

# HIGH PERFORMANCE RESEARCH COMPUTING

## Introduction to Containers Advanced Content

featuring Singularity on the FASTER cluster

by Richard Lawrence

February 10, 2023



High Performance  
Research Computing

DIVISION OF RESEARCH

**Spring 2023**



# Outline

- More Container Examples
- Build your own Containers

# Learning Resources

- HPRC Wiki <https://hprc.tamu.edu/wiki/SW:Singularity>
- HPRC on Youtube <https://www.youtube.com/c/TexasAMHPRC>  
(video of this course will be posted)
- Singularity Manual <https://apptainer.org/user-docs/3.8/>
- Docker Manual <https://docs.docker.com/>
- Other container courses:
  - NBIS  
<https://nbis-reproducible-research.readthedocs.io/en/latest/singularity/>
  - Arizona <https://learning.cyverse.org/projects/Container-camp-2020/>
  - TACC <https://learn.tacc.utexas.edu/mod/page/view.php?id=95>

# More Container Examples

With exercises

# Singularity with GPU

- Containers should be built with CUDA version compatible with local GPUs (CUDA  $\geq$  11)
- Just add the `--nv` flag to your singularity command

Many repositories on Docker Hub have GPU-ready images. Search for images with “gpu” in tags

The nvidia cloud also provides GPU-ready images. See:  
[https://hprc.tamu.edu/wiki/SW:Singularity:Examples#NVIDIA\\_GPU\\_Cloud](https://hprc.tamu.edu/wiki/SW:Singularity:Examples#NVIDIA_GPU_Cloud)

# TensorFlow GPU Exercise

Image file: `tensorflow_2.4.1-gpu.sif` from  
`docker://tensorflow/tensorflow:2.4.1-gpu`

Located at `/scratch/data/Singularity/images/`

Also: `tf-gpu-test.py` in the same location.

```
$ srun --mem=512m --time=01:00:00 \
--gres=gpu:1 --partition=gpu --pty bash -i
$ singularity exec --nv tensorflow_2.4.1-gpu.sif \
python3 tf-gpu-test.py
Num GPUs Available:  1
```

*Following along live? add*  
`--reservation=training`

If error “couldn't find `tf-gpu-test.py`” try using `--bind` and/or specify full path to the file.

# Containers in HPRC Portal

Container support for **Interactive Apps** on the HPRC Portal

Currently:

- Jupyter Notebook (a popular Python IDE)

Will extend container support to other Apps. If you would like something, please ask [help@hprc.tamu.edu](mailto:help@hprc.tamu.edu)

# Jupyter Exercise

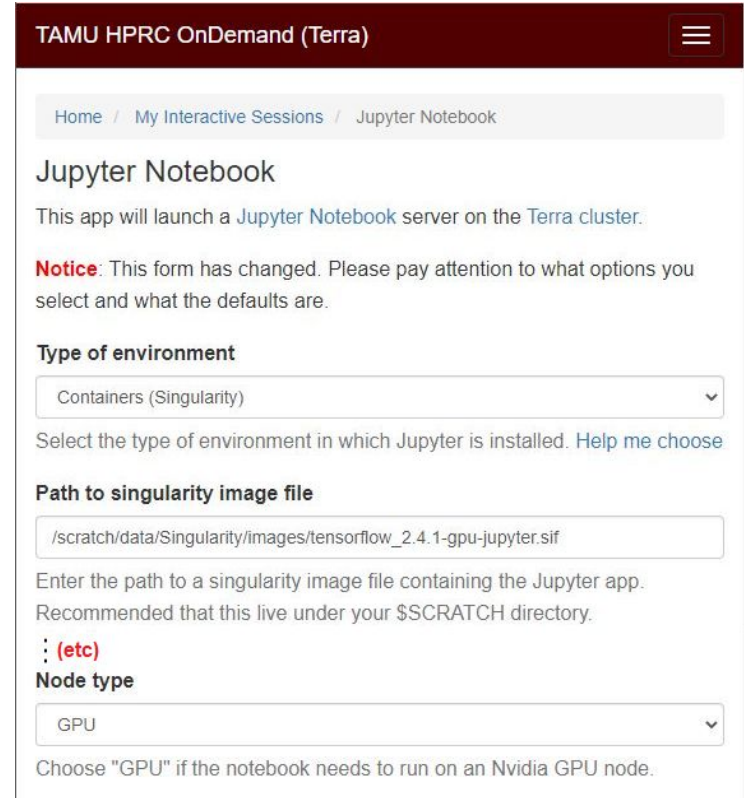
- Choose an image file on disk that contains Jupyter
- Image file: `tensorflow_2.4.1-gpu-jupyter.sif` from `docker://tensorflow/tensorflow:2.4.1-gpu-jupyter`

