WILLIAM & MARY

CHARTERED 1693

Update on Containers

Guidelines, Singularity, Batch Processing

David Blyth, Wouter Deconinck, Markus Diefenthaler,
Andrea Dotti, David Lawrence, Alexander Kisselev

Overview

- Container guidelines
- Containerization status updates
  - ANL
  - BNL
  - JLab
- Containers for batch processing

Guidelines

- Targeted at persons with limited time to invest in building the software.
- Help ensure that software developed at different sites (ANL, BNL, JLab) is packaged in a similar way simplifying cross comparisons.
- Evolve to reflect changes in the software or in the requirements of the community.

Guidelines

- https://gitlab.com/ESC/containers/blob/master/Guidelines.md

Guidelines.md 7.48 KB

Edit Web IDE

ESC Container Guidelines

Edited by D. Blyth, W. Deconinck, M. Diefenthaler, A. Dotti, A. Kiselev, D. Lawrence

> The current version of this document will focus on an interactive VM style container. A future version will include guidelines for an appliance style container appropriate for use on GRID sites, commercial cloud systems, and HPC resources. This version will include guidelines for deriving Singularity and Shifter images from the Docker images.

Introduction

This document provides guidelines for producing container images that provide EIC software. The primary goal of the images is to provide an easy means for scientists to start running EIC software. Specifically targeted are persons with limited time to invest in building the software and all of its prerequisites. The guidelines will help ensure that software developed at different sites (ANL, BNL, JLab) is packaged in a similar way simplifying cross comparisons and will evolve over time to reflect

Goals

- Make it easier for users by minimizing the "installation overhead" for new users.
- Allow users to run the software interactively on any computer through the use of containers.
- Provide consistency between software generated at different facilities.

Technology

- Docker generates initial image; all the content should go into Dockerfile and associated input files if needed.
- Consideration should be given to repeatability of builds.
- Consideration should be given to using the image with Singularity or Shifter by simple import from Docker.

Deployment

Images will be deployed using the electronioncollider swarm on Docker Hub. In addition, any Dockerfiles should be committed to the ESC/containers repository on GitLab. Any required supporting files should be included or clear instructions given on where to obtain them.

Deployment (2)

Singularity images can be deployed using the electronioncollider GitHub organization on Singularity Hub.

wdconinc Create Singularity.1.0.0		Latest commit a387f63 on Apr 12
README.md	Create README.md	a month ago
Singularity.1.0.0	Create Singularity.1.0.0	a month ago
Singularity.1.0.1	Create Singularity.1.0.1	a month ago
Singularity.1.0.2	Create Singularity.1.0.2	a month ago
Singularity.1.0.3	Rename Singularity to Singularity.1.0.3	a month ago

Overview

- ✔ Container guidelines
- Containerization status updates
  - ANL
  - BNL
  - JLab
- Containers for batch processing

