THE
perfSONAR
POWER HOUR
C L
PULL-ON/VOL
V-HOLD
COLOR
TINT
BRIGHT
PICTURE

WITH

ANDY LAKE - ESNET

SARAH LARSEN - ESNET

MARK FEIT - INTERNET2

LÆTITIA DELUAUX - PSNC

UP NEXT
ANDY
LAKE

perfSONAR

The perfSONAR Power Hour

Andy Lake – ESnet – andy@es.net

perfSONAR is developed by a partnership of ESnet GÉANT INDIANA UNIVERSITY INTERNET2 RNP UNIVERSITY OF MICHIGAN

What is perfSONAR?

- An **open source software collaboration** led by ESnet, GEANT, Indiana University, Internet2, RNP and the University of Michigan.

- **Goal is to provide network measurements between organizations** to help identify and troubleshoot network issues. Most commonly these include (but are not limited to):
  - Throughput
  - Packet Loss
  - One-way latency
  - Traceroute

perf5.0NAR

- perfSONAR 5.0 released **April 17th, 2023.**

- **Over 50% of perfSONAR deployments** currently running 5.0

- **Enables greater visualization and analysis capabilities** through the replacement of the backend measurement storage database with OpenSearch

- **New pScheduler test plugins** to support WiFi BSSID, 802.1X authentication, DHCP response time and more

- **Ubuntu 20 support added** with additional OSes like
  - EL8 and EL9 added in early summer
  - Ubuntu 22 and Debian 11 coming soon

- **Looking ahead, 5.1 will focus on improving UI** and leverage the changes put in place by 5.0 to add new capabilities

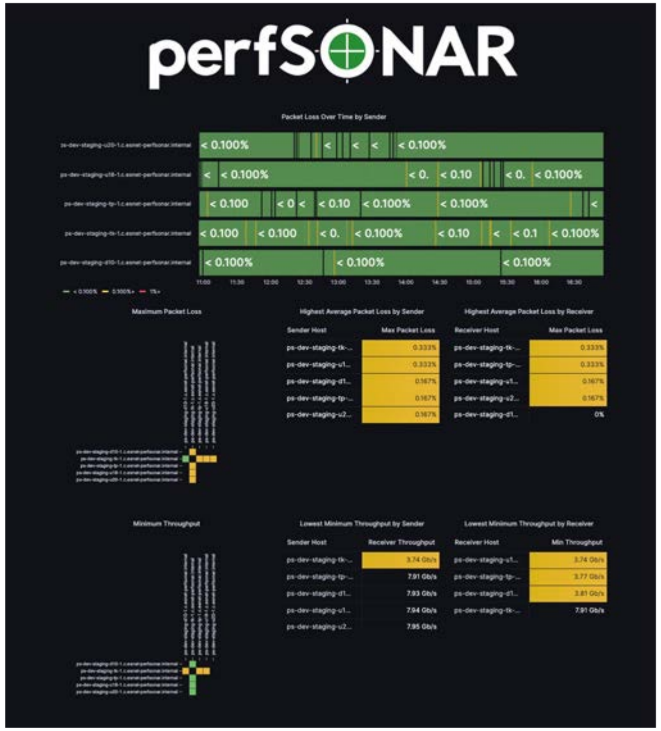

Highest Average Packet Loss by Sender

Sender Host	Max Packet Loss
ps-dev-staging-tk-...	0.333%
ps-dev-staging-u1...	0.333%
ps-dev-staging-d1...	0.167%
ps-dev-staging-tp-...	0.167%
ps-dev-staging-u2...	0.167%

Highest Average Packet Loss by Receiver

Receiver Host	Max Packet Loss
ps-dev-staging-tk-...	0.333%
ps-dev-staging-tp-...	0.333%
ps-dev-staging-u1...	0.167%
ps-dev-staging-u2...	0.167%
ps-dev-staging-d1...	0%

Lowest Minimum Throughput by Sender

Sender Host	Receiver Throughput
ps-dev-staging-tk-...	3.74 Gb/s
ps-dev-staging-tp-...	7.91 Gb/s
ps-dev-staging-d1...	7.93 Gb/s
ps-dev-staging-u1...	7.94 Gb/s
ps-dev-staging-u2...	7.95 Gb/s

Lowest Minimum Throughput by Receiver

Receiver Host	Min Throughput
ps-dev-staging-u1...	3.74 Gb/s
ps-dev-staging-tp-...	3.77 Gb/s
ps-dev-staging-d1...	3.81 Gb/s
ps-dev-staging-tk-...	7.91 Gb/s

Image: Example of Grafana dashboard users can setup in 5.0 using our guide

ESnet GÉANT INDIANA UNIVERSITY INTERNET2 RNP UNIVERSITY OF MICHIGAN

Building on the Foundation

UI: Toolkit UI and MaDDash Beyond 5.0

UI: On-Demand Testing

Tool Enhancements: Multi-threaded iperf3

Deployment Models: perfSONAR on Internet2 Backbone

Toolkit UI and MaDDash Beyond 5.0

Andy Lake – ESnet – andy@es.net

perfSONAR is developed by a partnership of

Today’s Toolkit UI

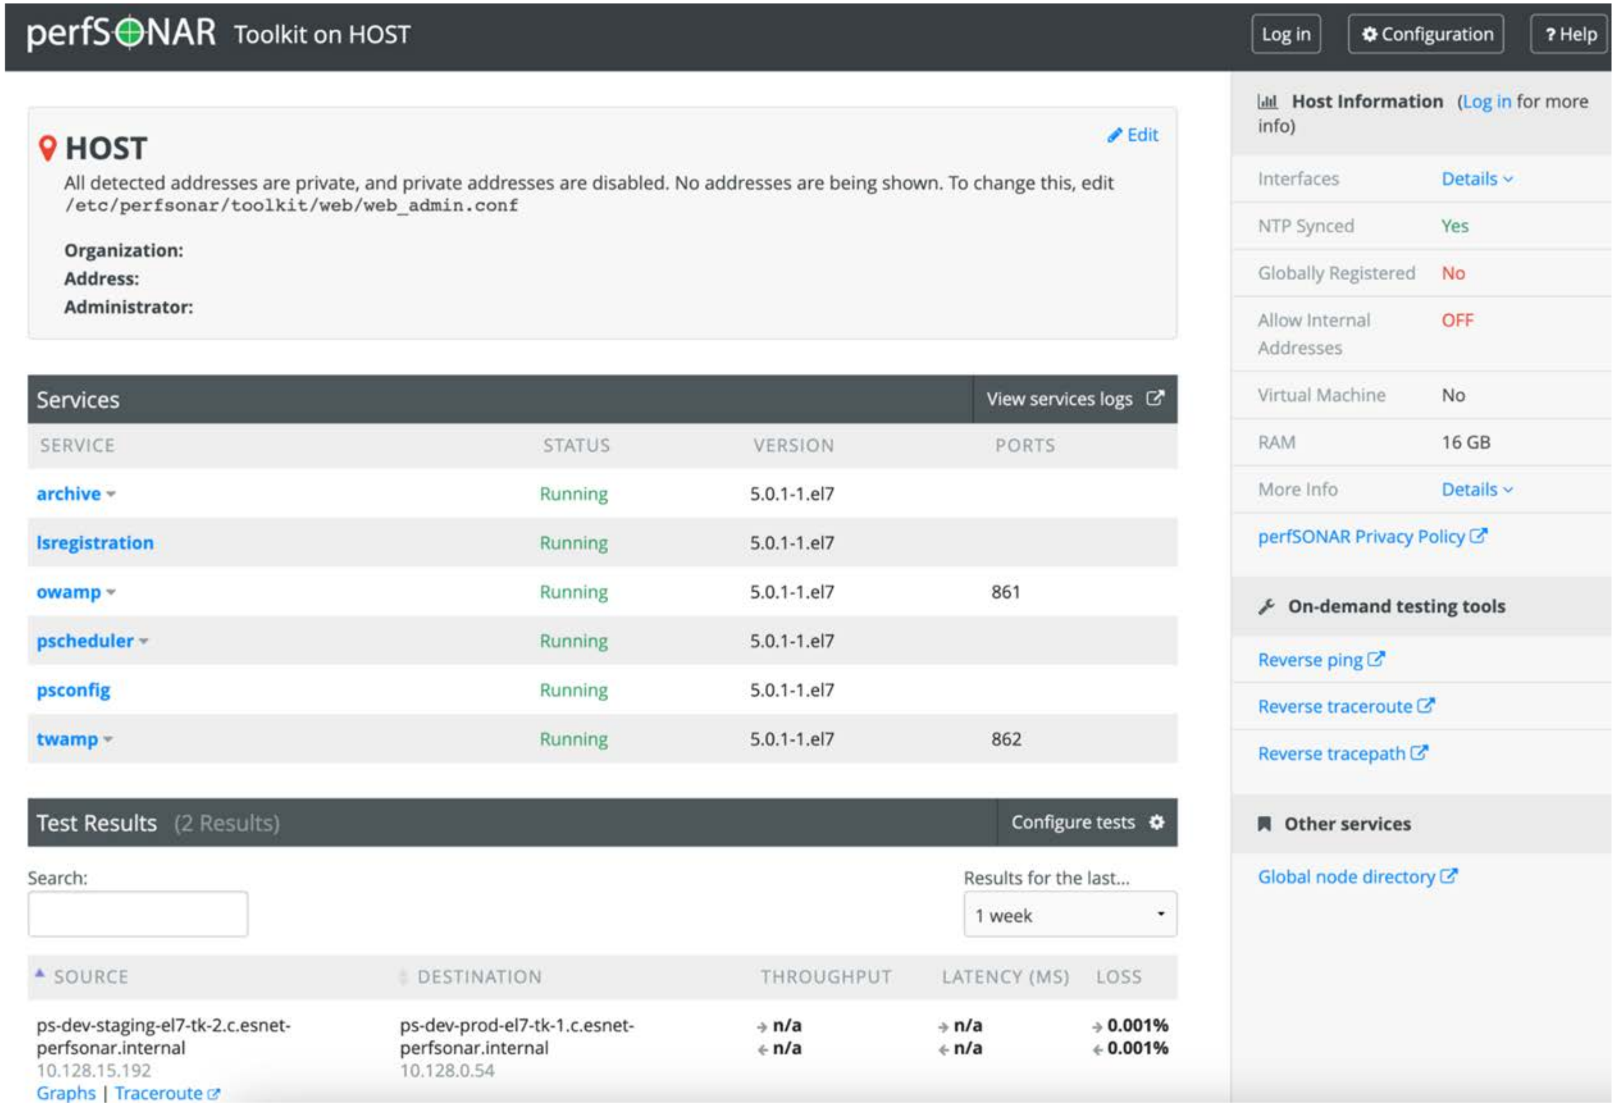

HOST

All detected addresses are private, and private addresses are disabled. No addresses are being shown. To change this, edit `/etc/perfsonar/toolkit/web/web_admin.conf`

