

perfSONAR - 20 years monitoring Internet performance

Andrijana Todosijevic / Katarina Simonovic / Antoine Delvaux ▪ AMRES / GÉANT project ▪ perfsonar@lists.geant.org

RIPE 85 - Belgrade

October 2022

perfSONAR is developed by a partnership of



ESnet



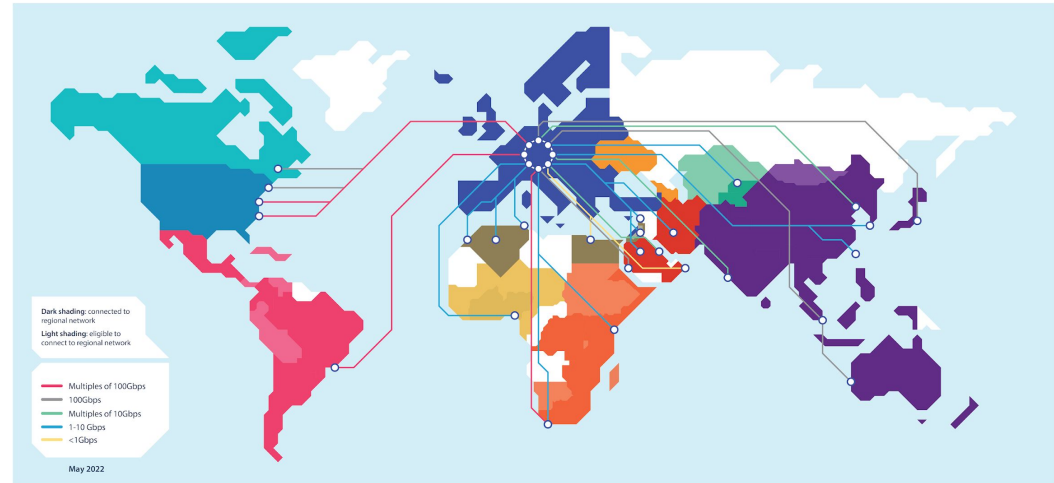
INDIANA UNIVERSITY



GÉANT network – international connectivity

- The global Research & Education network ecosystem
- Must operate seamlessly from “end to end”

AT THE HEART OF GLOBAL RESEARCH AND EDUCATION NETWORKING



This map is produced as part of the GÉANT Service Level Agreement (SLA) (SLA) (SLA), that has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska Curie Grant. The content of this document is the sole responsibility of GÉANT and can neither be considered nor regarded as endorsing the opinion of the European Union.


geant.org



What is perfSONAR?

performance Service-Oriented Network monitoring Architecture:

- Monitoring and assuring network performance
- Make optimal use of the network for apps
- Detect “soft failures” in the network
- Help fix these problems



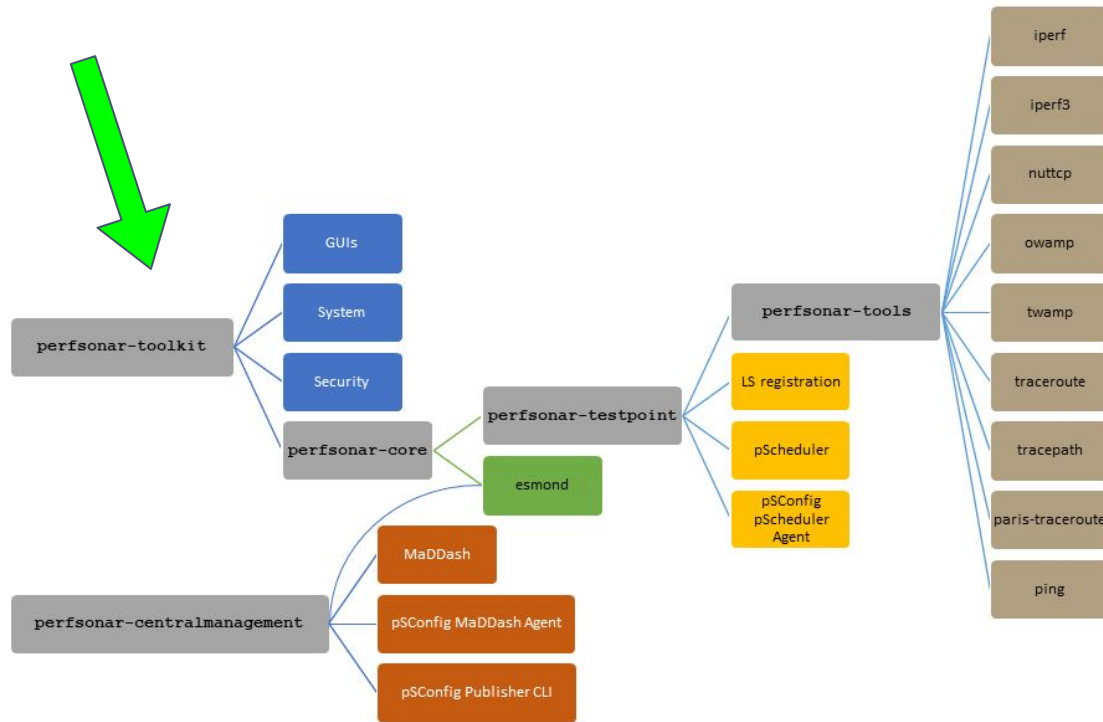
- Testing is done actively, by placing traffic onto the network(s) under test



