inside a running job, you can query SLURM environment variables to see what was allocated:

```

1 echo "Job ID: $SLURM_JOB_ID"
2 echo "Nodes allocated: $SLURM_NNODES"
3 echo "Node list: $SLURM_NODELIST"
4 echo "CPUs per node: $SLURM_CPUS_ON_NODE"
5 echo "Tasks per node: $SLURM_TASKS_PER_NODE"
6 echo "Job partition: $SLURM_JOB_PARTITION"

```

Check GPU Allocation

For GPU jobs, it is important to verify which GPUs are visible:

```

1 # GPUs assigned by SLURM (logical IDs)
2 echo "CUDA_VISIBLE_DEVICES: $CUDA_VISIBLE_DEVICES"
3
4 # List visible GPUs and utilization
5 nvidia-smi

```

Check CPU and Memory Usage

You can inspect CPU and memory characteristics on the node:

```

1 # CPU info
2 lscpu
3
4 # Memory info
5 free -h
6
7 # Per-node memory summary (first lines of /proc/meminfo)
8 cat /proc/meminfo | head

```

Check Resources Used by Completed Jobs

After a job finishes, accounting information can be retrieved with `sacct`:

```

1 # Accounting info for a finished job
2 sacct -j <JOBID> --format=JobID,JobName,Partition,AllocCPUS,Elapsed,MaxRSS,State

```

This shows how many CPUs were allocated, how long the job ran, peak memory usage (`MaxRSS`), and the final state (e.g., `COMPLETED`, `FAILED`, `TIMEOUT`).

Minimal Resource Debug Job

The following SLURM script does nothing except report which resources it received. It is useful when testing new partitions or debugging allocation issues:

```

1 #!/bin/bash
2 #SBATCH --job-name=resource_debug
3 #SBATCH -N 1
4 #SBATCH -p kimg
5 #SBATCH --gres=gpu:1
6 #SBATCH --time=00:10:00
7
8 echo "Job ID: $SLURM_JOB_ID"
9 echo "Node list: $SLURM_NODELIST"
10 echo "CPUs/node: $SLURM_CPUS_ON_NODE"
11 echo "GPUs visible: $CUDA_VISIBLE_DEVICES"
12
13 echo "==== nvidia-smi ===="
14 nvidia-smi
15
16 echo "==== lscpu ===="
17 lscpu
18
19 echo "==== memory ===="
20 free -h

```

Aspect	CPU	Single GPU	DataParallel (DP)	DistributedDataParallel (DDP)
Hardware	CPU cores only	1 GPU	Multiple GPUs (1 node)	Multiple GPUs (1+ nodes)
Processes	1	1	1	1 per GPU
GPU usage	None	Full single GPU	Automatic batch split	Explicit per-rank GPU
Multi-node support	No	No	No	Yes
Performance	Slowest	Fast	Medium	Fastest
Scalability	Poor	Limited	Limited	Excellent
Communication	None	None	CPU-based gather/scatter	NCCL all-reduce
Gradient sync	N/A	N/A	Implicit (inefficient)	Explicit (efficient)
Memory efficiency	High	Medium	Low (replication overhead)	High
Fault isolation	N/A	N/A	Poor	Good
Code complexity	Very low	Low	Low	Medium
Launch method	<code>python</code>	<code>python</code>	<code>python</code>	<code>torchrun</code> / <code>srun</code>
Best use case	Debugging, tests	Small models	Prototyping	Large-scale training
Recommended for HPC	No	Limited	No	Yes

Table 2.3: Comparison of CPU, single-GPU, DataParallel (DP), and DistributedDataParallel (DDP) execution modes in PyTorch.

