Improving WLCG Networks Through Monitoring and Analytics

Edoardo Martelli, Marian Babik, Shawn McKee
on behalf of WLCG Network Throughput WG

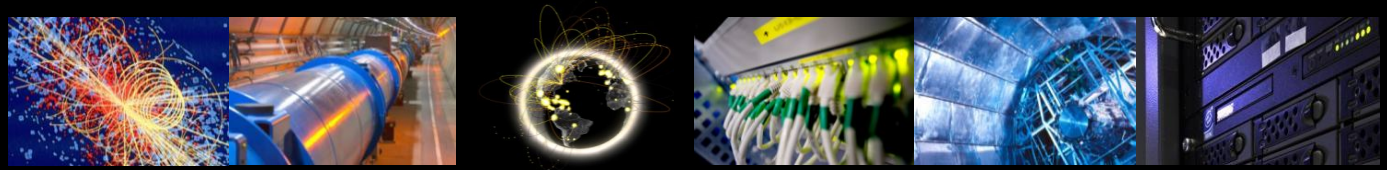

WLCG
Worldwide LHC Computing Grid

Open Science Grid

Outline

- Network Performance
  - Importance of measuring our networks
  - OSG/WLCG activities and WLCG Network Throughput WG
- Platform Overview
  - WLCG perfSONAR Infrastructure
  - OSG Network Monitoring Platform
- Platform Use
  - Activities and collaborations
  - Network Throughput Support Unit
  - Operations
- Evolution
- Summary

Importance of Measuring Our Networks

- End-to-end network issues are difficult to spot and localize
  - Network problems are multi-domain, complicating the process
  - Performance issues involving the network are complicated by the number of components involved end-to-end
  - Standardizing on specific tools and methods focuses resources more effectively and provides better self-support.
- **Network problems can severely impact experiments workflows and have taken weeks, months and even years to get addressed!**
- perfSONAR provides a number of standard metrics we can use
  - Latency, Bandwidth and Traceroute
  - These measurements are critical for network visibility
- **Without measuring our complex, global networks we wouldn't be able to reliably use those network to do science**

OSG/WLCG Networking Activities

- OSG is in its 6th year of supporting WLCG/OSG networking focused on:
  - Assisting its users and affiliates in identifying and fixing network bottlenecks
  - **Developing and operating a comprehensive Network Monitoring Platform**
  - Improving our ability to manage and use network topology and network metrics for analytics

- WLCG Network Throughput Working Group was established to ensure sites and experiments can better understand and fix networking issues:
  - Oversees the **WLCG perfSONAR infrastructure**
    - Core infrastructure for taking network measurements and performing low-level debugging activities
  - **Coordinates WLCG network performance incidents** - runs a dedicated support unit which involves sites, network experts, R&Es and perfSONAR developers
    - Many issues are potentially resolvable within the working group