Partners and collaboration

- Developed for and by the Research & Education community but useful to any network operator
- Current development partners are:
 - ESnet - US
 - GÉANT - Europe
 - Indiana University - US
 - Internet 2 - US
 - RNP - Brazil
 - University Michigan - US

Packages and bundles



- CentOS ISO image
- RPM packages
 - CentOS 7
 - Alma 8
- Deb packages
 - Debian 10
 - Ubuntu 18
 - Ubuntu 20
- Docker container for testpoint
- Ansible playbooks and roles


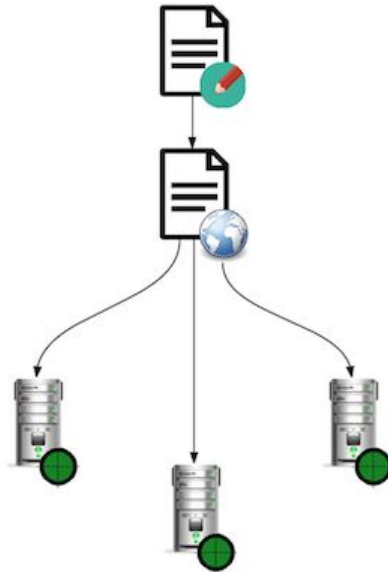
A coordinated measurement mesh

The most complete deployment type

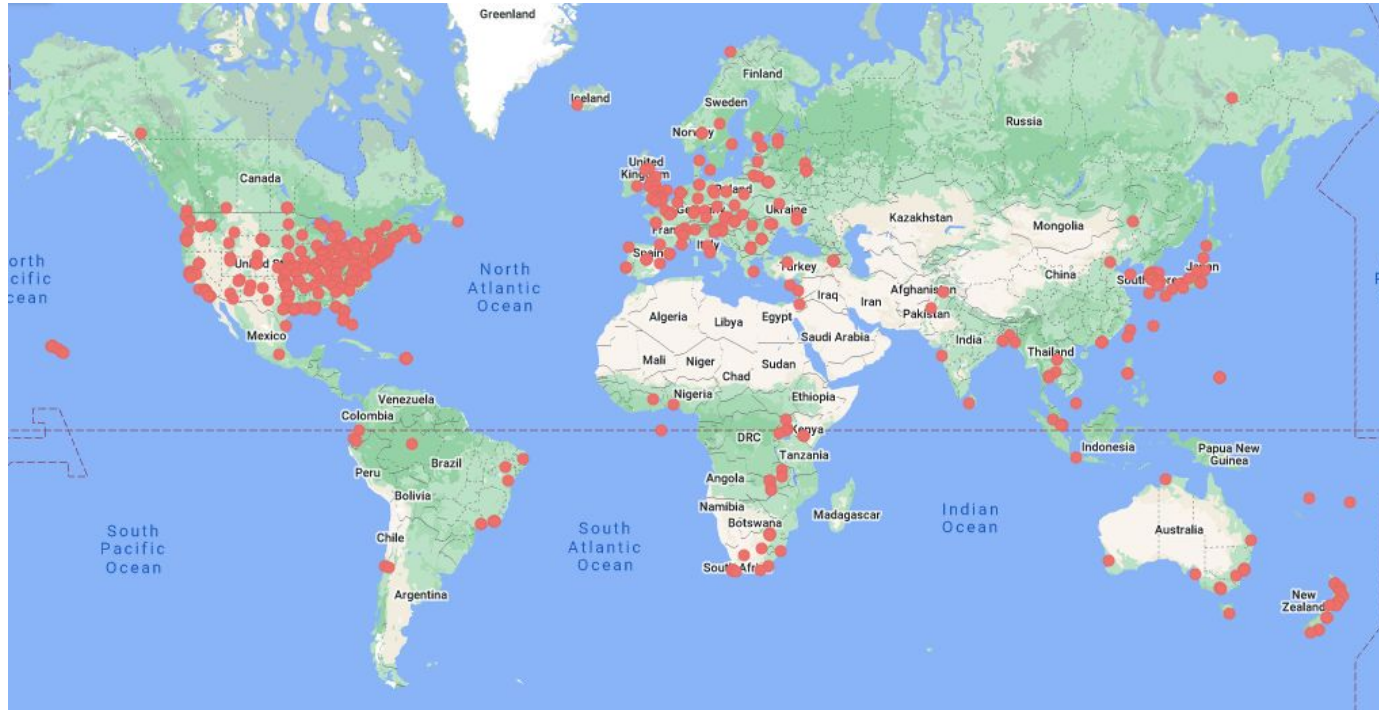
Step 1: Create Central Configuration File

