

# Networking Topics for WLCG

**Shawn McKee (University of Michigan Physics), Marian Babik (CERN)**

**Fall 2024 HEPiX Meeting, University of Oklahoma**

**November 07, 2024**

# Introduction & Overview

Today we will provide updates on various WLCG network related activities including:

- Network monitoring including perfSONAR
- SciTags (traffic identification)
- Analytics work

Much of this work is being driven by and coordinated with the WLCG Network Data Challenges and associated mini-challenges.

# Network Monitoring and Testing

# perfSONAR 5.1.4

Recently [released](#) (Oct 17th) (~40% of perfSONARs on 5.1.4) Includes additional fixes for OSG-LHC identified issues

Fixes for configuration and error handling to increase robustness of toolkits

## Main features (5.1):

- New Grafana interface (replaces MaDDash!!)
- Threaded iperf3 support
- Conversion of pSConfig from Perl to Python (better support; easier features)
- **Removal of EL7, Debian 10 and Ubuntu 18**
- Supports EL9, Debian 11-12, Ubuntu 20, 22
- Multiple fixes for Elmond, Logstash, registration daemon, pscheduler, ...

- Install help: `curl -s`

```
https://raw.githubusercontent.com/perfsonar/project/master/install-perfsonar | sh -s -  
--help
```

# perfSONAR Infrastructure

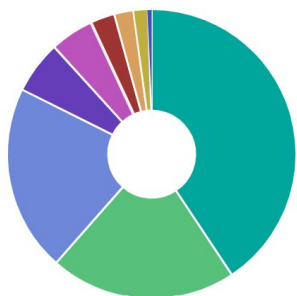
181

Active perfSONARs

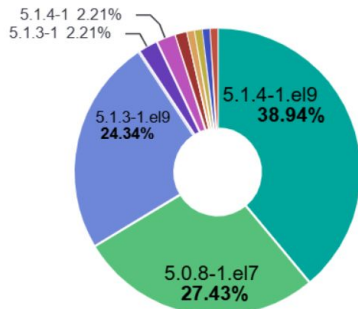
68

Communities in Use

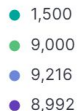

NIC speed



perfSONAR distributions

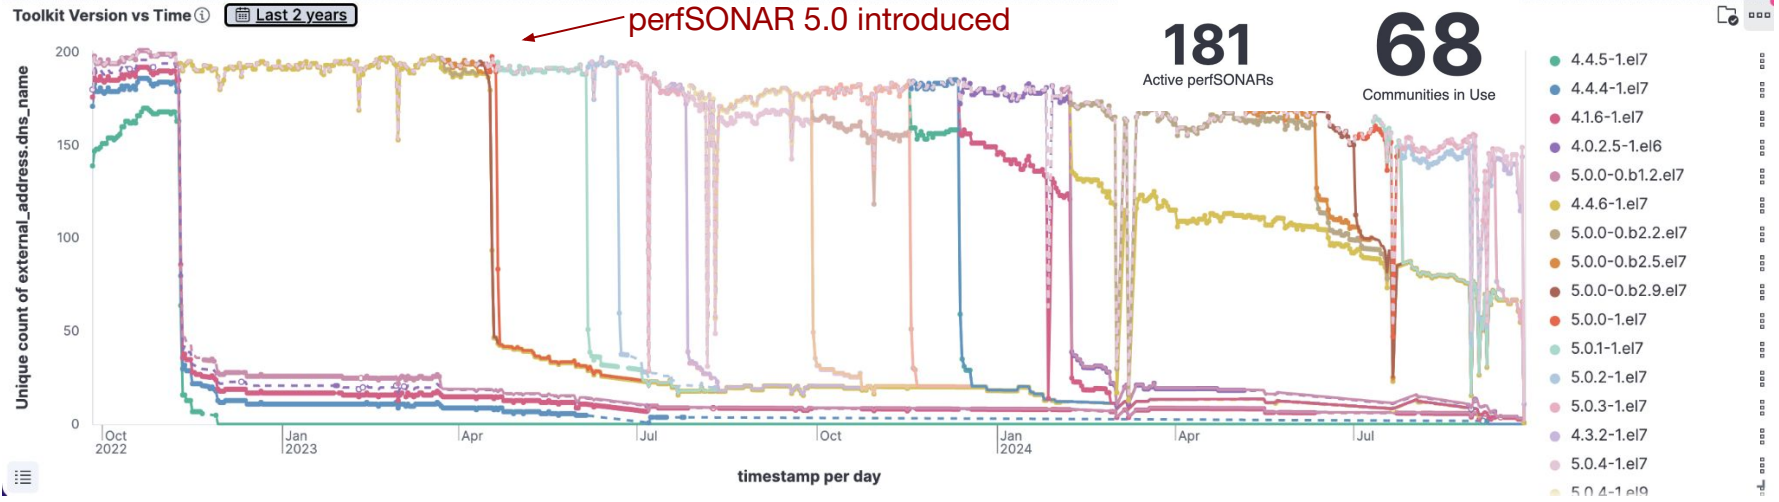


pS NIC MTU



- 69% toolkits now on 5.1 and EL9
  - 27% still on EL7 !
- Core deployments are still on 10Gbps, but we have about 20% with 100Gbps
  - For WLCG/OSG testing purposes 10Gbps is still sufficient
  - **Important** to refresh HW along with the update to EL9
- MTU - around 36% on jumbo frames (9000), rest is on standard frames (1500)
- We have small testbed with about 10 perfSONARs with BBRv3 enabled

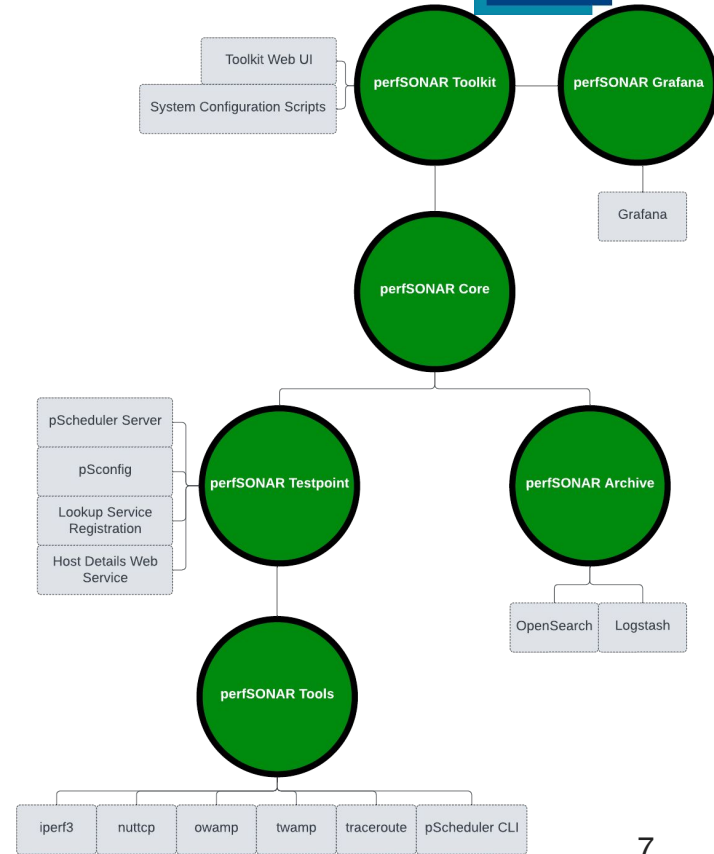
# perfSONAR Infrastructure Evolution



- Long-term trend of perfSONAR releases adoption shows a steady downtrend
  - Only reachable “toolkit version” installations are shown
  - Stability and reliability of the releases clearly playing a role
- New strategy and deployment models will need to be considered
  - Exploring simple deployments (without complex components such as ES)
  - Providing means to easily reset/restart deployment (to avoid accumulation of issues)
  - Easy to co-locate with storages (or even co-hosted on storages with multiple NICs)

