

Conversion of docker to singularity images



CSC – Suomalainen tutkimuksen, koulutuksen, kulttuurin ja julkishallinnon ICT-osaamiskeskus

Outline



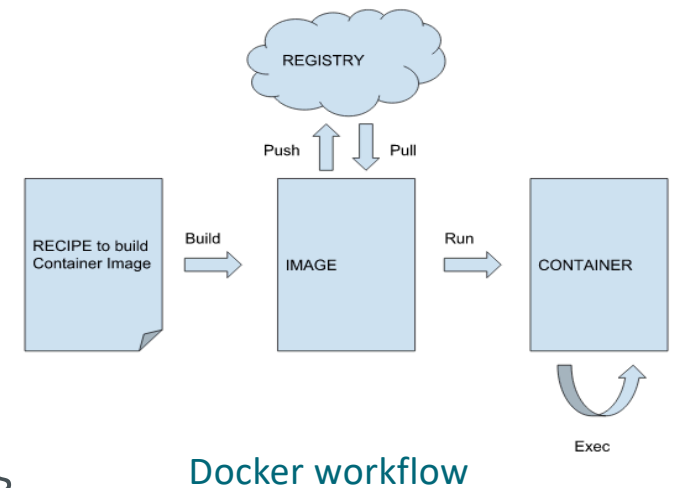
- Brief introduction to container images
- Running singularity containers from docker images available in registries
- Running singularity containers from local docker images
- Converting images efficiently on CSC HPC environment
- Good practices while converting images

Brief Introduction to container Images

Docker Images in HPC



- The docker is a leading container platform and is widely adopted in industry/cloud applications
- Large number of images are archived in docker registries
- GPU-accelerated solutions to the HPC community
- One needs root privileges to create docker containers from docker images
- Good news: Singularity nicely integrates with docker images
- Not all docker images will work smoothly with singularity



Docker vs. Singularity Images

- Docker image
 - Image is a *layered* structure (OCI blobs)
 - Stored in the local image cache when using docker client
 - Saves some space by sharing some layers across different local images when storing / caching the build layers can speed up the build process
- Singularity image
 - Single image file (.sif) / images are *flat*
 - Stored a normal file
 - Easy portability : transferring and sharing them across a cluster very easy
 - Image can be used as binary (./image.sif)

Running Singularity Containers from Docker Images from Registries

Learn to Run Singularity Containers from Docker Images



- Use Singularity subcommands
 - singularity pull ...
 - singularity build ...
 - Build is functionnally equivalent to pull but does more (like make writable image)

```
singularity pull docker://ilumb/mylolcow
```

```
singularity build docker://ilumb/mylolcow
```

```
singularity build --sandbox mylolcow_latest_sandbox  
docker://godlovedc/lolcow
```

Running Singularity Containers from Docker Images



- In private repo, one has to get login authentication
- singularity pull --docker-login docker://ilumb/mylolcow
 - export SINGULARITY_DOCKER_USERNAME=ilumb
 - export SINGULARITY_DOCKER_PASSWORD=<redacted>
- Finally test that image works : ./imagename.sif

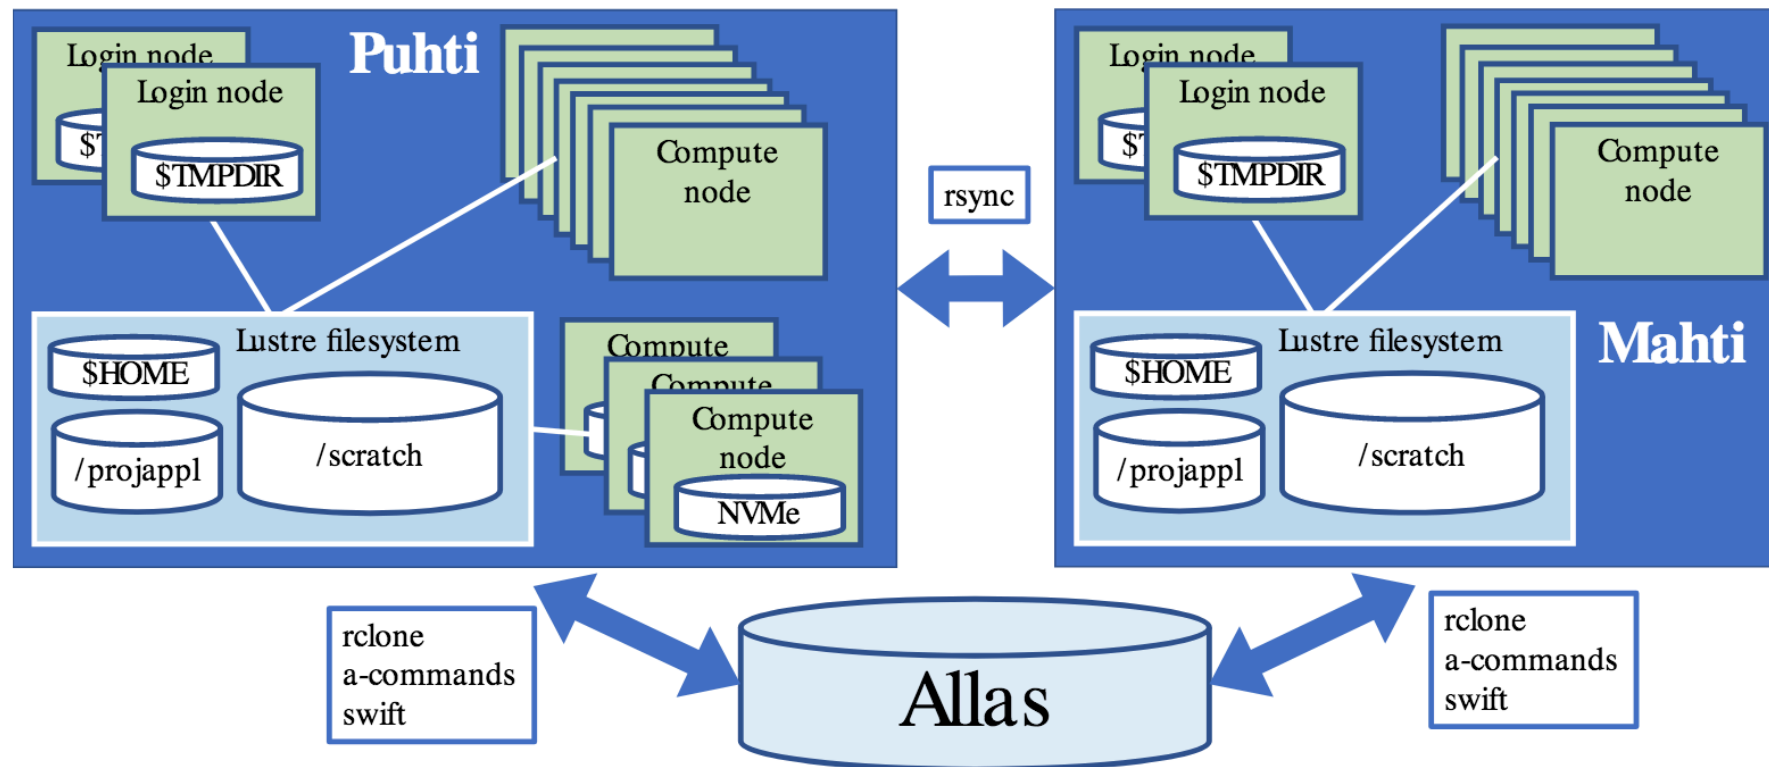
Running Singularity Containers from Local Docker Images



