

Magic Tools To Install and Manage Software



Siva Prasad Kasetti

HPC User Services

LSU HPC / LONI

sys-help@loni.org

Louisiana State University
Baton Rouge

March 18, 2026

Magic Tools To Install and Manage Software



CONDA Virtual Environment

1. Why Container?

- 1) Problems
- 2) Container & Singularity

2. Run an Existing Container Image

- 1) What you need
- 2) Basic commands
- 3) Running jobs with Singularity

3. Get More Container Images

- 1) What you need
- 2) Where to get
- 3) How to get

4. Build Your Own Container Image

- 1) What you need
- 2) Typical workflow
- 3) Make it easier - Recipe

1. Why Container?

- 1) Problems
- 2) Container & Singularity

2. Run an Existing Container Image

- 1) What you need
- 2) Basic commands
- 3) Running jobs with Singularity

3. Get More Container Images

- 1) What you need
- 2) Where to get
- 3) How to get

4. Build Your Own Container Image

- 1) What you need
- 2) Typical workflow
- 3) Make it easier - Recipe

1. Why Container?

- 1) Problems
- 2) Container & Singularity

2. Run an Existing Container Image

- 1) What you need
- 2) Basic commands
- 3) Running jobs with Singularity

3. Get More Container Images

- 1) What you need
- 2) Where to get
- 3) How to get

4. Build Your Own Container Image

- 1) What you need
- 2) Typical workflow
- 3) Make it easier - Recipe

- Core problem:

Installing software on an HPC system

- **Traditional Linux solution:**
 - Compiling from source code

1

Dependencies

Most software depends on other libraries or tools. If those are missing, the software may not install or run.

1. Dependencies (Welcome to Linux!)

Third-party components

A full installation of BUSCO requires *Python 3.3+* (2.7 is not supported from v4 onwards), *BioPython*, *pandas*, *BBMap*, *tBLASTn 2.2+*, *Augustus 3.2+*, *Prodigal*, *Metaeuk*, *HMMER3.1+*, *SEPP*, and *R + ggplot2* for the plotting companion script. Some of these tools are necessary only for analysing certain type of organisms and input data, or for specific run modes.

- <https://biopython.org/>
- <https://pandas.pydata.org/>
- <https://jgi.doe.gov/data-and-tools/software-tools/bbtools/>
- <https://ftp.ncbi.nlm.nih.gov/blast/executables/blast+/LATEST>
- <http://bioinf.uni-greifswald.de/augustus/>
- <https://github.com/soedinglab/metaeuk>
- <https://github.com/hyattpd/Prodigal>
- <http://hmmer.org/>
- <https://github.com/smirarab/sepp/>
- <https://www.r-project.org/>

Please make sure that each software package listed above works INDEPENDENTLY of BUSCO before attempting to run any BUSCO assessments.

1. Dependencies (Welcome to Linux!)

Third-party components

A full installation of BUSCO requires *Python 3.3+* (2.7 is not supported from v4 onwards), *BioPython*, *pandas*, *BBMap*, *tBLASTn 2.2+*, *Augustus 3.2+*, *Prodigal*, *Metaeuk*, *HMMER3.1+*, *SEPP*, and *R + ggplot2* for the plotting companion script. Some of these tools are necessary only for analysing certain type of organisms and input data, or for specific run modes.

- <https://biopython.org/>
- <https://pandas.pydata.org/>
- <https://jgi.doe.gov/data-and-tools/software-tools/bbtools/>
- <https://ftp.ncbi.nlm.nih.gov/blast/executables/blast+/LATEST>
- <http://bioinf.uni-greifswald.de/augustus/>
- <https://github.com/soedinglab/metaeuk>
- <https://github.com/hyattpd/Prodigal>
- <http://hmmer.org/>
- <https://github.com/smirarab/sepp/>
- <https://www.r-project.org/>

Please make sure that each software package listed above works INDEPENDENTLY of BUSCO before attempting to run any BUSCO assessments.

1. Dependencies (Welcome to Linux!)

Third-party components

A full installation of BUSCO requires *Python 3.3+* (2.7 is not supported from v4 onwards), *BioPython*, *pandas*, *BBMap*, *tBLASTn 2.2+*, *Augustus 3.2+*, *Prodigal*, *Metaeuk*, *HMMER3.1+*, *SEPP*, and *R + ggplot2* for the plotting companion script. Some of these tools are necessary only for analysing certain type of organisms and input data, or for specific run modes.

- <https://biopython.org/>
- <https://pandas.pydata.org/>
- <https://jgi.doe.gov/data-and-tools/software-tools/bbtools/>
- <https://ftp.ncbi.nlm.nih.gov/blast/executables/blast+/LATEST>
- <http://bioinf.uni-greifswald.de/augustus/>
- <https://github.com/soedinglab/metaeuk>
- <https://github.com/hyattpd/Prodigal>
- <http://hmmer.org/>
- <https://github.com/smirarab/sepp/>
- <https://www.r-project.org/>

Please make sure that each software package listed above works INDEPENDENTLY of BUSCO before attempting to run any BUSCO assessments.



```
• Dependencies
The following dependencies are required for AUGUSTUS:
  ◦ for gzip compressed input: (set ZIPINPUT = false in common.mk if this feature is not available)
    ▪ libboost-iostreams-dev
    ▪ zlib1g-dev
  ◦ for comparative AUGUSTUS (multi-species, CGP): (set COMPGENEPRED = false in common.mk if this feature is not available. Augustus can then only be run in single-genome mode, which is what most users need.)
    ▪ libgsl-dev
    ▪ libboost-all-dev
    ▪ libsuitesparse-dev
    ▪ liblsolve55-dev
    ▪ libsqlite3-dev (add SQLITE = false to common.mk if this feature is not required or the required library is not available)
    ▪ libmysql++-dev (add MYSQL = false to common.mk if this feature is not required or the required library is not available)
  ◦ for compiling utilities bam2hints and filterBam:
    ▪ libbamtools-dev zlib1g-dev
  ◦ for compiling utility utrnanseq:
    ▪ libboost-all-dev (version must be > Boost_1_49_0)
  ◦ for compiling utility bam2wig:
    ▪ Follow these instructions. Note that it shouldn't be a problem to compile AUGUSTUS without bam2wig. In practice, you can simply use bamToWig.py to accomplish the same task.
  ◦ For compiling homgenemapping (set BOOST = FALSE in auxprogs/homgenemapping/src/Makefile if the option --printHomologs is not required or the required libraries are not available)
    ▪ libboost-all-dev
  ◦ for scripts:
    ▪ Perl
    ▪ Python3
  ◦ for the python3 script bamToWig.py:
    ▪ twoBitInfo and faToTwoBit from http://hgdownload.soe.ucsc.edu/admin/execute . bamToWig.py will automatically download these tools to the working directory during execution if they are not in your $PATH.
    ▪ SAMtools (available e.g. via package managers or here - see notes below)
```

1

Dependencies

Most software depends on other libraries or tools. If those are missing, the software may not install or run.

2

Version Conflicts

Different packages may need different versions of the same dependency, causing conflicts in a shared environment.

2) Conflicted packages

- Many scientific and ML applications require specific versions of Python, CUDA and libraries like TensorFlow, or PyTorch.
- For example:
 - TensorFlow 2.10 supports Python 3.7–3.10, not Python 3.11,
 - Combine with newer software in a working environment creates version conflicts.

1

Dependencies

Most software depends on other libraries or tools. If those are missing, the software may not install or run.

2

Version Conflicts

Different packages may need different versions of the same dependency, causing conflicts in a shared environment.

3

root permission

Installing software often requires admin access, which most HPC users do not have.



3) Permission denied (Welcome to HPC!)

```
[kasetti@mike003 E3RC]$ module load python  
[kasetti@mike003 E3RC]$ pip install gdal
```

3) Permission denied (Welcome to HPC!)

```
[kasetti@mike003 E3RC]$ module load python  
[kasetti@mike003 E3RC]$ pip install gdal
```

```
writing dependency_links to gdal-utils/GDAL.egg-info/dependency_links.txt  
writing entry points to gdal-utils/GDAL.egg-info/entry_points.txt  
writing requirements to gdal-utils/GDAL.egg-info/requires.txt  
writing top-level names to gdal-utils/GDAL.egg-info/top_level.txt  
Traceback (most recent call last):  
  File "<string>", line 92, in fetch_config  
  File "/usr/local/packages/python/3.9.7-anaconda/lib/python3.9/subprocess.py", line 951, in __init__  
    self._execute_child(args, executable, preexec_fn, close_fds,  
  File "/usr/local/packages/python/3.9.7-anaconda/lib/python3.9/subprocess.py", line 1821, in _execute_chi  
ld  
    raise child_exception_type(errno_num, err_msg, err_filename)  
FileNotFoundError: [Errno 2] No such file or directory: 'gdal-config'
```

3) Permission denied (Welcome to HPC!)

```
[kasetti@mike003 E3RC]$ module load python  
[kasetti@mike003 E3RC]$ pip install gdal
```

```
writing dependency_links to gdal-utils/GDAL.egg-info/dependency_links.txt  
writing entry points to gdal-utils/GDAL.egg-info/entry_points.txt  
writing requirements to gdal-utils/GDAL.egg-info/requires.txt  
writing top-level names to gdal-utils/GDAL.egg-info/top_level.txt  
Traceback (most recent call last):  
  File "<string>", line 92, in fetch_config  
  File "/usr/local/packages/python/3.9.7-anaconda/lib/python3.9/subprocess.py", line 951, in __init__  
    self._execute_child(args, executable, preexec_fn, close_fds,  
  File "/usr/local/packages/python/3.9.7-anaconda/lib/python3.9/subprocess.py", line 1821, in _execute_chi  
ld  
  raise child_exception_type(errno_num, err_msg, err_filename)  
FileNotFoundError: [Errno 2] No such file or directory: 'gdal-config'
```

libgdal

3) Permission denied (Welcome to HPC!)

- If you ask Google / ChatGPT...

```
$ sudo yum install libgdal-devel      # On Red Hat
$ sudo apt-get install libgdal-dev    # On Ubuntu
```

3) Permission denied (Welcome to HPC!)

- If you ask Google / ChatGPT...

```
$ sudo yum install libgdal-devel      # On Red Hat
```

```
$ sudo apt-get install libgdal-dev    # On Ubuntu
```

3) Permission denied (Welcome to HPC!)

- If you ask Google / ChatGPT...

```
$ sudo yum install libgdal-devel      # On Red Hat
$ sudo apt-get install libgdal-dev    # On Ubuntu
```

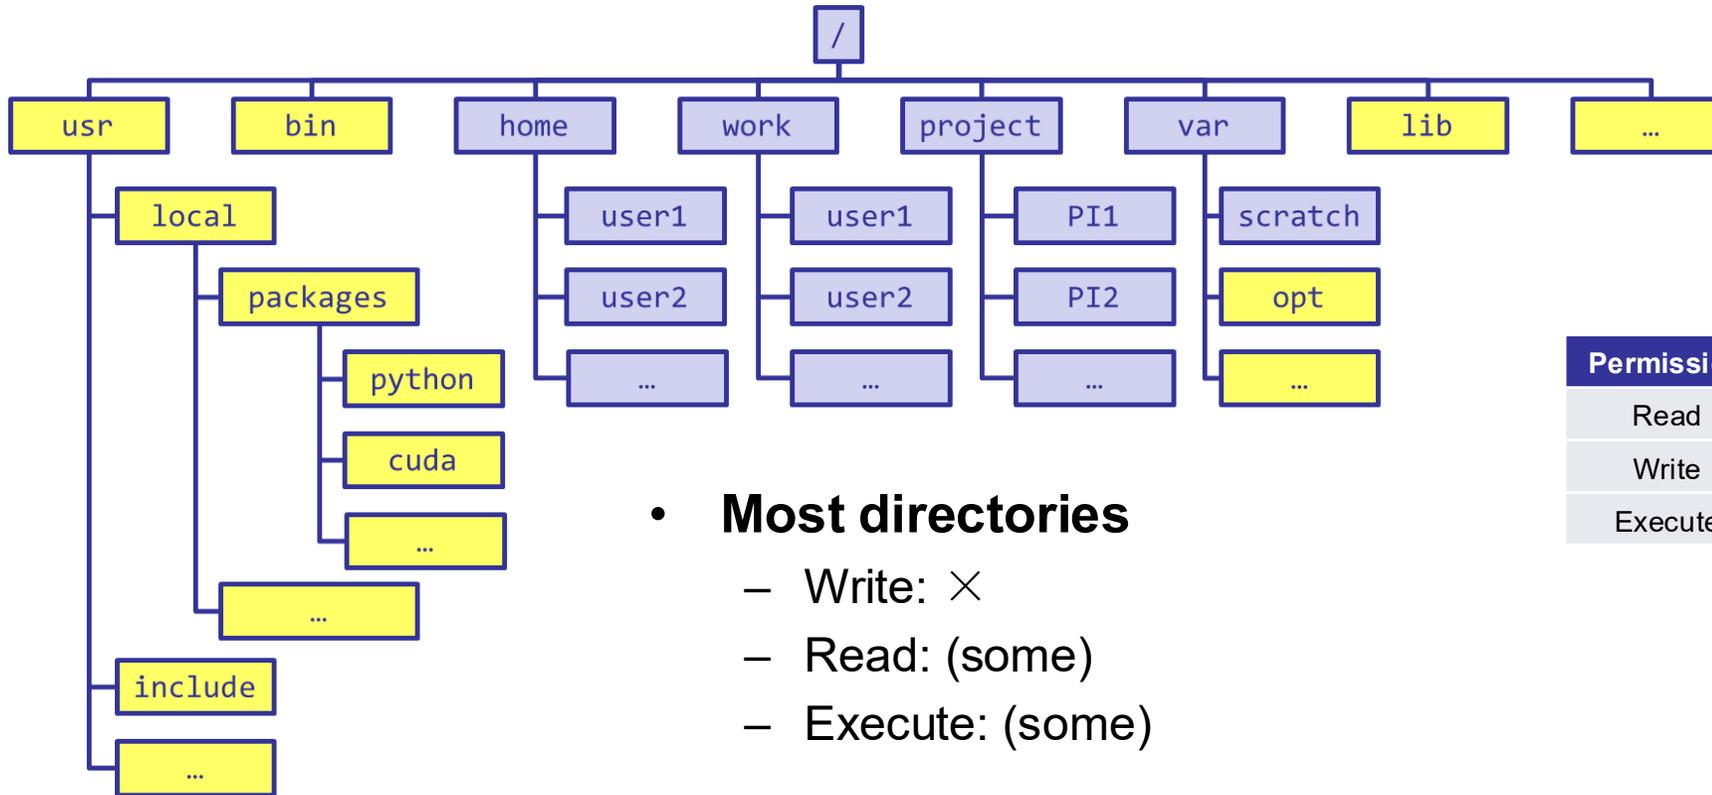
3) Permission denied (Welcome to HPC!)

- If you ask Google / ChatGPT...

```
$ sudo yum install lldpd-devel # On Red Hat  
$ sudo apt-get install lldpd-dev # On Ubuntu
```



3) Permission denied (Welcome to HPC!)



- **Most directories**
 - Write: ✗
 - Read: (some)
 - Execute: (some)

Permission	Green	Red	Yellow
Read	√	✗	√*
Write	√	✗	✗
Execute	√	✗	√*

1

Dependencies

Most software depends on other libraries or tools. If those are missing, the software may not install or run.

2

Version Conflicts

Different packages may need different versions of the same dependency, causing conflicts in a shared environment.

3

root permission

Installing software often requires admin access, which most HPC users do not have.

4

Reproducibility

Reproducing the same setup or sharing it with others can be surprisingly difficult.

4. Reproducibility - Sharing / Migrating your software

- Huge effort & large disk quota to install
 - What if my colleagues want to use?
 - What if I want to migrate a different cluster?

Any of those apply to you?

