



ATLAS Software Installation in the Grid Sites

Alessandro De Salvo Alessandro.DeSalvo@roma1.infn.it
2-12-2010

Outline

- **Current status**
- **Efficiency of the system**
- **Recent problems**
- **Future improvements and roadmap**



Current status of the Installation Facilities in the ATLAS Grids

■ LCG

- Autoinstaller running for all the sites
- Releases are pushed automatically by the system in all the sites as soon as an admin defines and sets them to autoinstall mode
 - Usually takes up to 1 day to install a release in all the sites, except when a problem occurs
 - Time needed for a full reinstallation of a site depending on the number of production releases and patches (currently ~180)
 - Manual operations only when an installation fails to install
 - Handling software releases (including nightlies), DQ2, cctools, PFC, Frontier setup
 - User interface (LJSFi), with registration to the site
 - (incomplete) documentation
 - <https://atlas-install.roma1.infn.it/twiki/bin/view/Main/LJSFiEndUserGuide>

■ NorduGrid

- Local installation

■ OSG

- Moving from the pilot-based installation to LJSFi, to use a single system and simplify the deployments and directory structure (impacting also on the pilot code)



Deployment status in the grid sites

- **Several removals started on Nov 15th**
 - 8 releases removed
 - More than 100GB should have been freed

- **Another bunch of removals starting now**
 - 6 releases to be removed
 - Additional > 70GB should have been freed

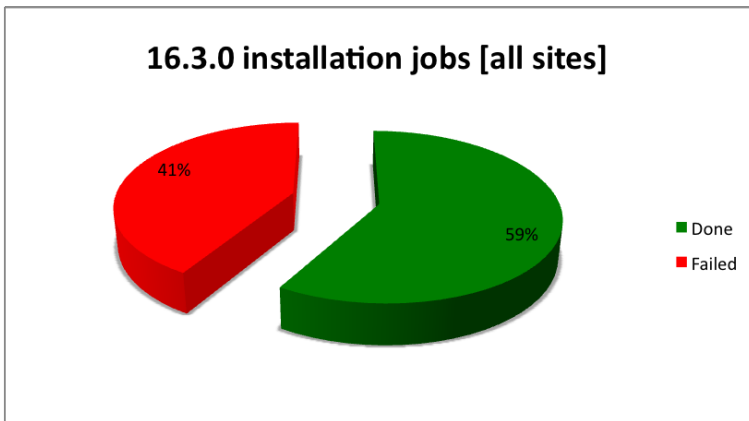
- **Production set**
 - 8 major releases + patches + utility & tools

- **Sw area size**
 - We'll probably be able to live with 500GB sw areas, for the moment
 - ...unless we use CVMFS, in which case the sw area needed size would be a few MB

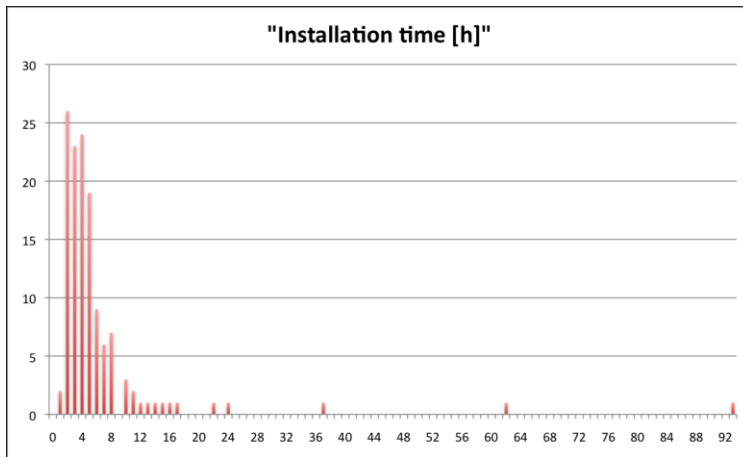
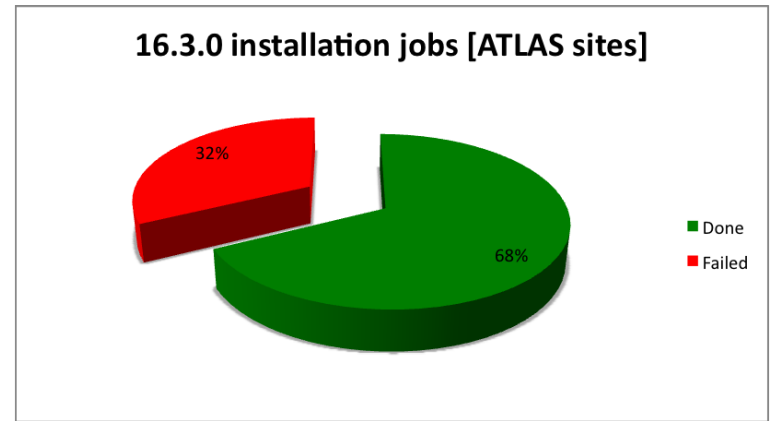


Efficiency

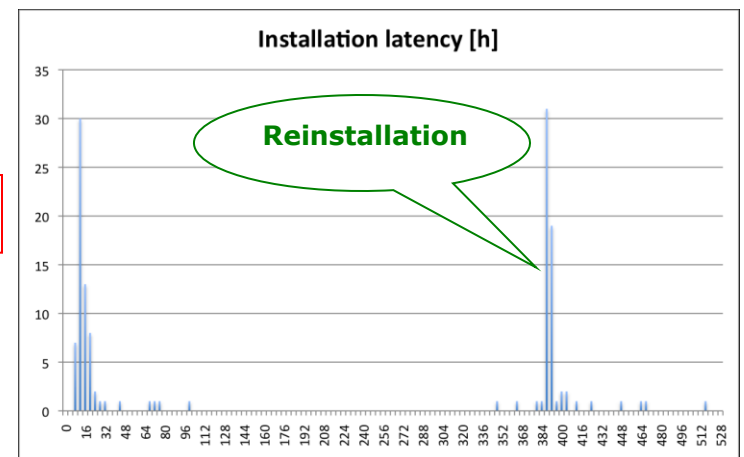
■ Efficiency on the installation of release 16.3.0



