perfSONAR at your fingertips
Open. Extensible. Worldwide

Szymon Trocha, Poznań Supercomputing and Networking Center, PL

UKNOF48, Manchester, 18-19 November 2021

Public

www.geant.org

Heterogeneous world

- The global Research & Education network ecosystem is comprised of multiple networks
- All interconnect but owned and operated by separate organizations
- This complex, heterogeneous set of networks must operate seamlessly from “end to end”
- To support science and research collaborations that are distributed globally

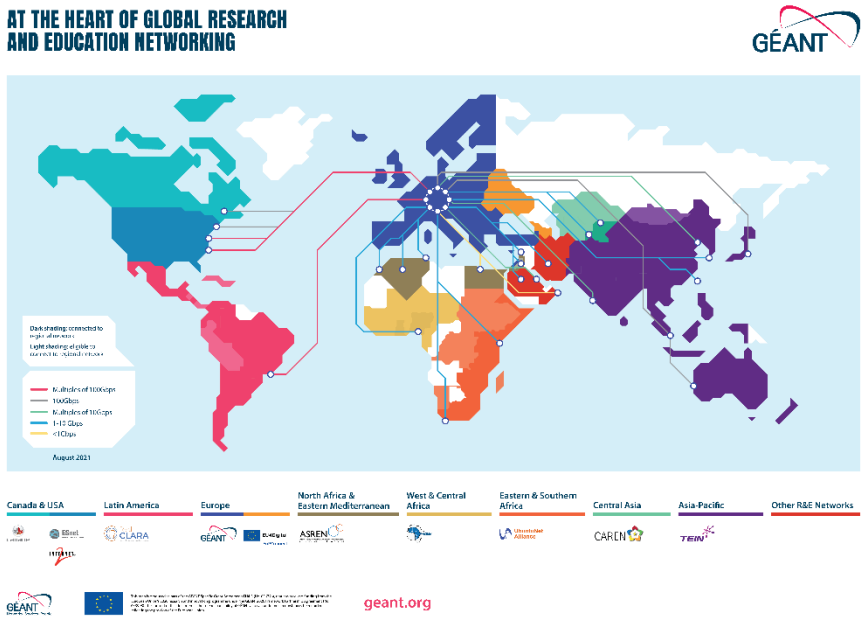

Map source: www.geant.org

What is perfSONAR?

- perfSONAR is a tool to:
  - Set network performance expectations
  - Make optimal use of the network for applications
  - Find network problems (“soft failures”)
  - Help fix these problems
- All in multi-domain environments as problems are all harder when multiple networks are involved
- perfSONAR provides a standard way to publish monitoring data
- Part of the Science DMZ model for supporting efficient data transfers
- This data is interesting to network researchers as well as network operators

perfSONAR Toolkit

- Network performance comes down to a couple of key metrics:
  - Throughput (e.g. “how much can I get out of the network”)
  - Latency (time it takes to get to/from a destination)
  - Packet loss/duplication/order (do they all make it to the other side?)
- But we measure more. And we can get these from a selection of measurement tools – the perfSONAR Toolkit

