THE UNIVERSITY OF QUEENSLAND
AUSTRALIA
CREATE CHANGE

Behind the scenes of NeuroDesk: What we learned so far in building and distributing containers in a community driven fashion

Steffen Bollmann

The University of Queensland

Acknowledgement of **Country**

The University of Queensland (UQ) acknowledges the Traditional Owners and their custodianship of the lands on which we meet.

We pay our respects to their Ancestors and their descendants, who continue cultural and spiritual connections to Country.

We recognise their valuable contributions to Australian and global society.

Acknowledgements

Australian Research Container Orchestration Service (ARCOS)

nectar

National Imaging Facility

Australian Brain Alliance | EMCR Network

OHBM 2020
A Virtual Experience for Engaging Minds & Empowering Brain Science

ORACLE Cloud

ARDC Australian Research Data Commons

Martin Grignard

Peter Winn

Jakub Kaczmarzyk

Niruhan Viswarupan

Dao Thanh Thuy

Aswin Narayanan

Yorguin Jose Mantilla

Will Woods

Steffen Bollmann

Tom Johnstone

Angela Renton

David White

Thomas Shaw

Oren Civier

Megan Campbell

Paris Lyons

Ryan Sullivan

Collaborators:
- Thomas Close, USyd
- Wojtek Goscinski, Monash/NIF
- Tony Hannan, Florey
- Gary Egan, Monash
- Linda Richards, UQ
- Paul Sowman, Macquarie
- Marta Garrido, UoM
- Patrick Johnston, QUT
- Aina Puce, Indiana University
- Franco Pestilli, Indiana University
- Levin Kuhlmann, Monash
- Gershon Spitz, Monash
- David Abbott, Florey
- Megan Campbell, Newcastle
- Nigel Rogasch, University of Adelaide
- Satrajit Ghosh, MIT

Why Neurodesk?

https://youtu.be/2ATgTOsiGdY

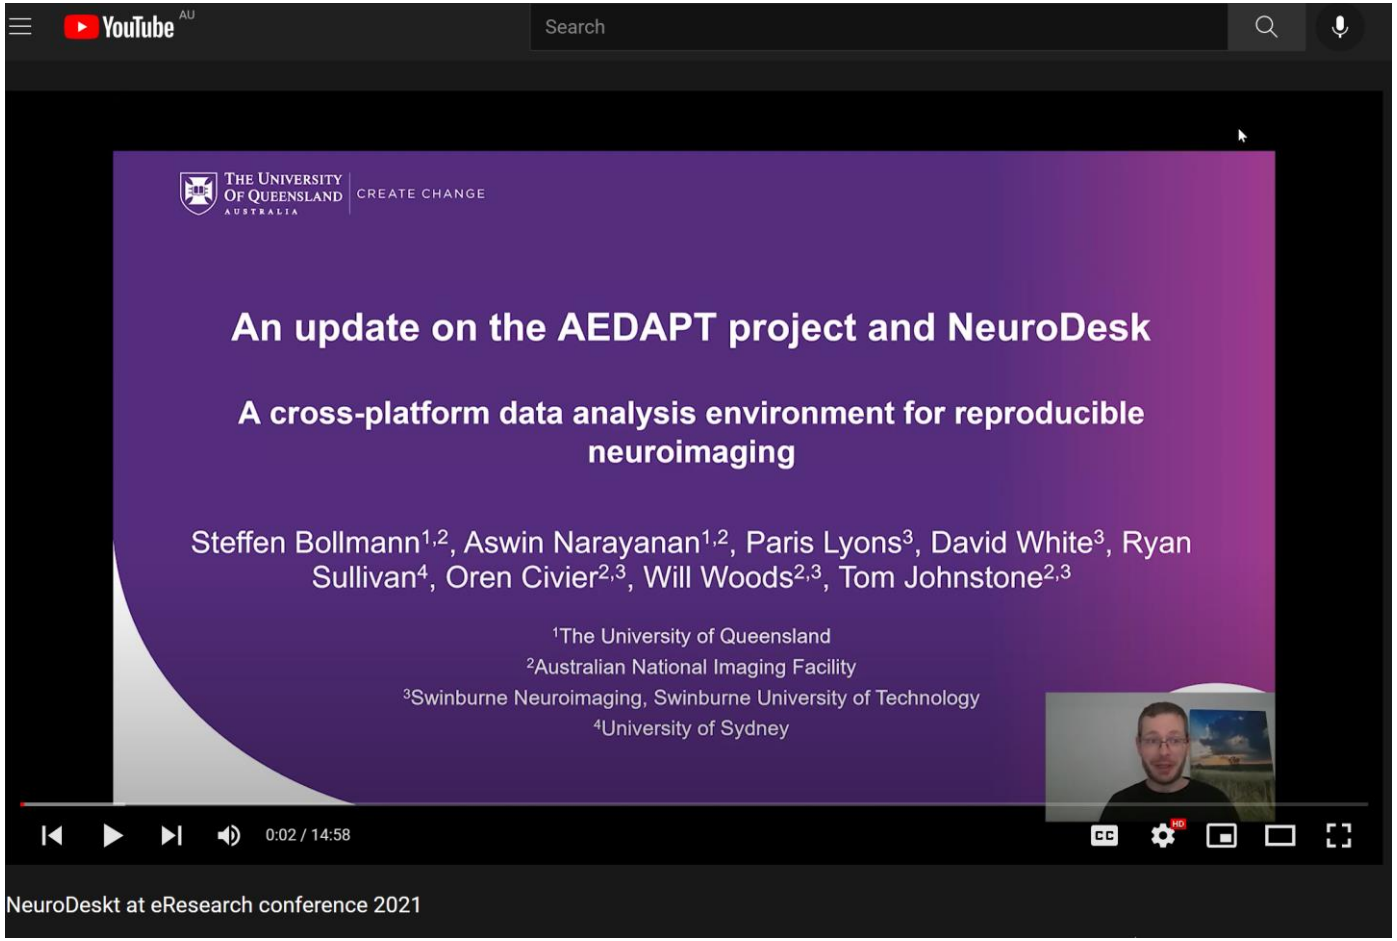

Why Neurodesk - in a Nutshell

- Most neuro imaging tools require Linux
- Tools are not available in standard package systems
- Compiling from source often a nightmare
- Conflicting dependencies
- Reinstalling tools on different computing platforms takes time
- Differing results between software versions

