

# Working with Containers

Please open this in your browser to follow along:

<https://goo.gl/hhkKSP>

v.3.0 (2020-06-23)

# Agenda

1. The problem we're solving
2. Virtual machines vs containers
3. Docker: theory
4. Docker: practice
5. Composition and orchestration
6. Docker vs Singularity
7. Installing and testing Singularity
8. Creating and working with containers
9. Writing a Singularity definition file
10. Using host resources
11. Distributing Singularity containers
12. Cloud resources
13. Docker <-> Singularity interoperability
14. Extra credits (if time allows)

**The problem we're solving**

# Problem (for developers, part 1)

Suppose you're writing some software. To do that, you need the tooling for your language of choice (compiler / interpreter), language-specific libraries with development versions, and native system libraries.

You may need extra things depending on your project, such as a web server stack, a database, etc.

If you're lucky, it's easy to install at least some versions of all that on your OS, and get your development project going.

If you aren't, there are complex actions required to install your stack; Ones that you want to write down, or better yet - automate.

# Problem (for developers, part 2)

Okay, you got 1 project working. But now you want to work on two projects at once.

And those projects requires different software stacks, which are incompatible with each other.

Some programming languages offer solutions in terms of virtual environments or other project-specific ways to track dependencies, but this is only a partial solution.

# Problem (for developers, part 3)

Suppose have written some software.  
It works great on your machine.

However, eventually it has to leave your machine: has to run on your colleague's machine, or deployed in its production environment.

It can be a completely different flavour of OS, with a different set of libraries and supporting tools.

It can be difficult to test if you accounted for all those variations on your own development system.  
You may have things in your environment you're not even aware of that make a difference.

Your users could also be less technically inclined to deal with dependencies. You may wish to decrease this friction.

# Problem (for users)

Suppose you want to run or deploy some piece of software.

First off, you really would like some sort of turn-key solution. Of course there's none, there's only the source code.

The build instructions indicate 5-years-old out of date libraries on top of a similarly old OS distribution.

And no, the original developer is most certainly no longer available.

You also don't trust this software fully not to mess up your OS.

Or, you want to run it on a remote server for which you don't even have the privileges to comfortably install all the dependencies.

# Problem (for researchers)

Suppose you have a piece of scientific software you used to obtain some result.

Then someone half across the globe tries to reproduce it, and can't get it to run, or worse - is getting different results for the same inputs. What is to blame?

Or, even simpler: your group tries to use your software a couple of years after you left, and nobody can get it to work.

For a reproducible way to do science with the help of software, packaging just the source code might not be enough; the environment should also be predictable.

# Problem (for server/cluster administrators)

Suppose you have a hundred of users, each requesting certain software.

Some of it needs to be carefully built from scratch, as there are no prebuilt packages.

Some of the software works with mutually-incompatible library versions. Possibly even known-insecure ones.

Any such software change has to be injected in a scheduled maintenance window, but users want it yesterday.

And finally, *you most certainly don't* trust any of this software not to mess up your OS. From experience.

# What would be a solution?

- **A turnkey solution**

A recipe that can build a working instance of your software, reliably and fast.

- **BYOE: Bring Your Own Environment**

A way to capture the prerequisites and environment together with the software.

- **Mitigate security risks**

Provide a measure of isolation between the software running on a system. No security is perfect, but some is better than none.

# The Solution(s)

# Solution: Virtual Machines?

A virtual machine is an isolated instance of a **whole other "guest" OS** running under your "host" OS.

A **hypervisor** is responsible for handling the situations where this isolation causes issues for the guest.

From the point of view of the guest, it runs under its own, dedicated hardware. Hence, it's called **hardware-level virtualization**.

Most\* guest/host OS combinations can run: you can run Windows on Linux, Linux on Windows, etc.

---

\* MacOS guests being a complicated case due to licensing.

# Virtual Machines: the good parts

- **The BYOE principle is fully realized**

Whatever your environment is, you can package it fully, OS and everything.

- **Security risks are truly minimized**

Very narrow and secured bridge between the guest and the host means little (but not zero) opportunity for a bad actor to break out of isolation.

- **Easy to precisely measure out resources**

The contained application, together with its OS, has restricted access to hardware: you measure out its disk, memory and allotted CPU.

# Virtual Machines: the not so good parts

- **Operational overhead**

For every software packaged as a VM, the full underlying OS has to be run, and corresponding resources allocated.

- **Setup overhead**

Starting and stopping a virtual machine is not very fast, and/or requires saving its state.

Changing the allocated resources can be hard too.

- **Hardware availability**

The isolation between the host and the guest can hinder access to specialized hardware on the host system.

# Solution: Containers (on Linux)?

If your software expects Linux, there's a more direct and lightweight way to reach similar goals.

Linux kernel allows to isolate processes from the rest of the system, presenting them with their own view of the system.

You can package entire other Linux distributions, and with the exception of the host kernel, all the environment can be different for a process.

From the point of view of the application, it's running on the same hardware as the host, hence containers are sometimes called **operating system level virtualization**.

# Containers: the good parts

- **Lower operational overhead**

You don't need to run a whole second OS to run an application.

- **Lower startup overhead**

Setup and teardown of a container is much less costly.

- **More hardware flexibility**

You don't have to dedicate a set portion of memory to your VM well in advance, or contain your files in a fixed-size filesystem.

Also, the level of isolation is up to you. You may present devices on the system directly to containers if needed.

# Containers: the not so good parts

- **Kernel compatibility**

Kernel is shared between the host and the container, so there may be some incompatibilities.

- **Security concerns**

The isolation is thinner than in VM case, and kernel of the host OS is directly exposed.

- **Linux on Linux**

Containers are mostly a Linux technology. You need a Linux host (or a Linux VM) to run containers, and only Linux software can run.

# Containers on different OSes

## Running Linux containers on other systems

Some container solutions (e.g. Docker) offer to run on host systems other than Linux.

It's important to know that this relies on a Linux VM as a "base" for the running containers, so it comes with VM problems.

# Containers on different OSes

## Running Linux containers on other systems

Some container solutions (e.g. Docker) offer to run on host systems other than Linux.

It's important to know that this relies on a Linux VM as a "base" for the running containers, so it comes with VM problems.

## Windows containers

Microsoft has recently embraced the idea of operating system level virtualization with [Windows Containers](#).

The principles are the same (a process is presented a different environment but still runs under the same kernel).

# History of containers

The idea of running an application in a different environment is not new to UNIX-like systems.

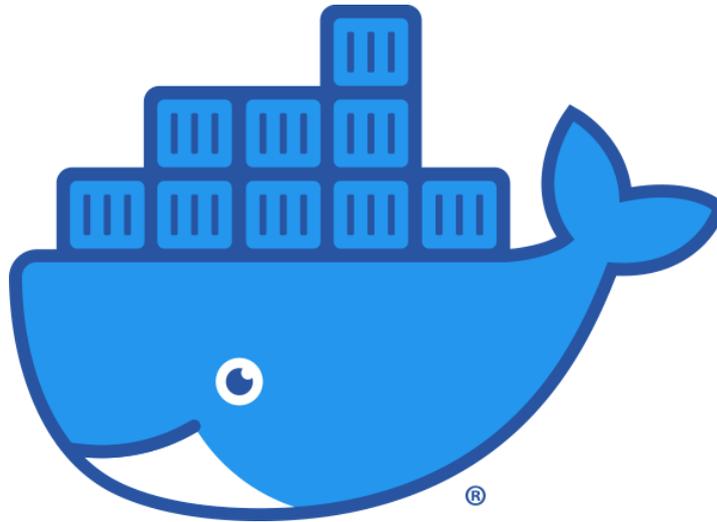
Perhaps the first effort in that direction is the `chroot` command and concept (1982): presenting applications with a different view of the filesystem (a different root directory `/`).

