


# Network and Transfer Metrics WG Meeting

Shawn McKee, Marian Babik

Network and Transfer Metrics WG Meeting  
8<sup>th</sup> April 2015



# Outline

- Next Meetings/Events

- 6 May, 3 June, 8 July, 2 Sept - all at 4pm CEST
- CHEP – talk attached in indico, please send your comments
- CHEP paper – deadline 17<sup>th</sup> May – input from experiments needed
  - Input to use cases document still missing !

## Today

- perfSONAR status
  - Mesh configuration changes
  - Operations status
- Network performance incidents
- esmond status/plans
- Integration projects



# perfSONAR Status

# Mesh Configuration changes

- Added DebugMesh – AGLT2 – SARA
- Added Belle II mesh – lead by Malachi Schram
- Added Global mesh – aimed at integrating Tier3s
- Removed BNL from certain meshes to decrease the load – now only in US ATLAS and WLCG
- Changed parameters for the default latency tests – increased sample count from 300 to 600
- Dual-stack mesh
  - Currently using same parameters for IPv4 as default – shall we merge ?
  - Test frequency can be lowered if necessary – any update ?
- UK mesh contains two nodes for Manchester
- DE mesh contains UKI-NorthGrid-Shef-HEP
- Configuration interface at <https://oim.grid.iu.edu/oim/meshconfig>
  - You need to be registered in OIM and authorized to access
  - Please contact me or Shawn if you have issues

# perfSONAR operations

- 3.4.2 release status
- [http://grid-monitoring.cern.ch/perfsonar\\_report.txt](http://grid-monitoring.cern.ch/perfsonar_report.txt):
  - WLCG perfSONAR service status report on 2015-04-08 04:02:24.048580 =====
  - Active perfSONAR instances: 233
  - Registered/monitored perfSONAR instances: 259
  - perfSONAR-PS versions deployed:
    - 3.4.1 : 33
    - 3.4.2 : 172
    - Unknown: 26
  - Incorrectly configured (failing >4 metrics): 26
- Please check status of sonars in your region

# perfSONAR operations

- For LHCOPN/LHCONE on-going investigation if sonars are consistently delivering metrics
  - Significant improvement observed after 3.4.2
- Wrote data completeness check
  - Active sonars (participating in a mesh): 233
  - Latency sonars (OPN mesh): 13
  - Theoretical size of full mesh: 156
  - Total number of working links (both directions): 156
  - Ratio: 100.00
  - Bandwidth sonars (WLCG mesh): 110
  - Theoretical size of full mesh: 11990
  - Total number of working traces (both directions): 9690
  - Ratio: 80.82
- Plan
  - Start with a planned top-k ramp up in the WLCG Latency

# perfSONAR operations

- Security – CVE
  - CVE released 2nd of April 2015 for cassandra, which is used by the perfSONAR measurement archive software, esmond. NO action required to protect perfSONAR Toolkit since vulnerable ports are both disabled and firewalled.
- Infrastructure monitoring
  - Added check\_memory (checks 4GB minimum requirement)
  - Added 3.4.2 as minimum version
  - Increased time-range for bw tests to 7200 seconds
  - Already in test, should move to production tomorrow



# Network Incidents Follow up



# Network Incidents Follow up

- Discussed at the WLCG ops coordination, agreed to start with it as is with possible modifications once we gain more experience
- Experiments
  - Report to mailing list (wlcg-network-throughput) or GGUS SU (TBD) – GGUS SU will be backed by the mailing list (initial members are experiments and transfer contacts and WLCG perfsonar support unit)
  - WLCG perfSONAR support unit to confirm if this is network issue - contacts sites
  - Concerned sites report to their R&E informing mailing list
  - List of on going incidents on WG page ([https://twiki.cern.ch/twiki/bin/view/LCG/NetworkTransferMetrics#Network\\_Performance\\_Incidents](https://twiki.cern.ch/twiki/bin/view/LCG/NetworkTransferMetrics#Network_Performance_Incidents))
- Sites
  - Report to their NREN or provider, inform mailing list
  - Escalate to WLCG operations coordination to resolve inter-site issues



# Datastore/esmond and Pilot Projects

# esmond Status/Plans

- Validation work on-going - getting metrics to check accuracy and coverage/completeness of the data collection
  - Check if content in esmond matches local MAs (via sampling)
  - Check if esmond content is accurate

# Experiments Interface to Esmond

- RSV probes (OSG) – collecting metrics
- esmond (OSG) – datastore
- esmond2mq – Henryk developed a prototype in python
  - Update from Henryk TBD
  - Retrieves all data (meta+raw) from esmond depending on existing mesh configs
  - Publishes to a topic
- Proximity/topology service
  - Work on-going to get initial version ([proximity.cern.ch](http://proximity.cern.ch)) – started with site-based mappings
  - Fetch active SEs from FTS and map them via SITEs to perfSONARs
  - Fetch SEs from site topology and map them to perfSONARs
  - Mapping/translation of services (service to service; storage to sonar, sonar to storage), service to site (sonar to site) to be accessible via API (JSON)
  - Plan: Test different algorithms (site mapping, traceroutes, geoip)
  - Evaluate if existing tools can be reused for this purpose

# FTS performance

- FTS - low level data movement service – used by majority of WLCG transfers
- Current granularity and coverage is a good match to perfSONAR network
- Integration use cases available from FTS, but also from experiments
- ATLAS is doing a study (FTS performance study) that already integrates FTS and network metrics (traceroutes) – provides excellent starting point and can be extended to CMS and LHCb
  - Contacted CMS and LHCb, CMS interested, still waiting for LHCb
- Contacted Saul, Tomas and Hassen
  - Work is on-going to get common source of data for both FTS dashboard and FTS performance study
  - Outcome of the study can be very useful as an input to the FTS dashboard roadmap
- To be discussed in detail at the next meeting (May 6<sup>th</sup> 4PM CEST; <https://indico.cern.ch/event/382623/>)

# AOB



# esmond2mq

- ----SUMMARY----
- Fetch metadata: 1177 sec, avg. 0.065830 sec/entry
- Fetch raw data: 69335 sec, avg. 2.557548 sec/event, 30 threads
  
- Metadata processed: 17893 (avg. 4.886909 entries/sec)
- Events processed: 27110 (avg. 7.404241 events/sec)
- Types: {'packet-trace': 12686, 'packet-count-sent': 3545,
- 'packet-count-lost': 3518, 'histogram-owdelay': 3394, 'throughput': 452,
- 'packet-loss-rate': 3515}
  
- Queues:
- Preprocessing: 316
- Postprocessing: 30
  
- Messages sent: 77
- Types: {'packet-trace': 64, 'packet-count-sent': 4,
- 'packet-count-lost': 3, 'histogram-owdelay': 2, 'throughput': 1,
- 'packet-loss-rate': 3}
  
- Run time: 3661 sec
- -----