Aspect	CPU	Single GPU	DataParallel (DP)	DistributedDataParallel (DDP)
Hardware	CPU cores only	1 GPU	Multiple GPUs (1 node)	Multiple GPUs (1+ nodes)
Processes	1	1	1	1 per GPU
GPU usage	None	Full single GPU	Automatic batch split	Explicit per-rank GPU
Multi-node support	No	No	No	Yes
Performance	Slowest	Fast	Medium	Fastest
Scalability	Poor	Limited	Limited	Excellent
Communication	None	None	Gather/scatter overhead	NCCL all-reduce
Gradient sync	N/A	N/A	Implicit (less efficient)	Explicit (efficient)
Memory efficiency	High	Medium	Lower (replication overhead)	Medium (replicated per GPU)
Fault isolation	N/A	N/A	Poor	Better
Code complexity	Very low	Low	Low	Medium
Launch method	<code>python</code>	<code>python</code>	<code>python</code>	<code>torchrun/srun</code>
Best use case	Debugging, tests	Small models	Prototyping	Large-scale training
Recommended for HPC	No	Limited	No	Yes

Table 2.4: Comparison of CPU, single-GPU, DataParallel (DP), and DistributedDataParallel (DDP) execution modes in PyTorch.

2.8 PyTorch DataParallel (DP) on SLURM

This section describes PyTorch **DataParallel (DP)**, a simpler multi-GPU approach compared to Distributed Data Parallel (DDP). DP uses a *single Python process* and automatically splits each batch across multiple GPUs on the same node.

Important: DP is suitable for quick experiments on a *single node* only. For serious training, large models, or multi-node jobs, prefer DDP (Section 2.9).

2.8.1 When to Use DataParallel

- Single node with multiple GPUs.
- Prototyping or debugging multi-GPU behavior.
- Minimal code changes from single-GPU training.

2.8.2 When *Not* to Use DataParallel

- Multi-node training (DP does not scale across nodes).
- Large models or long training runs (DDP is faster and more stable).
- When precise performance benchmarking matters.

2.8.3 Minimal DataParallel Training Skeleton

Save as `train_dp.py`. This example assumes a single node with multiple GPUs.

Listing 2.1: Minimal PyTorch DataParallel training skeleton (train_dp.py)

```
1 import torch
2 import torch.nn as nn
3 from torch.utils.data import DataLoader
4
5 def build_model():
6     # Replace with your FCN / VAE model
7     return nn.Identity()
8
9 def build_dataset(split="train"):
10    # Replace with your dataset
11    raise NotImplementedError
12
13 def main():
14    device = torch.device("cuda" if torch.cuda.is_available() else "cpu")
15
16    model = build_model()
17
18    if torch.cuda.device_count() > 1:
19        print(f"Using DataParallel on {torch.cuda.device_count()} GPUs")
20        model = nn.DataParallel(model)
21
22    model = model.to(device)
23
24    train_ds = build_dataset("train")
25    train_loader = DataLoader(
26        train_ds,
27        batch_size=16,
28        shuffle=True,
29        num_workers=4,
30        pin_memory=True,
31        drop_last=True,
32    )
33
34    optim = torch.optim.Adam(model.parameters(), lr=1e-3)
35
36    model.train()
37    for epoch in range(5):
38        for batch in train_loader:
39            X, Y = batch
40            X = X.to(device, non_blocking=True)
41            Y = Y.to(device, non_blocking=True)
42
43            optim.zero_grad(set_to_none=True)
44            Yhat = model(X)
45            loss = torch.mean((Yhat - Y) ** 2)
46            loss.backward()
47            optim.step()
48
49            print(f"Epoch {epoch} done. loss={loss.item():.6f}")
50
51 if __name__ == "__main__":
52    main()
```

2.8.4 SLURM Script for DataParallel (single node)

DP requires **one Python process only**. Do *not* use `torchrun` or multiple tasks.

Listing 2.2: SLURM script for DataParallel (`dp_1node.slurm`)

```

1  #!/bin/bash
2  #SBATCH --job-name=dp_1node
3  #SBATCH --output=dp_1node.out.%j
4  #SBATCH --error=dp_1node.err.%j
5  #SBATCH -N 1
6  #SBATCH -p kimg
7  #SBATCH --gres=gpu:4
8  #SBATCH --cpus-per-task=8
9  #SBATCH --time=04:00:00
10
11  echo "Host: \$(hostname)"
12  echo "CUDA_VISIBLE_DEVICES: \${CUDA_VISIBLE_DEVICES}"
13
14  source ~/myEnvs/tf_env/bin/activate
15
16  export OMP_NUM_THREADS=\${SLURM_CPUS_PER_TASK}
17
18  # Single process only
19  python train_dp.py
20
21  deactivate

```

2.8.5 DP vs DDP: Practical Comparison

Feature	DataParallel (DP)	DistributedDataParallel (DDP)
Processes	1 total	1 per GPU
Multi-node support	No	Yes
Performance	Slower	Faster
Scalability	Limited	Excellent
Failure isolation	Poor	Good
Recommended for	Prototyping	Production / research

Table 2.5: Comparison between DataParallel and DistributedDataParallel.