This minimal isolation was improved in in FreeBSD with `jail` (2000), separating other resources (processes, users) and restricting how applications can interact with each other and the kernel.

Linux developed facilities for isolating and controlling access to some processes with namespaces (2002) and `cgroups` (2007).

Those facilities led to creation of solutions for containerization, notably `LXC` (2008), `Docker` (2013) and `Singularity` (2016).

# Docker



# Docker

- Docker came about in 2013 and since has been on a meteoric rise as the golden standard for containerization technology.
- A huge amount of tools is built around Docker to build, run, orchestrate and integrate Docker containers.
- Many cloud service providers can directly integrate Docker containers. Docker claims x26 resource efficiency improvement at cloud scale.

# Docker architecture

Docker has a client-server architecture. Each system using Docker has to run a **Docker daemon**: a persistent, privileged service that can manage containers and other resources.

The daemon exposes an API that clients can use to control it. The classic example of a client is the `docker` command line utility.

Together, this server+client setup is called the **Docker Engine**.

Several Docker daemons can work together as a cluster, spreading workloads across several host machines.

# Dockerfiles, images and containers

A **Dockerfile** is a step by step recipe on how to build and run a container. It uses already-existing images, local and remote files to populate the container.

An **image** is a read-only template with instructions for creating a container. This includes the base filesystem for a container. An image is a result of executing instructions in a Dockerfile.

A **container** is a runnable instance of an image. A container can be created from an image, started (according to instructions in the image), stopped, attached to networks, have storage attached to it, etc.

# Where do images come from?

Most of the images are built as further modifications to some base image. E.g. you would take a base Ubuntu image and add commands to install your software.

How to bootstrap an image from nothing is outside the scope of this course.

So, where do base images come from? Docker uses special servers called registries to store images, which will be downloaded as needed.

# Where do images come from?

Most of the images are built as further modifications to some base image. E.g. you would take a base Ubuntu image and add commands to install your software.

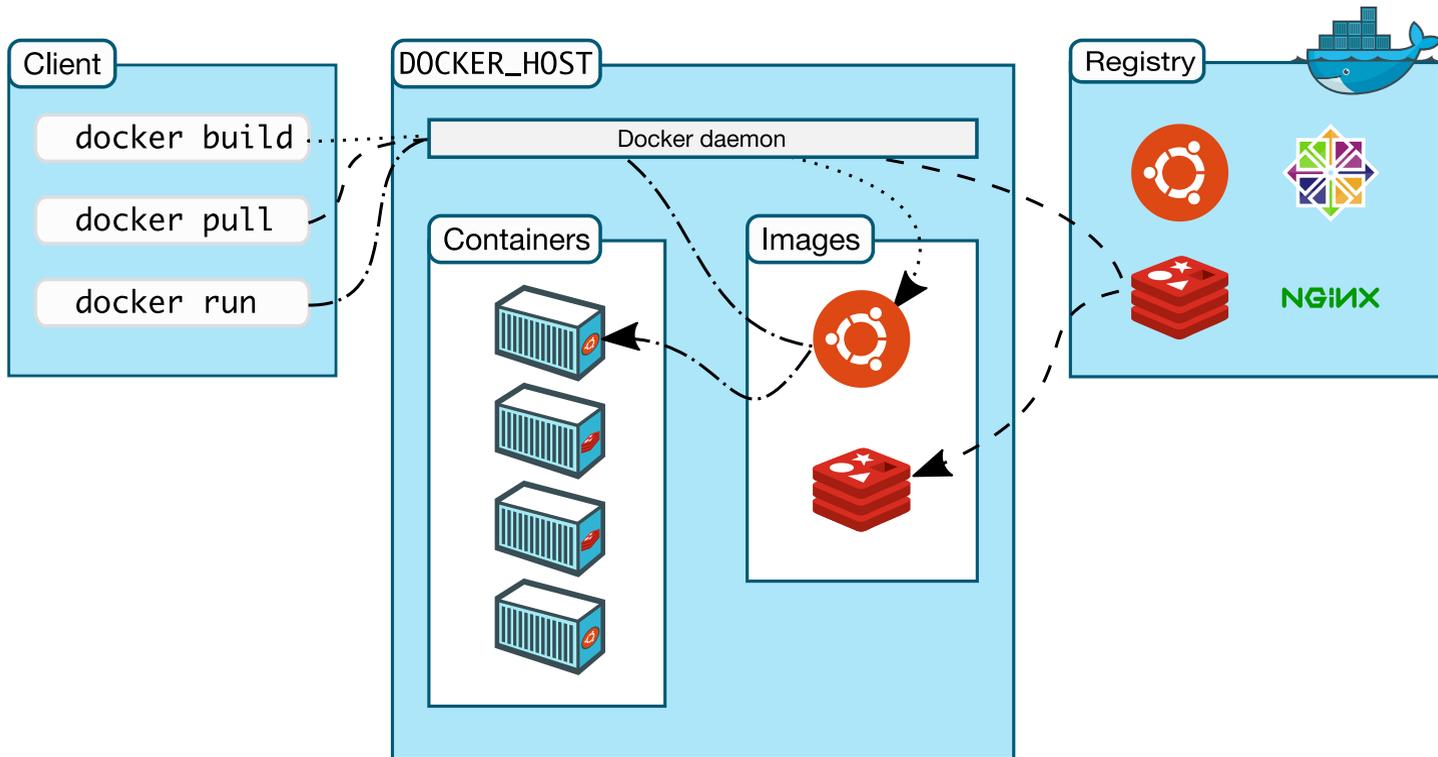
How to bootstrap an image from nothing is outside the scope of this course.

So, where do base images come from? Docker uses special servers called registries to store images, which will be downloaded as needed.

The central public registry is [Docker Hub](#), where anyone can put public images and has a freemium model for private images.

It's possible to deploy your own registry servers.

# Docker: a diagram



Source: [Docker overview](#)

# Docker Engine versions

Docker Engine comes in two editions:

- **Docker Enterprise Edition (EE)** is the paid version of Docker, with extensive official support, certifications, priority updates, long term support, extra features, and private image hosting. A "batteries included" version.
- **Docker Community Edition (CE)** is the open-source, free version of Docker provided directly by Docker that anyone can install on compatible systems, but has limited support. A "your own risk" version.

Your Linux distribution may have its own maintained version of Docker (based on the open-source version). This has no direct support for Docker, instead it's on your distro's maintainers to keep it up to date and solve issues.

# Docker Desktop

For Windows and MacOS, Docker provides a package called **Docker Desktop** that simplifies running Docker under those systems.

It's important to stress out that those packages run a base Linux system as a VM using one of the OS-appropriate methods as the base for running Docker.

On Windows, this requires either Hyper-V (compatible hardware + Pro version or higher) or WSL2 (requires latest Windows 10 2004).

Also, Docker Desktop on Windows allows managing Windows containers (under Hyper-V).

# Using Docker

# Installing Docker

This course assumes a Docker Engine CE installation (on Linux). I will provide the standard procedure for Ubuntu Linux.

If you're using your own Linux, please proceed with the installation of Docker CE as per [documentation](#).

If your distribution isn't listed, you'll need to look up installation instructions yourself.

If you do not want to use your own Linux, please ensure at this point that you can connect to the training system.

# Installing Docker (Ubuntu)

To install Docker on Ubuntu, we'll follow the documentation at <https://docs.docker.com/engine/install/ubuntu/>

There's also a convenience script at <https://get.docker.com/>

## *Exercise:*

If you're using your own Ubuntu Linux, follow the instructions.