Goal of talk today

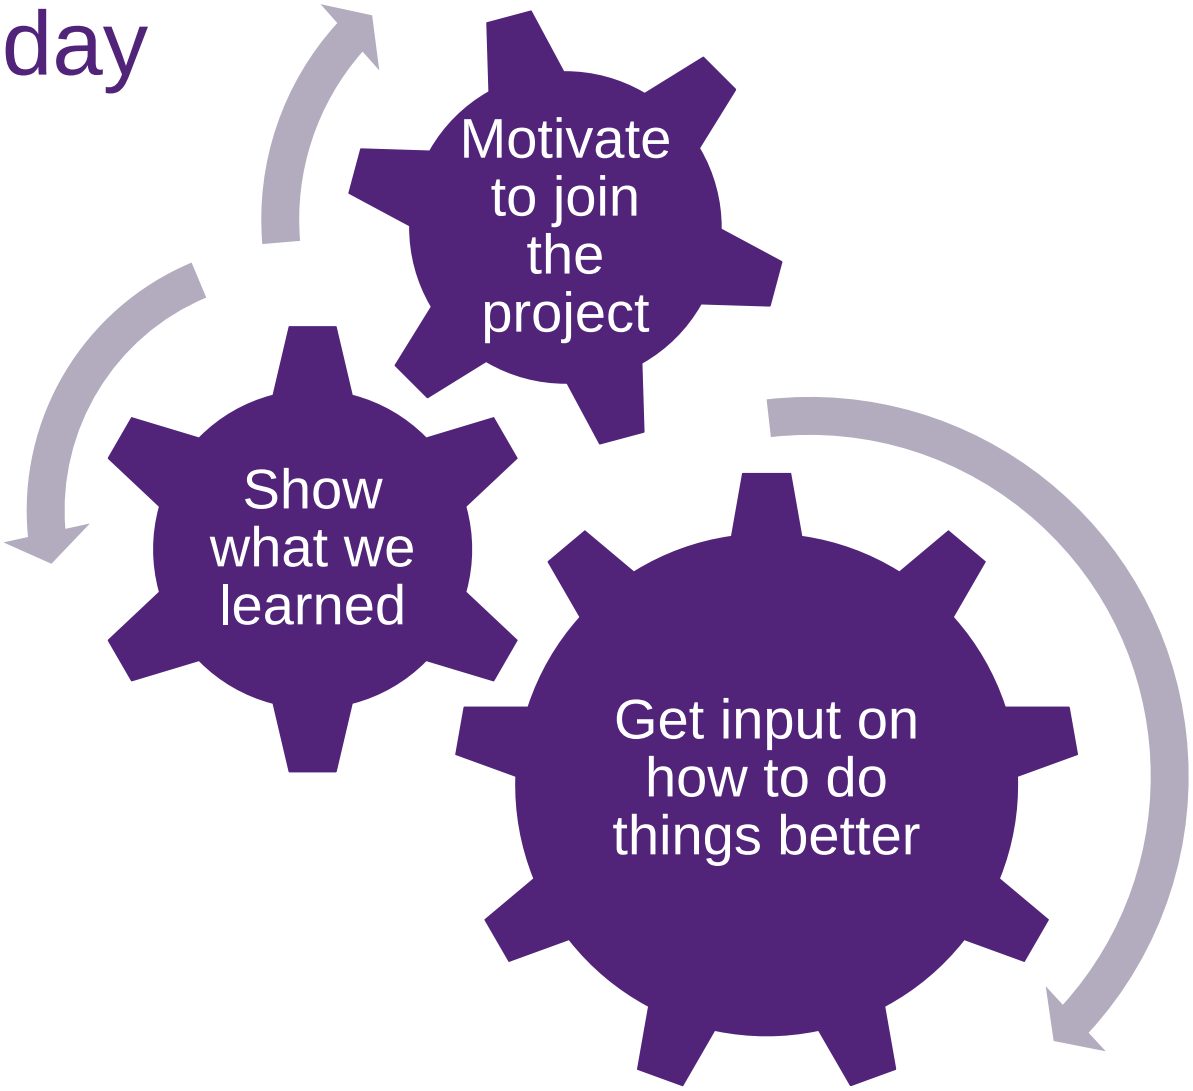

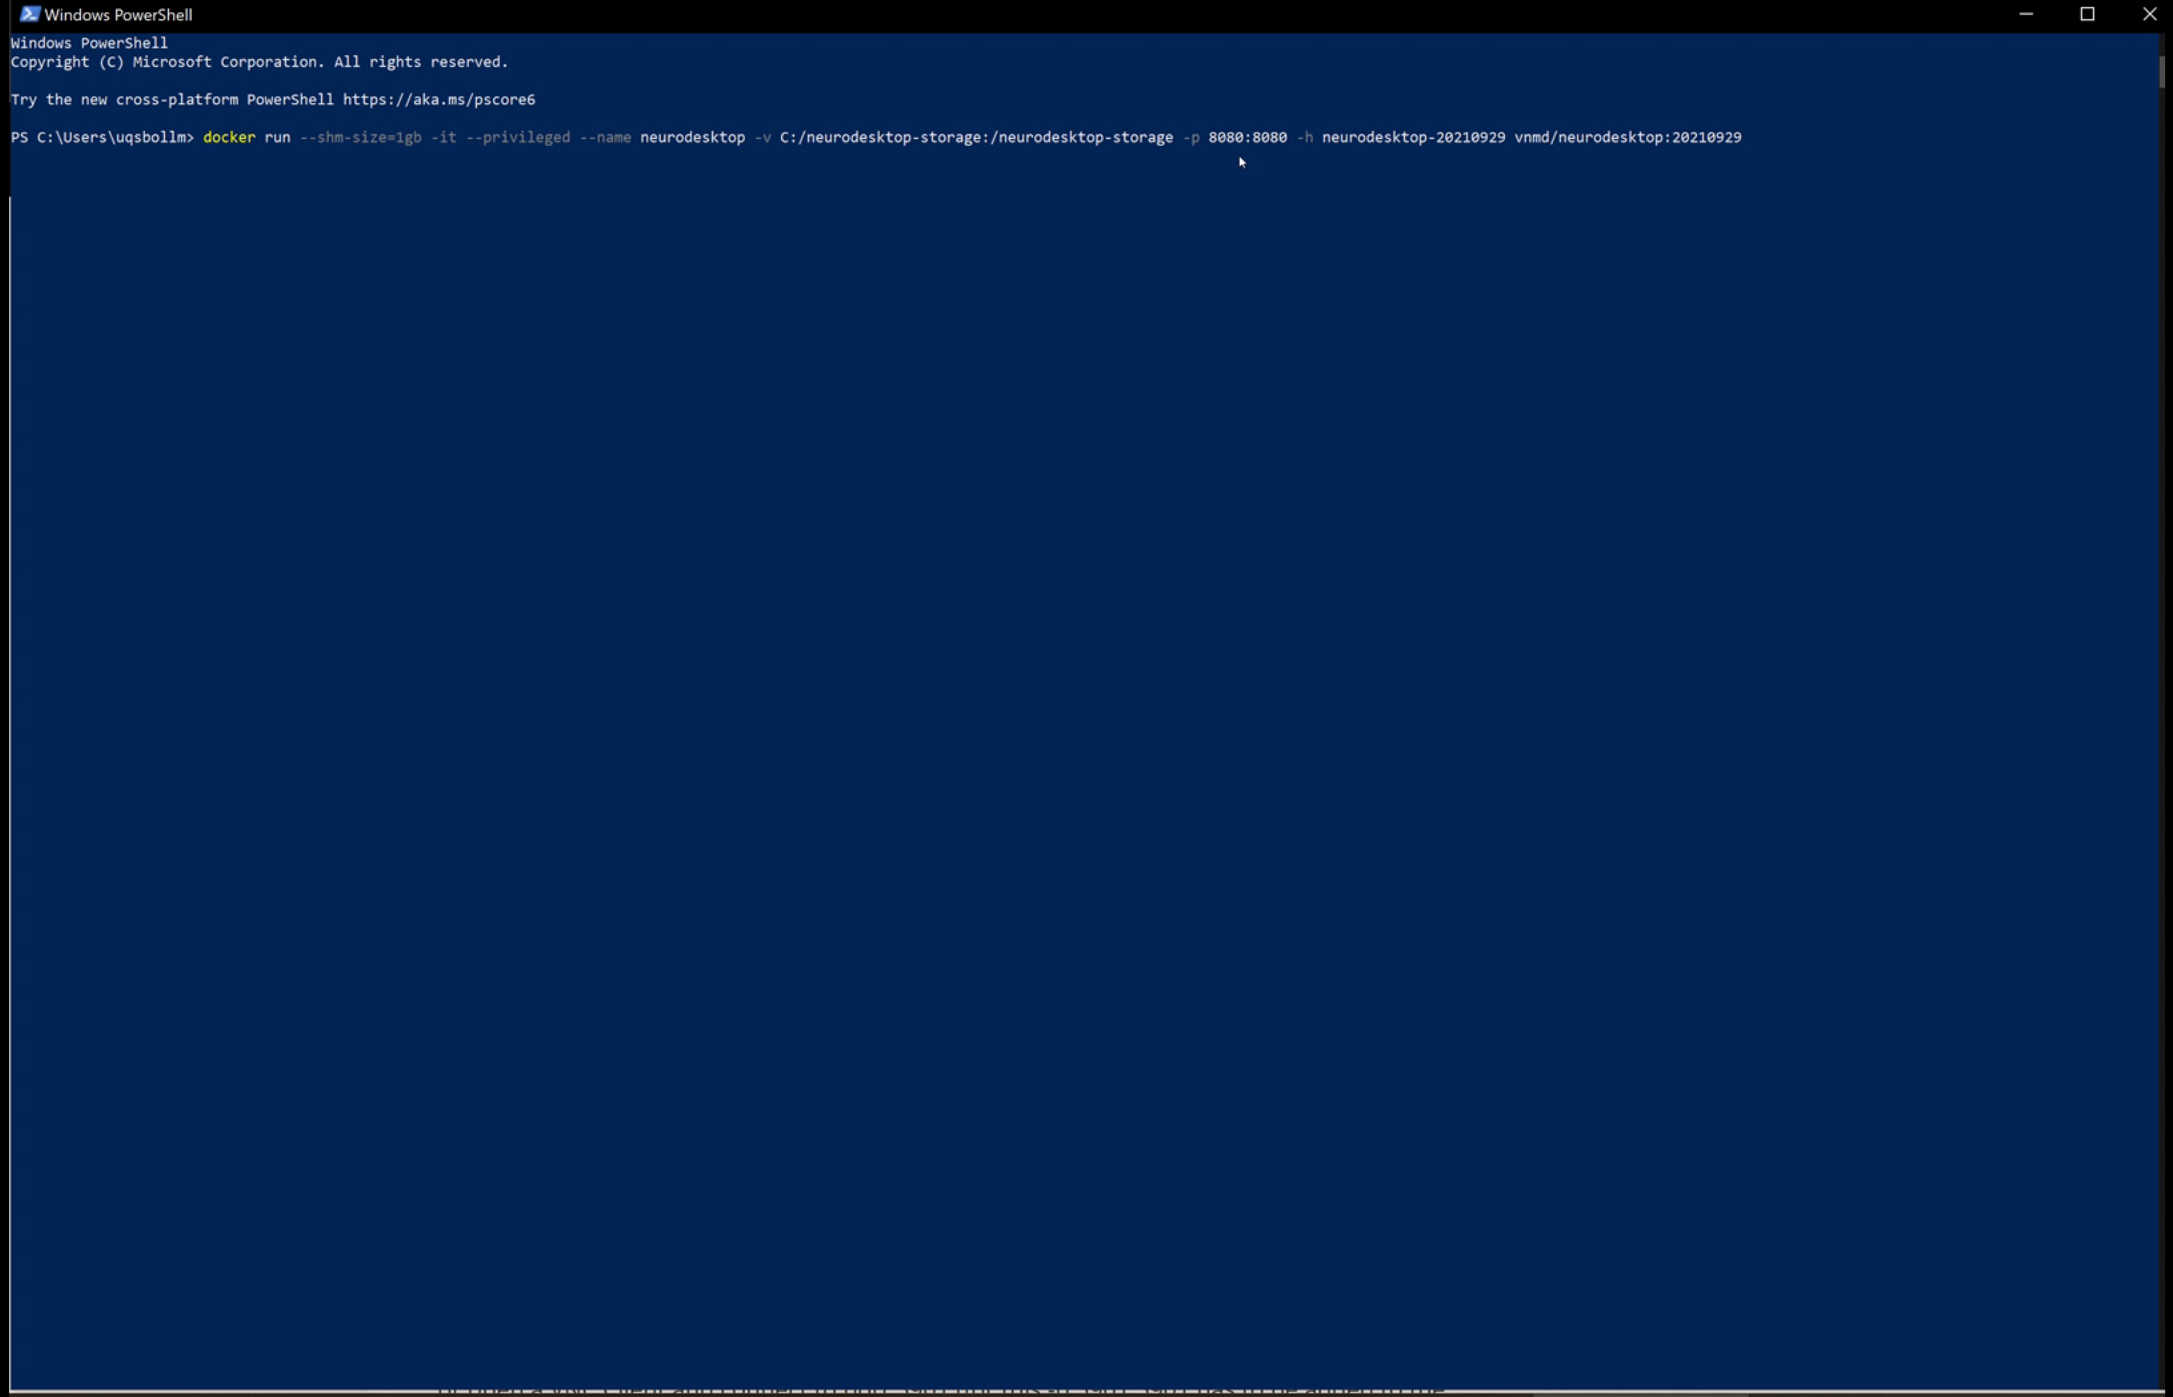
Windows PowerShell
Windows PowerShell
Copyright (C) Microsoft Corporation. All rights reserved.
Try the new cross-platform PowerShell https://aka.ms/pscore6
PS C:\Users\uqsbollm> docker run --shm-size=1gb -it --privileged --name neurodesktop -v C:/neurodesktop-storage:/neurodesktop-storage -p 8080:8080 -h neurodesktop-20210929 vnmd/neurodesktop:20210929
Windows | NeuroD...
neurodesktop-stora...
Desktop
neurodesk_and_AE...
Windows PowerShell
8:58 AM