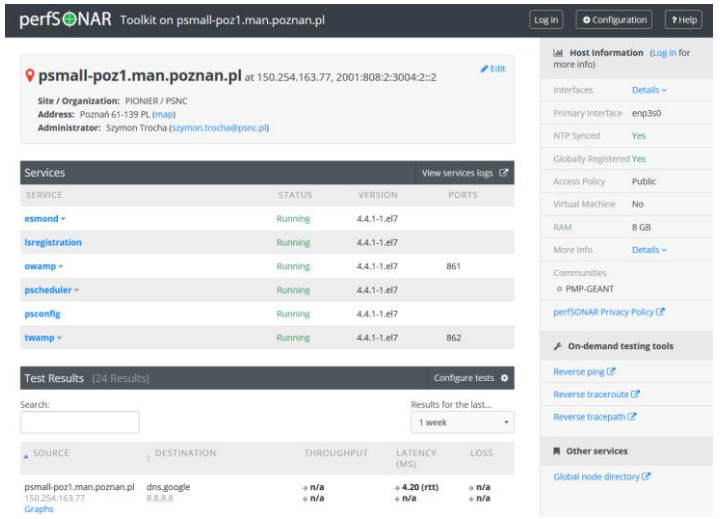

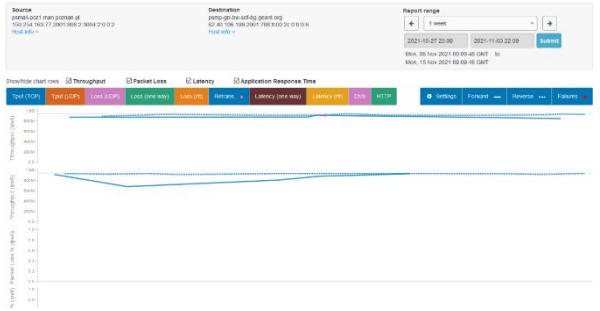

- All components are available as RPMs, DEBs, and CentOS ISO
- Easy to install and configure

- perfSONAR is developed by a partnership of ESnet, Indiana University, Internet2, RNP, University of Michigan
  - And GÉANT community under GN4-3 EU project

Bulding meshes

- Mesh deployment style involves coordinating several nodes
- Nodes can potentially be maintained in different networks
- Nodes share mesh configuration

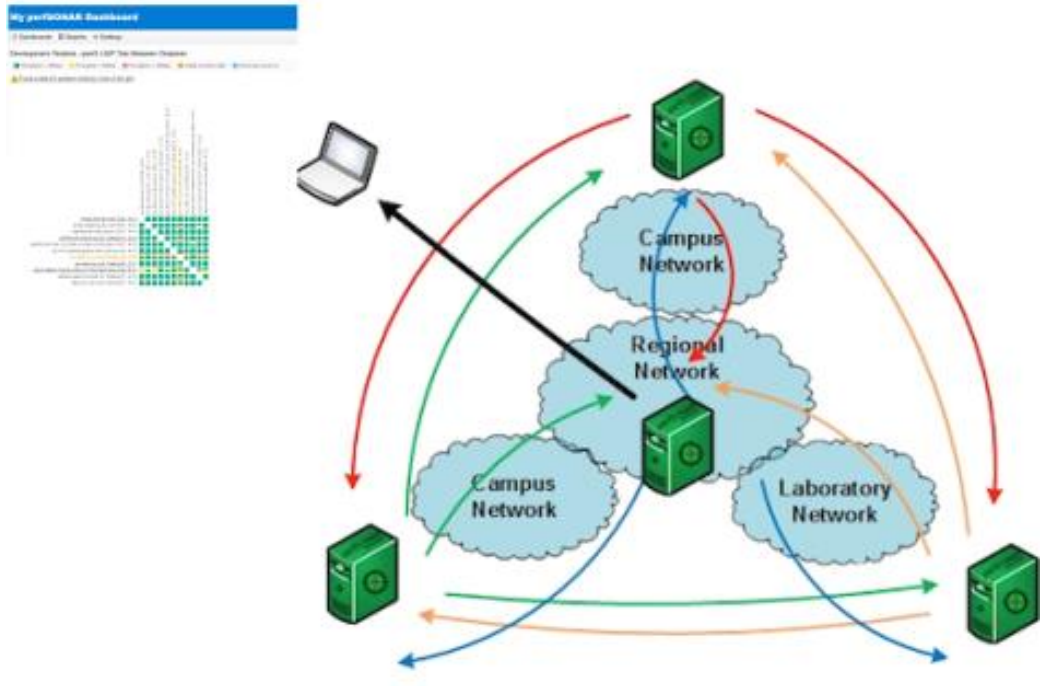

Source: www.perfsonar.net

The importance of regular testing

- We can’t wait for users to report problems and then fix them
- Important to continually collect, archive, and alert on active test results
- perfSONAR includes tools to
  - describe and configure a topology of tasks
  - define and publish configuration of meshes
  - collect and present monitoring data grids