ANL EIC Containers

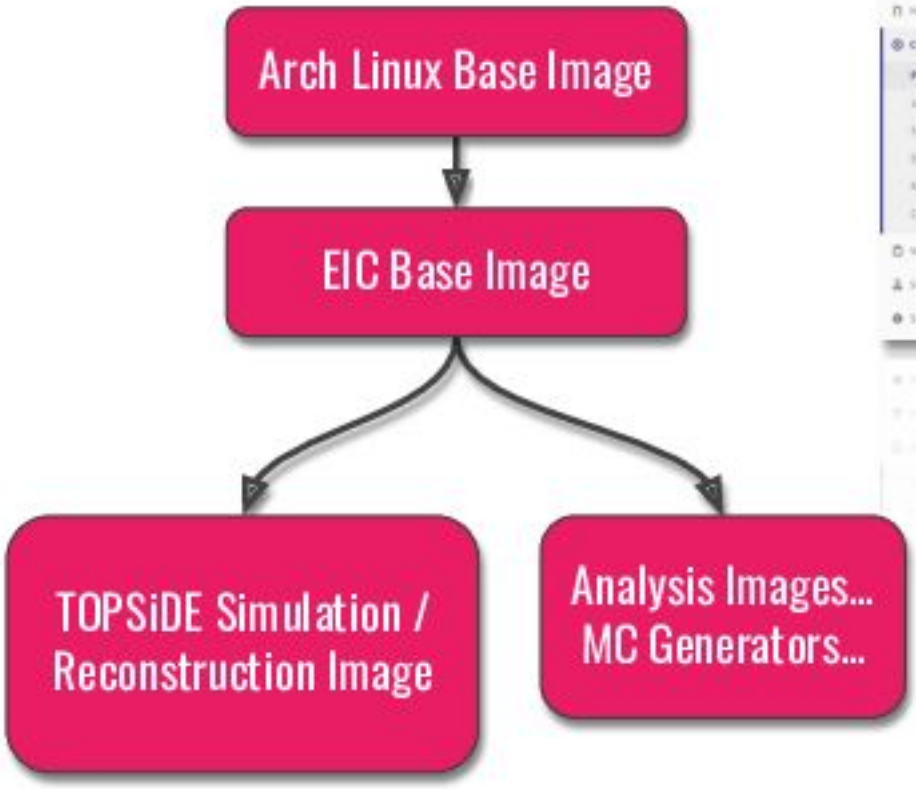

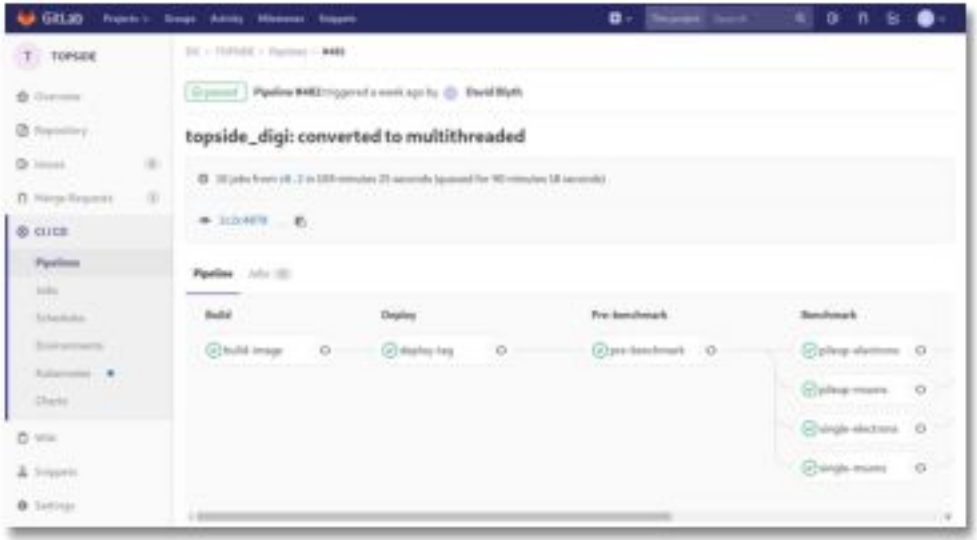

- Clean, normative layout (FHS)
- Addition of `/image` directory to guarantee reproducibility (ESC container guidelines)
- Validated and deployed with GitLab CI
  - Automated TOPSiDE benchmarks pass inspection before deployment
- EIC base image factors out I/O software and common tools for use by analyses and MC generators

Containers at BNL

EicRoot bundle (Docker)

- Grown on top of the base image from David Lawrence
- Exists in two flavors: graphics provided either via David's Web interface or via X11 forwarding through SSH connection
- Contents: cut-down EicRoot distribution (sufficient to start with: import generator ASCII files and run basic scripts)
- Users: 3 & myself; activity is in general low

STAR software (Shifter)

- Data productions at NERSC
- ~50 million CPU*hours per year

sPHENIX software (Singularity) -> the next two slides

- Since RACF batch farm at BNL is switching to SL7, consider to run analysis scripts embedded in SL6 Singularity image in order to avoid SL6-vs-SL7 inconsistencies

PHENIX

Containers

sPHENIX

Why now?

