


Network Monitoring and Metrics

Shawn McKee, Marian Babik

WLCG workshop 2014, Barcelona



Outline

- Network monitoring status
 - perfSONAR TF summary
 - Current deployment and challenges
- Network and transfer metrics WG
 - Mandate, Team
 - Objectives
 - Use cases
 - Challenges

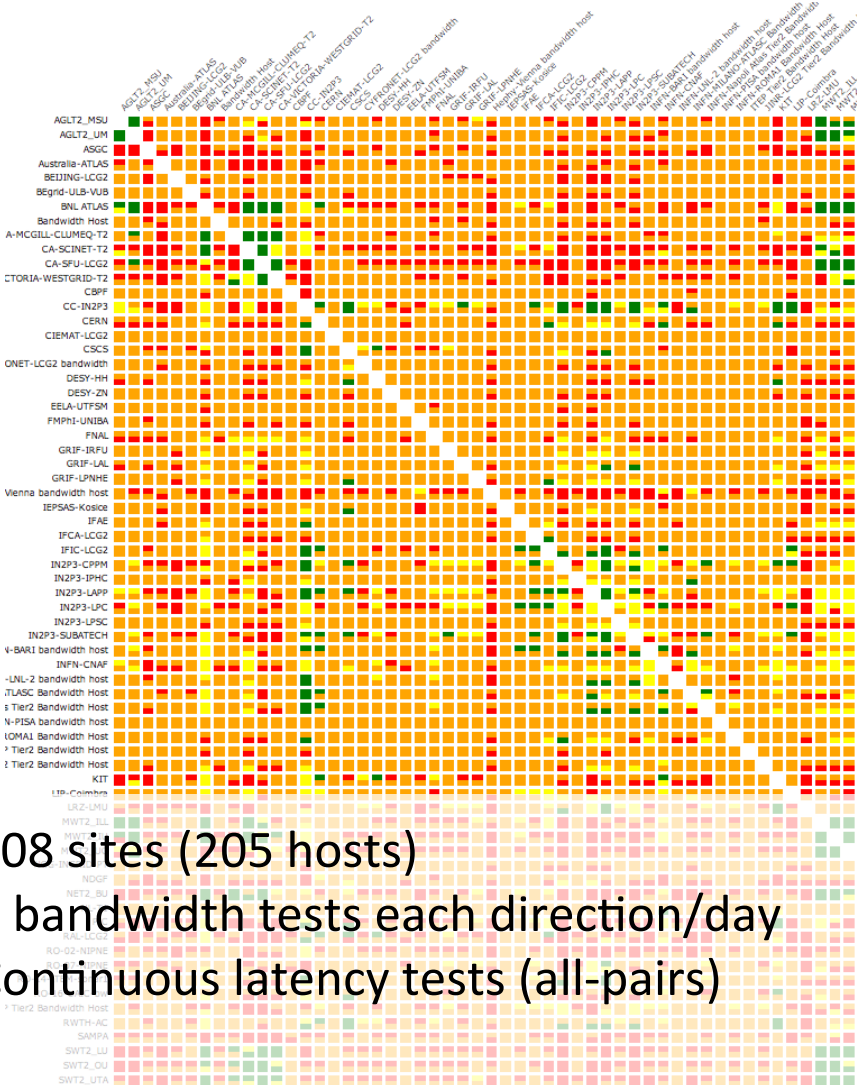
Network Monitoring Status

- Network monitoring based on [perfSONAR](#)
 - Measuring bandwidth and latency
 - End to end network measurements
- Deployment in WLCG followed by perfSONAR TF*
 - Started October 2013, by the end of April 2014 successfully deployed at most of the WLCG sites
 - 205 hosts running (only 8 sites missing)
 - Mesh-based configurations
 - Multiple meshes exists (per region, experiment)
 - Established perfSONAR dashboard tracking status and availability of the instances

*Credits: Simone Campana, Shawn McKee and perfSONAR TF team

Network Monitoring Status

WLCG sites - BWCTL Test Between WLCG Bandwidth Hosts



Servicegroups (Summary)

3 30s Edit View

Name	Alias	O	W	C	U	P
Bandwidth	Bandwidth Test Controller	91	0	7	0	0
NDT	Network Diagnostic Tester	85	0	8	0	0
NPAD	Network Path and Application Diagnosis	85	0	8	0	0
OWAMP	One-Way Ping Service OWAMP	98	0	11	0	0
PS-Admins	PS Toolkit Administrator Configured, cached and checked every hour	160	0	43	0	0
PS-Homepage	PS-Homepage access checked every 6 hours(port 443)	163	2	31	0	0
PS-Homepage-No-Cert	PS-Homepage access checked every 6 hours(port 80)	165	3	33	0	0
PS-LatLong	PS Toolkit Latitude/Longitude Configured, cached and checked every hour	132	0	71	0	0
PS-Version	PS Toolkit Version, cached and checked every hour	136	0	67	0	0
PingER	PingER Measurement Archive	89	0	20	0	0
TracerouteMA	Traceroute Measurement Archive	89	0	20	0	0
WLCG-Mesh-Updates	Check for WLCG mesh updates	1	0	0	0	0
perfSONAR-BUOY-MA	perfSONAR-BUOY Measurement Archive	182	0	22	0	0

