


Integrating Network and Transfer Metrics to Optimize Transfer Efficiency and Experiment Workflows

Shawn McKee, Marian Babik for the
WLCG Network and Transfer Metrics Working Group

CHEP 2015, Okinawa, Japan
13th April 2015



Network Integration Motivations

- A crucial contributor to the success of the massively scaled global computing system that delivers the analysis needs of the LHC experiments is the **networking infrastructure** upon which the system is built.
 - The LHC experiments have been able to exploit excellent high-bandwidth networking in adapting their computing models for the most efficient utilization of resources.
 - However there are still challenges and opportunities to make this even more effective for our needs.

Evolving Network Capabilities

- We already have deployed End-to-end monitoring of our networks using perfSONAR.
 - perfSONAR, combined with data flow performance metrics, further allows our applications to adapt based on real time conditions.
- New, advanced networking technologies are slowly becoming available as production equipment is upgraded and replaced
 - Software Defined Networking(**SDN**) holds the potential to further leverage the network to optimize workflows and dataflows
- We foresee eventual proactive control of the network fabric on the part of high level applications such as experiment workload management and data management systems. **We must prepare for this...**

Network Monitoring in WLCG/OSG

- Goals:
 - Find and isolate “network” problems; alerting in time
 - Characterize network use (base-lining)
 - Provide a source of network metrics for higher level services
- Choice of a standard open source tool: perfSONAR
 - Benefiting from the R&E community consensus
- Tasks achieved:
 - Monitoring in place to create a baseline of the current situation between sites
 - Continuous measurements to track the network, alerting on problems as they develop
 - Developed test coverage and made it possible to run “on-demand” tests to quickly isolate problems and identify problematic links

perfSONAR Deployment

http://grid-monitoring.cern.ch/perfsonar_report.txt for stats

259 perfSONAR instances
registered in GOCDB/OIM
233 Active perfSONAR instances
172 Running latest version (3.4.2)



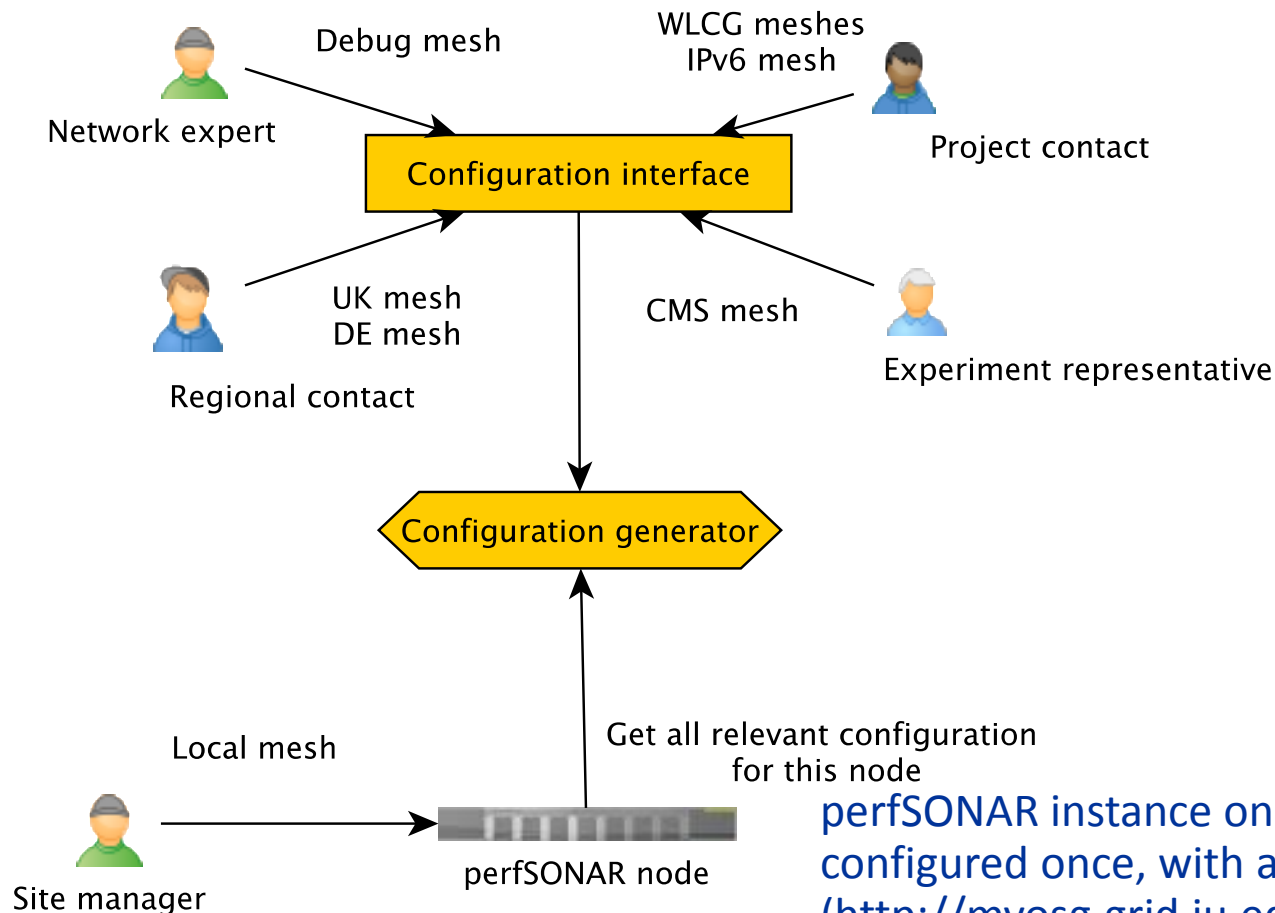
- Initial deployment coordinated by WLCG perfSONAR TF
- Commissioning of the network followed by WLCG Network and Transfer Metrics WG

