[639] Docker

Meenakshi Syamkumar

Learning Objectives

Use existing docker images	to launch containers
Define new docker images	using Dockerfiles
Troubleshoot	common issues with running Docker containers

Virtualization

- Definition: the illusion of private resources, provided by software
- Our focus:
  - Virtual machines (hardware)
  - Virtual Operating System (container)

virtualized resources include CPU, RAM, disks, network devices, etc

VMs rarely use all their allocated resources, so overbooking is possible

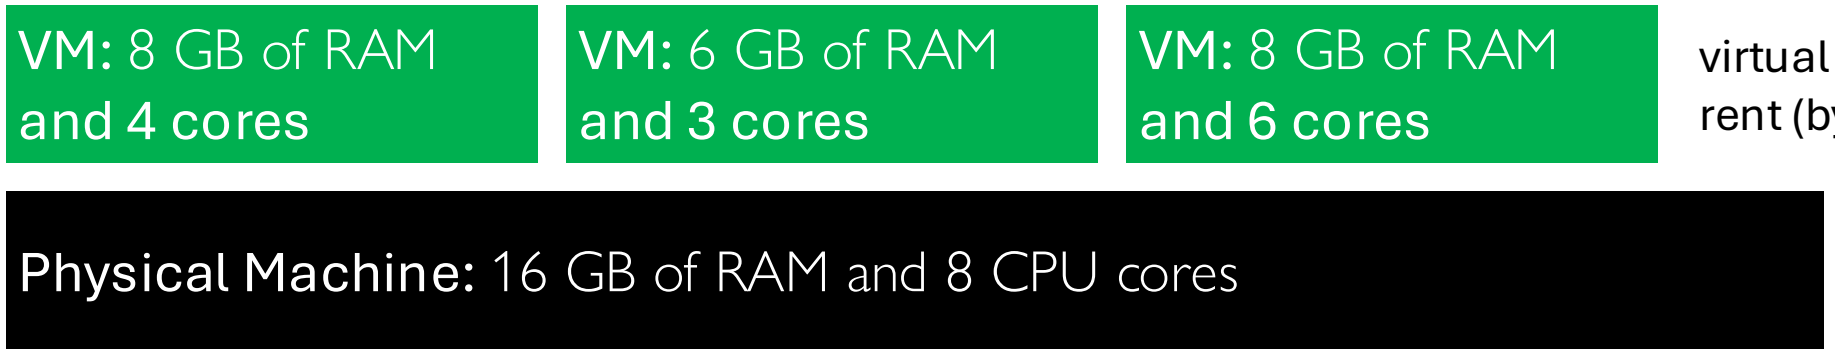

virtual machines for rent (by you)

actual hardware bought by cloud provider (like Google GCP) for their cloud services

Virtualization

- Definition: the illusion of private resources, provided by software
- Our focus:
  - Virtual machines (hardware)
  - Virtual Operating System (container)

these operating systems are mostly unaware that they run on VMs instead of physical hardware