108 sites (205 hosts)

8 bandwidth tests each direction/day

Continuous latency tests (all-pairs)

Network Monitoring Challenges

- perfSONAR is a “special” service
 - It tests a multi-domain network path, involving a service at the source and at the destination
 - It requires dedicated hardware and comes in a bundle with the OS.
- Installing a service at every site is one thing, but commissioning a NxN system of links is squared the effort.
 - This is why we have perfSONAR-PS installed but not all links are monitored.
- Tracking versions and updates is very challenging
- Firewalls issues. There are 2 kinds of firewalls rules to be considered:
 - For the hosts to be able to run the tests among themselves
 - For the hosts to be able to expose information to the monitoring tools.
- Test coverage and parameters
 - We need to ensure we’re consistently getting the needed metrics
 - Should we have more VO-specific meshes, i.e. WLCG-CMS, WLCG-ATLAS ?
 - What frequency of testing for traceroute (currently 1/hour)?
- Documentation (as usual) needs to be improved



Network and Transfer Metrics WG

Network and Transfer Metrics

- With the current challenges in mind, we proposed to form a new WG in May:
 - Network and Transfer Metrics WG
- Proposed 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

Motivation

- Optimized experiment workflows
 - Network resource allocation along with CPU and storage in data and job placement
- Integration of network-awareness at different levels of experiment workflows
 - Planning, design phase – network bandwidth reservations, algorithm design and optimization
 - Transfer routing phase – routing based on correlated network and transfer metrics
 - File-transfer phase – consulting network bandwidth for long transfer queues
- Preserve investments already made in network and transfer monitoring

Areas

- **Monitoring**
 - Allows re-active use – react to events or situations in the networks based on network and transfer throughputs
 - Raise alarms and/or continue
 - Abort re-start transfers
 - Choose different source
 - **Topology monitoring**
 - Detect site isolation cases
 - Influence source selection
- **Network control**
 - **Allow pro-active use**
 - Network bandwidth allocations
 - Resource allocations based on CPU, storage and network
 - Optimize network topologies based on operational conditions

Use case: Find/Localize Problems

- Identifying and localizing a network problem
 - Often this is very difficult and time-consuming for Wide-Area Network (WAN) problems
- Scheduled perfSONAR bandwidth and latency metrics monitor WLCG network paths
 - Significant **packet-loss** or **consistent large deviation from baseline bandwidth** indicate a potential network problem (see in GUI or via alarms).
 - On-demand tests to perfSONAR instances can verify the problem exists. Different test points along the path can help pin-point the location.
 - Correlation with other paths sharing common segments can be used to localize the issue.
 - **The time** things change is also very useful to find the root causes. Scheduled tests provide this.

Use case 2: Optimize Net Use

- Both ATLAS and CMS are working together on an NSF project: **Advanced Network Services for Experiments (ANSE)**, adding network awareness into **PANDA** and **PhEDEx**
 - In both cases they can benefit from network info
- Use network information for:
 - AAA/FAX brokerage
 - Job/task assignment
 - Improve flow of activated jobs
 - Improved accounting and diagnosis for transferring jobs
 - Site selection (data sources vs available CPU)
- The plan is to utilize network metrics to make intelligent decisions regarding our use of the network

WG Objectives

- Identify and continuously make available relevant transfer and network metrics
 - Ensure we can consistently publish all the relevant metrics
 - Common metric attributes semantics needed for analytics/correlations
 - Improve our understanding on what metrics are needed and how we can get them
- Document metrics and their use
- Facilitate their integration in the middleware and/or experiment tool chain
 - Work with experiments on their use cases
- Coordinate commissioning and maintenance of WLCG network monitoring
 - Ensure all links continue to be monitored and sites stay correctly configured
 - Verify coverage and optimize test parameters

WG Team

- Chairs: Shawn McKee, Marian Babik
- Members: invited previous members of the perfSONAR-PS TF responsible for different clouds
 - Current status at <https://twiki.cern.ch/twiki/bin/view/LCG/MeshLeaders>
 - Very good feedback so far
 - Added contacts for missing clouds (NDGF, CMS)
- Invited members knowledgeable about FAX/AAA, FTS, PhEDEx, Panda and Rucio

Summary

- Successful deployment of perfSONAR network
- New WG formed to work on both network and transfer metrics
 - Potential use in different areas/tools of experiment workflows
 - Work foreseen in identifying and making available metrics based on experiment use cases
 - Commissioning of the perfSONAR network vital to achieve its potential
- Kick-off meeting planned
 - Agree on the organization of work
 - Discuss objectives and define tasks to be done
- Ideas and suggestions welcome
- Contact: wlcg-ops-coord-wg-metrics@cern.ch