WG Use Cases / Responsibilities

- The experiments and various infrastructure projects have been asked for use-cases for network monitoring and working group responsibilities
- **Summary of responses:**
 - Defining and understanding slow transfers.
 - perfSONAR can debug, once you know that a link is "slow".
 - Identifying and localizing real network problems
 - Alerting the right people when there are network problems
 - Defining procedures to contact the sites and to "assign" them the issues
 - Producing a "best source" table, on the line of what we have done with the Cost-Matrix, which can also improve our brokering.
 - Commissioning new sites and storage
 - Enabling network-aware tools

perfSONAR Configuration Interface

- perfSONAR instance can participate in more than one mesh
- Configuration interface and auto-URL enables dynamic re-configuration of the entire perfSONAR network



perfSONAR instance only needs to be configured once, with an auto-URL (http://myosg.grid.iu.edu/pfmesh/hostname/<node_id>)

pS Infrastructure Monitoring

- Auto-summaries are available per mesh
- Service summaries per metric type
- **GOAL:** To ensure we continue to reliably obtain ALL network metrics

https://psomd.grid.iu.edu/WLCGperfSONAR/check_mk/

Check **MK** 1.2.4p5 Hostgroup OPN

Tactical Overview

Hosts	Problems	Unhandled
251	30	30

Services

Services	Problems	Unhandled
3084	720	720

Quicksearch

Views

- Dashboards
 - Host & Services Problems
 - Main Overview
 - Network Topology
- Hosts
 - All hosts
 - All hosts (Mini)
 - All hosts (tiled)
 - Favourite hosts
 - Host search

