

Quattor Asset Database

*QWG Meeting
4-5 November 2009, Brussels, BE
Christos Triantafyllidis ctria@grid.auth.gr*



What is Quattor Asset Database

- This was the “best” name we could think about an application that
 - will interact with Quattor templates
 - will aggregate profile information to others (local DB cache)
 - visualize the stored information
 - ...

The need

- At first (as it was presented last year)
 - A web interface to allow (junior) admins to do basic operations
 - Adding new users
 - Configuring new nodes based on existing machine-types
 - Changing an IP or a node role
 - A level of separation of the “operations” part of the infrastructure from the “development”

The need

- Next was fulfillment of the requests
 - Where is the X node located?
 - What guests does this XEN host?
 - Which switch port does this node use?

- Can i have a view of this Rack?

The need

- And lately was the need to feed our monitoring system with “preprocessed” data
 - This was the real need that pushed the development
- more on this on my (last?) presentation regarding monitoring

The implementation

- We used Ruby on Rails as development language
- Very good interface with XML (profiles)
- Perfect database management
- Rest interface to interact with clients
- Supported by our machine-type/webserver :)

Schema extensions

- (Re)introduction of the “/monitor” structure
 - more on this on the monitoring presentation
- Added at structure_rack:
 - position
 - size
- Added structure_switch (switch) at structure_nic:
 - name
 - port

Subversion access to QAD?

- This was a “hot” topic for us
 - If we allow svn access then we will be able to parse configuration files...
 - ... or even allow modifications
- How will it be in sync?

Other solutions?

- Probably the best solution to achieve this would be:
 - have the QAD modify a database
 - have a SCDB ant task that will sync changes from the database to the trunk

- A lot of effort for us

Subversion for now

- We've chosen to allow SVN access to our QAD installation
 - Sync is done as a post-deploy hook at our SCDB server
- Works fine for now as our focus is on processing the profiles for the monitoring tool

Proof of concept...

- Demo

Thank you

Questions?