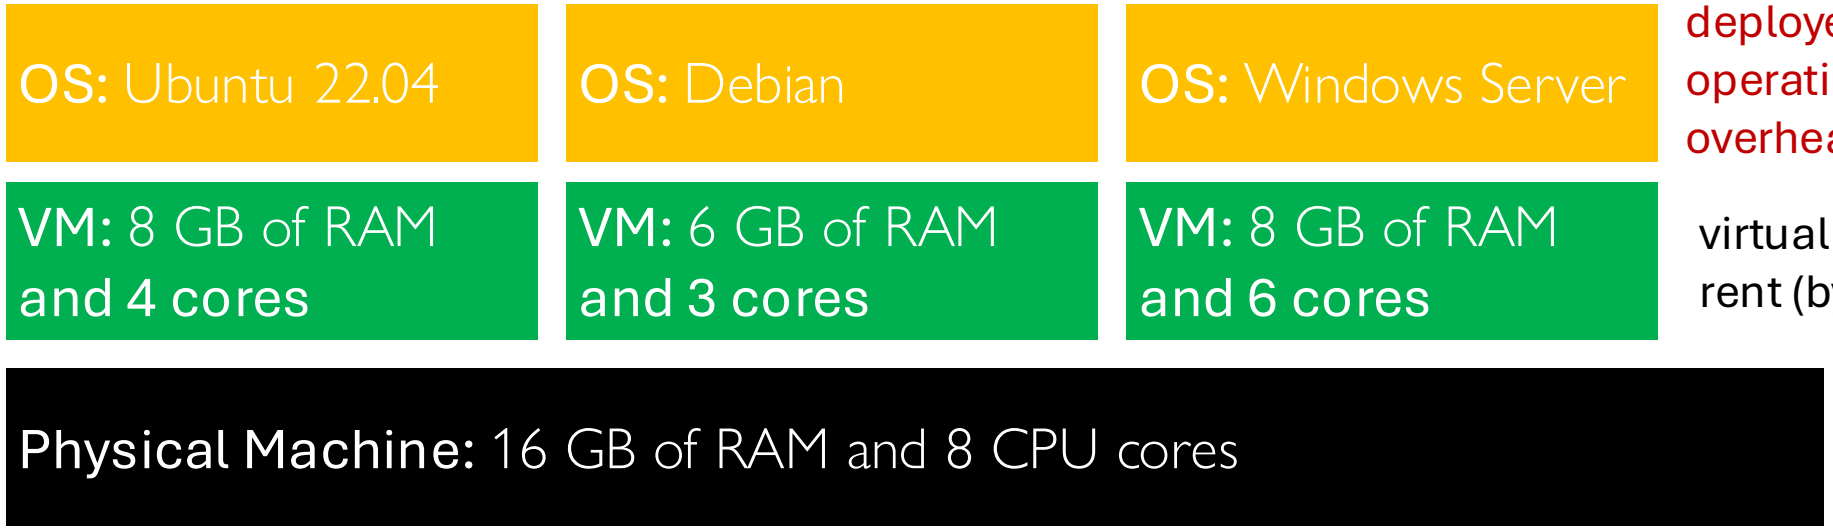

problem: if each program is deployed to a different VM, operating system overheads will dominate

virtual machines for rent (by you)

actual hardware bought by cloud provider (like Google GCP) for their cloud services

Virtualization

- Definition: the illusion of private resources, provided by software
- Our focus:
  - Virtual machines (hardware)
