



# WLCG perfSONAR Update

Marian Babik, CERN IT  
on behalf of WLCG Network Throughput WG



# Outline

- WLCG perfSONAR infrastructure status
- 100Gbps Testing
- OSG/WLCG Network Monitoring Platform
- New Analytics and Tools
- Summary

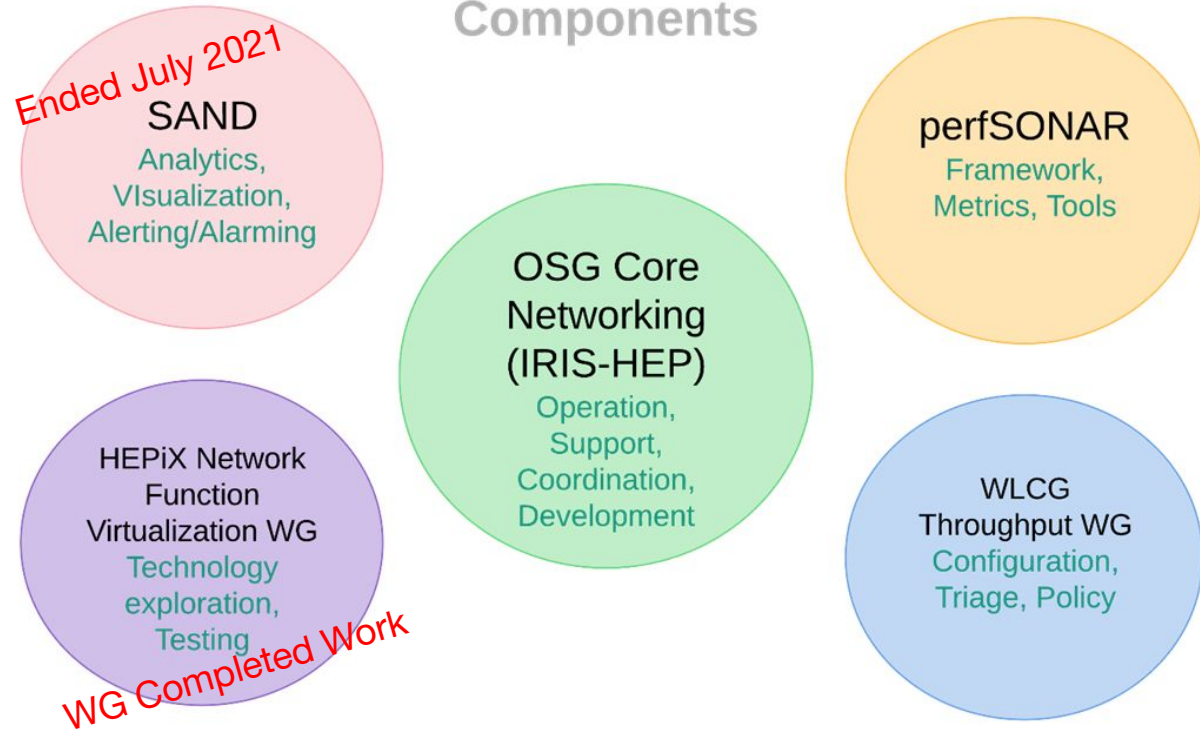


# OSG/WLCG networking projects

There have been 4 coupled projects around the core OSG Net Area

1. **SAND** (NSF) project for analytics (ended)
2. **HEPiX NFV WG** (finished work)
3. **perfSONAR** project
4. **WLCG Network Throughput WG**

## OSG Networking Components



# perfSONAR deployment




**238 Active** perfSONAR instances

- **207 production endpoints**
- T1/T2 coverage
- Dedicated latency and bandwidth nodes at each site
- Testing coordinated and managed from central place
- Continuously testing over 5000 links
- LHC experiments, DUNE, BelleII, LSST
- LHCOPN/LHCONE, ARCHIVER, StashCache, SLATE, CC\*