If you're using your own other Linux, install Docker independently.

If you're using the training VM, you don't need to do anything here other than try to connect to the system with SSH.

# Hello, Docker!

If Docker is properly installed, we should be able to run our first container.

```
$ sudo docker run hello-world
```

*Exercise:*

Run this on your system.

# Hello, Docker!

If Docker is properly installed, we should be able to run our first container.

```
$ sudo docker run hello-world
```

## *Exercise:*

Run this on your system.

If successful, this will give some technical output, then:

```
Hello from Docker!  
This message shows that your installation appears to be working correctly.
```

Followed by an explanation of what just happened. Let's go over that in more detail.

## Aside: why sudo?

Docker daemon runs with high privileges, and interacting with it requires similarly high privileges. By default, you need to be effective root (using sudo) to interact with it.

Docker installation normally creates a group `docker` which is authorized to run `docker` without access to normal `sudo`.

# Aside: why sudo?

Docker daemon runs with high privileges, and interacting with it requires similarly high privileges. By default, you need to be effective root (using sudo) to interact with it.

Docker installation normally creates a group docker which is authorized to run docker without access to normal sudo.

However, it must be stressed: **Having access to the Docker daemon can allow privilege escalation to root.** So assigning the docker group to a user is as dangerous as allowing sudo and is mostly just a convenience.

## *Note:*

To switch to using the docker group after using sudo, you probably need to fix permissions:

```
$ sudo chown -R username:username ~/.docker
```

# Hello, Docker! (in slo-mo, pt. 1)

Running the `docker` command runs a Docker **client** that talks to a locally running Docker daemon.

`docker run hello-world` is a command to create a new **container** using an **image** specified by `hello-world`.

Assuming a fresh install of Docker, no such image exists locally (in the images cached for the current user).

Therefore, Docker looks for the image in configured registries, which is Docker Hub by default.

# Image name specification

Generally, images in a registry are organized by two components: a **name** and a **tag**:

```
name:tag
```

For Docker Hub, except for a few handpicked images, the name must be **namespaced** with the Docker user ID, e.g. `godlovedc/lolcow` means an image `lolcow` from user `godlovedc`.

Tags are a way to version multiple images of the same name.

Specifying a tag is optional: it defaults to a tag named `latest` (which does not have any special semantics, just a fixed name, like `master` in Git).

So `docker run hello-world` is looking for an image named `hello-world` and its tag `latest`.

# Hello, Docker! (in slo-mo, pt. 2)

Not finding an image `hello-world:latest` in the local cache, Docker contacts the registry at Docker Hub for it, and finds this:

[https://hub.docker.com/\\_/hello-world](https://hub.docker.com/_/hello-world)

It downloads some objects (layers, more on that later) from Dockerhub, after which the local image repository has the image `hello-world:latest`

## *Exercise:*

Verify this with `docker image ls` to list available images.

# Hello, Docker! (in slo-mo, pt. 3)

Now that the image `hello-world:latest` is available, `docker run` creates a fresh **container** based on it.

It then runs the specified default executable within that container.

By default, `docker run` runs the container **attached**: the output of the program will be piped to the standard output of the shell.

Therefore, you're seeing the output of that program. After the program terminates, the container stops, and you're back at the shell.

## *Note:*

The exact specifics of attachment can be configured; for example, for interactive programs one can use `-it` flags to attach input and pseudo-TTY.

# New container every time?

If `docker run` creates a new container, what happens to that container after it's done?

If we try to run `docker container ls`, it will show nothing. But in fact, by default it only shows running containers.

## *Exercise:*

Run `docker run hello-world` a few more times, then examine the full container list with `docker container ls -a`

How to clean up this mess?

# Cleanup after run

Containers (and images) can be deleted with a corresponding `rm` command, e.g. `docker image rm hello-world`

## *Exercise:*

Try to remove the `hello-world` image. Why did it fail?

Try to remove a container. What kind of name does it expect?

Besides `rm`, you can use `prune` to delete all unused images and stopped containers, respectively.

## *Note:*

`docker run --rm` will delete the container (but not the image) automatically after a container stops.

# Interactive container

Let's try running an Ubuntu container interactively, as `hello-world` suggests.

## *Exercise:*

Execute `docker run -it ubuntu bash` to get an interactive terminal.

As a result, you have a root terminal in a minimal Ubuntu installation. Its filesystem comes from the image and is completely separate from the host OS.

## *Exercise:*

Create a file in the root folder of the container.

If you exit from the container (e.g. with `exit`), the container will be stopped, but not deleted.

# Running an existing container

## *Exercise:*

Exit from the container, locate the container ID and/or name with `docker container ls -a`.

To run the same container again, use `docker start -ia <name>`.

Verify that your custom file still exists.

## *Note:*

-a flag stands for "attach". It is equivalent to -t for run.

# Persistent changes

How does Docker maintain changes?

A Docker image is **immutable**: once created, it can be deleted but not modified.

So what allows a container, an **instance** of an image, to have unique changes?

# Persistent changes

How does Docker maintain changes?

A Docker image is **immutable**: once created, it can be deleted but not modified.

So what allows a container, an **instance** of an image, to have unique changes?

Taking the image as a base, Docker **overlays** any changes on top, storing them as a separate **layer** unique to the container.

If a container is deleted, those changes are lost; it's possible to make a new image out of the modified state of a container.

# Docker layers

Docker frequently operates with filesystem layers.

A base layer is just a filesystem with files; further layers store only changes to the base layer (added, modified and deleted files).

An image can (and often does) consist of multiple layers corresponding to build stages of the image, as we'll see in a moment.

When a file is queried in a container, all layers are examined in order, and the last layer that provides that file is used.

As mentioned, image layers are immutable; a container has one extra read-write layer added on top of its source image.

# Making a new container interactively

## *Exercise:*

Run a new Ubuntu container, specifying a name for it:

```
$ docker run -it --name my-ubuntu ubuntu bash
```

Follow the other commands

Let's install some software that is not present in the base image. Running inside the container:

```
# apt-get update  
# apt-get install -y figlet
```

Test it:

```
# figlet Hello Docker
```

# Creating a new image

Let's exit the container and create a new image out of it:

```
# exit  
$ docker container commit my-ubuntu my-figlet:1
```

This creates a new image named `my-figlet` with tag `1` out of the current container state.

We can now run `figlet` from this image:

```
$ docker run --rm my-figlet:1 figlet Testing
```



Testing

# Using a Dockerfile

Let's reproduce our manual steps in a Dockerfile recipe.

Create a new file, `Dockerfile`, with the following contents:

```
FROM ubuntu
RUN apt-get update
RUN apt-get install -y figlet
```

Let's build a version 2 of `my-figlet` image:

```
$ docker build -t my-figlet:2 .
[...]
```

```
$ docker run --rm my-figlet:2 figlet Testing version 2
```

It behaves identically to our first version.

# Defining a default command

Modify the Dockerfile:

```
FROM ubuntu
RUN apt-get update
RUN apt-get install -y figlet
CMD ["figlet", "Default command"]
```

Build a new version and note that we can invoke `run` without any arguments:

```
$ docker build -t my-figlet:3 .
[...]
$ docker run --rm my-figlet:3
```

Same applies to `start` without arguments.

# Where do arguments go?

Why does specifying a shell command after `run` actually execute it?

The arguments are passed to an **entrypoint** of a container, which by default invokes a shell inside the container.

This can be overridden:

```
FROM ubuntu
RUN apt-get update
RUN apt-get install -y figlet
ENTRYPOINT ["figlet"]
CMD ["Default command"]
```

## *Exercise:*

Try building a version 4 with this dockerfile, and try running it with and without arguments.