Step 2: Publish to Web

Step 3: perfSONAR Hosts Download Configuration



perfSONAR Lookup Service Directory



<https://stats.es.net>
<https://www.perfsonar.net>
<https://docs.perfsonar.net>



Example Use Case

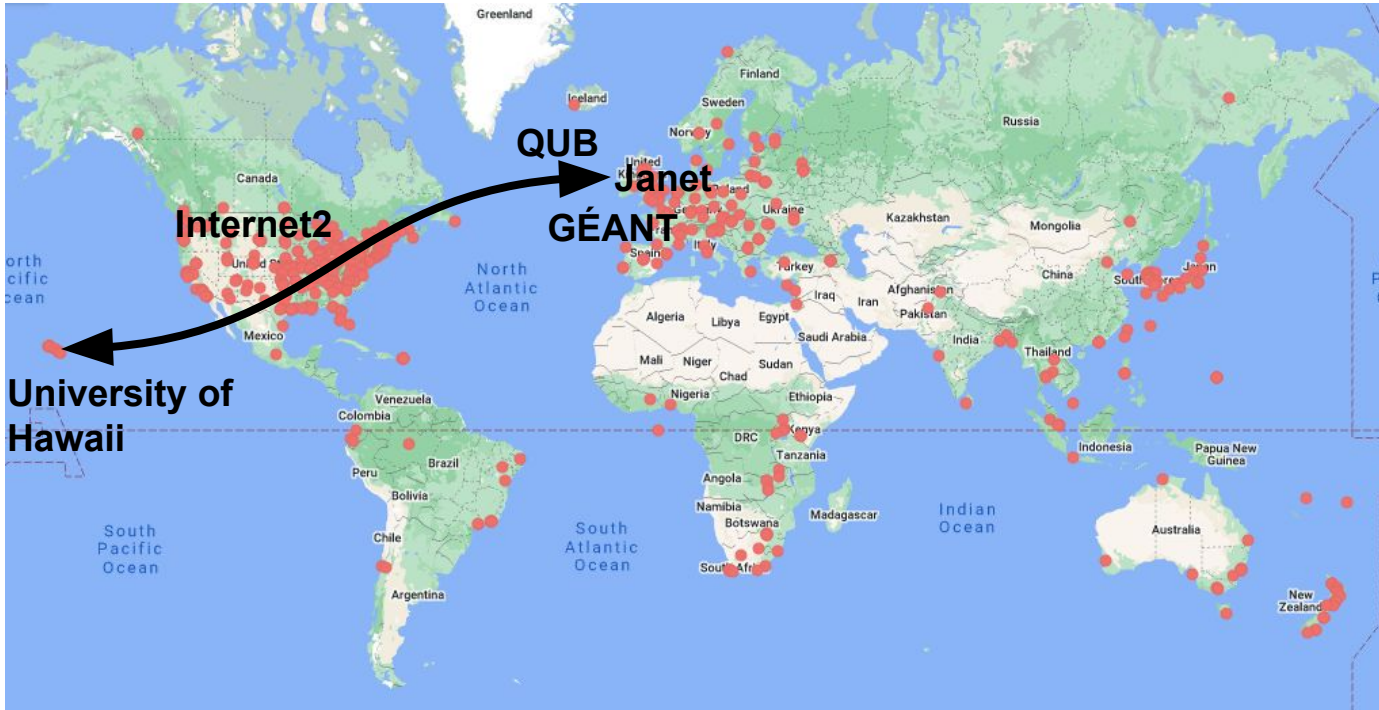
- **Hawaii to Queen's University Belfast**
- Main actors:
 - Queens University Belfast, UK
 - Institute for Astronomy, the University of Hawaii, USA - ATLAS Project (<https://fallingstar.com/>)
- Large data transfers from experiments and measurements



