Open Science Grid

Containers and GPUs

Christina Koch (ckoch5@wisc.edu)

Research Computing Facilitator

University of Wisconsin - Madison

GPUS

What is a GPU?

- GPU = Graphical Processing Unit
- Has hundreds to thousands of “cores” that can be used to parallelize work.

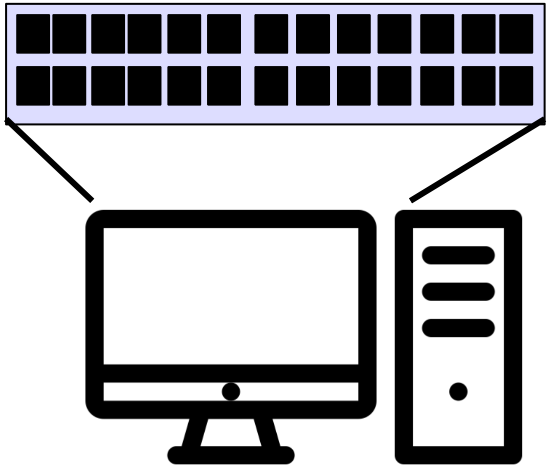

Created by Idealogo Studio
from Noun Project

GPU Use Cases

- Programs that map well to GPUs include:
  - Deep learning
  - Molecular dynamics
  - Anything with lots of number crunching (like matrix operations) and low(er) data load.

GPUs on the OSG

- Scale: 10s (vs 100s-1000s of CPUs)
- Variety of available GPU cards
- Same restrictions as always: not sure what you’ll get, jobs can be interrupted
- May take longer to start
- Free!

Making robust GPU jobs

- Use a software strategy that can run on different GPU types:
  - Container
  - Install inside the job
- OR use job requirements to request certain kind of GPU (more limiting)

Submit File options

- Request GPUs with “request_gpus”
- Can use custom requirements

```
request_gpus = 1

requirements = (CUDACapability >= 6.0)
```

CONTAINERS

Containers

- Containers are a tool for capturing an entire job “environment” (software, libraries, operating system) into an “image” that can be used again.

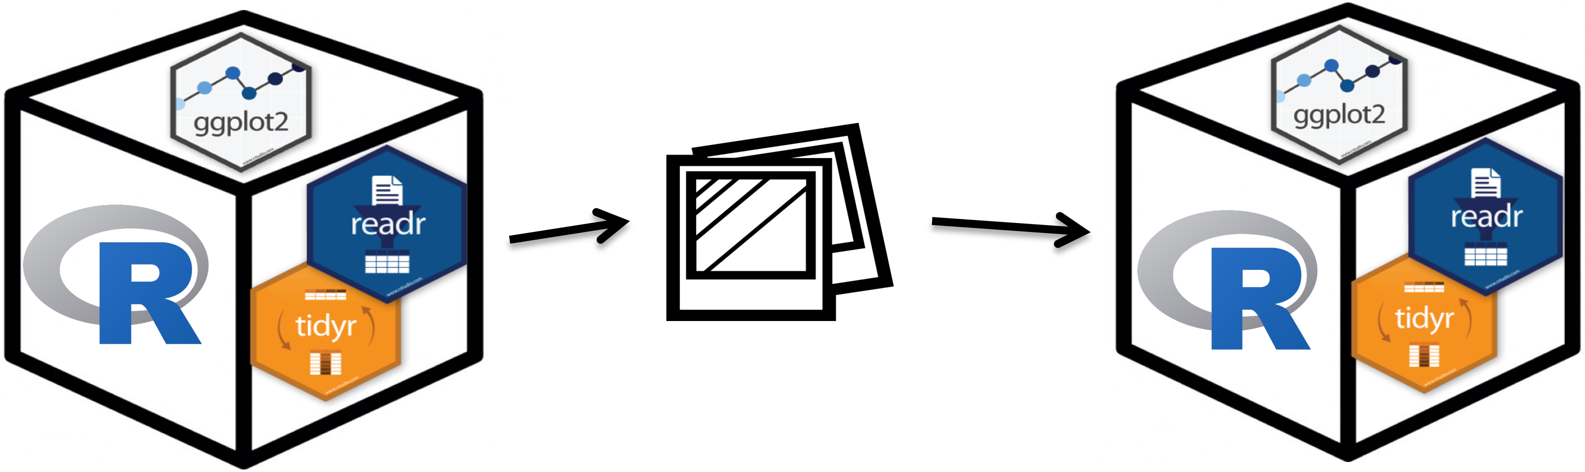

polaroid photos by Nick Bluth from the Noun Project

Container Types

- Two common container systems:

Docker

https://www.docker.com/

Singularity

https://sylabs.io/

The container itself will always be some version of Linux - but can be run on Linux / Mac / Windows if Docker or Singularity is installed