Magic Tools to Install / Manage Software



1. Why Container?

- 1) Problems
- 2) Container & Singularity

2. Run an Existing Container Image

- 1) What you need
- 2) Basic commands
- 3) Running jobs with Singularity

3. Get More Container Images

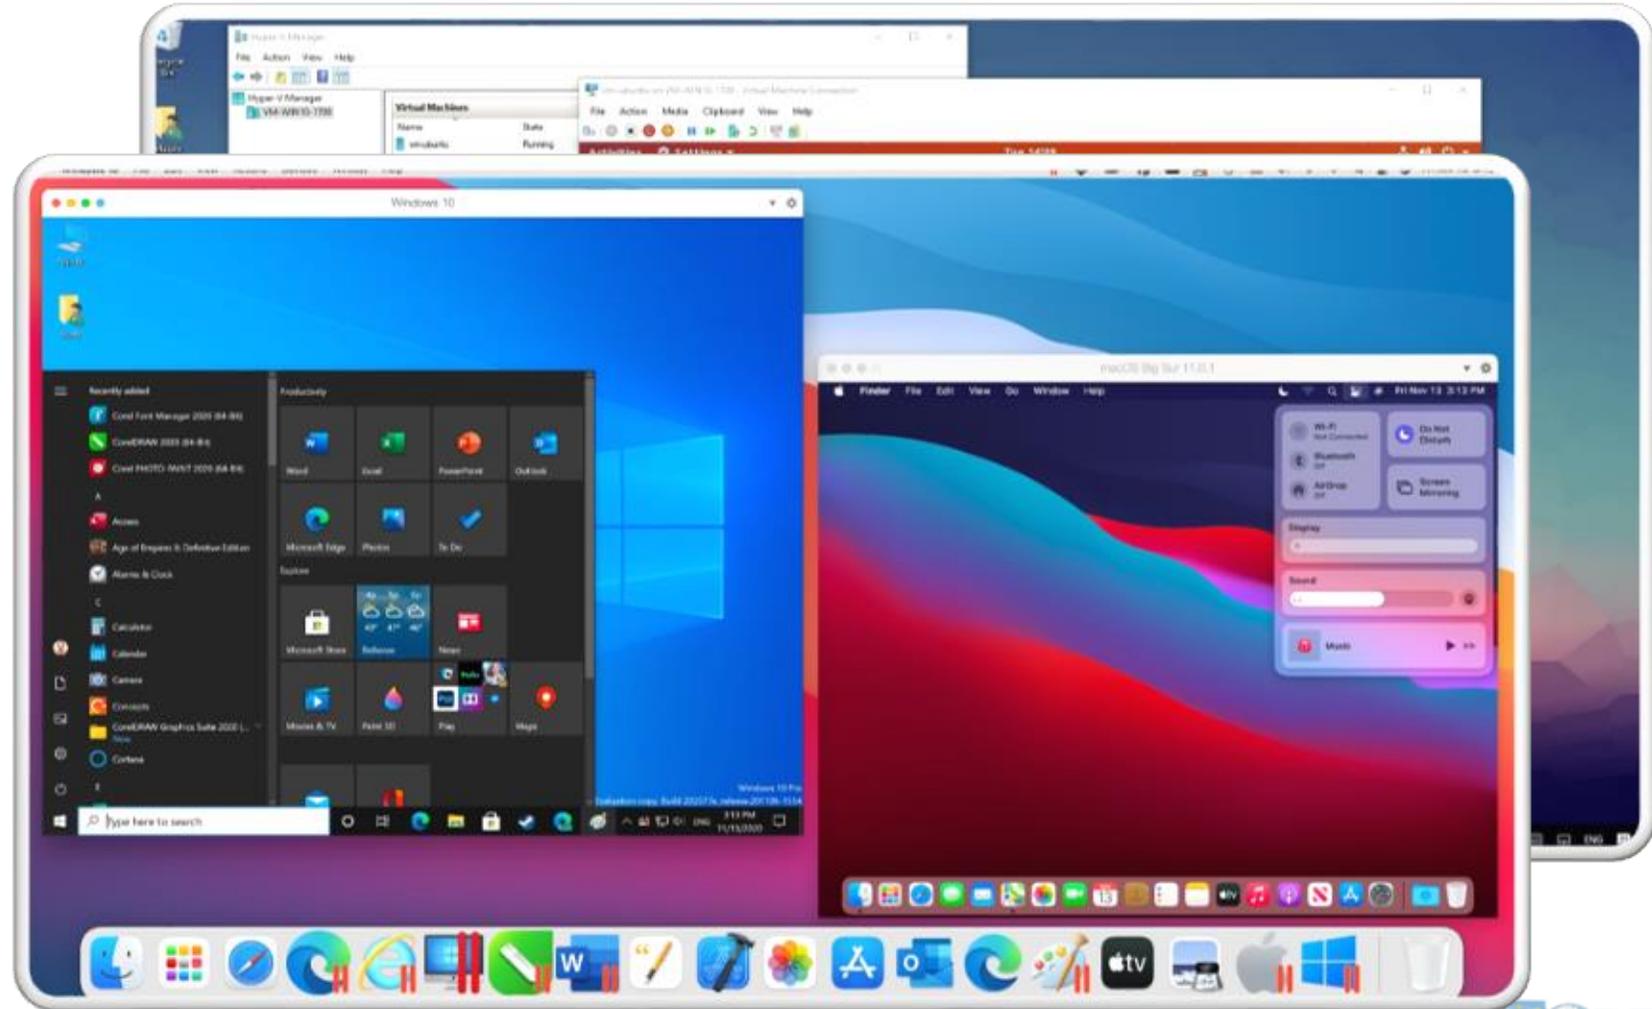
- 1) What you need
- 2) Where to get
- 3) How to get

4. Build Your Own Container Image

- 1) What you need
- 2) Typical workflow
- 3) Make it easier - Recipe

a) What is a **container**?

- **Virtual machine**
 - “Virtualize” / “mimic” an **entire computer** on another computer
 - Virtualize both **hardware** and **software**

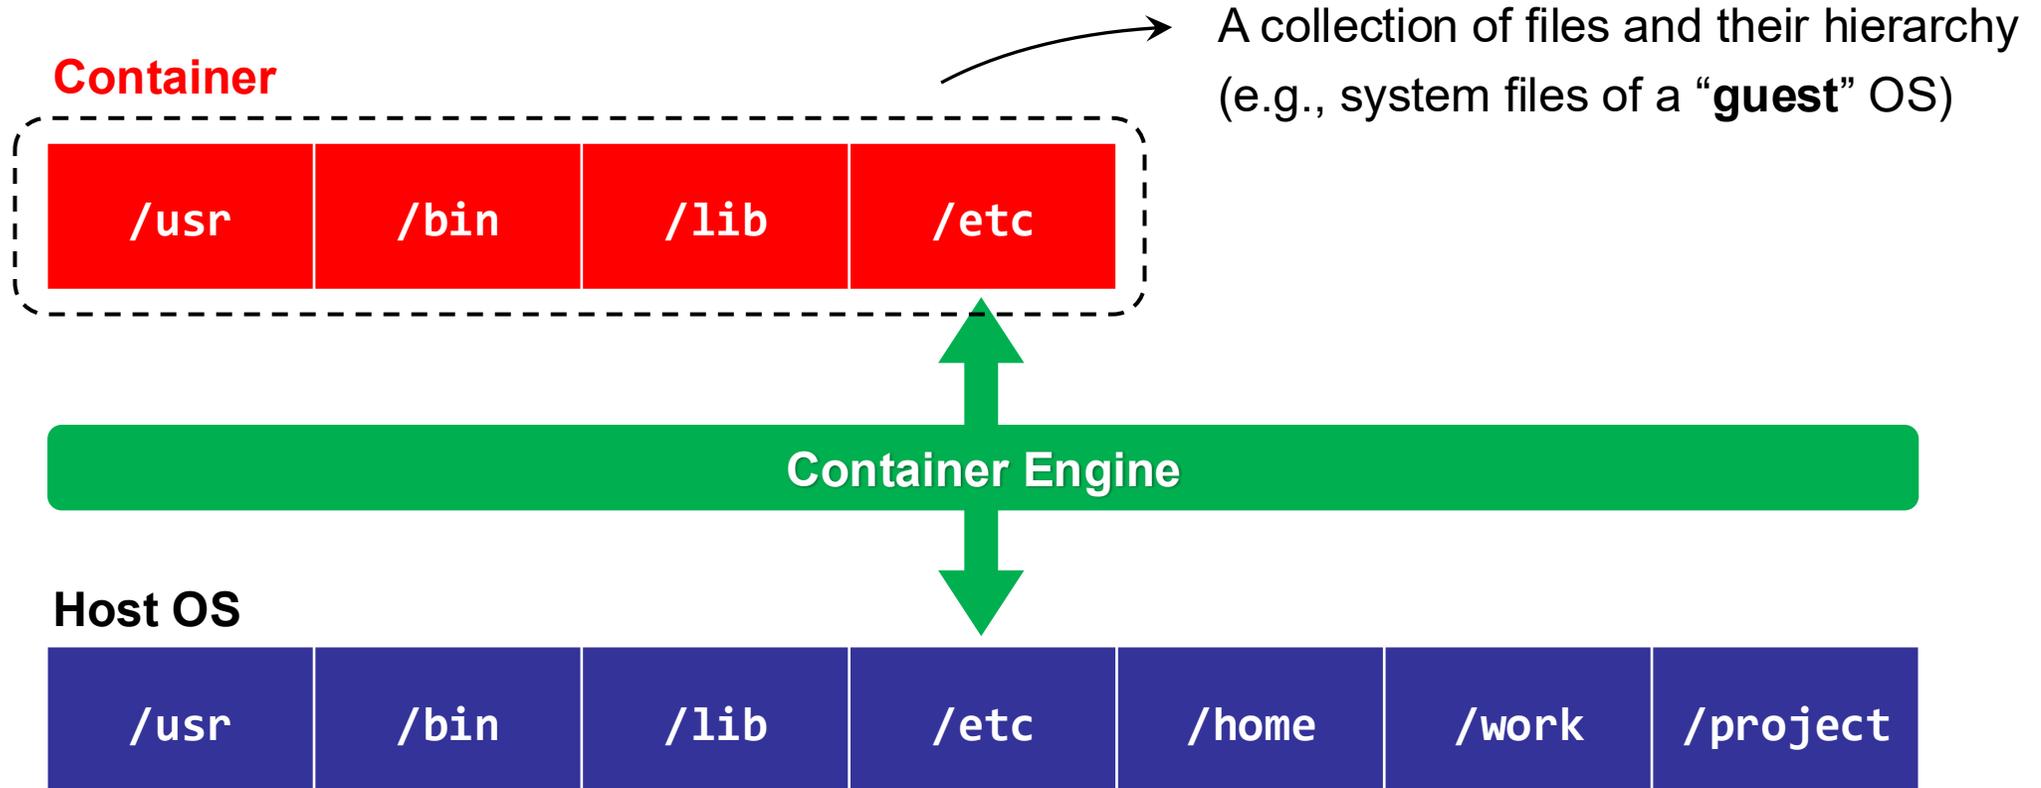


a) What is a **container**?

- **Container:**

- A **lightweight** and **fast** virtual machine
- Only virtualize the **Operation System** (meaning, does not virtualize hardware)
- Only virtualize **Linux** on **Linux**

a) What is a **container**?



A collection of files and their hierarchy (e.g., system files of a "guest" OS)

Container



Host OS



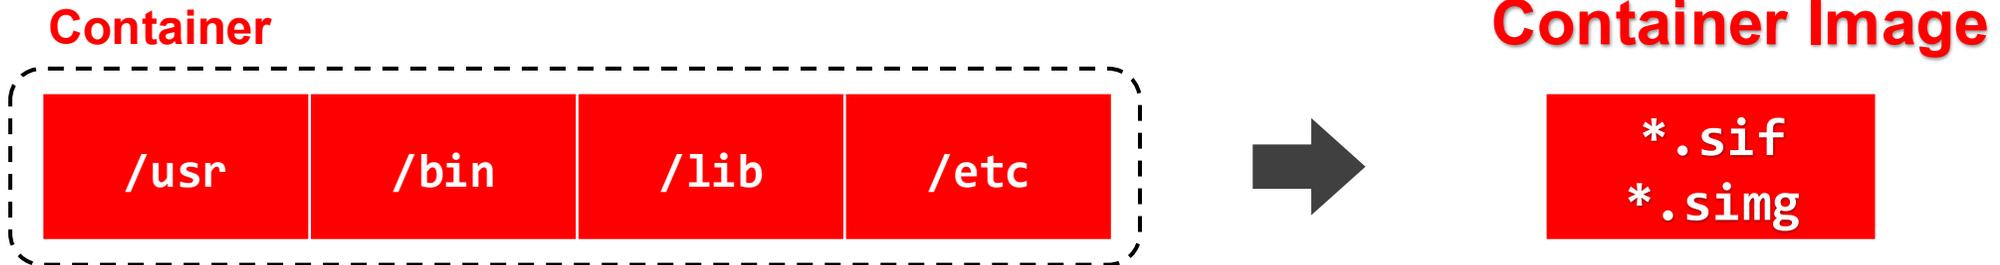
a) What is a **container**?



- A “chimera” system:
 - Can virtualize **an entirely different OS** !
 - Can contain other **software packages** (inc. dependencies, environment settings, etc.) installed in the guest OS



a) What is a **container**?



a) What is a **container**?

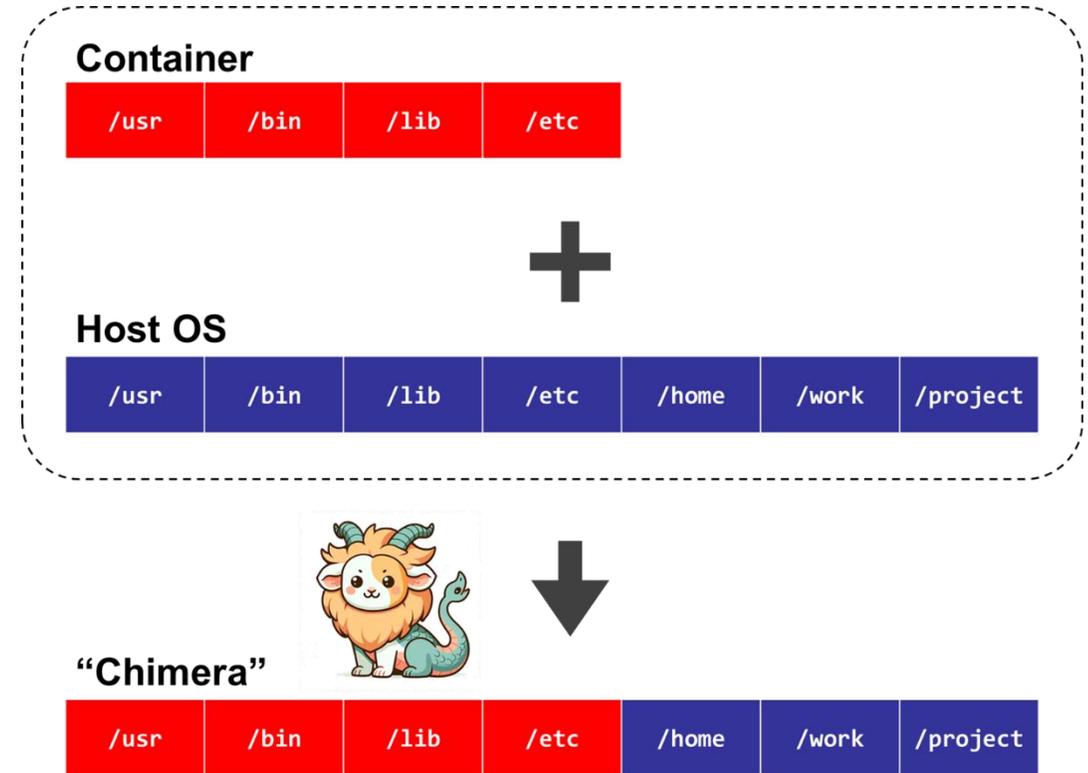
- **Properties:**

- **Self-contained**

All dependencies can be installed within the container

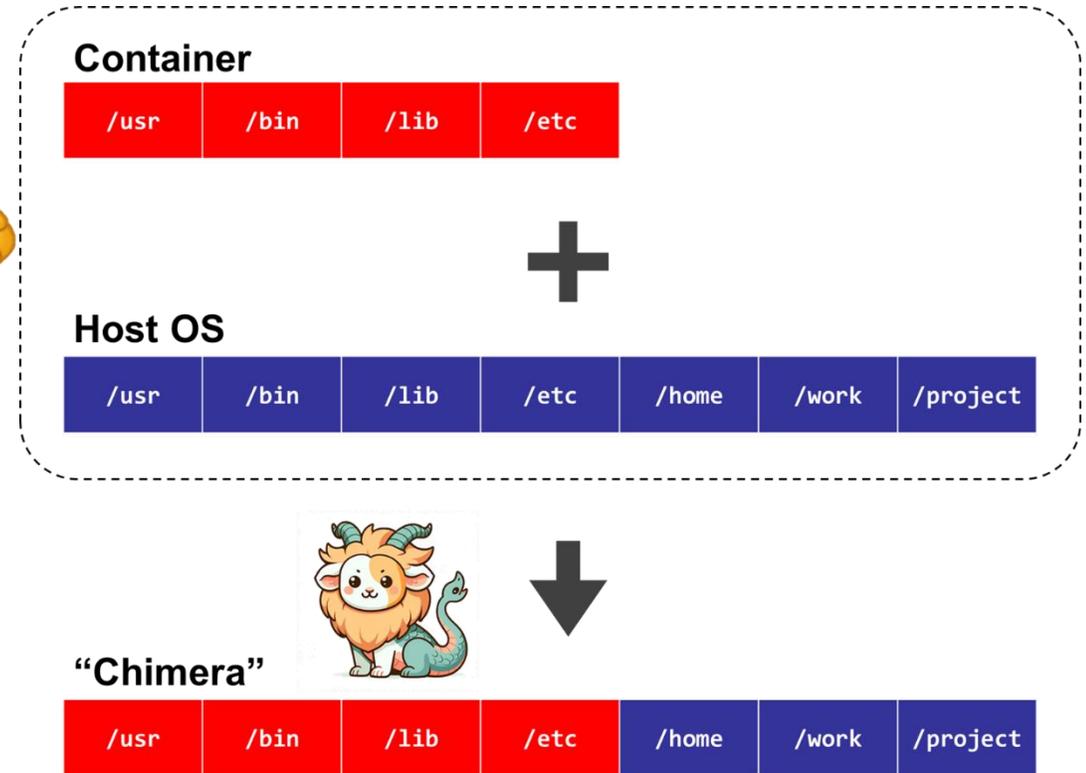
- **Isolated**

Whatever happens in a container stays in that container...



