


Middleware Deployment

Simone Campana CERN IT-ES

Shawn McKee AGLT2

December 2012 WLCG GDB



Where to install perfSONAR

- Ultimately, all WLCG sites should install and configure perfSONAR
 - Two hosts, one for Latency, one for Bandwidth tests
- But pragmatically, we started from sites indicated by experiments
- A list of sites can be found here
 - <https://docs.google.com/spreadsheet/ccc?key=0AmoqjvJA--ajdGlZdzZjSUNfbXJxUDlZ0c4b0FQZnc#gid=0>
- Sites are grouped in “regions” or “clouds” (experiment specific)
 - Some or a lot of overlap, depending on the regions.

perfSONAR configuration

- So far, tests were configured in each perfSONAR server
 - OK for a pre-production service, very complicated to maintain in the long run
- We tested successfully the “mesh” configuration method in US and IT
 - Central configuration files managed by few experts
 - An example:
 - <https://grid-deployment.web.cern.ch/grid-deployment/wlwg-ops/perfsonar/conf/itcloud.conf>
- The mesh configuration is not natively built into the current perfSONAR release
 - The installer should run a simple bash script, no issues reported by beta testers
 - The next release (Q1 2013) will support the mesh conf natively
- The perfSONAR mesh configuration is being further improved by Internet2

perfSONAR configuration

- For deployment strategy we intend to:
 - Ask sites where perfSONAR is already installed to move to the mesh configuration
 - Ask sites where perfSONAR is not already installed to go directly for the mesh configuration
- For the timescale we intend to
 - Ask sites NOT to wait for the new release
 - So, use the current one and complement with the script to complement for mesh config
- Sites should register perfSONAR in GOCDDB/OIM
 - Two service types for Latency and Bandwidth

Scope of the tests

- First priority scenario: intra-region tests
 - One mesh per region, Bandwidth, Traceroute and Latency
 - Plus the mesh of OPN sites
- Second priority scenario: T2s to T1s tests
 - Each regional mesh against the OPN mesh
- Third priority scenario: inter-region tests
 - The WLCG mesh: Each T2 against each T2
- Parameters of each scenario will be different
 - First priority tests will be more “demanding”

Parameters of the test

- Will configure Latency, Bandwidth and Traceroute tests
 - Recommending a "standard" version of all three tests
- Latency tests
 - 10 UDP packets/sec, 14Bytes/packet
- Traceroute
 - 1 test every 10 minutes
- Bandwidth
 - 30 second tests every N hours
 - N will depend on the scenario of the previous slide
 - Currently N=6h for intra-cloud tests

Manpower

- Some regions very involved in the TF
 - IT, UK, US
- And the result is obvious:
 - Basically all sites deployed perfSONAR
 - IT and US already using mesh
- We need involvement from **all** regions
 - Building/Maintaining config files
- So, we ask 1 name from each region/cloud
 - 1 person for all experiments is OK
 - if there is a good coordination at national level across experiments
 - Otherwise 1 person per region per experiment
 - At least ATLAS and CMS. ATLASsites+CMSSites approx includes all LHCbSites

Next Steps

- Documentation being streamlined
 - All there, but need to condensate in a single page
 - Timescale is a couple of weeks
- From January, the large deployment campaign will start
 - Broadcasts will be sent to sites
 - By then we expect to have the names of contact people we asked for
 - Those people should help with the first priority scenario
 - A few members of the TF will centrally start configuring meshes for second and third priority scenarios

Contacts

- perfSONAR deployment
 - wlcg-ops-coord-tf-perfsonar@cern.ch