

Fundamentals of Containers

Singularity and **Charliecloud**

on **ACES**

Richard Lawrence
10/21/2025



High Performance
Research Computing
DIVISION OF RESEARCH



developed for



Outline

- Overview of Containers
- Overview of Charliecloud
- Overview of Singularity
- Getting Started
- Working with Images
- Container Recipes


Learning Resources

- Slides on the course web page
https://hprc.tamu.edu/training/aces_containers.html
(highly recommended for working along)
- HPRC's Knowledge Base
<https://hprc.tamu.edu/kb/Software/Charliecloud/>
<https://hprc.tamu.edu/kb/Software/Singularity/>
- HPRC on YouTube
<https://www.youtube.com/@TexasAMHPRC/playlists>
- ACCESS Links
<https://support.access-ci.org/ci-links>

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: the container (⚙️ on the right) sees the **image** instead of the physical filesystem




Why Use Containers?

- **Shareability:**
 - Share your container image file by uploading to a public repository
 - Use images shared by others
- **Portability:**
 - Use images on any computer with the same architecture (x86-64)
- **Reproducibility:**
 - Container users are largely unaffected by changes to the cluster environments


What Goes In Container Images?

- Unlike in VMs, the OS Kernel is not duplicated
- Container images are smaller than VM images


Local Build, or “Bare metal”



Virtual Machine




Container




Popular Container Runtimes

Instant deployment to users on different devices!



LXC
2008



Docker
2013




Singularity
2015



SHIFTER

Shifter 2016



Charliecloud

Charliecloud
2017



Podman
2018

Overview of Charliecloud

Charliecloud

- A lightweight, fully-unprivileged container solution



Presented by



Charliecloud Features

- Charliecloud is a container runtime and an image builder.
- Charliecloud can read and convert Docker images.
- Filesystem inside container is isolated.
- User inside container is isolated.
- Works with high-performance cluster technologies

Read more in the Charliecloud manual on github
<https://hpc.github.io/charliecloud/>

Charliecloud on ACES

- Charliecloud is available from our module system.
 - `module load charliecloud`
- Charliecloud images can be large on disk. Be aware of your storage quota. (`/scratch` > `/home`)
- Some container activities may be too cpu-intensive for the *shared* login node. Be courteous to others and use a compute node for large image operations.
- Some container activities may be too I/O-intensive for the *shared* network filesystem. Be courteous to others and use a local filesystem for large image operations.

Overview of Singularity

Singularity

- An easy-to-use, high-performance container solution



Presented by



**Where Next-level
Container Performance
and Security Converge**

<https://sylabs.io/solutions/>

Singularity is Apptainer



Singularity Features

- Singularity is a container runtime and an image builder.
- Singularity can read and convert Docker images .
- Filesystem inside container is isolated .
- User inside container is the same as the user outside.
- Works with high-performance cluster technologies

Read more in the Apptainer manual


<https://apptainer.org/user-docs/3.8/>

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-intensive 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.



Accessing ACES via the Portal (ACCESS)



The screenshot shows the ACCESS portal interface. At the top is a navigation bar with links: ALLOCATIONS, SUPPORT, OPERATIONS, METRICS, and a Login link. Below the navigation bar is the ACCESS logo. A section titled "Consent to Attribute Release" is expanded, showing a message from TAMU ACES ACCESS OIDC requesting access to user information. Below this is a "Select an Identity Provider" dropdown menu, which is highlighted with a red rectangle. The dropdown shows "ACCESS CI (XSEDE)" as the selected option. Below the dropdown is a checkbox labeled "Remember this selection". A yellow "LOG ON" button is positioned below the dropdown menu, with a red line pointing from the dropdown to the text below.

ALLOCATIONS SUPPORT OPERATIONS METRICS Login

ACCESS

Consent to Attribute Release

TAMU ACES ACCESS OIDC requests access to the following information. If you do not approve this request, do not proceed.

- Your CILogon user identifier
- Your name
- Your email address
- Your username and affiliation from your identity provider


Select an Identity Provider

ACCESS CI (XSEDE)

☐ Remember this selection

LOG ON

Select the Identity Provider appropriate for your account.



The screenshot shows the ACCESS portal login page. At the top is the ACCESS logo. Below the logo is a message: "If you had an XSEDE account, please enter your XSEDE username and password for ACCESS login." Below this are two input fields: "ACCESS ID" and "ACCESS Password". Below the password field is a yellow "LOGIN" button. At the bottom of the page are three links: "Register for an ACCESS ID", "Forgot your password?", and "Need Help?".

