


# Update on Network and Transfer Metrics WG

Shawn McKee, Marian Babik

GDB

8<sup>th</sup> October 2014



# Network and Transfer Metrics WG

- Mandate
  - Ensure all relevant network and transfer metrics are identified, collected and published
  - Ensure sites and experiments can better understand and fix networking issues
  - Enable use of network-aware tools to improve transfer efficiency and optimize experiment workflows

# Recent meetings

- Last report at WLCG workshop in Barcelona
- 8<sup>th</sup> Sept: Kick-off meeting
  - organization, tasks, review
- 15<sup>th</sup>-16<sup>th</sup> Sept: LHCONe, LHCOPN joint meeting
- 3<sup>rd</sup> Oct: perfSONAR operations
  - deployment and operations changes, shell shock vulnerability

# Team

Shawn McKee (chairperson), Marian Babik (co-chair), ATLAS (Simone Campana), CMS (Tony Wildish), LHCb (Stefan Roiser, Joel Closier), Alice (Latchezar Betev, Costin Grigoras), FAX (Ilija Vukotic), FTS (Mikael Salichos, Oliver Keeble), CMS/PhEDEx (Tony Wildish), Panda (Kaushik De), Rucio (Vincent Garonne)

perfSONAR contacts: US-ATLAS (Shawn McKee), US-CMS (Jorge Alberto Diaz Cruz), UK-ALL (Alessandra Forti, Chris Walker, Duncan Rand), IT-ATLAS (Alessandro de Salvo), IT-CMS (Enrico Mazzoni), CA-ALL (Ian Gable), FR-ALL (Frederique Chollet, Laurent Caillat, Frederic Schaer), TW-ALL (Hsin-Yen Chen, Felix Lee), ND-ALL (Ulf Tigerstedt), DE-ALL (Guenter Duckeck, Andreas Petzold, DE-KIT: Bruno Hoefl, Aurelie Reymund), ES-ALL (Fernando Lopez, Josep Flix), CERN (Stefan Stancu), LHCOPN/LHCONE (John Shade, ESNet: Mike O'Connor), RU-ALL (Victor Kotlyar), ESnet Science Engagement group (Jason Zurawski)

Mailing lists:

[wlcg-ops-coord-wg-metrics@cern.ch](mailto:wlcg-ops-coord-wg-metrics@cern.ch) (ALL)

[wlcg-perfsonar-support@cern.ch](mailto:wlcg-perfsonar-support@cern.ch) (perfSONAR contacts)

Twiki:

<https://twiki.cern.ch/twiki/bin/view/LCG/NetworkTransferMetrics>

# WG Organization

- Areas of work
  - Metrics area – research oriented, review of the current status, correlations of network and transfer metrics
  - perfSONAR area – operations, deployment and configuration
- Meeting schedule
  - Topical meetings, called in advance via doodle
- Reporting and Task tracking
  - Reporting progress at WLCG ops coordination
  - JIRA for task tracking
  - Still looking for volunteers to lead specific tasks

# Metrics Area Tasks

- T1.1: Gather requirements and use cases from experiments, data and workload management systems
- T1.2: Review existing transfer and network metrics and determine how they overlap (or can be extended wrt semantics), identify gaps (missing metrics)
  - Determine current operational status of underlying systems, ensure we can consistently make available all needed metrics
- T1.3: Determine current test coverage and propose how to tune parameters and mesh configurations
- T1.4: Propose how to map perfSONAR topology to transfer systems topologies (and vice versa)

# Metrics

- perfSONAR
  - Network Path - perfSONAR's traceroute to track the network path between sites (once an hour)
  - Bandwidth - perfSONAR's Iperf tool to measure achievable bandwidth.
  - Latency – sending 10Hz of one-way delay measurement packets between all sites. Packet statistics are summarized every minute.
- Transfers
  - Transfer rates – bytes transferred between two endpoints at file-close event (volume/throughput)
    - Other events published (file open, state)
    - Experiments metadata can be included
  - Link status – error-rate for file transfers
  - Errors – detailed error messages, mainly used for debugging purposes