# Layer cache

In the previous build commands, you may notice the following line in the output:

```
---> Using cache
```

Docker skips build steps that didn't change between builds, by reusing layers from other (previous) images.

This is beneficial: it saves both disk space and computation time.

But in case you really need to rerun steps, you can use `--no-cache` for building.

## *Note:*

To force an update to the base image, which is also cached, use `docker pull`.

# Uploading an image to Docker Hub

## *Exercise:*

1. Create a Docker ID account at <https://hub.docker.com/>
2. Create a public repository named `my-figlet`
3. Rename the latest version of `my-figlet` to match the repository name:

```
$ docker tag my-figlet:4 username/my-figlet
```

1. Login into Docker Hub with `docker login`
2. Upload (push) your image

```
$ docker push username/my-figlet
```

1. Remove the local copy of `username/figlet`
2. Try to run it. It will be downloaded from Docker Hub.

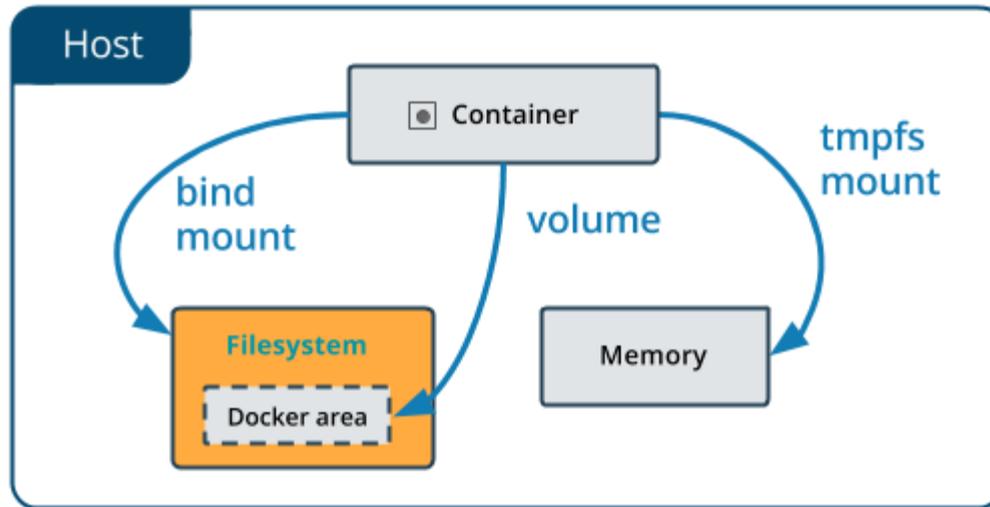
# Other ways of providing data

Docker allows mounting data into locations inside the container at runtime.

There are 3 main types of mounts:

- **Volumes** are Docker-managed directories that can be mounted inside a container. Docker keeps track of what containers rely on what volumes, allows sharing volumes between containers, and other nice features. The host OS is not supposed to modify these files.
- **Bind-mounts** directly map a folder on the host into the container. Changes made by the host are instantly reflected inside the container and vice versa.
- **tmpfs** mounts are temporary locations in memory that are cleared as soon as a container stops and do not affect the host file system.

# Mounts example



The following prints the host OS hostname:

```
$ docker run --rm --mount type=bind,source=/etc,target=/host_etc \
  ubuntu cat /host_etc/hostname
```

Documentation: <https://docs.docker.com/storage/bind-mounts/>

# Composition and orchestration

# Docker networking, pt. 1

Docker allows you to expose network ports inside the container on the host OS. This allows you to have web applications inside containers, for example.

## *Exercise:*

Run the following command to start an Nginx webserver, mapping the port 5000 on the host to 80 in the container.

```
$ docker run --rm -p 5000:80 nginx
```

Access the server with <http://localhost:5000/>

## *Note:*

If you're running on the training VM, add your username number to 5000, e.g. 5003 for training03, and access using the training machine's IP instead of localhost.

# Docker networking, pt. 2

Docker also allows containers to communicate on a dedicated network between them (**bridge networking**).

Multiple docker daemons can create an **overlay network** between them, allowing cross-host communication between containers.

Manually setting up such networks is an advanced topic not covered here; usually other tools take care of it.

Documentation: <https://docs.docker.com/network/>

# Multi-container applications and microservices

The ability of multiple containers to talk to each other allows to package independent parts of an application stack, e.g. a web server and a database, into separate independent containers.

This allows for easier maintenance, as those containers can be updated independently as long as they conform to protocols between them.

This approach encourages further splitting of components into independent **microservices** - modular components talking through established protocols.

# Container composition

To ensure that multi-container workload is running as expected, **composition** tools are used.

Provided an overall description of a service, composition tools ensure that all containers are running and configured as needed.

An example composition tool is Docker Compose, provided by Docker. It needs to be installed separately.

# Docker compose example

*Exercise:*

Follow along with the tutorial at  
<https://docs.docker.com/compose/gettingstarted/>

# Orchestration

Going beyond composition, **orchestration** allows container workloads to be spread between multiple systems, be scaled to multiple load-balanced instances, etc.

Examples of orchestration tools:

- Docker Swarm
- Kubernetes (or "k8s")

Orchestration is beyond the scope of this course.

# Docker vs Singularity

Why did another technology emerge?

# Docker concerns

- Docker uses a pretty complicated model of images/volumes/metadata, orchestrating swarms of those containers to work together, and it not always very transparent with how those are stored.
- Also, isolation features require superuser privileges; Docker has a persistent daemon running with those privileges and many container operations require root as well.

# Docker concerns

- Docker uses a pretty complicated model of images/volumes/metadata, orchestrating swarms of those containers to work together, and it not always very transparent with how those are stored.
- Also, isolation features require superuser privileges; Docker has a persistent daemon running with those privileges and many container operations require root as well.

Both of those issues make Docker undesirable in applications where you don't wholly own the computing resource - HPC environments.

Out of those concerns, and out of scientific community, came Singularity.

# Singularity

Singularity was created in 2016 as an HPC-friendly alternative to Docker. It is still in rapid development.

# Singularity

Singularity was created in 2016 as an HPC-friendly alternative to Docker. It is still in rapid development.

- It's usually straightforward to convert a Docker container to a Singularity image.

This gives users access to a vast library of containers.

# Singularity

Singularity was created in 2016 as an HPC-friendly alternative to Docker. It is still in rapid development.

- It's usually straightforward to convert a Docker container to a Singularity image.

This gives users access to a vast library of containers.

- Singularity uses a monolithic, image-file based approach. Instead of dynamically overlaid layers.

You build a single file on one system and simply copy it over or archive it.

This addresses the "complex storage" issue with Docker.

# Singularity and root privileges

The privilege problem was a concern from the ground-up, to make Singularity acceptable for academic clusters.

# Singularity and root privileges

The privilege problem was a concern from the ground-up, to make Singularity acceptable for academic clusters.

- Addressed by having a `setuid`-enabled binary that can accomplish container startup and drop privileges ASAP.

# Singularity and root privileges

The privilege problem was a concern from the ground-up, to make Singularity acceptable for academic clusters.

- Addressed by having a `setuid`-enabled binary that can accomplish container startup and drop privileges ASAP.
- Privilege elevation inside a container is impossible: `setuid` mechanism is disabled inside the container, so to be root inside, you have to be root outside.
  - Recent `--fakeroot` option allows for administrative actions inside the container, without affecting effective access to the host.

# Singularity and root privileges

The privilege problem was a concern from the ground-up, to make Singularity acceptable for academic clusters.