Organization:
Address:
Administrator:

Services

SERVICE	STATUS	VERSION	PORTS
archive	Running	5.0.1-1.el7	
lsregistration	Running	5.0.1-1.el7	
owamp	Running	5.0.1-1.el7	861
pscheduler	Running	5.0.1-1.el7	
psconfig	Running	5.0.1-1.el7	
twamp	Running	5.0.1-1.el7	862

Test Results (2 Results)

Search:

Results for the last... 1 week

SOURCE	DESTINATION	THROUGHPUT	LATENCY (MS)	LOSS
ps-dev-staging-el7-tk-2.c.esnet-perfsonar.internal 10.128.15.192 Graphs | Traceroute	ps-dev-prod-el7-tk-1.c.esnet-perfsonar.internal 10.128.0.54	→ n/a ← n/a	→ n/a ← n/a	→ 0.001% ← 0.001%

Host Information (Log in for more info)

Interfaces	Details
NTP Synced	Yes
Globally Registered	No
Allow Internal Addresses	OFF
Virtual Machine	No
RAM	16 GB
More Info	Details

perfSONAR Privacy Policy

On-demand testing tools

Reverse ping
Reverse traceroute
Reverse tracepath

Other services

Global node directory

Today’s MaDDash UI

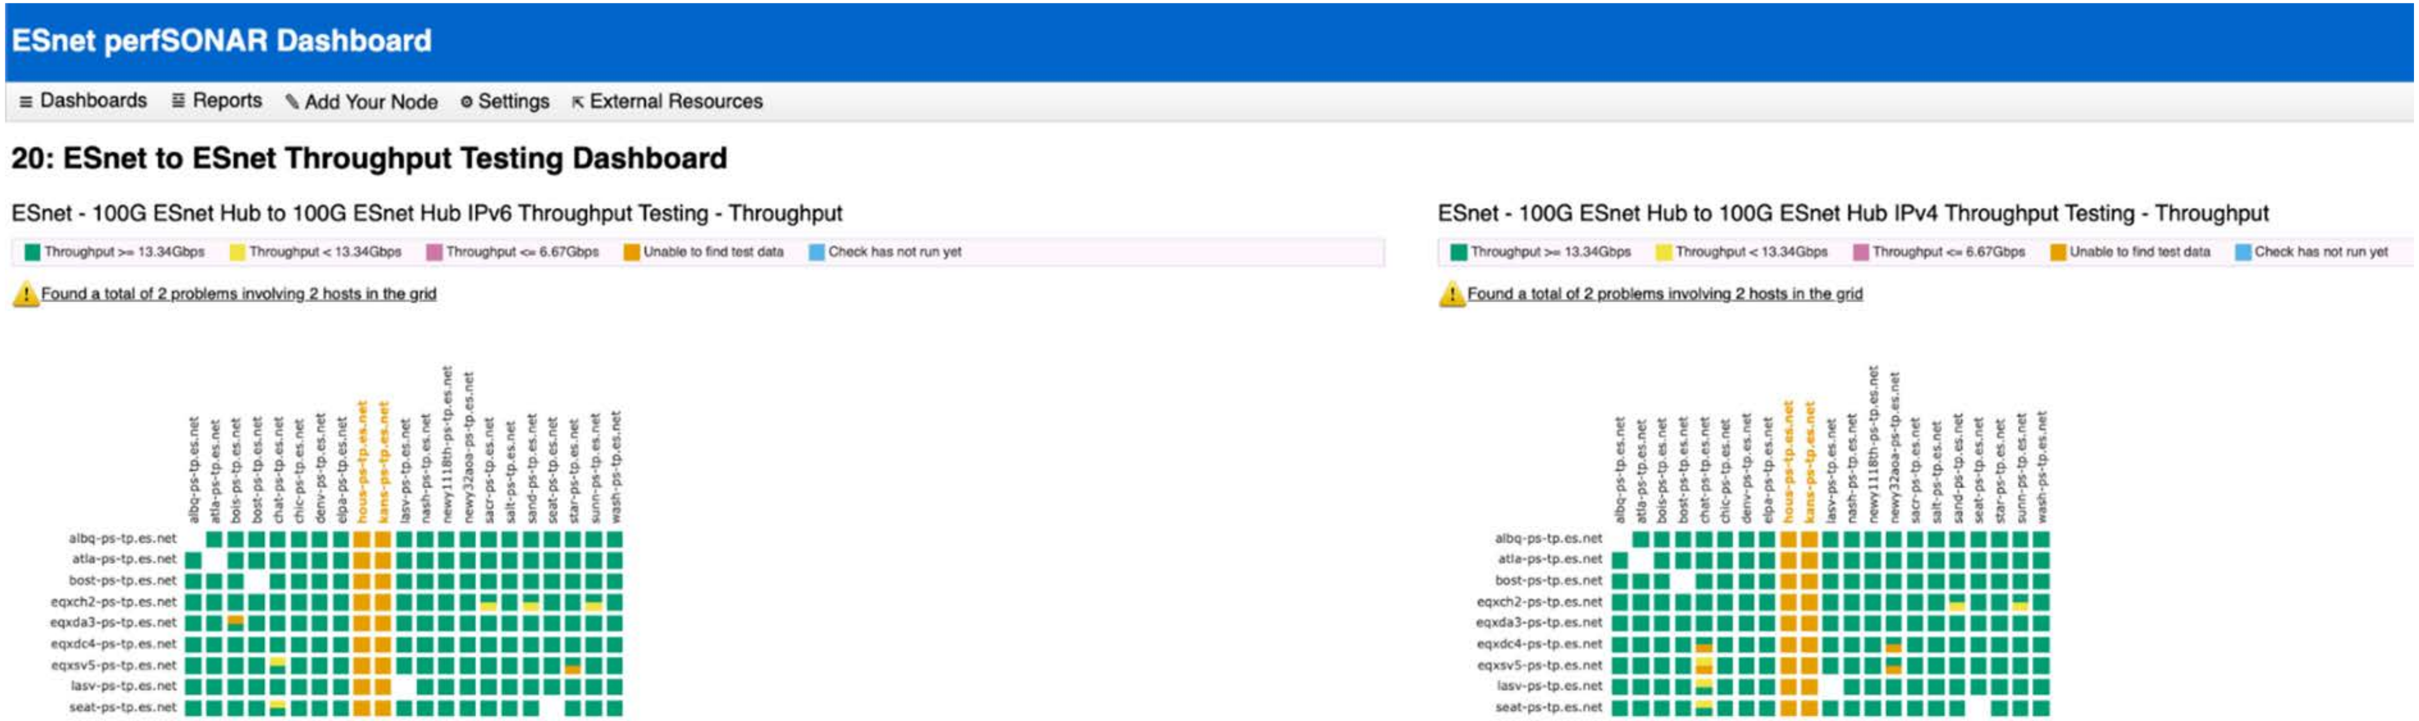

ESnet perfSONAR Dashboard

20: ESnet to ESnet Throughput Testing Dashboard

ESnet - 100G ESnet Hub to 100G ESnet Hub IPv6 Throughput Testing - Throughput

🟩 Throughput >= 13.34Gbps 🟨 Throughput < 13.34Gbps 🟪 Throughput <= 6.67Gbps 🟧 Unable to find test data 🟦 Check has not run yet

