Using Containers on ACES
for Simulations, Bioinformatics and AI/ML

Richard Lawrence and Joshua Winchell

7/20/2024

ACES Workshop, Providence RI

TACC
TEXAS ADVANCED
COMPUTING CENTER

High Performance
Research Computing
DIVISION OF RESEARCH

UNIVERSITY OF
ILLINOIS
URBANA-CHAMPAIGN

Outline

- Overview of Containers on ACES
- Getting Started
- Containerized Scientific Applications
  - Genomics
  - AI/ML
  - Molecular Dynamics

Overview of Containers

What Are Containers?

- A container is a process (⚙️) that has its own **view** of local resources:
  - Filesystem
  - User IDs
  - Network

  etc.
- Example: this container (⚙️on the right) sees the **image** instead of the physical filesystem

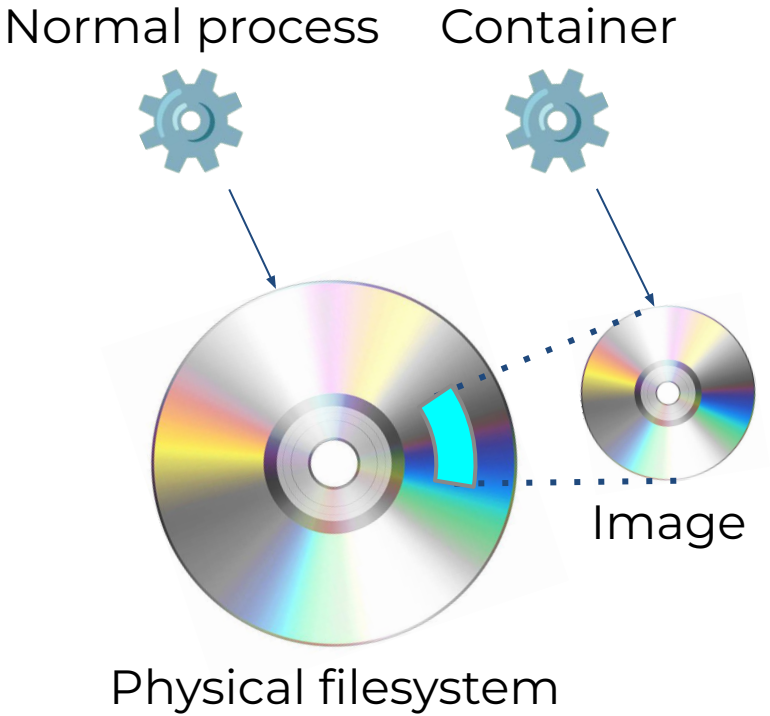

Container Runtimes on ACES

ACES offers a wide variety of Container Runtimes

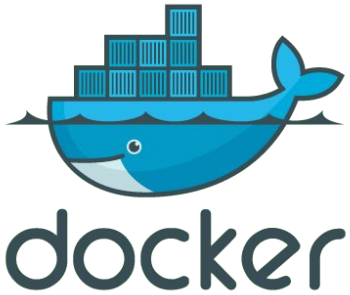

Docker

✔

*rootless

Singularity

✔

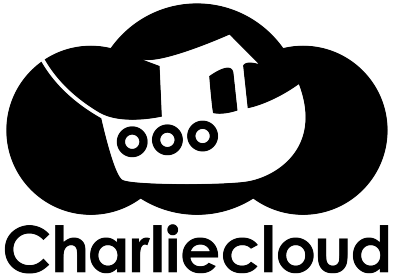

Charliecloud

✔

*rootless installations do not support all features of the runtime

Singularity on ACES

- Singularity is available on Compute nodes
  - Singularity activities are too cpu-intensive for login nodes.
- Singularity images can be large on disk. Be aware of your storage quota. (`/scratch` > `/home`)
- Some container activities may be too I/O-intense for the shared network filesystem. Be courteous to others and use a local filesystem for large image operations.