All jobs sent




Successful installations





Errors in 16.3.0





OSG migration to LJSFi

- **OSG gatekeepers now visible from the CERN WMS**
 - **Allows a simple integration with the LCG/EGEE system (LJSFi)**
 - Migrating to the same directory structure as LCG/EGEE, while keeping also compatibility links to allow a smooth transition
 - Does NOT need reinstallation of the sites, the releases will only be re-validated and tagged, and the links to the new directory structure will be added
 - **A few sites are already using the new system**
 - BNL-ATLAS
 - IllinoisHEP
 - OU_OCCHP_SWT2
 - OU_OSCER_ATLAS
 - WISC-ATLAS
 - UTD-HEP
 - **BNL was migrated very quickly without any interruption of the service**
 - **Now moving to the T2s, work in progress**
 - A few problems with some sites
 - The decision was to migrate one site by one, to avoid interruptions of the jobs in all the sites, in case of problems
 - **OSG sites using a dedicated instance of the autoinstaller**



Support for CVMFS [1]

- **CVMFS mounts are fully supported by the system**
 - **Currently a mixture of shared fs and pure CVMFS**
 - Shared fs needed for
 - Local setup (PFC, Frontier, etc.), DQ2 clients, cctools
 - But in the next future both DQ2 and cctools will be available from CVMFS too
 - CVMFS used for
 - DBReleases, software releases, pool condition files
 - Need to have a few links in place in the software area to CVMFS
 - A wiki with instructions on how to migrate T2 and T1 sites to CVMFS is under preparation
 - Directory structure, links, best practices, ...
 - **The installation agent (sw-mgr) detects which type of filesystem the software is and perform different actions**
 - Normal shared FS
 - Installation, test, tag
 - CVMFS
 - Test, tag
 - Completely transparent to the sites



Support for CVMFS [2]

- **CVMFS support at CERN**
 - CVMFS updated regularly each time a new release is out
 - ...but it's **NOT** yet a production service for CERN
 - Pushing in this direction
 - **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
 - **NOT a replacement of the current Installation System**
 - 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
 - <https://twiki.cern.ch/twiki/bin/viewauth/Atlas/Tier3CVMFS2SLC5>
- **A few sites already migrated, others expressed interest in migrating**
 - RAL, QMUL, Wuppertal



Support for x86_64 releases

- **The installation agent is already able to install x86_64 releases, but there are problems setting up a pure x86_64 release**
 - **Trying with 16.0.3, apparently there are no x86_64 executables in it**
 - Is it normal?
 - Anyone else see the same problem?
 - **To be followed up before we can deploy any x86_64 release in the grid sites**
- **Will install i686 and x86_64 executables in the same directory, but will create different entry points for the users**
 - **Currently in the grid you can auto setup via \$VO_ATLAS_SW_DIR/software/<version>/setup.sh**
 - Can be used to setup both i686 and x86_64 releases
 - Defaults to setting up i686 AtlasOffline <version> via AtlasLogin or AtlasSetup, depending on the release number
 - **Will keep the same for i686 releases and will create other setups for i686 and x86_64 releases, separately**
 - \$VO_ATLAS_SW_DIR/software/i686_slc5_gcc43_opt/<version>/setup.sh
 - \$VO_ATLAS_SW_DIR/software/x86_64_slc5_gcc43_opt/<version>/setup.sh
- **To be tested in the coming weeks**



Recent problems & solutions

■ 16.2.1 & 16.0.3

■ 16.0.3 was installed by mistake in the same directory as 16.2.1

- This should be in principle supported, but in practice it's not
- Both the setup of 16.2.1 and 16.0.3 was not working in this situation
- Tracked to be a problem with AtlasSetup, solved by David Quarrie, but not yet in the grid sites
- The solution was to (try to) install 16.0.3 in a different directory

■ ...but then again problems when installing 16.0.3 on a different directory

- This has been working so far, so it was a big surprise
- Tracked to another bug of AtlasSetup (or maybe a variant to the same as before)
- Only affecting the installations done via LJSFi, given the different directory structure
- NOT affecting CVMFS
- Fixed again in AtlasSetup, but still need to understand how to fix it in the grid site without reinstalling 16.2.1
- For the moment 16.0.3 cannot be installed until we find a solution

■ Old problems solved

■ WMS instabilities

- The WMS is working fine, even at “high” load (2k jobs/day)

■ Autoretry engine for failed installs/removals

- The autoretry and garbage collector minimize the manual interventions
- A few race conditions to be fixed for sites with high number of failures



Improving the system

- **Explore the possibility to add Panda as an additional backend system, as an alternative to the WMS**
 - **Need to sort out a few features that can be only done via the WMS at the moment**
 - Dynamic discovery of sites
 - Full brokering based on the published software tags
 - **Panda can be configured as an additional plugin to LJSFi**

- **Extend the support of the installation agent (sw-mgr) to alternatives to pacman/pacballs**
 - **sw-mgr can use whatever installation backend is available**
 - **At the moment direct pacman packages (via snapshots) and pacballs are supported**
 - **If we really need to replace pacman we need to find a valid alternative**
 - Handling the dependencies
 - Working in all the OS variants used in ATLAS (including debian and ubuntu)
 - Having a “snapshot” feature (to reproduce the pacballs behavior), if we want to keep this feature