

Efficient unpacking of required software from CernVM-FS

Samuel Teuber, EP-SFT Openlab Summer Student

Nicholas Hazekamp, Jakob Blomer, Gerardo Ganis

Efficient unpacking of required software from CernVM-FS

Samuel Teuber, EP-SFT Openlab Summer Student

Nicholas Hazekamp, Jakob Blomer, Gerardo Ganis

**HOW MANY GERMANS DOES IT TAKE
TO CHANGE A LIGHT BULB?**



**ONE. WE ARE EFFICIENT AND
DON'T HAVE HUMOR.**



Shrink Wrapping

A method for efficiently packaging required software from CVMFS into standalone images

What is CVMFS?

The CernVM File System

- Software delivery network filesystem
- Used for computing in the WLCG
- Blazingly fast for its use case

```
[~]$ ls /cvmfs/cms.cern.ch
bin                               glidein                slc5_amd64_gcc434  slc6_amd64_gcc630
bootstrap.sh                      grid                   slc5_amd64_gcc461  slc6_amd64_gcc700
bootstrapmp                        latex                  slc5_amd64_gcc462  slc6_mic_gcc481
CMS@Home                          lhpdf                 slc5_amd64_gcc470  slc7_aarch64_gcc530
cmsset_default.csh                oo77                  slc5_amd64_gcc472  slc7_aarch64_gcc700
...
```

2 ways of delivering software


CVMFS

- On demand delivery
- No preparation time
- Allows release management through versioning system

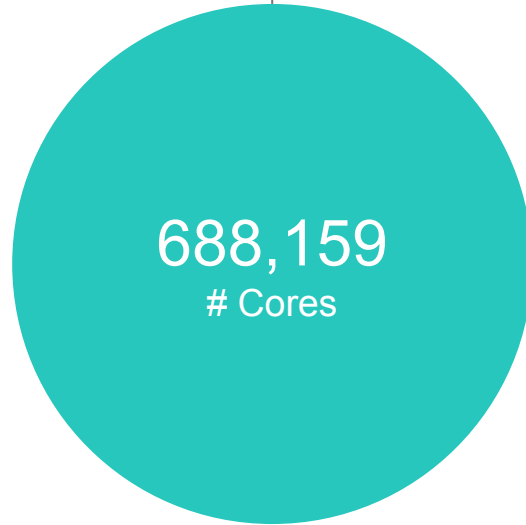
Standalone Image

- No need for:
 - Filesystem in Userspace (FUSE)
 - Internet
 - Hard disk cache
 - Usable as container image
-


Entire **WLCG**



Entire **WLCG**



**One modern
supercomputer
(NERSC, USA)**



- No FUSE
- No internet
- No hard disk

Automated Specification building based on trace file

```
^/bar/etc/*  
/bar/Modules/setup.sh  
/foo/Packages/ROOT/*  
^/foo/Packages/AliRoot/*
```



- No **internet**
- No **disk** cache
- No **fuse** client

Trace by enabling
CVMFS_TRACEFILE
during workflow

Export with
cvmfs_shrinkwrap
(tar, squash, docker, ...)

```
"1532006774980.384", "4", "/lcg/releases/gcc/7.1.0", "readdir"  
"1532006774980.468", "4", "/lcg/releases/gcc/7.1.0/x86_64-centos7", "lookup"  
"1532006774980.624", "3", "/lcg/releases/gcc/7.1.0/x86_64-centos7", "readlink"  
"1532006774980.650", "4", "/lcg/contrib", "lookup"
```

~70 MB/s

Export data rate with warm cache from CVMFS to POSIX folder

Conclusion

Easy tracing utility

Automated specification building

Efficient export

To posix, tar, squashfs, Docker, ...

Experiments are about to test it

In talks with US-CMS, US-ATLAS and LHCb