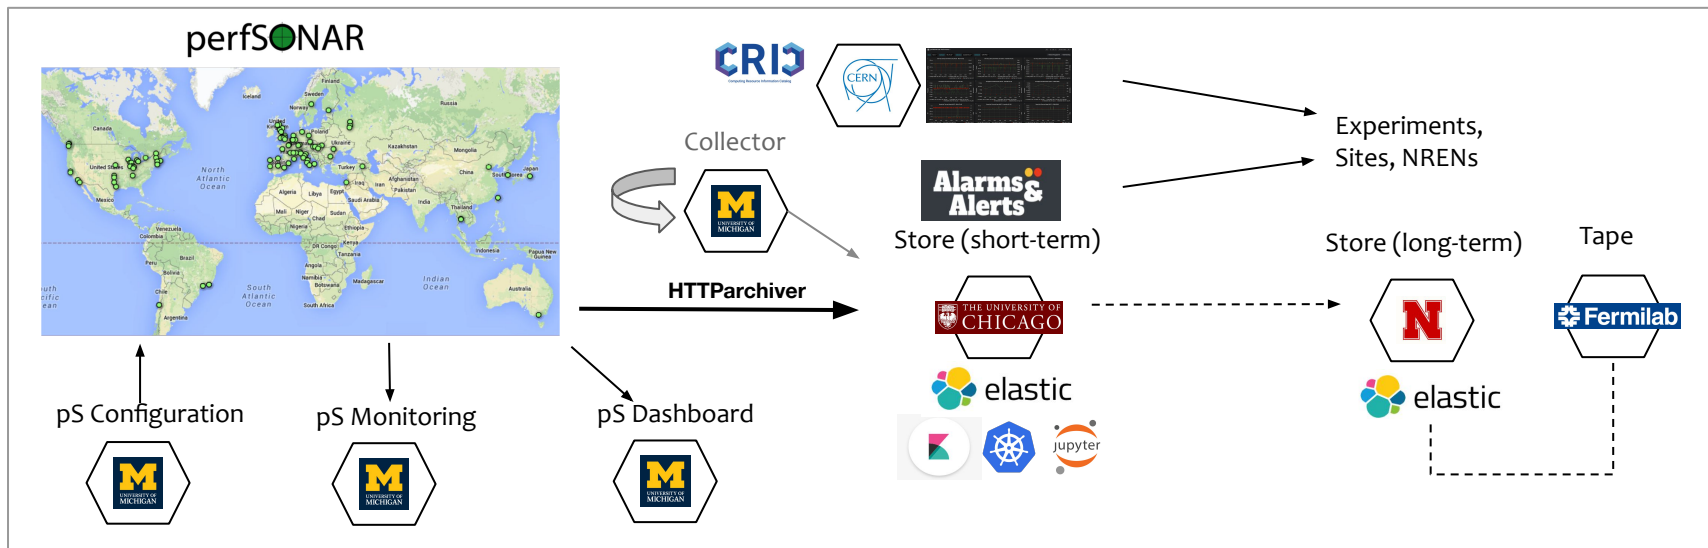
# perfSONAR Testpoint

- Motivated unreliability of full toolkit deploy
- perfSONAR Testpoint:
  - Automatically runs tests on a regular schedule
  - Participates in a centrally managed set of tests
  - Publishes the existence of a measurement node
- Can run in a container, i.e. starting testpoint as easy as
  - `docker run -d --name perfsonar-testpoint --net=host perfsonar/testpoint`
- Main concerns
  - No local cache that would sustain an outage of central services (48-72 hours)
  - No way to remotely monitor - needs extra packages/containers to run node\_exporter and perfsonar exporter
  - Missing auto-updates and persistent configuration
  - Performance and integration with OS needs testing
- Looking for volunteers willing to deploy, test and develop the missing components



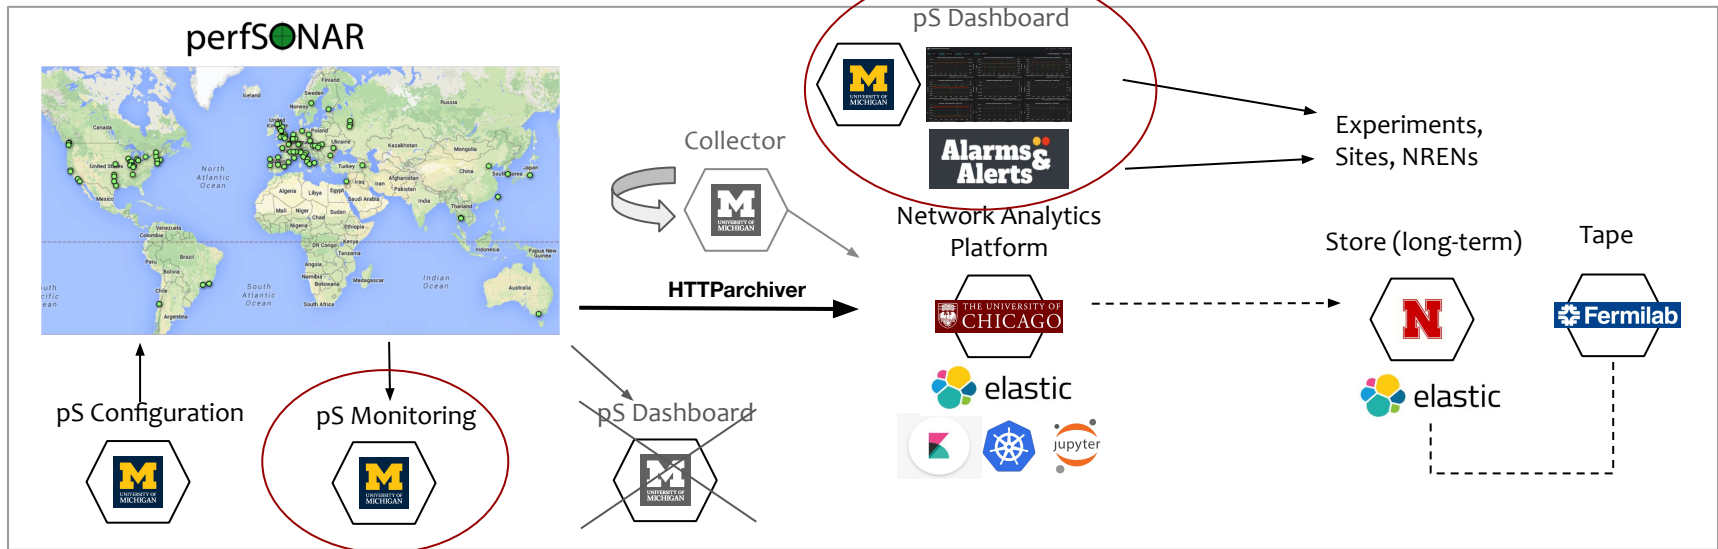
# Network Measurement Platform Status

- Our platform collects, stores, configures and transports all network metrics
- Evolution based on the perfSONAR 5 already partially implemented.
  - Now directly publishing results from perfSONARs to ES@UC
  - Collector used only as a fallback;
  - **WLCG CRIC** now used for topology



# Network Measurement Platform Plans

- Evolution based on the perfSONAR 5 already partially implemented.
  - Forwarding to UNL and backup to FNAL still to be implemented
  - **pS Monitoring** - update to latest Checkmk and enable SSO authentication
  - **ps Dashboard** - integrate with Analytics Platform/Grafana (**retire maddash**)
  - **ps Configuration** - clarify development roadmap and support



# perfSONAR Infrastructure Monitoring

- Updated to Checkmk 2.3.0 (from 1.6.0)
- Integration with CILogon (single-sign on) - moving away from x509 certs
- New tests
  - Node diagnostics based on “pscheduler troubleshoot” command
  - Tracking measurements in central ElasticSearch
- Now in pre-production at [psetf-itb.aglt2.org](https://psetf-itb.aglt2.org)

**Main dashboard**  
Monitor > Overview > Main dashboard

Dashboard Add Dashboards Display Help

**Host statistics**  
308 Up  
0 In downtime  
0 Unreachable  
0 Down  
308 Total

**Service statistics**  
1854 OK  
0 In downtime  
0 On down host  
276 Warning  
1224 Unknown  
1661 Critical  
5015 Total

**Host Problems (unhandled)**

State	Host	Icons	Age	Summary
CRIT	perfsonar.dur.scotgrid.ac.uk	📧	191 s	connect to address perfsonar.dur.scotgrid.ac.uk at 443: Connection timed out
CRIT	perfsonar.nersc.gov	📧	8 m	TCP CRITICAL - Invalid hostname, address or socket address
CRIT	ps-development.bnl.gov	📧	23 m	connect to address ps-development.bnl.gov and port 443: Connection timed out
CRIT	ps-latency.clumeq.mcgill.ca	📧	31 m	connect to address ps-latency.clumeq.mcgill.ca at 861: Connection timed out

**Service Problems (unhandled)**

Your query produced more than 1000 results. Repeat query and allow more results. Note: the shown results are incomplete a

**Events of recent 4 hours**

Time	Host	Service	Summary
191 s	perfsonar1.nipne.ro	perfsONAR services: web/https IPv6	connect to address perfsonar1.nipne.ro and port 443: Connection timed out
8 m	perfsonar1.nipne.ro	perfsONAR services: web/https IPv6	connect to address perfsonar1.nipne.ro and port 443: Connection timed out
23 m	psonartest2.fnal.gov	perfsONAR services: owamp	TCP OK - 127.300 second response time on psonartest2.fnal.gov port 861
31 m	t2-psn2.jinr.ru	perfsONAR services: pscheduler	UNKNOWN - Exception caught while executing plugin (invalid literal for int() with base 10: b'<DOCTYPE HTML PUBLIC "-//IETF/DTD HTML 2.0/EN">
47 m	ps100.farm.particle.cz	perfsONAR configuration: meshes	OK - Auto-URL configured correctly

