Part four Tips and tricks to optimize container creation.

Nicola Malavasi - Marie Curie Fellow @ MPE - 16/04/2024









Main problems of image creation

nmalavasi@ga-	-lt7982 tutorial_n	naterial % docke	er image ls	
REPOSITORY	TAG	IMAGE ID	CREATED	SIZE
debian	latest	5c2e61c12a03	39 hours ago	139MB
python	3.11.9-bookworm	c9a42713a959	8 days ago	1.01G
ubuntu	latest	2b7cc08dcdbb	6 weeks ago	69.2M
nmalavasi@ga-	-lt7982 tutorial_n	naterial %		

- The two main bottlenecks of image creation from Dockerfile are the time required to build the image and the memory size of the image.
 - Memory size can be up to a few GB.
 - Both are dependent on what is installed inside the image at creation.





Consider the following Dockerfile. End result is python installation with a few packages.

Users > nmalavasi > Desktop > PUNCH_useful 3

syntax=docker/dockerfile:1 1 2 3 **#Start from python image** FROM python:3.11.9-bookworm 4 5 6 #Use pip to install matplotlib RUN pip install matplotlib 7 8 9 RUN pip install astropy 10 RUN pip install numpy 11 12 RUN pip install scipy 13





- Consider the following Dockerfile. End result is python installation with a few packages.
 - python image
 - python version
 - OS version



Consider the following Dockerfile. End result is python installation with a few packages.

Users > nmalavasi > Desktop > PUNCH_useful 3

syntax=docker/dockerfile:1 1 2 3 **#Start from python image** 4 FROM python:3.11.9-bookworm 5 6 #Use pip to install matplotlib RUN pip install matplotlib 7 8 9 RUN pip install astropy 10 RUN pip install numpy 11 12 RUN pip install scipy 13

Modules installed sequentially



Users > nmalavasi > Desktop > PUNCH_useful 3

syntax=docker/dockerfile:1 1 2 #Start from python image 3 FROM python:3.11.9-bookworm 4 5 6 #Use pip to install matplotlib RUN pip install matplotlib 7 8 RUN pip install astropy 9 10 RUN pip install numpy 11 12 13 RUN pip install scipy

Consider the following Dockerfile. End result is python installation with a few packages.

Final size is about 1.5 GB

 Python is installed on top of full OS installation: we don't need all of this if we are just going to use python.

• Each RUN command is a layer (i.e. an additional intermediate step in image creation). Each layer takes memory.







Example Package installation in Python Compare with the following. End result is same as before.

syntax=docker/dockerfile:1

#Start from python image FROM python:3.11.9-slim-bookworm

#Use pip to install matplotlib RUN pip install matplotlib astropy numpy scipy

RUN python -m pip cache purge

- We use "slim" version of OS: only necessary packages for python installation.
- Packages are installed in same command: only one layer.
- Clean pip cache: once packages are installed they can be used, no need to keep them in cache (installation won't be repeated within same container use).

Example Package installation in Python Compare with the following. End result is same as before.

syntax=docker/dockerfile:1

#Start from python image FROM python:3.11.9 slim-bookworm

#Use pip to install matplotlib RUN pip install matplotlib astropy numpy scipy

RUN python -m pip cache purge

- We use "slim" version of OS: only necessary packages for python installation.
- Packages are installed in same command: only one layer.
- Clean pip cache: once packages are installed they can be used, no need to keep them in cache (installation won't be repeated within same container use).

Example **Comparison of light and heavy images**

The difference in memory use can be large.

[nmalavasi@ga-lt7982 tutorial_ma REPOSITORY TAG IMAGE python_small latest 724bbb python_large latest 82da34 nmalavasi@ga-lt7982 tutorial_mat

terial	%	doo	cker :	image	ls
ID		CRI	EATED		SIZE
cac001		16	hour	s ago	612MB
5bcbe4		16	hour	s ago	1.47GB
terial	%				

Common practices to optimize containers Essentially: memory use reduction

- Be mindful of installed code: install only needed libraries, reduce what you install to the bare minimum needed.
- Number of layers increases size: bundle commands when possible (e.g. aptget update, install, clean).
- Choose correct starting image: use "slim" images if you don't need full OS functionality.
- Clean up after installation: clear apt-get and pip caches, remove tarballs.