Getting Started

ACES Portal

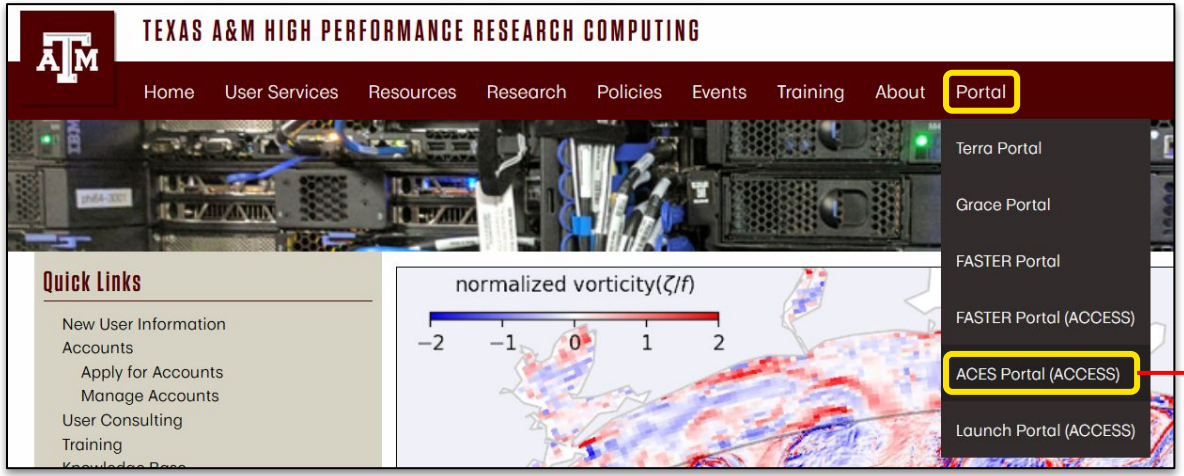

ACES Portal portal-aces.hprc.tamu.edu
is the web-based user interface for the ACES cluster

Open OnDemand (OOD) is an advanced web-based graphical interface framework for HPC users