**Overview**

Hosts	Unhandled p.
308	0

**Services**

Services	Unhandled p.
5015	3161

**Events**

Events	Unhandled p.
0	0

**Bookmarks**

**Master control**

- Notifications
- Service checks
- Host checks
- Flap detection
- Event handlers
- Performance data

**Quicksearch**

# perfSONAR Infrastructure Monitoring



checkmk



Monitor



Customize



Setup

## Services of host ps-latency.lhcmn.triumf.ca

Monitor > Overview > All hosts > ps-latency.lhcmn.triumf.ca > Services of host

Commands Host Services Export Display Help ↻

Acknowledge problems

Schedule downtimes

Filter


Show checkboxes

ps-latency.lhcmn.triumf.ca

### ps-latency.lhcmn.triumf.ca

State	Service	Icons	Summary	Age	Checked	Perf-O-Meter
OK	perfSONAR configuration: contacts	☰	OK - Contact and organization found	2024-07-02 17:26:53	47 m	
OK	perfSONAR configuration: location	☰	OK - Location: -122.5428/49.1646	2024-07-02 17:26:53	47 m	
OK	perfSONAR configuration: meshes	☰	OK - Auto-URL configured correctly	2024-07-02 17:26:53	47 m	
WARN	perfSONAR ES freshness: owamp	☰	WARNING - ES 52.78 % complete for event type: histogram-owdelay	6 h	6 h	
CRIT	perfSONAR esmond freshness: owamp	☰	CRITICAL - esmond 0.0 % complete for event type: histogram-owdelay	2024-07-01 16:22:33	112 m	
OK	perfSONAR hardware check	☰	OK - CPU:1/24cores/3700.0000Mhz RAM:62GB NIC:10.0Gbps/9000MTU/IPv6 enabled	2024-07-02 17:26:53	47 m	
OK	perfSONAR json summary	☰	OK - Toolkit metadata successfully retrieved	2024-07-02 17:26:50	47 m	
OK	perfSONAR services: ntp	☰	OK - NTP synchronized	2024-07-02 17:26:53	47 m	
OK	perfSONAR services: owamp	☰	TCP OK - 0.056 second response time on ps-latency.lhcmn.triumf.ca port 861	2024-08-25 15:14:32	7.15 s	
OK	perfSONAR services: pscheduler	☰	OK - pscheduler stats retrieved	2024-07-02 17:35:25	39 m	
WARN	perfSONAR services: pscheduler diags	☰	Performing basic troubleshooting of ps-latency.lhcmn.triumf.ca.	19 h	19 h	
OK	perfSONAR services: regular testing/pscheduler	☰	OK - pscheduler is enabled and running	2024-08-31 15:26:54	47 m	
OK	perfSONAR services: versions	☰	OK - Toolkit version found: 5.0.8-1.e17	2024-07-02 17:26:53	47 m	
OK	perfSONAR services: web/https IPv6	☰	TCP OK - 0.056 second response time on ps-latency.lhcmn.triumf.ca port 443	2024-08-25 15:14:41	14 m	

# perfSONAR Infrastructure Monitoring: Freshness



## Service perfSONAR ES freshness: owamp, ps-latency.lhcomn.triumf.ca

Monitor > Overview > All hosts > ps-latency.lhcomn.triumf.ca > Services of host > Service

Commands Service Host Export Display Help

Acknowledge problems Schedule downtimes Filter Show checkboxes Services of host

1 row

Site alias	Local site etf
Host name	ps-latency.lhcomn.triumf.ca
Service description	perfSONAR ES freshness: owamp
Service labels	
Service icons	
Service state	<b>WARN</b>
Summary	<p>WARNING - ES 52.78 % complete for event type: histogram-owdelay</p> <p>Time-range: 3600 Even-types checked: histogram-owdelay</p> <p>Mesh (Event-type): Canadian Tier1-2 Latency (histogram-owdelay) Missing destinations: 1 Missing destinations: lcg-bw.sfu.computeCanada.ca</p> <p>Mesh (Event-type): LHCONE Latency (histogram-owdelay) Missing destinations: 6 Missing destinations: lat.scinet.utoronto.ca, perfsonar-ps-01.desy.de, wash-ps-ps-lat.es.net, atlas-npt1.bu.edu, perfsonar2.ultralight.org, perfsonar-latency-v471-portal.gridka.de</p> <p>Mesh (Event-type): OPN Latency (histogram-owdelay) Missing destinations: 2 Missing destinations: lhperfmon.bnl.gov, t1-ps01-lt.gridpp.rl.ac.uk</p>
Details	<p>Mesh (Event-type): WLCG ATLAS Latency (histogram-owdelay) Missing destinations: 28 Missing destinations: sonar2.itim-cj.ro, ps1.oceph.ou.edu, perfsonar01.lcg.cscs.ch, perfsonar-lt.tier2.hep.manchester.ac.uk, psonar1.lal.in2p3.fr, mwt2-ps01.campuscluster.illinois.edu, perfsonar2.na.infn.it, perfsonar2.cc.kek.jp, ps-latency.atlas.unimelb.edu.au, ps-latency.clumeq.mcgill.ca, perfsonar-ps-01.desy.de, t2ps-latency.physics.ox.ac.uk, atlas-npt1.bu.edu, lhperfmon.bnl.gov, perfsonar-latency.grid.pub.ro, perfsonar2.roma1.infn.it, lcg-lat.sfu.computeCanada.ca, psuta01.atlas-sw2.org, lcgperfradar.dnp.fmph.uniba.sk, atrog007.nipne.ro, perfsonar-latency.esc.qmul.ac.uk, hepInx129.pp.rl.ac.uk, lcgps01.gridpp.rl.ac.uk, gridpp-ps-lat.ecdf.ed.ac.uk, psnr-farm10.slac.stanford.edu, psiflc01.ific.uv.es, perfson1.pggrid1.rhul.ac.uk, perfsonar2-grid.uaic.ro</p> <p>Working destinations: argus-11.net2.mghpcc.org, perfsonar-latency.grid.surfsara.nl, perfsonar01.ft.uam.es, perfsonar2.nipne.ro, ps01-l.farm.particle.cz, psum02.itep.ru, llrpsonar1.in2p3.fr, ps-gsd01.sdfarm.kr, marper02.in2p3.fr, psmsu01.aglt2.org, psum01.aglt2.org, ps002.gla.scotgrid.ac.uk, perfsonar01.datagrid.cea.fr, pygrid-sonar2.lancs.ac.uk, perfsonar-ps.ndgf.org, lapp-ps02.in2p3.fr, ps0002.m45.ihep.su, ps-latency.hepnetcanada.ca, t2ps-latency2.physics.ox.ac.uk, hepsonar2.ph.liv.ac.uk, lcg-pslat.uw.computeCanada.ca, t1-pfsn2.jlnr-t1.ru, perfsonar2.ihep.ac.cn, lpnhe-psl.in2p3.fr, tech-ps.hep.technion.ac.il, perfsonar02-iep-grid.saske.sk, btw-lat.grid.kiae.ru, uct2-net1.mwt2.org, iut2-net09.iu.edu, echolot1.zeuthen.desy.de, ccperfsnar2.in2p3.fr, perfsonar1.icepp.jp, btw-lat.t1.grid.kiae.ru, psI01-gva.cern.ch, perfsonar2-de-kit.gridka.de, perfsonar-ow.cnaf.infn.it, psonar9-191-2.fnal.gov, picperfsonar-latency.pic.es</p> <p>Documentation for this check can be found at </p>