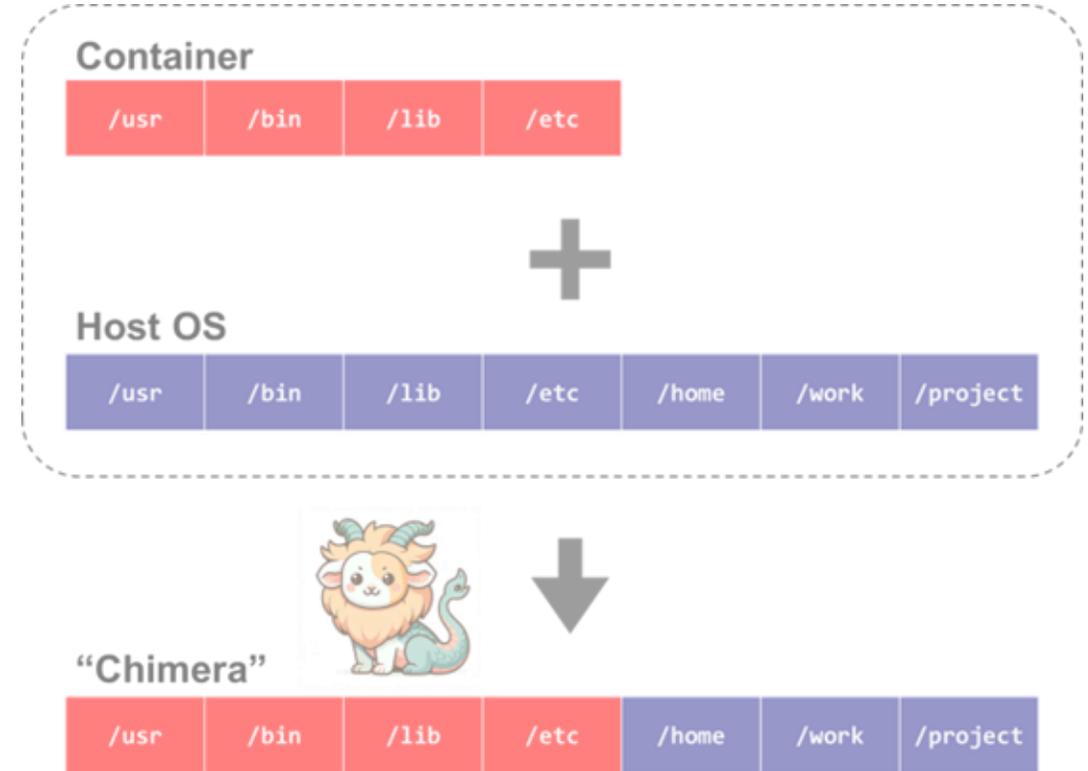
b) How does it solve my problems?

- **Dependency issue**
 - Pack all dependencies (even OS) in container
 - Can use `apt-get` or `yum`
 - **Developers now release containers!**
- **Permission issue**
 - Can't write to certain paths on HPC, but **CAN** write to them in container
- **Conflicted packages**
 - Install in different containers.
- **Share / Migrate**
 - Copy-paste a container image!



c) What is Singularity?

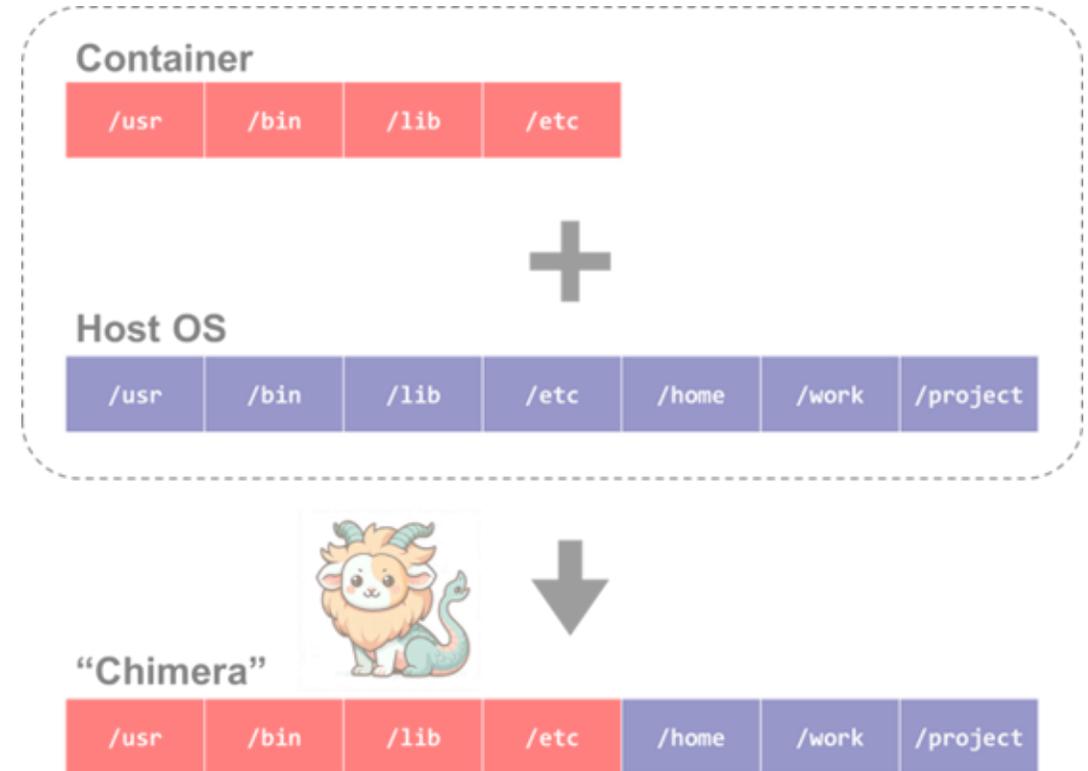
Technology →



c) What is Singularity?



↑ **Software** system that implements the technology



c) What is **Singularity**?



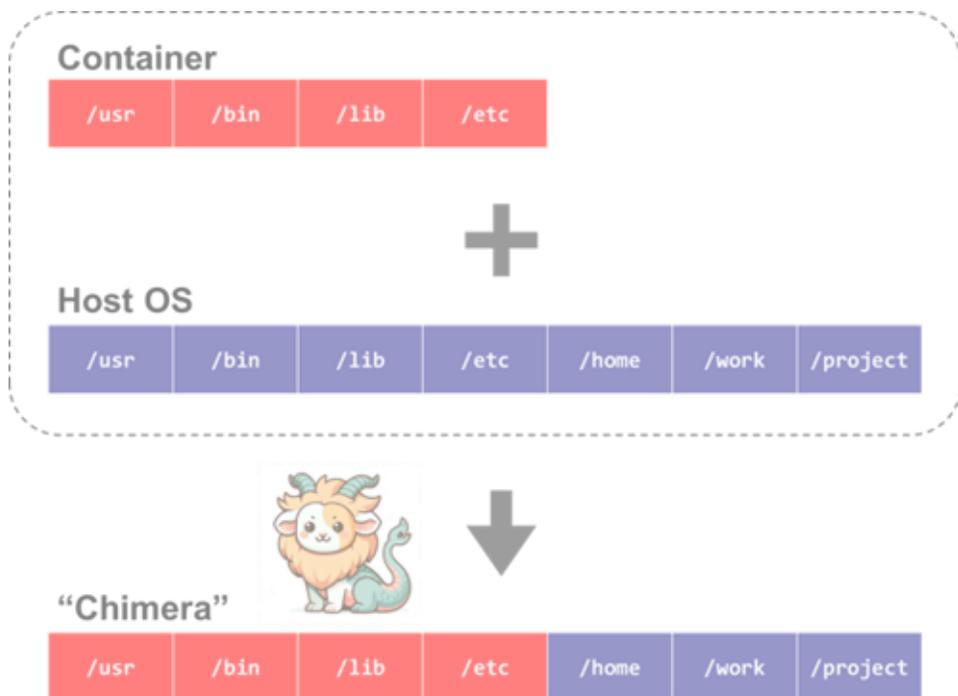
c) What is Singularity?



- Does **NOT** need root privileges
- “**Container for HPC**”



- **Needs** root privileges



Technology that helps with software installation

Software system that implements the technology

1. Why Container?

- 1) Problems
- 2) Container & Singularity

2. Run an Existing Container Image

- 1) What you need
- 2) Basic commands
- 3) Running jobs with Singularity

3. Get More Container Images

- 1) What you need
- 2) Where to get
- 3) How to get

4. Build Your Own Container Image

- 1) What you need
- 2) Typical workflow
- 3) Make it easier - Recipe

1. Why Container?

- 1) Problems
- 2) Container & Singularity

2. Run an Existing Container Image

- 1) What you need
- 2) Basic commands
- 3) Running jobs with Singularity

3. Get More Container Images

- 1) What you need
- 2) Where to get
- 3) How to get

4. Build Your Own Container Image

- 1) What you need
- 2) Typical workflow
- 3) Make it easier - Recipe

- **Singularity availability**
 - a) On **all clusters**
 - ✓ **LSU HPC:** SMIC, Deep Bayou, SuperMike 3
 - ✓ **LONI:** QB3, QB4
 - b) Only on **computing nodes**
 - × Unavailable on head nodes
 - ✓ Must start a job (interactive & batch)
 - c) To **all users**
 - × No additional action required

To ...	What you need ...
Run an Existing Container Image	<ul style="list-style-type: none">• Access to our HPC systems<ul style="list-style-type: none">- An active account- An active allocation

1. Why Container?

- 1) Problems
- 2) Container & Singularity

2. Run an Existing Container Image

- 1) What you need
- 2) Basic commands
- 3) Running jobs with Singularity

3. Get More Container Images

- 1) What you need
- 2) Where to get
- 3) How to get

4. Build Your Own Container Image

- 1) What you need
- 2) Typical workflow
- 3) Make it easier - Recipe

- Available images

- On all clusters: `/project/containers/images`

```
[kasetti@qbd1 ~]$ ls /project/containers/images/  
adcirc.5bc04d6.centos8.sif          mesos_1.9.0_ubuntu16.04.simg  
adcirc.5bc04d6.rky8.5.sif          minknow22.sif  
adcirc.5bc04d6.sif                 nco-4.9.8.sif  
agat-1.4.0.sif                     nextflow-23.10.1.sif  
alphafold-catgumag-2.2.sif         nvidia-dli.sif  
alps-2.3.0-dockerhub.simg          octopus-14.0.sif  
alps-2.3.0-dockerhub-v2.simg       ollamaondemand-1.0.0.sif  
bcftools-1.18.sif                 ollamaondemand-1.0.0.sif.old  
beast2-2.7.7.sif                  openbabel-3.1.1.sif  
blast-2.14.1.sif                  openfoam10-openmpi.4.0.3-pmi2.sif  
blender-2.79b-cuda-8.0-ubuntu-16.04.simg openfoam9.sdfibm-openmpi.4.0.3-pmi2.sif  
bowtie2-2.5.1.sif                 openfoam-v2212-openmpi.4.1.2-pmi2.sif  
braker-3.0.8.sif                  openfoam.v9.sdfibm.sif  
busco-5.7.1.sif                   paraviewweb-5.12.1.sif
```

a) Common usage 1: Open a shell in the image

Syntax	Description
<code>singularity shell <container></code>	Starts an interactive shell in the image

Try me: `/project/containers/images/singularity-training-ubuntu.sif`

a) Common usage 1: Open a shell in the image

Syntax		Description
singularity shell <i>[options]</i> <container>		Starts an interactive shell in the image
<i>[Options]</i>	-B /path/to/bind	Bind a path(s) <ul style="list-style-type: none">• Access host files with bind mounts• /home, /work is bound by default
	--nv	Enable Nvidia GPU

b) Common usage 2: Execute a single command in the image

Syntax	Description
<code>singularity exec <container> <command></code>	Execute a command in the image

Try me: `/project/containers/images/singularity-training-ubuntu.sif`

b) Common usage 2: Execute a single command in the image

Syntax		Description
singularity exec <i>[options]</i> <container> <command>		Execute a command in the image
<i>[Options]</i>	-B /path/to/bind	Bind a path(s) <ul style="list-style-type: none">• /home is bound by default
	--nv	Enable Nvidia GPU

c) Another (less) common usage: Run a prewritten script

Syntax		Description
singularity run <i>[options]</i> <container>		Run a prewritten script
<i>[Options]</i>	-B /path/to/bind	Bind a path(s) <ul style="list-style-type: none">• /home is bound by default
	--nv	Enable Nvidia GPU

- Quick recap

Syntax	Description
singularity shell <i>[options]</i> <container>	Starts an interactive shell in the image
singularity exec <i>[options]</i> <container> <command>	Execute a command in the image
singularity run <i>[options]</i> <container>	Run a prewritten script

1. Why Container?

- 1) Problems
- 2) Container & Singularity

2. Run an Existing Container Image

- 1) What you need
- 2) Basic commands
- 3) Running jobs with Singularity

3. Get More Container Images

- 1) What you need
- 2) Where to get
- 3) How to get

4. Build Your Own Container Image

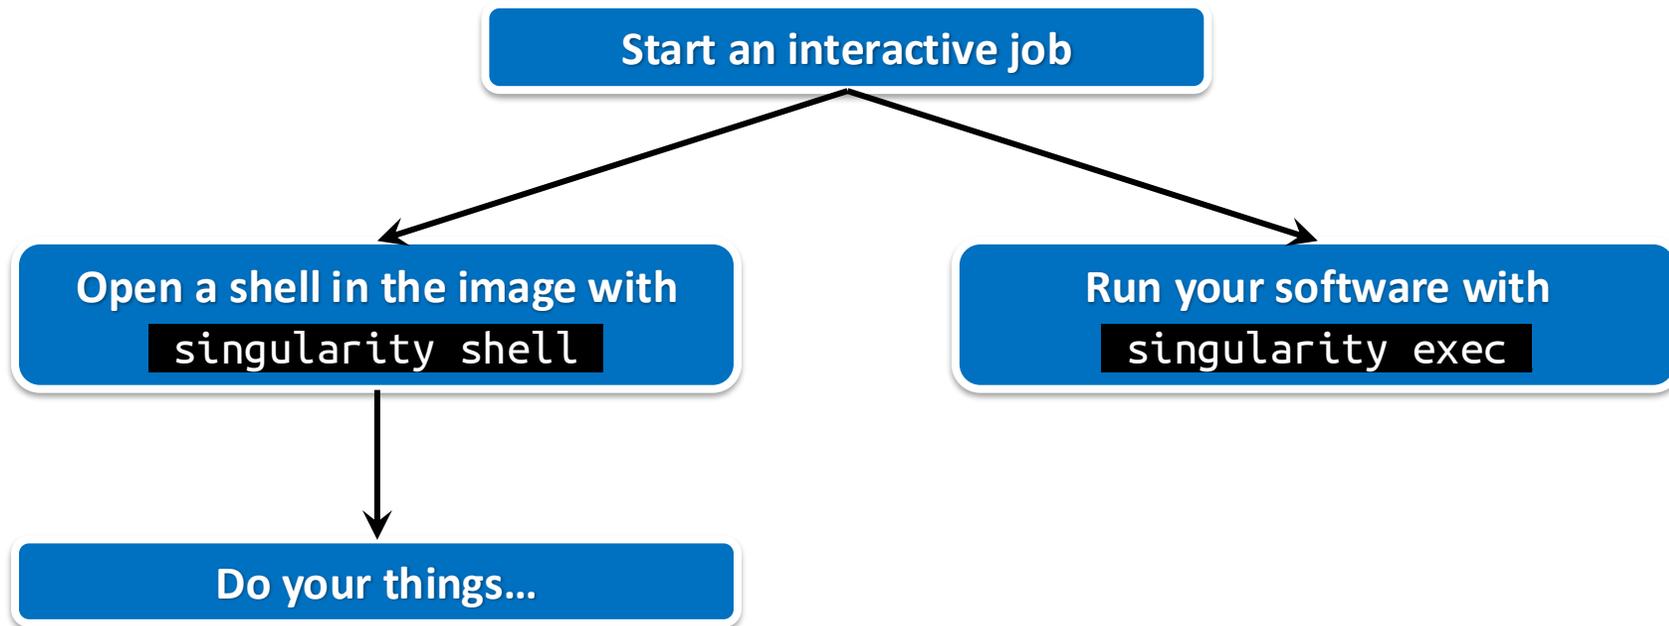
- 1) What you need
- 2) Typical workflow
- 3) Make it easier - Recipe