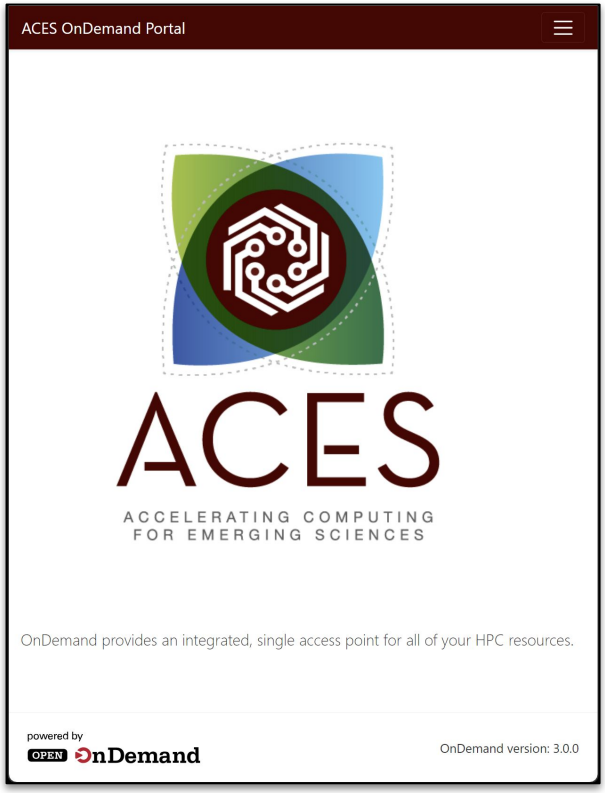

Get a Shell on ACES

Click on “Clusters” menu → _aces Shell Access

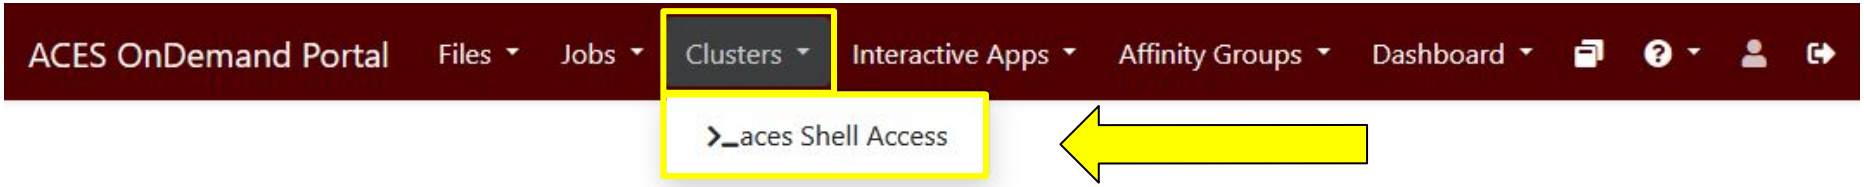

Success!

Welcome to the ACES login node.

Set Up Your Tutorial Environment

```
cd $SCRATCH
mkdir s_tutorial
cd s_tutorial
pwd

export TRAINING=/scratch/training/singularity
cd $TRAINING
pwd
```

Set Up Your Singularity Environment

- Get to a compute node from the login node
  `srun --time=120 --mem=4G --pty bash -i`
- Return to your tutorial directory (if necessary)
  `cd $SCRATCH/s_tutorial`
- Set your singularity cache directory for temporary files
  `export SINGULARITY_CACHEDIR=$TMPDIR`
- Connect to the internet for fetching images
  `module load WebProxy`

Your First Singularity Container

- Singularity can fetch an image *and* launch a shell in one line.
  `singularity shell --help`

Activity:

- Fetch an image and launch a shell from it
  ```
  singularity shell docker://almalinux:8
  cat /etc/redhat-release
  exit
  ```

Compare:

- The ACES compute nodes also have Red Hat linux installed.
  `cat /etc/redhat-release`


Congratulations!
Welcome to containers
WWW.FUNIMADA.COM

Containerized Scientific Applications

Biocontainers Registry Example

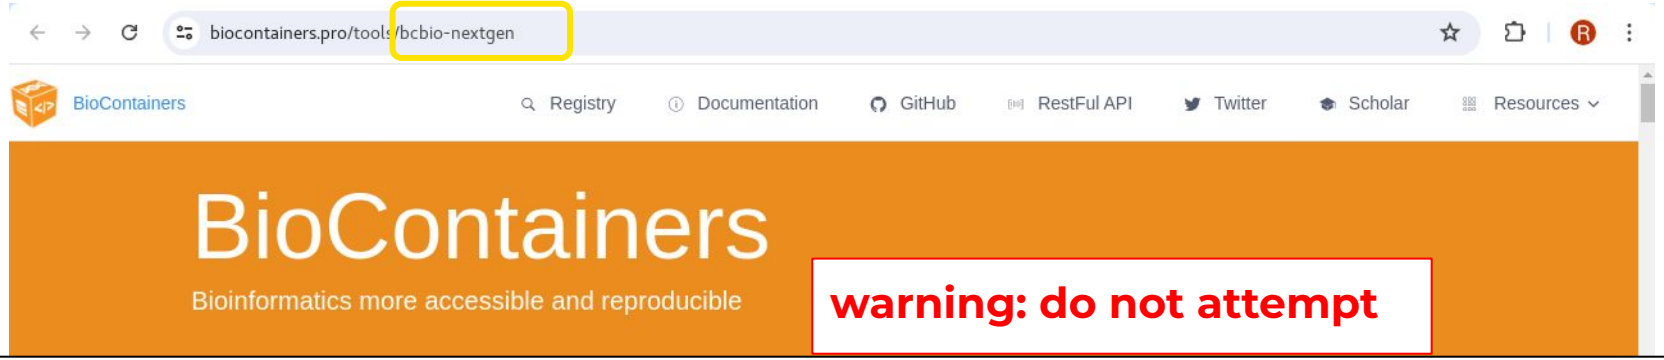


```
singularity pull docker://quay.io/biocontainers/bcbio-nextgen:1.2.9--pyh5e36f6f_3
```

