

## Using CVMFS for software distribution for FGI, the Finnish Grid Infrastructure

Ulf Tigerstedt



The Finnish Grid Initiative was created out of the ashes of mGrid, an old pre-EGI collaboration in Finland for materials sciences. CSC, the IT center for science coordinates the NGI functions, with the sites just having CEs.

- 10 sites, of which 9 are universities.
- Hardware mostly bought at the same time.
- Standardized on SL6, SLURM and ARC.

mGrid used nfs-shares to distribute software, creating a manual process where the admins of the clusters needed to be reminded to install the software. It also required large amounts of shared disk space.

For FGI we had time to plan the deployment before actually getting the hardware to the sites and tested CVMFS that at the time (2011) was quite new.

The tests were carried out with a secondhand bladeserver with fibre channel storage for the data. The only thing noted was that the layout of the server assumes that everything is one big partition.

At the same time we also debated different methods of installing the workernodes and how to keep them patched. We realized that a cluster local squid solved both the problems for kickstart+yum and cvmfs, as the alternative had been local mirrors of SL+EPEL for all clusters. IdCmuWe

We are still on the old version. As long as 2.1 is considered testing it won't be installed on the server.

The replica is synced every hour, and will enable easier downtimes of the master server + eventual upgrade to SL6.

- Multiple compilers (gcc,intel,pgi,g95)
- Multiple Python versions (2.7, 3.0, 3.1, 3.2)
- Multiple MPI versions, compiled with different compilers
- Root, gpaw, Dalton, Geant4, R, shrimp..



The software comes from three different sources:

- Grid support staff
- Site admins
- Users

Some software requires extensive knowledge or dark rituals to compile. We let the experts do it and just deploy it.

- Software deployed and tested
- module-file for environment-modules deployed and tested.
- RTE for grid-runnable software deployed, and copied to the ARC-CE by a cronjob.

- One site has deployed, another is deploying SandyBridge cores. Some software have improved performance by being compiled for that.
- SL7 will arrive at some point.
- Lua-based environment-modules