



CSC

Expertise from Knowledge

Managing multidisciplinary software repositories for grid with CernVM-FS



20.5.2014 EGI CF 2014
Ulf Tigerstedt, CSC, Finland

FGI : Second Finnish grid

- FGI, the NGI_FI local name was born from the ashes of an earlier grid that had lost momentum a few years earlier.
- We wanted find a better way to distribute software to the clusters in the grid than using rpms or other packaged software.

CernVM-FS

- Was in preproduction at CERN at the time, but seemed stable enough.
- fgi.csc.fi became the first non-CERN cvmfs repository.
- Tricky to set up, since it requires the partitioning to be suitable, the software is not really configurable. This has gotten a bit better, but still requires everything to be just right or it won't work at all.

FGI when it started

- 10 clusters, same environment, same OS (SL6)
- A lot of requested software, some only used by local cluster users.
- Some non-opensource software.

Software

- Compilers (gcc/PGI/Intel/g95)
- Interpreters/script languages (Python, R)
- Libraries (fftw3, Matlab runtime)
- Other systems (Gromacs, BLAST, ORCA)
- Tools (scripts, firmwares)
- Module-environments files

Software distribution when we started

- Software has never been built on stratum 0 server
- Started by having different build machines without cvmfs, but it was hard to build software that required special libraries
- Software mostly arrived as tarballs of varying quality.
- One metapackage rpm needed to pull in libraries from EPEL.

FGI is an national grid, noone is left out (except partially HEP)

- Biosciences: BLAST, tophat, bowtie
- Statistics: R
- Astronomy: Custom python code
- Physical chemistry: GPAW, Gromacs
- Physics: Geant4, fluka
- Engineering: Matlab runtime

Introducing cvmfs-tool

- cvmfs-tool bind-mounts a unique temporary directory from `/tmp/cvmfs.xxxx` over any directory, making it writable to the user.
- It requires limited sudo-rights for the user
- The user can then compile, install and test the software installed in the correct position.
- After the user is happy, it can type “exit” to unravel the bindmount.

Introducing cvmfs-tool

- The contents of the `/tmp/cvmfs.xxxxx` directory can then be sent to the stratum 1 for easy deployment.
- Problems: Sometime the hierarchy needs to be created on the stratum 0, especially for new software. New versions of already published software is easier.

```
/cvmfs/fgi.csc.fi/tools/cvmfs/cvmfs-helper
#!/bin/bash
# Ulf Tigerstedt 2014
```

```
if [ $# -ne 1 ]; then
    echo Only give one parameter: The directory on /cvmfs you want to shadow
    exit -1
fi
if [ -d $1 ]; then
    echo OK: $1 is a directory
else
    echo $1 needs to be a directory
    exit -1
fi
TMP_DIR=`mktemp -d --tmpdir=/tmp cvmfs.XXXXXX`

echo Mounting $TMP_DIR over $1
sudo mount --bind $TMP_DIR $1
echo Now install software into $1, then type exit
bash
echo unmounting $1
sudo umount $1
echo Software is now in $TMP_DIR
```