Help

# perfSONAR Infrastructure Monitoring: pscheduler

checkmk

## Service perfSONAR services: pscheduler diags, psb02-gva.cern.ch

Monitor > Overview > All hosts > psb02-gva.cern.ch > Services of host > Service

30

Commands Service Host Export Display Help

Acknowledge problems Schedule downtimes Filter Show checkboxes Services of host

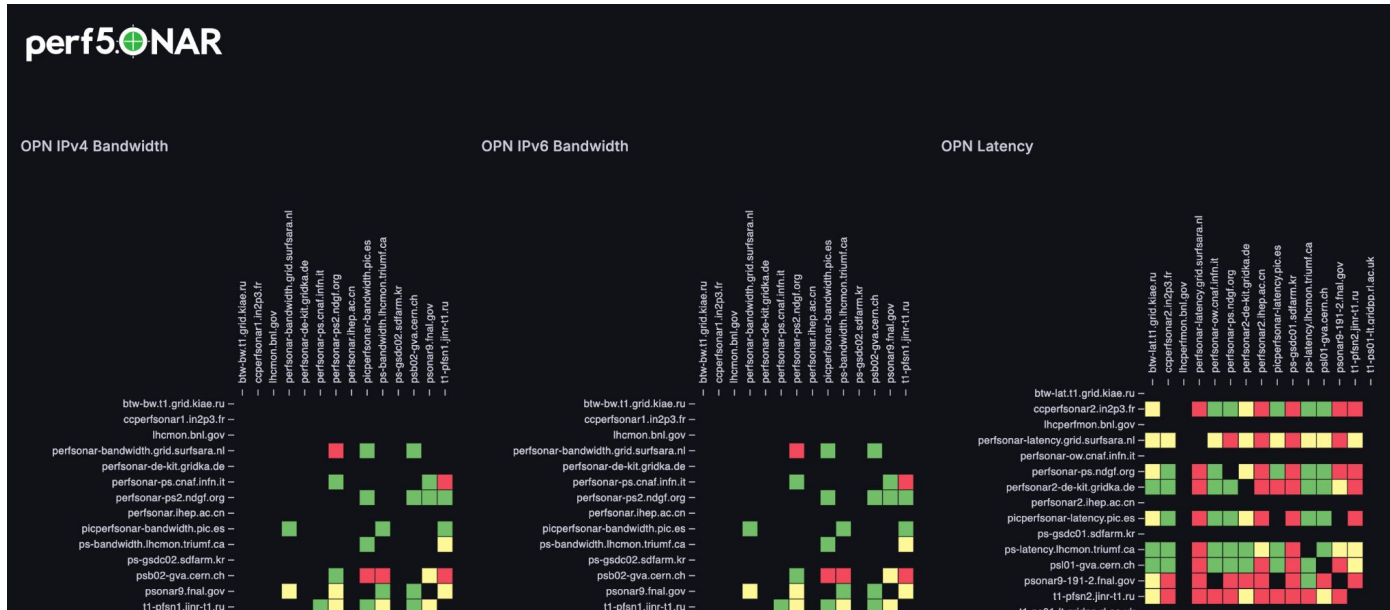
1 row

Site alias	Local site etf
Host name	psb02-gva.cern.ch
Service description	perfSONAR services: pscheduler diags
Service labels	
Service icons	☰
Service state	OK
Summary	Performing basic troubleshooting of psb02-gva.cern.ch.
Details	<pre>psb02-gva.cern.ch: Checking that host "psb02-gva.cern.ch" resolves... 2001:1458:301:86::100:b Looking for pScheduler... OK. Fetching API level... 6 Checking clock... Unsynchronized (Not considered fatal) Exercising API... Archivers... Contexts... Tests... Tools... OK. Fetching service status... OK. Checking services... Ticker... Scheduler... Runner... Archiver... OK. Checking limits... OK. Last run scheduled... 5 seconds ago Last run completed... 6 seconds ago Server Statistics: Archiving Backlog ..... 0 Upcoming ..... 0 HTTP Queue Backlog ..... 0 Length ..... 1 Runs Pending ..... 17767 On Deck ..... 0 Running ..... 28 Cleanup ..... 3 Finished ..... 91390 Overdue ..... 0 Missed ..... 1 Failed ..... 16064 Preempted ..... 0 Non-Starting ... 84  pScheduler appears to be functioning normally.</pre>

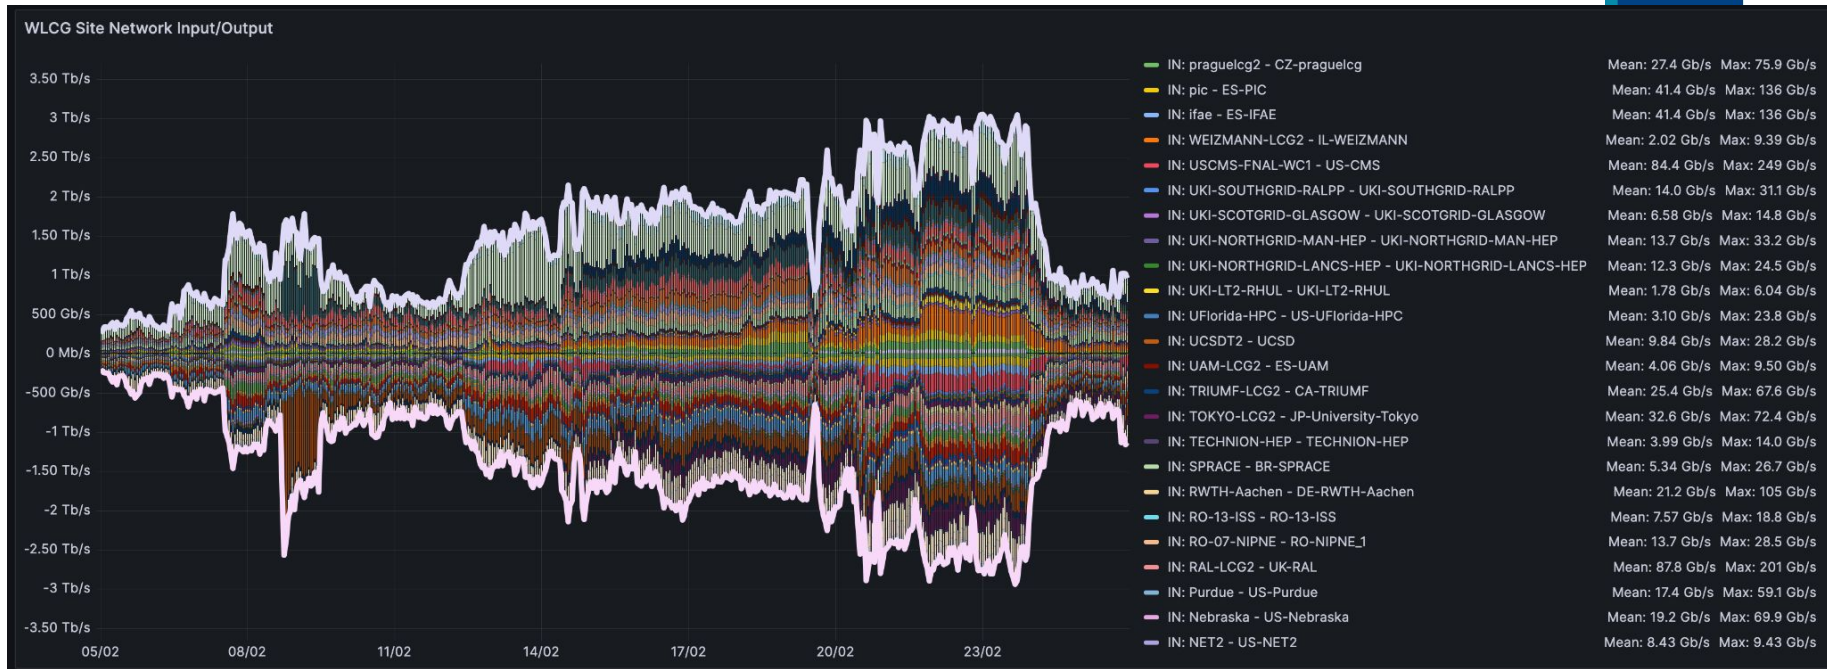
Monitor  
Customize  
Setup  
Help  
User

# Replacement perfSONAR Dashboard

- New dashboard service replacing previous maddash
- Based on Grafana - based on perfSONAR 5 code base
  - Modified to use the **central ElasticSearch**, which uses different schema
  - Generation of dashboards required some new code which is now upstream



# WLCG Site Network Monitoring



**Site Network Utilisation - computed from aggregated utilisation (SNMP counters) provided by sites via simple API. Screenshot shows network utilisation during DC24 as seen by the sites. We still have sites that have NOT completed configuring this. If your site is not shown, please follow up using the guidance at <https://gitlab.cern.ch/wlcg-doma/site-network-information>**