Focus on Docker

- Docker has well-established and well-documented ways to build container images.
- If you have a Docker image:
  - Can run with Docker
  - Can run with Singularity
  - Can convert to a Singularity image

Running Containers

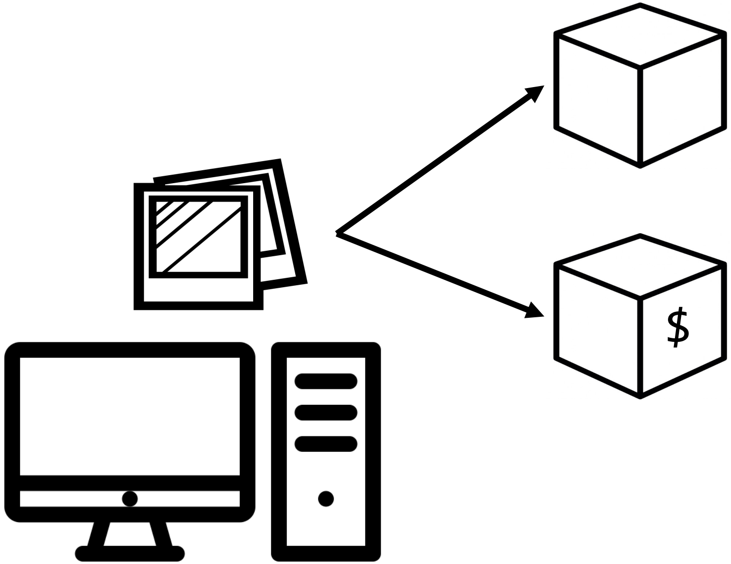

```
docker run <container> <command>
```

```
docker run -it <container> /bin/sh
```