Located at `/scratch/data/Singularity/images/`



# Jupyter Exercise

- Navigate to HPRC Portal → FASTER OnDemand → Interactive Apps → Jupyter Notebook
- Select “Containers (Singularity)” from “Type of environment” drop-down menu
- Enter the path to the image file.



The screenshot shows the TAMU HPRC OnDemand (Terra) interface for configuring a Jupyter Notebook. The page title is "TAMU HPRC OnDemand (Terra)" and the breadcrumb trail is "Home / My Interactive Sessions / Jupyter Notebook". The main heading is "Jupyter Notebook". Below the heading, it states: "This app will launch a Jupyter Notebook server on the Terra cluster." A notice in red text says: "Notice: This form has changed. Please pay attention to what options you select and what the defaults are." There are two main configuration sections: "Type of environment" and "Path to singularity image file". The "Type of environment" section has a dropdown menu currently set to "Containers (Singularity)". Below this dropdown is the text: "Select the type of environment in which Jupyter is installed. Help me choose". The "Path to singularity image file" section has a text input field containing the path: "/scratch/data/Singularity/images/tensorflow\_2.4.1-gpu-jupyter.sif". Below the input field is the text: "Enter the path to a singularity image file containing the Jupyter app. Recommended that this live under your \$SCRATCH directory." Below the text is a red colon and the text "(etc)". The "Node type" section has a dropdown menu currently set to "GPU". Below this dropdown is the text: "Choose 'GPU' if the notebook needs to run on an Nvidia GPU node."

# Singularity with MPI

There are several ways to containerize MPI.

Check with your container source to see how it is meant to be used.

- MPI installed and run inside container.  
`singularity exec <sif> mpirun <exe>`
- MPI installed and run outside container.  
`mpirun singularity exec <sif> <exe>`
- MPI installed outside, run inside.  
`singularity exec --bind /dir <sif> /dir/mpirun <exe>`

See <https://sylabs.io/guides/3.8/user-guide/mpi.html>

# LAMMPS MPI Exercise

Image file: `lammeps-stable_29Oct2020[...].sif` from  
`docker://lammeps/lammeps:stable_29Oct2020[...]`

Located at `/scratch/data/Singularity/images/`

Also: `lammeps-test-input.txt` in the same location.

```
$ srun --ntasks=4 --nodes=1 --mem=2560m --time=01:00:00 \  
--pty bash -i
```

```
$ singularity run --bind "/scratch,$TMPDIR:/work" \  
lammeps-stable_29Oct2020_ubuntu20.04_openmpi_py3.sif mpirun -np 4 \  
/bin/lmp_mpi -in lammeps-test-input.txt
```

# Biocontainers Exercise

Image file: `bcbio-nextgen-1.2.8.sif`

From docker://quay.io/biocontainers/bcbio-nextgen

<https://biocontainers.pro/tools/bcbio-nextgen>

Located at `/scratch/data/Singularity/images/`

```
$ srun --mem=512m --time=01:00:00 --pty bash -i
```

```
$ singularity exec bcbio-nextgen-1.2.8.sif bcbio_nextgen.py --help
```

```
usage: bcbio_nextgen.py [...]
```

```
$ singularity exec bcbio-nextgen-1.2.8.sif python -c \
```

```
"import pysam; print(pysam.__version__)"
```

```
0.16.0.1
```

# Singularity with Writable Filesystem

Problem:

- You want to save your work or add things *inside* the container  
- but not the real, local filesystem
- Example: an application creates millions of files

Reminder:

- Singularity images are **immutable**.
- Singularity filesystem is **ephemeral** (does not *persist*)

Solution:

- A second file called an **overlay** that is user writable.

