Using containers for HPC workloads (but avoiding Docker)

Eric Burgueño

A day in the life...

```
$ ./ToolA file1 file2
$ ls
file1 file1.out file2 file2.out ToolA
```

Processing multiple files (local test)

```
$ export PATH=$(pwd):${PATH}
$ cd ~/input/files/to/process/samples
$ ls -l | wc -l
38
$ qsub=""
$ for files in *; do $qsub ToolA $files; done
...
$ ls -l | wc -l
75
```

Using Docker

```
$ cd ~/input/files/to/process/samples
$ rm -f *.out
$ ls -l | wc -l
38
$ docker run -d registry.example.com/ToolA:latest
file1
e61d12292d69556eabe2a4...f95438e504afb7b02810
```

Using Docker (try 1)

```
$ ls -l | wc -l
38
$ ls *out
$
```

Using Docker (try 2)

```
$ docker run -d -v $(pwd):/mnt
registry.example.com/ToolA:latest /mnt/file1
653e785339099e374b...f393ee0e4adbb795a3935060acb
$ ls -l | wc -l
38
$ ls *out
$
$ docker logs 653e785339
ToolA: /mnt/file1: Permission denied
```

Using Docker (try 3)

```
$ docker run -d -v $(pwd):/mnt:z
registry.example.com/ToolA:latest /mnt/file1
8ebfcbcb31bea0696e0a7...19c9623e1846d8185972dc3b
$ ls -l | wc -l
38
$ ls *out
$
$ docker logs 8ebfcbcb31
ToolA: /mnt/file1: Permission denied
```

Using Docker (try 4)

```
$ chmod 777 file1
$ docker run -d --name=test -v $(pwd):/mnt:z
registry.example.com/ToolA:latest /mnt/file1
0b61185ef4a78dce988bb3...d5aea2b6a583d7ffbceca16
$ ls *out
$
$ docker logs test
ToolA: cannot create regular file
'/mnt/file1.out': Permission denied
```

Using Docker (try 5)

```
$ docker run -d -u $(id -u):$(id -g) --name=test
-v $(pwd):/mnt:z
registry.example.com/ToolA:latest /mnt/file1
```

Using Docker (try 5)

```
$ docker run -d -u $(id -u):$(id -g) --name=test
-v $(pwd):/mnt:z
registry.example.com/ToolA:latest /mnt/file1

docker: Error response from daemon: Conflict. The
container name "/test" is already in use by
container "0b61185ef4a78...a2b6a583d7ffbceca16".
You have to remove (or rename) that container to
be able to reuse that name.
```

Using Docker (try 5 - final)

```
$ docker rm test
$ docker run -d -u $(id -u):$(id -g) --name=test
-v $(pwd):/mnt:z
registry.example.com/ToolA:latest /mnt/file1
06d5b3d52e1167cde50c2...4672cd3ca91043ad23c1fe09
$ ls *out
file1.out
$
```

Using Docker (try 5 - final)

```
$ docker rm test
$ docker run -d -u $(id -u):$(id -g) --name=test
-v $(pwd):/mnt:z
registry.example.com/ToolA:latest /mnt/file1
06d5b3d52e1167cde50c2...4672cd3ca91043ad23c1fe09
$ ls *out
file1.out
$
```

Putting it all together

```
$ cd ~/input/files/to/process
$ ls -l | wc -l
934752984
$
$ for files in *; do qsub -q short_jobs -N
"toola_${files}" docker run -d -u $(id -u):$(id
-g) --name="toola_${files}" -v $(pwd):/mnt:z
registry.example.com/ToolA:latest /mnt/${files};
done
```

Using Singularity

```
$ cd ~
$ singularity pull --name ToolA.simg
docker://registry.example.com/ToolA:latest
$ ls
input ToolA.simg
$ ./ToolA.simg
Usage: ToolA [-o {json|yaml}] <file1>
[file2...fileN]
```

Using Singularity

```
$ cd ~/input/files/to/process

$ for files in *; do qsub -q short_jobs -N
"toola_${files}" ~/ToolA.simg ${files}; done
```

In conclusion

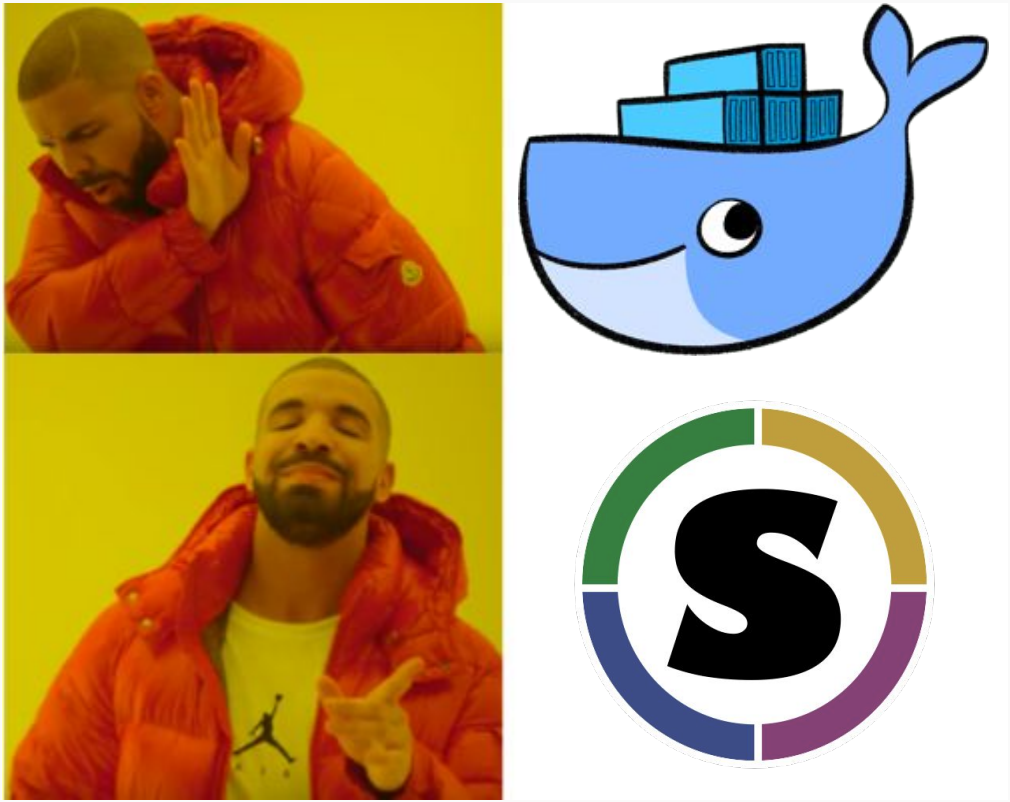

I get no commission from this

https://www.sylabs.io/

Thank you

https://ebur.co