  - Virtual Operating System (container)

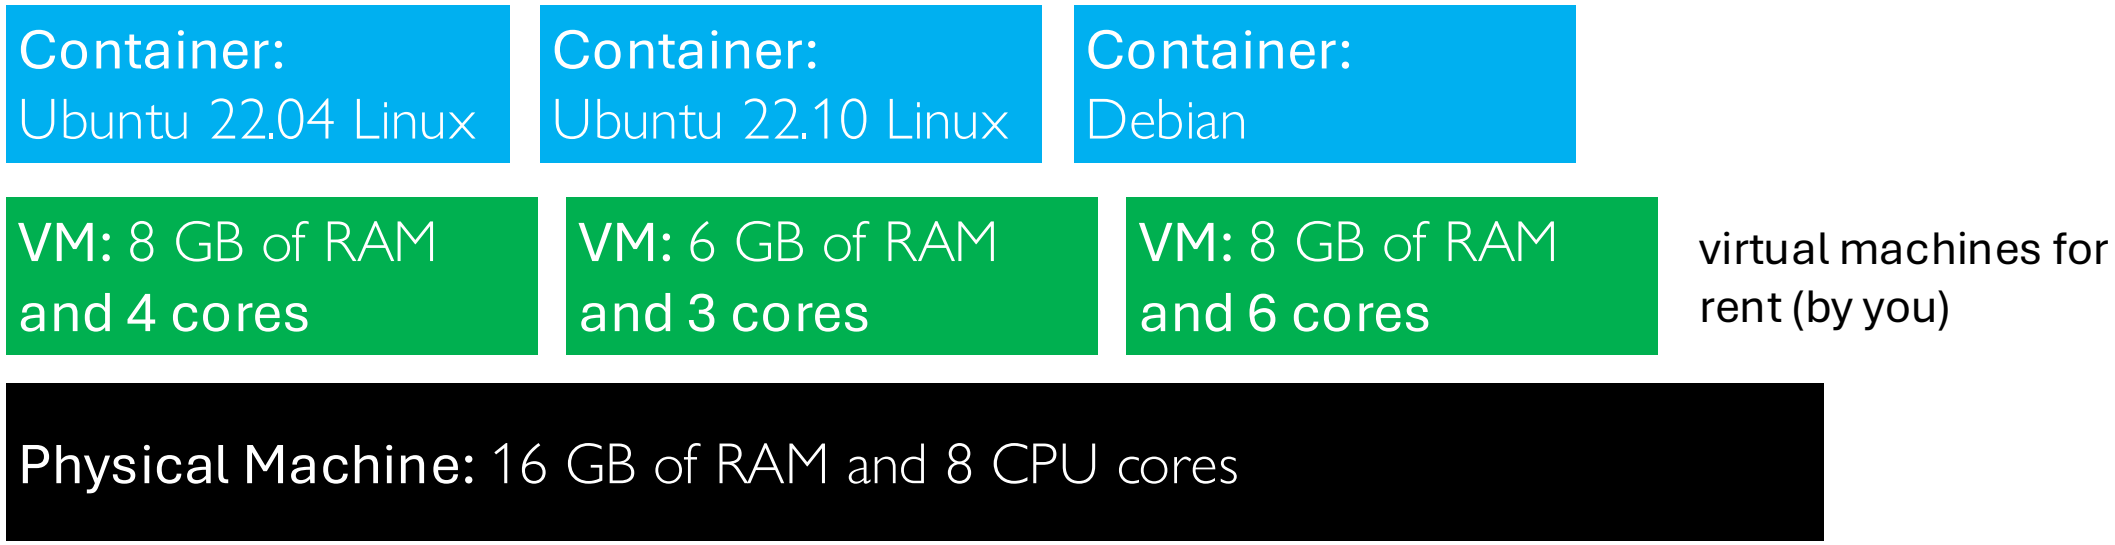

Containers

Linux containers are a lightweight alternative to virtual machines

- Docker makes creation easy
- The "physical" OS is shared, which is very efficient
- Programs in different containers can uses different flavors of Linux
- Cannot have a Windows container on Linux

