

Monitoring Site Perspective

Alessandra Forti
Ops TEG
23 January 2012

Fabric Monitoring

- Sites use a variety of tools
 - Nagios, ganglia and lemon
 - Cfengine, puppet, cron jobs
- Not all sites have good or enough probes
 - Sharing/repository would be good
 - It was attempted in 2007 on request of Hepix and WLCG.
 - <http://www.sysadmin.hep.ac.uk>
 - sys admins liked the idea of having it but didn't contribute
 - DPM has also an admin contrib
 - Not sure how much that is used either

Network Monitoring

- There is currently no infrastructure
 - Atlas sonar tests were the first push to do something for many sites
- Perfsonar should be pushed at all sites
 - At the moment it seems to be an option
- But also other tools should be encouraged
 - “last mile” or even the last meter problem
 - Backbone problems
 - Many tools all of them useful for slightly different things
 - From monitoring each interface to aggregated traffic
- Avoid the ping pong effect

Middleware monitoring

- There are too many pages
- Currently done by submitting jobs to verify the functionalities
 - This tries to test the site but actually uses the whole infrastructure
 - Difficult to separate site from infrastructure
- Some errors due to middleware bugs
 - DPM misreporting space if a data server is down
 - Space is actually available but reported as negative

Middleware monitoring

- More monitoring should be done also on the machines themselves.
 - There is very little to do that
- (Error) logging is a never ending story
 - Even when there are log files it is difficult to parse them
 - Tools provided by the developers to parse their log files
 - Not only for errors but also for statistics and security.
 - DPM tools and cvmfs have nagios probes
 - Others should follow the good example
 - And site should use and give feedback

Experiment Monitoring

- There are way too many pages
- Can they really be unified?
 - Not easy, for some reason each page was developed with a slightly different technology
 - Not only between different experiments but also within the same experiment
 - Historical dashboard and DDM2 are two completely different beasts.
- Too difficult to go back for the existing
- Any future development there should be more coordinated

Experiment Monitoring

- Site error, infrastructure error, or unknown?
 - If it is not a site error it shouldn't count in its efficiency
 - Not easy I know but it is important to separate the two
 - Site problems are more visible because they have peaks but there a variety of other errors whose cumulative effect is important.
 - 57% of atlas prod errors unknown or not-site errors last year
 - There should at least be 4 categories
 - Site, Infrastructure, Cancelled, Unknown
- Currently this deficiency is covered by asking sites to do a report and remove the errors that were not theirs
 - It is a good exercise for the site admins but it requires someone's time

Conclusions

- There are too many pages (have I said that?)
 - Coordination for the future is the key
 - SSB is a good effort at having a common display
- Mandatory the content of any tool is reliable and validated
 - Problem is not only visualization
- It is important that whatever monitoring reflects site issues and not the whole infrastructure
 - Particularly if used to calculate efficiency/availability/usability of the site
- Network monitoring infrastructure becoming more urgent