# SciTags (Data Visibility on the Network)

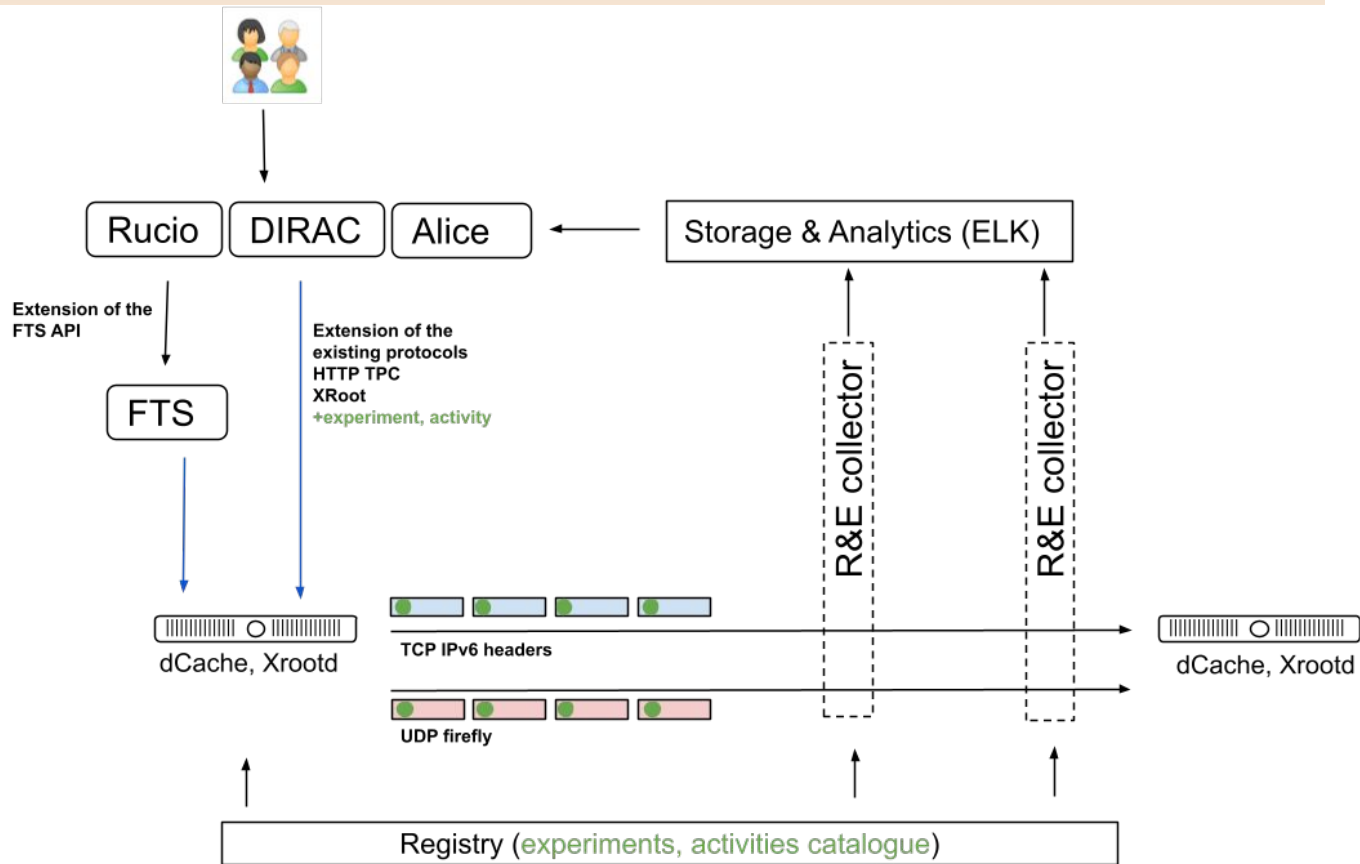
# Network Visibility and Scitags

- Scientific Network Tags (scitags) is an initiative promoting identification of the science domains and their high-level activities at the network level.



- Enable tracking and correlation of our transfers with **Research and Education Network Providers (R&Es)** network flow monitoring
  - Utilizing packet and flow marking to identify traffic owner/purpose.
- Experiments** can better understand how their network flows perform along the path
  - Improve visibility into how network flows perform (per activity) within R&E segments
  - Get insights into how experiments are using the networks, get additional data from R&Es on behaviour of our transfers (bottlenecks, troubleshooting, optimisation, etc.)
- Sites** could get visibility into how different network flows perform

# How scitags work



# Finding More Information: <https://scitags.org>


Code

Technical Spec

Mailing List

scitags.org

Network Flow and Packet Marking for  
Global Scientific Computing



Scientific network tags (scitags) is an initiative promoting identification of the science domains and their high-level activities at the network level.

It provides an open system using open source technologies that helps *Research and Education (R&E) providers* in understanding how their networks are being utilised while at the same time providing feedback to the *scientific community* on what network flows and patterns are critical for their computing.

Our approach is based on a network tagging mechanism that marks network packets and/or network flows using the science domain and activity fields. These tags can then be captured by the *R&E providers* and correlated with their existing netflow data to better understand existing network patterns, estimate network usage and track activities.

The initiative offers an **open collaboration on the research and development of the packet and flow marking prototypes** and works in close collaboration with the scientific storage and transfer providers to enable the marking capability. The project is currently in the prototyping phase and is open for participation from any science domain that require or anticipate to require high throughput computing as well as any interested *R&E providers*.

#### Participants



#### Upcoming and Past Events

- March 2022: LHCOPN/LHCONE workshop
- November 2021: GridPP Technical Seminar (slides)
- November 2021: ATLAS ADC Technical Coordination Board
- October 2021: LHCOPN/LHCONE workshop (slides)
- September 2021: 2nd Global Research Platform Workshop (slides)

Hosted on GitHub Pages — Theme by orderedlist

Presentations

# Current Status: DDM & Storages

## Implementation status:

### Propagation:

- **Rucio** supports Scitags from 32.4.0
- **FTS/gfal2** support Scitags from 3.2.10/2.21.0

### Storages:

- **XRootD** provides [Scitags implementation](#) (from 5.0+)
- **EOS** provides Scitags support from 5.2.19+
  - Project on production rollout at CERN for WLCG has been approved
- **dCache** prototype exists, roadmap for release pending
- [StoRM](#) prototype exists, planning release and pre-production deployment

### Collectors:

- Production deployments at ESnet and Jisc

# SciTags & Data Challenge 24

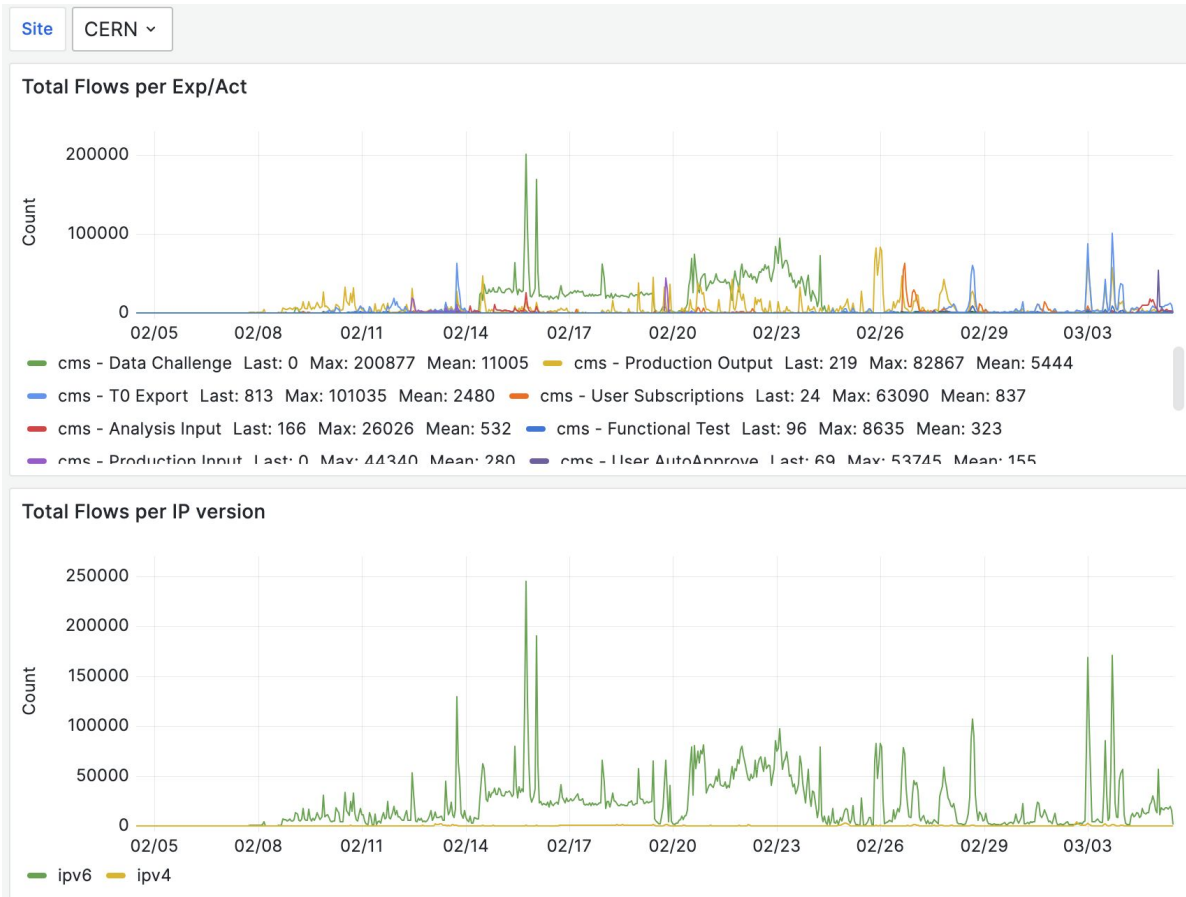
- **Scitags Deployment**

- 80% of EOS CMS (production), UNL production storage
- Flow labeling functionality (fireflies)