Neurodesk Architecture

Researchers on Windows, Mac, Linux

Advanced users on HPC/Linux

Developers

Neurodesk Architecture

Neurodesktop

Neurocommand

Transparent-Singularity

Neurocontainers

Neurodocker

Researchers on Windows, Mac, Linux

Advanced users on HPC/Linux

Developers

What is neurodocker?

- It's a docker/singularity recipe generator for neuroimaging software:

```
neurodocker generate ${neurodocker_buildMode} \
   --base-image ubuntu:16.04 \
   --pkg-manager apt \
   --ants version=${toolVersion} \
  > ${imageName}.Dockerfile
```

```
FROM ubuntu:16.04
RUN printf '#!/bin/bash\nls -la' > /usr/bin/ll
RUN chmod +x /usr/bin/ll
RUN mkdir /afm01 /afm02 /90days /30days /QRISdata /RDS /data /short /proc_temp /TMPDIR /nvme /local /gpfs1 /working
ENV ANTSPATH="/opt/ants-2.3.4/" \
    PATH="/opt/ants-2.3.4:$PATH"
RUN apt-get update -qq \
    && apt-get install -y -q --no-install-recommends \
           ca-certificates \
           curl \
    && rm -rf /var/lib/apt/lists/* \
    && echo "Downloading ANTs ..." \
    && mkdir -p /opt/ants-2.3.4 \
    && curl -fsSL https://dl.dropbox.com/s/gwf51ykkk5bifyj/ants-Linux-centos6_x86_64-v2.3.4.tar.gz \
    | tar -xz -C /opt/ants-2.3.4 --strip-components 1
ENV DEPLOY_PATH="/opt/ants-2.3.4/"
COPY ["README.md", \
      "/README.md"]
```

https://github.com/ReproNim/neurodocker

Why neurodocker and what we learned

- Writing efficient Dockerfiles is tedious (chaining commands, cleaning up after apt/yum …)
- figuring out how to install neuroimaging software is a lot of work -> it needs a village to solve some of the dependency issues
- Neurodocker has a big user base
  (232 stars, 82 forks)
- we are one of the forks and we provide pull requests back upstream for fixes