state Host Icons Alias OK Wa Un Cr Pd

UP	ccperfsnar1.in2p3.fr		ccperfsnar1.in2p3.fr	10	2	0	0	0
UP	ccperfsnar2.in2p3.fr		ccperfsnar2.in2p3.fr	10	2	0	0	0
UP	logps01.gridpp.rl.ac.uk		logps01.gridpp.rl.ac.uk	10	2	0	0	0
UP	logps02.gridpp.rl.ac.uk		logps02.gridpp.rl.ac.uk	10	2	0	0	0
UP	lhcbandwidth.twgrid.org		lhcbandwidth.twgrid.org	10	2	0	0	0
UP	lhclatency.twgrid.org		lhclatency.twgrid.org	9	0	0	3	0
UP	lhcomon.bnl.gov		lhcomon.bnl.gov	9	2	0	1	0
UP	lhperfmon.bnl.gov		lhperfmon.bnl.gov	10	2	0	0	0
UP	perfsnar-bw.cern.ch		perfsnar-bw.cern.ch	9	2	0	1	0
UP	perfsnar-de-kit.gridka.de		perfsnar-de-kit.gridka.de	9	2	0	1	0
UP	perfsnar-it.cern.ch		perfsnar-it.cern.ch	10	2	0	0	0
	perfsnar-ow.cnaf.infn.it		perfsnar-ow.cnaf.infn.it	10	2	0	0	0
	perfsnar-ps.cnaf.infn.it		perfsnar-ps.cnaf.infn.it	9	2	0	1	0
	perfsnar-ps.ndgf.org		perfsnar-ps.ndgf.org	10	2	0	0	0
	perfsnar-ps2.ndgf.org		perfsnar-ps2.ndgf.org	10	2	0	0	0
	perfsnar2-de-kit.gridka.de		perfsnar2-de-kit.gridka.de	10	2	0	0	0
	ps-bandwidth.lhcomon.triumf.ca		ps-bandwidth.lhcomon.triumf.ca	10	2	0	0	0
	ps-gsdc01.sdfarm.kr		ps-gsdc01.sdfarm.kr	10	2	0	0	0
	ps-gsdc02.sdfarm.kr		ps-gsdc02.sdfarm.kr	10	2	0	0	0
	ps-latency.lhcomon.triumf.ca		ps-latency.lhcomon.triumf.ca	10	2	0	0	0
	ps.lhcopn-ps.sara.nl		ps.lhcopn-ps.sara.nl	10	2	0	0	0
	ps2.lhcopn-ps.sara.nl		ps2.lhcopn-ps.sara.nl	10	2	0	0	0
	psb01.pic.es		psb01.pic.es	10	2	0	0	0
	psl01.pic.es		psl01.pic.es	10	2	0	0	0
UP	psnar3.fnal.gov		psnar3.fnal.gov	10	2	0	0	0
UP	psnar4.fnal.gov		psnar4.fnal.gov	10	2	0	0	0


t2ps-bandwidth.physics.ox.ac.uk

State	Service	Status detail	Age	Checked	Icons	Perf-O-Meter
OK	perfSONAR 3.4+ Toolkit Version	OK toolkit version found 3.4.1	2015-02-17 07:22:26	27 sec		
OK	perfSONAR Administrator Details	OK - Administrator is Ewan Mac Mahon, email e.macmahon1@physics.ox.ac.uk (cached:0)	2014-12-11 19:57:58	3 hrs		
OK	perfSONAR BWCTL Bandwidth Test Controller	TCP OK - 0.139 second response time on 163.1.5.211 port 4823	2014-12-11 19:58:23	29 min		139.213 ms
WARN	perfSONAR esmond Freshness Bandwidth Direct	WARNING Found stale hosts for certain events, time-range: 3700	2015-02-17 21:47:47	3 hrs		
WARN	perfSONAR esmond Freshness Bandwidth Reverse	WARNING Found stale hosts for certain events, time-range: 3700	2015-02-17 21:48:10	3 hrs		
OK	perfSONAR esmond Measurement Archive	OK esmond reachable	2014-12-11 19:56:42	3 hrs		
OK	perfSONAR Homepage	OK homepage reachable	2015-01-27 19:58:50	3 hrs		
OK	perfSONAR Latitude/Longitude Configured	OK - Latitude is 51.81806, Longitude is -1.30489 (cached:1)	2014-12-11 19:54:37	3 hrs		
OK	perfSONAR Mesh Configuration	OK auto-URL configured	2015-01-26 13:55:06	3 hrs		
OK	perfSONAR NTP Service	OK NTP synchronized	2015-01-29 20:16:35	28 min		
OK	perfSONAR Regular Testing Service	OK Regular Testing enabled and running	2015-01-29 20:17:03	28 min		
OK	perfSONAR Toolkit Version	OK - Version 3.4.1 OK (cached:1)	2014-12-11 19:56:17	3 hrs		

