

gLite WMS


status & near, mid-term plans for a sustainable future

Marco Cecchi

INFN - CNAF




EGEE-I



WMS 1.5 for SL3
WMS 3.0 for SL3
WMS 3.1 for SL3

EGEE-II



Patch #1726: WMS 3.1 for SLC4 ia32 (so as a savannah patch) – certification took place at Imperial College

- general performances and robustness improvements
- bottlenecks analyzed and fixed
- improved job submission rate and service stability

Patch #1841: WMS 3.1 for SLC4 ia32 (mega-patch, first clean ETICS build)

EGEE-III

Patch #2562/2923: WMS 3.1 for SLC4 ia32 (in production)

Patch #2597: WMS 3.2 for SL4 ia32 – performance, stability, race conditions... - in certification at IN2P3/ALICE – also builds on x86_64

Patch #2875: WMS-UI 3.2 for SL5 x86_64 (in certification) - first 64-bit ever

- **Bug fixing (savannah)**
 - ~40 Ready for test/certified/closed savannah bugs (i.e. understood, applied fix and/or definitely closed) from patch 2562 on.
- **Support (GGUS)**
 - ~50 solved/answered tickets (<10 are open at the moment)
- **Direct support to experiments**
 - frontline requests from SA1, experiments etc.
- **Created Website...**
 - <http://web.infn.it/gLiteWMS/>
 - as an official meeting place for support, information, documentation, papers etc.
- **...twikis, and...**
 - Deployment & configuration HOWTOs
 - Documentation
 - Test reports

- **GPC09: Grid and Pervasive Computing 2009**

- IWWM09: **Workflow management**

- WGCV09: **Grids, clouds & virtualization**

- SSDU09: Symposium on service, security & data management technologies

Invited speakers:

- Ian Foster

- David Anderson: seti@home & BOINC

- Bernhard Schott: Platform LSF EU-Research Program Manager

Virtualization, clouds and green scheduling

- David Bernstein: VP & GM of Network Applications Infrastructure for Cisco Systems

“Cloud Interoperability at Internet Scale”

It will either be **history repeating**, or our collective manifest destiny, to evolve Cloud Computing to a worldwide, interoperable, transparent platform. In other words, *Cloud will become to Computing just what the Internet is for Data.*

History repeating or reinventing the wheel? Wasn't the Grid supposed to be the 'computing Internet'?

Dorian Gorgan: <http://www.envirogrids.net/> (FP7) is starting using gLite and our WMS, but is he looking for?

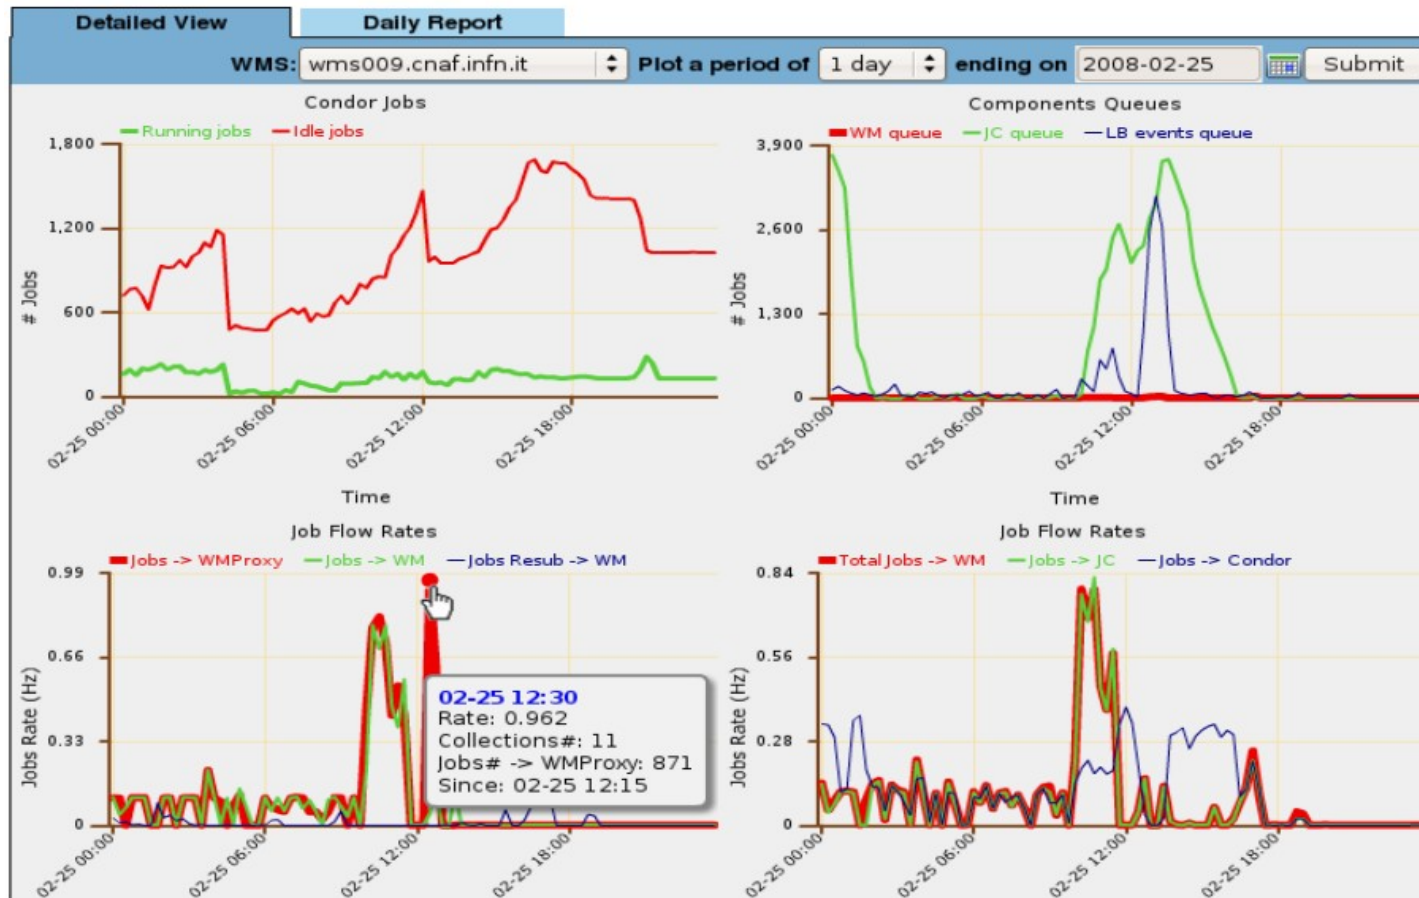
- Moving towards UMD
 - Inputs from the latest JRA1/SA3 all-hands
 - ***“Additional challenges: Do the most with the resources we will have!”***
 - ***Professional middleware for production use***