- Addressed by having a `setuid`-enabled binary that can accomplish container startup and drop privileges ASAP.
- Privilege elevation inside a container is impossible: `setuid` mechanism is disabled inside the container, so to be root inside, you have to be root outside.
  - Recent `--fakeroot` option allows for administrative actions inside the container, without affecting effective access to the host.
- Users don't need explicit root access to operate containers (at least after the initial build).

# Singularity and HPC

Thanks to the above improvements over Docker, HPC cluster operators are much more welcoming to the idea of Singularity support.

As a result of a joint Pipeline Interoperability project between Swiss Science IT groups, the UniBE Linux cluster UBELIX started to support Singularity.

Once your software is packaged in Singularity, it should work across all Science IT platforms supporting the technology.

# Singularity niche

When is Singularity useful over Docker?

- The major use case was and still is **shared systems**: systems where unprivileged users need the ability to run containers.

However, an admin still needs to install Singularity for it to function.

- Singularity is useful as an alternative to Docker. If you have admin privileges on the host, Singularity can do more than in unprivileged mode.

It doesn't have the same level of ecosystem around it, but currently gaining features such as OCI runtime interface, native Kubernetes integration and own cloud services.

# Singularity "sales pitch"

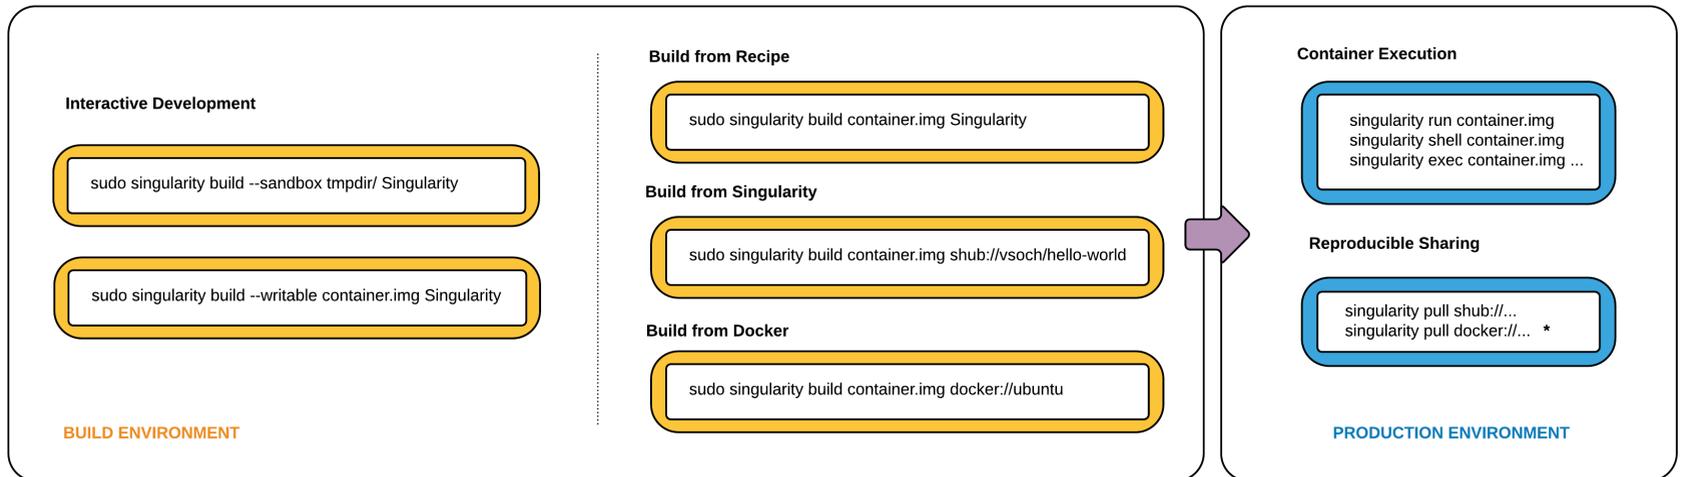
Quoting from Singularity Admin documentation:

*Untrusted users (those who don't have root access and aren't getting it) can run untrusted containers (those that have not been vetted by admins) safely.*

This won over quite a few academic users; for a sampling:

<https://sylabs.io/case-studies>

# Singularity workflow



\* Docker construction from layers not guaranteed to replicate between pulls

1. Interactively develop steps to construct a container.
2. Describe the steps in a recipe.
3. Build an immutable container on own machine.
4. Deploy this container in the production environment.

# Working with Singularity

Installation and basic use

# Singularity versions

There are two major branches of Singularity:

- 2.x branch (ended at 2.6.1): legacy branch with no active development, but still deployed in places.
- 3.x branch (currently at 3.5.3): actively developed branch, with most of the code completely rewritten in Go.

# Singularity versions

There are two major branches of Singularity:

- 2.x branch (ended at 2.6.1): legacy branch with no active development, but still deployed in places.
- 3.x branch (currently at 3.5.3): actively developed branch, with most of the code completely rewritten in Go.

Due to freshness of code and new Go dependency, 3.x adoption is slow. This course will cover 3.4.2 release (for UBELIX compatibility).

# Singularity versions

There are two major branches of Singularity:

- 2.x branch (ended at 2.6.1): legacy branch with no active development, but still deployed in places.
- 3.x branch (currently at 3.5.3): actively developed branch, with most of the code completely rewritten in Go.

Due to freshness of code and new Go dependency, 3.x adoption is slow. This course will cover 3.4.2 release (for UBELIX compatibility).

Singularity aims to be backwards-compatible: containers built with earlier versions should work with newer ones.

# Installing Singularity

Installing Singularity from source is probably preferred, as it's still a relatively new piece of software.

Instructions at: <https://sylabs.io/guides/3.4/user-guide/installation.html#install-on-linux>

# Installing Singularity

Installing Singularity from source is probably preferred, as it's still a relatively new piece of software.

Instructions at: <https://sylabs.io/guides/3.4/user-guide/installation.html#install-on-linux>

It is required to install Golang compiler  $\geq 1.12$  as a **build** dependency. It is not required to run the compiled software.

On Ubuntu, Go can be installed with

```
sudo snap install --classic go
```

## *Exercise:*

If you want to try installing Singularity on your Linux system, follow the build instructions and version 3.4.2

If you're using the remote training machine, skip this step.

# Using Singularity

If you followed build instructions, you should now have singularity available from the shell.

```
user@host:~$ singularity --version  
singularity version 3.4.2-1
```

# Using Singularity

If you followed build instructions, you should now have singularity available from the shell.

```
user@host:~$ singularity --version
singularity version 3.4.2-1
```

The general format of Singularity commands is:

```
singularity [<global flags>] <command> [<command flags>] [<arguments>]
```

Singularity is pretty sensitive to the order of those.

Use `singularity help [<command>]` to check built-in help.

You can find the configuration of Singularity under `/usr/local/etc/singularity` if you used the default prefixes.

# Container images

A Singularity image is, for practical purposes, a filesystem tree that will be presented to the applications running inside it.

# Container images

A Singularity image is, for practical purposes, a filesystem tree that will be presented to the applications running inside it.

A Docker container is built with a series of *layers* that are stacked upon each other to form the filesystem. Layers are collections of updates to files, and must be inspected to find the latest version of the file.

Singularity collapses those into a single, portable file.

# Container images

A Singularity image is, for practical purposes, a filesystem tree that will be presented to the applications running inside it.

A Docker container is built with a series of *layers* that are stacked upon each other to form the filesystem. Layers are collections of updates to files, and must be inspected to find the latest version of the file.

Singularity collapses those into a single, portable file.

A container needs to be somehow bootstrapped to contain a base operating system before further modifications can be made.