- RHIC Atlas Computing Facility (RACF) upgrades from SL6 to SL7.
- Even running SL6 compiled libs under SL7 results in small numerical differences – the dark art of floating point implementations
- I prefer for our ongoing production not to have to deal with this, a test with an RACF provided SL6 singularity container provides identical results.
- Singularity image mounts our disks, our software is loaded via afs

→Welcome to a use case of containers in a large scale production environment (PB sized dataset, 15000 cpu cores)

Docker was not an option because of security concerns

Some thoughts about containers

- RACF (and other facilities I suppose) will move to a model where the farm OS is a separate entity and jobs will run in containers (no more coordination between users and facility needed for OS upgrades/changes)
- Static containers are useless in a fast paced development environment, our libraries are outdated the minute they are published
- Need a mechanism to pull in our setup and software (currently afs, cvmfs coming)
- The above works on a farm – how about desktops or the (in)famous laptop on a plane?

Containers at JLab

- Container exists that includes multiple software packages that make up the JLab Common Environment (see right)
- JLEIC Geometry in format used by GEMC for simulation
- Example documentation including sample file of Pythia generated events
- Minimal desktop environment (xfwm4) w/ OpenGL graphics
- noVNC server to allow Web browser to be used for graphics interface
- Still Needed:
  - Detailed instructions for changing geometry / adding detectors
  - Instructions for Fast MC mode
  - Reconstruction