2.8.6 Common DP Pitfalls

- Do not combine DP with `torchrun` or `srun -n > 1`.
- Expect slower scaling beyond 2–4 GPUs.
- Avoid DP for multi-node jobs.
- Always move the model to GPU *after* wrapping with `DataParallel`.

2.8.7 Rule of Thumb

- **1 node, quick test** → DataParallel
- **Anything serious or multi-node** → DDP

2.9 PyTorch Distributed Data Parallel (DDP) on SLURM

This section shows a minimal, practical recipe for PyTorch DDP on an HPC cluster. DDP runs one Python process per GPU and synchronizes gradients efficiently.

2.9.1 What You Need (Checklist)

- Request **multiple GPUs** in SLURM (e.g., `-gres=gpu:4`).
- Launch with `torchrun` (recommended) or `srun` + environment variables.
- In Python: call `init_process_group`, set the per-rank GPU, wrap model with `DistributedDataParallel`.
- Use `DistributedSampler` for training (and validation if you want correct global metrics quickly).

2.9.2 Minimal DDP Training Skeleton (single file)

Save as `train_ddp.py`. Replace `build_model()`, `build_dataset()`, and the loss/optimizer with your own code.

Listing 2.3: Minimal PyTorch DDP training skeleton (train_ddp.py)

```

1 import os
2 import torch
3 import torch.nn as nn
4 import torch.distributed as dist
5 from torch.nn.parallel import DistributedDataParallel as DDP
6 from torch.utils.data import DataLoader
7 from torch.utils.data.distributed import DistributedSampler
8
9 # -----
10 # 1) DDP setup / teardown
11 # -----
12 def ddp_setup():
13     """
14     Works with: torchrun (sets RANK, LOCAL_RANK, WORLD_SIZE, MASTER_ADDR, MASTER_PORT
15     ).
16     """
17     if dist.is_available() and not dist.is_initialized():
18         dist.init_process_group(backend="nccl") # GPUs -> nccl
19
20     rank = int(os.environ.get("RANK", "0"))
21     local_rank = int(os.environ.get("LOCAL_RANK", "0"))
22     world_size = int(os.environ.get("WORLD_SIZE", "1"))
23
24     if torch.cuda.is_available():
25         torch.cuda.set_device(local_rank)
26         device = torch.device(f"cuda:{local_rank}")
27     else:
28         device = torch.device("cpu")
29
30     return device, rank, local_rank, world_size
31
32 def ddp_cleanup():
33     if dist.is_available() and dist.is_initialized():
34         dist.destroy_process_group()
35
36 # -----
37 # 2) Your builders (replace)
38 # -----
39 def build_model():
40     # Replace with your FCN/VAE model builder
41     return nn.Identity()
42
43 def build_dataset(split="train"):
44     # Replace with your dataset (e.g., CleanImageFolder / PairedImageFolder)
45     # Must return a torch.utils.data.Dataset
46     raise NotImplementedError
47
48 # -----
49 # 3) Main training loop
50 # -----
51 def main():
52     device, rank, local_rank, world = ddp_setup()
53     is_ddp = (world > 1)
54
55     model = build_model().to(device)
56     if is_ddp:
57         model = DDP(model, device_ids=[local_rank], output_device=local_rank)

```

2.9.3 SLURM Script for DDP (single node, 4 GPUs)

This launches **one process per GPU** on the same node.