# Transfer Systems

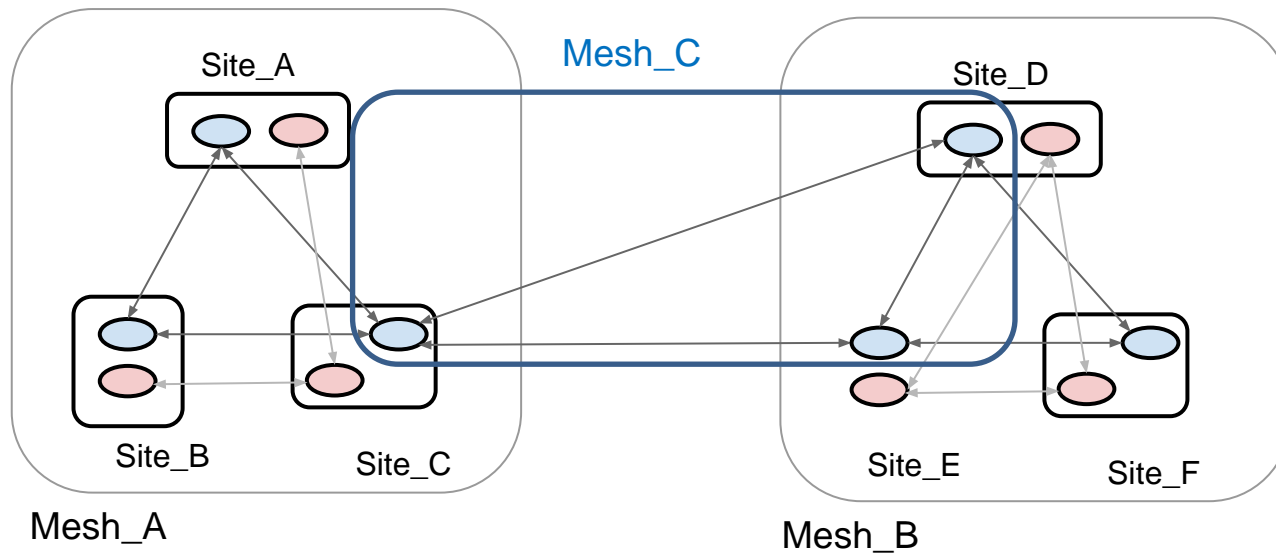
	FTS	FAX/AAA	PhEDEx
VOs	ATLAS, CMS, LHCb	ATLAS, CMS	CMS
Sources	Event messages from FTS service Command line client FTS3 monitoring service FTS dashboard service	XrdMon (GLED) FAX/AAA dashboards HammerCloud/SSB/FSB - synthetic tests	File download agents CMS PhEDEx dashboard (throughput and link)
APIs	FTS dashboard API Event messages in ActiveMQ	FAX/AAA dashboard APIs, SSB/AGIS APIs GLED publishes to ActiveMQ	Oracle DB with API
Status	Production	Production	Production
More details at <a href="https://indico.cern.ch/event/336520/material/slides/0.pdf">https://indico.cern.ch/event/336520/material/slides/0.pdf</a>			




# perfSONAR Area Tasks

- T2.1 perfSONAR Commissioning/Operations
  - Tracking status of the network and developing procedures to follow up incidents/requests
  - Improve docs, create troubleshooting guide with a clear mapping to status tracking
  - Plan deployment campaigns
- T2.2 perfSONAR Storage
  - WLCG/OSG data store testing and deployment
- T2.3 perfSONAR Configuration
  - WLCG configuration interface – testing and deployment
  - Parameter tuning, host allocations tuning