⚠ Found a total of 2 problems involving 2 hosts in the grid

	albq-ps-tp.es.net	atla-ps-tp.es.net	bois-ps-tp.es.net	bost-ps-tp.es.net	chat-ps-tp.es.net	chic-ps-tp.es.net	denv-ps-tp.es.net	elpa-ps-tp.es.net	hous-ps-tp.es.net	kans-ps-tp.es.net	lasv-ps-tp.es.net	nash-ps-tp.es.net	newy1118th-ps-tp.es.net	newy32aoa-ps-tp.es.net	sacr-ps-tp.es.net	salt-ps-tp.es.net	sand-ps-tp.es.net	seat-ps-tp.es.net	star-ps-tp.es.net	sunn-ps-tp.es.net	wash-ps-tp.es.net
albq-ps-tp.es.net		🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟧	🟧	🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟩
atla-ps-tp.es.net	🟩		🟩	🟩	🟩	🟩	🟩	🟩	🟧	🟧	🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟩
bost-ps-tp.es.net	🟩	🟩	🟩		🟩	🟩	🟩	🟩	🟧	🟧	🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟩
eqxch2-ps-tp.es.net	🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟧	🟧	🟩	🟩	🟩	🟩	🟩🟨	🟩	🟩🟨	🟩	🟩	🟩🟨	🟩
eqxda3-ps-tp.es.net	🟩	🟩	🟧🟩	🟩	🟩	🟩	🟩	🟩	🟧	🟧	🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟩
eqxdc4-ps-tp.es.net	🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟧	🟧	🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟩
eqxsv5-ps-tp.es.net	🟩	🟩	🟩	🟩	🟨🟩	🟩	🟩	🟩	🟧	🟧	🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟩🟧	🟩	🟩
lasv-ps-tp.es.net	🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟧	🟧		🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟩
seat-ps-tp.es.net	🟩	🟩	🟩	🟩	🟨🟩	🟩	🟩	🟩	🟧	🟧	🟩	🟩	🟩	🟩	🟩	🟩	🟩		🟩	🟩	🟩

ESnet - 100G ESnet Hub to 100G ESnet Hub IPv4 Throughput Testing - Throughput

🟩 Throughput >= 13.34Gbps 🟨 Throughput < 13.34Gbps 🟪 Throughput <= 6.67Gbps 🟧 Unable to find test data 🟦 Check has not run yet

⚠ Found a total of 2 problems involving 2 hosts in the grid

	albq-ps-tp.es.net	atla-ps-tp.es.net	bois-ps-tp.es.net	bost-ps-tp.es.net	chat-ps-tp.es.net	chic-ps-tp.es.net	denv-ps-tp.es.net	elpa-ps-tp.es.net	hous-ps-tp.es.net	kans-ps-tp.es.net	lasv-ps-tp.es.net	nash-ps-tp.es.net	newy1118th-ps-tp.es.net	newy32aoa-ps-tp.es.net	sacr-ps-tp.es.net	salt-ps-tp.es.net	sand-ps-tp.es.net	seat-ps-tp.es.net	star-ps-tp.es.net	sunn-ps-tp.es.net	wash-ps-tp.es.net
albq-ps-tp.es.net		🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟧	🟧	🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟩
atla-ps-tp.es.net	🟩		🟩	🟩	🟩	🟩	🟩	🟩	🟧	🟧	🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟩
bost-ps-tp.es.net	🟩	🟩	🟩		🟩	🟩	🟩	🟩	🟧	🟧	🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟩
eqxch2-ps-tp.es.net	🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟧	🟧	🟩	🟩	🟩	🟩	🟩	🟩	🟩🟨	🟩	🟩	🟩🟨	🟩
eqxda3-ps-tp.es.net	🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟧	🟧	🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟩
eqxdc4-ps-tp.es.net	🟩	🟩	🟩	🟩	🟩🟧	🟩	🟩	🟩	🟧	🟧	🟩	🟩	🟩	🟩🟧	🟩	🟩	🟩	🟩	🟩	🟩	🟩
eqxsv5-ps-tp.es.net	🟩	🟩	🟩	🟩	🟨🟧	🟩	🟩	🟩	🟧	🟧	🟩	🟩	🟩	🟩🟧	🟩	🟩	🟩	🟩	🟩	🟩	🟩
lasv-ps-tp.es.net	🟩	🟩	🟩	🟩	🟨🟩	🟩	🟩	🟩	🟧	🟧		🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟩
seat-ps-tp.es.net	🟩	🟩	🟩	🟩	🟨🟩	🟩	🟩	🟩	🟧	🟧	🟩	🟩	🟩	🟩	🟩	🟩	🟩		🟩	🟩	🟩

Thinking about a new interface

- User Requirements
  - Users need to still have ability to view test results of local tests
  - Authentication required to see additional system stats, configure tests, and configure LS registration
    - Whole subset of requirements around configuring test that won't cover in this deck
  - System information needs to be available via an API (see current JSON service)
- Dev Requirements
  - Devs need to get rid of old perl cgi scripts that are increasingly difficult to maintain
  - Ditto old Javascript
  - Consolidate redundant interfaces (e.g. PWA and Toolkit Test Config)
  - Stop discovering system stats that differ across OSes, when other things have this solved (eg. node_exporter)
- Other nice to haves
  - Users have greater flexibility in building own views of data
  - Tighter integration with MaDDash and main toolkit UI
  - We don't have a lot of UI devs, if there is a way to make easier to maintain without being JS expert, that is good thing.

What is Grafana?

- Grafana is an **open-source platform** for for exploring data from a variety of sources
- It has a few key features
  - It's **multi-data source**
  - It has a bunch of **built-in visualizations** that **don't require you to be a Javascript developer** to use
  - It has a **plugin framework** for all of the above so they can be extended and a process for becoming official plugins
- More Info: https://grafana.com/docs/grafana/latest/introduction/