<https://outerspace.stsci.edu/display/PANSTARRS/>



Hawaii to Belfast (11.000 km, 180ms RTT)



Example Use Case 1

Problem:


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

- Troubleshoot: perfSONAR traceroute and loss output
- Observations:
 - Tests results showed drop in performance
 - Slow, steady increase in loss within a couple of days



Example Use Case 1


- 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
 - Faulty optic on one interface needed replacement
- (Interim) solution –Taking the faulty link out of the aggregate



Example Use Case 2

- **Max (Baltimore) to Jisc (Slough)**

Not just Hawaii to Belfast, also present between completely different perfSONAR hosts




Example Use Case 2

UH IfA - 10G Bandwidth Tests - Throughput




✓ No problems found in grid

Average throughput from Institute for Astronomy (IfA) Hawaii to Jisc host in Slough is 3.3 Gbps



Example Use Case 2

- ITC (Hawaii) to Jisc (Slough)



Example Use Case 2

- **GÉANT (London) to Jisc (Didcot)**

Also noticed continuous packet loss



Example Use Case 2

- **Max (Baltimore) to Jisc (Slough)**

Diagnosis: faulty optic in 6 x 100G aggregate between Janet London Harbour Exchange and London Powergate

Janet NOC removed the faulty link from the aggregate and drops in throughput disappear

- **GÉANT (London) to Jisc (Didcot)**

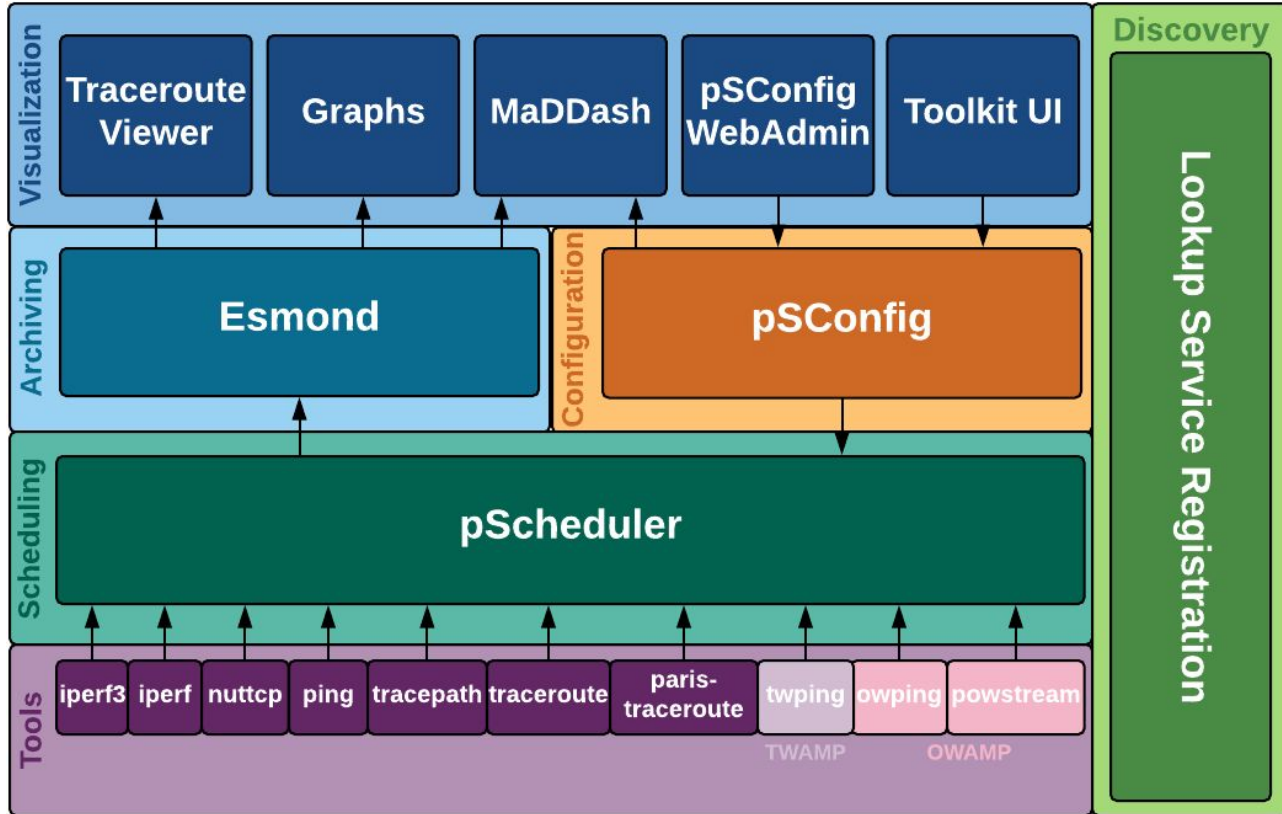
Janet NOC removed the faulty link from the aggregate and packet loss disappears