ACCESS

If you had an XSEDE account, please enter your XSEDE username and password for ACCESS login.

ACCESS ID

ACCESS Password

LOGIN

[Register for an ACCESS ID](#)


[Forgot your password?](#)

[Need Help?](#)

Log-in using your ACCESS or institutional credentials.

Get a Shell on ACES

Click on “Clusters” menu → _aces Shell Access



Success!

Welcome to the
ACES login node.

```
portal-aces.hprc.tamu.edu/pun/sys/shell/ssh/login.aces
Host: login.aces Themes: Default

=====
| Texas A&M University High Performance Research Computing |
| Website: https://hprc.tamu.edu |
| Consulting: help@hprc.tamu.edu (preferred) or (979) 845-0219 |
| ACES Documentation: https://hprc.tamu.edu/kb/User-Guides/ACES |
| FASTER Documentation: https://hprc.tamu.edu/kb/User-Guides/FASTER |
| Grace Documentation: https://hprc.tamu.edu/kb/User-Guides/Grace |
| Terra Documentation: https://hprc.tamu.edu/kb/User-Guides/Terra |
| YouTube Channel: https://www.youtube.com/texasamhprc |
=====

*****
*      === IMPORTANT POLICY INFORMATION ===      *
* - Unauthorized use of HPRC resources is prohibited and subject to *
*   criminal prosecution. *
* - Use of HPRC resources in violation of United States export control *
*   laws and regulations is prohibited. Current HPRC staff members are *
*   US citizens and legal residents. *
* - Sharing HPRC account and password information is in violation of *
*   Texas State Law. Any shared accounts will be DISABLED. *
* - Authorized users must also adhere to ALL policies at: *
*   https://hprc.tamu.edu/policies/ *
*****

!! WARNING: THERE ARE ONLY NIGHTLY BACKUPS OF USER HOME DIRECTORIES. !!

Please restrict usage to 8 CORES across ALL login nodes.
Users found in violation of this policy will be SUSPENDED.

To see these messages again, run the motd command.
[u.rl117197@aces-login2 ~]$
```

Set Up Your Charliecloud Environment

On the login node:

```
module load charliecloud  
module list
```

Note. *light green highlight* means please perform this exercise in your ACES shell terminal.

Your First Charliecloud Image

The Charliecloud image tool helps you build and organize your images.

```
ch-image --help
```

Let's fetch a small, basic linux distro: Almalinux.

```
ch-image pull almalinux:8  
ch-image list
```

The image is in your personal temporary local image repository.

```
echo $CH_IMAGE_STORAGE  
ls $CH_IMAGE_STORAGE/img/
```

Your First Charliecloud Container

The ACES login node has Red Hat Enterprise linux installed.

```
cat /etc/redhat-release
```

The charliecloud-run tool launches containers out of existing images.

```
ch-run --help
```

Launch a bash shell, investigate linux, and stop the container.

```
ch-run almalinux:8 bash  
cat /etc/redhat-release  
exit
```

Set Up Your Singularity Environment

Get to a compute node (from the login node).

```
srun --time=90 --mem=4G --pty bash -i
```

attending live? add:
--reservation=training

The compute node should be similar to the login node.

```
cat /etc/redhat-release
```

Set your singularity cache directory.

```
export SINGULARITY_CACHEDIR=$TMPDIR
```

Connect to the internet.

```
module load WebProxy
```

Your First Singularity Container

Singularity can fetch an image *and* launch a shell in one line.

```
singularity shell --help
```

Fetch that same image (again) and launch a shell from it (again).

```
singularity shell docker://almalinux:8  
cat /etc/redhat-release  
exit
```

Don't forget to return to the login node.

```
exit
```



Congratulations!

Welcome to containers

WWW.FUNIMADA.COM

Working with Images

Charliecloud Image Formats

- Charliecloud container images come in two main formats:
 1. Directory
 2. Single file. HPRC supports the squashfs filesystem format for single file images. (more about that on a later slide)
- The `ch-convert` tool copies images into different formats.

```
ch-convert --help
```

Directory Image Format

- The image name should end in /.
- Directory images are writable.
- Directory read/write operation are slow, so put directory images on the high-speed `/tmp` filesystem.
- Images in `$CH_IMAGE_STORAGE` are also directory images, but you refer to them by name without the trailing slash.