3) Run jobs with Singularity

- Job types and commands

Job Type	Commands	Purpose
Interactive	<ul style="list-style-type: none">• singularity shell <i>[options]</i> <container>• singularity exec <i>[options]</i> <container> <command>	<ul style="list-style-type: none">• Debugging & testing
Batch	<ul style="list-style-type: none">• singularity exec <i>[options]</i> <container> <command>	<ul style="list-style-type: none">• Production

a) Interactive job



b) Batch job

```
#!/bin/bash
#SBATCH -A <Allocation name>
#SBATCH -p workq
#SBATCH -N 1
#SBATCH -n 64
#SBATCH -t 24:00:00
```

Example

```
cd /to/work/directory
```

```
IMG=/home/admin/singularity/singularity-  
training-ubuntu.sif
```

```
singularity exec -B /work,/project $IMG \  
python myjob.py
```

Syntax	Description
singularity shell <i>[options]</i> <container>	Run a prewritten script
singularity exec <i>[options]</i> <container> <command>	Execute a command in the image
singularity run <i>[options]</i> <container>	Run a prewritten script

1. Why Container?

- 1) Problems
- 2) Container & Singularity

2. Run an Existing Container Image

- 1) What you need
- 2) Basic commands
- 3) Running jobs with Singularity

3. Get More Container Images

- 1) What you need
- 2) Where to get
- 3) How to get

4. Build Your Own Container Image

- 1) What you need
- 2) Typical workflow
- 3) Make it easier - Recipe

3. Get More Container Images

- Available images

- On all clusters: `/project/containers/images`

```
[kasetti@qbd1 ~]$ ls /project/containers/images/  
adcirc.5bc04d6.centos8.sif          mesos_1.9.0_ubuntu16.04.simg  
adcirc.5bc04d6.rky8.5.sif          minknow22.sif  
adcirc.5bc04d6.sif                 nco-4.9.8.sif  
agat-1.4.0.sif                     nextflow-23.10.1.sif  
alphafold-catgumag-2.2.sif         nvidia-dli.sif  
alps-2.3.0-dockerhub.simg          octopus-14.0.sif  
alps-2.3.0-dockerhub-v2.simg       ollamaondemand-1.0.0.sif  
bcftools-1.18.sif                  ollamaondemand-1.0.0.sif.old  
beast2-2.7.7.sif                   openbabel-3.1.1.sif  
blast-2.14.1.sif                   openfoam10-openmpi.4.0.3-pmi2.sif  
blender-2.79b-cuda-8.0-ubuntu-16.04.simg openfoam9.sdfibm-openmpi.4.0.3-pmi2.sif  
bowtie2-2.5.1.sif                  openfoam-v2212-openmpi.4.1.2-pmi2.sif  
braker-3.0.8.sif                   openfoam.v9.sdfibm.sif  
busco-5.7.1.sif                    paraviewweb-5.12.1.sif
```

1. Why Container?

- 1) Problems
- 2) Container & Singularity

2. Run an Existing Container Image

- 1) What you need
- 2) Basic commands
- 3) Running jobs with Singularity

3. Get More Container Images

- 1) What you need
- 2) Where to get
- 3) How to get

4. Build Your Own Container Image

- 1) What you need
- 2) Typical workflow
- 3) Make it easier - Recipe

To ...	What you need ...
Run an Existing Container Image	<ul style="list-style-type: none">• Access to our HPC systems<ul style="list-style-type: none">- An active account- An active allocation
BEFORE Get More Container Images	<ul style="list-style-type: none">• (Everything above)• Added to Singularity group<ul style="list-style-type: none">- Needed to change group ownership- Send email to sys-help@loni.org to request

To ...	What you need ...
Run an Existing Container Image	<ul style="list-style-type: none">• Access to our HPC systems<ul style="list-style-type: none">- An active account- An active allocation
NOW Get More Container Images	<ul style="list-style-type: none">• (Everything above)• Added to Singularity group<ul style="list-style-type: none">— Needed to change group ownership— Send email to sys-help@loni.org to request

1. Why Container?

- 1) Problems
- 2) Container & Singularity

2. Run an Existing Container Image

- 1) What you need
- 2) Basic commands
- 3) Running jobs with Singularity

3. Get More Container Images

- 1) What you need
- 2) Where to get
- 3) How to get

4. Build Your Own Container Image

- 1) What you need
- 2) Typical workflow
- 3) Make it easier - Recipe

2) Where to get

- You can get container images from a lot of places

➤ **Not that you should!**

- **Concerns?**

- **Reliability**

- Some third-party or untested images may not work

- **Security risk**

- Some untrustworthy publishers may pack something you don't know about

- **Solution**

- Always get from a source that **you can trust !**



[1] <https://www.techradar.com/pro/security/malware-attacks-on-docker-hub-spread-millions-of-malicious-repositories>

- **Tier 1: Developer release (official release)**
 - On software’s official website, look for “**Docker**” / “**Singularity**” / “**Container**” / etc.
 - E.g., [Tensorflow](#), [Trinity](#), [Salmon](#)
- **Tier 2: Trustworthy third party**

Name	Notes
Biocontainers	<ul style="list-style-type: none">• https://biocontainers-edu.readthedocs.io/en/latest/• For biology
Nvidia NGC	<ul style="list-style-type: none">• https://catalog.ngc.nvidia.com/containers• For Nvidia GPU
Bitnami	<ul style="list-style-type: none">• https://bitnami.com/stacks/containers• By VmWare
Docker Hub Quay.io	<ul style="list-style-type: none">• https://hub.docker.com/ & https://quay.io/• BUT! Do NOT just trust everything you see there!• Look for trustworthy icons like Docker Official Image or Verified Publisher• Avoid third-party publishers that you don’t know

1. Why Container?

- 1) Problems
- 2) Container & Singularity

2. Run an Existing Container Image

- 1) What you need
- 2) Basic commands
- 3) Running jobs with Singularity

3. Get More Container Images

- 1) What you need
- 2) Where to get
- 3) How to get

4. Build Your Own Container Image

- 1) What you need
- 2) Typical workflow
- 3) Make it easier - Recipe

- **Steps:**

- a) Step 1: Pull the image

- ~~b) Step 2: Change group ownership~~

3) How to get

a) Step 1: Pull the image

Syntax		Description
singularity pull <code><source></code>		Pull an image from source
<code><source></code>	<code>docker://container[:tag]</code> <ul style="list-style-type: none">(Compare to Docker command) <code>docker pull container[:tag]</code>	Pull a Docker container and convert to Singularity <ul style="list-style-type: none">Many software's official container release is in Docker form (may or may not on Docker Hub)
	<code>http://www.myexample.com/container_image.sif</code>	Download an image file from a web source

a) Step 1: Pull the image

Syntax		Description
singularity build	<code><target></code> <code><source></code>	Build an image from source (Advanced)
	<code>docker://container[:tag]</code>	Build from a Docker container
<code><source></code>	<code>container_image.sif</code>	Build from a local image file
	<code>container_sandbox/</code>	Build from a local sandbox (A directory form of a container)
	<code>container_recipe.def</code>	Build from a recipe (an instruction script of how to build an image)

a) Step 1: Pull the image

Syntax	Description
<code>singularity pull [options] [target] <source></code>	Simple pull
<code>singularity build [options] <target> <source></code>	Advanced build command

3) How to get

- **BONUS: Hot packages!**
 - i. **PyTorch** (w/ GPU support)

```
$ singularity pull docker://pytorch/pytorch:2.10.0-cuda13.0-cudnn9-devel
```

- ii. **Tensorflow** (w/ GPU support)

```
$ singularity pull docker://tensorflow/tensorflow:2.21.0-gpu-jupyter
```

3) How to get

- **BONUS: Hot packages!**
 - i. **PyTorch** (w/ GPU support)

```
$ module load pytorch
```

- ii. **Tensorflow** (w/ GPU support)

```
$ module load tensorflow
```

- Pull the image or/and build the image

Syntax	Description
<code>singularity pull [options] [target] <source></code>	Simple pull
<code>singularity build [options] <target> <source></code>	Advanced build command

1. Why Container?

- 1) Problems
- 2) Container & Singularity

2. Run an Existing Container Image

- 1) What you need
- 2) Basic commands
- 3) Running jobs with Singularity

3. Get More Container Images

- 1) What you need
- 2) Where to get
- 3) How to get

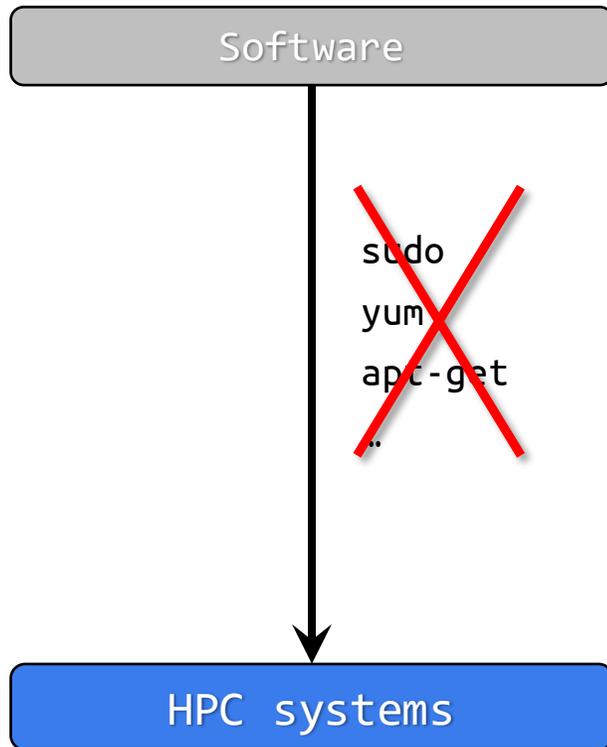
4. Build Your Own Container Image

- 1) What you need
- 2) Typical workflow
- 3) Make it easier - Recipe

- **Scenarios:**

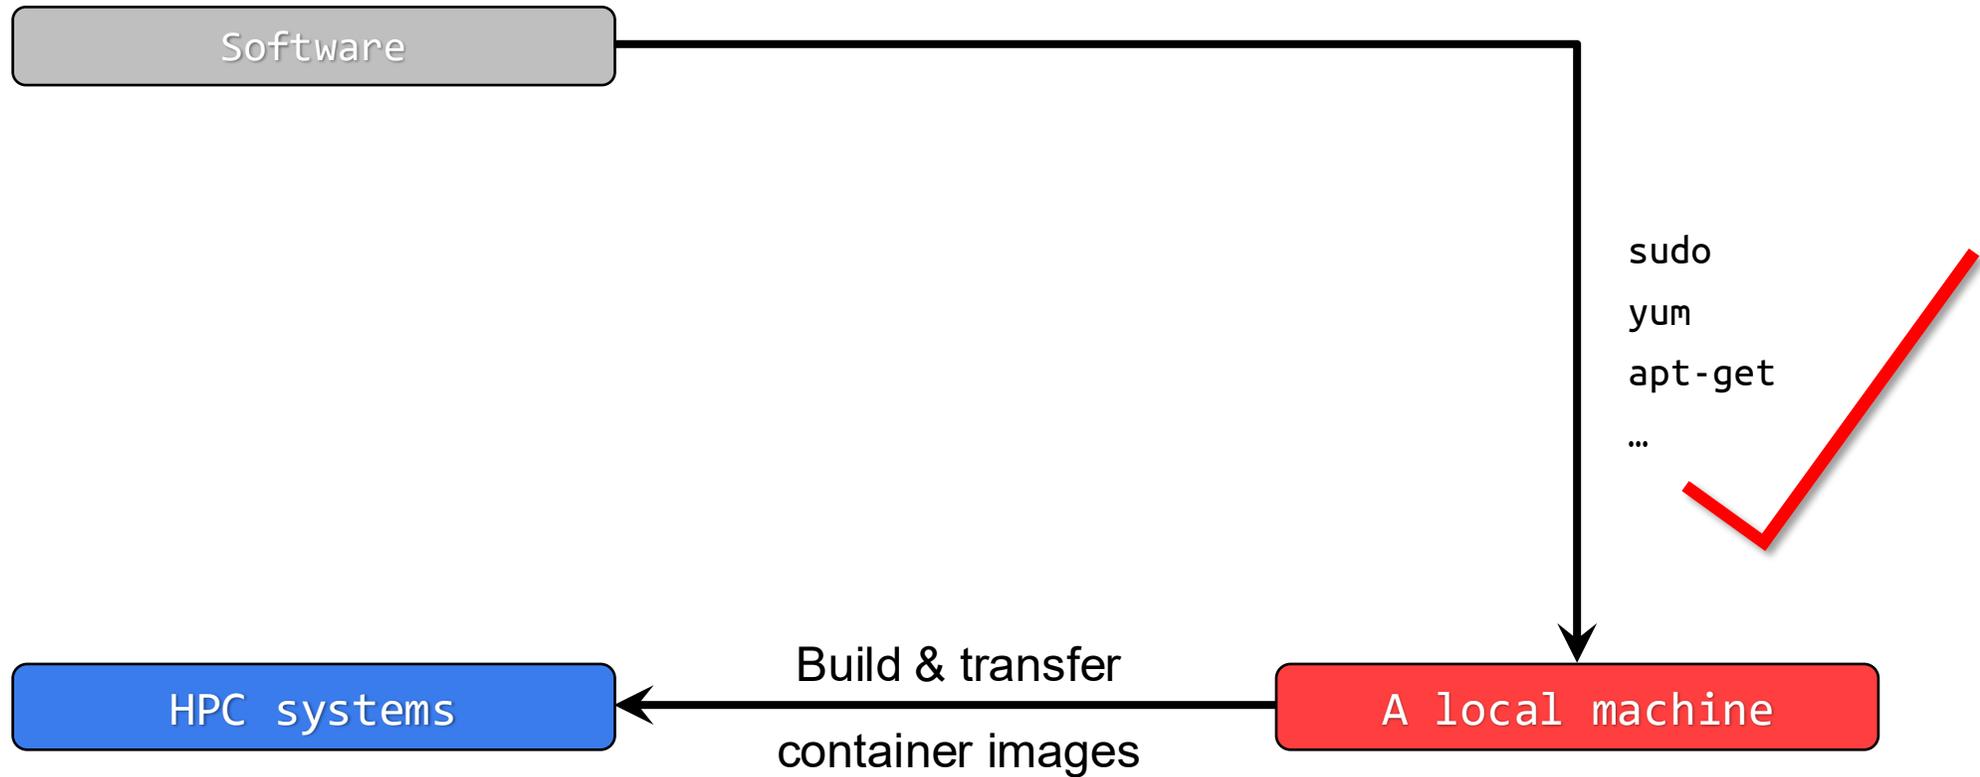
- I did not find any container release. Need to DIY.
- Installation can be easily done using `sudo apt` or `sudo yum` if I have the permission.
- I found a container, but need to make changes (e.g., adding something else).