- Savannah bugs posted since project day 1 by category
 - *Assigned to: any, open/closed: any, status: any*
 - BDII : 60, Information Providers: 150
 - BLAH : 72
 - CEMon: 48
 - CREAM: 166
 - Data Management: 784
 - LB: 197, JP: 8, **LB+JP: 205**
 - UI: 58, WMS: 858, **WMS+UI: 916, WMS+UI+LB+JP=1121**
 - Security (VOMS, but also WMS, DataMgmt, LCAS, GLEEXEC, LCMAPS, CREAM): 909
 - Build system: 192
 - Glite middleware (general): 208
 - ...

- Moving towards UMD
 - Inputs from JRA1/SA3 all-hands
 - Additional challenges: Do the most with the resources we will have!
 - Professional middleware for production use
 - **Implement project priorities: portability, IP v6, ...**

- Plots taken from WMSMon of a production instance of a WMS 3.1 patch 2562
 - typical CMS use-case. Job submission based on collections (1-1 WMS-LB).

- Neither LB 2.1 nor LB 2.2 will address the issue of optimizing how to forward events to the LB server
- Bottleneck causing huge backlogs to JC
- Jobdir as input for JC (WMS 3.2) should help
- How does ICE deal with this?



- **WMS 3.3 for SL5 x86_64 (and ia32)**
 - Restructuring of the purchaser(s) (removed Idif2classad/ii_attrutils wrapper libraries/system Idap)
 - Migration to the new LB 2.0 (basically an API change)
 - Support for GridSite (1.5) delegation 2.0
- **Testing submission to CREAM via ICE**
 - Stress tests (ICE+JC/LM separately, together) are very important, and need being documented
- **Define (&automate) more formal testing procedures (not yet as a savannah 'task')**
 - This is a key issue which should be given a higher priority
- **Implement 'cancel' for collections**

Oddly enough the most used/encouraged job type do not support a 'cancel'
- **Support IPv6**
- **Address platform portability**
 - Very important, especially for the UI. But we also need to evaluate how the increasing role of virtualization will deal with this, especially for the WMS

- **Integration with the new gLite authorisation framework**
 - It will only affect authZ to submit to a given WMS
- **DNS-based load balancing and client failover**
- **Unifying & synchronizing WMS&ICE (yaim) releases**
- **Mechanisms to forward requirements/parameters from the JDL to the LRM (not yet as a savannah 'task')**
 - via Globus RSL for lcg-ce, but especially via classads to ICE for CREAM/BLAH
 - tune up the execution environment at the LRM side by adding/forwarding #BSUB / #PBS directives (LFS/PBS)
 - *Virtualization*
 - *Consumable resources (i.e. licensing)*
 - *Interoperability (i.e. with ARC, hard-coded at the moment)*
 - *MPI (hard-coded at the moment)*
 - *HPC?*

– Moving towards UMD

- Inputs from JRA1/SA3 all-hands
 - *“Additional challenges: Do the most with the resources we will have!”*
 - *Professional middleware for production use*
 - *Implement project priorities: portability, IP v6, ...*
 - *Change the logging format: break operations*
 - *Change an interface (WSDL, API, CLI): break applications*
 - *Change a library: force unexpected upgrades on partners*
 - *Clean up error codes & messages in all components*
 - **Are we sure this would not break something either?**
 - *Reduce dependencies (review the build system components/VCS)*
 - **What have we all being doing during EGEE-II? What about the dependency challenge?**
 - *Must be able to make releases with JUST bug fixes*
 - *Critical issues resolved by upgrading to ‘old’ release*
 - but also think of more "elastic" ways of handling similar issues

- End of EGEE-III: work at our best to make of the WMS a stable, performant, usable, service for production-quality Grids
 - Focus on finalizing missing functionality for the end of EGEE-III (bottlenecks, delegation 2, Ipv6, load balancing, ICE/CREAM)
 - Support (CMS in primis)
 - Extend platform portability
- **Moving towards gLite consortium/UMD will require working in group even more and professional way. We suggest customer/supplier model even between ourselves. Quality criteria are more than needed.**
- Seriously start thinking of workflows (simply as DAGs for CREAM?)
- Keep an eye to how virtualization and clouds will evolve