Listing 2.4: SLURM script for DDP on 1 node (`ddp_1node.slurm`)

```

1  #!/bin/bash
2  #SBATCH --job-name=ddp_1node
3  #SBATCH --output=ddp_1node.out.%j
4  #SBATCH --error=ddp_1node.err.%j
5  #SBATCH -N 1
6  #SBATCH -p kimg
7  #SBATCH --gres=gpu:4
8  #SBATCH --cpus-per-task=4
9  #SBATCH --time=04:00:00
10
11 echo "Host: \$(hostname)"
12 echo "Node list: \${SLURM_NODELIST}"
13 echo "CUDA_VISIBLE_DEVICES: \${CUDA_VISIBLE_DEVICES}"
14
15 source ~/myEnvs/tf_env/bin/activate
16
17 # One process per GPU:
18 export OMP_NUM_THREADS=\${SLURM_CPUS_PER_TASK}
19
20 # torchrun is recommended.
21 # --standalone works for single-node. For multi-node, use rendezvous settings shown
   below.
22 srun torchrun \
23   --standalone \
24   --nproc_per_node=4 \
25   train_ddp.py
26
27 deactivate

```

2.9.4 SLURM Script for DDP (2 nodes, 4 GPUs per node)

For multi-node DDP, you need a rendezvous endpoint (master address/port). A common pattern is to use the first node in `\${SLURM_NODELIST}` as the master.

Listing 2.5: SLURM script for DDP on 2 nodes (ddp_2node.slurm)

```

1  #!/bin/bash
2  #SBATCH --job-name=ddp_2node
3  #SBATCH --output=ddp_2node.out.%j
4  #SBATCH --error=ddp_2node.err.%j
5  #SBATCH -N 2
6  #SBATCH -p kimg
7  #SBATCH --gres=gpu:4
8  #SBATCH --cpus-per-task=4
9  #SBATCH --time=06:00:00
10
11  source ~/myEnvs/tf_env/bin/activate
12
13  # Pick master node and port
14  MASTER_ADDR=$(scontrol show hostnames "$SLURM_NODELIST" | head -n 1)
15  MASTER_PORT=29500
16
17  echo "MASTER_ADDR=$MASTER_ADDR"
18  echo "MASTER_PORT=$MASTER_PORT"
19
20  export OMP_NUM_THREADS=$SLURM_CPUS_PER_TASK
21
22  # Total processes = nodes * gpus_per_node = 2 * 4 = 8
23  srun torchrun \
24  --nnodes=2 \
25  --nproc_per_node=4 \
26  --rdzv_backend=c10d \
27  --rdzv_endpoint=$MASTER_ADDR:$MASTER_PORT \
28  train_ddp.py
29
30  deactivate

```

2.9.5 Common DDP Pitfalls (Fast Debug List)

- **Do not use DataParallel with DDP.** Use one or the other.
- **Call DistributedSampler(...)** for training and call `set_epoch(epoch)` each epoch.
- **Set the per-process GPU** via `LOCAL_RANK` and `torch.cuda.set_device(local_rank)`.
- **Print/checkpoint only on rank 0** to avoid duplicates and file corruption.
- If validation metrics must be global, either:
 - use `DistributedSampler` for validation and `all_reduce` metric sums, or
 - run validation only on rank 0 (simpler but slower).

Chapter 3

Example Commands

3.1 Example Commands for File Transfer

This chapter provides practical, copy-pasteable examples for transferring files and directories between your local machine and the CRADLE cluster. These commands are typically run from a local terminal, not inside a SLURM job.

Useful HPC Commands for SLURM, Modules, Conda, and Remote Workflows

The following commands are commonly used when working on an HPC cluster. They help users inspect cluster resources, manage software modules, configure Conda environments, transfer files, and monitor SLURM jobs.

