

WLCG Network Throughput WG

Shawn McKee, Marian Babik
pre-GDB - Networking

Introduction

- [WLCG Network Throughput WG](#) formed in the fall of 2014 within the scope of WLCG operations with the following objectives:
 - Ensure sites and experiments can better understand and fix networking issues
 - Measure end-to-end network performance and use the measurements to single out on complex data transfer issues
 - Improve overall transfer efficiency and help us determine the current status of our networks
- Core activities:
 - Deployment and operations of perfSONAR infrastructure
 - Gain visibility into how our networks operate and perform
 - Network performance incidents response team
 - Provide support to help debug and fix network performance issues
 - Network analytics
 - Improve our ability to fully utilize the existing network capacity

perfSONAR

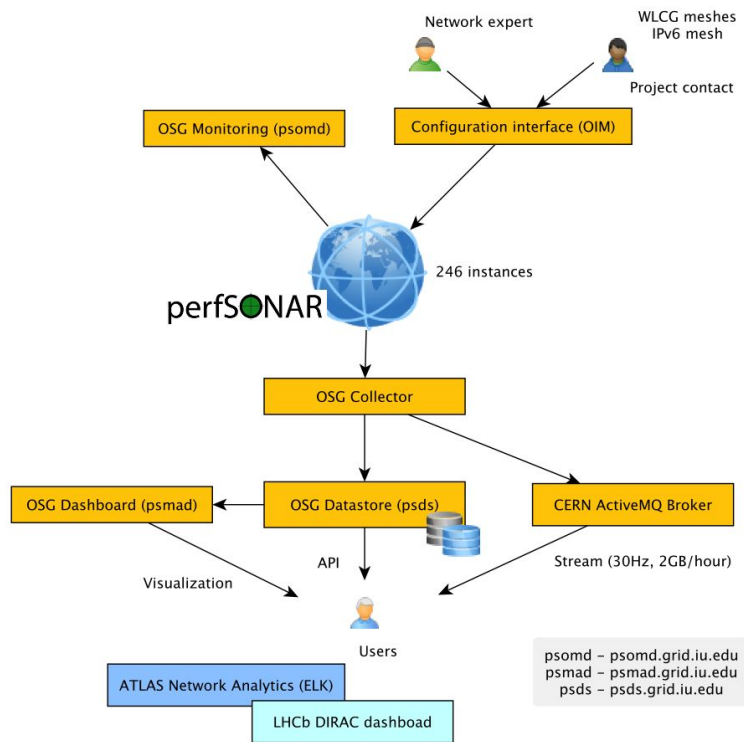
- Consortium of organizations building an interoperable network monitoring middleware
 - ESnet, Internet2, GEANT, RNP, etc.
- Measurement middleware to test network performance
 - Supporting many different tools: iperf3, nuttcp, traceroute, tracepath, ow-ping, etc.
 - Operates test scheduler - ensuring one test at a time
 - Open platform and infrastructure
 - Supports direct, reverse as well as 3rd party tests (from CERN testing BNL to PIC)
- Different versions available
 - Tools - just tools, as mentioned above
 - Testpoint - tools + scheduler
 - Toolkit - testpoint + local data store (esmond) + web interface (incl. admin)
 - Central services - central (mesh) configuration, dashboards

WLCG perfSONAR Infrastructure

- perfSONAR has been widely deployed in WLCG
 - 249 active instances, deployed at 120 sites including major network hubs at ESNet, GEANT
 - Measuring many different network metrics on all existing LHCOPN/LHCONE links
 - Deployment and support coordinated by the [WLCG Network Throughput WG](#)
 - In addition there are more than 1600 perfSONAR deployed in public
 - Also used by HEPiX IPv6 WG
- Close to 10k network links measured continuously
 - Metrics include: traceroute/tracepath, throughput (iperf3), one way latency, ping packet loss, packet reordering, packet duplicates, jitter, TTL, etc.