# perfSONAR deployment

238 Active perfSONAR instances - 207 production endpoints - T1/T2 coverage


perfSONAR Toolkit NIC Speed Histogram




perfSONAR Toolkit Host Memory Histogram




perfSONAR NIC Speed (Bps)




perfSONAR distributions



pS NIC MTU



pS NIC Speed vs Time



Our global toolkit deployment has a range of systems in terms of age and capability

## Dashboard in ELK

Sites should remember to not only upgrade perfSONAR software but also the underlying hardware, as nodes become too old or are unable to test at the site storage speed.



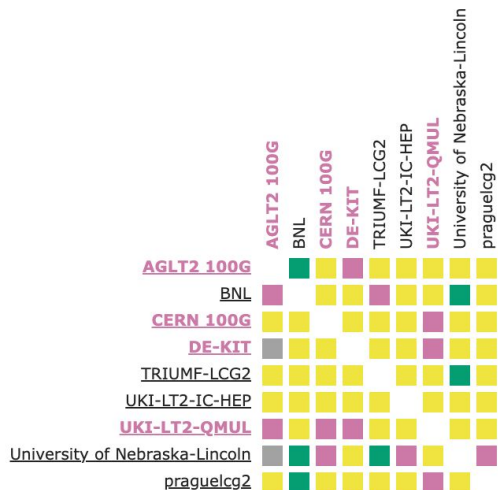
# 100Gbps Testing 24th May 2022

- WLCG 100Gbps mesh  
Thresholds updated


WLCG 100G Mesh - WLCG 100G IPv6 Bandwidth - Throughput



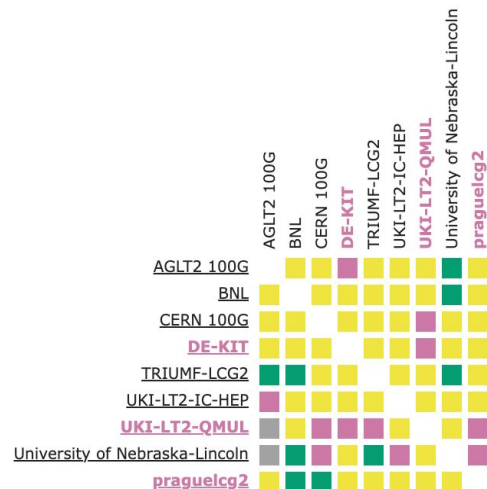
! Found a total of 5 problems involving 4 hosts in the grid



WLCG 100G Mesh - WLCG 100G IPv4 Bandwidth - Throughput



! Found a total of 4 problems involving 3 hosts in the grid



# 100Gbps Testing









- Monthly meetings since January
  - Aim to achieve 10% of avail. capacity (~10Gbps) on a regular basis
  - Discussing ways to tune the nodes and improve stability
  - wlcg-perfsonar-100g mailing list ([join](#))
- Tunings
  - Used CheckMK monitoring along with ES/Kibana dashboards to check status
  - TCP buffers and MTU appear to have made the biggest difference
    - TCP buffers by default at ~ 200MB, need to be increased to 1GB
  - References:
    - <https://fasterdata.es.net/host-tuning/linux/100g-tuning/>
  - Tried FQ but that actually decreased the throughput in tests (not work-conserving)
  - NIC interrupts/core sync only possible via manual tests
- maddash shows by default avg. over 24 hours - extended to 4 days
- New host-based Grafana [dashboard](#) available