You'll run Docker containers this semester to create a new development environment for each project / topic

Sandboxes

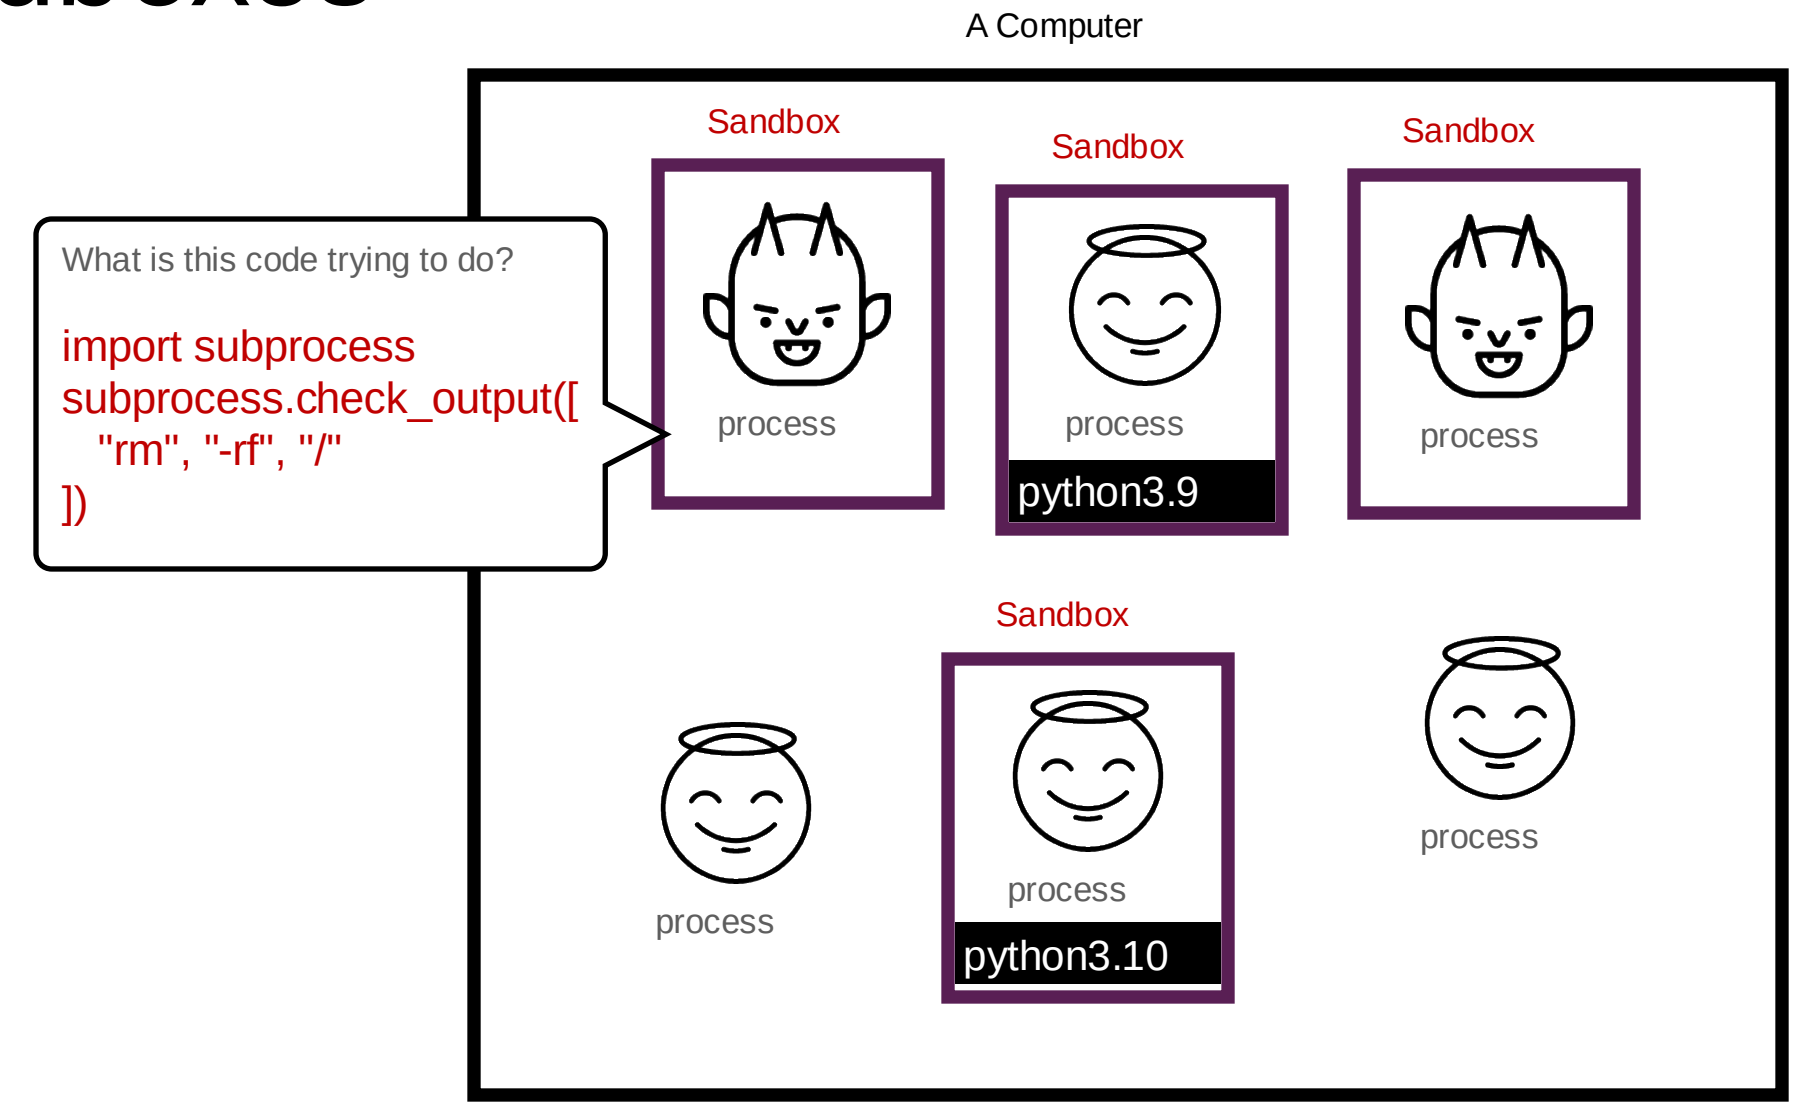
A Computer
Sandbox
process
Sandbox
process
python3.9
Sandbox
process
What is this code trying to do?
import subprocess
subprocess.check_output([
"rm", "-rf", "/"
])
process
Sandbox
process
python3.10
process