# Pulling Docker images

The simplest way of obtaining a working Singularity image is to pull it from either Docker Hub or Singularity Hub.

Let's try it with CentOS 6:

```
user@host:~$ singularity pull docker://centos:6
```

This will download the layers of the Docker container to your machine and assemble them into an image.

The result will be stored as `centos_6.sif`

# Pulling Docker images

```
user@host:~$ singularity pull docker://centos:6
INFO:   Starting build...
Getting image source signatures
Skipping fetch of repeat blob sha256:ff50d722b38227ec8f2bbf0cdbce428b66745077c1
Copying config sha256:5d1ece75fd80b4dd0e4b2d78a1cfebbabad9eb3b5bf48c4e1ba7f9dd2
 1.51 KiB / 1.51 KiB [=====] 6
Writing manifest to image destination
Storing signatures
INFO:   Creating SIF file...
INFO:   Build complete: centos_6.sif
```

Note that this **does not require sudo or Docker!**

## *Exercise:*

Pull the CentOS 6 image from Dockerhub with the above command

# Entering shell in the container

To test our freshly-created container, we can invoke an interactive shell to explore it with **shell**:

```
user@host:~$ singularity shell centos_6.sif  
Singularity centos_6.sif:~>
```

At this point, you're within the environment of the container.

We can verify we're "running" CentOS:

```
Singularity centos_6.sif:~> cat /etc/centos-release  
CentOS release 6.10 (Final)
```

# User/group within the container

Inside the container, we are the same user:

```
Singularity centos_6.sif:~> whoami  
user  
Singularity centos_6.sif:~> exit  
user@host:~$ whoami  
user
```

We will also have the same groups.

That way, if any host resources are mounted in the container, we'll have the same access privileges.

# Root within the container

If we launched singularity with sudo, we would be root inside the container.

```
user@host:~$ sudo singularity shell centos_6.sif
Singularity centos_6.sif:/home/user> whoami
root
```

# Root within the container

If we launched singularity with sudo, we would be root inside the container.

```
user@host:~$ sudo singularity shell centos_6.sif
Singularity centos_6.sif:/home/user> whoami
root
```

**Most importantly:** setuid mechanism will not work within the container. Once launched as non-root, no command can elevate your privileges.

A `--fakeroot` option gives us root inside the container without root outside, but does not affect certain privileged operations and access to host resources.

# Default mounts

In addition to the container filesystem, by default:

- user's home folder,
- /tmp,
- /dev,
- the folder we've invoked Singularity from

are accessible inside the container.

# Default mounts

In addition to the container filesystem, by default:

- user's home folder,
- /tmp,
- /dev,
- the folder we've invoked Singularity from

are accessible inside the container.

The idea is to provide minimal friction working with software inside the container: no need for extra mounts to access data or store preferences.

It is possible to override this default behavior.

# Default mounts

```
user@host:~$ singularity shell centos_6.sif
Singularity centos_6.sif:~> ls ~
[..lists home folder..]
Singularity centos_6.sif:~> touch ~/test_container
Singularity centos_6.sif:~> exit
user@host:~$ ls ~/test_container
/home/user/test_container
```

The current working directory inside the container is the same as outside at launch time.

# Running a command directly

Besides the interactive shell, we can execute any command inside the container directly with **exec**:

```
user@host:~$ singularity exec centos_6.sif cat /etc/centos-release  
CentOS release 6.10 (Final)
```

## *Exercise:*

Invoke the python interpreter with `exec`.

Compare the version with the host system.

# Modifying containers

Let's make our own

# Modifying the container

Let's try to install some software in the container.

```
user@host:~$ singularity shell centos_6.sif
Singularity centos_6.sif:~> fortune
bash: fortune: command not found
```

fortune is not part of the base image. Let's try installing it.

# Modifying the container

Let's try to install some software in the container.

```
user@host:~$ singularity shell centos_6.sif
Singularity centos_6.sif:~> fortune
bash: fortune: command not found
```

fortune is not part of the base image. Let's try installing it.

```
Singularity centos_6.sif:~> exit
user@host:~$ sudo singularity shell centos_6.sif
Singularity centos_6.sif:~> whoami
root
Singularity centos_6.sif:~> yum -y --enablerepo=extras install epel-release
[...]
[Errno 30] Read-only file system: '/var/lib/rpm/.rpm.lock'
[...]
^C
```

Despite having root, we can't write to the filesystem.

# Images and overlays

Singularity image files are read-only squashfs filesystems.

Singularity can use an **overlay**: a layer on top of the image that holds changes to it.

# Images and overlays

Singularity image files are read-only squashfs filesystems.

Singularity can use an **overlay**: a layer on top of the image that holds changes to it.

Overlays can be persistent (stored in a folder) or temporary. Singularity 2.x uses a temporary overlay by default.

```
user@host:~$ singularity shell --writable-tmpfs centos_6.sif
Singularity centos_6.sif:~> touch /test
Singularity centos_6.sif:~> ls /test
/test
```

```
user@host:~$ mkdir persistent_overlay
user@host:~$ sudo singularity shell --overlay persistent_overlay centos_6.sif
Singularity centos_6.sif:~> touch /test
Singularity centos_6.sif:~> ls /test
/test
```

# Sandbox containers

A more conventional way to write to a container is to use **sandbox** format, which is just a filesystem tree stored in a folder.

```
$ singularity build --sandbox centos-writable docker://centos:6
$ ls centos-writable/
bin dev environment etc home lib lib64 lost+found media mnt opt
proc root sbin selinux singularity srv sys tmp usr var
```

Building sandbox containers used to require root, but not anymore.

# Sandbox containers

A more conventional way to write to a container is to use **sandbox** format, which is just a filesystem tree stored in a folder.

```
$ singularity build --sandbox centos-writable docker://centos:6
$ ls centos-writable/
bin dev environment etc home lib lib64 lost+found media mnt opt
proc root sbin selinux singularity srv sys tmp usr var
```

Building sandbox containers used to require root, but not anymore.

Passing `--writable` to `shell` or `exec` will now enable changes:

```
$ singularity shell --fakeroot --writable centos-writable
Singularity centos-writable:~> touch /test
Singularity centos-writable:~> ls /test
/test
Singularity centos-writable:~> exit
$ ls centos-writable/test
centos-writable/test
```

# Writing to a container, finally:

We should now be able to enter it **in writable mode** and install software:

```
user@host:~$ singularity shell --fakeroot --writable centos-writable
Singularity centos-writable:~> yum -y --enablerepo=extras install epel-release
[...]
Singularity centos-writable:~> yum -y install fortune-mod
[...]
Singularity centos-writable:~> exit
user@host:~$ singularity exec centos-writable fortune
[some long-awaited wisdom of a fortune cookie]
```

# Default run script

A container can have a "default" command which is run without specifying it.

Inside the container, it's `/singularity`. Let's try modifying it:

```
user@host:~$ nano centos-writable/singularity
```

By default you'll see a sizeable shell script.

```
#!/bin/sh
OCI_ENTRYPOINT=''
OCI_CMD="/bin/bash"
CMDLINE_ARGS=""
# [...] #
```

# Custom run script

We installed fortune, so let's use that instead:

```
#!/bin/sh
exec /usr/bin/fortune "$@"
```

Now we can invoke it with **run**:

```
user@host:~$ singularity run centos-writable
[..some wisdom or humor..]
```

# Converting to final container

One way to produce a "final" container is to convert it from the sandbox version:

```
user@host:~$ singularity build fortune.sif centos-writable  
[...]
```

Now we can test our container:

```
user@host:~$ singularity run fortune.sif  
[..some more wisdom..]
```

# Running a container directly

Note that the container file is executable:

```
user@host:~$ ls -lh fortune.sif
-rwxr-xr-x 1 root root 99M Feb 30 13:37 fortune.sif
```

If we run it directly, it's the same as invoking `run`:

```
user@host:~$ ./fortune.sif
[..a cracking joke..]
```

This does require to have `singularity` installed on the host, however, and is just a convenience.