# Grafana dashboard

General / perfSONAR H2H Performance


Last 7 days

Src Bandwidth pse01-gva-100g.cern.ch ▾ Src Latency psl01-gva.cern.ch ▾ Dst1 Bandwidth lhcmn3.bnl.gov ▾ Dst1 Latency lhperfmon.bnl.gov ▾ Dst2 Bandwidth perfsonar-test-v1683-portal.gridka.de ▾ Dst2 Latency perfsonar-2-de-kit.gridka.de ▾  
Dst3 Bandwidth lt2ps00-bw.grid.hep.ph.ic.ac.uk ▾ Dst3 Latency lt2ps00-lat.grid.hep.ph.ic.ac.uk ▾ [Network Throughput WG](#) [OSG Networking](#)













# Grafana dashboard



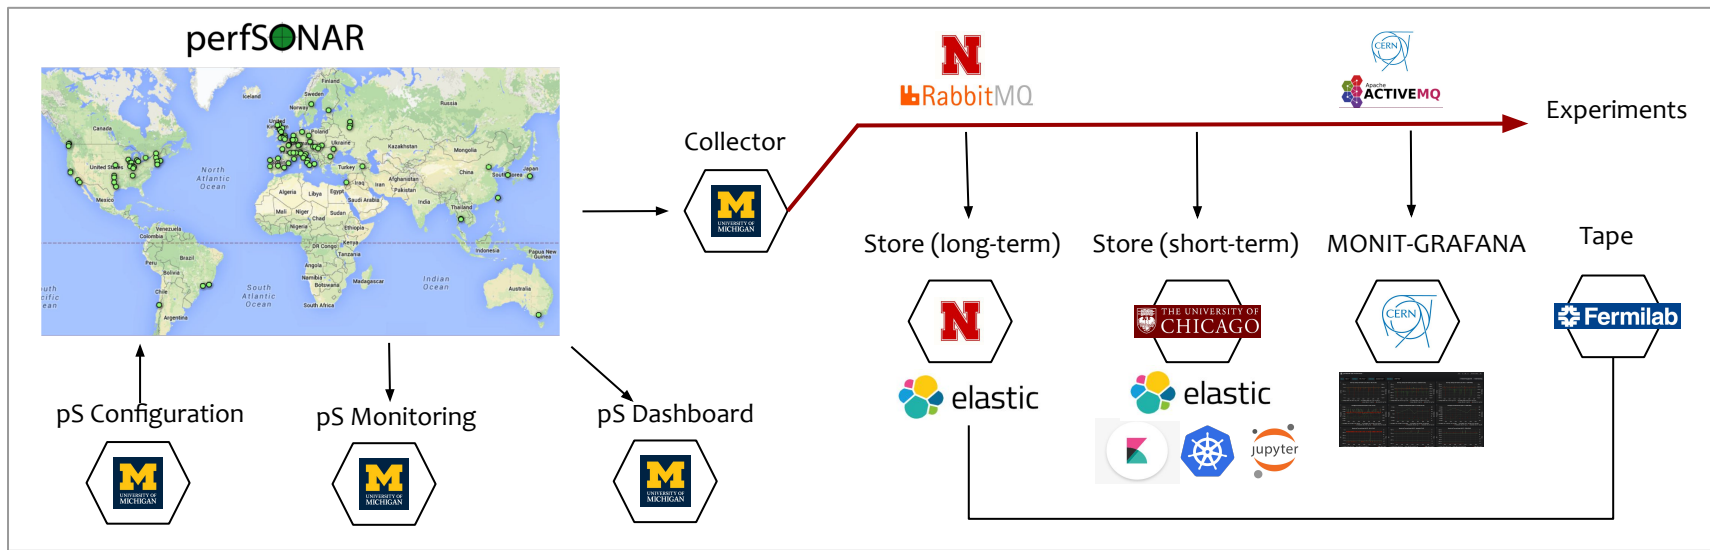
General / perfSONAR H2H Performance

Src Bandwidth: pse01-gva-100g.cern.ch | Src Latency: psl01-gva.cern.ch | Dst1 Bandwidth: lhcmn3.bnl.gov | Dst1 Latency: lhcpfmon.bnl.gov | Dst2 Bandwidth: perfsonar-test-v14 | Dst3 Bandwidth: lt2ps00-bw.grid.hep.ph.ic.ac.uk | Dst3 Latency: lt2ps00-lat.grid.hep.ph.ic.ac.uk