perfSONAR deployment

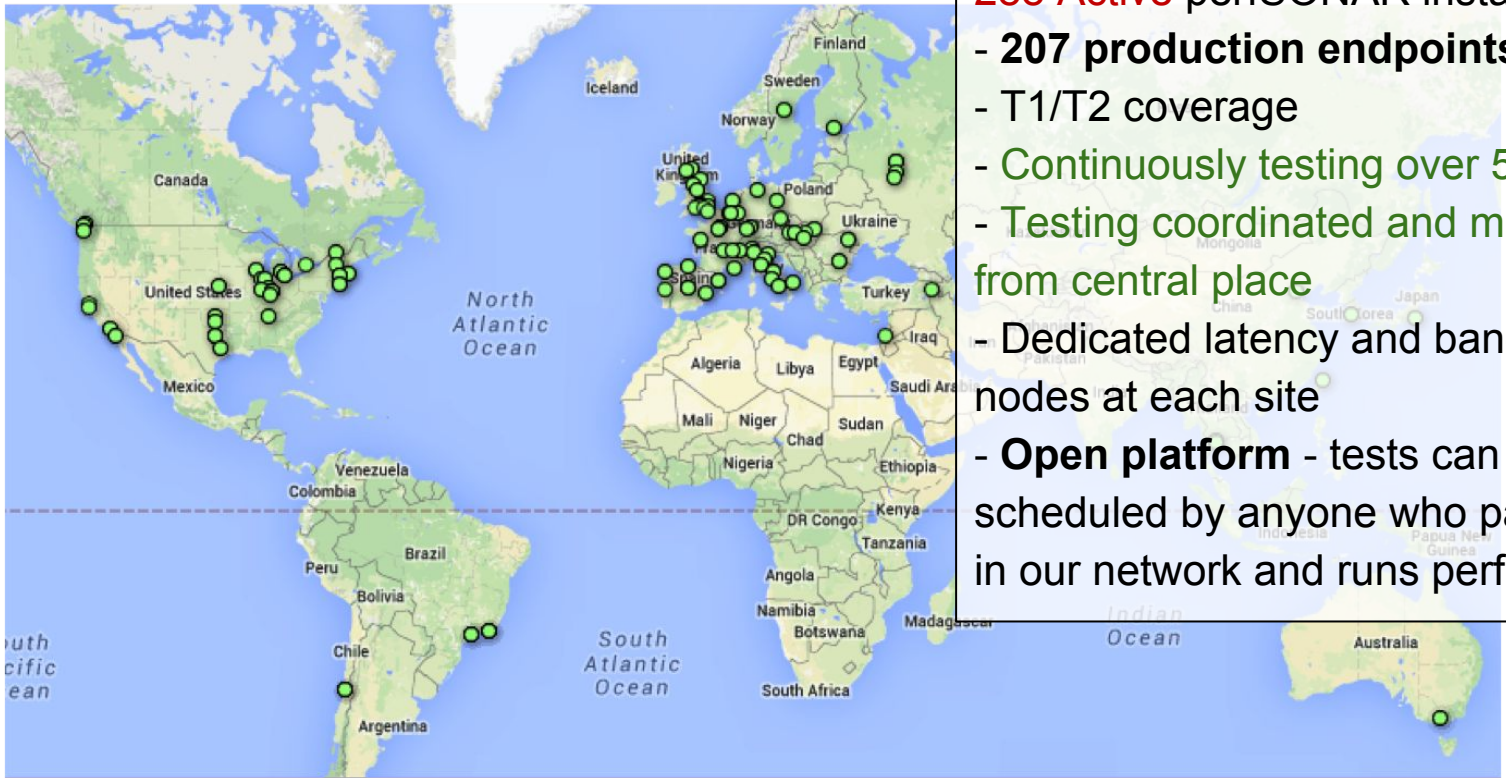

288 Active perfSONAR instances
- **207 production endpoints**
- T1/T2 coverage
- Continuously testing over 5000 links
- Testing coordinated and managed from central place
- Dedicated latency and bandwidth nodes at each site
- **Open platform** - tests can be scheduled by anyone who participates in our network and runs perfSONAR

Platform Overview

- Collects, stores, configures and transports all network metrics
  - Distributed deployment - operated in collaboration
- All perfSONAR metrics are available via API, live stream or directly on the analytical platforms
  - Complementary network metrics such as ESNet, LHCOPN traffic also via same channels