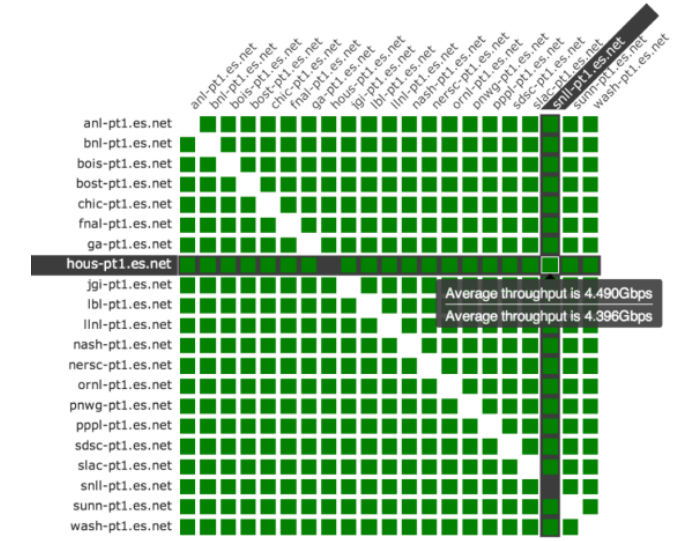

Extensible architecture

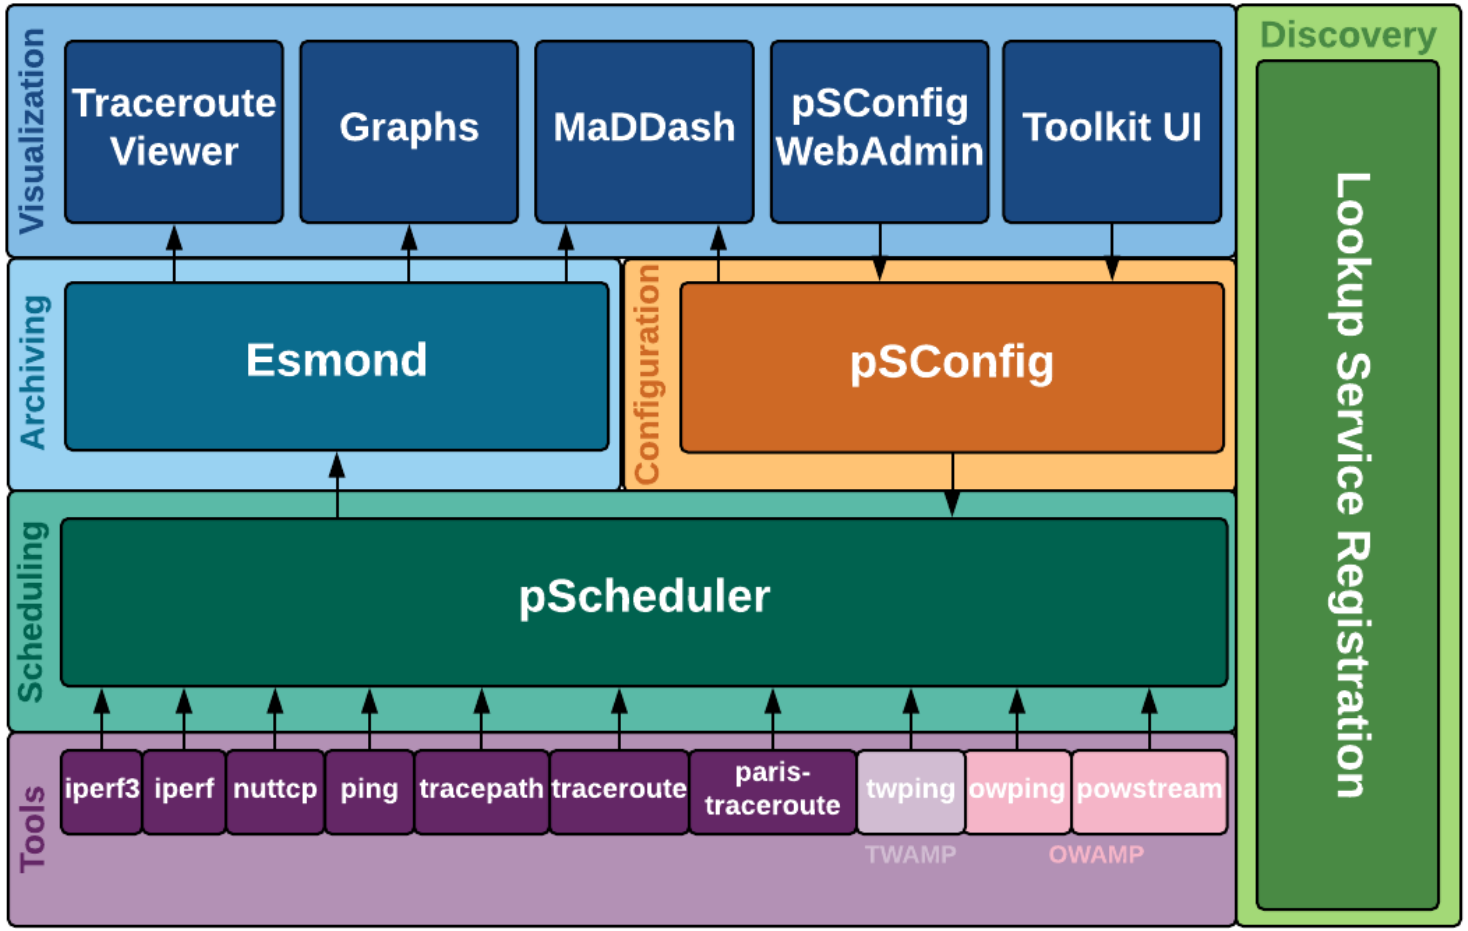

Source: www.perfsonar.net

Plug-in architecture

- Opens community involvement in system's extensions
- pScheduler allows integration of new:
  - Tests (ways to describe measurements)
  - Tools (applications to do the measurements)
  - Archivers (ways to store test results)
  - Contexts (measurement environments)
- Well documented REST API with JSON data format
- Plugin development toolkit (PDK)
  - Supports integration of 3rd party tools with pScheduler API
  - Automates building the environment to develop plugins
  - Reduces time and effort

Worldwide

- ~2000 advertised instances in the world
- A component of NRENs and Virtual Organisations
- Many of which available for open testing

Example use case

- Main actors
  - Queens University, Belfast, UK
  - ATLAS Project (Institute for Astronomy, University of Hawaii, USA)
- Application
  - Astronomy – detecting comets - https://panstarrs.stsci.edu/
  - Large data transfers from experiments / measurements
- Networks involved
  - QUB
  - Janet
  - GÉANT
  - Internet2
  - University of Hawaii

Source: panstarrs.stsci.edu

Thanks to T.Chown, D.Rand from Jisc for their contribution to this part of presentation

Example use case (2)

- Problem
  - Approximately 1 in 8 transfers, which typically ran at 4.8 MB/s from Hawaii, were running very slowly, down to about 50-100 KB/s, which was causing transfers to become backlogged
- Investigation -> use perfSONAR traceroute and loss output to troubleshoot
- Observations
  - Test results show drop in performance
  - Loss plots shows slow, steady increase in loss from 24th October to fix around noon on 29th October