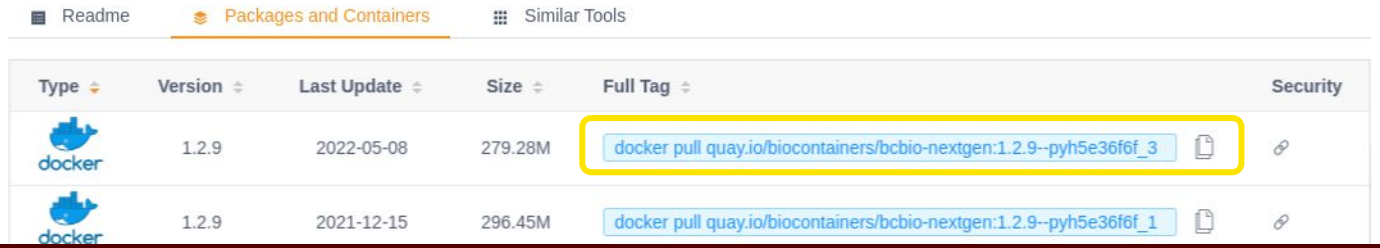

Biocontainers Exercise

- Image file: `bcbio-nextgen-1.2.8.sif`
  From `docker://quay.io/biocontainers/bcbio-nextgen`
  Located at: `/scratch/training/singularity`

Activity:

- From the login node:
  `srun --mem=4G --time=60 --pty bash -i`
- From the compute node:
  `cd /scratch/training/singularity` *or* `cd $TRAINING`
- (all on one line)

```
singularity exec bcbio-nextgen-1.2.8.sif python -c
"import pysam; print(pysam.__version__)"
```

Singularity with NVIDIA GPU

- Containers should be built with CUDA version compatible with local GPUs (CUDA ≧ 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/kb/Software/Singularity/Examples/#nvidia-gpu-cloud