Custom Plugins by NetSage and ESnet

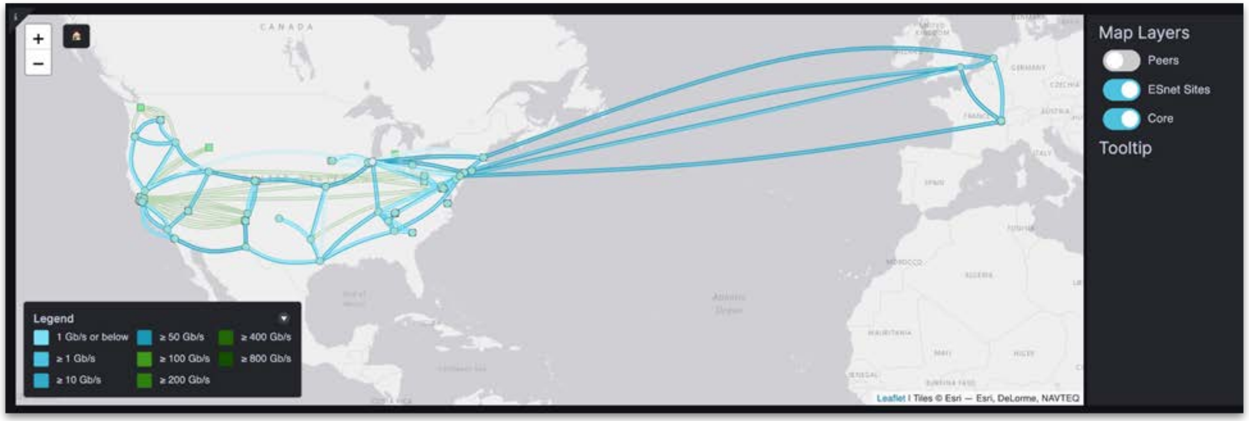

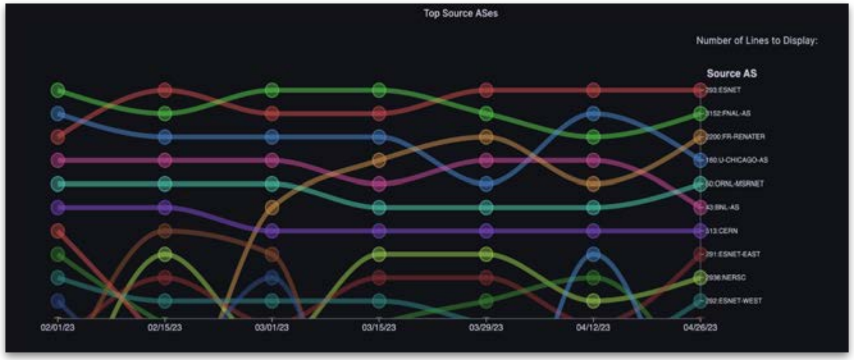

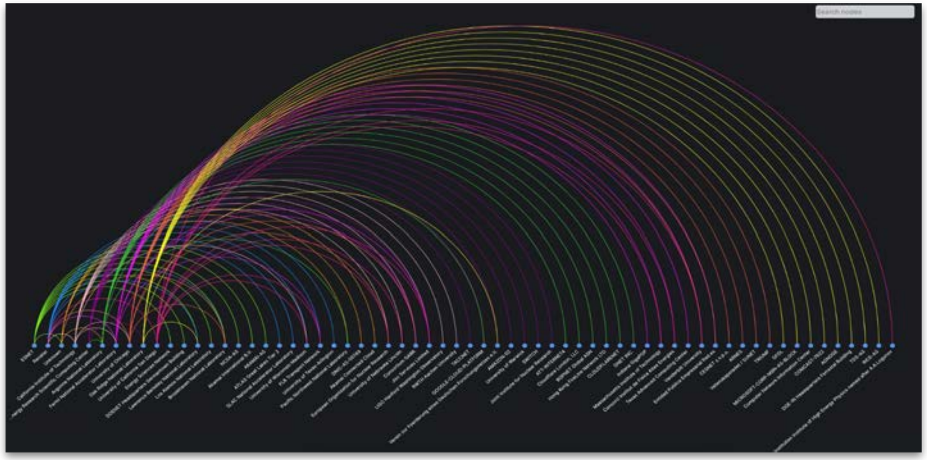

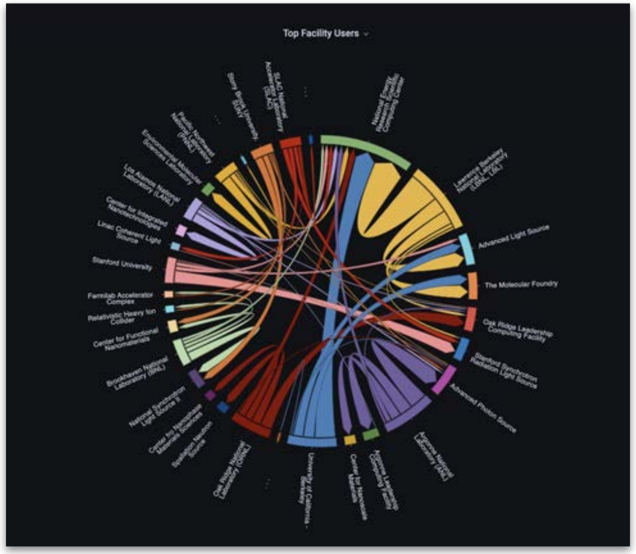

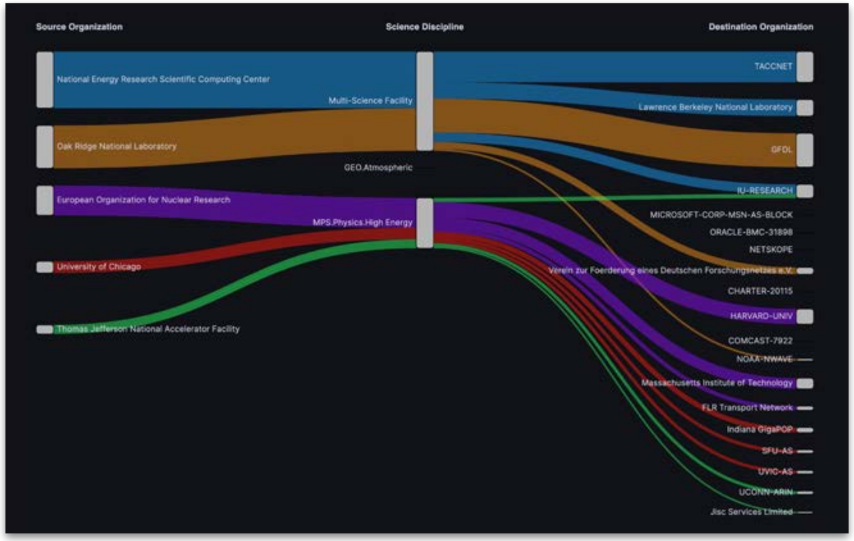

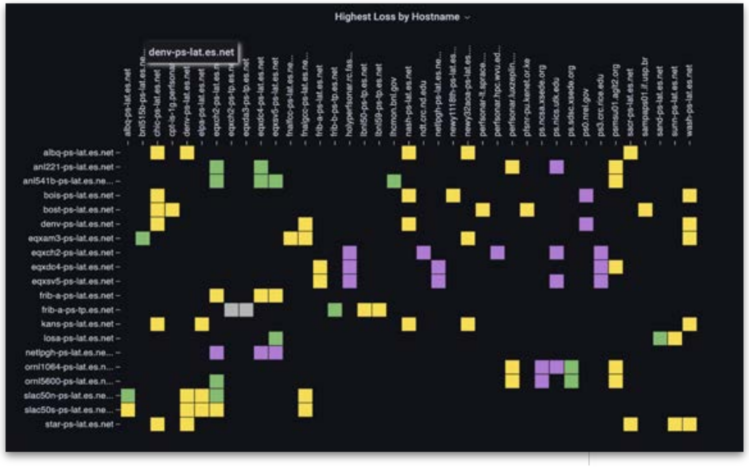

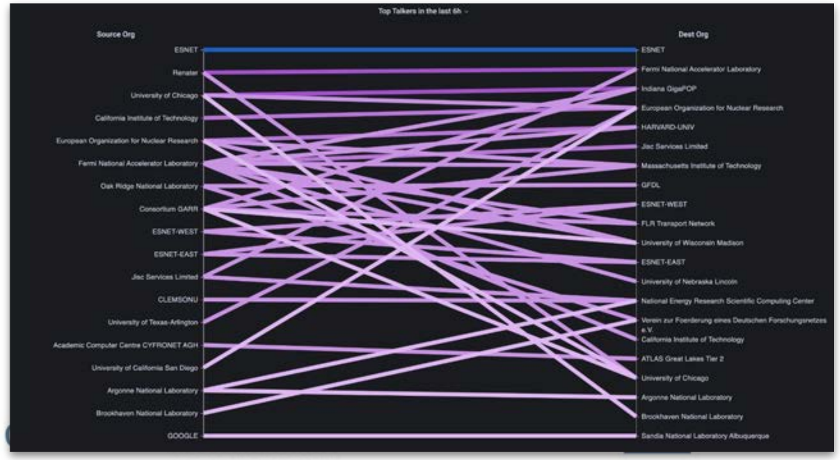

New UI: Focus on Measurements

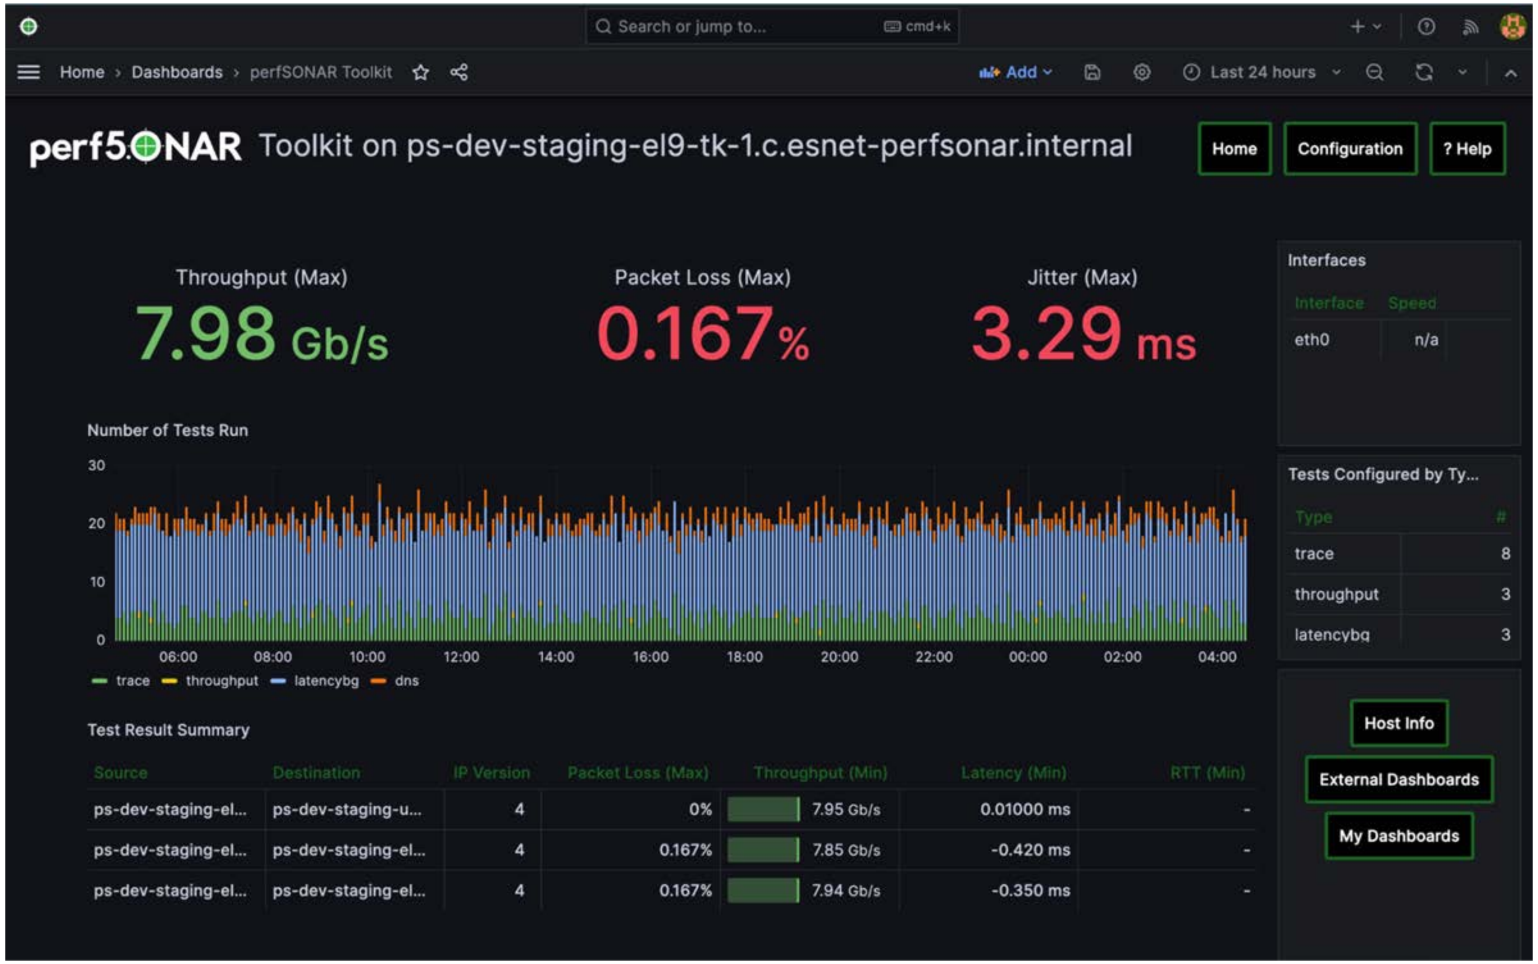

perfSONAR Toolkit on ps-dev-staging-el9-tk-1.c.esnet-perfsonar.internal

Throughput (Max) 7.98 Gb/s

Packet Loss (Max) 0.167%

Jitter (Max) 3.29 ms

Number of Tests Run

Interfaces

Interface	Speed
eth0	n/a

Tests Configured by Ty...