ESnet




Architecture

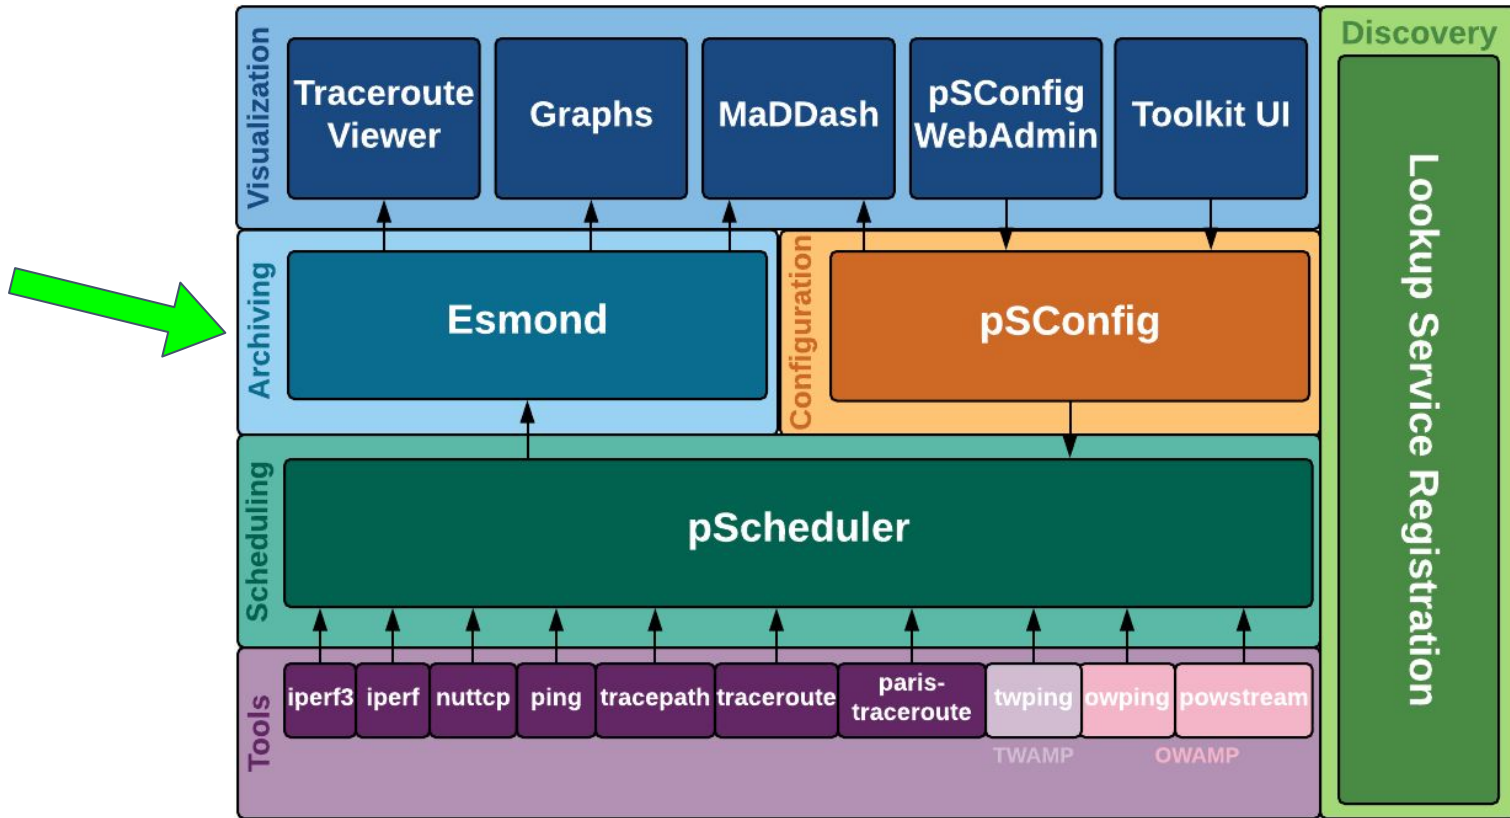


What is Esmond?