NVIDIA Container Registry Example

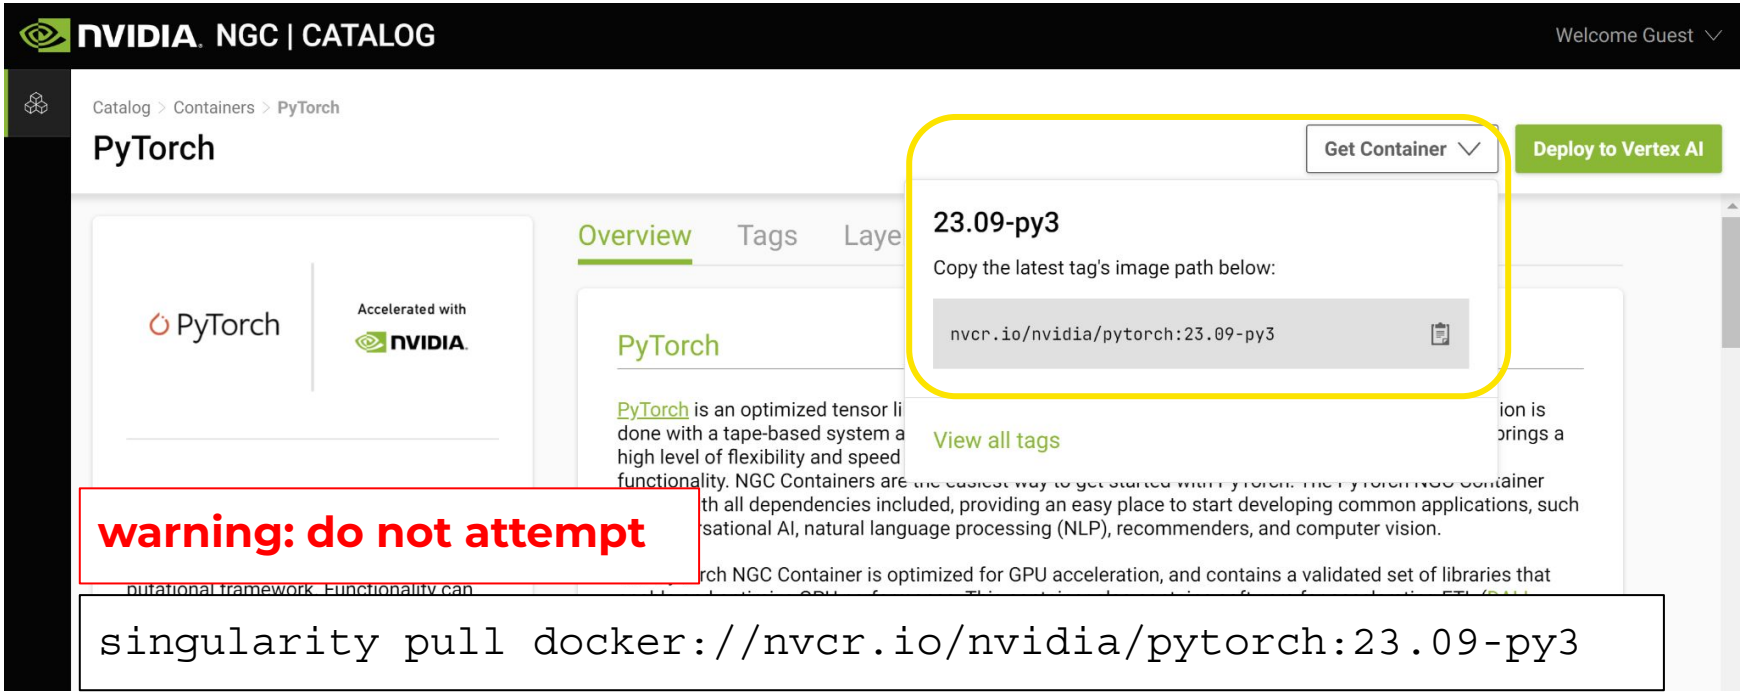

PyTorch NVIDIA GPU Exercise

- Image file: `pytorch_23.09-py3.sif`
  from: `docker://nvcr.io/nvidia/pytorch:23.09-py3`
  Located at: `/scratch/training/singularity/`

Activity:

- From the login node: (all on one line)
  `srun --mem=4G --time=60 --gres=gpu:1 --partition=gpu --pty bash -i`

Following along live? add
`--reservation=workshop_h100`

- On the compute node (if necessary):
  `cd /scratch/training/singularity` *or* `cd $TRAINING`
- From the compute node: (all on one line)

```
singularity exec --nv pytorch_23.09-py3.sif
python3 -c "import torch;
print(torch.cuda.device_count())"
```

PyTorch with Intel GPU

- Containers should be built with OneAPI version compatible with local GPUs
- No additional steps

Many repositories on Docker Hub have GPU-ready images. Search for images with “xpu” in the tag