- Idea



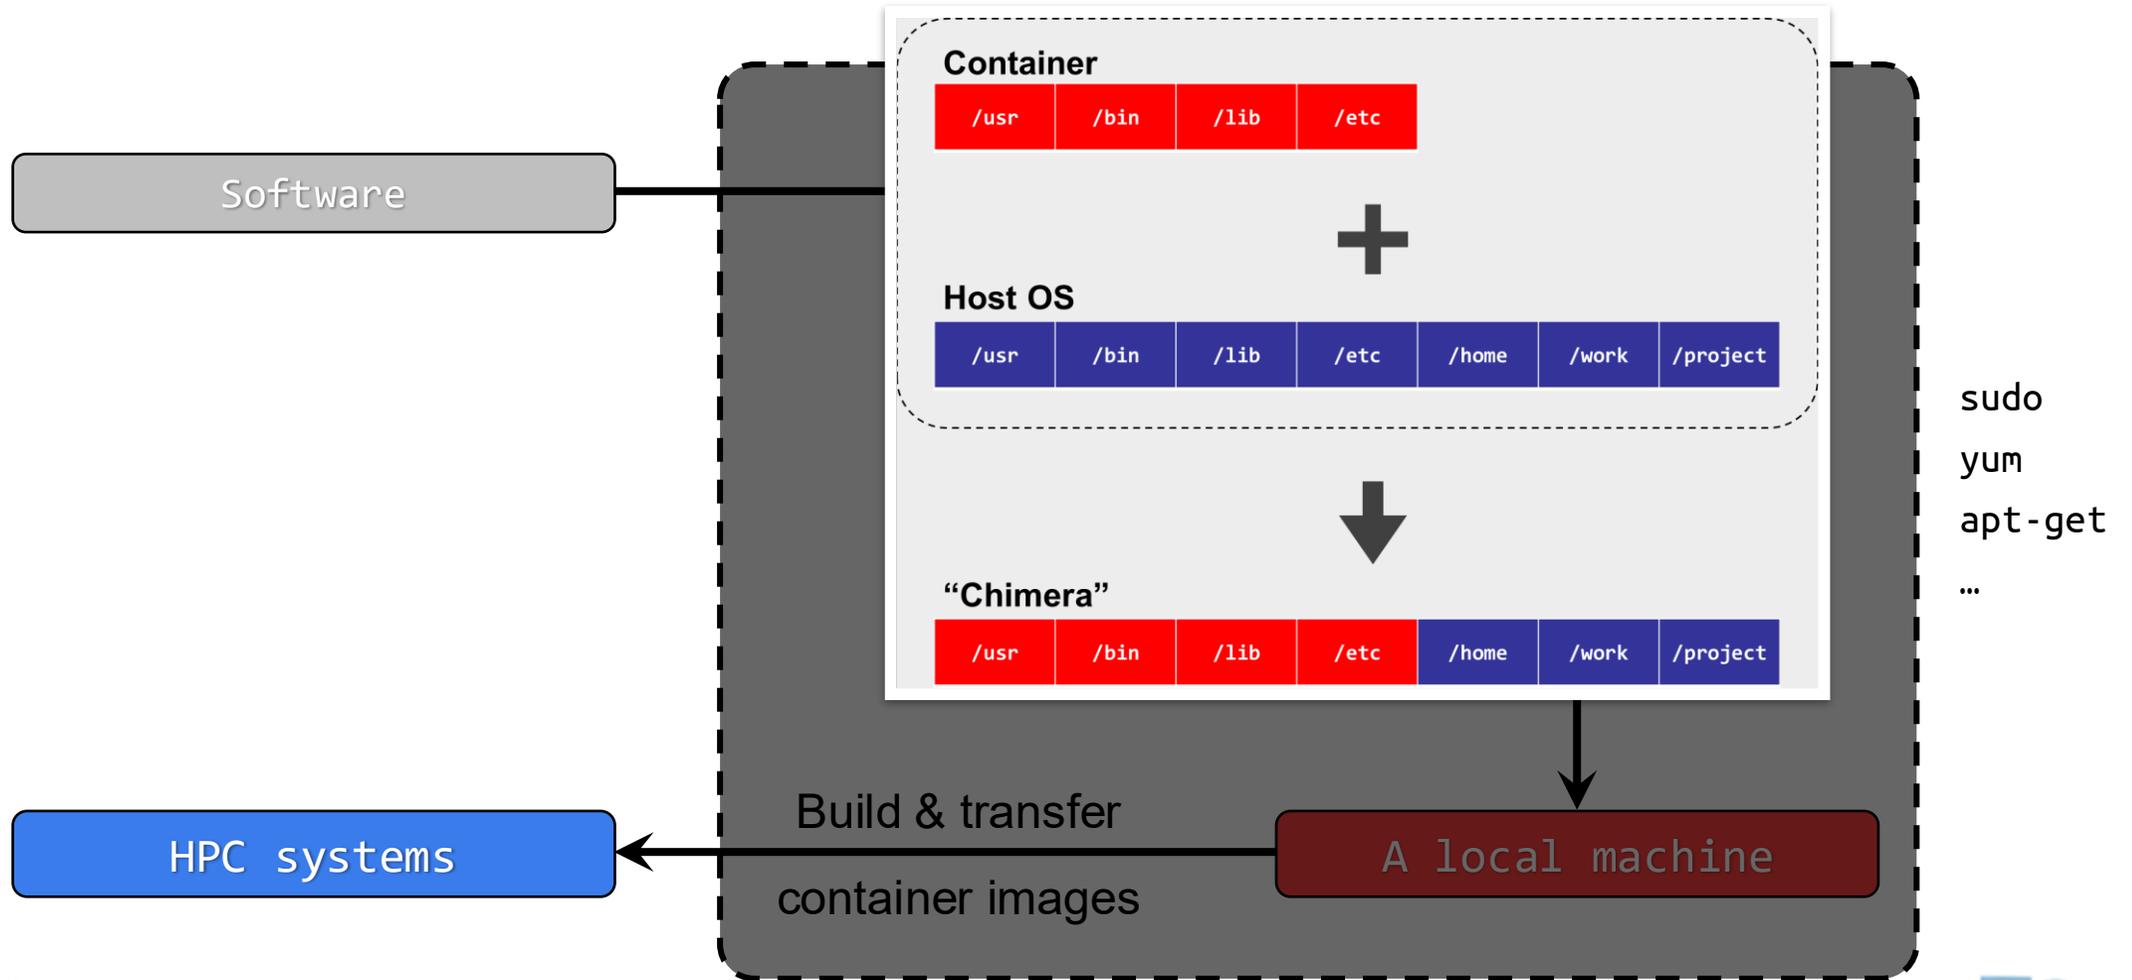
4. Build Your Own Container Image

- Idea



4. Build Your Own Container Image

- Idea



1. Why Container?

- 1) Problems
- 2) Container & Singularity

2. Run an Existing Container Image

- 1) What you need
- 2) Basic commands
- 3) Running jobs with Singularity

3. Get More Container Images

- 1) What you need
- 2) Where to get
- 3) How to get

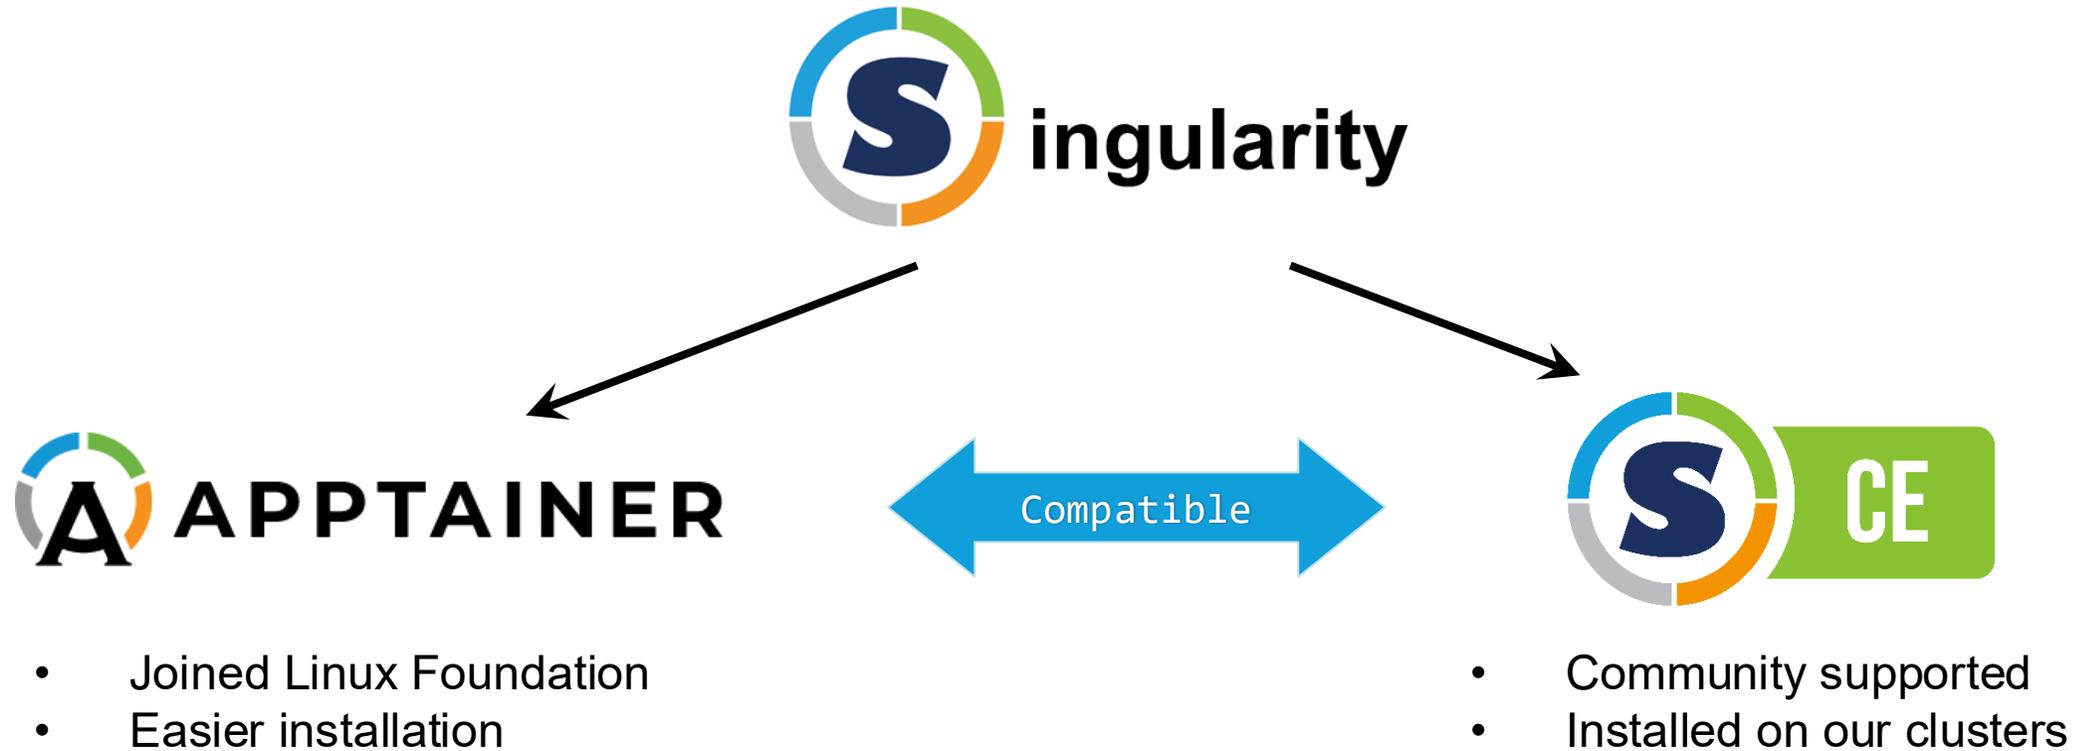
4. Build Your Own Container Image

- 1) What you need
- 2) Typical workflow
- 3) Make it easier - Recipe

To ...	What you need ...
Run an Existing Container Image	<ul style="list-style-type: none">• Access to our HPC systems<ul style="list-style-type: none">- An active account- An active allocation
Get More Container Images	<ul style="list-style-type: none">• (Everything above)• Added to Singularity group<ul style="list-style-type: none">— Needed to change group ownership— Send email to sys-help@loni.org to request
Build Your Own Container Image	<ul style="list-style-type: none">• (Everything above)• A local Linux machine, where:<ul style="list-style-type: none">- You have root permission

1) What you need

- Install Singularity



1. Why Container?

- 1) Problems
- 2) Container & Singularity

2. Run an Existing Container Image

- 1) What you need
- 2) Basic commands
- 3) Running jobs with Singularity

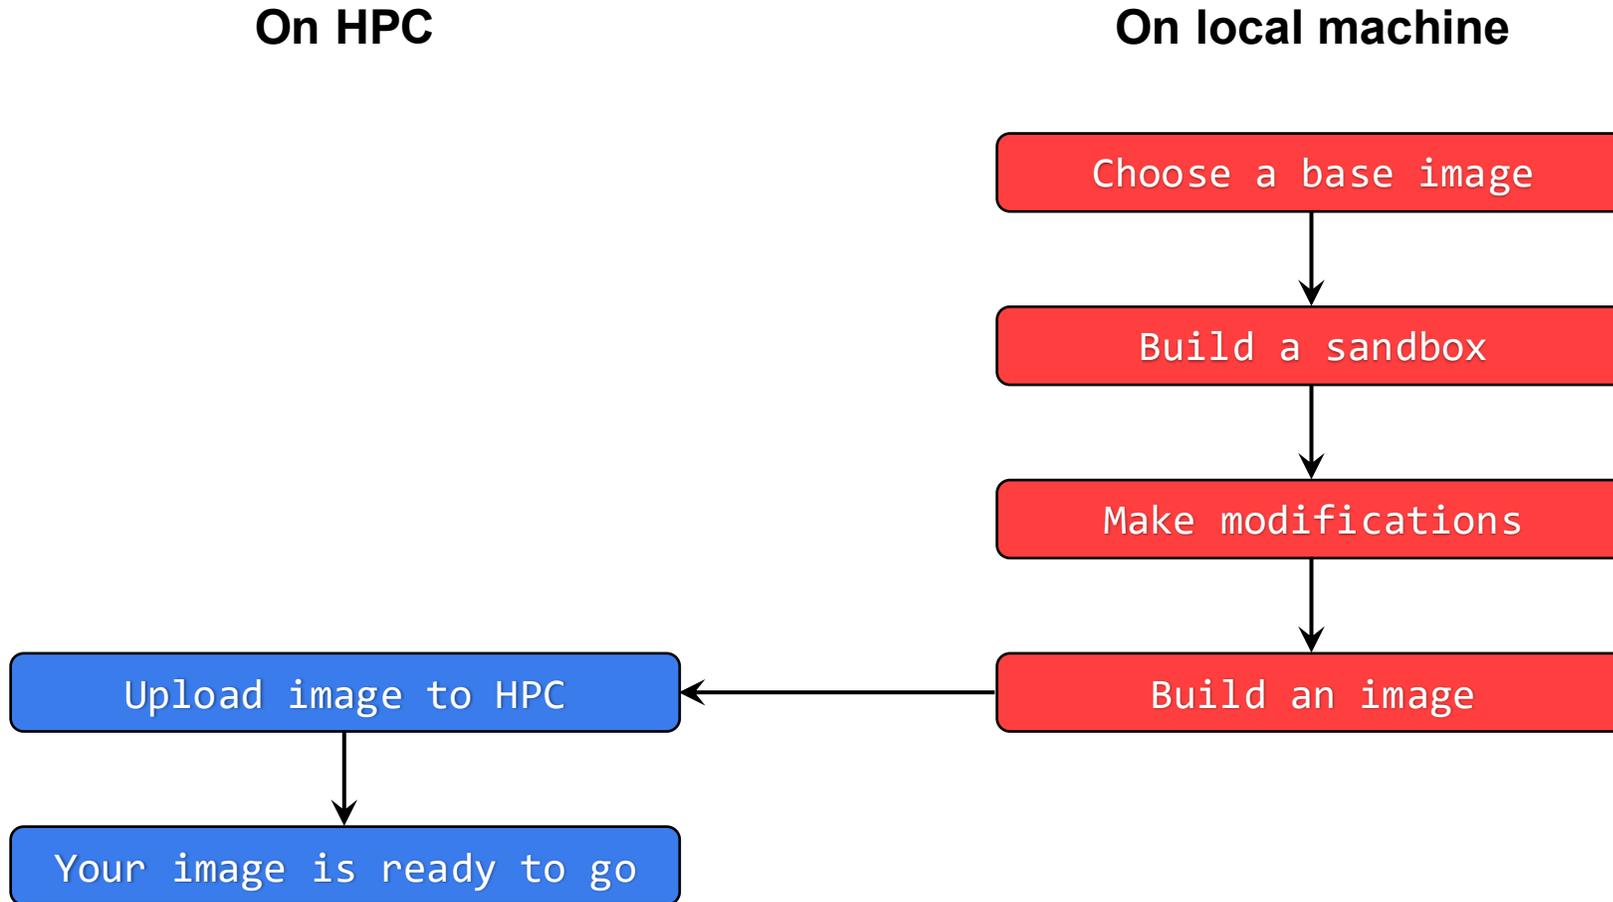
3. Get More Container Images

- 1) What you need
- 2) Where to get
- 3) How to get

4. Build Your Own Container Image

- 1) What you need
- 2) Typical workflow
- 3) Make it easier - Recipe

2) Typical workflow



a) Choose a base image

Common choices	Typical scenarios
<p>A minimum, “mint” OS (e.g., Ubuntu, Rocky, Debian, ...)</p>	<ul style="list-style-type: none">• You cannot find an existing image with the software you need, and need to install from the scratch.• You need to build a minimum size image
<p>A container with software already installed (e.g., TensorFlow, PyTorch, ...)</p>	<ul style="list-style-type: none">• You need to modify an existing working image (e.g., add a Python module to Tensorflow image)

b) Build a sandbox

- What's a **sandbox** ?
 - A **directory** form (unlike a single image file) of a container
 - Allows modification

b) Build a sandbox

```
$ singularity build [options] <target> <source>
```



<source>	<code>docker://container[:tag]</code>	Build from a Docker container
	<code>container_image.sif</code>	Build from a local image file
	<code>container_sandbox/</code>	Build from a local sandbox (A directory form of a container)
	<code>container_recipe.def</code>	Build from a recipe (an instruction script of how to build an image)

b) Build a sandbox

```
$ singularity build --sandbox [options] <target> <source>
```