Registries, Images, Containers, and Dockerfiles

Today's lecture worksheet – for taking notes

docker **SOME-COMMAND** arg1, arg2, …

Registries, Images, Containers, and Dockerfiles

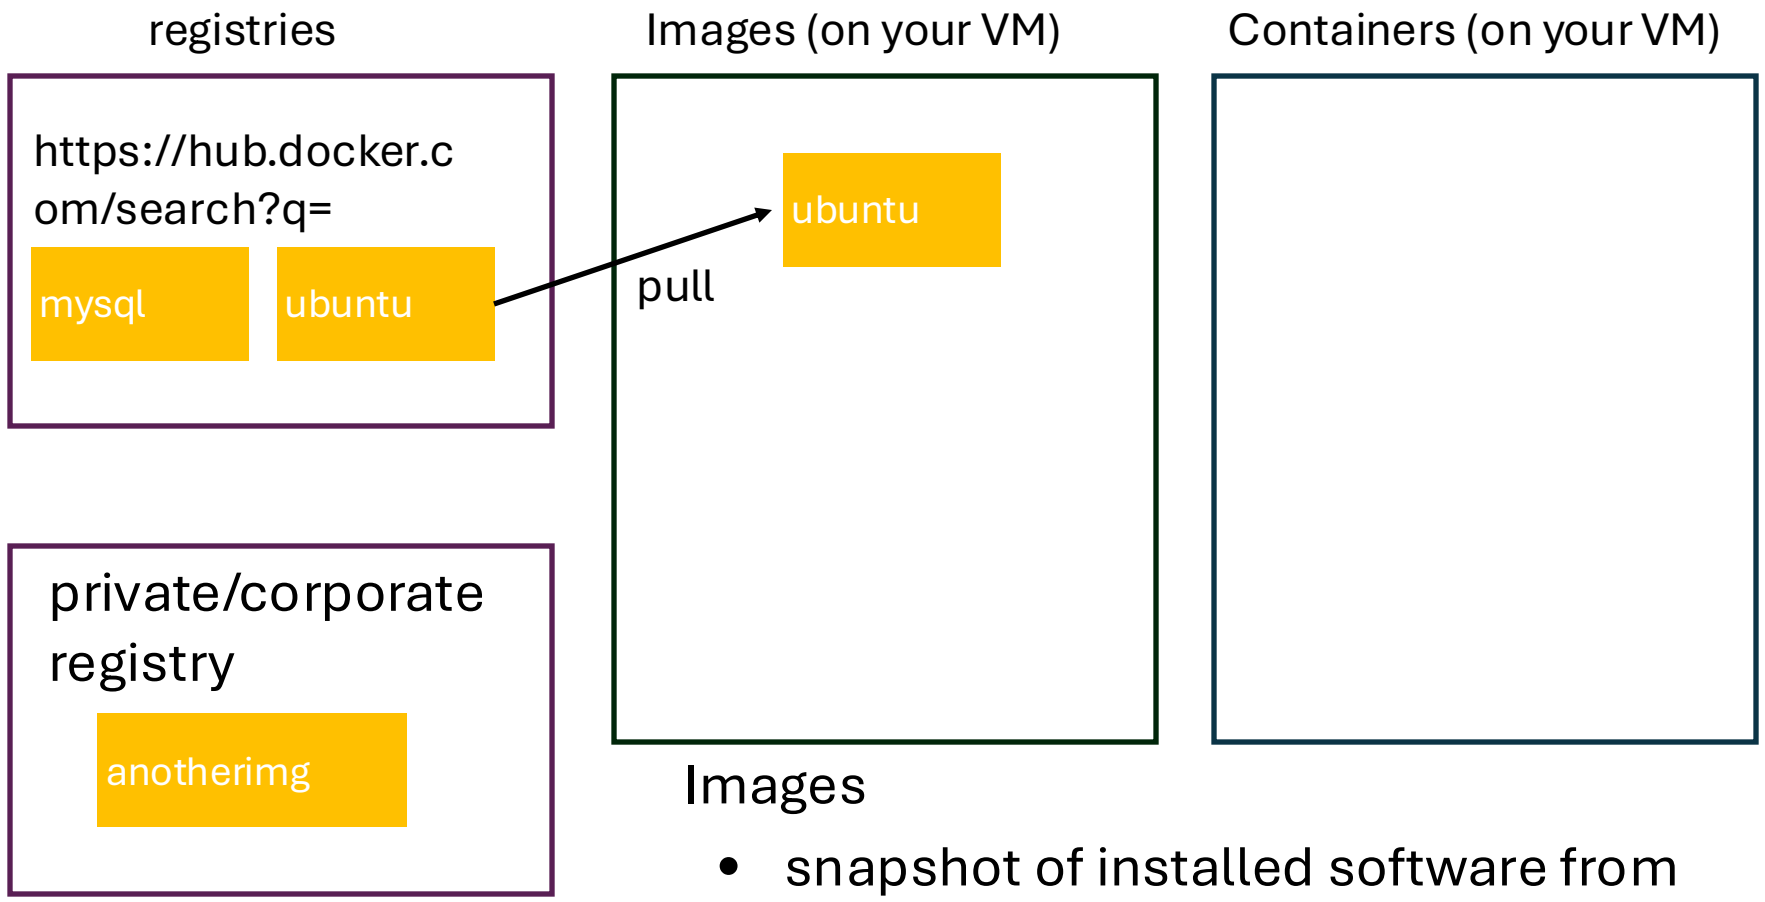

Images

- snapshot of installed software from which to create a container
- docker **pull** ubuntu:22.04