Convert to Directory Exercise

Create a space on the login node for yourself.

```
mkdir /tmp/$USER
```

Convert our image in the cache to a directory image.
(note the order of the arguments)

```
ch-convert almalinux:8 /tmp/$USER/almalinux/
```

What did we make?

```
ls /tmp/$USER/almalinux/
```

Editing Images Exercise

Directory images can be modified by adding the `--write` flag to `ch-run`. Any changes you make will be saved.

```
ch-run --write /tmp/$USER/almalinux/ bash  
mkdir /my_dir  
exit
```

Are the changes still there?

```
ch-run /tmp/$USER/almalinux/ bash  
ls /  
exit
```

Squashfs Image Format

- Squashfs is an open-source file format for filesystem images.
- The whole filesystem becomes one single file.
- The image name should end in `.sqfs`
- Squashfs images are read-only.
- Squashfs read operations are fast, so put squashfs images on the network filesystem `/scratch`.

Set Up Your Environment

Create a workspace in your **scratch** directory.

```
cd $SCRATCH  
mkdir c_tutorial  
cd c_tutorial  
pwd
```

Convert to Squashfs Exercise

Make sure you are still in your `c_tutorial` directory in `$SCRATCH`.

```
pwd
```

Then convert (*note the order of the arguments*).

```
ch-convert /tmp/$USER/almalinux/ almalinux.sqfs
```

Are your changes still there?

```
ch-run almalinux.sqfs /bin/bash
```

```
ls
```

```
exit
```

Singularity Image Formats

- Singularity container images come in two main formats:
 1. Directory
 2. Single file. Singularity uses the SIF format for single file images. This is the default.
- The `singularity build` tool can convert images in both formats.
`singularity build --help`
- The `--sandbox` option is used to create directory-format images.

Set Up Your Singularity Environment

Get to a compute node (from the login node).

```
srun --time=90 --mem=4G --pty bash -i
```

attending live? add:
--reservation=training

Set your singularity cache directory.

```
export SINGULARITY_CACHEDIR=$TMPDIR
```

Connect to the internet.

```
module load WebProxy
```

Return to your scratch area.

```
cd $SCRATCH/c_tutorial
```

Singularity Image Exercise

Singularity pull can fetch an image and write to either file format.
(note the order of the arguments)

```
singularity pull almalinux.sif docker://almalinux:8
```

Singularity can convert an image to the directory file format.
Use the --sandbox argument to specify the directory type.
(note the order of the arguments)

```
singularity build --sandbox $TMPDIR/almalinux almalinux.sif
```

Singularity Write Exercise

Directory images are writable. Simply add the `--writable` flag to your container command.

```
singularity shell --writable $TMPDIR/almalinux  
mkdir /my_dir  
exit
```

Are the changes still there?

```
singularity shell $TMPDIR/almalinux  
ls /  
exit
```

Are Directory Images All Compatible?

Let's try an experiment. Still on your compute node:

```
module load charliecloud  
ch-run $TMPDIR/almalinux bash  
ls /  
exit
```

Directory images are universal.

They can come from a variety of sources, for example:

<https://github.com/alpinelinux/docker-alpine/>

This github repo contains a Linux package in a directory format that is designed for containers.

Container from Tarball Exercise (1/3)

Navigate to your tmp space.

```
cd $TMPDIR
```

Fetch the distro's tarball:

```
wget -O alpine.tar.xz  
'https://github.com/alpinelinux/docker-alpine/blob/v3.18/x86_64/alpine-minirootfs-3.18.9-x86_64.tar.gz?raw=true'
```

or

```
cp /scratch/training/containers/alpine.tar.xz .
```

Container from Tarball Exercise (2/3)

Oh no! the tarball is a tar *bomb* (too many files).

```
tar tf alpine.tar.xz | head -10
```

Unpack it in a new subdirectory.

```
mkdir alpine
cd alpine
tar xf ../alpine.tar.xz
ls
cd ..
```

Container from Tarball Exercise (3/3)

Now we can run a shell in the container!

```
ch-run ./alpine -- /bin/sh
cat /etc/alpine-release
exit
```

```
singularity shell ./alpine
cat /etc/alpine-release
exit
```

(alpine doesn't have bash.)

(singularity knew what to do; charliecloud needed a hint.)

Container Recipes

Container Recipes

- Modifying containers by hand is bad in practice. The information about what steps were taken is lost.
- It is better to write down those steps in a recipe file.
- Docker uses a recipe file named Dockerfile.
- Chariecloud supports Dockerfiles.
- Singularity uses a different recipe file called a Definition file.