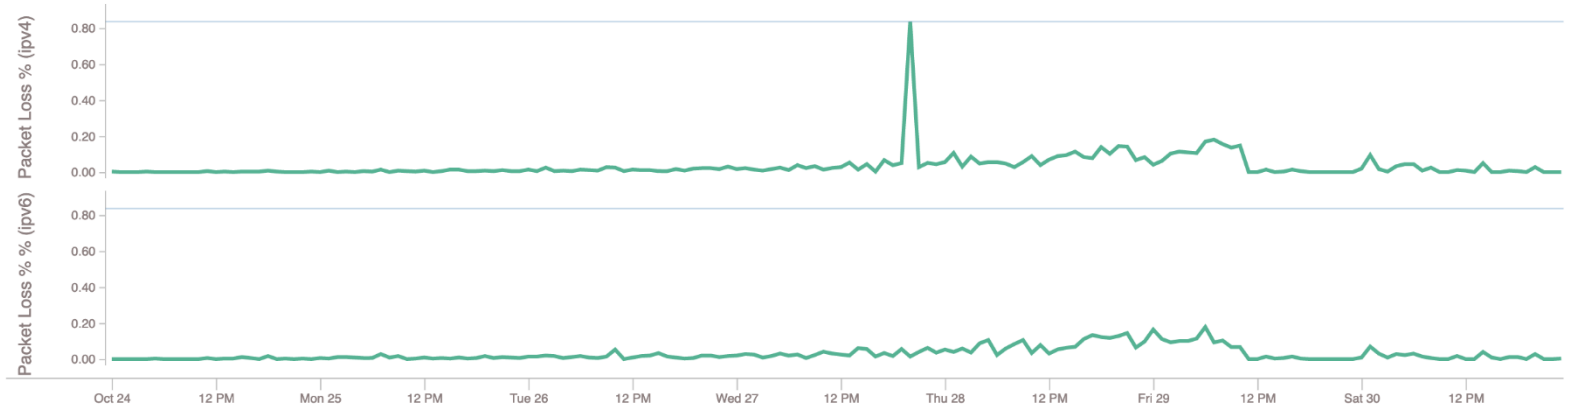

Example use case (3)

- More observations
  - This was an intermittent or ”soft” fault
- Problem found
  - One of eight aggregated 100 Gb/s links between London and Birmingham was faulty
  - (very low) error rate not initially seen by NOC, but enough to affect TCP transfers that were hashed onto that link
  - Faulty optic on one interface needed replacement
- (Interim) solution
  - Taking the faulty link out of the aggregate

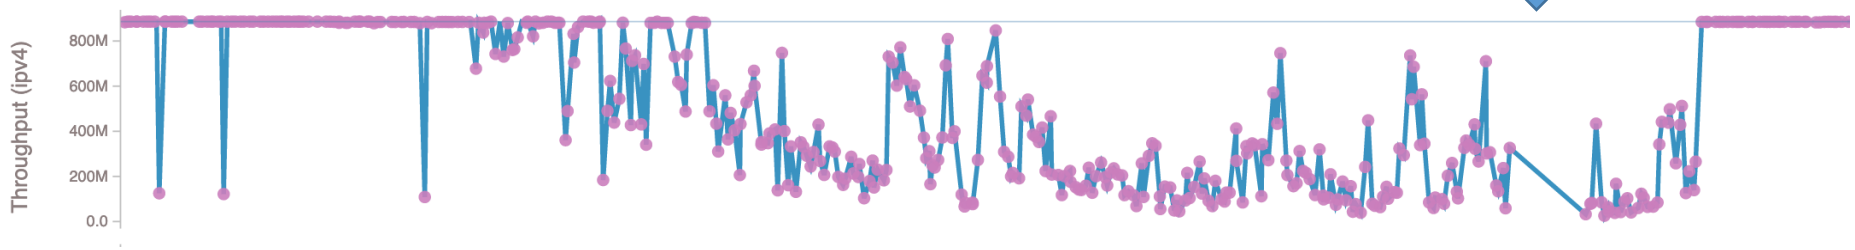

perfSONAR in the future

- Archived data integration
  - An archive is a place where visualization retrieves data
  - We have Esmond but there are many good open source alternatives for storing time-series data
  - Move to OpenSearch (open source derived fom Elasticsearch)
- AI
  - Data anomaly analysis
- OS support change
  - Due to CentOS release strategy change

More info

- www.perfsonar.net
- docs.perfsonar.net
- www.youtube.com/perfSONARProject/

- www.geant.org/Services/Connectivity_and_network/Pages/perfSONAR.aspx
- pmp-central.geant.org/maddash-webui/

Thank you

Any questions?

Szymon.trocha@psnc.pl

www.geant.org

The research leading to these results has received funding from the European Union's Horizon 2020 research and innovation programme under Grant Agreement No. 856726 (GN4-3).

perfSONAR at your fingertips
Open. Extensible. Worldwide

Szymon Trocha, Poznań Supercomputing and Networking Center, PL

UKNOF48, Manchester, 18-19 November 2021

Public

www.geant.org

The scientific work is published for the realization of the international project co-financed by Polish Ministry of Science and Higher Education in the years 2019 - 2022 from financial resources of the programme entitled "PMW"; Agreement No. 5023/H2020/2019/2