Registries, Images, Containers, and Dockerfiles

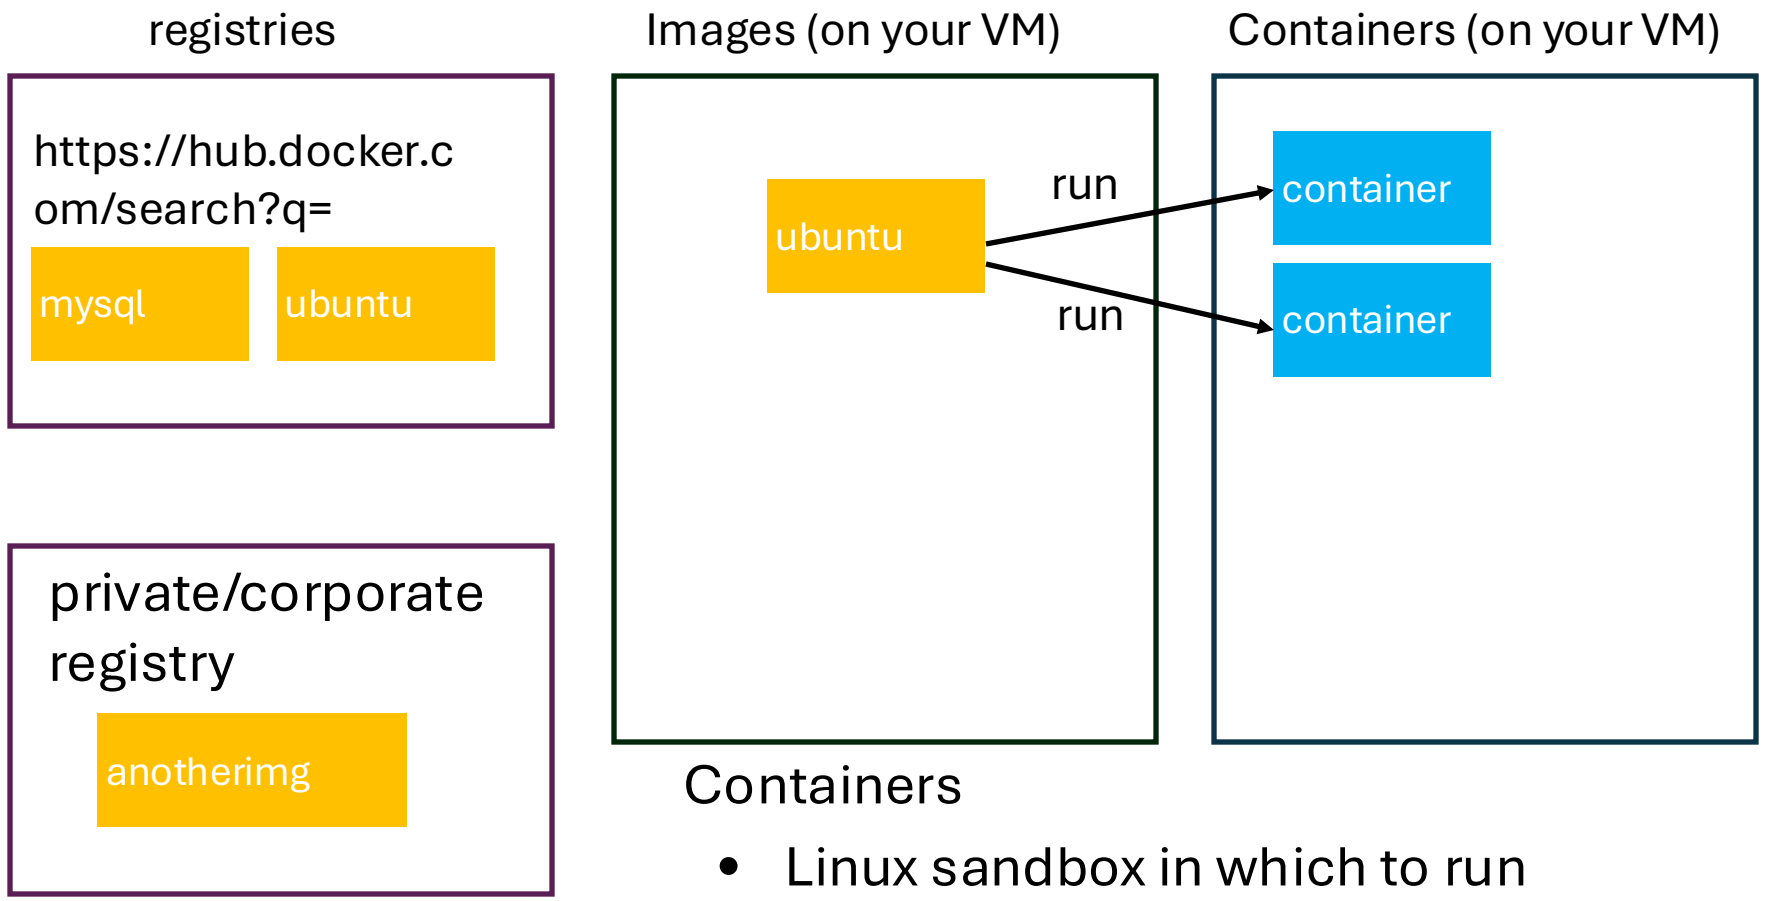

Containers

- Linux sandbox in which to run processes
- docker **run** ubuntu

