



January 31st 2012 – LHCONe Meeting

Shawn McKee – University of Michigan

John Shade – CERN

Jason Zurawski – Internet2

Monitoring Status/Wish List

Where we are

- Deployed infrastructure
 - At the T0, T1s, many T2s
 - Supporting Networks
 - Many metrics
- Dashboard to show status
 - Between sites
 - Groups (OPN vs ONE vs clouds [US, IT])
- Getting 'used' to operational use cases
 - What it tells us about the network
 - How to use the tools for 'at a glance' performance vs using them for diagnostics*

* N.B. Jason has been making extensive use of all machines in debugging international performance issues – it has cut down debugging from months to weeks.

Where we are going

- Fully configured infrastructure
 - Changes to one (e.g. adding a new host) need to be reflected elsewhere
 - Upgrading should not break things
- Free of defects
 - Firewalls are still a factor, at many sites
 - Hosts 'may' be (Dell config vs 'spare' machines)
 - Where are these in the infrastructure? Want them on the border device, is that true for all?
- On the paths that matter
 - Need backbones, regional networks, XPs, NRENs involved. We are on our way, still not 100%
- With tools that make sense for the community
 - The dashboard is popular – we should release this as something that can be used installed by all (Question to BNL)

Where we need to be

- Debugging real problems
 - Already close this being done – these tools are useful for fixing problems
 - Help solve around 5 problems a year in USATLAS since we started using them
 - Fixed TA links (twice)
 - Currently doing things between US T2s and EU T1s.
 - Teach others how to use them (?)
 - Workshops for the operators? WLCG involvement?
- Ensuring performance
 - Alarms/GUIs for the pro-active positions. Maybe the current dashboard, maybe something new (what does the WLCG say – and could they help?)
- Enabling science
 - Monitoring is not exciting – so lets fix the problems and get out of the way ☺

Where we could be

- Where to report problems?
 - Where do they go now, the WLCG? Local networking staff? Other? It's a hard thing to answer...
 - Is a central location valuable? Is it something the LHC should run?
- Who fixes them?
 - Local staff ideally, but on an end to end problem, there are multiple domains, all pointing at each other
 - Does the LHC need a 'pert'? Where could it be funded and who would staff it?
 - Something like this in the EU already
 - Nothing like this in the US (ESnet/Internet2 often collaborate with regional partners to fix problems reported to us).