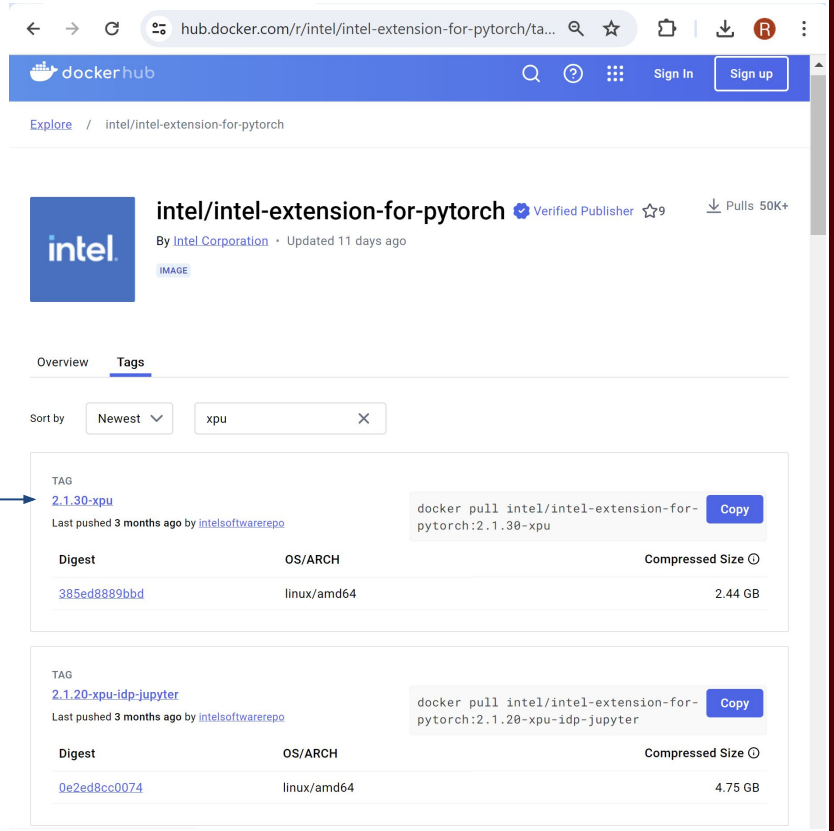

Learn more about Intel extension for PyTorch at
https://intel.github.io/intel-extension-for-pytorch/xpu/latest/tutorials/getting_started.html

PyTorch Intel GPU Exercise

- Image file: `pytorch_2.1.20-xpu-idp-jupyter.sif`
  from: `docker://intel/intel-extension-for-pytorch`
  Located at: `/scratch/training/singularity/`

Activity:

- From the login node: (all on one line)
  `srun --mem=4G --time=60 --gres=gpu:1 --partition=pvc --pty bash -i`

  Following along live? add `--reservation=workshop_pvc`
- On the compute node (if necessary):
  `cd /scratch/training/singularity` *or* `cd $TRAINING`
- From the compute node: (all on one line)

```
singularity exec pytorch_2.1.20-xpu-idp-jupyter.sif python3
-c "import torch; import intel_extension_for_pytorch as ipex;
print(torch.xpu.device_count())"
```

LAMMPS Molecular Dynamics on GPUs

- LAMMPS is a classical MD code
- https://www.lammps.org/ has a cool animated logo.
- NVIDIA provides GPU-ready container images for lammps. https://catalog.ngc.nvidia.com/orgs/hpc/containers/lammps