Type	#
trace	8
throughput	3
latencybg	3

Test Result Summary

Source	Destination	IP Version	Packet Loss (Max)	Throughput (Min)	Latency (Min)	RTT (Min)
ps-dev-staging-el...	ps-dev-staging-u...	4	0%	7.95 Gb/s	0.01000 ms	-
ps-dev-staging-el...	ps-dev-staging-el...	4	0.167%	7.85 Gb/s	-0.420 ms	-
ps-dev-staging-el...	ps-dev-staging-el...	4	0.167%	7.94 Gb/s	-0.350 ms	-

New UI: Enhancing the Fundamentals

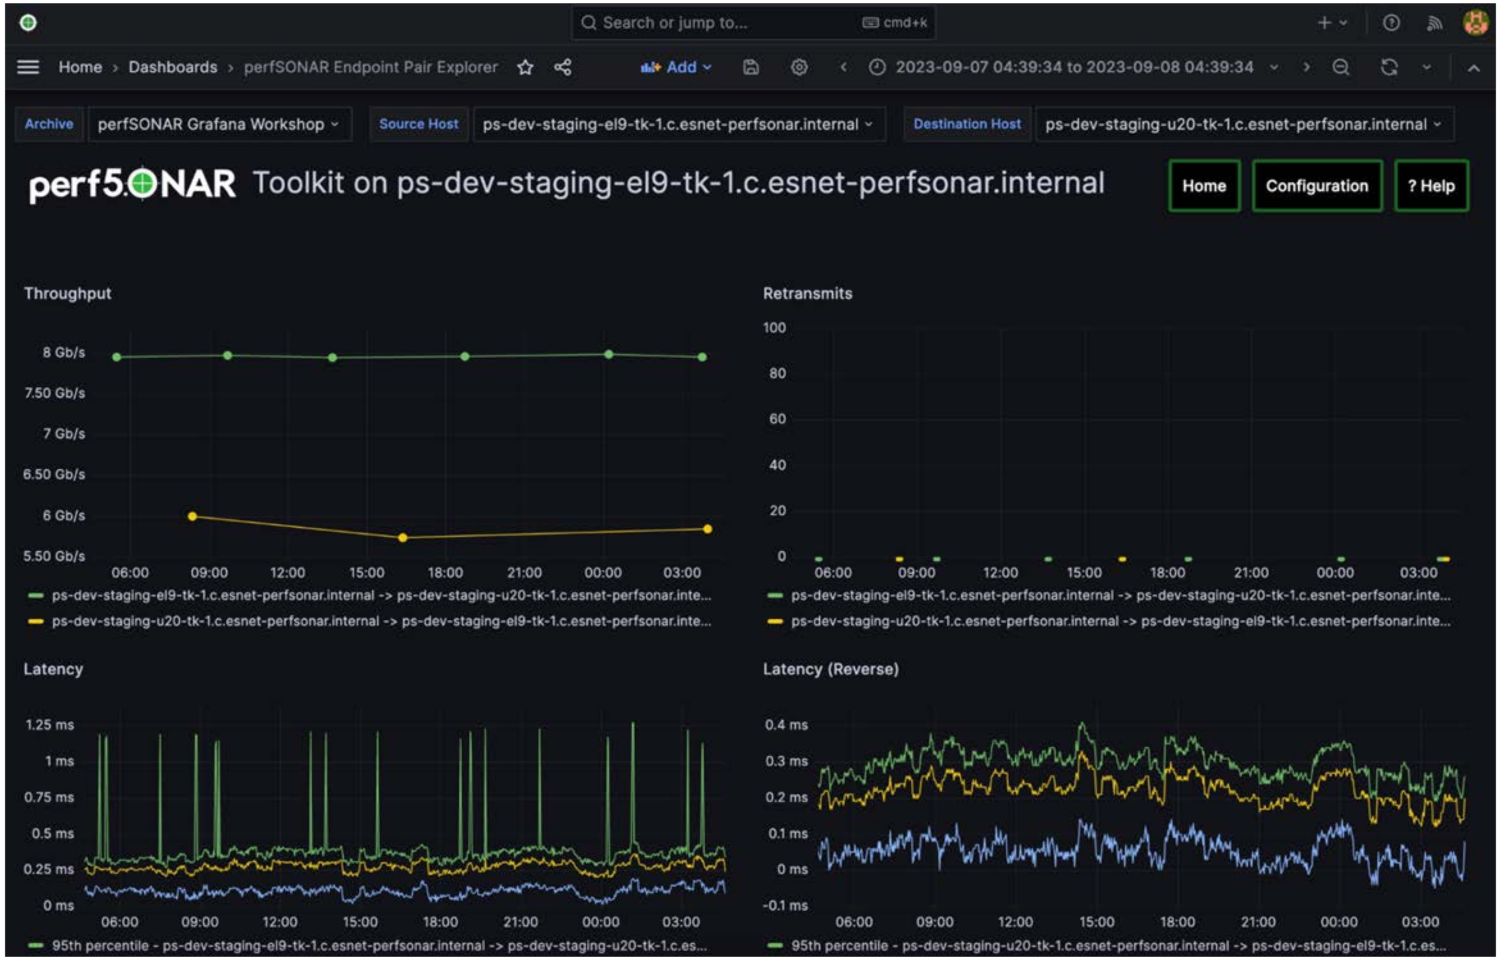

Search or jump to...
cmd+k
Home › Dashboards › perfSONAR Endpoint Pair Explorer
Add
2023-09-07 04:39:34 to 2023-09-08 04:39:34

Archive: perfSONAR Grafana Workshop

Source Host: ps-dev-staging-el9-tk-1.c.esnet-perfsonar.internal

Destination Host: ps-dev-staging-u20-tk-1.c.esnet-perfsonar.internal

perfSONAR Toolkit on ps-dev-staging-el9-tk-1.c.esnet-perfsonar.internal

Home
Configuration
? Help
Throughput
8 Gb/s
7.50 Gb/s
7 Gb/s
6.50 Gb/s
6 Gb/s
5.50 Gb/s
06:00
09:00
12:00
15:00
18:00
21:00
00:00
03:00
ps-dev-staging-el9-tk-1.c.esnet-perfsonar.internal -> ps-dev-staging-u20-tk-1.c.esnet-perfsonar.inte...
ps-dev-staging-u20-tk-1.c.esnet-perfsonar.internal -> ps-dev-staging-el9-tk-1.c.esnet-perfsonar.inte...
Retransmits
100
80
60
40
20
0
06:00
09:00
12:00
15:00
18:00
21:00
00:00
03:00
ps-dev-staging-el9-tk-1.c.esnet-perfsonar.internal -> ps-dev-staging-u20-tk-1.c.esnet-perfsonar.inte...
ps-dev-staging-u20-tk-1.c.esnet-perfsonar.internal -> ps-dev-staging-el9-tk-1.c.esnet-perfsonar.inte...
Latency
1.25 ms
1 ms
0.75 ms
0.5 ms
0.25 ms
0 ms
06:00
09:00
12:00
15:00
18:00
21:00
00:00
03:00
95th percentile - ps-dev-staging-el9-tk-1.c.esnet-perfsonar.internal -> ps-dev-staging-u20-tk-1.c.es...
Latency (Reverse)
0.4 ms
0.3 ms
0.2 ms
0.1 ms
0 ms
-0.1 ms
06:00
09:00
12:00
15:00
18:00
21:00
00:00
03:00
95th percentile - ps-dev-staging-u20-tk-1.c.esnet-perfsonar.internal -> ps-dev-staging-el9-tk-1.c.es...