- First, save your local docker image (why?) from your machine (VM, laptop)
 - Check using : `sudo docker images`
 - `sudo docker save image_id -o local_docker.tar`
- Then copy to another machine (e.g., HPC environment)
- Finally, build a singularity image from the local_docker.tar
 - `singularity build local_singularity_image.sif docker-archive://local_docker.tar`
 - Using docker-archive bootstrap agent.

Converting images efficiently in CSC HPC environment

Disk and Storage Overview



How to use NVMe disks in CSC environment



- Efficient when you are converting bigger images
- Not all compute nodes have these solid state disks
- Request these resources in batch script /do it in interactive node
- In batch jobs, you have to copy the files to back to Lustre files to preserve the results



Singularity Cache in HPC environment



- Default location: `$HOME/.singularity`
 - Amount of available space in `$HOME` in HPC system is limited by a quota
 - Your home can be quickly filled up and end up in disk space errors
- CSC users can configure the Singularity cache using two environment variables:
 - `SINGULARITY_CACHEDIR`: Cache folder for images from a container registry.
 - `SINGULARITY_TMPDIR`: Temporary directory to build container file-systems.
- Friendly tips:
 - `singularity cache list #` show storage capacity used by the cache
 - `singularity cache clean #` clean up everything

Good Practices While Converting Images

Good Practices in Image Conversion

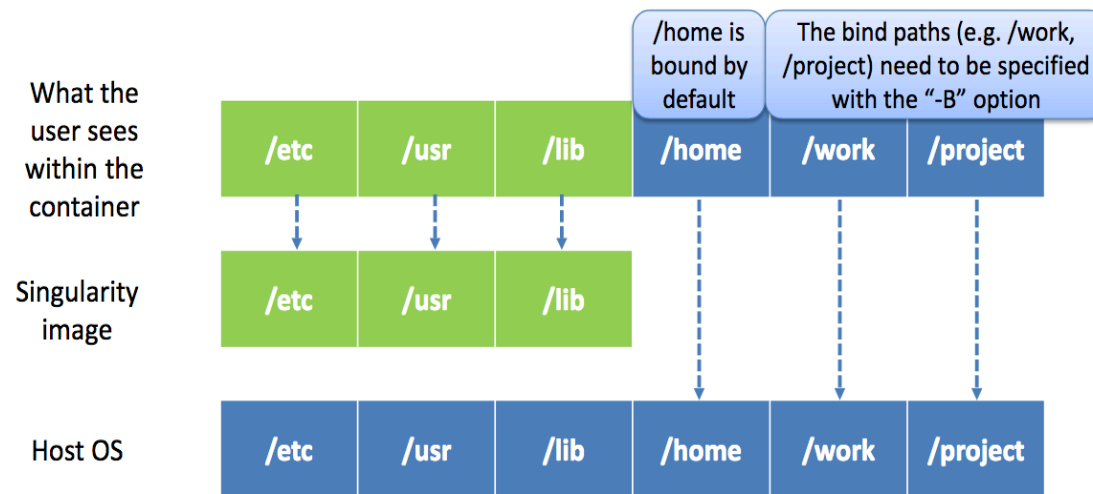


- Maintain tags of the image (reproducibility !)
 - Generic URI for docker : `docker://<user>/<repo-name>[:<tag>]`
 - Developers release several different versions of the same container with different tags
- Hashes are even more unique to image
 - Always pull same image: `singularity pull library://debian:sha256.b92c7...`
- Pay attention to special tag: latest
 - Today's latest is different from the tomorrow's latest (avoid using it)

Good Practices in Image Conversion



Avoid using docker images that install to \$HOME or \$TMP



Overlay file system in singularity

Good Practices in Image Conversion

- Some containers will not work if kernels of host system are too old
- Pay attention to Entrypoints in dockerfile (usually broken in singularity conversion) :
 - WORKDIR /code
 - ENTRYPOINT ["entrypoint_script.sh"]