# Container definition files

Making the container reproducible

# Making the container reproducible

Instead of taking some base image and making changes to it by hand, we want to make this build process reproducible.

This is achieved with definition files called **Definition files**, historically also called "recipes".

Let's try to retrace our steps to obtain a fortune-telling CentOS.

## *Exercise:*

Open a file called `fortune.def` in an editor, and prepare to copy along.

# Bootstrapping

The definition file starts with a header section.

The key part of it is the `Bootstrap`: configuration, which defines how we obtain the "base" image.

There are multiple types of bootstrap methods:

- pull an image from a cloud service such as `docker`
- using `yum/debootstrap` on the host system to bootstrap a similar one
- `localimage` to base off another image on your computer

We'll be using the `Docker` method.

```
Bootstrap: docker
From: centos:6
```

# Setting up the container

There are 2 sections for setup commands (essentially shell scripts):

1. **%setup** for commands to be executed **outside the container**.

You can use `$SINGULARITY_ROOTFS` to access the container's filesystem, as it is mounted on the host during the build.

2. **%post** for commands to be executed **inside** the container.

This is a good place to set up the OS, such as installing packages.

# Setting up the container

Let's save the name of the build host and install fortune:

```
Bootstrap: docker
From: centos:6

%setup
  hostname -f > $$SINGULARITY_ROOTFS/etc/build_host

%post
  yum -y --enablerepo=extras install epel-release
  yum -y install fortune-mod
  yum clean all
```

# Adding files to the container

An additional section, **%files**, allows to copy files or folders to the container.

We won't be using it here, but the format is very similar to `cp`, with sources being outside and the final destination being inside the container:

```
%files
some/file /some/other/file some/path/
some/directory some/path/
```

Note that this happens **after** `%post`. If you need the files earlier, copy them manually in `%setup`.

# Setting up the environment

You can specify a script to be sourced when something is run in the container.

This goes to the **%environment** section. Treat it like `.bash_profile`.

```
%environment
export HELLO=World
```

Note that by default, the host environment variables are passed to the container.

To disable it, use `-e` when running the container.

# Setting up the runscript

The runscript (`/singularity`) is specified in the `%runscript` section.

Let's use the file we copied at `%setup` and run `fortune`:

```
%runscript
read host < /etc/build_host
echo "Hello, $HELLO! Fortune Teller, built by $host"
exec /usr/bin/fortune "$@"
```

# Testing the built image

You can specify commands to be run at the end of the build process inside the container to perform sanity checks.

Use `%test` section for this:

```
%test
test -f /etc/build_host
test -x /usr/bin/fortune
```

All commands must return successfully or the build will fail.

# The whole definition file

```
Bootstrap: docker
From: centos:6

%setup
  hostname -f > $SINGULARITY_ROOTFS/etc/build_host
%post
  yum -y --enablerepo=extras install epel-release
  yum -y install fortune-mod
  yum clean all
%environment
  export HELLO="World"
%runscript
  read host < /etc/build_host
  echo "Hello, $HELLO! Fortune Teller, built by $host"
  exec /usr/bin/fortune "$@"
%test
  test -f /etc/build_host
  test -x /usr/bin/fortune
```

## *Exercise:*

Check that your fortune.def is the same as above.

# Building a container from definition

To fill a container using a definition file, we invoke `build`:

```
user@host:~$ rm fortune.sif
user@host:~$ singularity build --fakeroot fortune.sif fortune.def
[...]
```

## *Exercise:*

1. Bootstrap the image as shown above.
2. Test running it directly.

# Inspecting a built container

Container has some metadata you can read:

```
user@host:~$ singularity inspect fortune.sif
==labels==
org.label-schema.build-date: Tuesday_10_September_2019_11:1:10_CEST
org.label-schema.schema-version: 1.0
org.label-schema.usage.singularity.deffile.bootstrap: docker
[...]
```

You can inspect the original definition file:

```
user@host:~$ singularity inspect -d fortune.sif
Bootstrap: docker
From: centos:6
%setup
  hostname -f > $SINGULARITY_ROOTFS/etc/build_host
[...]
```

See `singularity help inspect` for more options, and `/.singularity.d/` inside the container to see how it's all stored.

# Runtime options

Fine-tuning container execution

# Host resources

A container can have more host resources exposed.

For providing access to more directories, one can specify bind options at runtime with `-B`:

```
$ singularity run -B source[:destination[:mode]] container.sif
```

where **source** is the path on the host, **destination** is the path in a container (if different) and **mode** is optionally `ro` if you don't want to give write access.

Of course, more than one bind can be specified.

Note that you can't specify this configuration in the container!

System administrators may specify binds that apply to all containers (e.g. `/scratch`).

# Host resources

Additionally, devices on the host can be exposed, e.g. the GPU; but you need to make sure that the guest has the appropriate drivers. One solution is to bind the drivers on the container.

For Nvidia CUDA applications specifically, Singularity supports the `--nv` flag, which looks for specific libraries on the host and binds them in the container.

---

OpenMPI should also work, provided the libraries on the host and in the container are sufficiently close.

If set up correctly, it should work normally with `mpirun`:

```
$ mpirun -np 20 singularity run mpi_job.sif
```

# Network

Historically, Singularity defaulted to no network isolation, with an option of full isolation.

With 3.x, Singularity implements in-between options through Container Network Interface:

<https://github.com/containernetworking/cni>

# Network

Historically, Singularity defaulted to no network isolation, with an option of full isolation.

With 3.x, Singularity implements in-between options through Container Network Interface:

<https://github.com/container networking/cni>

Port remapping example:

```
$ sudo singularity instance start --writable-tmpfs \  
  --net --network-args "portmap=8080:80/tcp" docker://nginx web2  
$ sudo singularity exec instance://web2 nginx  
$ curl localhost:8080  
[...]  
$ singularity instance stop web2
```

# Fuller isolation

By default, a container is allowed a lot of "windows" into the host system (dictated by Singularity configuration).

For an untrusted container, you can further restrict this with options like `--contain`, `--containll`.

In this case, you have to manually define where standard binds like the home folder or `/tmp` point.

See `singularity help run` for more information.

# Distributing the container

Using the container after creation on another Linux machine is simple: you simply copy the image file there.

Note that you can't just run the image file on a host without Singularity installed!

## *Exercise:*

Test the above, by trying to run `fortune.sif` inside itself.

# Distributing the container

Using the container after creation on another Linux machine is simple: you simply copy the image file there.

Note that you can't just run the image file on a host without Singularity installed!

## *Exercise:*

Test the above, by trying to run `fortune.sif` inside itself.

This approach makes it easy to deploy images on clusters with shared network storage.

You can easily integrate Singularity with the usual scheduler scripts (e.g. Slurm).

# Cloud services

Current and upcoming ecosystem

# Using Singularity Hub

Singularity Hub allows you to cloud-build your containers from Bootstrap files, which you can then simply pull on a target host.

<https://singularity-hub.org/>

# Using Singularity Hub

Singularity Hub allows you to cloud-build your containers from Bootstrap files, which you can then simply pull on a target host.

<https://singularity-hub.org/>

This requires a GitHub repository with a Singularity definition file. After creating an account and connecting to the GitHub account, you can select a repository and branches to be built.