- **Results:**

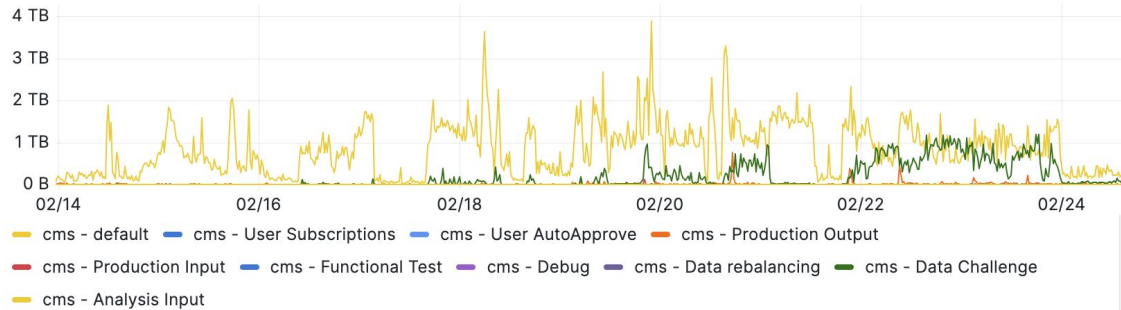
- **Confirmed the capability to propagate Scitags all the way to the storages (for both ATLAS and CMS)**
- Sending fireflies (from XRootd, EOS storages)
- **Collection and visualisation at ESnet collector**
  - Results shown in [live dashboard](#)
- With limited deployment we were able to get valuable insights into flow durations, their characteristics (splits by exp/activity), sources of IPv4 traffic (split by applications) and potential impact of new TCP congestion algorithms (performance correlated with flow data)

# SciTags & Data Challenge 24



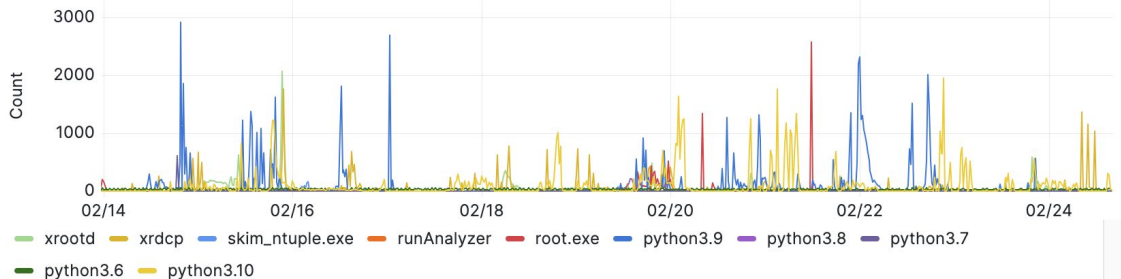
# DC24: University of Nebraska (UNL)

Bytes Sent per Exp/Act



Bytes sent per Exp/Activity - shows non-FTS traffic was dominant

Total Flows per Application (top10)



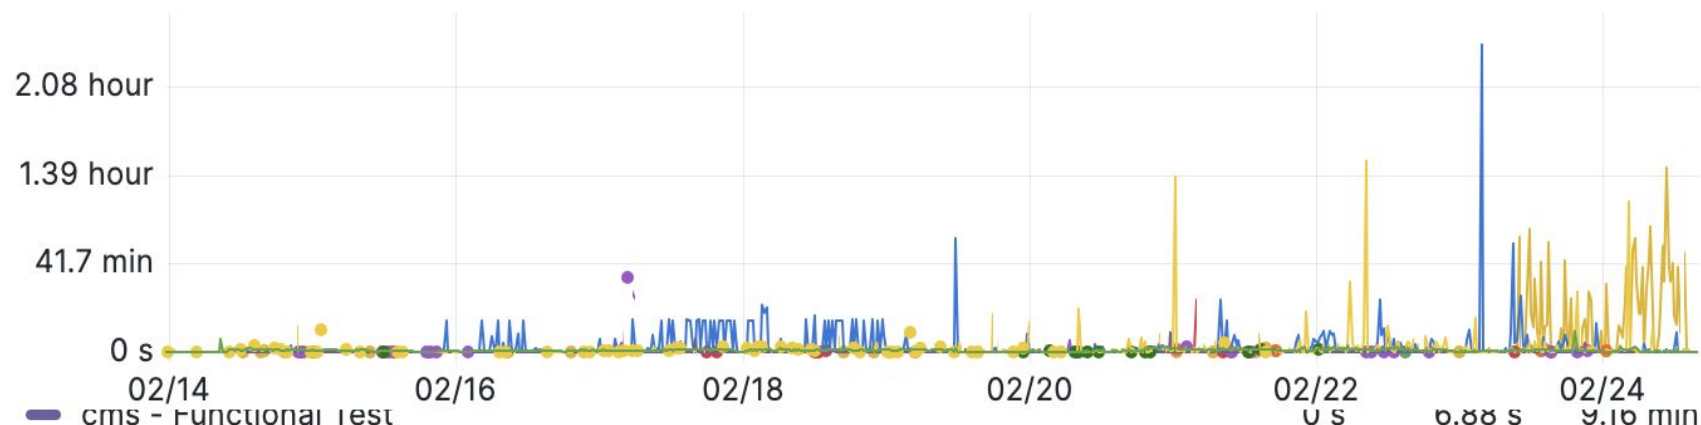
Capability to show a split by application (as reported by xrootd)

# DC24: CERN EOS CMS

Median duration of flows split by Exp/Activity

Duration of DC flows was quite short wrt. production/rebalancing

Median Duration Received per Exp/Act

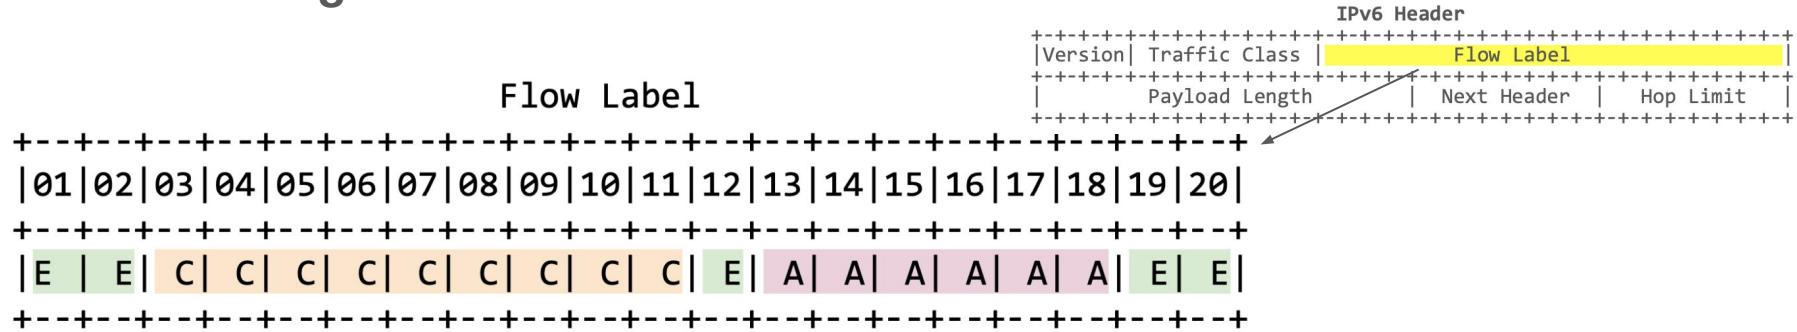


cms - Functional test	0 s	0.88 s	9.16 min
cms - Debug	0 s	11.6 s	2.08 min
cms - Data rebalancing	0 s	1.97 min	1.50 hour
cms - Data Challenge	0 s	46.2 s	9.93 min