Multi-stage builds

Image 1 **Install libraries and** software.

> Bring over only the software needed to install software in part 2.

Image 2 Install further software.

Only image 3 is kept, others are deleted.

It is possible to use an image as a base for further image creation also within the same Dockerfile.



Image 3 Install further software.





```
Users > nmalavasi > Desktop > arcane_dockerfiles > OTF_pointing_correction_container > 🗇 Dockerfile > ...
       # syntax=docker/dockerfile:1
  2
       #Builder image: lighter to download, we use it ot install some stuff- python image
  3
       FROM python:3.8.16-slim-bullseye as builder
  -4
  5
  6
       command, git, wget, xz-utils, pip3
       RUN apt-get update && apt-get -y install wget xz-utils git time
  8
       #Download CASA tarball
  9
 10
 11
       #Extract in home directory
 12
       RUN tar -xvf casa-6.5.2-26-py3.8.tar.xz
 13
 14
       #Make directory for python packages
 15
       RUN mkdir requirements_folder
 16
 17
       #Git clone arcane suite
 18
       RUN git clone https://github.com/rstofi/arcane_suite.git
 19
 20
       #Change folder to the arcane_suite
 21
       WORKDIR /arcane_suite
 22
 23
       #Set specific commit
 24
       RUN git checkout d7d394fc03375b3efa7cc2711834462c812473ad
 25
 26
       #Install all requirements
 27
 28
       RUN pip3 install ---target=/requirements_folder -r ./requirements.txt
 29
```

#Update apt-get and install useful packages: python-is-python3 to symlink python command to python3

RUN wget https://casa.nrao.edu/download/distro/casa/release/rhel/casa-6.5.2-26-py3.8.tar.xz

#Use kernsuite image as it contains python-casacore and chgcentre FROM kernsuite/base:7 as runner

#Install python-casacore and chgcentre RUN docker-apt-install python3-casacore chgcentre

#Copy what previously installed

COPY --- from=builder /requirements_folder /usr/local/src/ COPY --- from=builder /requirements_folder/bin /usr/bin/ COPY --- from=builder /casa-6.5.2-26-py3.8 /usr/local/src/casa_from_container/ COPY --- from=builder /arcane_suite /usr/local/src/arcane_suite/

```
command pip3
```

#Update PYTHONPATH ENV PYTHONPATH=\$PYTHONPATH:/usr/local/src

#Create symbolic link to point to the correct python RUN ln -s /usr/bin/python3.8 /usr/local/bin/python

```
#Create symbolic link for casa
RUN ln -s /usr/local/src/casa_from_container/bin/casa /usr/local/bin/casa
```

#Change folder to the arcane_suite WORKDIR /usr/local/src/arcane_suite

#Install the arcane arcane_suite RUN pip3 install -e ./

#Update apt-get and install useful packages: python-is-python3 to symlink python command to python3

RUN apt-get update && apt-get -y install python-is-python3 graphviz python3-pydot python3-pip

Part five Apptainer. **Container use in C4P.**

Nicola Malavasi - Marie Curie Fellow @ MPE - 16/04/2024











Apptainer

- Docker is not the only container engine. Another widely used is Apptainer.
 - Effectively available only on linux.
- Same principle as Docker: can pull/create images, has own version of Dockerfile and syntax, runs processes in isolated way.



ΑΡΡΤΑΙΝΕR





Integration to host

Integration to host

• Process is run as current user. USER keyword of Dockerfile is ignored.

Complete isolation from host

• User inside container is completely independent of user outside.

Integration to host

- Process is run as current user. USER keyword of Dockerfile is ignored.
- Inside container host file system is still visible (less need for bind mounts).

- User inside container is completely independent of user outside.
- Filesystem inside and outside container are disconnected. Strong need for bind mounts.

Integration to host

- Process is run as current user. USER keyword of Dockerfile is ignored.
- Inside container host file system is still visible (less need for bind mounts).
- Image saved as .sif file on host system, easy to transfer.