For additional examples of secure file transfer using `scp`, see the following reference:
<https://www.geeksforgeeks.org/linux-unix/scp-command-in-linux-with-examples/>

```
1
2 #####
3 # Inspect Cluster Resources
4 #####
5
6 # Show available partitions (queues), nodes, and their status
7 sinfo
8
9 # Example output fields:
10 # PARTITION -> queue name
11 # AVAIL -> whether it is available
12 # NODES -> number of nodes
13 # STATE -> idle, mix, alloc, drain
14
15
16 #####
17 # Environment Modules Management
18 #####
19
20 # Show all available software modules
21 module avail
22
23 # Show modules currently loaded in your environment
24 module list
```

```
25
26 # Load a software module (example: CUDA toolkit)
27 module load cuda/12.3
28
29 # Unload a specific module
30 module unload cuda/12.3
31
32 # Remove all loaded modules and reset the module environment
33 module purge
34
35
36 #####
37 # Creating Project Directories
38 #####
39
40 # Create a main project directory
41 mkdir -p ~/project_workspace
42
43 # Create a directory for storing datasets
44 ## -p means create parent directories as needed
45 mkdir -p ~/project_workspace/Data_Source
46
47
48 #####
49 # Installing Conda (Optional)
50 #####
51
52 # Download the full Anaconda distribution
53 wget https://repo.anaconda.com/archive/Anaconda3-2025.12-2-Linux-x86_64.sh
54
55 # OR download the lightweight Miniconda installer
56 wget https://repo.anaconda.com/miniconda/Miniconda3-latest-Linux-x86_64.sh
57
58
59 #####
60 # Conda Environment Setup
61 #####
62
63 # Enable conda in the current shell session
64 # (required if conda is installed but not automatically initialized)
65 eval "$(conda shell.bash hook)"
66
67 # Verify conda installation
68 conda --version
69
70 # Check which Python interpreter is currently active
71 # Useful for confirming the active Conda environment
72 which python
73 which python3
74
75 # Make conda initialization permanent for future logins
76 conda init bash
77
78
```

```
79 #####
80 # Disable Automatic Base Environment
81 #####
82
83 # Prevent the base environment from activating automatically
84 conda config --set auto_activate_base false
85
86 # Reload shell configuration
87 source ~/.bashrc
88
89
90
91
92
93
94
95 #####
96 # Running Jupyter Notebook on HPC
97 #####
98
99 # Start a Jupyter notebook on a remote machine
100 jupyter notebook --no-browser --port=8888
101
102
103 #####
104 # Create SSH Tunnel from Local Machine
105 #####
106
107 # Forward local port 8888 to the remote HPC node
108 ssh -L 8888:localhost:8888 username@login.cluster.edu
109
110 # If Jupyter is running on a compute node, first identify that node name
111 hostname
112
113 # Then create an SSH tunnel to that specific node through the login node
114 ssh -L 8888:localhost:8888 -J username@login.cluster.edu username@compute-node
115
116 # Open the URL printed by Jupyter in your local browser
117 http://localhost:8888/tree?token=YOUR_TOKEN_HERE
118
119
120
121 #####
122 # Copy Files and Folders to HPC with scp
123 #####
124
125 # Step 1: Log in to the HPC cluster and create destination directories
126 ssh yourUserName@login.cradle.university.edu
127
128 mkdir -p ~/project_workspace
129 mkdir -p ~/project_workspace/envs
130
131 # Step 2: From your local computer, copy files to the cluster
132 # Quotes are used because Windows paths often contain spaces
```

```
133
134 # Copy an entire local directory to the HPC cluster
135 # The -r option means "recursive", so the full folder and its contents are copied
136 scp -r "C:\Users\...\Data_Source" \
137     yourUserName@login.cradle.university.edu:/home/yourUserName/project_workspace
138
139 # Copy a single local text file
140 scp "C:\Users\...\mask_rcnn\install_steps.txt" \
141     yourUserName@login.cradle.university.edu:/home/yourUserName/project_workspace/
142     envs
143
144 # Copy another single local file
145 scp "C:\Users\...\Projects\HPC\install_hpc.txt" \
146     yourUserName@login.cradle.university.edu:/home/yourUserName/project_workspace/
147     envs
148 #####
149 # Fix Windows Line Ending Issues
150 #####
151
152 # Convert Windows CRLF line endings to Linux format
153 dos2unix job_script.sh
154
155 # Alternative method if dos2unix is unavailable
156 sed -i 's/\r$/\n/' job_script.sh
157
158 #####
159 # SLURM Job Monitoring
160 #####
161
162 # Check whether a specific job is still in the queue
163 squeue -j JOB_ID
164
165 # Show accounting information for a finished job
166 # (may be disabled on some clusters)
167 sacct -j JOB_ID
168
169 #####
170 # Inspect Job Output Files
171 #####
172
173 # List files in the current directory
174 ls
175
176 # Display the job's standard output
177 cat job_output.out
178
179 # Display the job's error output
180 cat job_error.err
181
182 # View large output files interactively
```

```
185 less job_output.out
```

These commands form a typical workflow when interacting with an HPC cluster: checking available resources, managing software modules and environments, running interactive applications such as Jupyter, and monitoring SLURM jobs.