<source>	<code>docker://container[:tag]</code>	Build from a Docker container
	<code>container_image.sif</code>	Build from a local image file
	<code>container_sandbox/</code>	Build from a local sandbox (A directory form of a container)
	<code>container_recipe.def</code>	Build from a recipe (an instruction script of how to build an image)

c) Make modifications

```
$ singularity shell [options] <container>
```

c) Make modifications

```
$ singularity shell --writable [options] <container>
```

- i. Allows **writing** to the sandbox
 - Read-only without it

c) Make modifications

```
$ sudo singularity shell --writable [options] <container>
```

- ii. Run the container as **root**
 - Grants root privilege in container
 - Needed in most cases
 - Technically not required, but cannot run things like `sudo apt` or `sudo yum` without it

- i. Allows **writing** to the sandbox
 - Read-only without it

c) Make modifications

```
$ sudo singularity shell --writable [options] <container>  
Singularity>  
Singularity> apt update  
Singularity> apt install ...
```

d) Build an image from sandbox

```
$ singularity build [options] <target> <source>
```



<source>	<code>docker://container[:tag]</code>	Build from a Docker container
	<code>container image.sif</code>	Build from a local image file
	<code>container_sandbox/</code>	Build from a local sandbox (A directory form of a container)
	<code>container_recipe.def</code>	Build from a recipe (an instruction script of how to build an image)

d) Build an image from sandbox

```
$ sudo singularity build [options] <target> <source>
```



Modify with “**sudo**”? Build with “**sudo**”!

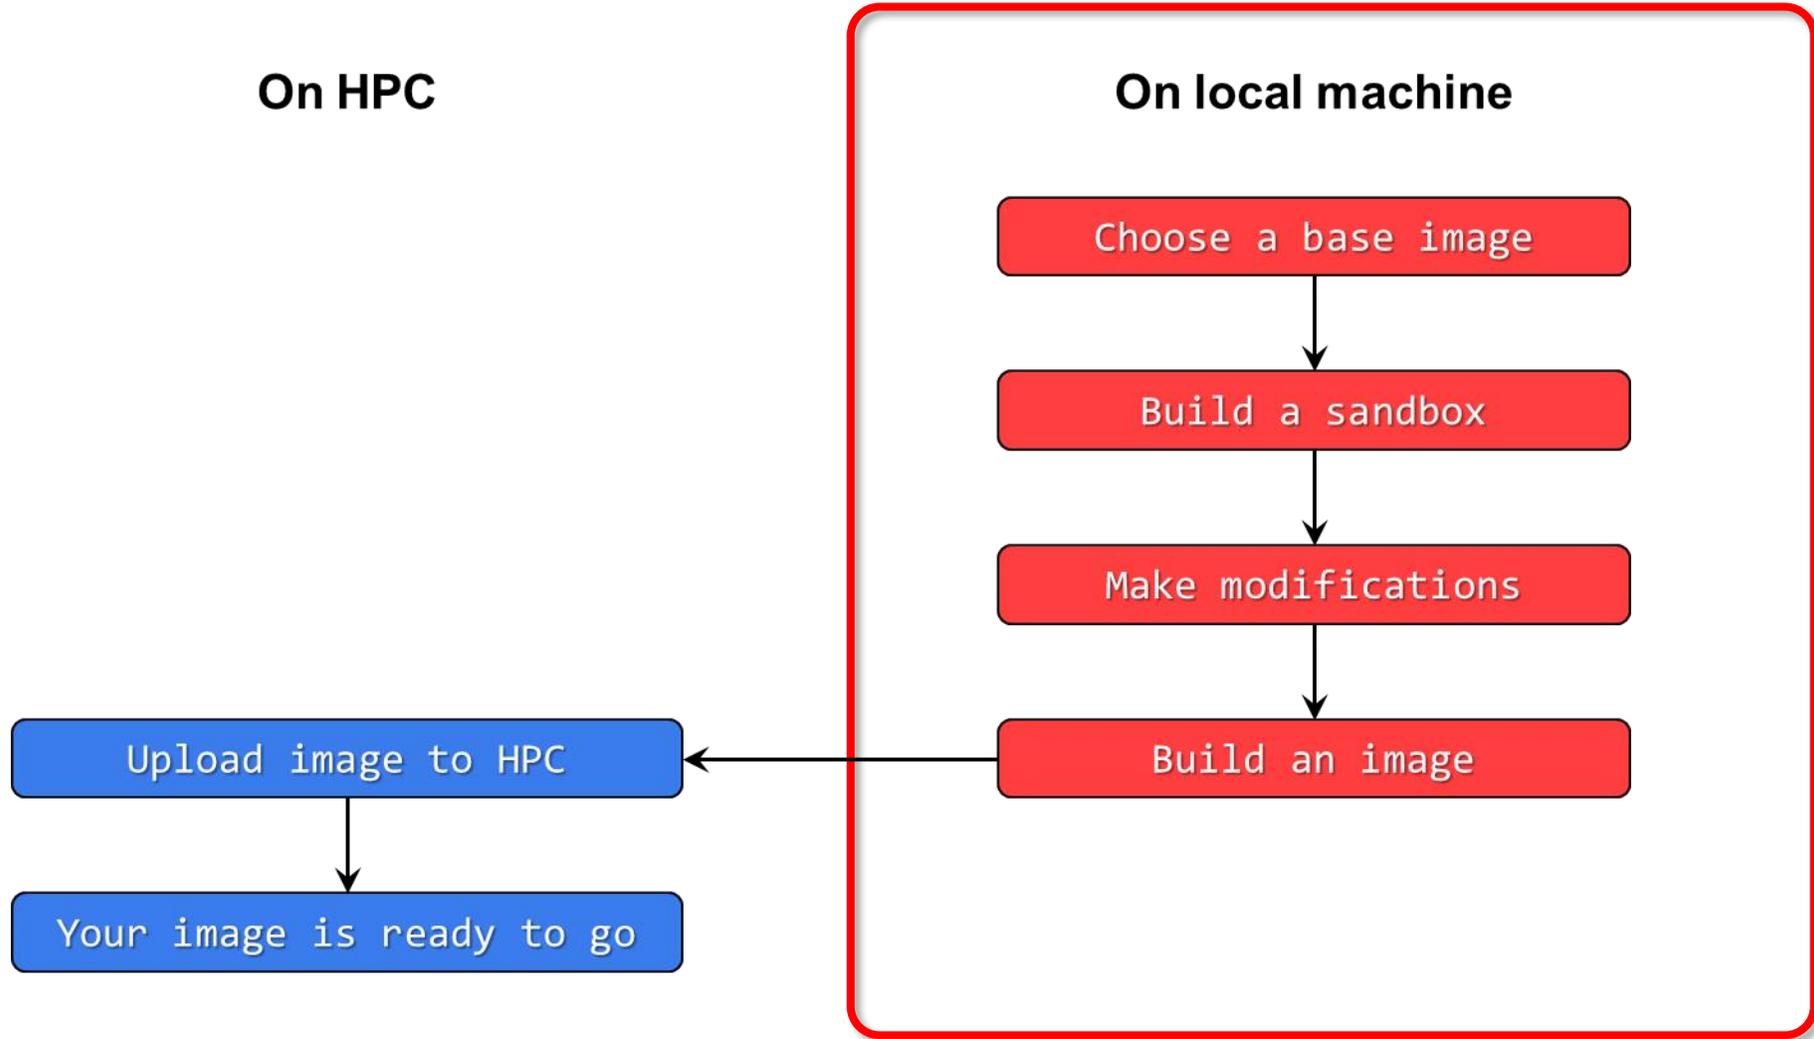
- Quick recap

To ...	You need to ...
Build a sandbox	\$ singularity build --sandbox ...
Modify a sandbox	\$ sudo singularity shell --writable ...
Build an image from sandbox	\$ sudo singularity build ...

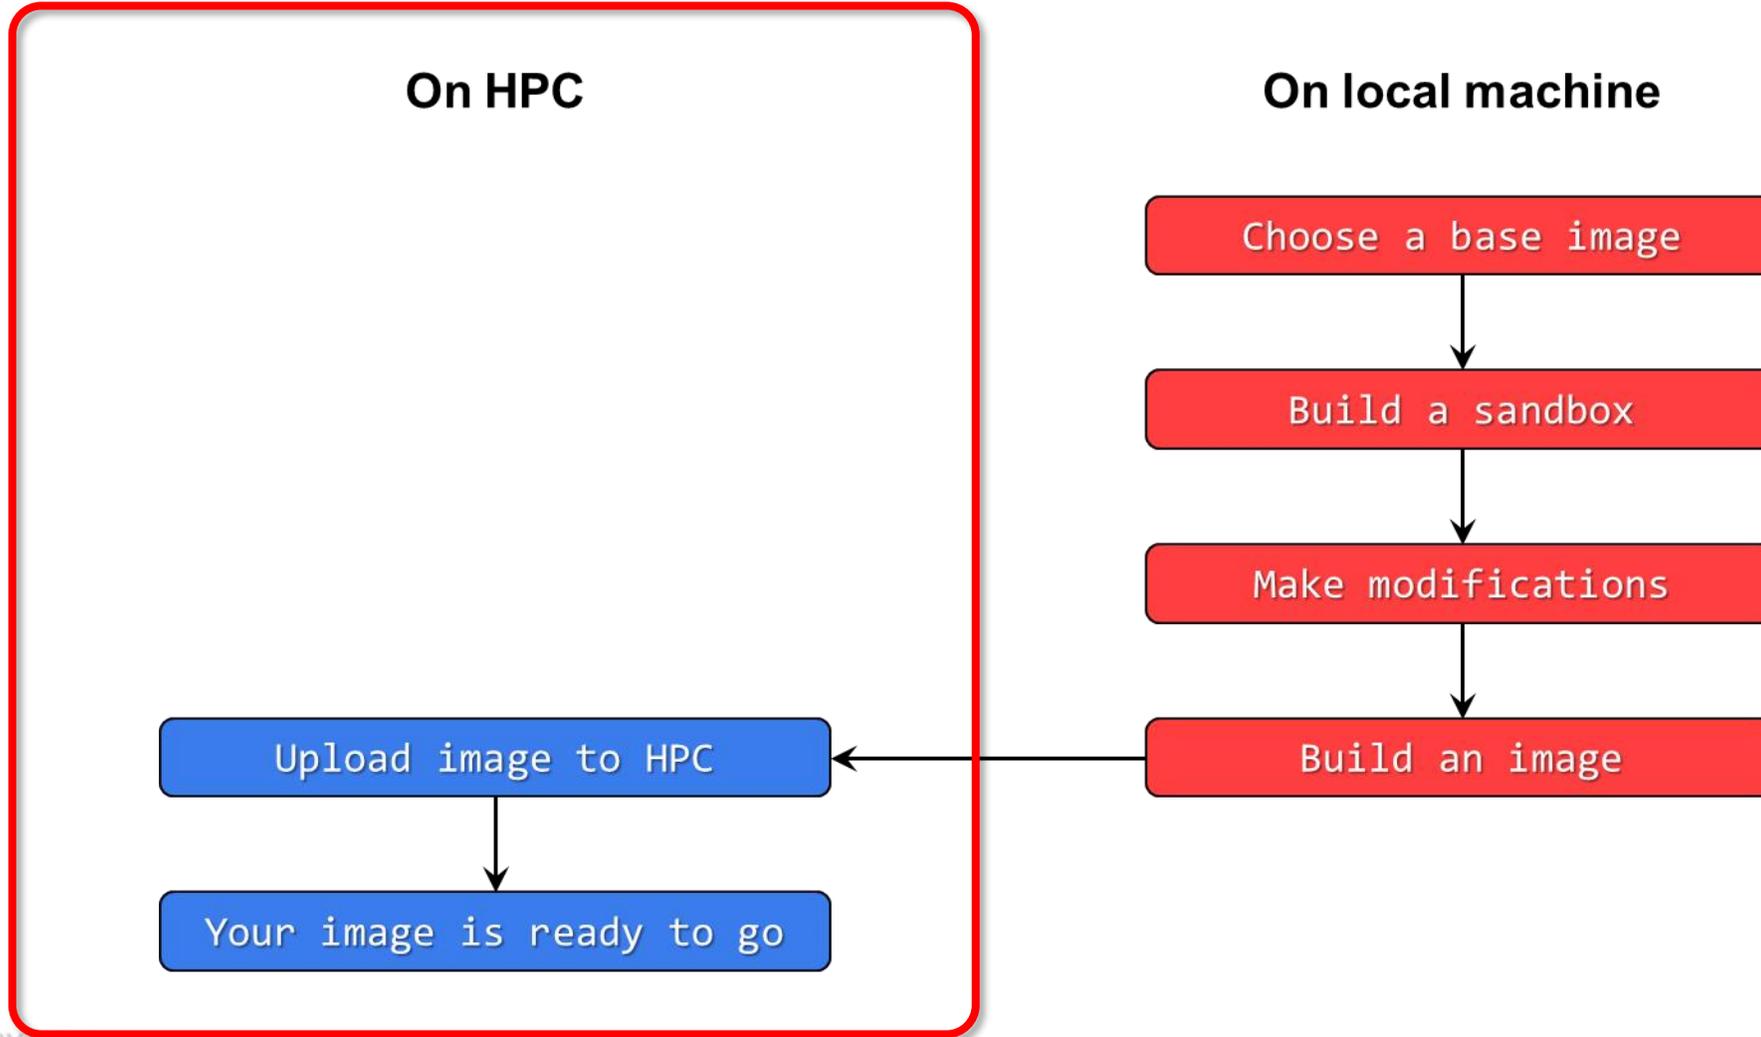
e) Upload image to HPC and run

Now! The moment of truth!