- Default archive that most users run
- Django app with custom REST API
- Use two backend databases
 - PostgreSQL
 - Cassandra



Architecture change for 5.0



Elasticsearch, Logstash, Kibana (ELK) and Grafana

- **Elasticsearch** - Stores and indexes documents and lets you do searches
- **Logstash** - Accepts input from lots of different sources, enriches with location data and more, can output it to different places (like Elasticsearch)
- **Kibana** - Visualizes data in elasticsearch
- **Grafana*** - Visualization platform for ElasticSearch and more



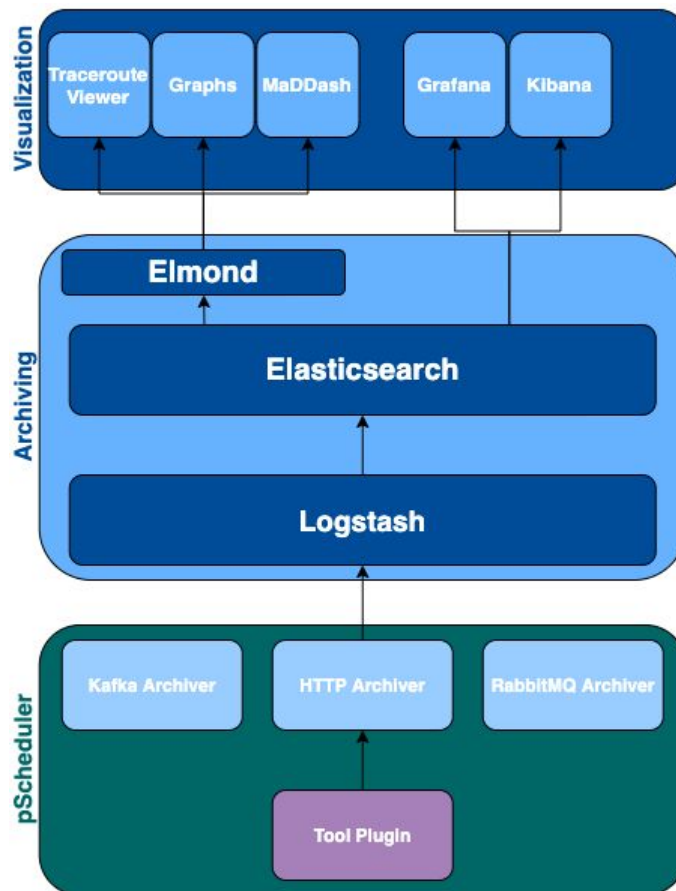
Moving to OpenSearch

- Because of licensing changes
- Default perfSONAR bundles will rely on OpenSearch

- Will maintain compatibility with Elasticsearch for those with existing installations



The Software Pieces



Other features of 5.0

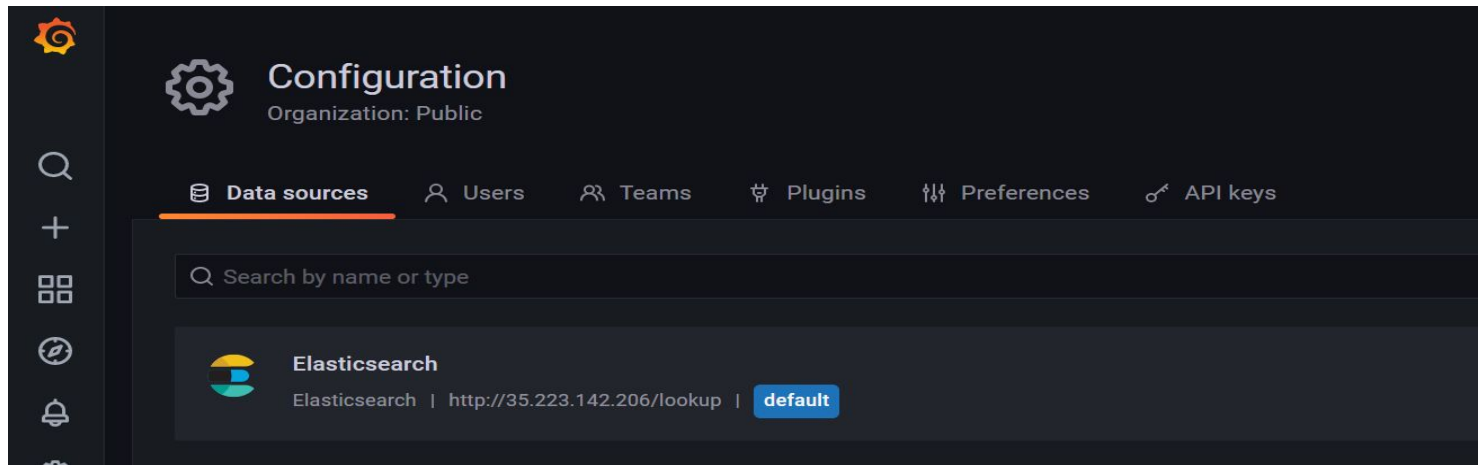
- Various pScheduler improvements and some new plugins
- Some Toolkit UI improvements
- pSConfig Web Admin (PWA) changes
- Optional packages if you want
 - to keep Esmond
 - to use Kibana