Copying a Single File from Local Machine to CRADLE

```
1 scp my_script.py \  
2 yourUserName@login.cradle.university.edu:/home/yourUserName/projects/
```

This copies `my_script.py` from your current local directory into `/home/yourUserName/projects/` on the cluster.

Copying a Single File from CRADLE to Local Machine

```
1 scp yourUserName@login.cradle.university.edu:/home/yourUserName/projects/results.txt \  
2 .
```

The dot (`.`) represents your current local directory.

Copying an Entire Directory to CRADLE

To recursively copy a directory (for example, a project folder):

```
1 scp -r myProjectFolder/ \  
2 yourUserName@login.cradle.university.edu:/home/yourUserName/projects/
```

This creates `/home/yourUserName/projects/myProjectFolder/` on the cluster.

Copying an Entire Directory from CRADLE to Local Machine

```
1 scp -r yourUserName@login.cradle.university.edu:/home/yourUserName/projects/  
myProjectFolder/ \  
2 ~/Desktop/
```

This downloads the entire project directory to your local desktop.

Transferring Job Output After Completion

After a SLURM job finishes, you often want to retrieve output files such as logs or trained models.

```
1 scp yourUserName@login.cradle.university.edu:/home/yourUserName/slurm_outputs/myJob.  
out.123456 \  
2 .
```

Here, `123456` is the SLURM job ID.

Copying Files Using Wildcards

You can use wildcards to transfer multiple files at once:

```
1 scp yourUserName@login.cradle.university.edu:/home/yourUserName/slurm_outputs/*.out.* \  
2 .
```

This copies all SLURM output files in the directory.

Recommended Workflow with SLURM Jobs

A typical workflow looks like this:

1. Copy scripts to CRADLE using `scp` or `git clone`.
2. Submit the job using `sbatch`.
3. Monitor the job with `squeue`.
4. Copy results back to your local machine using `scp`.

Important Note: File transfers should generally be performed from the *login node*, not from inside running compute jobs, unless explicitly required.

Installing Anaconda or Miniconda on CRADLE

Download the installer from the login node:

```
1 cd /home/yourUserName
2
3 # Full Anaconda distribution
4 wget https://repo.anaconda.com/archive/Anaconda3-2025.12-2-Linux-x86_64.sh
5
6 # OR lightweight Miniconda installer
7 wget https://repo.anaconda.com/miniconda/Miniconda3-latest-Linux-x86_64.sh
```

Run the installer:

```
1 # Example for Anaconda
2 bash Anaconda3-2025.12-2-Linux-x86_64.sh
3
4 # Example for Miniconda
5 bash Miniconda3-latest-Linux-x86_64.sh
```

Activating Conda in the Current Session

If you chose not to initialize Conda automatically during installation, enable it manually in the current shell:

```
1 eval "$(/home/yourUserName/anaconda3/bin/conda shell.bash hook)"
2 conda --version
```

Making Conda Available for Future Logins

To make Conda automatically available in every login session:

```
1 /home/yourUserName/anaconda3/bin/conda init bash
2 source ~/.bashrc
3 conda --version
```

