

# ECP Tutorial: Getting Started with Containers on HPC

Shane Canon<sup>1</sup>, Sameer Shende<sup>2</sup>, Andrew J. Younge<sup>3</sup>

<sup>1</sup>Lawrence Berkeley National Lab  
scanon@lbl.gov

<sup>2</sup>University of Oregon  
sameer@cs.uoregon.edu

<sup>3</sup>Sandia National Labs  
ajyoung@sandia.gov



NERSC



UNIVERSITY OF  
OREGON



Sandia  
National  
Laboratories

# Outline


- 14:00 – 14:30 Introduction to Containers in HPC (Younge)
- 14:30 – 15:00 How to build your first Docker container (Canon)
- 15:00 – 15:30 How to deploy a container on a supercomputer (Canon)
- 15:30 – 15:45 Best Practices (Canon)
- 16:00 – 16:30 -- Break --
- 16:30 – 17:00 Running an HPC app on the E4S container (Shende)
- 17:00 - 17:30 How to build a Singularity container image (Younge)
- 17:30 - 18:00 Success Stories & Summary (Canon)

# Introduction to Containers in HPC




# What are Containers

- A lightweight collection of executable software that encapsulates everything needed to run a single specific task
  - Minus the OS kernel
  - Based on Linux only
- Processes and all user-level software is isolated
- Creates a portable\* software ecosystem
- Think chroot on steroids
- Docker most common tool today
  - Available on all major platforms
  - Widely used in industry
  - Integrated container registry via Dockerhub



## 5 Hypervisors and Containers

- Type 1 hypervisors insert layer below host OS
- Type 2 hypervisors work as or within the host OS
- Containers do not abstract hardware, instead provide “enhanced chroot” to create isolated environment
- Location of abstraction can have impact on performance
- All enable custom software stacks on existing hardware



## 6 Background

- Abstracting hardware and software resources has had profound impact on computing
- Virtual Machines to Cloud computing in the past decade
  - Early implementations limited by performance
  - HPC on clouds: FutureGrid, Magellan, Chameleon Cloud, Hobbes, etc
  - Some initial successes, but not always straightforward
- OS-level virtualization a bit different
  - User level code packaged in container, can then be transported
  - Single OS kernel shared across containers and provides isolation
  - Cgroups traditionally multiplexes hardware resources
  - Performance is good, but OS flexibility is limited

# Containers in Cloud Industry

- Containers are used to create large-scale loosely coupled services
- Each container runs just 1 user process – “micro-services”
  - 3 httpd containers, 2 DBs, 1 logger, etc
- Scaling achieved through load balancers and service provisioning
- Jam many containers on hosts for increased system utilization
- Helps with dev-ops issues
  - Same software environment for developing and deploying
  - Only images changes are pushed to production, not whole new image (CoW).
  - Develop on laptop, push to production servers
  - Interact with github similar to developer code bases
  - Upload images to “hub” or “repository” whereby they can just be pulled and provisioned

# Containers

- Containers are gaining popularity for software management of distributed systems
- Enable way for developers to specify software ecosystem
- US DOE High Performance Computing (HPC) resources need to support emerging software stacks
  - Applicable to DevOps problems seen with large HPC codes today
  - Support new frameworks & cloud platform services
- But HPC systems are very dissimilar from cloud infrastructure
  - MPI-based bulk synchronous parallel workloads are common
  - Scale-out to thousands of nodes
  - Performance is paramount

# Container features in HPC

- **BYOE - Bring-Your-Own-Environment**
  - Developers define the operating environment and system libraries in which their application runs.
- **Composability**
  - Developers explicitly define how their software environment is composed of modular components as container images,
  - Enable reproducible environments that can potentially span different architectures.
- **Portability**
  - Containers can be rebuilt, layered, or shared across multiple different computing systems
  - Potentially from laptops to clouds to advanced supercomputing resources.
- **Version Control Integration**
  - Containers integrate with revision control systems like Git
  - Include not only build manifests but also with complete container images using container registries like Docker Hub.

# Container features not wanted in HPC

- **Overhead**

- HPC applications cannot incur significant overhead from containers

- **Micro-Services**

- Micro-services container methodology does not apply to HPC workloads
  - 1 application per node with multiple processes or threads per container

- **On-node Partitioning**

- On-node partitioning with cgroups is not necessary (yet?)

- **Root Operation**

- Containers allow root-level access control to users
  - In supercomputers this is unnecessary and a significant security risk for facilities

- **Commodity Networking**

- Containers and their network control mechanisms are built around commodity networking (TCP/IP)
  - Supercomputers utilize custom interconnects w/ OS kernel bypass operations

# HPC Containers

- Docker not good fit for running HPC workloads
  - Security issues
    - Can't allow root on shared resources
  - Lack of HPC architecture support
    - No batch integration
    - Assumes local resources
    - Assumes commodity TCP/IP
- Many different container options in HPC



Shifter




SHIFTER

Singularity



Charliecloud



Charliecloud

...

# Developing Container Vision

- Support software dev and testing on laptops
  - Working builds that then can run on supercomputers
  - Dev time on supercomputers is expensive
  - May also leverage VM/binary translation
- Let developers specify how to build the env AND app
  - Import and run container on target platform
  - Many containers, but can have different code “branches”
  - Not bound to vendor and sysadmin software
- Focus on Interoperability
- Provide containerized services coupled with simulations
  - Developing mechanisms to support services
- Performance matters
  - Want to manage permutations of architectures and compilers
  - Ensure container implementations on HPC are performant
  - Keep features to support future complete workflows

# This tutorial will show you:

- How to build your first Docker container.
- How to run a Docker container on a supercomputer with Shifter.
- How to build your first Singularity container.
- How to use the Extreme-scale Scientific Software Stack (E4S) container image.
  - And a bit about Spack
- And maybe some best practices and lessons learned.

## Tutorial Link

[tinyurl.com/ecp20tut](https://tinyurl.com/ecp20tut)

<https://supercontainers.github.io/ecp2020-tutorial/>

# Tutorial Training Accounts

1. EC2 instance login
2. Cori training account



# Questions?

Next: learn how to work with your first container!