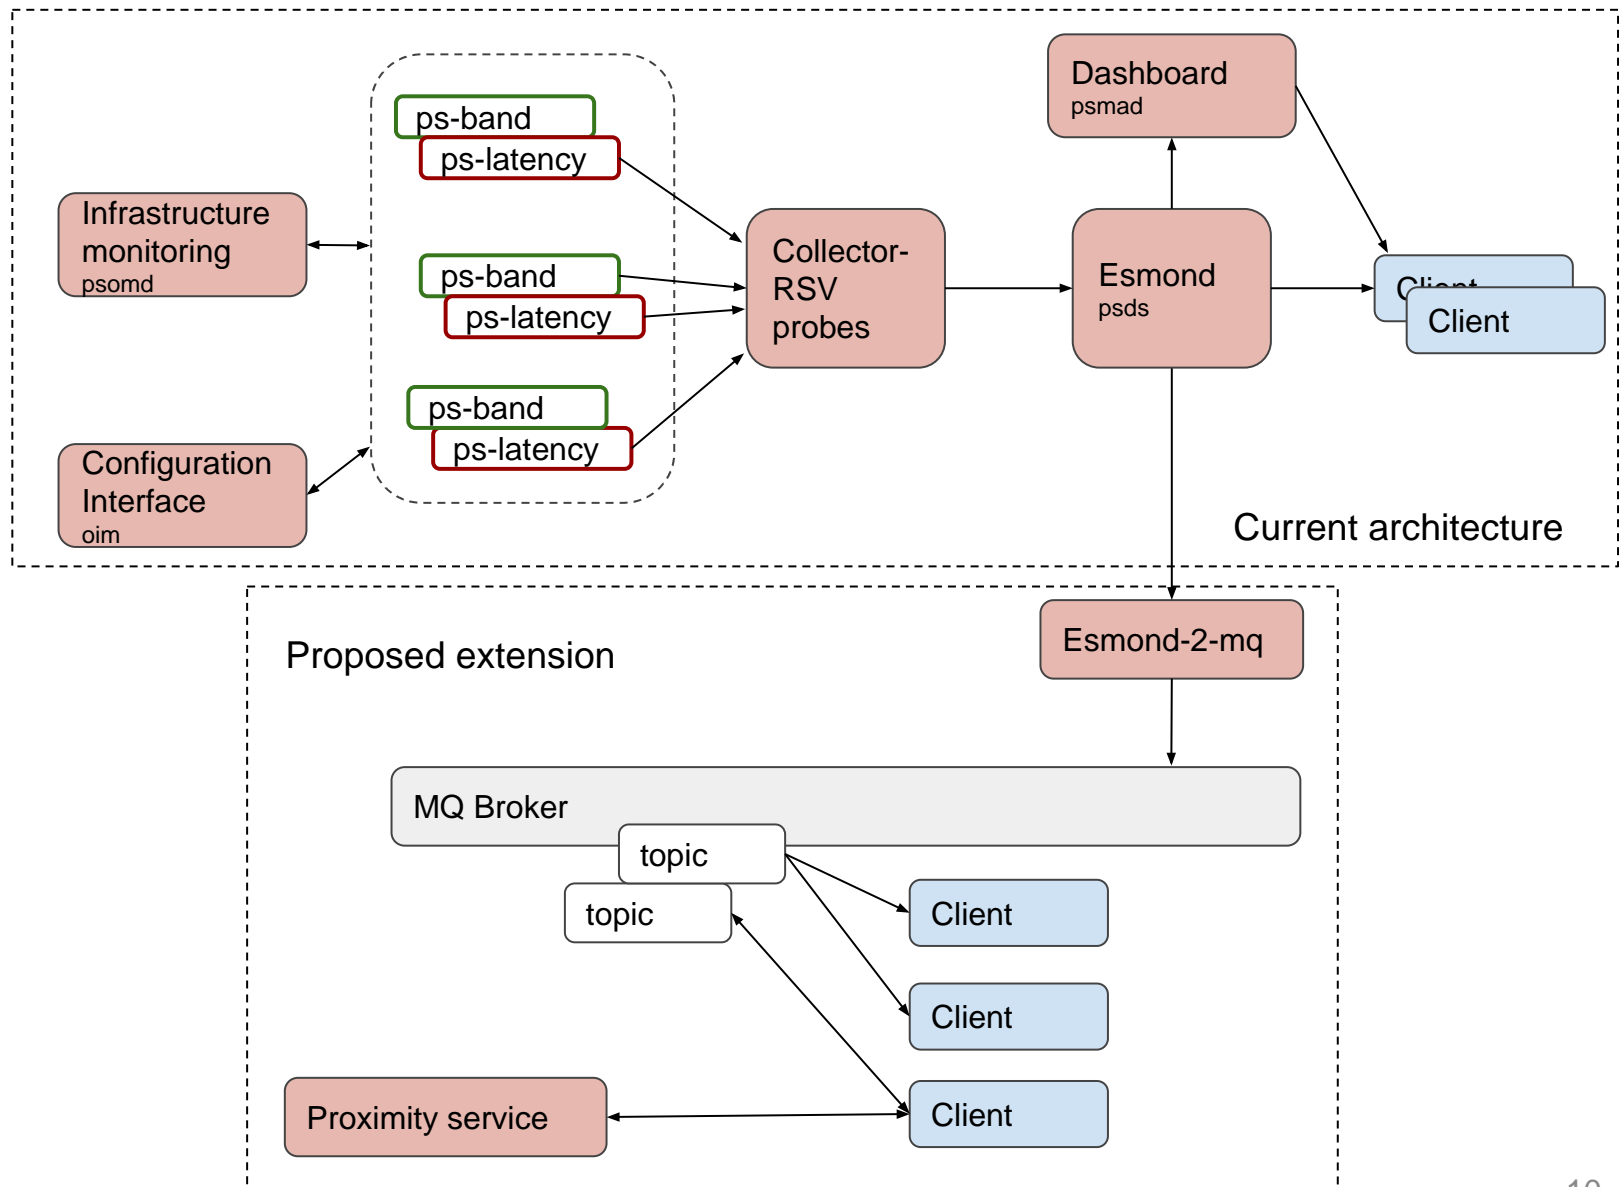
Stale services

- Addons
- Search Graphs
- Other
- Comments

-



Experiments Interface to Datastore



Integration Projects

- Goal
 - Provide platform to integrate network and transfer metrics
 - Enable network-aware tools (see ANSE <http://cern.ch/go/M9Sj>)
 - Network resource allocation along CPU and storage
 - Bandwidth reservation
 - Create custom topology
- Plan
 - Provide latency and trace routes and test how they can be integrated with throughput from transfer systems
 - Provide mapping between sites/storages and sonars
 - Uniform access to the network monitoring
- Pilot projects
 - FTS performance – adding latency and routing to the optimizer
 - Experiment's interface to datastore

Closing remarks

- perfSONAR widely deployed and already showing benefits in troubleshooting network issues
 - Additional deployments within R&E/overlay networks still needed
- Significant progress in **configuration** and **infrastructure monitoring**
 - Helping to reach full potential of the perfSONAR deployment
- **OSG datastore** – community network data store for all perfSONAR metrics – planned to enter production in Q3