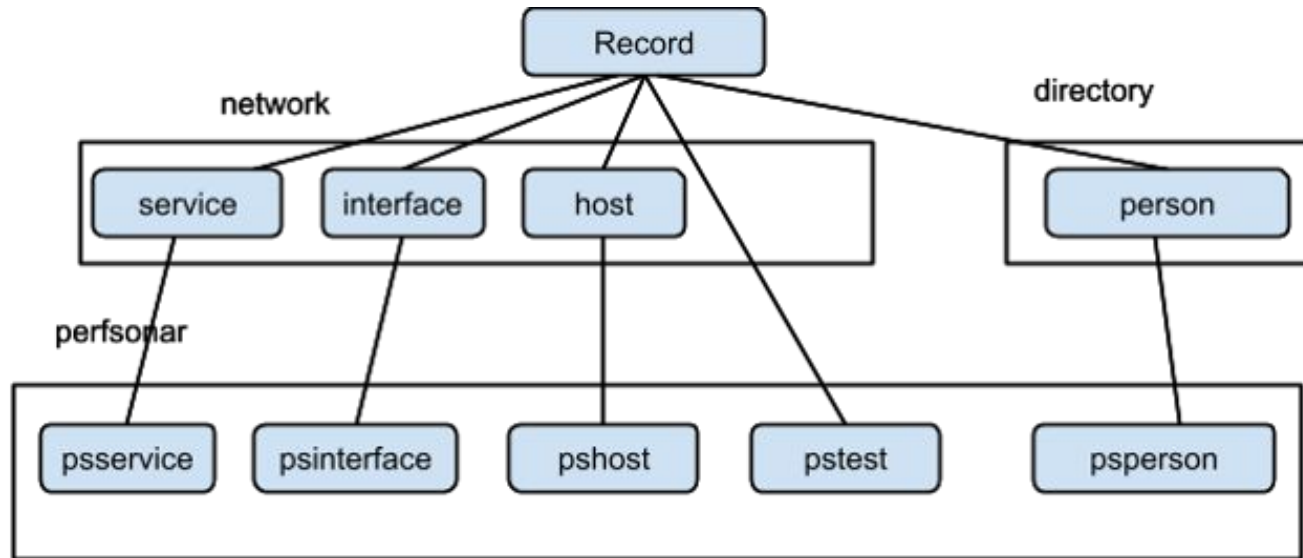
perfSONAR Elasticsearch API

GET/POST http://35.223.142.206/lookup/_search

```
{  
  query....  
}
```



perfSONAR Lookup Service records





Search Enter variable value Domain ja.net Communities pSmall-GEANT + PMP-GEANT + GEANT

- > Service Information (6 panels)
- > Host Information (9 panels)
- > Service Map (1 panel)