New UI: Instrumentation

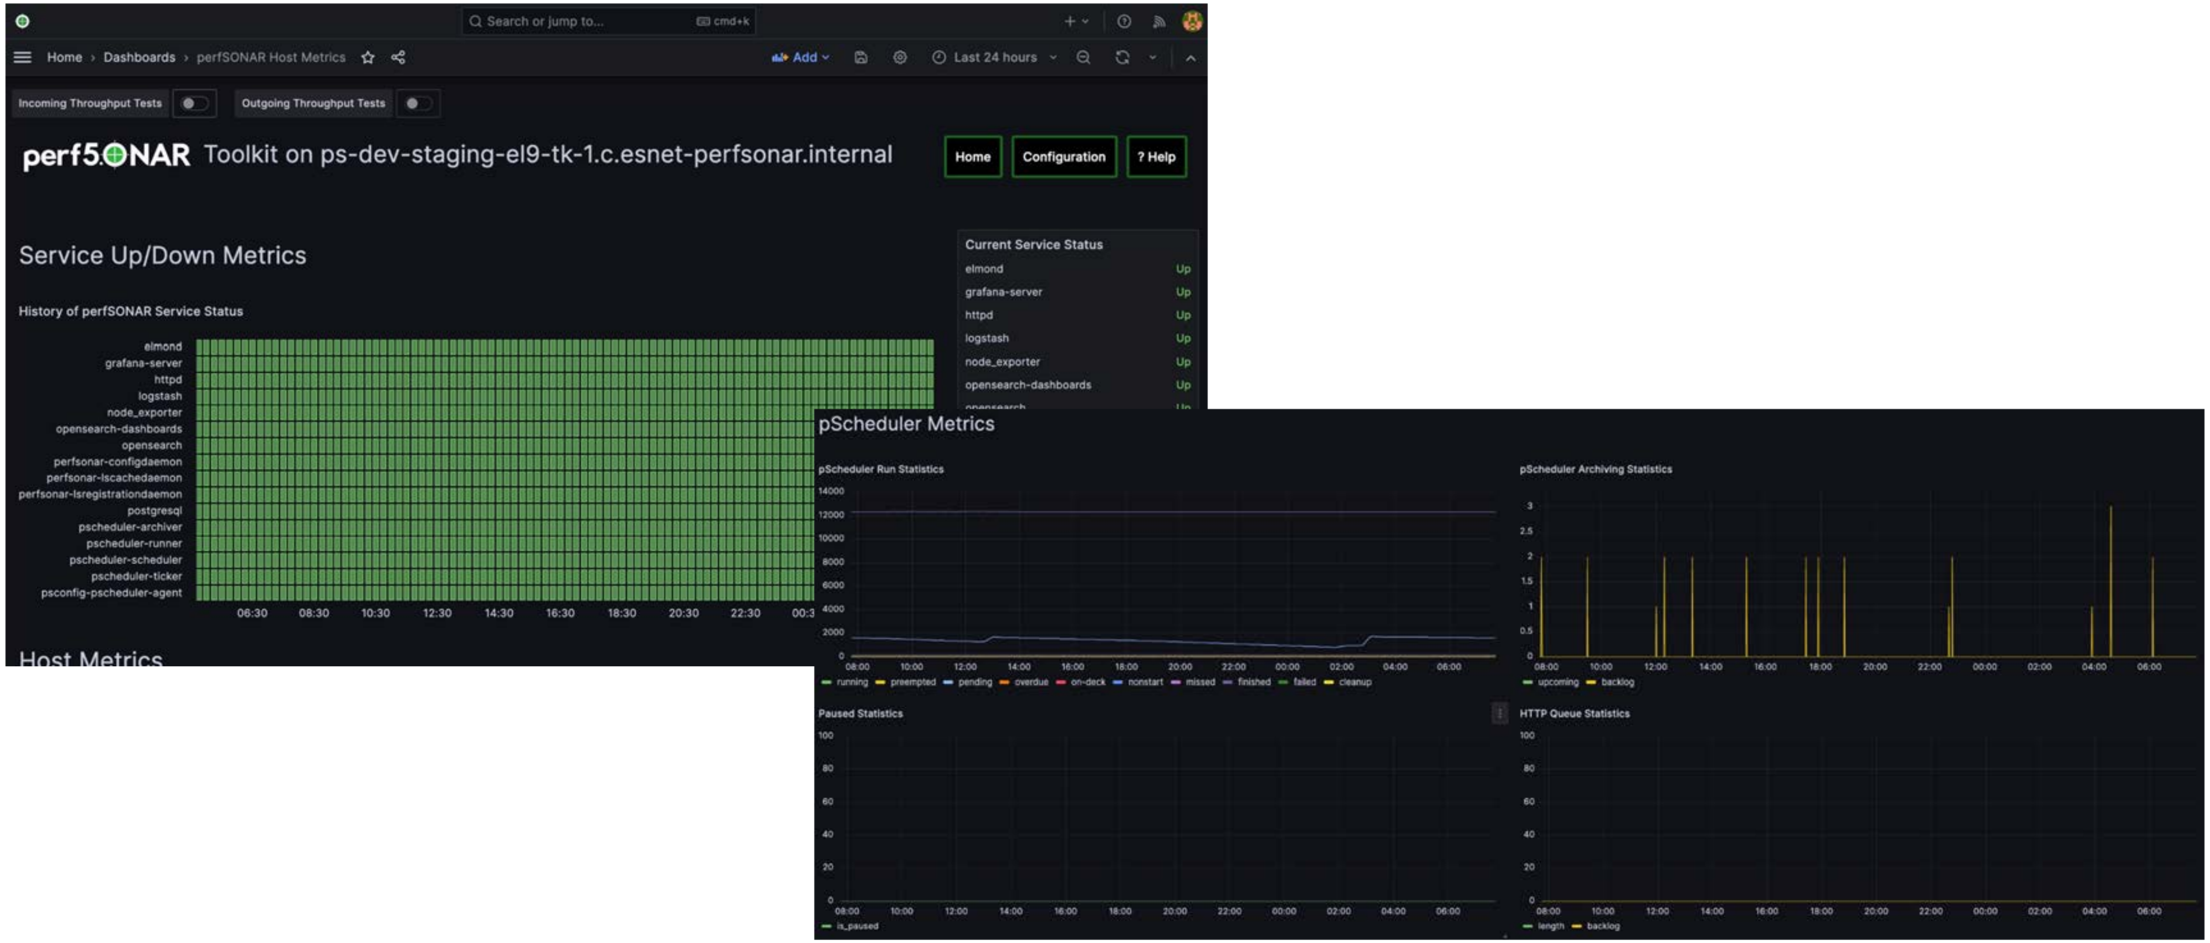

perfSONAR Toolkit on ps-dev-staging-el9-tk-1.c.esnet-perfsonar.internal

Service Up/Down Metrics

History of perfSONAR Service Status

Current Service Status

elmond	Up
grafana-server	Up
httpd	Up
logstash	Up
node_exporter	Up
opensearch-dashboards	Up
opensearch	Up

pScheduler Metrics

New UI: Data Correlation

iperf3 overlay
on host
metrics

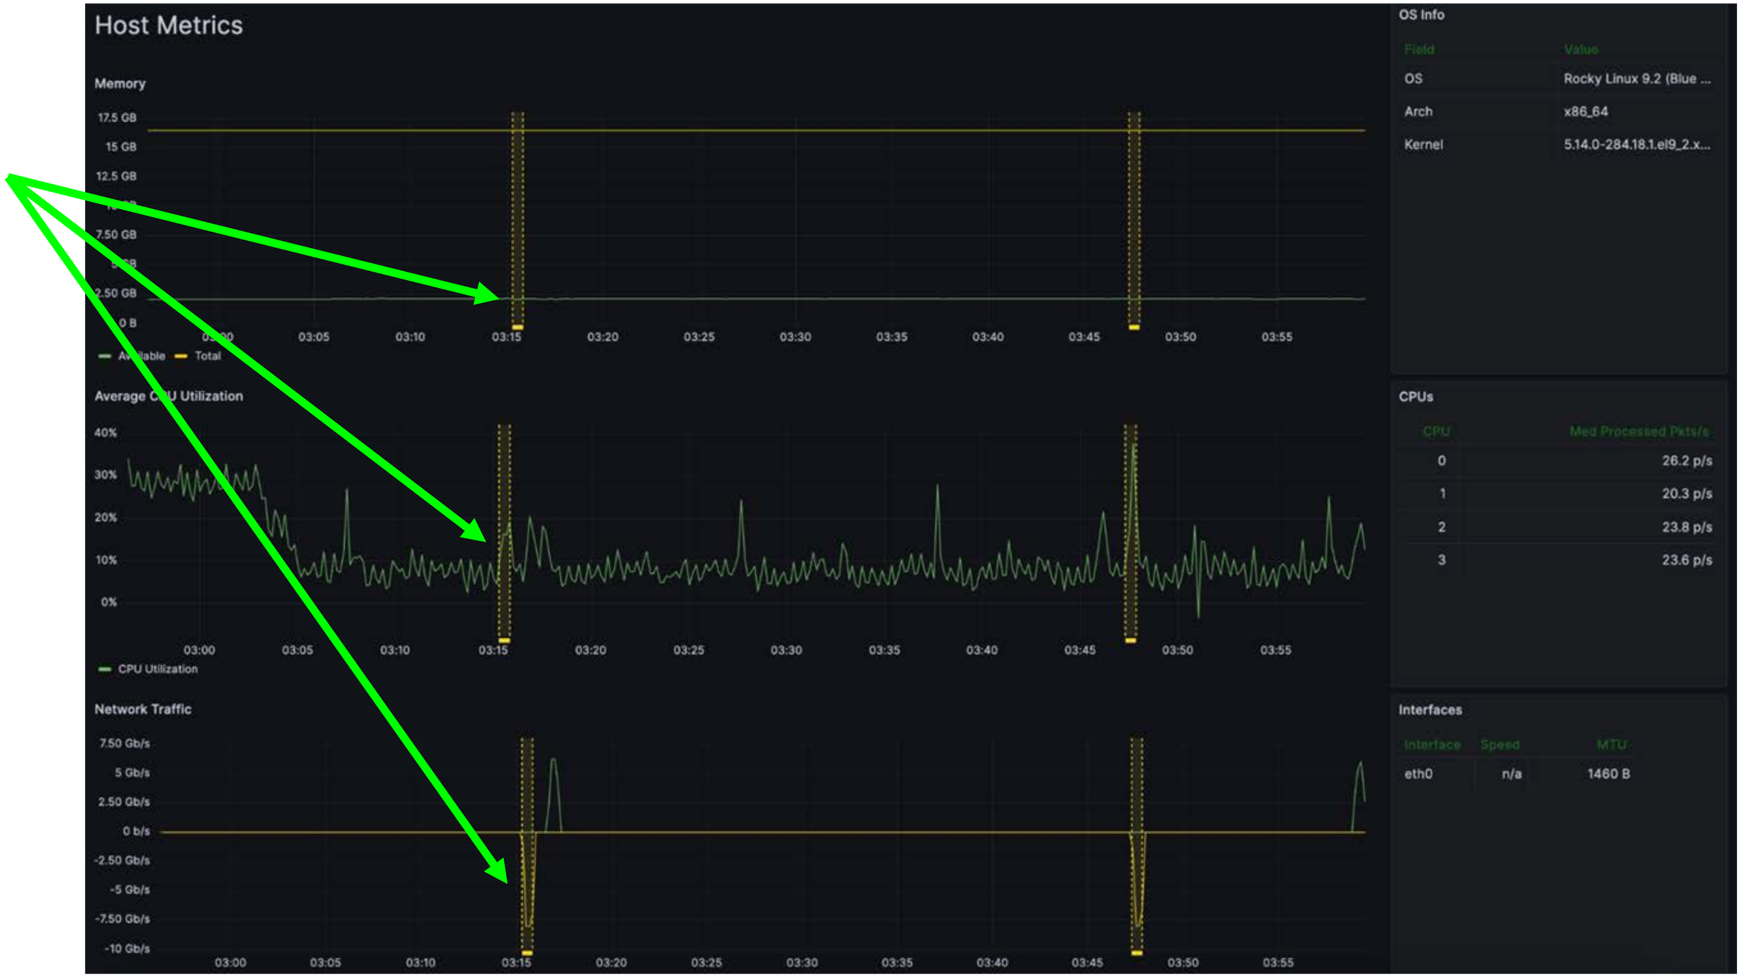

OS Info

Field	Value
OS	Rocky Linux 9.2 (Blue ...
Arch	x86_64
Kernel	5.14.0-284.18.1.el9_2.x...