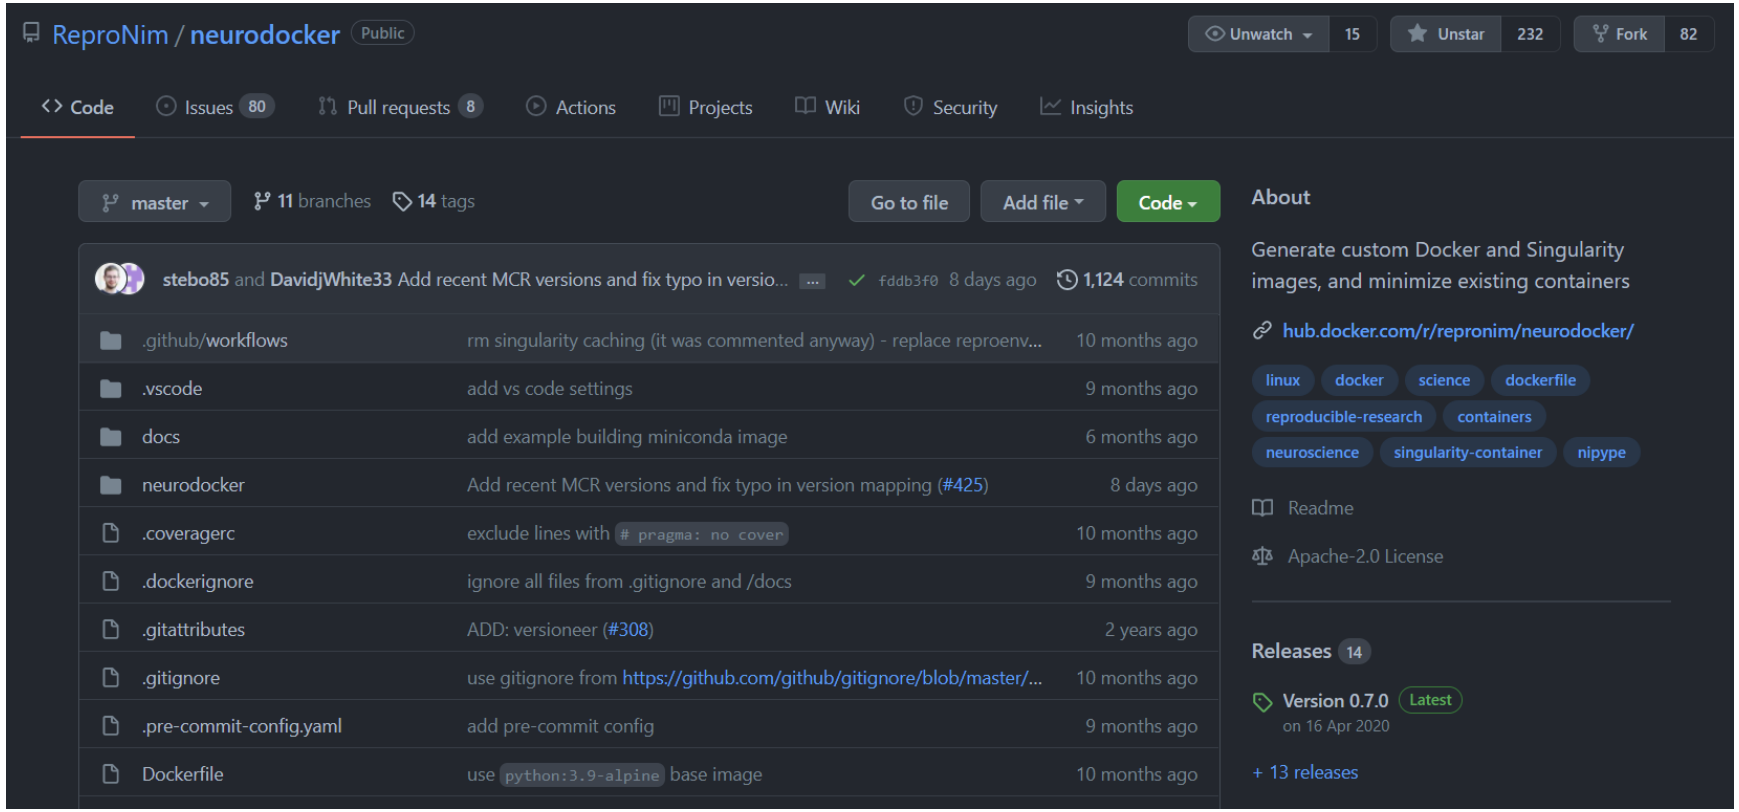

https://github.com/ReproNim/neurodocker

Neurodesk Architecture

Neurodesktop

Neurocommand

Transparent-Singularity

Neurocontainers

Neurodocker

Researchers on Windows, Mac, Linux

Advanced users on HPC/Linux

Developers

Every application is a singularity sub-container!