Example SLURM Job Script Using a Conda Environment

The following example shows how to run a GPU-enabled Python program on the CRADLE cluster using a Conda environment.

```

1  #!/bin/bash
2  #SBATCH --job-name=maskrcnn
3  #SBATCH --output=maskrcnn.out.%j
4  #SBATCH --error=maskrcnn.err.%j
5  #SBATCH -N 1
6  #SBATCH -p gpua30q
7  #SBATCH --gres=gpu:1
8  #SBATCH --cpus-per-task=2
9  #SBATCH --time=02:00:00
10
11 echo "Host: $(hostname)"
12 echo "CUDA_VISIBLE_DEVICES: $CUDA_VISIBLE_DEVICES"
13
14 # initialize conda
15 # source ~/anaconda3/etc/profile.d/conda.sh
16 # OR
17 source \$(conda info --base)/etc/profile.d/conda.sh
18
19 # activate environment
20 conda activate mask_rcnn
21
22 echo "Python path: \$(which python)"
23
24 python train_maskrcnn.py

```

This script requests one GPU from the gpua30q partition and activates the mask_rcnn Conda environment before running the training script.

Launching Jupyter Notebook on CRADLE

Start Jupyter Notebook on the remote machine (login or compute node):

```
1 jupyter notebook --no-browser --port=8888
```

The output will display a URL containing an authentication token, for example:

```
1 http://localhost:8888/tree?token=2c23d0
```

Creating an SSH Tunnel from Your Local Machine

On your local computer (Windows, macOS, or Linux), create a secure tunnel:

```
1 ssh -L 8888:localhost:8888 user@login.cradle.university.edu
```

This forwards local port 8888 on your computer to port 8888 on the CRADLE server.

Accessing Jupyter in Your Web Browser

After the SSH tunnel is established, open your browser and navigate to:

```
1 http://localhost:8888/tree?token=2c23d0
```

Paste the token exactly as shown in the Jupyter output on CRADLE.

Single Node, CPU-Only Job

```
1 #!/bin/bash
2 #SBATCH --job-name=cpu_only_job
3 #SBATCH --output=cpu_only_job.out.%j
4 #SBATCH --error=cpu_only_job.err.%j
5 #SBATCH -N 1
6 #SBATCH -p compute # CPU-only partition
7 #SBATCH --time=01:00:00 # 1 hour
8
9 echo "Host: $(hostname)"
10 echo "Using CPU-only environment"
11 source ~/myEnvs/cpu_env/bin/activate
12
13 echo "Running CPU-only Python script"
14 python ~/projects/my_cpu_script.py
15
16 echo "Deactivating environment"
17 deactivate
```

Single Node, Single GPU Job

```
1 #!/bin/bash
2 #SBATCH --job-name=single_gpu_job
3 #SBATCH --output=single_gpu_job.out.%j
4 #SBATCH --error=single_gpu_job.err.%j
5 #SBATCH -N 1
6 #SBATCH -p kimg # GPU partition
7 #SBATCH --gres=gpu:1 # 1 GPU
8 #SBATCH --time=02:00:00 # 2 hours
9
10 echo "Host: $(hostname)"
11 echo "CUDA_VISIBLE_DEVICES: $CUDA_VISIBLE_DEVICES"
12
13 echo "Activating GPU TensorFlow environment"
14 source ~/myEnvs/tf_env/bin/activate
15
16 echo "Running single-GPU training script"
17 python ~/projects/train_single_gpu.py
18
19 echo "Deactivating environment"
20 deactivate
```

Single Node, Multiple GPUs (e.g., 4 GPUs)

```
1 #!/bin/bash
2 #SBATCH --job-name=multi_gpu_job
3 #SBATCH --output=multi_gpu_job.out.%j
4 #SBATCH --error=multi_gpu_job.err.%j
5 #SBATCH -N 1
6 #SBATCH -p kimg # GPU partition
7 #SBATCH --gres=gpu:4 # 4 GPUs on the same node
8 #SBATCH --time=04:00:00 # 4 hours
9
10 echo "Host: $(hostname)"
11 echo "CUDA_VISIBLE_DEVICES: $CUDA_VISIBLE_DEVICES"
12
13 echo "Activating GPU TensorFlow environment"
14 source ~/myEnvs/tf_env/bin/activate
15
16 echo "Running multi-GPU training script (e.g., data parallel)"
17 python ~/projects/train_multi_gpu.py --num-gpus 4
18
19 echo "Deactivating environment"
20 deactivate
```

