

UAM issues with Quattor

Luis Fernando Muñoz Mejías

Universidad Autónoma de Madrid

NIKHEF, 10/2008

- I applied for a fellowship at CERN
- I'll know my future in November

- Upstream documentation is hand-written
 - The process is slow and boring
 - Often not consistent with Pan and Perl code
- The problem is much worse when configuring a site
 - Wiki pages get outdated really fast
- We'd love to see a tool to automatically extract documentation from comments
 - What about a Pan parser for Doxygen?

- They are too bloated
- On a normal amd64 WN, over 1,200 packages get installed
 - This makes OpenVZ installations fail because of the transaction size
- Can we use checkdeps.py to reduce the amount of packages?
 - On the same WN, I reduced the number of packages to 411, including Quattor and custom packages

SCDB uses centralized VCS

- We grab QWG via `svn:externals`
 - It's great for updates
 - But if LAL's SVN goes down, we cannot compile anything
 - If QWG is moved to Sourceforge the performance may be too bad
- We'd like to implant new ant targets
 - If we modify the `build.xml` we'll have to manually merge changes when upgrading
- The staging strategy fits better on a distributed model
 - Each stage is a tree to pull/push changes
- How M-S-specific is AQDB?
- How difficult would it be to have a CDB based on Git or Mercurial?
 - Is this a good project for SoC?