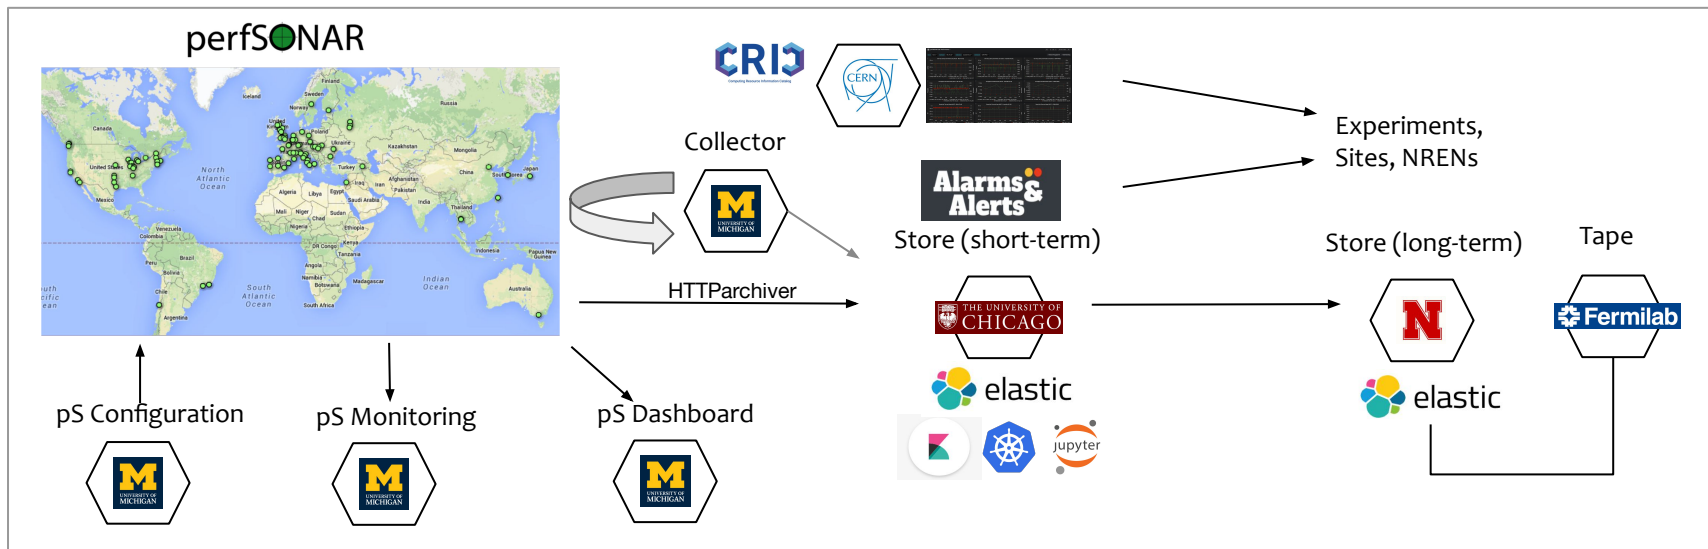
# Reminder: Network Measurement Platform Overview

- 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 ESNat, LHCOPN traffic also via same channels




# Network Measurement Platform Evolution

- **Collects, stores, configures and transports all network metrics**
  - Distributed deployment - operated in collaboration
- Planned evolution based on the perfSONAR 5
  - Directly publishing results from perfSONARs to ES@UC
  - High-level services provided to the experiments/users