# Deployment Organization



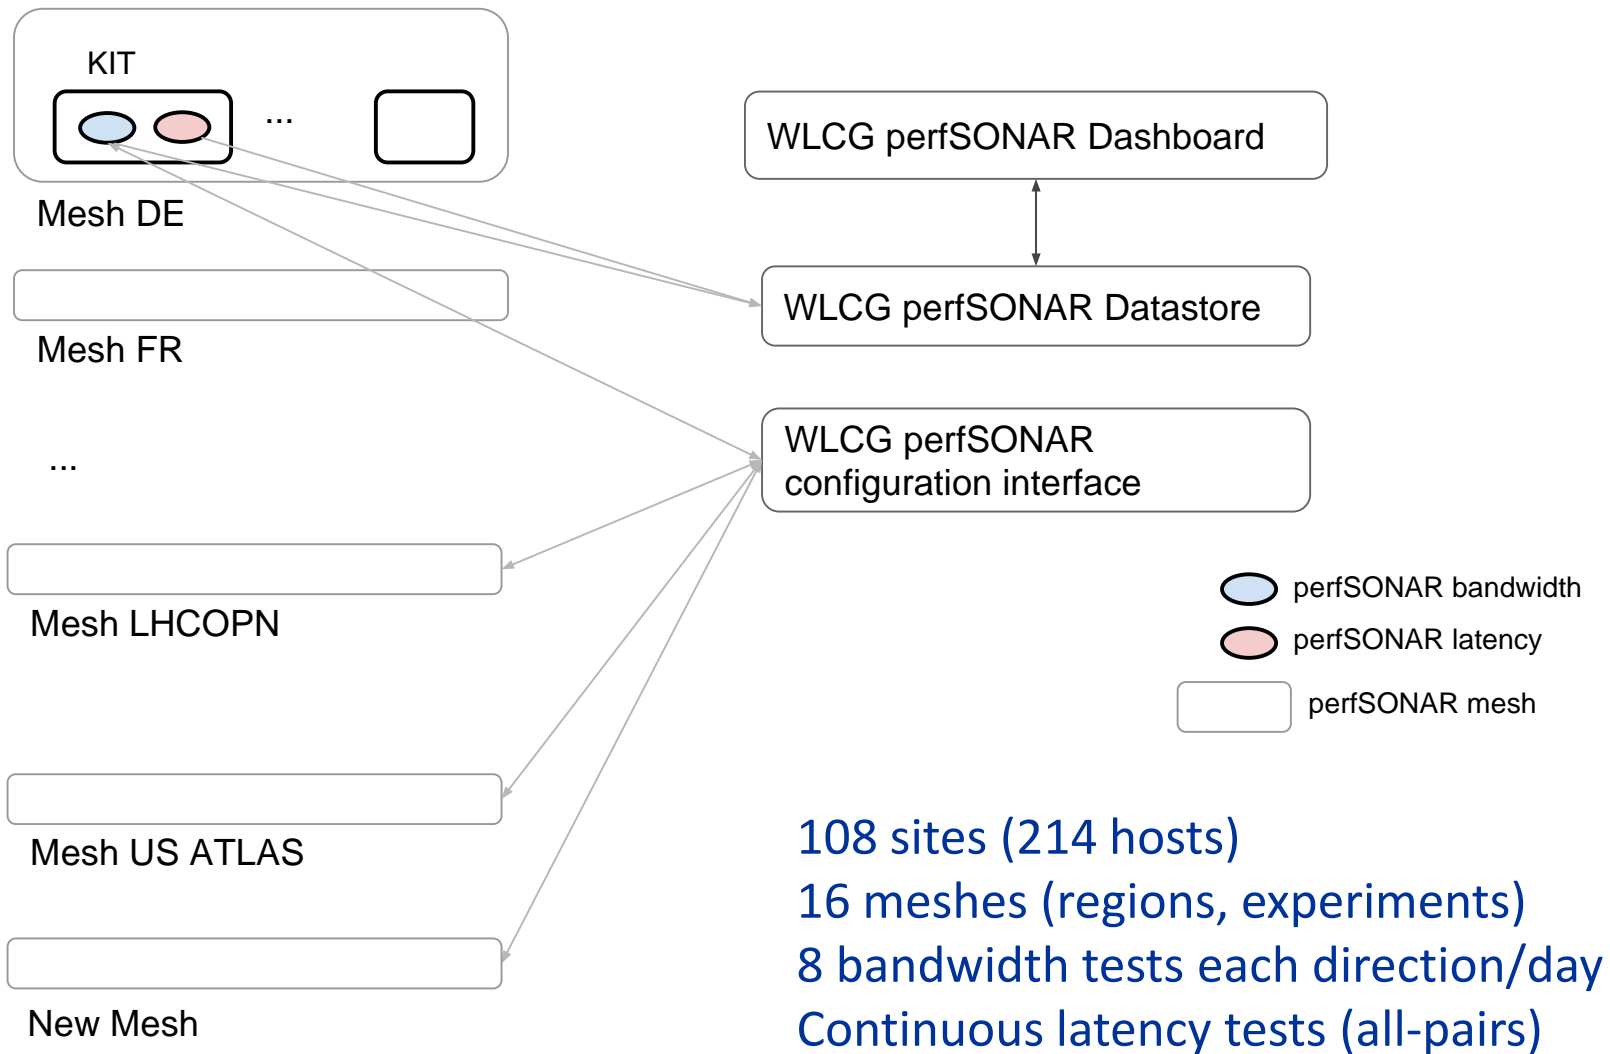
We use groupings of perfSONAR instances (by region, experiment or both) to organize the testing for WLCG. These are referred to as “meshes” and we provide a central mesh-configuration JSON file to control tests, members and test parameters.