# SciTags Research & Development

# Technical Spec for Packet Marking

## Packet Marking via the use of the IPv6 Flow Label



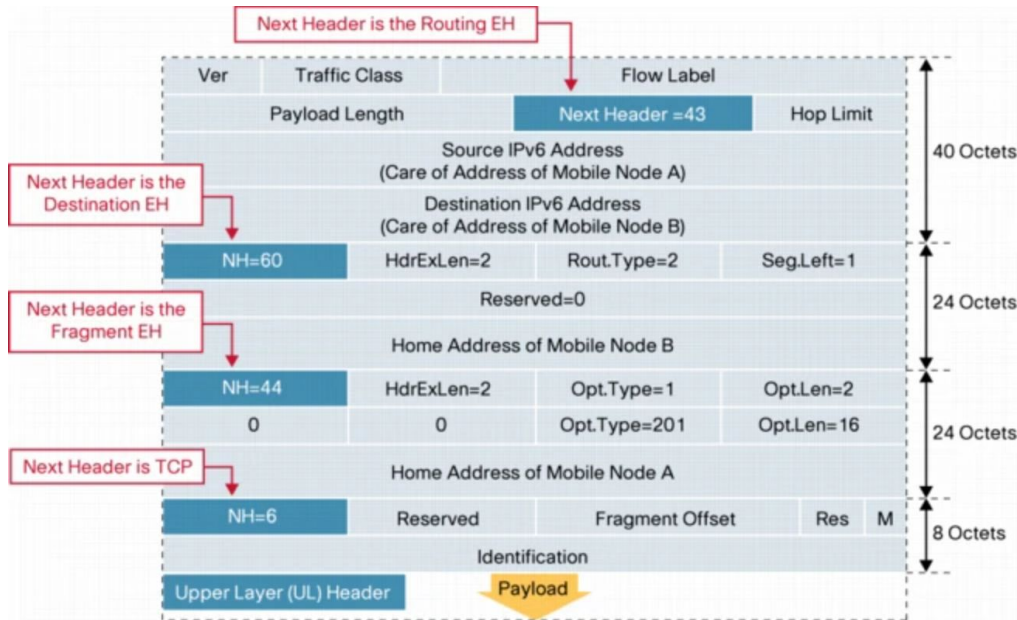
- (C) Community identifier: "Who are you affiliated with?"
- (A) Activity identifier: "What are you doing within your community?"
- (E) Entropy bits sprinkled throughout

[IETF RFC-Informational Draft](#) is available with more details  
Started exploring HbH option as an alternative ([eBPF-PDM](#), [eBPF-extHeaders](#))

# Destination Option and Hop-by-Hop Option

Based on feedback from IETF v6ops community

Started exploring Destination Option (DO) and Hop-by-Hop Option (HbH) as alternatives ([eBPF-PDM](#), [eBPF-extHeaders](#)) to flow label.



We plan to showcase packet marking with DO and HbH during SC24

Prototype implementation developed as part of **flowd** service

# Flow Daemon Library and Service (flowd)

**Service and library to prototype and test various approaches to packet marking and flow labelling**

Core features (part of scitags library):

- Plugins - determine which flows to mark and with what data (exp/act)
- Backends - implementations for packet marking and flow labelling
- Linux service - runs as systemd service

DevOps:

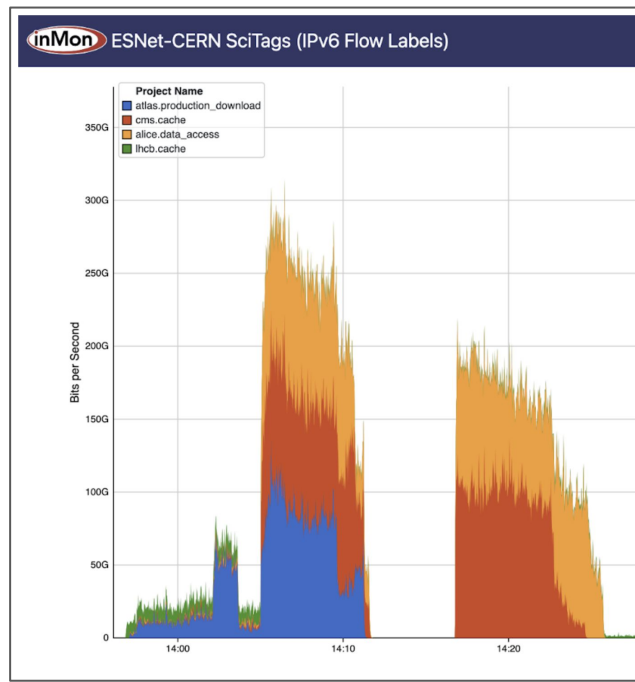
- Container, package distribution
  - Support for Alma8, Alma9, RHEL8 and RHEL9 (via public scitags repository)
  - Docker container

New “go-based” release from Pablo Collado Soto called “glowd” available **today**:

- <https://github.com/pcolladosoto/glowd/releases/tag/v1.0>

During Supercomputing 23 in Denver, we demonstrated a number of aspects of our packet and flow marking work.

- Showed **packet marking** at **300 Gbps** rates using **xrootd/iperf3** (with just two nodes; using eBPF).
- Integration with **ESnet's High-Touch Service**
  - Enabling analytics at the packet-level
- In collaboration with inMon Corp, set up packet collectors [via sflow](#) and demonstrated **real-time monitoring of flows by community/activity**.
- SC23 demo was run in collaboration with Starlight, ESnet, KIT, University of Victoria, University of Nebraska and CERN



In 2 weeks we plan to participate in Supercomputing 2024 and will demonstrate:

- Packet marking at 400 Gbps
- Packet marking using Hop-by-hop

These demonstrations help verify capability at leading edge scales.

# Analytics

## Investigate ML models/methods to process network measurements

- **Data-preprocessing, e.g.**
  - Train neural networks to predict network paths, e.g. help us fill the gaps in traceroute(s)
- **Build model(s) that represents our network(s)**
  - Network measurements are inherently noisy and therefore require robust models
- **Use ML models for anomaly detection (for alerts & alarms)**
  - Neural networks, Bayesian/probabilistic approaches
  - Detect anomalies in network paths and bandwidth measurements
  - Compare with the existing heuristic algorithms that we have developed
- **Correlate with other data**
  - Traceroutes with throughput for example, but also outside of perfSONAR, e.g. FTS
  - New types of data appearing (high-touch, scitags, in-band telemetry, etc.)

Petya Vasileva presented our recent work on analytics at CHEP 2024 in Krakow.

Details are [available](#) in the presentation

The core of the work is on correctly identifying the network path and correlating it with our metrics

The new methodology uses AI/ML to identify missing components on the network path

This is critical for then correlating problems with particular paths/locations.



# Reminder - Alarming: Why and What?

Alarming for us is defined as the identification of some issue that needs fixing.

## **An alarm:**

- Identifies a specific issue
- Involving one or more sites
- Occurs at a point in time or exists in a time range

We have been exploring various algorithms to define alarms of interest. Since it can be very difficult to identify certain types of issues from our complex data, we are trying various Artificial Intelligence (AI) & Machine Learning (ML) techniques.

We recruited an IRIS-HEP fellow (Yana Holoborodko) to work on this and we will report on her project results below.

# Student Summer Project

For 2024, we have identified one student to work with us on Alerting & Alarming:

- Yana Holoborodko (Ukraine), **Project:** Enabling Advanced Network and Infrastructure Alarms (see <https://docs.google.com/document/d/1ZjLHI7-CpALAUi3WlwQ4PP8Sg4o4qZftfT938gH4dY/edit> )

Yana had a challenging task in that she had to not only learn about our tools, data formats and previous results, but also had to come up to speed on the basics of networking, data transfers and how they can go wrong!

She completed her project at the beginning of October and she reported on it at a [Fellows talk](#) on October 7, 2024

- Her code is at <https://github.com/hooloobooroodkoo/aaas/tree/main>
- We have two new alarms related to reachability and activity for our perfSONARs that we are integrating into the [AAAS service](#)

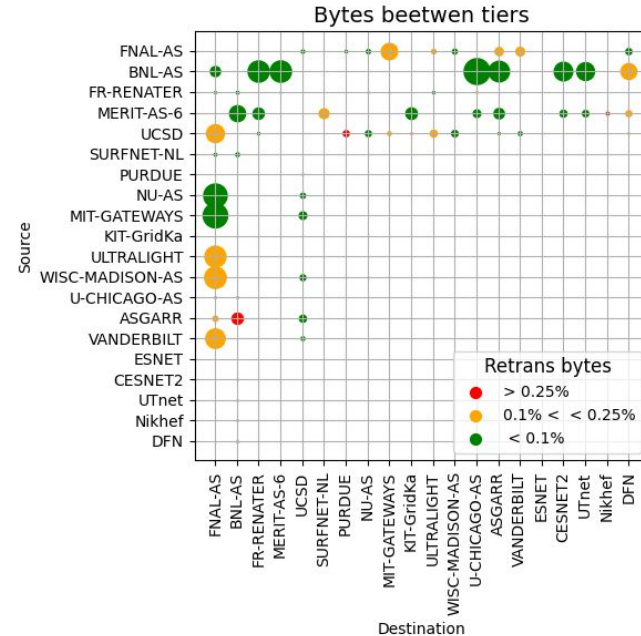
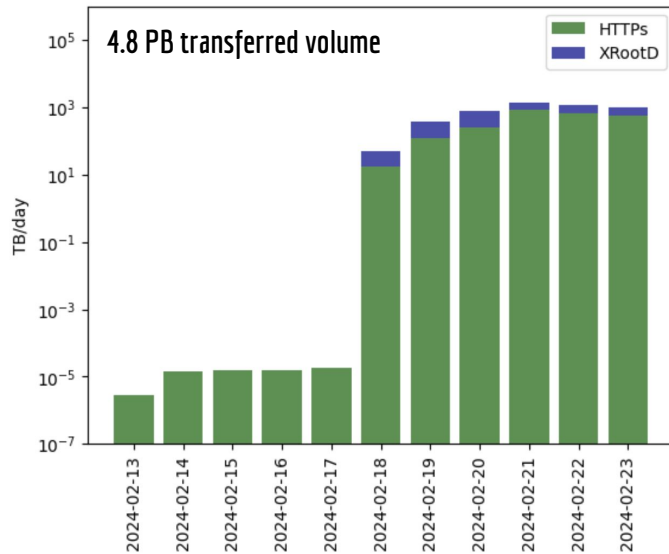
We may have the option to provide a follow-on Fellowship at CERN for 6-9 months!