CPUs

CPU	Med Processed Pkts/s
0	26.2 p/s
1	20.3 p/s
2	23.8 p/s
3	23.6 p/s

Interfaces

Interface	Speed	MTU
eth0	n/a	1460 B

New UI: MaDDash Integration

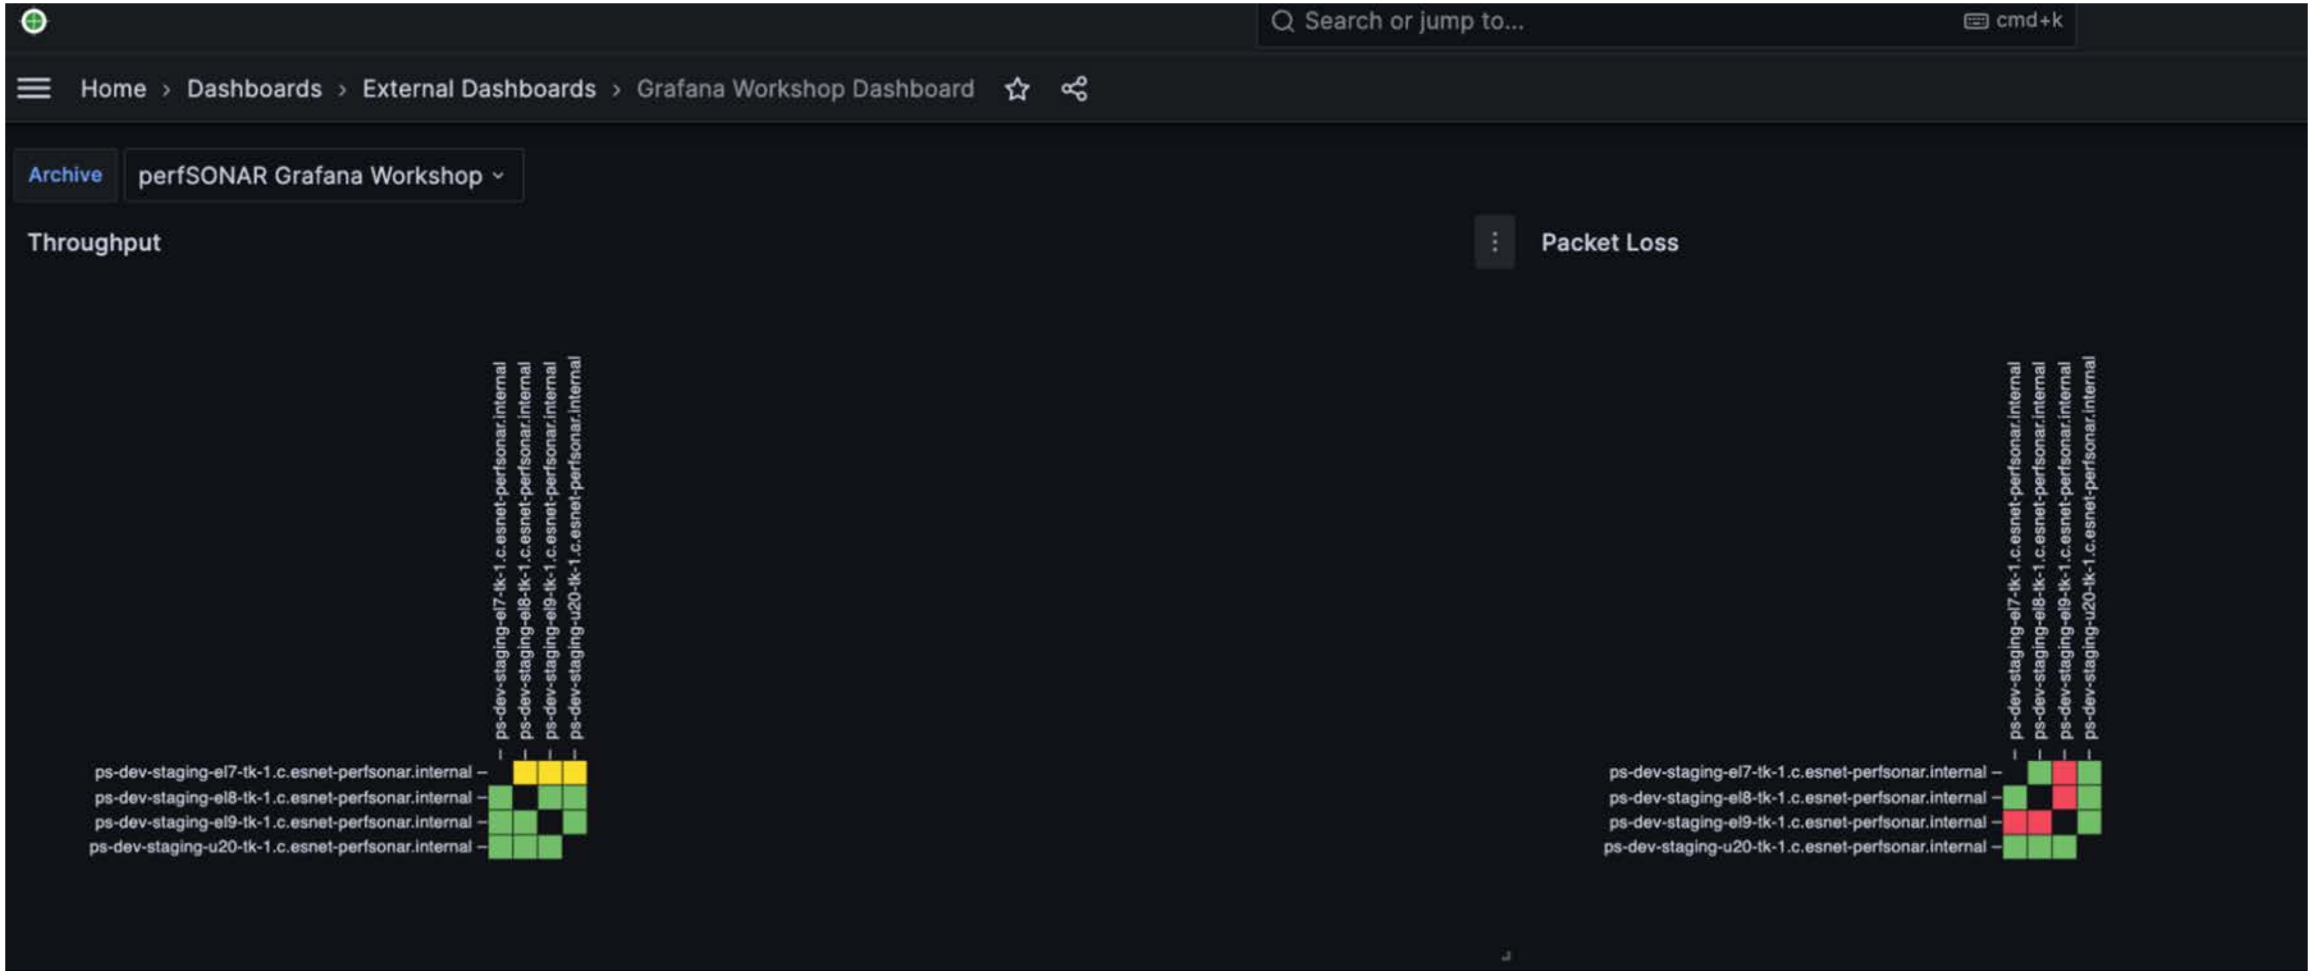
Search or jump to...
cmd+k
Home › Dashboards › External Dashboards › Grafana Workshop Dashboard
Archive
perfSONAR Grafana Workshop

Throughput

	ps-dev-staging-el7-tk-1.c.esnet-perfsonar.internal	ps-dev-staging-el8-tk-1.c.esnet-perfsonar.internal	ps-dev-staging-el9-tk-1.c.esnet-perfsonar.internal	ps-dev-staging-u20-tk-1.c.esnet-perfsonar.internal
ps-dev-staging-el7-tk-1.c.esnet-perfsonar.internal		🟨	🟨	🟨
ps-dev-staging-el8-tk-1.c.esnet-perfsonar.internal	🟩		🟩	🟩
ps-dev-staging-el9-tk-1.c.esnet-perfsonar.internal	🟩	🟩		🟩
ps-dev-staging-u20-tk-1.c.esnet-perfsonar.internal	🟩	🟩	🟩	