We can additionally test “between” meshes by adding subsets of sites to a new mesh (Mesh\_C above)

Our working group can customize and tune the type of meshes and the testing details as needed.

# perfSONAR deployment



# Datastore and configuration

- Network Datastore based on [esmond](#) [1]
  - Tested and deployed by OSG
  - Common API to access all perfSONAR data
  - Storage backend for the dashboard interfaces
- perfSONAR mesh configuration interface
  - Designed and developed from scratch
  - Web based interface supporting auto-configuration of instances
  - Instance details retrieved from GOCDDB and OIM
  - Extensive validation and testing required – volunteers are needed (e-mail Shawn)

# perfSONAR operations

- Introduction of GGUS SU (backed by wlcg-perfsonar-support)
- Complete re-write of the documentation needed
  - Integration (new sites) – requirements, procedure
  - Installation/upgrade guide
  - Troubleshooting/FAQ
  - Guide for experiments (storage APIs, dashboard, etc.)
- Infrastructure monitoring
  - Based on Nagios/Check\_mk, review of current metrics
  - SSB operations dashboard planned
- Better engage infrastructures in our activities
- perfSONAR 3.4 released yesterday
  - Brings critical new improvements and bug fixes
  - Please await instructions to be broadcasted via EGI/WLCG channels before upgrading

# Shell Shock Vulnerability

- Details at [ShellShock-perfSONAR](#) [2]
- Remediation plan
  - Focus on having just one campaign to reinstall, upgrade to 3.4 and move to the new mesh configurations
  - Perform internal security audit and give clear instructions on how to secure re-installed instances
  - Communicate to infrastructures our new perfSONAR security contact list
  - Review the current documentation
  - Send WLCG and EGI broadcast with the re-installation instructions to ALL sites

# Summary

- Network and Transfer Metrics WG – tasks overview and progress ([here](#))
- Next meetings:
  - Metrics area meeting 13-17<sup>th</sup> October
  - perfSONAR operations meeting 20-24<sup>th</sup> October
- Twiki
  - <https://twiki.cern.ch/twiki/bin/view/LCG/NetworkTransferMetrics>
- Contact
  - [wlcg-ops-coord-wg-metrics@cern.ch](mailto:wlcg-ops-coord-wg-metrics@cern.ch)

# References

[1] <http://software.es.net/esmond/>  
<https://twiki.opensciencegrid.org/bin/view/Production/OSGNetworkDatastorePlan>

[2]  
<https://twiki.cern.ch/twiki/bin/view/LCG/ShellShockperfSONAR>