# Tools and Applications for Network Data

- To organize access to all the various resources we have NEW homepage (<https://toolkitinfo-nextjs.vercel.app/>)
- We already have Kibana dashboards looking at
  - [Bandwidth](#)
  - [Traceroute](#)
  - [Packetloss](#) / [Latency](#)
  - [Infrastructure](#)
- With the completion of the SAND project, we have a few prototype tools that help us analyze and utilize our net data
  - We have a new perfSONAR focused dashboard: **ps-dash**
  - We have added a self-subscribe tool for network alarms call **AAAS**
  - *Next two pages have the details on these two apps*



# pS (perfSONAR) Dash

perfSONAR Toolkit Information

Kibana: Packet Loss in OSG/WLCC

Kibana: Packet Loss Tracking

MEPHI Tracer: Traceroute explorer

ps dash

SITES

LINKS

PLOTS

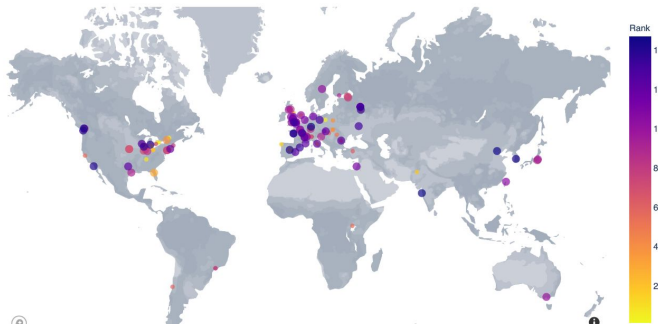
Sites' ranking based on their measures

The darker the color, the worse their performance. Smaller points indicate missing set of measures.

Click on a site in the map to see an overview over the past days

RRC-KI

Latency hosts		Throughput hosts					
IPv4	IPv6	IPv4	IPv6				
1	0	1	0				
PACKETLOSS (packets)		THROUGHPUT (Mbps)					
TODAY IN	TODAY OUT	TODAY IN	TODAY OUT				
0.02	0	35.77	109.75				
Change over the past 3 days (%)		Change over the past 3 days (%)					
05/10	06/10	07/10	05/10	06/10	07/10		
IN	+0.73	+9.46	+12.13	IN	-0.06	-0.16	-0.91
OUT	+0.37	-0.91	+0.31	OUT	+0.12	-0.38	-0.51
OWD (ms)		RETRANSMITS (packets)					
TODAY IN	TODAY OUT	TODAY IN	TODAY OUT				
52.26	6.25	3986.69	288.47				
Change over the past 3 days (%)		Change over the past 3 days (%)					
05/10	06/10	07/10	05/10	06/10	07/10		
IN	+0.09	+0.01	+0.03	IN	+2.91	-0.07	+2.09
OUT	-0.13	-0.67	-0.01	OUT	-0.07	-0.3	+0.49



GSI-LCG2 as destination of measures



<https://ps-dash.uc.ssl-hep.org/>

**Purpose:** provides a user dashboard to explore analyzed and summarized perfSONAR data.

**Currently:**

- Allows users to monitor their sites
- Provides tools for detecting basic problems

**Future plans:**

- Add today's Alarms
- Add traceroute data & plots
- Refine ranks
- Deduct possible cause for found issues



# ATLAS Alarms & Alerts Service

Alarms & Alerts

Home

## Alarms

- Analytics
- Networking
  - Perfsonar
    - bad owd measurements
    - indexing
    - complete packet loss
    - Firewall issue
    - large clock correction
    - high packet loss on multiple links
    - Bandwidth increased from/to multiple sites
    - Bandwidth decreased from/to multiple sites
    - Bandwidth increased
    - Bandwidth decreased
    - destination cannot be reached from multiple
    - destination cannot be reached from any
    - source cannot reach any
    - high packet loss
- Virtual Placement
- SLATE
- WFMS

<https://aaas.atlas-ml.org/>

**Purpose:** provides user-subscribable alerting for specific types of network issues found by analyzing perfSONAR data

Currently available:

- Main packet loss issues
- Main throughput issues

Future plans:

- Add traceroute alarms:
  - Destination never reached
  - **Network path changes**
  - Node causes issues with multiple sites

# Bandwidth Alarms

Detecting changes in measured throughput wrt. 21-day average (ipv4, ipv6)  
Currently working on creating high-level alarms (aggregating multiple alarms and running correlations with latencies and path alarms)

## Example: Alarms generated for Sat 26th March

Herewith a list of alarms you subscribed to. You may change preferences by visiting <https://aaas.atlas-ml.org>.