Multiple Nodes, CPU-Only (e.g., MPI or Distributed CPU Job)

```
1 #!/bin/bash
2 #SBATCH --job-name=cpu_multi_node_job
3 #SBATCH --output=cpu_multi_node_job.out.%j
4 #SBATCH --error=cpu_multi_node_job.err.%j
5 #SBATCH -N 2 # 2 nodes
6 #SBATCH -p compute # CPU-only partition
7 #SBATCH --time=03:00:00 # 3 hours
8
9 echo "SLURM_NODELIST: $SLURM_NODELIST"
10 echo "SLURM_NNODES: $SLURM_NNODES"
11
12 echo "Activating CPU-only environment"
13 source ~/myEnvs/cpu_env/bin/activate
14
15 # Example: launch an MPI or distributed job
16 # (replace with your actual launcher, e.g., srun mpirun, torchrun, etc.)
17 echo "Running distributed CPU job"
18 srun python ~/projects/distributed_cpu_job.py
19
20 echo "Deactivating environment"
21 deactivate
```

Multiple Nodes, Multiple GPUs (Distributed GPU Training)

```

1  #!/bin/bash
2  #SBATCH --job-name=multi_node_multi_gpu
3  #SBATCH --output=multi_node_multi_gpu.out.%j
4  #SBATCH --error=multi_node_multi_gpu.err.%j
5  #SBATCH -N 2 # 2 nodes
6  #SBATCH -p kimq # GPU partition
7  #SBATCH --gres=gpu:4 # 4 GPUs per node (total 8 GPUs)
8  #SBATCH --time=06:00:00 # 6 hours
9
10 echo "SLURM_NODELIST: $SLURM_NODELIST"
11 echo "SLURM_NNODES: $SLURM_NNODES"
12 echo "CUDA_VISIBLE_DEVICES: $CUDA_VISIBLE_DEVICES"
13
14 echo "Activating GPU TensorFlow environment"
15 source ~/myEnvs/tf_env/bin/activate
16
17 # Example using srun to launch a distributed training job
18 # Replace with your actual distributed launcher (e.g., torchrun, horovodrun)
19 echo "Running multi-node, multi-GPU training"
20 srun python ~/projects/train_distributed_gpu.py \
21     --nodes $SLURM_NNODES \
22     --gpus-per-node 4
23
24 echo "Deactivating environment"
25 deactivate

```

3.1.1 Quick Mapping: Which Example to Use?

- **Create new CPU environment:** Listing in Section 2.7.1.
- **Clone shared GPU TF environment:** Listing in Section 2.7.2.
- **CPU-only job:** Section 3.1.
- **Single-GPU job:** Section 3.1.
- **Single node, multiple GPUs:** Section 3.1.
- **Multi-node CPU job:** Section 3.1.
- **Multi-node, multi-GPU job:** Section 3.1.

3.2 Summary

In this tutorial, you have seen:

- How to write and submit a basic SLURM job script that runs a TensorFlow program.
- How to interpret key #SBATCH directives: number of nodes, partition, GPU requests, and time limits.

- Why shared environments on HPC clusters are read-only and how to safely extend them by cloning to your home directory.
- How to create your own Python virtual environments from scratch and how to connect them to SLURM job scripts.

This workflow lets you run fast, GPU-enabled jobs using cluster-optimized software while keeping your custom Python packages isolated and safe.

Index

- cpus-per-task, 7
- gres, 7
- time, 8
- N, 7
- p, 7
- #SBATCH directives, 5
- Conda environment
 - IR, 18
- CRADLE, 5
- CUDA, 8
- cuDNN, 8
- environment
 - shared, 6
- GPU
 - partition, 7
 - request, 7
- GPU workload, 5
- GRES, 7
- HPC cluster, 1
- Micromamba, 18
- node
 - compute, 7
- pandas, 14
- partition, 7
- Python environment, 15
- sacct, 24
- scheduler
 - directives, 6
- scp, 11, 18
- SLURM, 1
 - batch script, 5
 - environment variables, 23
 - resources, 23
- ssh, 5
- TensorFlow, 5
 - shared installation, 8
- time limit, 8
- university cluster, 5
- venv, 15
- virtual environment, 5
- virtualenv-clone, 13
- VPN, 5