Registries, Images, Containers, and Dockerfiles

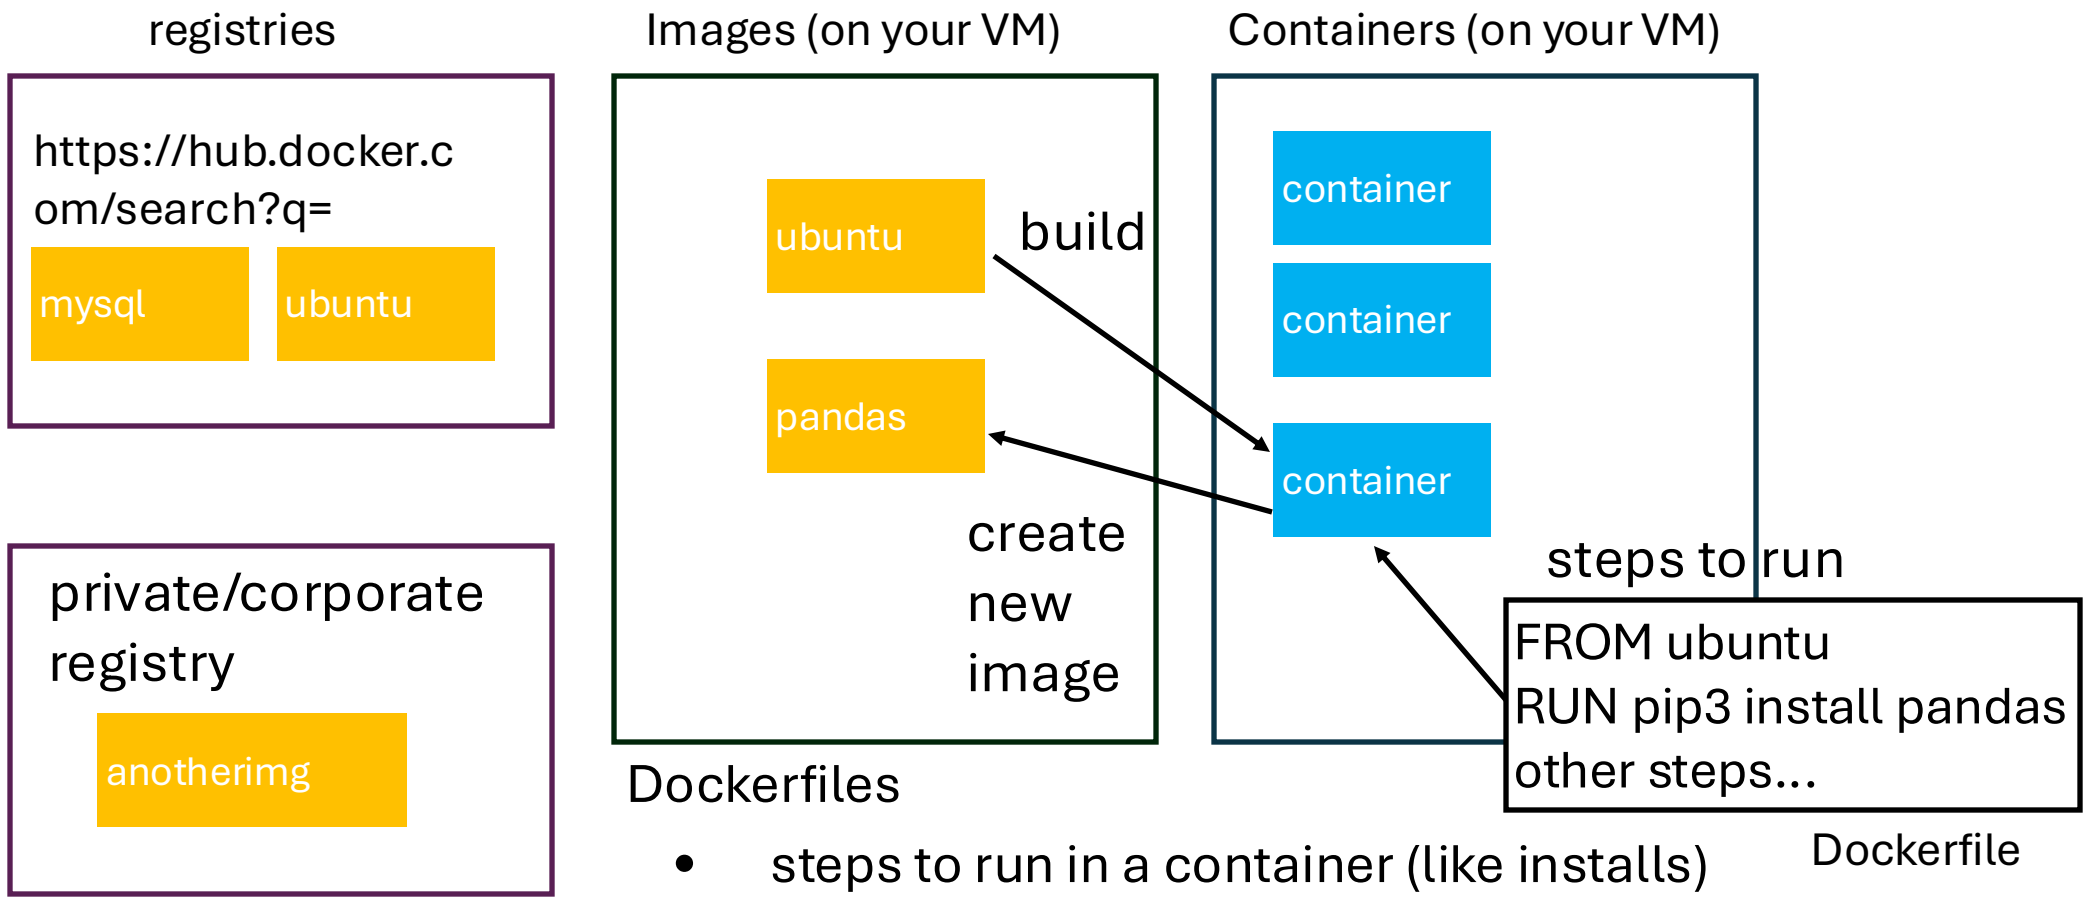

- steps to run in a container (like installs)
- creates a new image
- docker **build** myimg -t pandas