Sat, 26 Mar 2022 04:08:44    Networking/Perfsonar/Bandwidth decreased from/to multiple sites    Bandwidth decreased from/to multiple sites  
tags: IN2P3-CC

Bandwidth decreased for ipv4 links between site IN2P3-CC to sites: ['AGLT2', 'UFlorida-HPC'] change in percentages: [-55, -100]; and from sites: ['GLOW', 'IN2P3-LAPP', 'SiGNET', 'UTA\_SWT2'], change in percentages: [-72, -69, -27, -96] with respect to the 21-day average.

Sat, 26 Mar 2022 04:08:44    Networking/Perfsonar/Bandwidth decreased from/to multiple sites    Bandwidth decreased from/to multiple sites  
tags: IN2P3-CC

Bandwidth decreased for ipv6 links between site IN2P3-CC to sites: ['CA-VICTORIA-WESTGRID-T2', 'GLOW', 'SiGNET', 'pic'] change in percentages: [-16, -36, -11, -91]; and from sites: ['BEIJING-LCG2', 'CIT\_CMS\_T2', 'IN2P3-CPPM', 'IN2P3-LPSC', 'UAM-LCG2'], change in percentages: [-96, -49, -98, -14, -99] with respect to the 21-day average.

Sat, 26 Mar 2022 04:08:44    Networking/Perfsonar/Bandwidth decreased from/to multiple sites    Bandwidth decreased from/to multiple sites  
tags: RRC-KI-T1

Bandwidth decreased for ipv4 links between site RRC-KI-T1 to sites: ['BNL-ATLAS', 'IN2P3-LPSC', 'UKI-SCOTGRID-ECDF'] change in percentages: [-12, -30, -13]; and from sites: ['DESY-ZN', 'IN2P3-CPPM'], change in percentages: [-45, -81] with respect to the 21-day average.