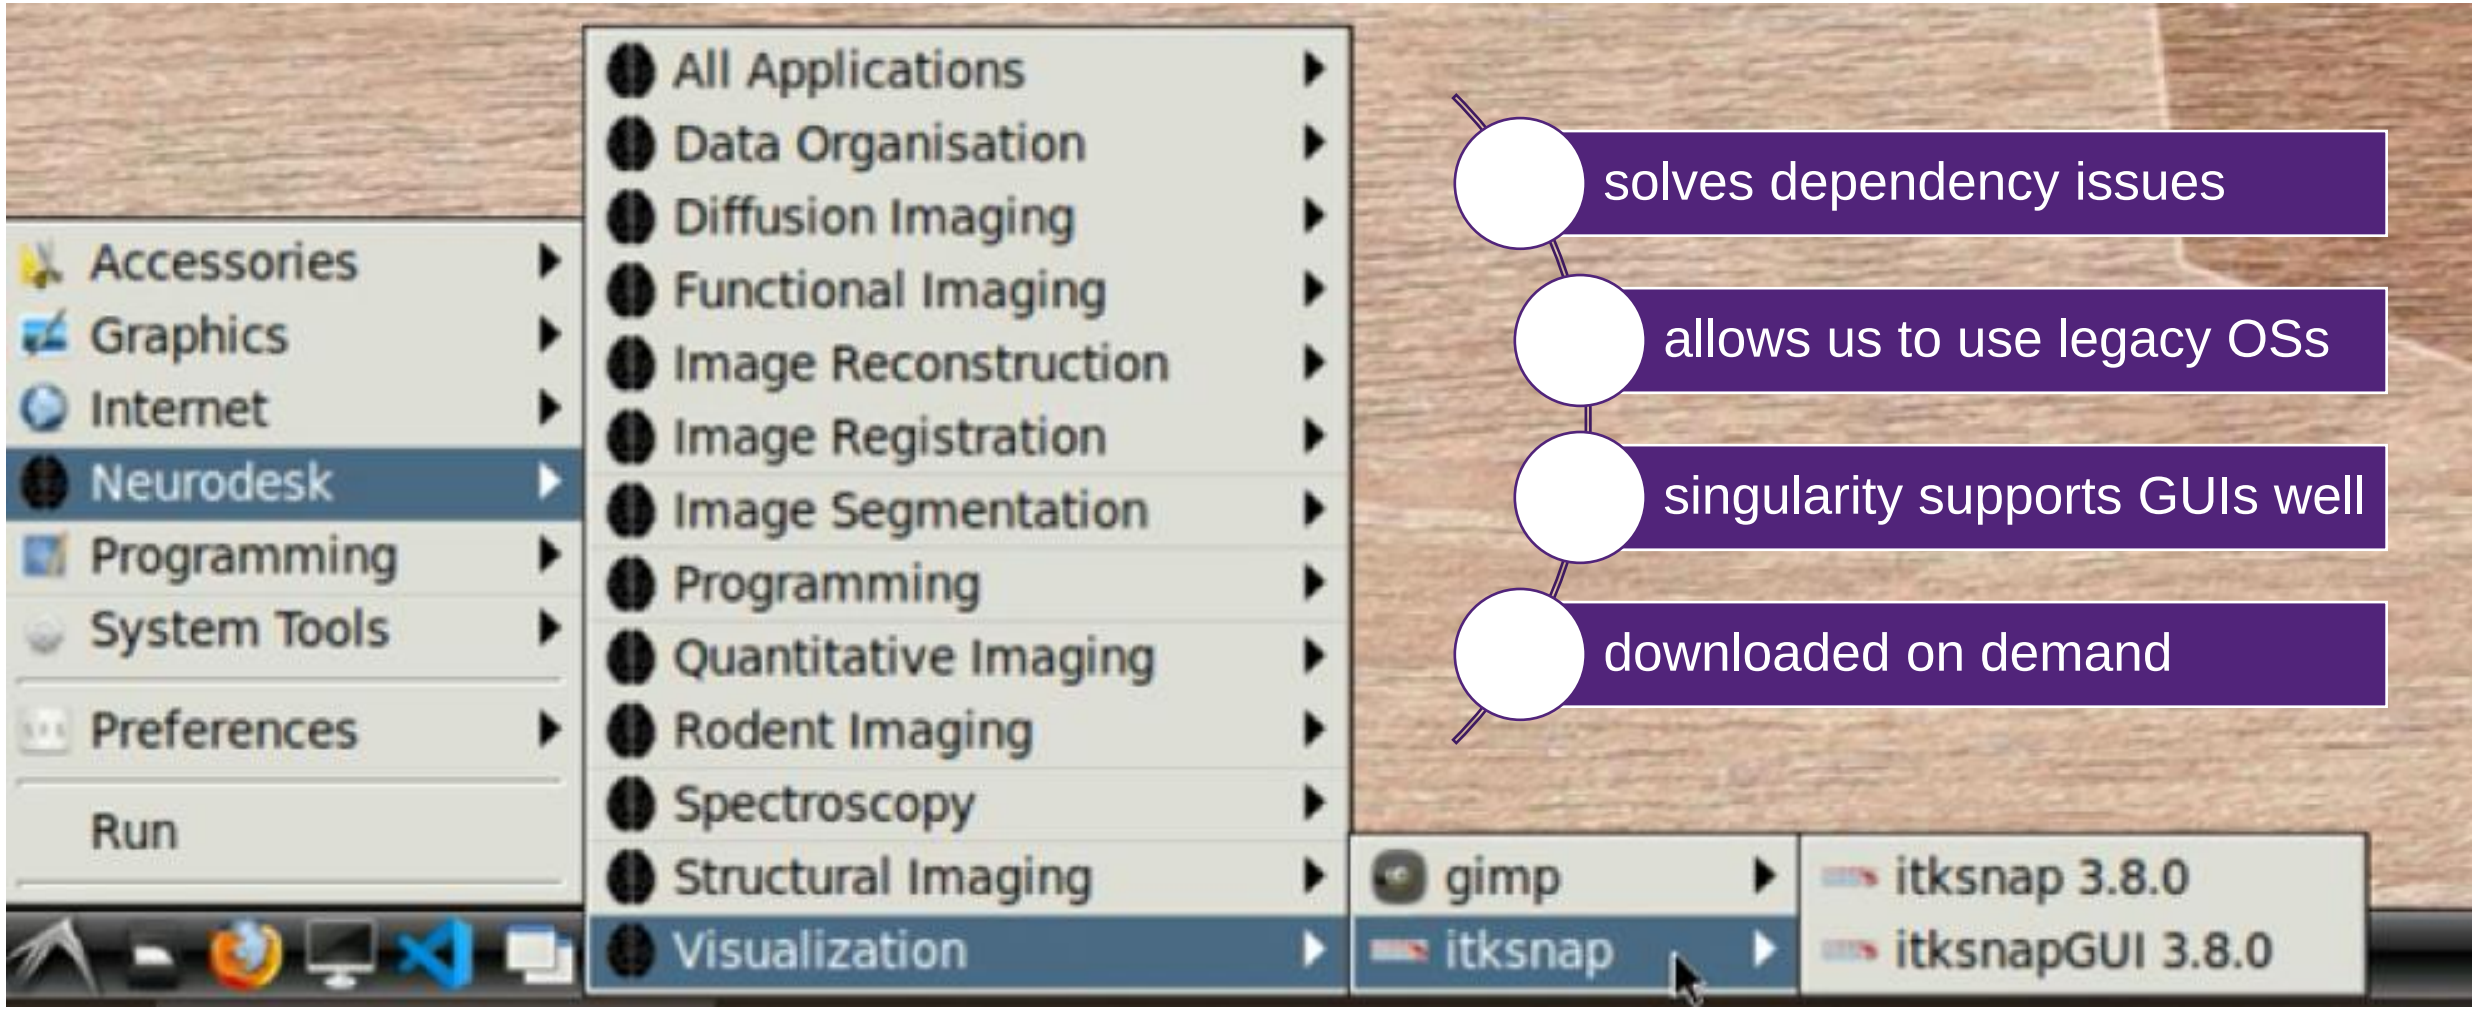

Github actions build the application containers

- Every application container has a .yml file that builds the container and pushes it to the registries
- example: https://github.com/NeuroDesk/neurocontainers/blob/master/.github/workflows/ants.yml

- runners have limited free diskspace -> cleanup action to the rescue

```yaml
- name: Free up space (optional)
  if: env.FREEUPSPACE
  uses: easimon/maximize-build-space@master
  with:
    root-reserve-mb: 40000
    swap-size-mb: 1024
    remove-dotnet: 'true'
    remove-android: 'true'
    remove-haskell: 'true'
    overprovision-lvm: 'true'
- name: Move docker installation (optional)
  if: env.FREEUPSPACE
  run: |
    sudo mv /var/lib/docker /home/runner/work/docker
    sudo ln -s /home/runner/work/docker /var/lib/docker
    sudo systemctl restart docker
```

What we learned

Github action runners can be too small	• clean up actions available
docker hub is rate-limiting	• use ghcr.io + hope for ARCOS registry
running singularity containers within docker containers …	• it sounds like a Christopher Nolan movie, but works well!

Neurodesk Architecture

Neurodesktop

Neurocommand

Transparent-Singularity

Neurocontainers

Neurodocker

Researchers on Windows, Mac, Linux

Advanced users on HPC/Linux

Developers

Transparent Singularity

- Problem: Users are familiar with running the application commands directly + workflow systems are often not aware of containers

❑ What users want:

```
fslmaths
```

❑ What users would need to run:

```
singularity exec --pwd $PWD /vnm/containers/fsl_6.0.3_20200820/fsl_6.0.3_20200820.sif fslmaths
```

- Our work-around: Automatically generate wrapper scripts for every application inside the container

```
neuro@5c2bff15d401:~$ cat /vnm/containers/fsl_6.0.3_20200820/fslmaths
#!/usr/bin/env bash
export PWD=`pwd -P`
singularity exec --pwd $PWD /vnm/containers/fsl_6.0.3_20200820/fsl_6.0.3_20200820.sif fslmaths $@
```

https://github.com/NeuroDesk/transparent-singularity

Transparent Singularity

- Using the lmod module system we can now combine the tools from different singularity-containers in a larger workflow ☺

```
neuro@5c2bff15d401:~$ module avail

---------------------------- /vnm/containers/modules ----------------------------
   freesurfer/7.1.0        fsl/6.0.3        mrtrix3/3.0.1
   freesurfer/7.1.1 (D)    itksnap/3.8.0
```

```
neuro@5c2bff15d401:~$ ml fsl
neuro@5c2bff15d401:~$ ml

Currently Loaded Modules:
  1) fsl/6.0.3

neuro@5c2bff15d401:~$ ml freesurfer
neuro@5c2bff15d401:~$ ml

Currently Loaded Modules:
  1) fsl/6.0.3   2) freesurfer/7.1.1
```

```
neuro@5c2bff15d401:~$ which freeview
/vnm/containers/freesurfer_7.1.1_20200924/freeview
neuro@5c2bff15d401:~$ which fslmaths
/vnm/containers/fsl_6.0.3_20200820/fslmaths
```

What we learned

wrapper scripts for singularity work well

- users can continue using their old scripts and haven't even noticed that they are running inside containers

lmod module system allows combining tools from different containers

- no dependency conflicts and full isolation between tools, but combination of different containers is straight forward