Afterwards, you can pull the result:

```
user@host:~$ singularity pull shub://kav2k/fortune
[...]
user@host:~$ ./fortune_latest.sif
Hello, World! Fortune Teller, built by shub-builder-1450-kav2k-fortune-[...]
```

# Singularity Hub quirks

- Singularity Hub is not like Docker Hub, or similar registry. You can't "push" an image there, it can only be built on their side.
- Singularity Hub is not an official Sylabs project, it's an academic non-profit project by other developers.
- Singularity Hub runs a modified version of Singularity 2.4, making some newer build-time features unavailable (but not runtime features).
- There are no paid plans. Users are allowed a single private project.

# Sylabs cloud offering

Starting with Singularity 3.0, the company behind Singularity aims to provide a range of cloud services to improve Singularity user experience.

- **Container Library** as a counterpart for Docker Hub, serving as an official image repository.
- **Remote Builder** service to allow unprivileged users to build containers in the cloud.
- **KeyStore** service to enable container signature verification.

# Sylabs Container Library

Container Library is the Singularity counterpart to Docker Hub: a cloud registry for both public and private containers.

<https://cloud.sylabs.io/library>

The Library allows direct upload of pre-built (and signed) containers, unlike Singularity Hub.

```
$ singularity push my.sif library://user/collection/my.sif:latest  
$ singularity pull library://user/collection/my.sif:latest
```

Use of Library is subject to quotas.

# Sylabs Remote Builder

Building a container from a recipe requires `sudo`, imposing a need for a separate container creation infrastructure.

## *Note:*

--fakeroot alleviates some of those problems

Sylabs provides a remote builder service that can build an image from a recipe file, then temporarily host it in Cloud Library to be downloaded.

```
user@host:~$ singularity build --remote output.sif fortune.def
searching for available build agent.....INFO: Starting build...
[...]
user@host:~$ ./output.sif
Hello, World! Fortune Teller, built by ip-10-10-30-146.ec2.internal
[..yet again, a funny quote..]
```

Caveat: all resources for a remote build must be accessible by the build node (i.e. over internet).

# Signing containers and Sylabs Keystore

To ensure safety of containers, SIF format allows them to be cryptographically signed.

```
user@host:~$ singularity sign output.sif
user@host:~$ singularity verify output.sif
```

This alone provides assurance of integrity (has not been modified).

For authentication, Sylabs provides a **keyserver** called Keystore, which can be used to check signatures of keys not locally available.

```
user@host:~$ singularity keys push <fingerprint>
user@host2:~$ singularity verify output.sif
```

# Sylabs commercial offering

Both the Container Library and Remote Builder are currently in free testing period. However, in future they will have a freemium model.

There will also be on-premise versions of both services (which are not open source).

Besides that, Sylabs offers Singularity PRO: a priority-supported version of Singularity with ready-built packages.

Pricing is "upon request", and is either based on number of hosts or is site-wide.

# Running on UBELIX

From a practical standpoint, we want to use the container technology on UBELIX.

Let's try with our toy container:

```
user@host:~$ ssh username@submit.unibe.ch
username@submit01:~$ singularity pull library://kav2k/default/fortune:latest
username@submit01:~$ sbatch -J fortune-test -t 00:00:10 \
--mem-per-cpu 100M --cpus-per-task 1 --wrap "./fortune_latest.sif"
```

# Docker and Singularity

Instead of writing a Singularity file, you may write a Dockerfile, build a Docker container and convert that.

Pros:

- More portable: for some, using Docker or some other container solution is preferable.
- Easier private hosting: there is no mature private registry tech for Singularity.

Cons:

- Blackbox: Singularity understands less about the build process, in terms of container metadata.
- Complexity: Extra tool to learn if you don't know Docker.

Advice on Docker compatibility: [Best Practices](#)

# Docker -> Singularity

If you have a Docker image you want to convert to Singularity, you have at least 4 options:

1. Upload the image to a Docker Registry (such as Docker Hub) and pull/Bootstrap from there.
2. Use a private Docker registry to not rely on external services
3. Directly pull from a local Docker daemon cache
4. Use intermediate format as generated by `docker save`

# "Extra credit" topics

# Reducing container size

Using traditional Linux distributions, even in minimal configurations, can still be an overkill for running a single application.

One can reduce container size by clearing various artifacts of the build process, such as package manager caches.

Alternatively, one can use minimal Linux distributions, such as Alpine Linux, as a base for containers, though compatibility needs extra testing.

```
$ ll -h
-rwxr-xr-x 1 user group 66M Jun 25 15:04 centos_6.sif*
-rwxr-xr-x 1 user group 2.0M Jun 25 16:08 alpine.sif*
```

# Singularity Instances

Running daemon-like persistent services with Singularity (such as a web server) can conveniently be done with the concept of Instances.

A `%startscript` section of the recipe describes what service to launch, which subsequently works with `instance` commands:

```
$ singularity instance start nginx.sif web
$ singularity instance list
INSTANCE NAME      PID      CONTAINER IMAGE
web                790      /home/user/nginx.sif
$ singularity instance stop web
```

While an instance is running, the standard commands like `shell` and `exec` work with an `instance:// namespace`.

# SCI-F

One of the approaches for building scientific pipelines is bundling several tools in a single "toolset" container.

SCI-F is a proposed standard for discovering and managing tools within such modular containers.

Definition file can have several sections, e.g.:

```
%appenv foo
  BEST_GUY=foo
  export BEST_GUY

%appenv bar
  BEST_GUY=bar
  export BEST_GUY

%apprun foo
  echo The best guy is $BEST_GUY

%apprun bar
  echo The best guy is $BEST_GUY
```

# SCI-F

You can then discover the apps bundled and run them:

```
$ singularity apps foobar.sif
bar
foo
$ singularity run --app bar foobar.sif
The best guy is bar
```

More sections can be made app-specific, including providing a help description:

```
$ singularity help --app fortune moo.sif
fortune is the best app
```

# Singularity Checks

A container check is a utility script that can verify a container.

Example uses:

- Making sure no leftover artifacts from the build process remains (e.g. root's bash history)
- Testing for common vulnerabilities
- Custom checks for your specific environment

```
$ singularity check --tag clean ubuntu.img
```

# Reproducibility going forward

Pinning a specific version of a base image makes it more probable that in future building the same recipe will be impossible.

Singularity allows for easy storage of resulting containers, and is good at providing backwards compatibility. This provides archival capability (but containers can be large).

# Reproducibility going forward

Pinning a specific version of a base image makes it more probable that in future building the same recipe will be impossible.

Singularity allows for easy storage of resulting containers, and is good at providing backwards compatibility. This provides archival capability (but containers can be large).

But a "frozen" container can get other compatibility problems down the line, especially if it needs some host-container interaction.

For example, compiled software in it is no longer optimized for newer hardware architectures.

# Reproducibility going forward

Pinning a specific version of a base image makes it more probable that in future building the same recipe will be impossible.

Singularity allows for easy storage of resulting containers, and is good at providing backwards compatibility. This provides archival capability (but containers can be large).

But a "frozen" container can get other compatibility problems down the line, especially if it needs some host-container interaction.

For example, compiled software in it is no longer optimized for newer hardware architectures.

Bottom line: containers are not a silver bullet to solve reproducibility problems, but they help.

# Further reading

- Singularity User Guide: <https://www.sylabs.io/guides/3.2/user-guide/>
- Singularity Admin Guide: <https://www.sylabs.io/guides/3.2/admin-guide/>
- Singularity White Paper: [link](#)
- *Extra credit:* <https://rootlesscontaine.rs/>

# Further resources

- **Play with Docker:** a [Docker playground](#) to experiment with Docker, including clustering.
- **Play with Kubernetes:** a [Kubernetes playground](#) for hands-on experimentation with orchestration.

Questions?