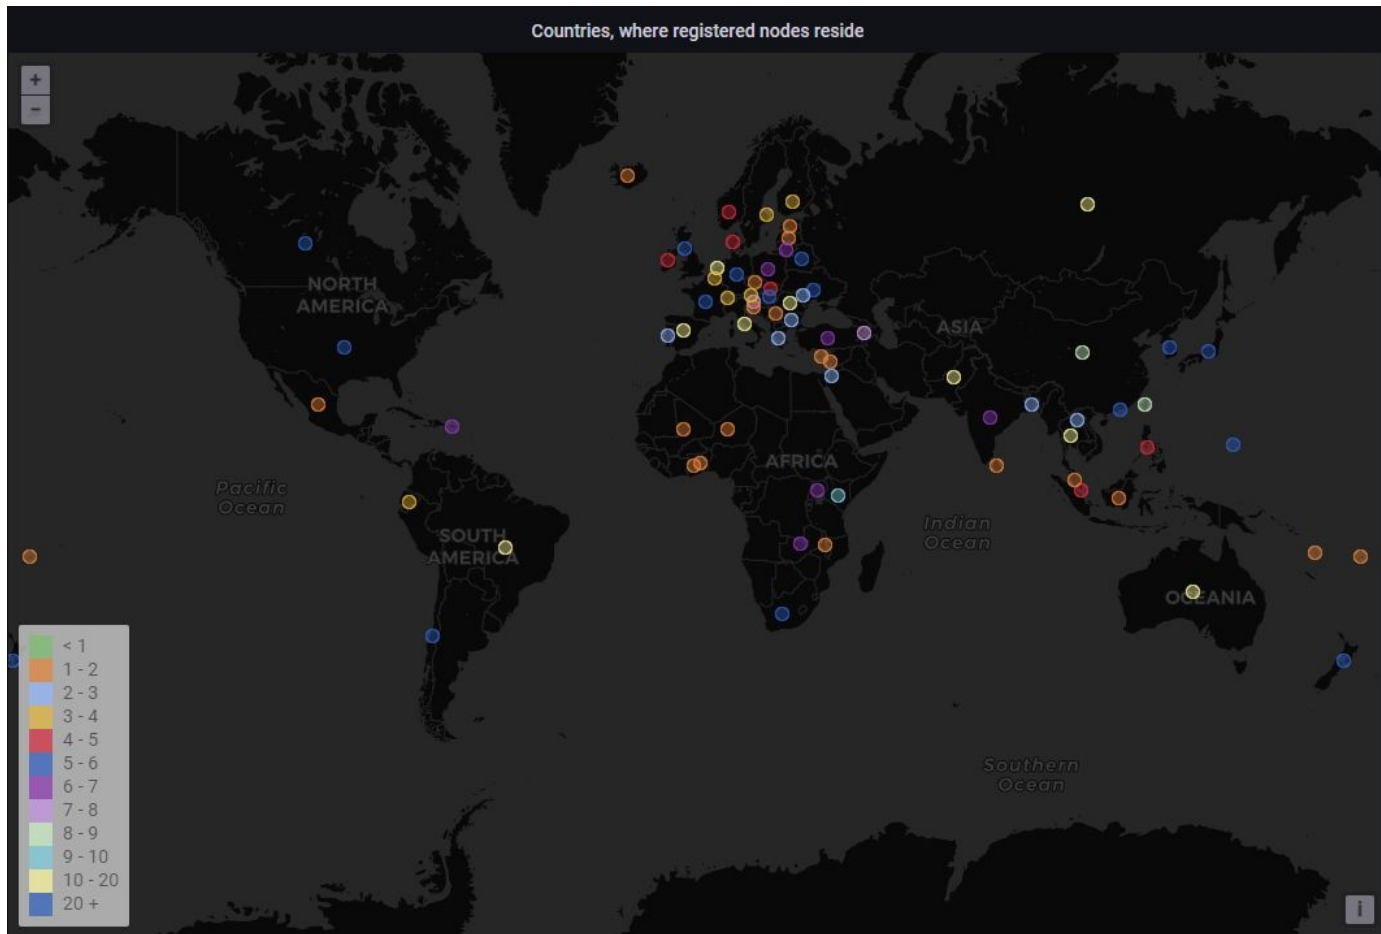
General

perfSONAR hosts perfSONAR services perfSONAR interfaces

2177 8537 2399



<https://stats.perfsonar.net>
<https://www.perfsonar.net>
<https://docs.perfsonar.net>



<https://stats.perfsonar.net>

<https://www.perfsonar.net>

<https://docs.perfsonar.net>



General / perfSONAR Public ☆ 🔗

Search

Domain

ja.net ▾

Communities

> Service Information (6 panels)

> Host Information (9 panels)

▾ Service Map

Map Ser



host-net-tcp-maxbuffer-recv:

host-net-tcp-maxbuffer-send:

host-os-architecture:

host-os-kernel:

host-os-name:

host-os-version:

host-productname:

host-vm:

location-city: Slough

location-country: GB

latitude: 51.4964

longitude: -0.1224

location-sitename: ps-small-slough

pscheduler-tests: clock,disk-to-disk,dns,http,idle,idlebgm,idleex,latency,latencybg,rtt,s3throughput,simplestreamcurl,dnsperf,iperf2,iperf3,nuttcp,owping,paris-traceroute,ping,powstream,psclbenchmark,simplestreamer,sleep,sleepbgm,snooze,tracepath,traceroute,twping

pscheduler-tools:

pshost-access-policy:

pshost-bundle:

pshost-bundle-version:

pshost-install-method:

pshost-role:

pshost-toolkitversion:



perfSONAR



Thanks icon by priyanka from The Noun Project

Thanks!

For more information,
please visit our web site:

<https://www.perfsonar.net>

perfSONAR is developed by a partnership of



ESnet



INDIANA UNIVERSITY



RNP

ORGANIZAÇÃO SOCIAL DO MCTI



UNIVERSITY OF MICHIGAN