



biomed VO

ongoing operational issues

EGI Technical Forum 2012
Franck Michel, Tristan Glatard
CNRS



- Average waiting jobs/running jobs = 3.5

(source Gstat2)

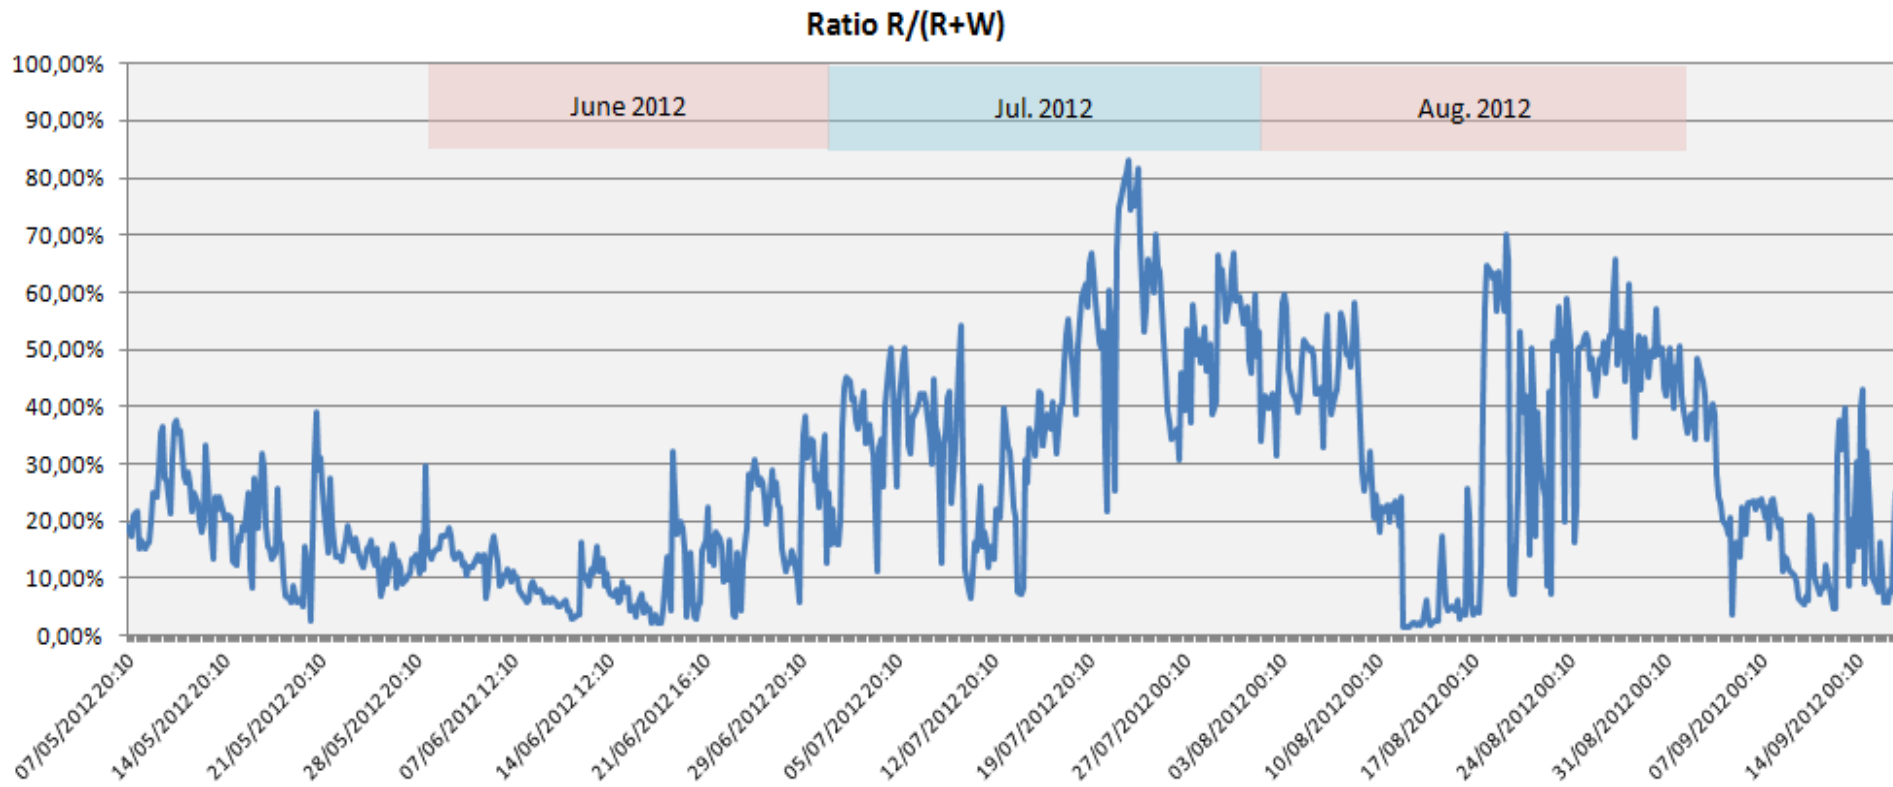


- Possible causes?

- Increase in number VOs' requests, possibly with higher priorities?
- Variability in utilization rate of resources centres?

- Better usage of resources
 - Analysis of Waiting/Running jobs continues
 - Promote the use of a DIRAC multi-VO instance (provided by french NGI)
 - Replanning functionality of WMS
 - VO-wide white list of resources
 - Inefficient applications: contacted users
- Increasing resources
 - EDGI has open 1 CE for biomed (3K to 4K cores)
 - 1 biomed-dedicated site being certified

- Analysis of the evolution of Waiting vs. Running jobs on all CE queues supporting biomed
- Average $W/R=3.9$ over last 12 months (GStat2)



- Ongoing work, need for per-CE measures
- Preliminary results:
 - 42% of queues with ratio $< 50\%$ ($W > R$)
 - 35% of queues with ratio $\geq 50\%$ ($R \geq W$)
- Outlined particular cases:
 - 22% of queues not used at all ($W \sim 0, R \sim 0$)
 - CE not working, does not fit with users' requirements, wrong capabilities published, black listed...
 - 24% of queues with jobs waiting forever ($W > 0, R \sim 0$)
 - Starvation in the queues, due to low priority?
- **Realistic/feasible to give opportunistic VOs a minimum priority?**

- User groups implement strategies to build white list of CEs that feed their job submission system
- GOCDDB vs. BDII CE status: ongoing discussions to figure out if inconsistency between both is actually a problem