Network Measurement Platform



OSG has developed an extensive **network measurement platform** using perfSONAR

- Tests can be centrally configured and are continuously gathered by the OSG collectors
- Long-term network measurements storage provided by OSG Datastore with public API
- All measurements are also available for subscriptions via ActiveMQ netmon brokers at CERN

Network incidents response

- WG established support unit to coordinate responses to potential network performance issues:
 - Tickets opened by sites or experiments triaged to the right destination
 - Many issues are potentially resolvable within the working group
 - Real network issues can be identified and directed to the appropriate network support centers
 - In collaboration with GEANT eduPERT
- Several major cases reported in the last year
 - Most of them due to issues at sites (configuration, insufficient or malfunctioning hardware, etc.), but some cases were also with NRENs
- Procedure and past cases reported at https://twiki.cern.ch/twiki/bin/view/LCG/NetworkTransferMetrics#Network_Throughput_Support_Unit

Network Analytics Activities

- Ilija Vukotic (Univ. of Chicago) has developed ELK/jupyter stack for ATLAS Analytics and worked with Xinran Wang on [anomaly detection and advanced alerting/notifications](#) for network problems
- Jerrod Dixon and Brian Bockelman (UNL) exploring network analytics in CMS
- Shawn McKee (Univ. of Michigan) working on real-time root cause analysis ([PuNDIT](#)) in collaboration with ESNet
- Henryk Giemza (NCBJ), Federico Stagni integrating perfSONAR in DIRAC for LHCb
- Hendrik and Marian working on developing models for [network cost-matrix](#) - determine performance of network paths

Network Anomaly Detection/Alerting

perfSONAR measurements can reveal network performance problems soon after their occurrence, but currently these events and indicators are just buried in a huge amount of data.


Univ. of Chicago is working on a system which continuously monitors all the data, and automatically sends notifications when a problem of interest just occurred.

- Prototype already available for testing
- The number of alarms based on simple thresholds still significant, work ongoing to evaluate machine learning approaches

Packet loss greater than 2% for a period of 3 hours on almost 5% of all LHCONE links

Network Cost Matrix

- We want to have the ability to understand which network paths perform better when choices are available.
 - Can we determine network path performance in real-time from perfSONAR measurements ?



We compared LHCOPN router traffic with latency and packet loss measurements (10Hz) from perfSONAR - looks promising.

Trained NN to estimate throughput based on one-way latency and packet loss - published to ActiveMQ for all LHCONE links

What's next ?

- perfSONAR infrastructure and network measurement platform
 - Campaign to upgrade all sites to perfSONAR 4.0 (and 4.1 which drops SL6 support)
 - Introduction of new deployment options (small-nodes, VMs)
 - New configuration system to be deployed which comes with perfSONAR 4.0
 - Upgrading the current Datastore to 24TB (4+ years in total)
 - Continue using message bus to stream data to potential users
- Network incident response
 - Better integration with the existing network analytics to help reduce time to resolution
- Network analytics
 - Start publishing router traffic for LHCOPN and ESNet
 - Bring alerting/notification platforms to production
 - Integrate analytics with FTS and other transfer systems to gain better insights into how we're using the network

Summary/Questions

- Deployment of perfSONARs at WLCG sites made it possible for us to see and debug end-to-end network problems
 - Shall we make perfSONARs mandatory for any site experiencing network problems ?
- A group focusing on helping sites and experiments with network issues using perfSONAR was formed - [WLCG Network Throughput](#)
 - It is increasingly important to focus on site-based network operations
 - Please CC the support unit if you think there is a network performance issue
- Recent network analytics of LHCOPN/LHCONE perfSONAR data also point out some very interesting facts:
 - Packet loss greater than 2% for a period of 3 hours on almost 5% of all LHCONE links
 - Please contact wlcg-network-throughput-wg@cern.ch to test
- Network telemetry (real-time network link usage) to become available for selected links (LHCOPN, ESNet)