Created by Idealogo Studio
from Noun Project

Docker Hub

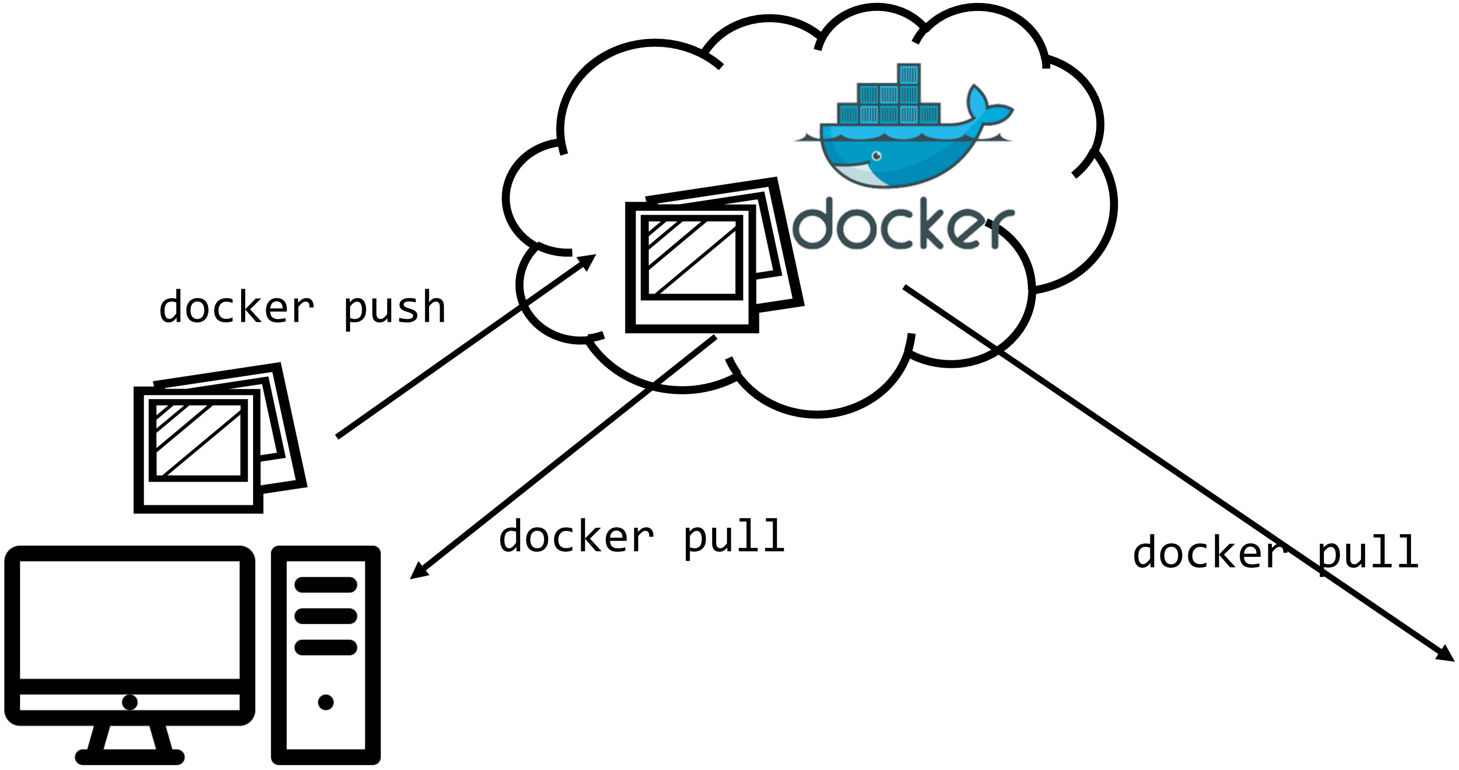

Created by Idealogo Studio
from Noun Project

Building Containers

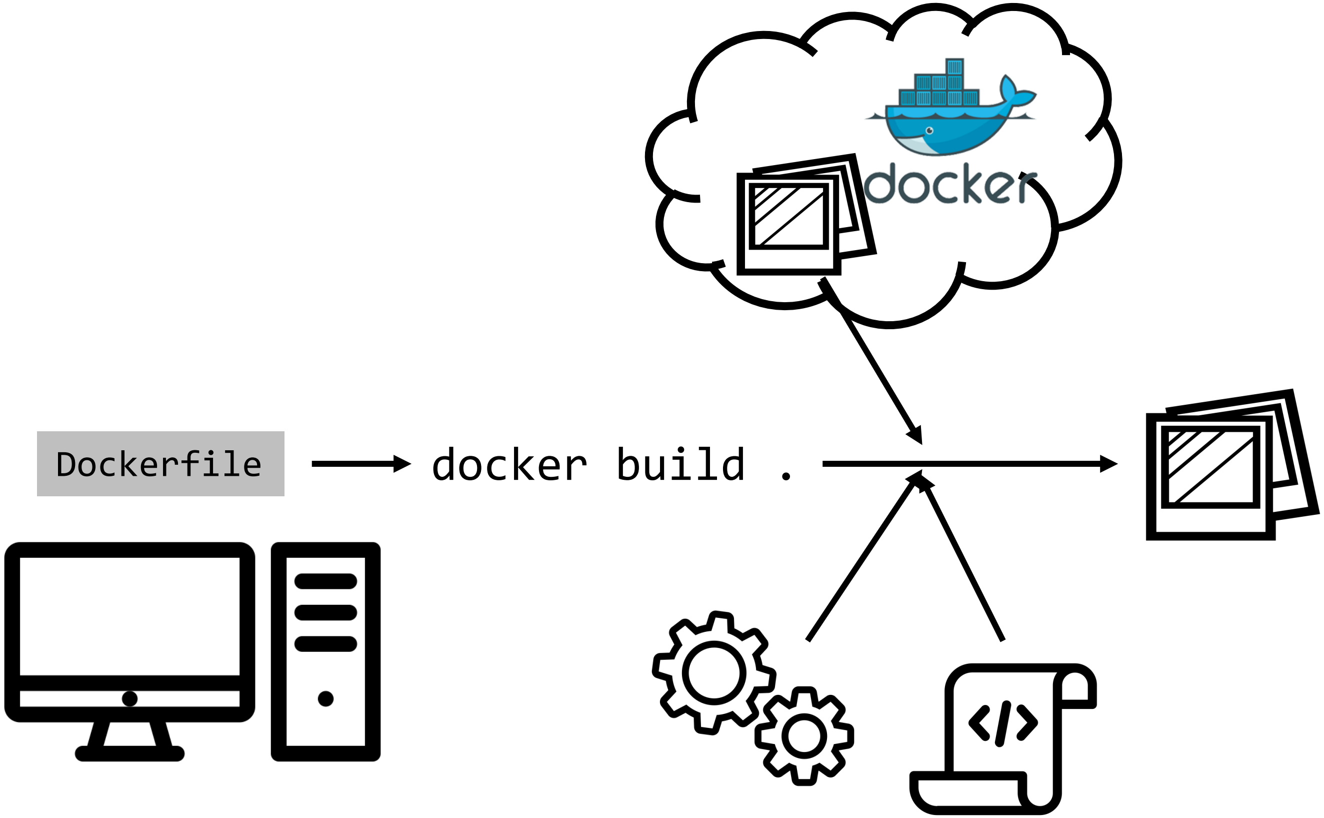

Created by Idealogo Studio
from Noun Project

Sample Dockerfile

```
# Start with this image as a "base".
# It's as if all the commands that created that image were inserted here.
FROM continuumio/miniconda:4.7.12

# Use RUN to execute commands inside the image as it is being built up.
RUN conda install --yes numpy

# Try to always "clean up" after yourself to reduce the final size of your image.
RUN apt-get update \
&& apt-get --yes install --no-install-recommends graphviz\
&& apt-get --yes clean \
&& rm -rf /var/lib/apt/lists/*
```