- 2 approaches:
 - Optimistic: assume that on most sites the BDII status can be reliably trusted, with the risk that a few jobs will fail when the status is wrong (considering this should be rare)
 - Pessimistic: VO-level approach, take into account any available information to build up a white list of CEs, with the risk of ignoring working CEs
- Decision still requires more inputs

- Ongoing discussion on how to deal with LFC/SE inconsistencies
 - Problems: zombies (aka. dark data), ghosts
 - SYNCAT (synchronization LFC/SE)
 - Prevents the creation of inconsistencies but not appropriate for curative solutions
 - No background research of inconsistencies
 - Investigation of LHC VOs existing solutions (LFC dumps, compare...)
- VO requirement to supporting sites: min. 1TB SEs
 - Dealing with full SEs = high operational cost
 - When a close SE is full, jobs efficiency will decrease

- LFC: WLCG plans to stop usage in 2013
 - LFC key component for LSGC VOs
 - LSGC willing to contribute to a join DPM/LFC community driven project
- Migration to EMI WNs
 - Test LSGC apps on EMI to be done: install of an EMI UI unsuccessful for now
 - Schedule for the migration of WNs to EMI: when? To which version?

A yellow sticky note is pinned to the white background with two red pushpins. The words "Thank you" are written in a black, cursive font on the note.

Thank
you

LSGC wiki	http://lsgc.org/en/LSGC:home
Support team email	biomed-technical-support [at] googlegroups [dot] com
VO Manager email	biomed-vo-managers [at] googlegroups [dot] com