2) Typical workflow



2) Typical workflow



1. Why Container?

- 1) Problems
- 2) Container & Singularity

2. Run an Existing Container Image

- 1) What you need
- 2) Basic commands
- 3) Running jobs with Singularity

3. Get More Container Images

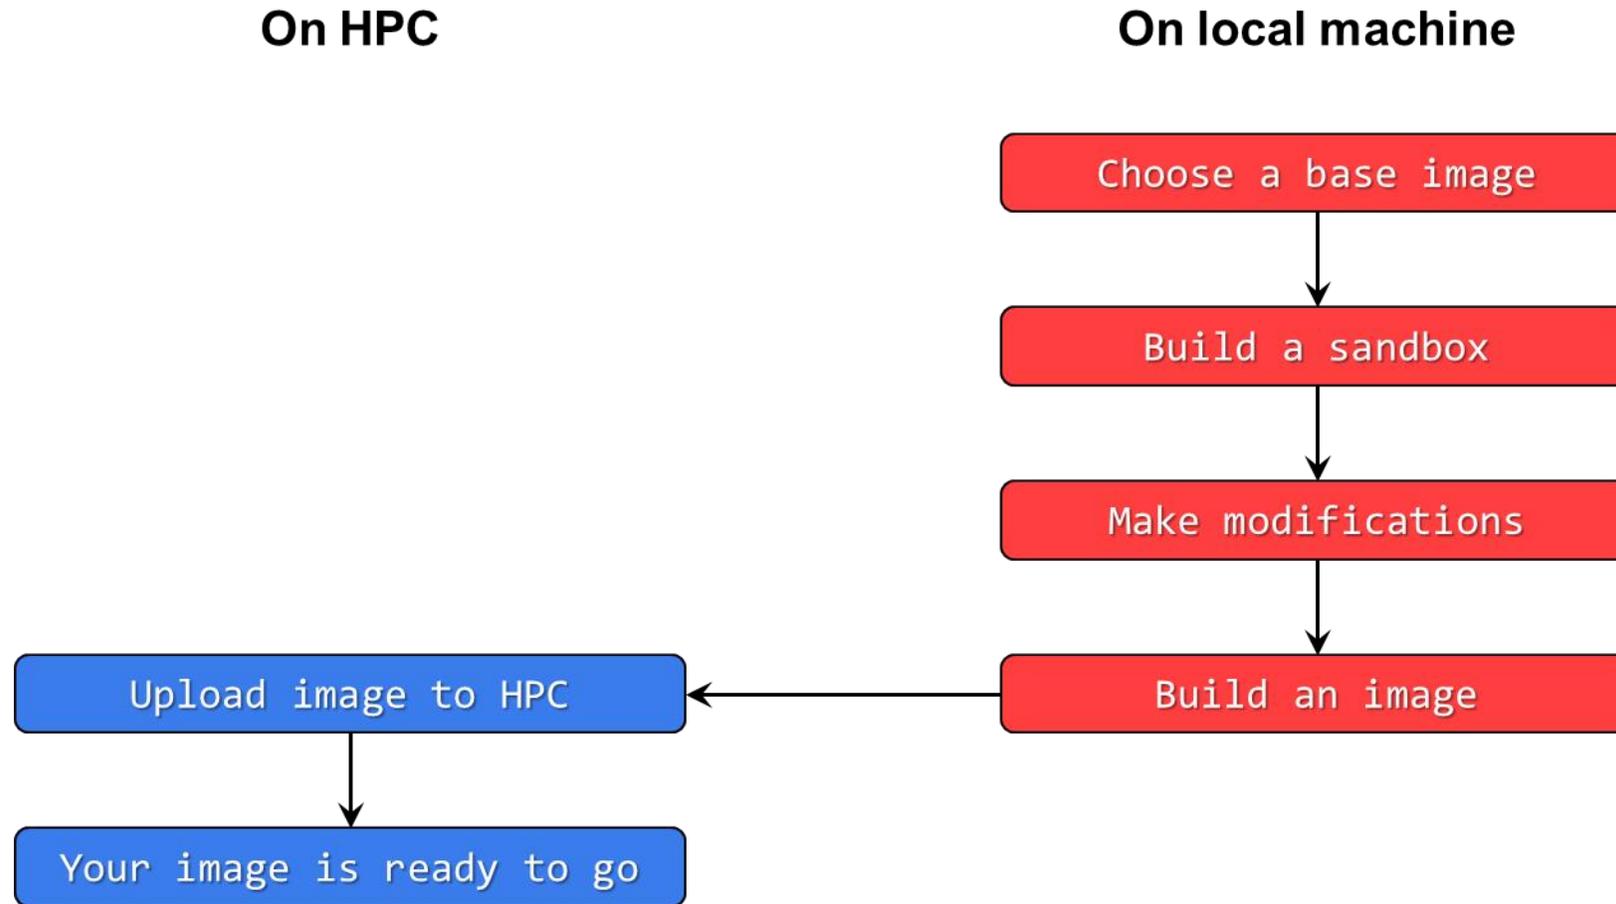
- 1) What you need
- 2) Where to get
- 3) How to get

4. Build Your Own Container Image

- 1) What you need
- 2) Typical workflow
- 3) Make it easier - Recipe

3) Make it easier - Recipe

- Why?



- Why?

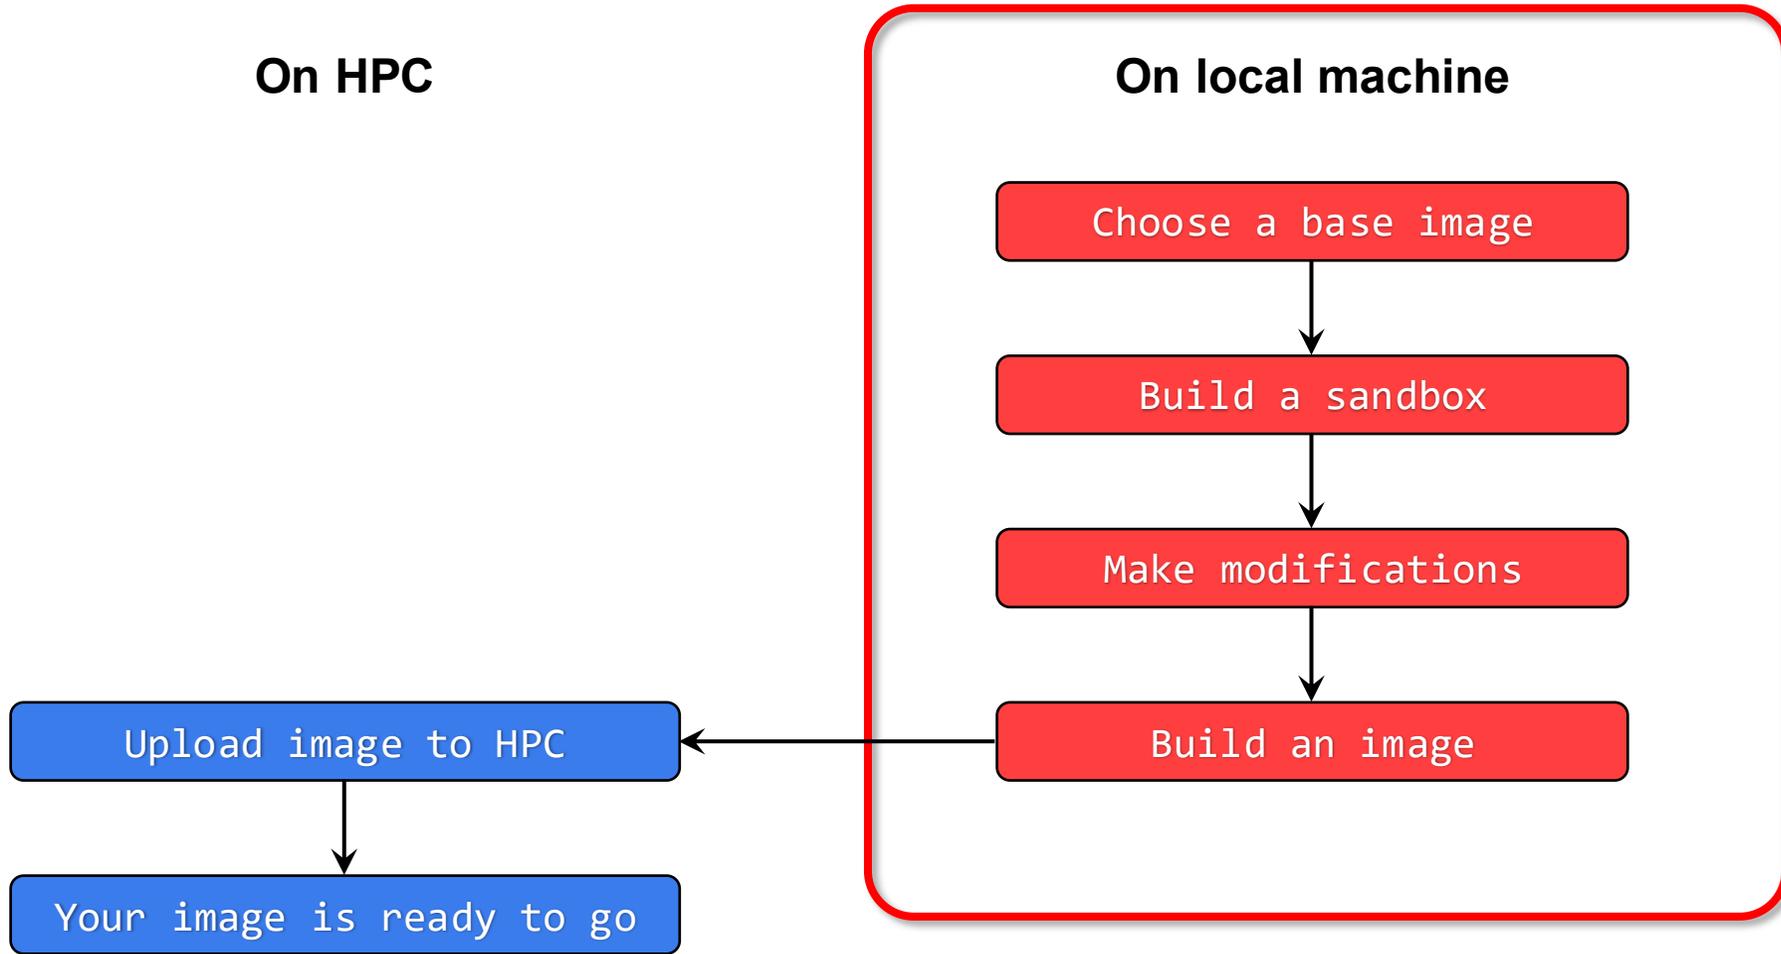
Pros	Cons
<ul style="list-style-type: none">• Flexibility	<ul style="list-style-type: none">• Repeatability• Minimizing image size

- Solution:

- **Recipe**: A text file containing instructions to build a container

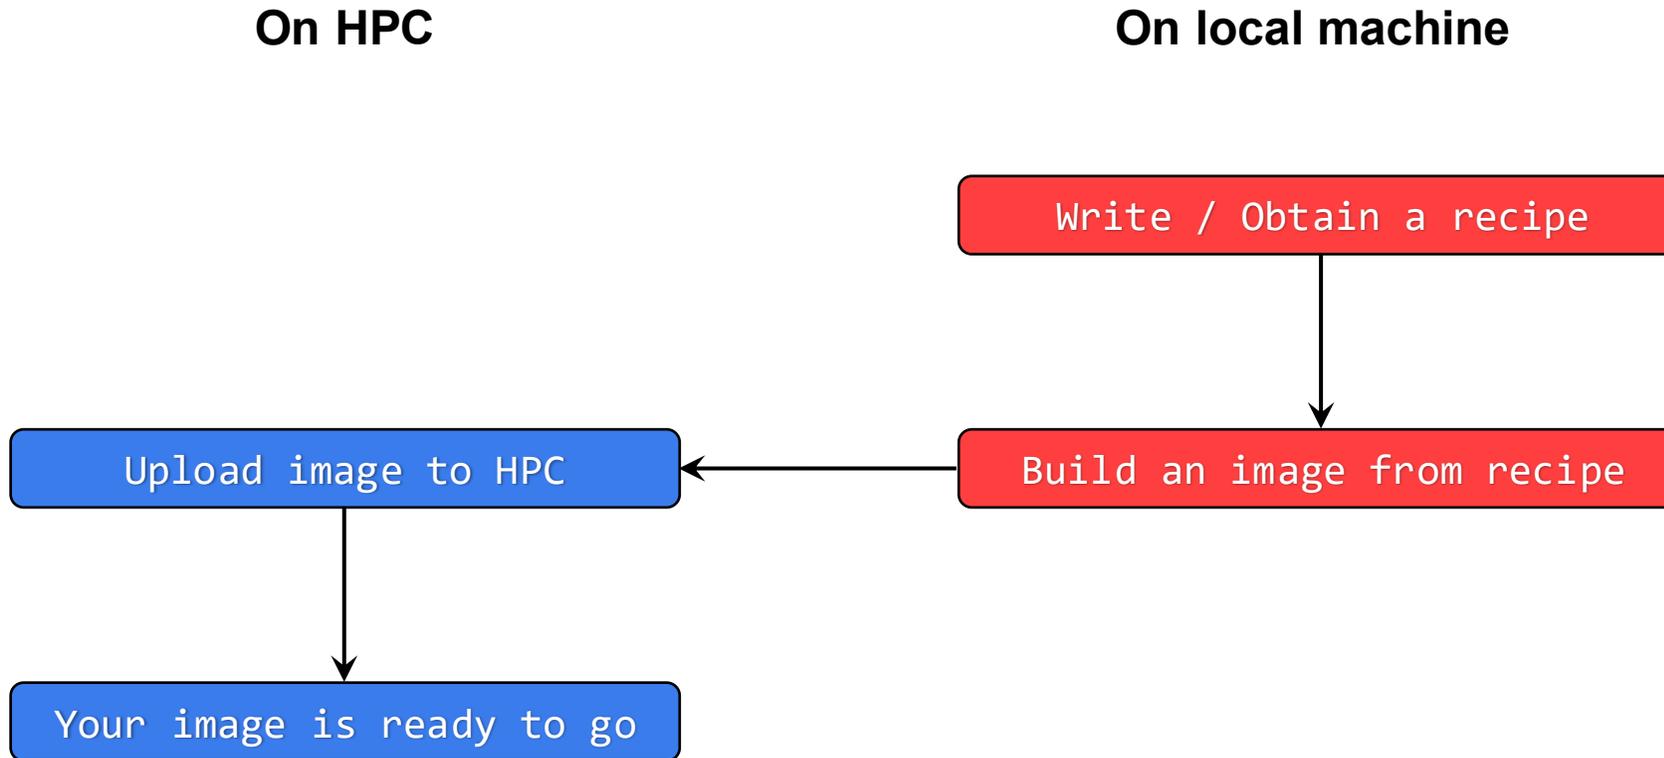
3) Make it easier - Recipe

- Why?



3) Make it easier - Recipe

- Why?



a) What does a recipe look like?

recipe.def

```
Bootstrap: docker
From: ubuntu:24.04
%labels
  Author Siva Prasad Kasetti
  Version 1.0
  Description Install curl
%post
  apt-get update -y
  apt-get install -y curl
%environment
  export MYENV="This is my environment variable."
%runscript
  echo "Print curl version:"
  curl --version
```

a) What does a recipe look like?

```
Bootstrap: docker
From: ubuntu:24.04

%labels
  Author Siva Prasad Kasetti
  Version 1.0
  Description Install curl

%post
  apt-get update -y
  apt-get install -y curl

%environment
  export MYENV="This is my environment variable."

%runscript
  echo "Print curl version:"
  curl --version
```

Header

- Base image info (how, where, what to pull)

3) Make it easier - Recipe

a) What does a recipe look like?

```
Bootstrap: docker
From: ubuntu:24.04

%labels
  Author Siva Prasad Kasetti
  Version 1.0
  Description Install curl

%post
  apt-get update -y
  apt-get install -y curl

%environment
  export MYENV="This is my environment variable."

%runscript
  echo "Print curl version:"
  curl --version
```

Label

- Container information (write whatever you want)

a) What does a recipe look like?

```
Bootstrap: docker
From: ubuntu:24.04
%labels
  Author Siva Prasad Kasetti
  Version 1.0
  Description Install curl
%post
  apt-get update -y
  apt-get install -y curl
%environment
  export MYENV="This is my environment variable."
%runscript
  echo "Print curl version:"
  curl --version
```

Post

- Commands to execute after the base image is pulled

3) Make it easier - Recipe

a) What does a recipe look like?

```
Bootstrap: docker
From: ubuntu:24.04

%labels
  Author Siva Prasad Kasetti
  Version 1.0
  Description Install curl

%post
  apt-get update -y
  apt-get install -y curl

%environment
  export MYENV="This is my environment variable."

%runscript
  echo "Print curl version:"
  curl --version
```

Environment

- Define environmental variables every time the container is executed

3) Make it easier - Recipe

a) What does a recipe look like?

```
Bootstrap: docker
From: ubuntu:24.04

%labels
  Author Siva Prasad Kasetti
  Version 1.0
  Description Install curl

%post
  apt-get update -y
  apt-get install -y curl

%environment
  export MYENV="This is my environment variable."

%runscript
  echo "Print curl version:"
  curl --version
```

Runscript

- Commands to be run with `singularity run`

a) What does a recipe look like?

```
Bootstrap: docker
From: ubuntu:24.04

%labels
  Author Siva Prasad Kasetti
  Version 1.0
  Description Install curl

%post
  apt-get update -y
  apt-get install -y curl

%environment
  export MYENV="This is my environment variable."

%runscript
  echo "Print curl version:"
  curl --version
```

a) What does a recipe look like?

```
Bootstrap: docker
From: ubuntu:24.04

%labels
  Author Siva Prasad Kasetti
  Version 1.0
  Description Install curl

%post
  apt-get update -y
  apt-get install -y curl

%environment
  export MYENV="This is my environment variable."

%runscript
  echo "Print curl version:"
  curl --version
```