Neurodesk Architecture

Neurodesktop

Neurocommand

Transparent-Singularity

Neurocontainers

Neurodocker

Researchers on Windows, Mac, Linux

Advanced users on HPC/Linux

Developers

Neurocommand: Bringing it all together

Define which application containers are available

Convert Docker containers to singularity containers and upload to object storage

Unpack singularity containers and store on CVMFS for distribution

https://github.com/NeuroDesk/neurocommand

Neurocommand: Bringing it all together

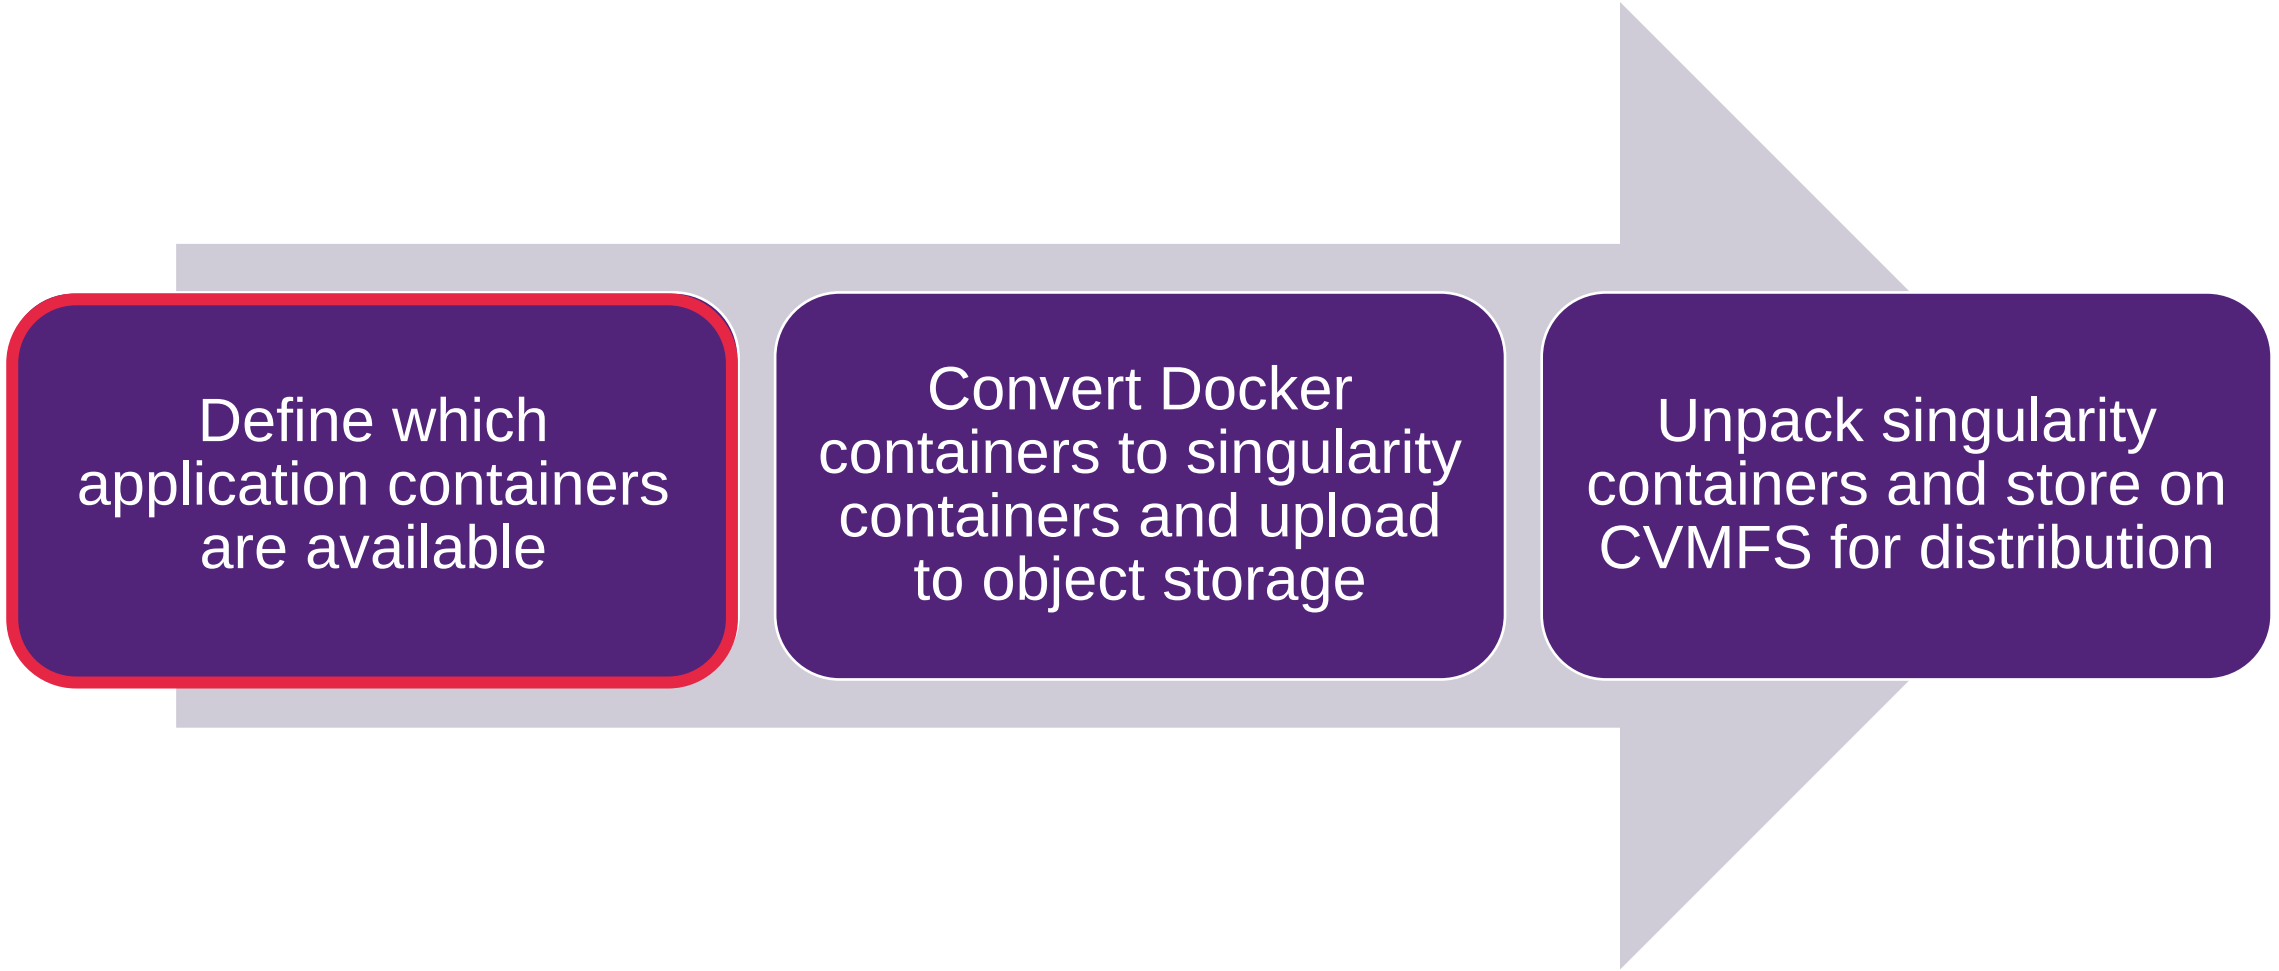

https://github.com/NeuroDesk/neurocommand

Json file defines which container versions are “live”

- https://github.com/NeuroDesk/neurocommand/blob/main/neurodesk/apps.json

```
"ants": {
    "apps": {
        "ants 2.3.4": {
            "version": "20210104",
            "exec": ""
        }
    },
    "categories": ["image registration"]
},
```

Neurocommand: Bringing it all together

Define which application containers are available

Convert Docker containers to singularity containers and upload to object storage

Unpack singularity containers and store on CVMFS for distribution

https://github.com/NeuroDesk/neurocommand

Converting Docker to Singularity containers

- Advantage: Use the build tools and caching of docker and the easy integration of GUIs of Singularity containers -> building automatically using github action

```
singularity build "$IMAGE_HOME/${IMAGENAME_BUILDDATE}.simg"  docker://vnmd/$IMAGENAME:$BUILDDATE
```

- Problem: There is currently no nice Singularity registry ☹
  - Workaround: store on object storage

```
curl -X PUT -u ${ORACLE_USER} --upload-file $IMAGE_HOME/${IMAGENAME_BUILDDATE}.simg $ORACLE_NEURODESK_BUCKET
```

  - and mirror:

```
rclone sync oracle-2021-us-bucket:/neurodesk nectar:/neurodesk/
rclone copy oracle-2021-us-bucket:/neurodesk oracle-2021-sydney-bucket:/neurodesk
```

  - and download using aria2:

```
aria2c https://objectstorage.us-ashburn-1.oraclecloud.com/n/sd63xuke79z3/b/neurodesk/o/$container https://objectstorage.eu-frankfurt-1
```

