



Containers

Stephen Bourke

Onsala Space Observatory

What is a container?

- It's a file system image (eg. tar archive) that gets used instead of the local file system
- An application to set things up (eg. singularity, docker)

When you run an application via a container that application is isolated from the rest of the OS

Possible by Linux kernel features

What can containers do

- Separate application from infrastructure
- Ease / Speed of deployment
- Archive / Repeatability
- Application isolation


Before:

The image shows a collection of terminal windows from a Linux environment, likely CentOS/RHEL 6, illustrating the initial setup of a software build system. The windows are titled with paths like `09_factor`, `00_root`, `04_astrol`, `90_syslibs`, `01_gcc`, `00_versions`, `99_tar_up`, and `05_casarest`.

- 09_factor**: Shows the installation of `00_versions` and `00_paths` directories, and the installation of `00_versions` and `00_paths` directories.
- 00_root**: Shows the installation of `00_versions` and `00_paths` directories, and the installation of `00_versions` and `00_paths` directories.
- 04_astrol**: Shows the installation of `00_versions` and `00_paths` directories, and the installation of `00_versions` and `00_paths` directories.
- 90_syslibs**: Shows the installation of `00_versions` and `00_paths` directories, and the installation of `00_versions` and `00_paths` directories.
- 01_gcc**: Shows the installation of `00_versions` and `00_paths` directories, and the installation of `00_versions` and `00_paths` directories.
- 00_versions**: Shows the installation of `00_versions` and `00_paths` directories, and the installation of `00_versions` and `00_paths` directories.
- 99_tar_up**: Shows the installation of `00_versions` and `00_paths` directories, and the installation of `00_versions` and `00_paths` directories.
- 05_casarest**: Shows the installation of `00_versions` and `00_paths` directories, and the installation of `00_versions` and `00_paths` directories.

The code snippets show the installation of various dependencies and the setup of a version control system for managing build instructions. The code is written in a shell script and includes comments in Swedish.

After



Introduction

2 ways to think of containers:

- As an application with all of its dependencies
- As an OS running alongside the native OS

A container may be limited in its use of system resources

- It can be restricted to a subset of cores / memory

Implementations

Singularity

- HPC Science & Engineering
- Containers used like any other application
- Native network comms
- Supported by many HPC centres

Docker

- PaaS / Micro Services
- Daemon runs and manages containers
- NAT
- Run on your own [cloud] resources

Singularity

You need admin right to create containers (mostly)

- Do this on your own machine

Copy container image to remote host

Run container as a normal user

Singularity commands

- **build**
 - `--sandbox`
 - Can put build specs in a text file
- **exec**
 - `--writable`
- **run**
 - Shortcut for most useful command in the container (defaults to bash)

Example

If you *really* need gvim on hebbe

```
sudo singularity build gvim recipe.gvim
```

```
scp gvim hebbe:
```

```
ssh hebbe
```

```
./gvim
```

```
# Hooray!
```

Real world examples

```
sudo singularity build --sandbox ubuntu/ docker://kernsuite/base
```

```
sudo singularity exec --writable ubuntu/ bash
```

(do stuff)

```
sudo singularity build ubuntu.simg ubuntu/
```

Part 2

Review: Containers

- A Container Image:
 - A file system that will be used as /
- A Running Container:
 - An application running with the image as its file system

Review: Container Image

- The Container Image (file system) can be very minimal
 - just the application, its libraries and a few low level files
- More common is that it will contain a base OS distribution
 - This is just because it's an easy way of making container images
 - Most of that stuff isn't needed

Containers: Process Isolation

- Linux kernel can isolate the container for rest of the OS
 - This is the default in docker
 - “instance” commands in Singularity
 - Multiple applications can be run in the same container

Inside the Container Image

- A typical linux filesystem
- Optionally some extra bits for convenience
 - “run” script, “exec” script, “environment” file, meta data file (Creator, date, etc)
- Default format is squashfs
 - Can be a plain directory structure
 - Can export an Image as a tar file

- You can also import a tar file as a container
 - But you don't really have to
 - Can just untar and use the directory as the container image

Archive Container Images

- Upload to singularity-hub or docker-hub
 - Run them on future versions of linux and reproduce results
 - Linux (the kernel) takes backwards compatibility very seriously
 - (If singularity, singularity-hub is still around :-P)

Singularity Summary

- Very straight forward and versatile
- Any idea I try seems to work
- GPU acceleration is a exception
 - Often possible but a bit messy
 - You're stuck with the host systems kernel module
 - Nvidia do not maintain backwards compatibility

More examples



CHALMERS
UNIVERSITY OF TECHNOLOGY