JLab Common Environment

docker hub: electronioncollider/jleic:1.0.3

singularity hub: electronioncollider/JEIC_singularity:1.0.3

JLEIC Desktop Environment via browser on host

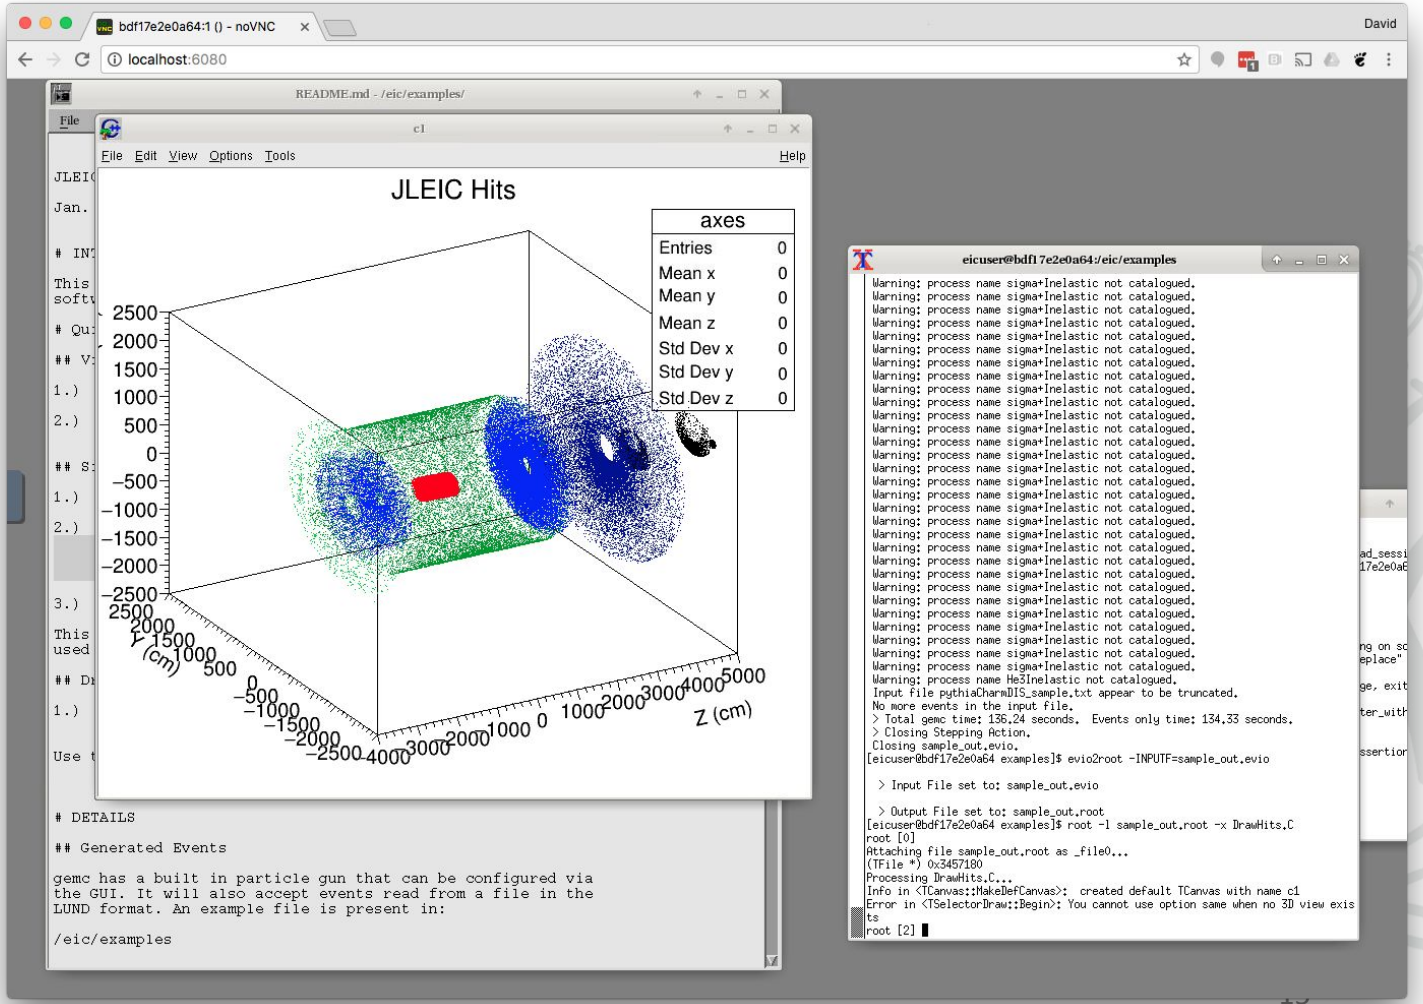

Containers at JLab

Also Singularity images available for interactive nodes:

```
ifarm1402.jlab.org> setenv https_proxy
https://jprox.jlab.org:8082
ifarm1402.jlab.org> singularity shell
shub://electronioncollider/jleic:1.0.3
Progress |==================================| 100.0%
Singularity: Invoking an interactive shell within container...

Singularity electronioncollider-jleic-master-1.0.3.simg:~>
```

Overview

- ✔ Container guidelines
- ✔ Containerization status updates
  - ANL
  - BNL
  - JLab
- Containers for batch processing

Containers for batch processing

- Intended use
  - Local batch farms
  - Grid e.g. OSG
- Local batch farms avoid Docker for security reasons
- OSG supports Docker or Singularity, but "singularity preferred containerization"
- NERSC uses Shifter

Containers for batch processing

- At this point, primary targeting Docker images which are then imported into Singularity (interactively or through singularity-hub.org)
- What are we losing by choosing this approach?

Documentation and Training

- SW Carpentry Tutorial on containers at JLab: May 23, 2018 (1 hour)

https://github.com/JeffersonLab/Swcarpentry-jlab-singularity

- Strongly singularity- and EIC-focused

Using Singularity at Jefferson Lab

Software containers are lightweight, stand-alone, executable packages that includes everything needed to run: code, runtime, system tools, system libraries, settings.

Containers are similar to virtual machines, however, the virtualization happens at a different point.

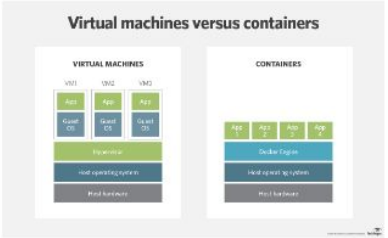

Image source: TechTarget

Although it can be hard to give a quick answer as to which one to use, Jack Wallen at TechRepublic boils it down to the following two questions:

- Do you need a full platform that can house multiple services? Go with a virtual machine.
- Do you need a single service that can be clustered and deployed at scale? Go with a container.

Prerequisites

You should have a basic familiarity with Unix shell navigation.

Schedule

	Setup	Download files required for the lesson
00:00	1. Introduction	Key question

Discussion and Next Steps

- We have condensed on a (somewhat) standard container layout
- From development of EIC containers to increasing user community of containers
  - Provide central place for users to get started
  - Could SW Carpentry framework form a basis?