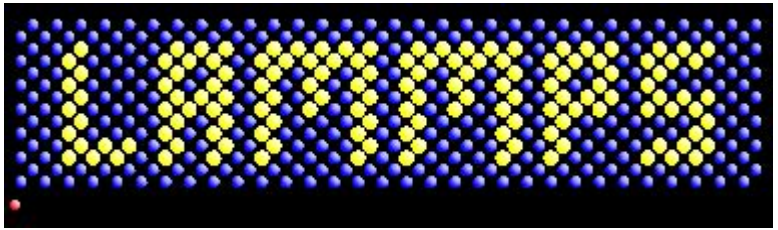

LAMMPS on H100 GPUs

- *This* specific build works with H100 GPUs

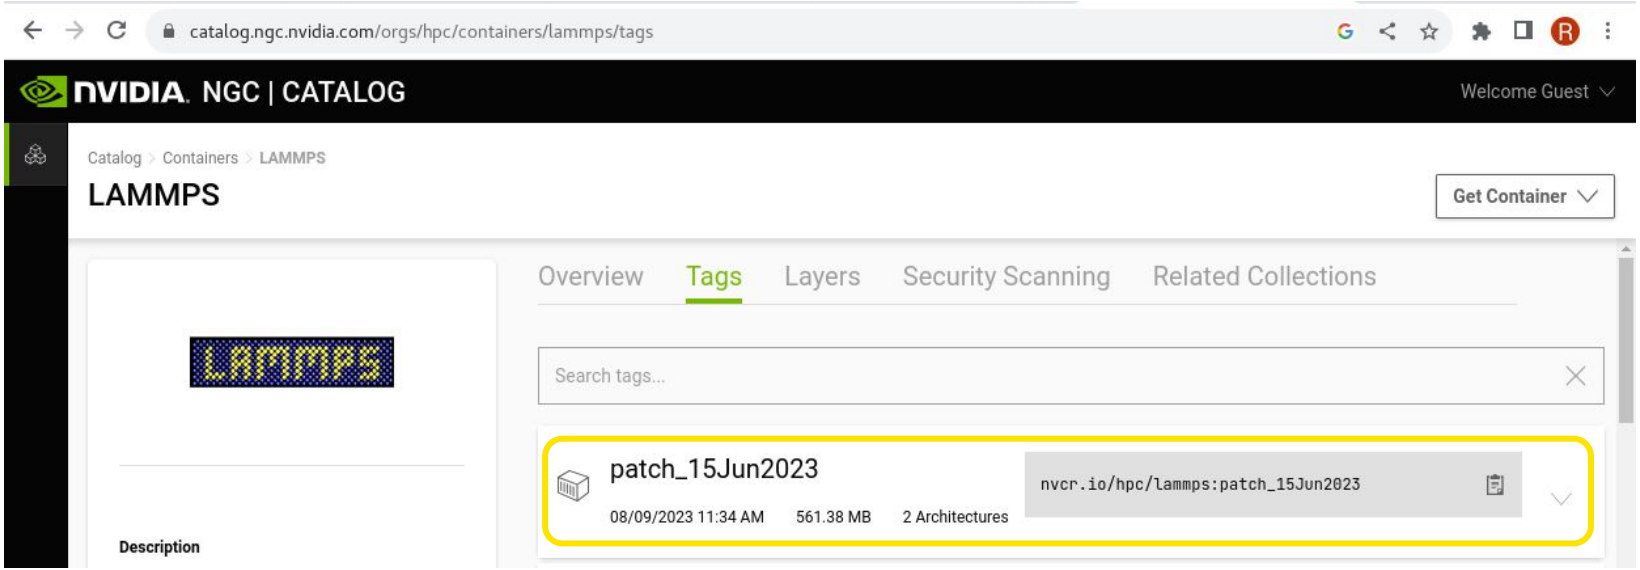

LAMMPS on NVIDIA GPUs

- Image file: `lammps-nv-patch_15Jun2023.sif`
  from: `docker://nvcr.io/hpc/lammps:patch_15Jun2023`
  Located at: `/scratch/training/singularity/`

Activity (step 1):

- From the login node: (all on one line)
  `srun --mem=4G --time=60 --gres=gpu:1 --partition=gpu --pty bash -i`
- From the compute node:
  `cd /scratch/training/singularity` *or* `cd $TRAINING`

Following along live? add `--reservation=workshop_h100`

LAMMPS on NVIDIA GPUs

- Image file: `lammps-nv-patch_15Jun2023.sif`
  from: `docker://nvcr.io/hpc/lammps:patch_15Jun2023`
  Located at: `/scratch/training/singularity/`

Activity (step 2):

- From the compute node: (all on one line)

```
singularity run --nv lammps-nv-patch_15Jun2023.sif mpirun
lmp -k on g 1 -sf kk -pk kokkos cuda/aware on neigh full
comm device binsize 2.8 -var x 4 -var y 4 -var z 4 -in
in.lj.txt -log $TMPDIR/log.lammps
```

Learn More About Containers

ACCESS Support

CCEP Community Grant Awarded

YouTube

ACES: Fundamentals of Containers

Prerequisites: Current ACCESS ID, basic Linux/Unix skills

This course introduces concepts of containers and covers common containerization tasks using the Charliecloud and Singularity container engines on the ACES cluster.