- User inside container is completely independent of user outside.
- Filesystem inside and outside container are disconnected. Strong need for bind mounts.
- Image saved in non user-friendly place and way.

Integration to host

- Process is run as current user. USER keyword of Dockerfile is ignored.
- Inside container host file system is still visible (less need for bind mounts).
- Image saved as .sif file on host system, easy to transfer.
- No root privilege is needed to create image.

- User inside container is completely independent of user outside.
- Filesystem inside and outside container are disconnected. Strong need for bind mounts.
- Image saved in non user-friendly place and way.
- Root privilege is needed to create image.

- Apptainer can run containers in Docker images
- Apptainer can create images from Dockerfiles

- 1. Write Dockerfile
- 2. Create image from Dockerfile with Docker
- 3. Have Apptainer read image from Docker daemon and convert it to .sif file
- 4. Run Apptainer container in Apptainer image

Only the first aspect is needed for C4P.

- 1. Write Dockerfile
- 2. Create image from Dockerfile with Docker
- 3. Have Apptainer read image from Docker daemon and convert it to .sif file
- 4. Run Apptainer container in Apptainer image apptainer run apptainer exec

apptainer build filename.sif docker-daemon:image_name:image_tag





- 1. Write Dockerfile
- 2. Create image from Dockerfile with Docker
- 3. Have Apptainer read image from Docker daemon and convert it to .sif file
- 4. Run Apptainer container in Apptainer image apptainer run apptainer exec

Same as Docker

apptainer build filename.sif docker-daemon:image_name:image_tag





- 1. Write Dockerfile
- 2. Create image from Dockerfile with Docker
- 3. Have Apptainer read image from Docker daemon and convert it to .sif file
- 4. Run Apptainer container in Apptainer image apptainer run apptainer exec







- 1. Write Dockerfile
- 2. Create image from Dockerfile with Docker
- 3. Have Apptainer read image from Docker daemon and convert it to .sif file
- 4. Run Apptainer container in Apptainer image apptainer run apptainer exec



- apptainer build filename.sif docker-daemon:image_name:image_tag
 - Image is currently sitting in the Docker image list





- 1. Write Dockerfile
- 2. Create image from Dockerfile with Docker
- 3. Have Apptainer read image from Docker daemon and convert it to .sif file
- 4. Run Apptainer container in Apptainer image apptainer run apptainer exec





Container use in C4P

- Process is automated, only needed input is Dockerfile.





Compute4PUNCH builds images via Docker and runs containers in them via Apptainer.

Add Dockerfile to C4P container stack repository. **Trigger image creation.**

Add image name to CVMFS image list. Jobs are run in container created from image.



Container use in C4P

Write Dockerfile

- Write Dockerfile, test image creation locally with Docker.
- Make sure it works in Apptainer too (isolation vs integration...).



Add Dockerfile to C4P container stack repository. **Trigger image creation.**

 Fork correct repository.

Add Dockerfile.

 Make sure CI/ CD pipeline runs to completion.

 Feedback from C4P developers! Add image name to CVMFS image list. Jobs are run in container created from image.

- Be mindful of resources used, keep container size small and simple.
- After the image is in CVMFS is very difficult to change it: test before you go through!





From C4P section in PUNCH website

Create your own Singularity Container

Currently the best way to get your software environment running on Compute4PUNCH is to build your own singularity container with your favorite OS and containing your software installation.

You have to do that in the following way.

- Fork <u>https://gitlab-p4n.aip.de/compute4punch/container-stacks</u> into your account
- Create a feature branch and put your Dockerfile and stuff into a meaningful sub-directory
- Submit a merge request, once the CI pipeline runs successful
- Let the Compute4PUNCH team know, when it is ready to be reviewed •

Afterwards the new container, which has been to automatically uploaded to the registry at AIP has to be put into https://github.com/cvmfs/images-<u>unpacked.cern.ch</u> via a pull request.

The recipe.yaml inside this repository contains already several available C4P containers. Once merged it will take sometime (~1 hour) until it will appear in

/cvmfs/unpacked.cern.ch/gitlab-p4n.aip.de\:5005/compute4punch/container-stacks/ on the login node and you can submit jobs.