Registries, Images, Containers, and Dockerfiles

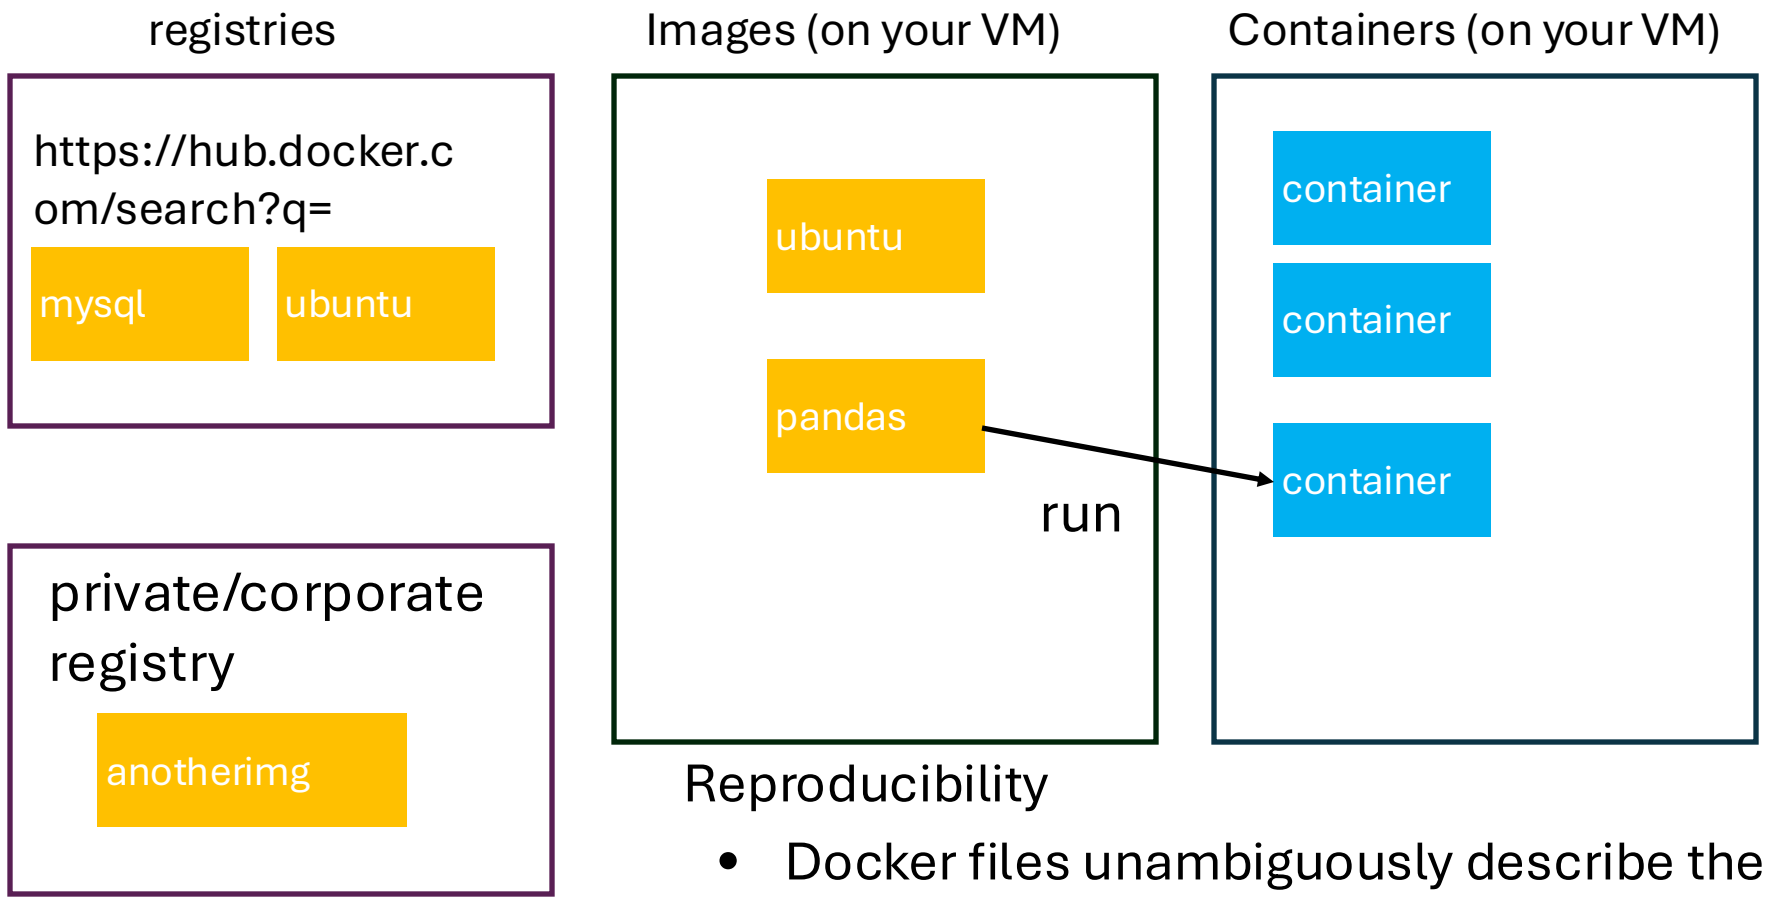

Reproducibility

- Docker files unambiguously describe the setup
- Others can get all the same version numbers

Docker cleanup

Make sure to cleanup docker containers once you stop using them

Containers & images cleanup

Clean up is very important

- docker **system prune**
- docker **kill <ID_or_name>**
- docker **rm <ID_or_name>**
- docker **rmi <Image>**
- docker **rm `docker ps –aq`**

If you don’t cleanup properly, you will run out of disk space on your VM!