Neurocommand: Bringing it all together

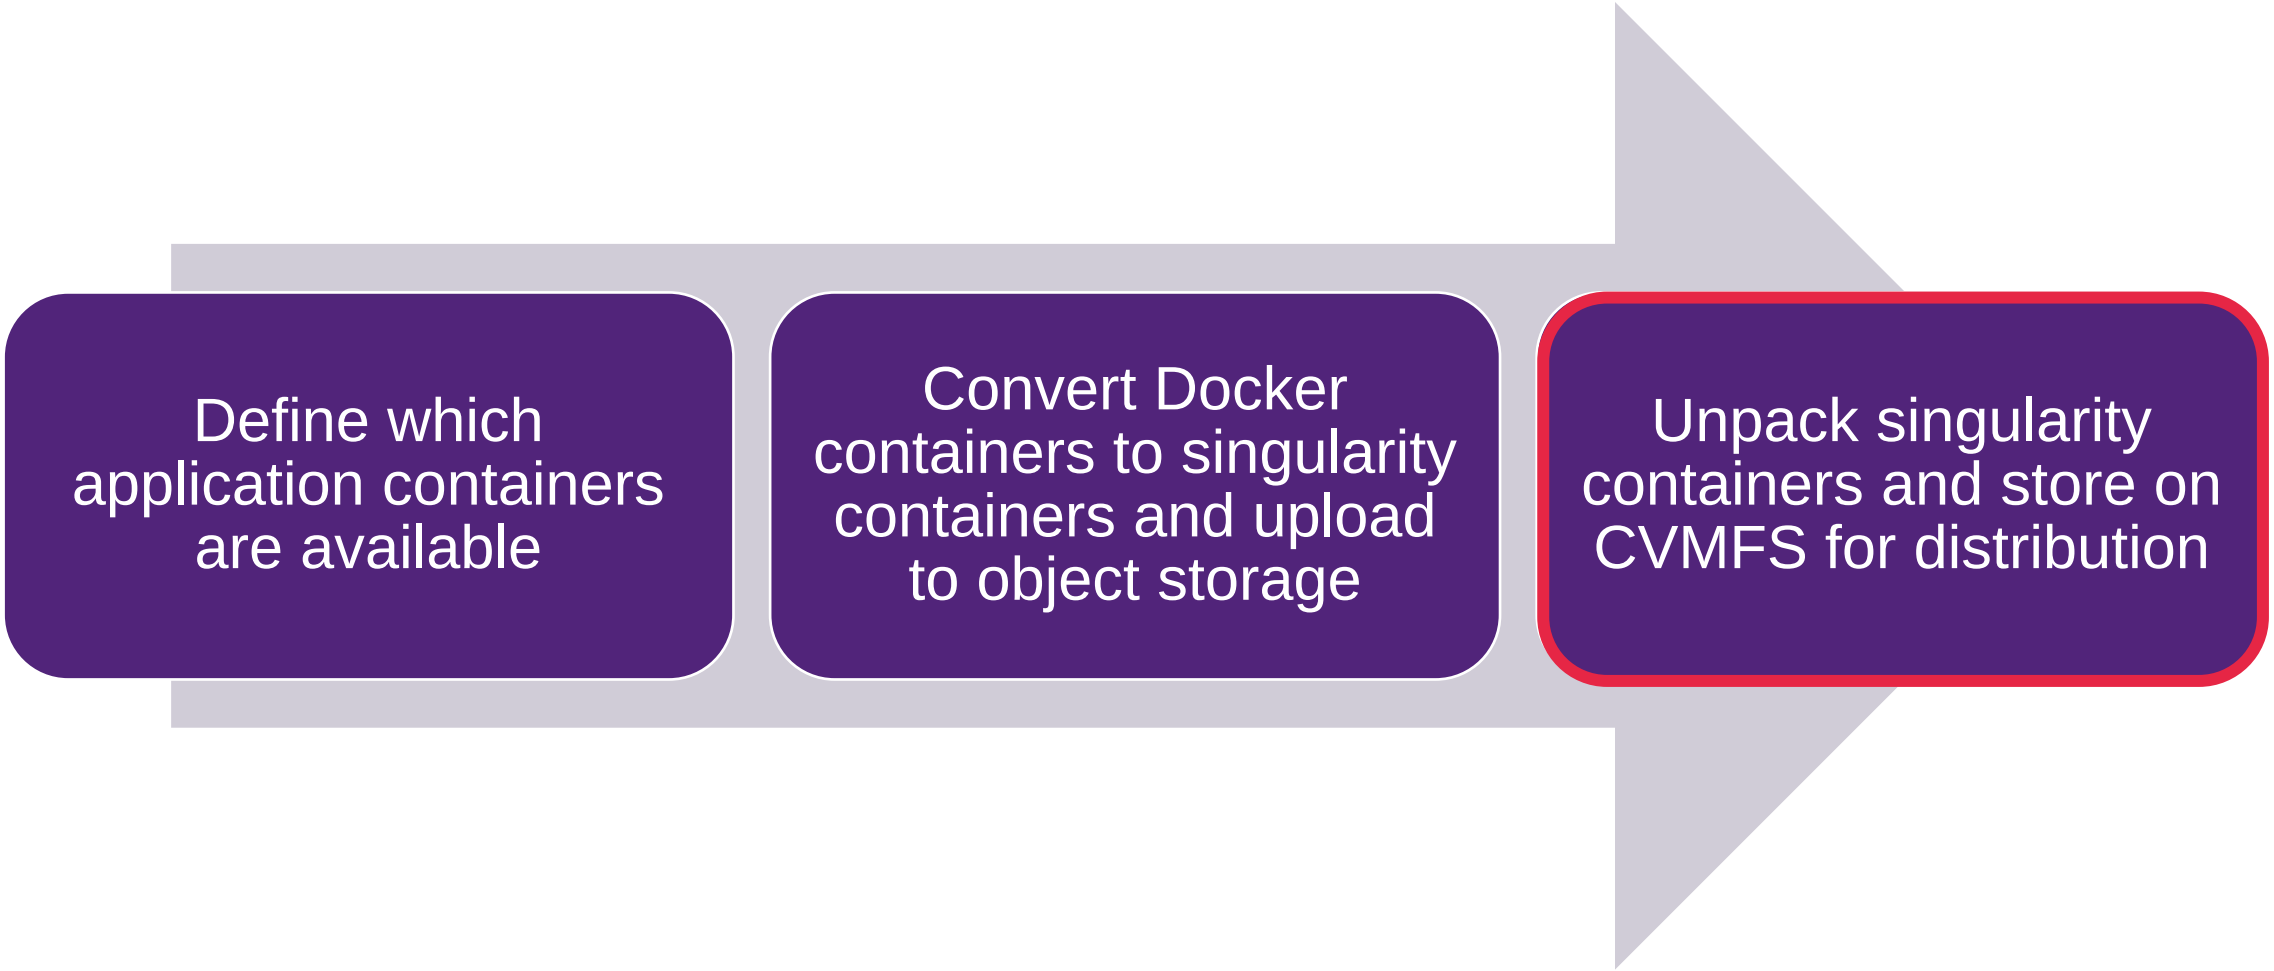

https://github.com/NeuroDesk/neurocommand

Distributing Singularity containers via CVMFS

- download and unpack singularity containers to CVMFS storage for distribution and on-demand access
- https://github.com/NeuroDesk/neurocommand/blob/main/cvmfs/sync_containers_to_cvmfs.sh

