



CVMFS and Grid Software Installation

*Alessandro De Salvo Alessandro.DeSalvo@roma1.infn.it
6-4-2011*

Outline

- **Why CVMFS in the grid sites?**
- **Support for CVMFS in the Grid Environment**
- **CVMFS setup in Grid Sites**
- **Roadmap**



Why CVMFS in the Grid Sites?

- **Very good option for sites with very limited space in the software area**
 - Only one possible problem, if you have machines (WNs) with limited amount of disk space
 - It's highly recommended to have 25-50 GB of disk cache local to the nodes mounting CVMFS, although you can set this to lower values (tested up to 8 GB per WN)
- **Easy installation in the nodes**
 - Need to provide a local squid in front of the nodes if running more than 20 cores
 - <https://twiki.cern.ch/twiki/bin/viewauth/Atlas/Tier3CVMFS2SLC5>
- **Integration with the Installation System for Grid Sites**
 - CVMFS is a filesystem, nodes still need to be validated
 - Validation (KV run) and tagging can be done with the same system for all the sites, regardless to which filesystem they use
- **A few sites already migrated, others expressed interest in migrating**
 - RAL, QMUL, Wuppertal



Support of CVMFS for the grid sites

- **CVMFS mounts are fully supported by the Installation System**
 - The installation agent (sw-mgr) detects which type of filesystem the software is using and performs different actions
 - Normal shared FS
 - Installation, test, tag
 - CVMFS
 - Test, tag
 - Completely transparent to the sites
 - When a site is tagged as CVMFS in the Installation DB the limit on the maximum number of parallel submissions of validation tasks (normally 15) is automatically removed
 - Sites are tagged automatically as using CVMFS via callback by the first job running with the new configuration
 - Can re-validate and tag mostly all the releases in parallel, provided that there are enough slots available to the sgm users
- **CVMFS support at CERN**
 - CVMFS updated regularly each time a new release is out (or declared as obsolete)
 - May define a new policy for releases in CVMFS (do we really need to remove software quickly from there?)
 - Using the Installation DB as global source of parameters for the deployment of the releases
 - New machine for CVMFS available at CERN as a production system
 - New mount point for CVMFS
 - Previous: /opt/atlas/...
 - New: /cvmfs/atlas.cern.ch/...
 - Need to reinstall the full set of software (production releases, database releases, condition pool files, ...) since the ATLAS software is (currently) relocatable only at installation time



CVMFS structure in ATLAS

■ Old CVMFS structure and contents

■ A mixture of shared fs and pure CVMFS

- Shared fs needed for
 - Local setup (PFC, Frontier, etc.), DQ2 clients, cctools
- CVMFS used for
 - DBReleases, software releases, pool condition files
 - Need to have a few links in place in the software area to CVMFS
 - Supporting i686 software only for grid sites

■ New CVMFS structure and contents

■ Using the same structure as in the current software area of the grid sites

- DQ2 clients, cctools, DBRelease, software releases
- Generic local customizations with possible overrides (see next slides) for
 - PFC
 - Frontier
 - DQ2 local site IDs
- No link needed anymore
- Supporting all software architectures
- The same software used for the grid sites will be used by the Tier3 tools, also in CVMFS
- Minimal customization needed by the site admins



CVMFS setup in grid sites

- **Sites migrating to CVMFS should provide**
 - A local squid cache for the CVMFS clients
 - A local disk cache for CVMFS on each node
 - Suggested minimum: 8 GB
 - Installation of the CVMFS clients on each node
 - No mount points already using `/cvmfs/atlas.cern.ch` !
 - One of the following setups for the local settings (to be discussed)
 - Manual settings (static model)
 - Site admins are in charge of setting the needed local settings in their environment and local disks
 - Frontier settings, dcache/DPM fixes, DQ2_LOCAL_SITE_ID set by the system admins
 - Centrally managed local setup area (semi-dynamic model)
 - Site admins define a variable `ATLAS_LOCAL_AREA` in the WNs, pointing to a directory writable by the SGM users where to store the overrides to the global local configuration in CVMFS
 - The Installation System will take care of updating the local settings each time a new PFC must be rebuilt
 - All jobs will still have to source the `<sw-area-path>/local/setup.sh` file, with overrides from the local area
 - Full dynamic model
 - Site admins do not care of anything and we obtain all the settings from AGIS
 - Probably an heavy load on the central system!
 - The preferred way would be to start with a semi-dynamic model, to possibly migrate to the fully dynamic model in the future



Roadmap

- **Start installing in the new server from next week**
 - A fully reinstallation could take up to 4 days, but this also depends on the time needed to compress and checksum the files and on the number of releases (which is always quite high!)
 - Will start with all the production releases
 - Nightlies will be included as well
 - Condition files to be included but as a separate entity with respect to the software server

- **Tests with the new servers can reasonably start by the end of April**