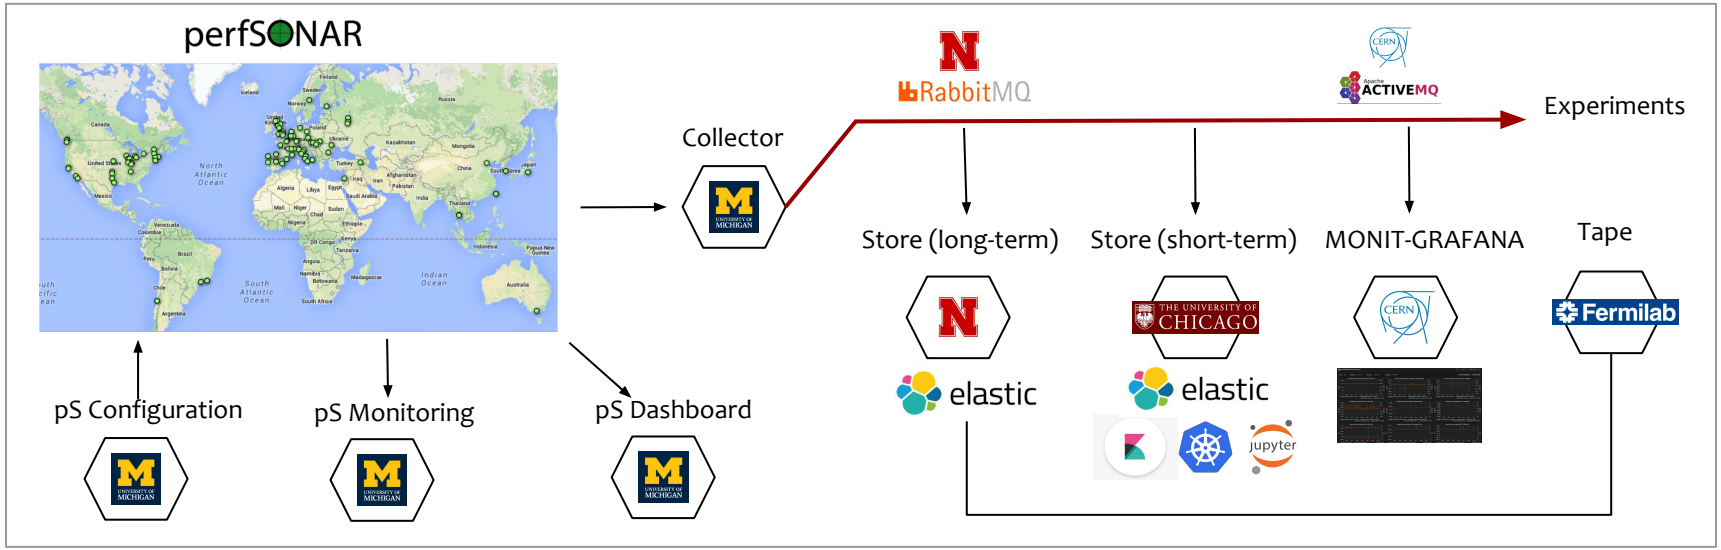

Platform Use

- **WLCG and OSG operations**
  - Baseline testing and interactive debugging for incidents reported via support unit
  - Regular reports at the WLCG operations coordination and WLCG weekly operations
  - Providing **Grafana dashboards** that help visualise the metrics
- Enabling analytical studies - data stored in the ATLAS Analytics platform
  - Providing an important source for network metrics (bandwidth, latency, path)
  - More on this in *Integrating Networking Into ATLAS* (this session)
- **Cloud testing - HNSciCloud** - testing commercial cloud providers
  - Baselining and evaluating network performance
- HEPiX IPv6 WG
  - Now testing bandwidth and paths over IPv6
- Collaboration with other science domains deploying perfSONAR
  - E.g., US Universities, Pittsburgh Supercomputer Center, European Bioinformatics Institute
  - Also close collaboration with (N)RENs who provide LHCONE perfSONAR coverage

WLCG Network Throughput Support Unit

Support channel where sites and experiments can report potential network performance incidents:

- Relevant sites, (N)RENs are notified and perfSONAR infrastructure is used to narrow down the problem to particular link(s) and segment. Also tracking past incidents.
- Feedback to WLCG operations and LHCOPN/LHCONE community

Most common issues: MTU, MTU+Load Balancing, routing (mainly remote sites), site equipment/design, firewall, workloads causing high network usage

As there is no consensus on the MTU to be recommended on the segments connecting servers and clients, LHCOPN/LHCONE working group was established to investigate and produce a recommendation.

Operations Dashboards

- One of many dashboards provided showing LHCOPN/ESNet, perfSONAR, FTS, etc.
- Above is a sample dashboard showing **side-by-side comparison of perfSONAR data, LHCOPN traffic and FTS transfers**
- Also available with ESNet traffic data

Evolution

- Platform evolution will be mainly driven by the next release of perfSONAR
  - Version 4.1 will become available soon
- New Endpoint control capabilities
  - Many more options on what can be centrally managed on a per-instance basis
- Improved pScheduler - better stability and performance
- Resource management – port pools
  - **Preemptive scheduling support** – improving client response time
- New plugins
  - **Network traffic capture (via SNMP)**
  - Application-level (e.g. http response time)
- TWAMP support (two-way active measurement)
  - ping alternative of owamp – routers/switches can participate in the tests
- **Docker support** for testpoint

Summary

- OSG in collaboration with WLCG are operating a comprehensive network monitoring platform
- Platform has been used in a wide range of activities from core OSG/WLCG operations to Cloud testing and IPv6 deployment
- Providing feedback to LHCOPN/LHCONE, HEPiX, WLCG and OSG communities
- Next version of perfSONAR will enable additional functionality as well as improve overall stability and performance
- Further analytical studies are planned to better understand our use of networks and how it could be improved
    - Potential use for experimental network activities, e.g. TCP BBR evaluation, understanding behaviour of mixed UDP/TCP flows, etc.

References

- OSG/WLCG Networking Documentation
  - https://opensciencegrid.github.io/networking/
- perfSONAR Stream Structure
  - http://software.es.net/esmond/perfsonar_client_rest.html
- perfSONAR Dashboard and Monitoring
  - http://maddash.opensciencegrid.org/maddash-webui
  - https://psetf.opensciencegrid.org/etf/check_mk
- perfSONAR Central Configuration
  - https://psconfig.opensciencegrid.org/
- Grafana dashboards
  - http://monit-grafana-open.cern.ch/
- ATLAS Analytics Platform
  - https://indico.cern.ch/event/587955/contributions/2937506/
  - https://indico.cern.ch/event/587955/contributions/2937891/

Networking Challenges

There are number of challenges in the networking, which will require improved collaboration with other sciences as well as HEP-focused R&D:

- Capacity/share for data intensive sciences
  - No issues wrt available technology, however
  - What if N more HEP-scale science domains start competing for the same resources ?
- Remote data access proliferating in the current DDM design
  - Promoted as a way to solve challenges within experiment's DDM
  - Different patterns of network usage emerging
    - Moving from large streams to a mix of large and small frequent event streams
- Integration of Commercial Clouds
  - Impact on funding, usage policies, security, etc.
- Technology evolution
  - Software Defined Networking (SDN)/Network Functions Virtualisation (NFV)

Network Evolution Areas

The following are some of the key areas for HEP Networking R&D:

- Improving efficiency of data transfers
  - TCP BBR - version 2 is in the works with promising improvements
  - Exploring alternative protocols for transfers (UDP)
- Caching
  - Data caches co-located with network hubs in a similar way as on commercial CDNs
- Federations/Clouds
  - Overlay networks spanning multiple domains
  - Multi-clouds - expanding DC networking via L3VPNs
- Technology
  - SDN/NFV approaches - currently looked at by HEPiX NFV WG
    - Compute - Agile service delivery on Cloud Infrastructures (OpenStack, Kubernetes)
    - Data Transfers - Network resource optimisation - dynamically optimising the network based on its load and state (more in Shawn/Ilija)