Target Recipe

1. Start with the almalinux:8 image.
2. Install Python 3.

```
yum -y install python39 (do not attempt)
```

3. Add a “hello.py” python script (and make it executable).

```
#!/usr/bin/python3  
print("Hello World!")
```

```
chmod 755 hello.py
```

And we shall name the image “hello”.

Elements of a Dockerfile

```
FROM almalinux:8
RUN yum -y install python39
COPY ./hello.py /
RUN chmod 755 /hello.py
```

1. FROM: We are extending the almalinux:8 base image.
2. RUN: Install the python39 RPM package.
3. COPY: Copy the file hello.py from outside the image into the root directory of the image.
4. RUN: Make that file executable.

Charliecloud Recipe

Take a moment to set up this workspace in \$SCRATCH:

```
c_tutorial/  
  hello.src.docker/  
    hello.py
```

```
#!/usr/bin/python3  
print("Hello World!")
```

Dockerfile

```
FROM almalinux:8  
RUN yum -y install python39  
COPY ./hello.py /  
RUN chmod 755 /hello.py
```

Build from Dockerfile

Make sure you have charliecloud loaded; then build the image:

```
cd $SCRATCH/c_tutorial #if necessary  
cd hello.src.docker  
ch-image build -t hello -f Dockerfile .  
ch-image list
```

ch-image build arguments:

- (-t hello) Build an image named (a.k.a. tagged) “hello”.
- (-f Dockerfile) Use the Recipe named “Dockerfile”.
- (.) Use the current directory as the context directory.

Testing the Recipe Image

Convert from image cache to flat image.

```
ch-convert hello hello.sqfs
```

Run the container and launch hello from python.

```
ch-run hello.sqfs -- /hello.py
```

Run the container and check for python.

```
ch-run hello.sqfs -- python3 --version
```

Does python3 exist on the local node?

```
python3 --version
```

Elements of a Definition File

```
Bootstrap: docker
From: almalinux:8
%files
    hello.py
%post
    yum -y install python39
    chmod 755 /hello.py
```

1. BOOTSTRAP: Base image comes from Docker Hub.
2. FROM: Base image is almalinux:8
3. %files:
Copy the file hello.py from outside the image into the root directory of the image.
4. %post:
 - a. Install the python39 RPM
 - b. Make /hello.py executable.

Singularity Recipe

Take a moment to set up this workspace in \$SCRATCH:

```
c_tutorial/  
  hello.src.def/  
    hello.py
```

Definition

```
#!/usr/bin/python3  
print("Hello World!")
```

```
Bootstrap: docker  
From: almalinux:8  
%files  
  hello.py  
%post  
  yum -y install python39  
  chmod 755 /hello.py
```

Build from Definition file

Make sure you are on the compute node, then build the image:

```
cd $SCRATCH/c_tutorial #if necessary  
cd hello.src.def  
singularity build --fakeroot hello.sif Definition
```

singularity build arguments:

- (--fakeroot) Needed for yum install permission.
- (Definition) Use the Recipe named "Definition".
- (hello.sif) Save the result in an SIF image.

Testing the Recipe Image

Singularity exec is used for running non-shell containers.

```
singularity exec --help
```

Run the container and launch hello from python.

```
singularity exec hello.sif /hello.py
```

Run the container and check for python.

```
singularity exec hello.sif python3 --version
```

Does python3 exist on the local node?

```
python3 --version
```

Acknowledgements

This work was supported by

- the National Science Foundation (NSF), award numbers:
 - 2112356 - ACES - Accelerating Computing for Emerging Sciences
 - 1925764 - SWEETER - SouthWest Expertise in Expanding, Training, Education and Research
 - 2019129 - FASTER - Fostering Accelerated Scientific Transformations, Education, and Research
- Staff and students at Texas A&M High-Performance Research Computing.
- ACCESS CCEP pilot program, Tier-II



**HIGH PERFORMANCE
RESEARCH COMPUTING**
TEXAS A&M UNIVERSITY

<https://hprc.tamu.edu>

HPRC Helpdesk:

help@hprc.tamu.edu

Phone: 979-845-0219

Take our short course survey!



HPRC Survey

https://u.tamu.edu/hprc_shortcourse_survey

Help us help you. Please include details in your request for support, such as, Cluster (ACES, FASTER, Grace, Launch), NetID (UserID), Job information (JobID(s), Location of your jobfile, input/output files, Application, Module(s) loaded, Error messages, etc), and Steps you have taken, so we can reproduce the problem.



High Performance Research Computing | NSF award #2112356