ACES: Containers for Scientific Workflows (Singularity / Apptainer)

Prerequisites: Current ACCESS ID, basic Linux/Unix skills

This course introduces the use of containers for scientific workflows using the Singularity container engine. Exercises will be performed using the ACES cluster, a composable accelerator testbed at Texas A&M University.

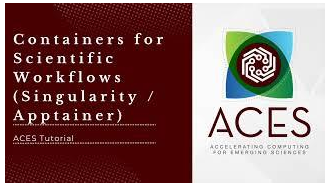

ACES: Containers for Scientific Workflows (Charliecloud)

Prerequisites: Current ACCESS ID, basic Linux/Unix skills

This course introduces the use of Containers using the Charliecloud software suite. Exercises will be performed using the ACES cluster, a composable accelerator testbed at Texas A&M University.

Texas A&M at PEARC24

Talk/Event	Date/Time	Room
Tutorial: Hands-on exercises on the Intel Data Center GPU Max 1100 (PVC-GPU) for AI/ML and Molecular	Mon, July 22, 2024 9:00 AM-12:30 PM ET	Room 553B
Seventh Workshop on Strategies for Enhancing HPC Education and Training (SEHET24)	Mon, July 22, 2024 9:00 AM-12:30 PM ET	Room 557
Workshop: Providing cutting-edge computing testbeds to the science and engineering community	Mon, July 22, 2024 1:30 PM-5:00 PM ET	Room 554A
Workshop: Engaging Secondary Students in Computing: K12 Outreach	Mon, July 22, 2024 1:30 PM-5:00 PM ET	Room 553A
Cultivating Cyberinfrastructure Careers through Student Engagement at Texas A&M University High Performance Research Computing	Tue, July 23, 2024 11:00 AM-11:25 AM ET	Junior Ballroom
Insight Gained from Migrating a Machine Learning Model to Intelligence Processing Units	Tue, July 23, 2024 11:00 AM-11:25 AM ET	Room 551 A&B
BOF 4: What's in it for me? How can we truly democratize the research computing and data community?	Tue, July 23, 2024 1:30 PM-2:30 PM ET	Room 551 A&B

Texas A&M at PEARC24

Talk/Event	Date/Time	Room
BRICCs: Building Pathways to Research Cyberinfrastructure at Under Resourced Institutions	Tue, July 23, 2024 3:25 PM-3:50 PM ET	Junior Ballroom
Memory Bandwidth Performance across Accelerators	Tue, July 23, 2024 3:25 PM-3:50 PM ET	Ballroom B
Container Adoption in Campus High Performance Computing	Wed, July 24, 2024 11:00 AM-11:25 AM ET	Ballroom B
Engaging Secondary Students in Computing and Cybersecurity	Wed, July 24, 2024 3:15 PM-3:30 PM ET	Room 557
Exploring the Viability of Composable Architectures to Overcome Memory Limitations in High Performance Computing Workflows	Wed, July 24, 2024 3:45 PM-4:00 PM ET	Room 553 A&B
Performance of Molecular Dynamics Acceleration Strategies on Composable Cyberinfrastructure	Wed, July 24, 2024 4:15 PM-4:30 PM ET	Room 551 A&B
BOF 17: Fantastic ACCESS Cyberinfrastructure Resources and Where to Find Them	Wed, July 24, 2024 4:45 PM-5:45 PM ET	Room 553 A&B
BOF 18: Recipes to build successful cross-institutional collaborative computing	Wed, July 24, 2024 4:45 PM-5:45 PM ET	Junior Ballroom

High Performance
Research Computing
DIVISION OF RESEARCH

Thank you

- Please visit our talks and BoFs at PEARC23
- Please join the ACES affinity group support.access-ci.org/affinity-groups/aces
- Questions?help@hprc.tamu.edu
- We gratefully acknowledge support from National Science Foundation awards #2112356 (ACES), #2019129 (FASTER) and #19257614(SWEETER)