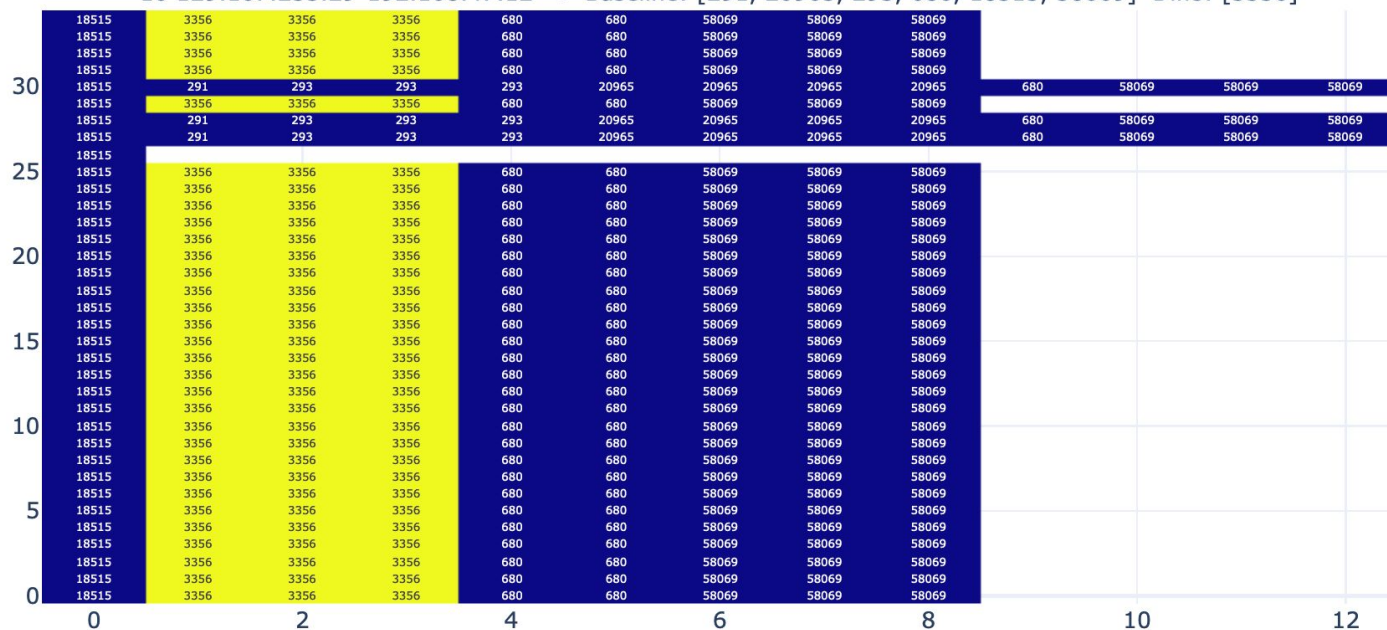
# Network Path Anomalies Detection

Detecting changes in ASNs sequences across all our traceroutes

Example: UTA\_SWT2 -> FZK-LCG2

(each row is a traceroute for this path, traces run every 30 minutes)

10 129.107.255.29-192.108.47.12 => Baseline: [291, 20965, 293, 680, 18515, 58069] Diffs: [3356]



# Summary

- OSG in collaboration with WLCG operates a comprehensive network monitoring platform
  - Provides data and feedback to LHCOPN/LHCONE, HEPiX, WLCG and OSG communities
- The IRIS-HEP and SAND projects have produced some new tools for exploring and utilizing our network data
- Developing high-level services based on perfSONAR measurements that will help sites, experiments and R&Es receive targeted alarms/alerts on existing issues in the infrastructure
- We have to continue to watch our network monitoring infrastructure as it is a complex system with lots of areas for issues to develop.

# Acknowledgements

We would like to thank the **WLCG**, **HEPiX**, **perfSONAR** and **OSG** organizations for their work on the topics presented.

In addition we want to explicitly acknowledge the support of the **National Science Foundation** which supported this work via:


- [OSG: NSF MPS-1148698](#)
- [IRIS-HEP: NSF OAC-1836650](#)





# Useful URLs

- OSG/WLCG Networking Documentation
  - <https://opensciencegrid.github.io/networking/>
- perfSONAR Infrastructure Dashboard
  - <https://atlas-kibana.mwt2.org:5601/s/networking/goto/9911c54099b2be47ff9700772c3778b7>
- perfSONAR Dashboard and Monitoring
  - <http://maddash.opensciencegrid.org/maddash-webui>
  - [https://psetf.opensciencegrid.org/etf/check\\_mk](https://psetf.opensciencegrid.org/etf/check_mk)
- perfSONAR Central Configuration
  - <https://psconfig.opensciencegrid.org/>
- Toolkit information page
  - <https://toolkitinfo.opensciencegrid.org/>
- Grafana dashboards
  - <http://monit-grafana-open.cern.ch/>
- ATLAS Alerting and Alarming Service: <https://aaas.atlas-ml.org/>
- The pS Dash application: <https://ps-dash.uc.ssl-hep.org/>
- ESnet WLCG DC Dashboard:  
<https://public.stardust.es.net/d/lkFCB5Hnk/lhc-data-challenge-overview?orgId=1>



**Backup Slides Follow**


# 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. (See coming [talk](#) :) )



# 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**