# Overlay Creation Tool

- The tools `dd` and `mkfs` create a virtual filesystem.
  - The version of `mkfs` matters.
  - FASTER has the `correct` `mkfs`
- Singularity can do the process automatically.

```
singularity overlay create <filename>
```

# Overlay Creation Exercise

```
$ srun --mem=512m --time=01:00:00 --pty bash -i  
$ singularity overlay create --size 256 $TMPDIR/my_overlay
```

Result: blank virtual filesystem `my_overlay` of size 256 MB  
(cp to `$SCRATCH` if you want to keep it)

# Overlay Creation Tool

- The tools `dd` and `mkfs` create a virtual filesystem.
  - The version of `mkfs` matters.
  - Some machines have the **wrong** version of `mkfs`.
- No problem! We have a *container* and a *script* for that case.
  - `docker://ubuntu:18.04`
  - `overlay-example.sh`



## Overlay Creation Exercise (Alternative)

Image file: `ubuntu-18.04.sif`

From docker://ubuntu:18.04

Script file: `overlay-example.sh`

Both located at `/scratch/data/Singularity/images`

```
$ srun --mem=512m --time=01:00:00 --pty bash -i
```

```
$ singularity exec ubuntu-18.04.sif bash overlay-example.sh
```

Result: blank virtual filesystem `my_overlay`, size 200 MB.

(cp to `$SCRATCH` if you want to keep it)

# Overlay Usage Exercise

Run your container image together with the overlay file.

```
singularity --overlay <overlay file>
```

Any files you create (outside the mounted filesystems) will go in the overlay file virtual filesystem and *persist*.

```
$ srun --mem=512m --time=01:00:00 --pty bash -i
$ singularity shell --overlay /path/to/my_overlay ubuntu-18.04.sif
> mkdir /new_dir
> touch /new_dir/new_file
> exit
$ singularity shell --overlay /path/to/my_overlay ubuntu-18.04.sif
> ls /new_dir
```

# Build your own Containers

Overview, no exercises

# Singularity Image Formats

Singularity has two image formats:

- A compressed read-only file (extension `.sif`)
- A writable directory structure called “sandbox”.

Singularity `build` command can convert from one to the other.

This is mainly useful for modifying images, such as installing new software.

Installing new software usually requires **root privileges**, which must be done on your personal machine.

# Modify Container Example

Starting with an existing image located at <source>

```
$ singularity build --sandbox <sandbox-name>  
<source>
```

This creates a directory <sandbox-name>/ that works the same as an image file, except it can be used in write mode.

```
$ sudo singularity shell --writable <sandbox-name>/
```

Now you install software in the root directories, as the root user. When you are done, you exit the sandbox and:

```
$ singularity build <final-name>.sif <sandbox-name>/
```

This creates the read-only file suitable for production use.

See [https://sylabs.io/guides/3.7/user-guide/build\\_a\\_container.html](https://sylabs.io/guides/3.7/user-guide/build_a_container.html)

# Container Recipe

Singularity `build` can build containers from a recipe.

A recipe might contain these facts:

- “Bootstrap”: what your build starts from. It could be an existing recipe or image, or it could be from scratch.
- What software to install in the container and how.
- “Runscript”: what your container does by default.

Once you have the recipe file defined:

```
$ sudo singularity build <container>  
<recipe-file>
```

See [https://sylabs.io/guides/3.8/user-guide/definition\\_files.html](https://sylabs.io/guides/3.8/user-guide/definition_files.html)

# Advanced Content Complete

# Conclusion

- Run Containers on clusters! It's easy.
- HPRC supports Singularity
- Convert Docker to Singularity!
- Charlie Cloud support also available.
- Ask for help!



# Survey

Please fill out the survey to let us know how you feel about this short course. This will help us improve.

# Questions



# Learning Resources

- HPRC Wiki <https://hprc.tamu.edu/wiki/SW:Singularity>
- HPRC on Youtube <https://www.youtube.com/c/TexasAMHPRC>  
(video of this course will be posted)
- Singularity Manual <https://apptainer.org/user-docs/3.8/>
- Docker Manual <https://docs.docker.com/>
- Other container courses:
  - NBIS <https://nbis-reproducible-research.readthedocs.io/en/latest/singularity/>
  - Arizona <https://learning.cyverse.org/projects/Container-camp-2020/>
  - TACC <https://learn.tacc.utexas.edu/mod/page/view.php?id=95>

# Thank you

Contact: [help@hprc.tamu.edu](mailto:help@hprc.tamu.edu)