- Integration projects underway to aggregate network and transfer metrics
 - FTS Performance
 - Experiment's interface to perfSONAR
- **ANSE project** is providing “hooks” in PANDA and PhEDEx to utilize & future SDN capabilities as they become available at our sites and in our networks.



Questions?

References

- Network Documentation
<https://www.opensciencegrid.org/bin/view/Documentation/NetworkingInOSG>
- Deployment documentation for OSG and WLCG hosted in OSG
<https://twiki.opensciencegrid.org/bin/view/Documentation/DeployperfSONAR>
- New 3.4 MA guide
http://software.es.net/esmond/perfsonar_client_rest.html
- Modular Dashboard and OMD Prototypes
 - <http://maddash.aglt2.org/maddash-webui>
 - https://maddash.aglt2.org/WLCGperfSONAR/check_mk
- OSG Production instances for OMD, MaDDash and Datastore
 - <http://psmad.grid.iu.edu/maddash-webui/>
 - https://psomd.grid.iu.edu/WLCGperfSONAR/check_mk/
 - <http://psds.grid.iu.edu/esmond/perfsonar/archive/?format=json>
- Mesh-config in OSG <https://oim.grid.iu.edu/oim/meshconfig>
- Use-cases document for experiments and middleware
<https://docs.google.com/document/d/1ceiNITUJCwSuOuvbEHZnZp0XkWkwdkPQTQic0VbH1mc/edit>