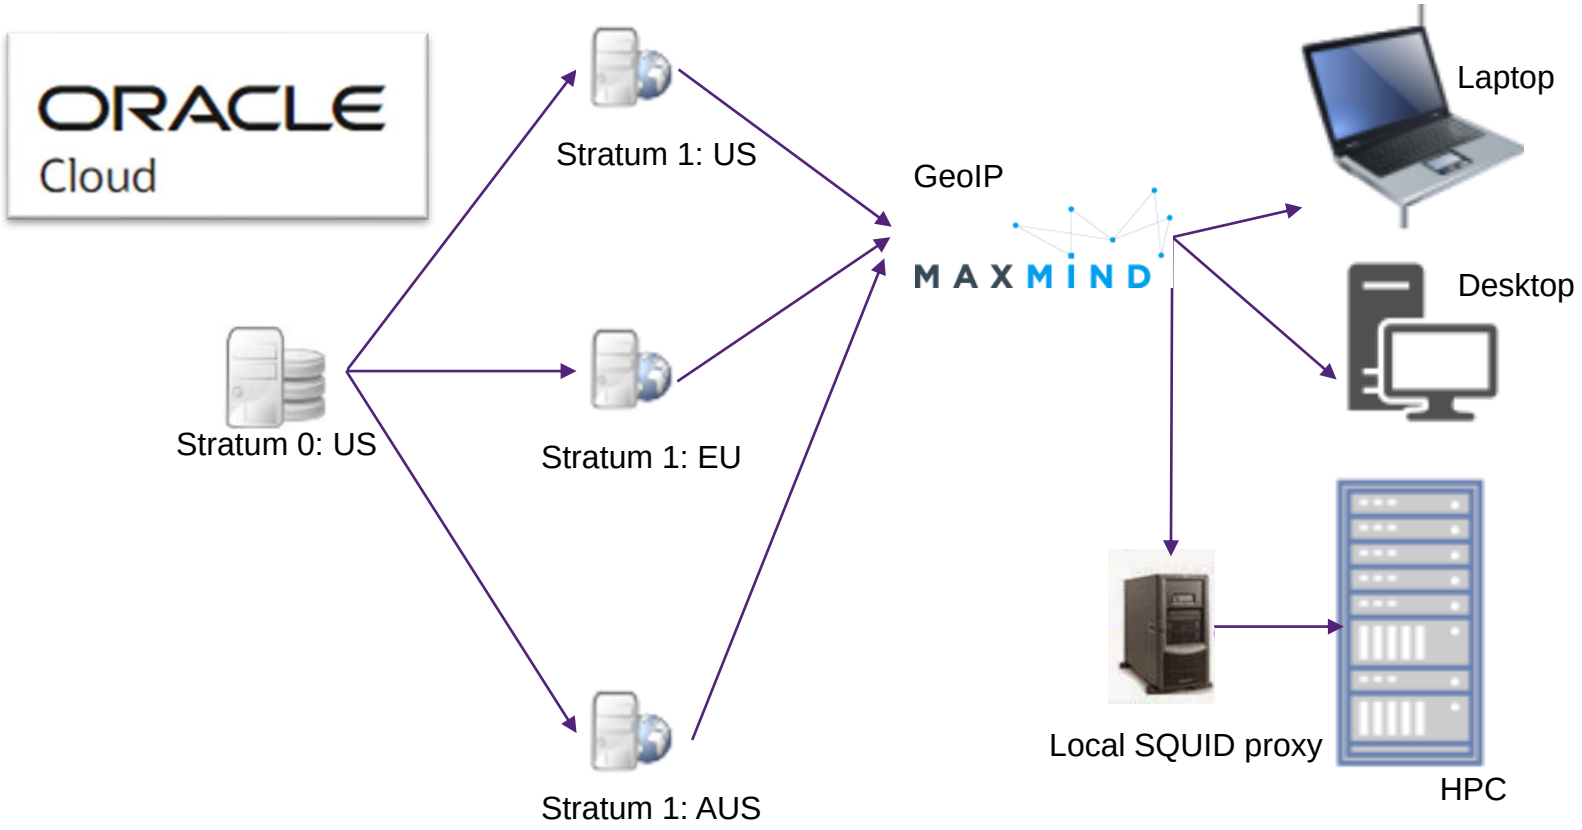


```
ll /cvmfs/neurodesk.ardc.edu.au/containers/

Aug 26 16:41 afni_21.2.00_20210714
Sep 29 12:35 aidamri_1.1_20210708
Aug 26 16:44 aidmri_1.1_20210330
Aug 26 16:49 ants_2.3.4_20210104
Aug 26 17:04 ashs_2.0.0_20210105
Aug 26 16:54 aslprep_0.2.7_20210323
Aug 26 17:05 bart_0.7.00_20210302
Aug 26 17:06 bidstools_1.0.0_20201208
Oct 18 18:01 clearswi_1.0.0_20211018
Oct 13 09:37 condaenvs_1.0.0_20211011
```

https://neurodesk.github.io/developers/cvmfs/

What we learned

a singularity registry would be nice …

- but current object storage workaround works

CVMFS is great and fast

- accessing >200GB of software on demand in our lightweight desktop container

Docker hub is rate-limiting our container pulls

- we couldn’t use the cvmfs ducc tool to pull containers, but build our own tool starting from singularity containers on object store

Neurodesk Architecture

Neurodesktop

Neurocommand

Transparent-Singularity

Neurocontainers

Neurodocker

Researchers on Windows, Mac, Linux

Advanced users on HPC/Linux

Developers

How we build NeuroDesktop

- Full LXDE desktop based on ubuntu:20.04
  - https://github.com/NeuroDesk/neurodesktop/blob/main/Dockerfile
  - Tomcat, Guacamole, Singularity, CVMFS, lmod, VS code, git, python, julia …
- Automatically build and deployed using GitHub Actions:
  - https://github.com/NeuroDesk/neurodesktop/blob/main/.github/workflows/build-neurodesktop.yml
  - deployed to Dockerhub and Github Packages via daily build
  - including test if CVMFS servers are working

```
ARG GO_VERSION="1.14.12"
ARG SINGULARITY_VERSION="3.8.2"
ARG TOMCAT_REL="9"
ARG TOMCAT_VERSION="9.0.52"
ARG GUACAMOLE_VERSION="1.3.0"

# Create final image
FROM ubuntu:20.04

# Install base image dependancies
RUN apt-get update \
    && DEBIAN_FRONTEND=noninteractive apt-get install --no-install-recommends -y \
        locales \
        sudo \
        wget \
        ca-certificates \
```

```
build:
  needs: [test_cvmfs]
  runs-on: ubuntu-latest
  steps:
  - uses: actions/checkout@v2
  - name: Pull latest neurodesktop build from GitHub packages
    run: |
      echo ${GITHUB_REF}
      echo "${{ secrets.GITHUB_TOKEN }}" | docker login ghcr.io -u $GITHUB_ACTOR --password-stdin
      IMAGEID=ghcr.io/$GITHUB_REPOSITORY/neurodesktop
      IMAGEID=$(echo $IMAGEID | tr '[A-Z]' '[a-z]')
      {
        docker pull $IMAGEID \
          && ROOTFS_CACHE=$(docker inspect --format='{{.RootFS}}' $IMAGEID) \
          && echo "ROOTFS_CACHE=$ROOTFS_CACHE" >> $GITHUB_ENV
      } || echo "$IMAGEID not found. Resuming build..."
      echo "IMAGEID=$IMAGEID" >> $GITHUB_ENV
  - name: Build new neurodesktop image
```

What we learned

seamless copy and paste between host and container in browser is important to people	• changed from novnc to guacamole
the neurodesktop container needs to be as lightweight as possible	• applications are in separate singularity containers that get downloaded on demand -> only 3.47 GB for desktop ☺
use docker caching to speed up builds	• daily build time reduced from 21 minutes to 2 minutes

Roadmap

More Neuroimaging containers

Support of GPUs

GUI for managing desktop container

Open Deployment on MyBinder + our own Kubernetes cluster

Support of M1/Arm processors

Acknowledgements

Australian Research Container Orchestration Service (ARCOS)

nectar

National Imaging Facility

Australian Brain Alliance | EMCR Network

OHBM 2020
A Virtual Experience for Engaging Minds & Empowering Brain Science

ORACLE Cloud

ARDC Australian Research Data Commons

Google Summer of Code

Martin Grignard

Peter Winn

Jakub Kaczmarzyk

Niruhan Viswarupan

Dao Thanh Thuy

Aswin Narayanan

Yorguin Jose Mantilla

Will Woods

Steffen Bollmann

Tom Johnstone

Angela Renton

David White

Thomas Shaw

Oren Civier

Megan Campbell

Paris Lyons

Ryan Sullivan

Collaborators:

- Thomas Close, USyd
- Wojtek Goscinski, Monash/NIF
- Tony Hannan, Florey
- Gary Egan, Monash
- Linda Richards, UQ
- Paul Sowman, Macquarie
- Marta Garrido, UoM
- Patrick Johnston, QUT
- Aina Puce, Indiana University
- Franco Pestilli, Indiana University
- Levin Kuhlmann, Monash
- Gershon Spitz, Monash
- David Abbott, Florey
- Megan Campbell, Newcastle
- Nigel Rogasch, University of Adelaide
- Satrajit Ghosh, MIT