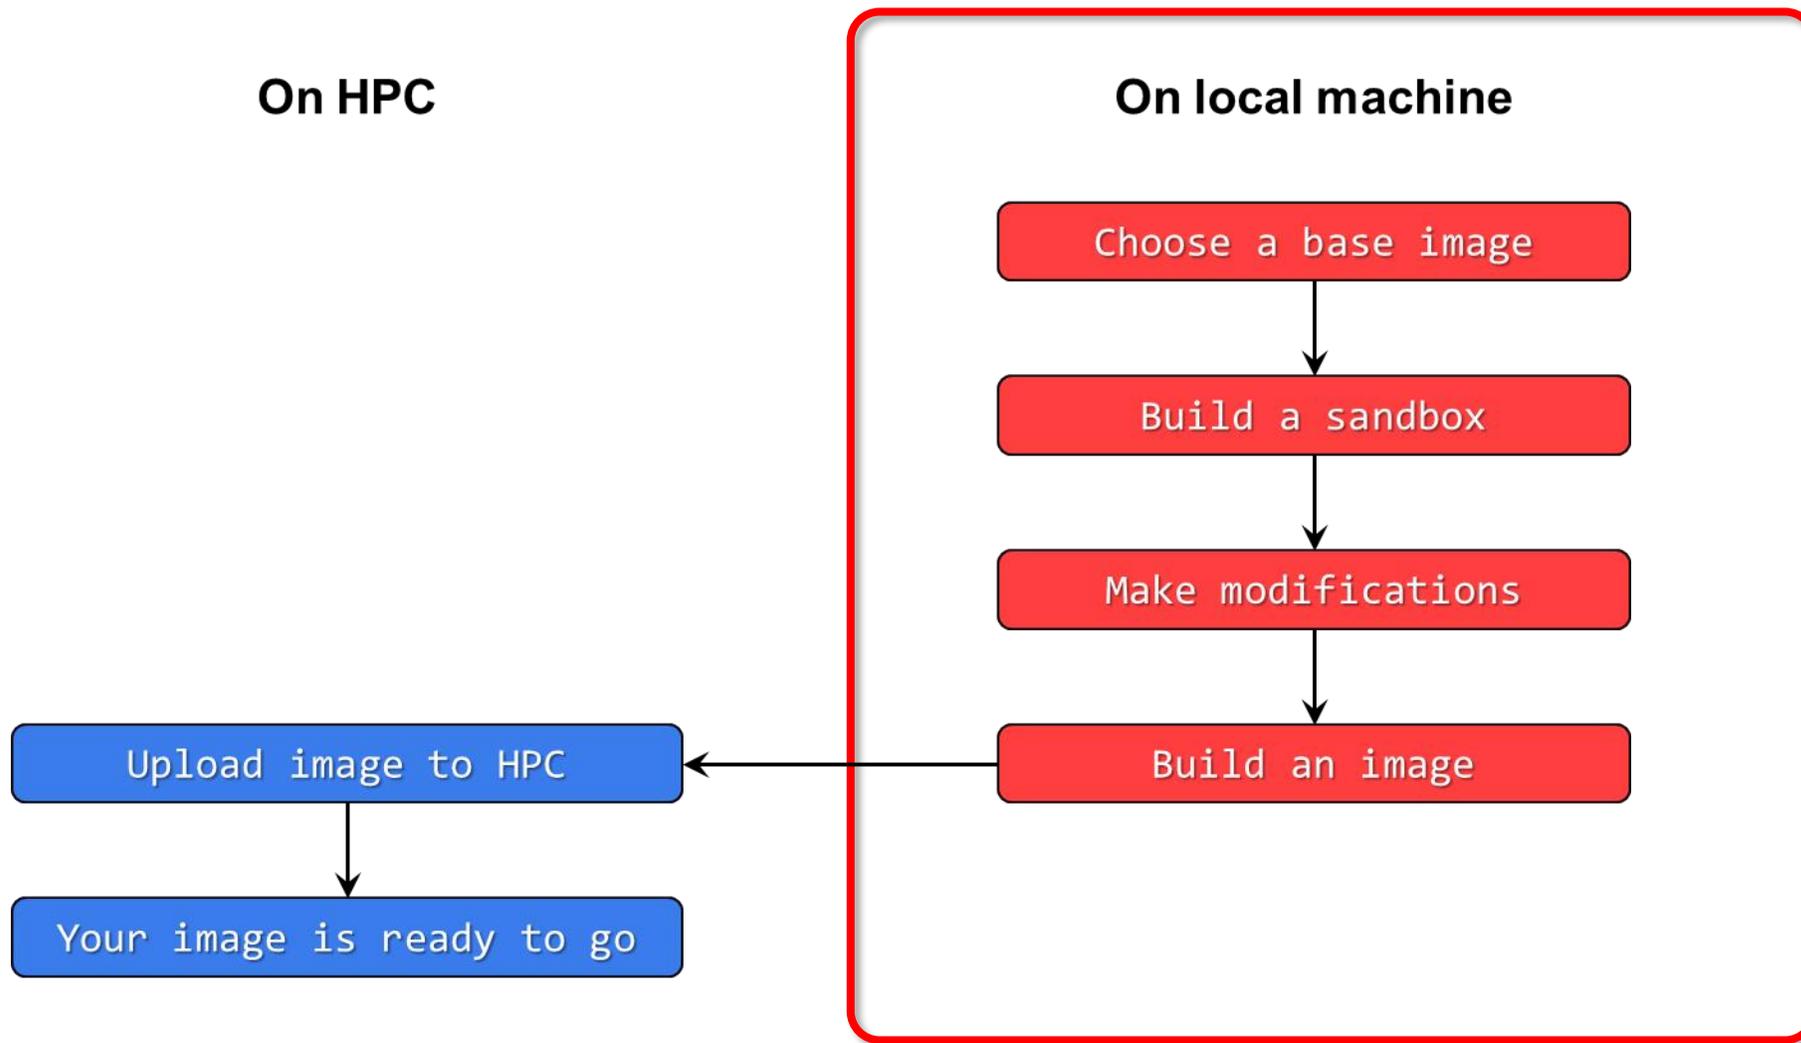
3) Make it easier - Recipe

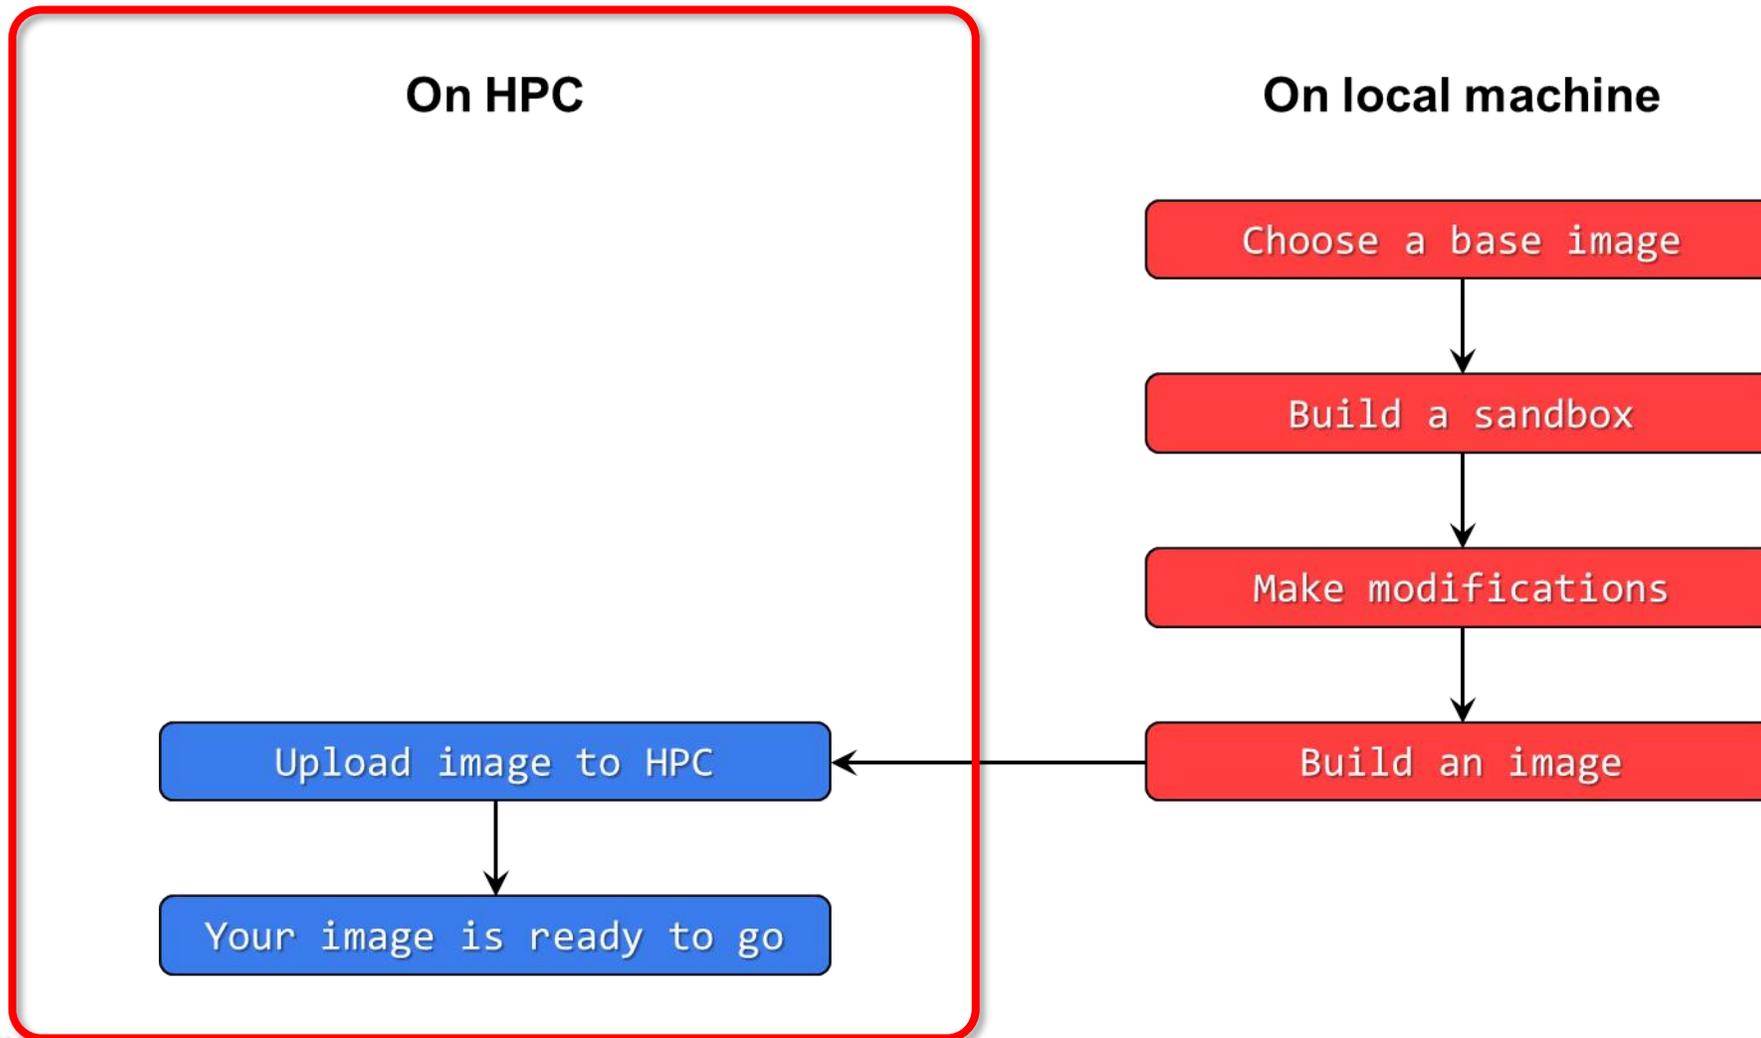
b) Build the recipe

```
$ singularity build [options] <target> <source>
```



<source>	<code>docker://container[:tag]</code>	Build from a Docker container
	<code>container_image.sif</code>	Build from a local image file
	<code>container_sandbox/</code>	Build from a local sandbox (A directory form of a container)
	<code>container_recipe.def</code>	Build from a recipe (an instruction script of how to build an image)





Conclusion

1. Why Container?

- 1) Problems
- 2) Container & Singularity

2. Run an Existing Container Image

- 1) What you need
- 2) Basic commands
- 3) Running jobs with Singularity

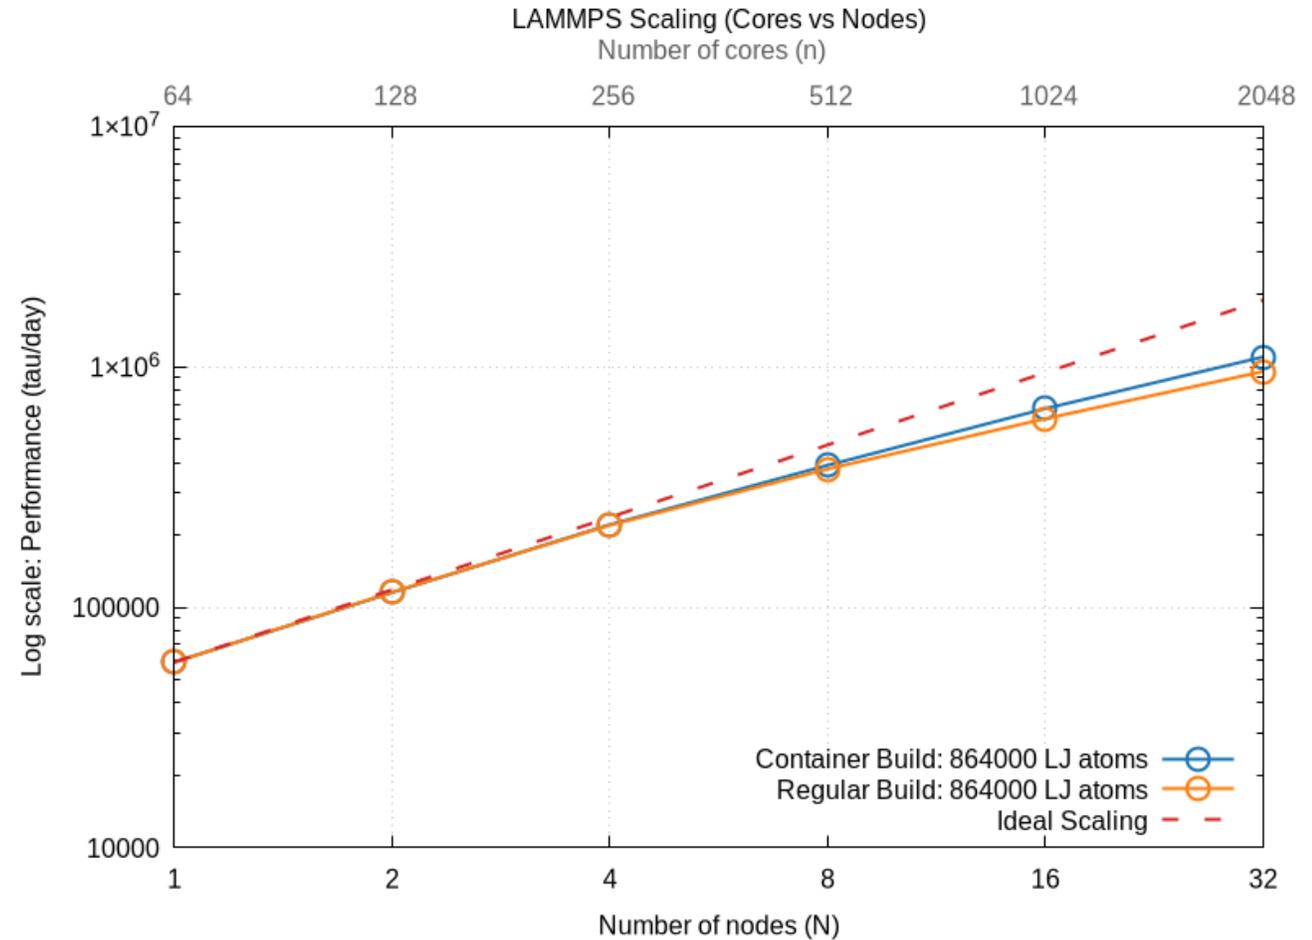
3. Get More Container Images

- 1) What you need
- 2) Where to get
- 3) How to get

4. Build Your Own Container Image

- 1) What you need
- 2) Typical workflow
- 3) Make it easier - Recipe

- ❖ The MD simulation package **LAMMPS** was used as the benchmark.
- ❖ Performance was compared between running LAMMPS directly from source and running it inside a container.
- ❖ Results showed identical performance between both approaches.



- **Contact user services**

- Email Help Ticket: sys-help@loni.org
- Telephone Help Desk: +1 (225) 578-0900

- Are you tired of writing the long, tedious singularity commands?

```
$ singularity exec --nv -B /work,/project,/usr/local/package \  
/home/admin/singularity/ubuntu-training.sif \  
python helloworld.py
```



- Try **SIMPLE-MOD** !

- <https://github.com/lshpche/p/SIMPLE-MOD>
- A GUI tool to create module key from container-based software.
- Using the software in containers is as easy as:

```
$ module load busco
$ busco --version
BUSCO 5.6.1
```