# Analyzing ESnet HighTouch Data

WLCG **DC24** data inserted into Hadoop/Hive platform @ PIC [ $2.3 \cdot 10^{10}$  NumRows / 13 TiB RawDataSize]

This platform offers **fast interactive exploration** (visualization) of the data

Preliminary analysis ongoing (M. Félix & J. Flix) - more ESnet HT datasets to be explored



# Outreach & Plans

# WLCG “Network” Data Challenges

A planned series of **end-to-end** tests to demonstrate the progress of **infrastructure**, **applications** and **middleware** of WLCG sites and experiments to reach the capacity and capability required for HL-LHC scale operations.

These data challenges provide many benefits, allowing sites, networks and experiments to evaluate their progress, motivate and validate their developments in hardware and software and show readiness of technologies at suitable scale.

After DC24 (Feb 2024) we had a lot of data to examine and correlate. Many site issues were identified and especially shortcomings in some central data distribution components like FTS, Rucio and associated monitoring.

We have new plans to enable “[mini-challenges](#)” that won’t require involvement of experts to allow us to “benchmark” parts of our infrastructure.

# DC26/27 Network Planning

The next WLCG data challenge will likely be in late 2026 or early 2027 (DC26/27) and we have a number of targets to work on.

- **Network planning:** make sure our **sites** and their **local and regional networks** are aware of WLCG requirements and timeline, and plan accordingly.
  - We need to socialize our requirements at sites and with regional networks
- Update and utilize **perfSONAR** to clean up links & fix problems before DC26/27.
- Instrument and document **site networks**, for at least our largest sites.
- **IPv6** should be enabled everywhere not just because of packet marking, but because it will allow us to get back to a single stack sooner!
- **Optimize** our ability to utilize the network (jumbo frames, protocols, pacing)
- Improve net traffic **visibility** (SciTags)
- Explore and demonstrate the value of **Network Orchestration/SDN**.

# OSG-LHC/IRIS-HEP Planning

At the [IRIS-HEP retreat](#) in September 2024, we discussed how to prepare for DC26  
As mentioned, mini-challenges are an important tool that we want to enable

Goals for the next DC:

- Move the majority of our data via IPv6 and have one or more sites **IPv6-only**
- Have 80%+ of our traffic identified by SciTags
- Have SENSE/Rucio used in production at one or more sites
- Improved site network monitoring to traffic traffic by LHCONe, LHCOPN, R&E and commodity

Rough plan:

- Before the end of 2024 re-run capacity tests for US sites to determine current values
- Around February 2025, execute a joint USATLAS-USCMS **capacity** mini-challenge for North America (identify current throughput limits by site and in aggregate)
- Early-to-mid Summer 2025, execute a joint USATLAS-USCMS **capability** mini-challenge for North America (storage tokens, traffic shaping, SciTags, SDN, etc.)

Will report on this at the upcoming [WLCG DOMA meeting](#)

# Summary and Conclusion

The **collaboration of OSG, WLCG and various research projects** have created an extensive infrastructure to monitor our networks via **perfSONAR** and provide associated analytics and visualization.

- We maintain our engagement with various groups, working on improving networking for scientific research.
- While we continue to monitor and maintain our infrastructure, we need to also develop and tune our tools and applications.
- We have lots of activity “pre DC26” to both incorporate improvements and fix issues

**Questions or Comments?**

# 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
-

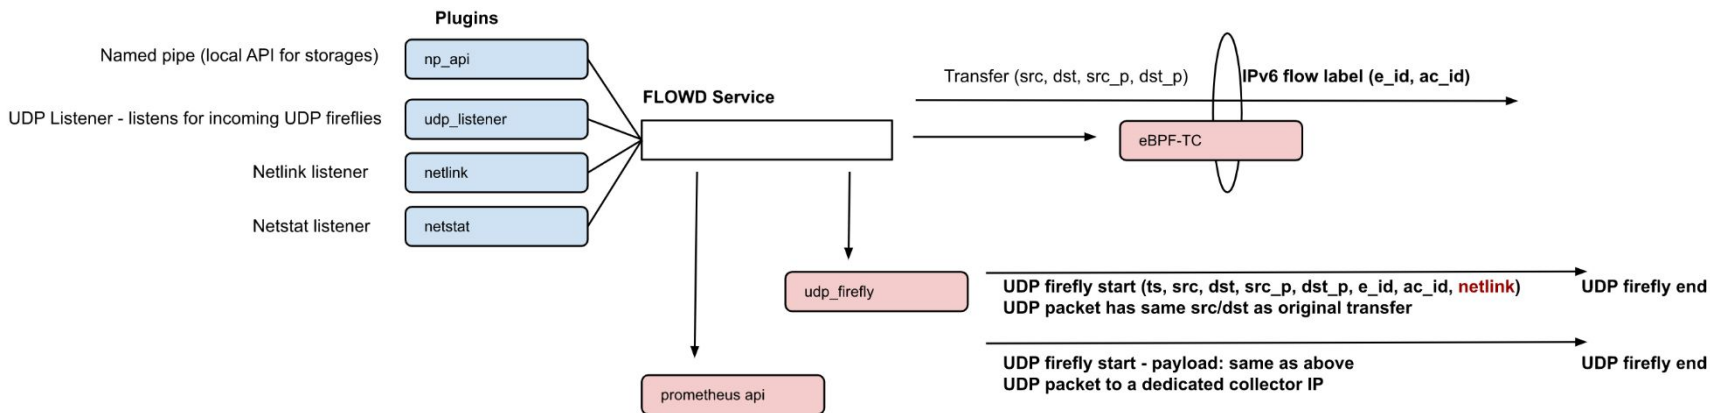
# Relevant URLs for Reference

- OSG/WLCG Networking Documentation
  - <https://opensciencegrid.github.io/networking/>
- WLCG DOMA DC24 plans
  - <https://indico.cern.ch/event/1225415/contributions/5155042/attachments/2593516/4476291/Data%20Challenge%202024.pdf>
- Recent presentations at LHCONE/LHCOPN
  - [perfSONAR Monitoring Update](#)
  - [SciTags](#) deployment, configuration and use
- Toolkit information page
  - <https://toolkitinfo.opensciencegrid.org/>
- ATLAS Alerting and Alarming Service: <https://psa.osg-htc.org/>
- The perfSONAR Dashboard application: <https://ps-dash.uc.ssl-hep.org/>
- ESnet WLCG Stardust Dashboard:  
<https://public.stardust.es.net/d/XkxDL5H7z/esnet-public-dashboards?orgId=1>

# Backup Slides

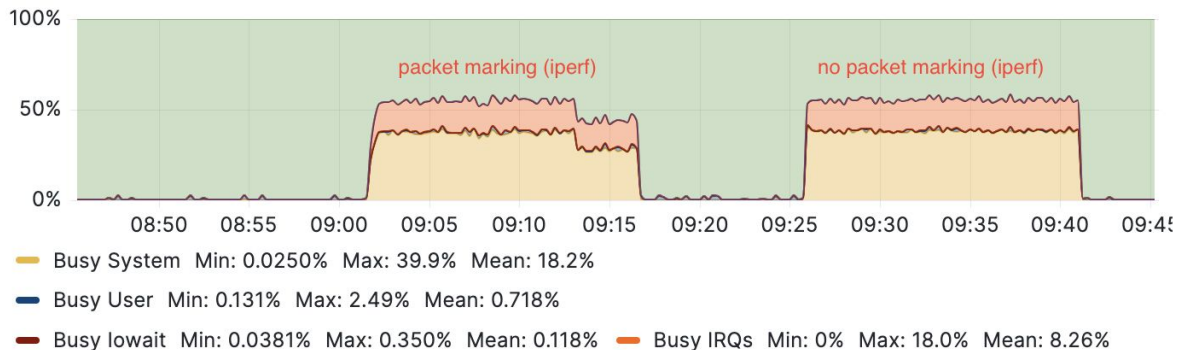
# End-system Utility: Flowd Service

- Architecture (based on event-driven architecture)



- Plugins provide different ways get connections to mark (or interact with storage)
  - New plugins were added to support netlink readout and UDP firefly consumer
- Backends are used to implement flow and/or packet marking
  - New backends were added to mark packets (via eBPF-TC) and expose monitored connection to Prometheus

## CPU Basic ⓘ



## Network Traffic Basic ⓘ