Packet Loss

	ps-dev-staging-el7-tk-1.c.esnet-perfsonar.internal	ps-dev-staging-el8-tk-1.c.esnet-perfsonar.internal	ps-dev-staging-el9-tk-1.c.esnet-perfsonar.internal	ps-dev-staging-u20-tk-1.c.esnet-perfsonar.internal
ps-dev-staging-el7-tk-1.c.esnet-perfsonar.internal		🟩	🟥	🟩
ps-dev-staging-el8-tk-1.c.esnet-perfsonar.internal	🟩		🟥	🟩
ps-dev-staging-el9-tk-1.c.esnet-perfsonar.internal	🟥	🟥		🟩
ps-dev-staging-u20-tk-1.c.esnet-perfsonar.internal	🟩	🟩	🟩	

New UI: Customization

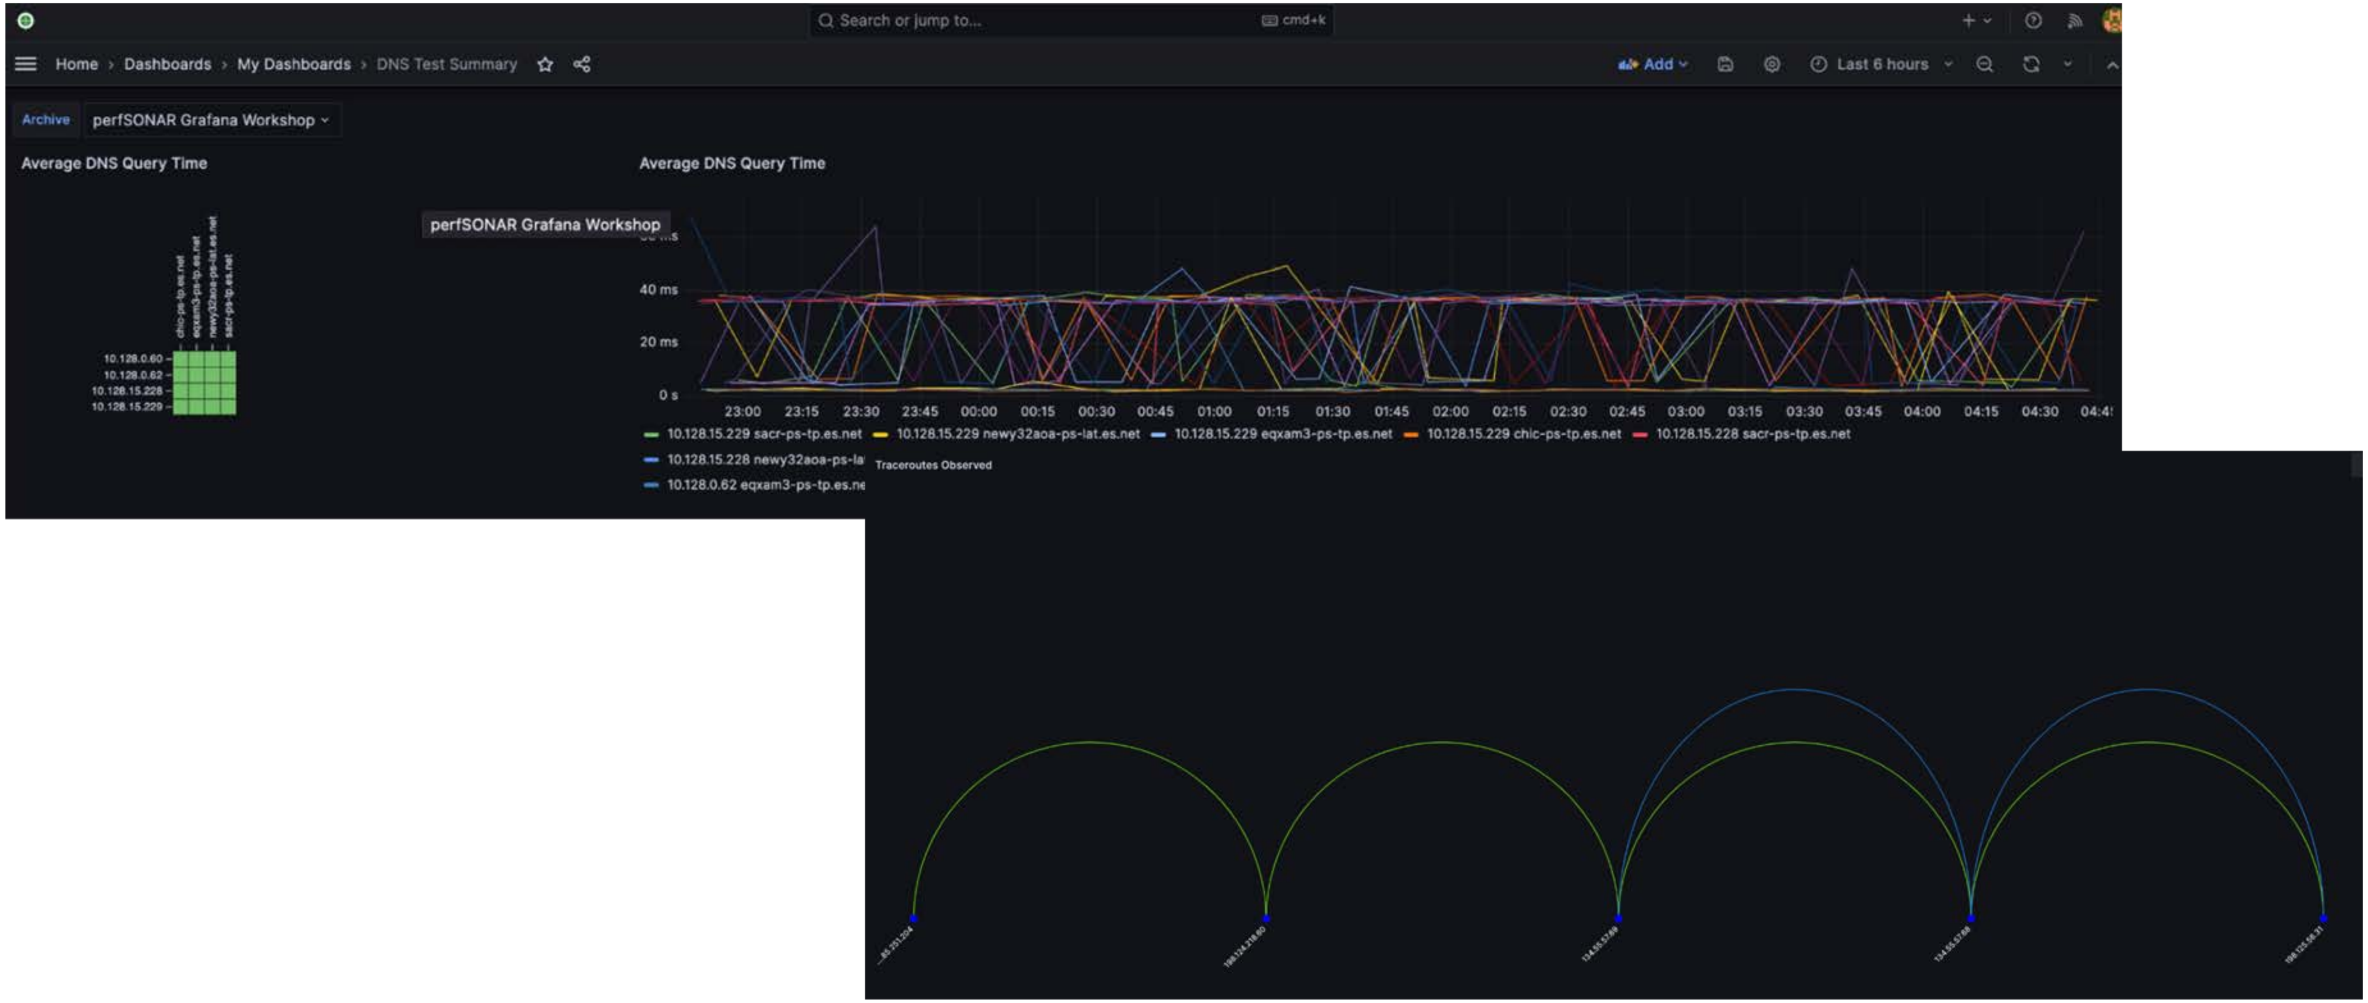

Check out our booth for a live demo!

UP NEXT
SARAH
LARSEN
C L
PULL-ON/VOL
V-HOLD
COLOR
TINT
BRIGHT
PICTURE

Multi-Threaded iperf3

First There was One, Now There are Many

Bruce Mah and Sarah Larsen, ESnet

perfSONAR is developed by a partnership of

What is iperf3?

- iperf3 is an open-source tool that measures network traffic performance between a client and server
- It was designed to be used in perfSONAR, but can also be used on its own
- Linux, MacOS, and FreeBSD are all officially supported

Current State

- Before multi-threading, iperf3 was capable of 30-50 Gbps, with single stream TCP, possibly more with tuning

- Many links in ESnet6 are faster than iperf3
  - Site connections: N x 100G
  - Backbone: 400G +

- So what if we want to support connections with higher bandwidth? How do we get more throughput?

- The problem: Adding more parallel connections doesn't increase throughput

What are we doing about it?

- Adding multi-threading to iperf3
  - Goal: Use multiple CPU cores

Example command: iperf3 –parallel 3

- Note: currently this refers to multiple parallel streams, but in the new multi-threaded iperf, this will represent the number of threads. This is because in the new version each stream will get its own thread.

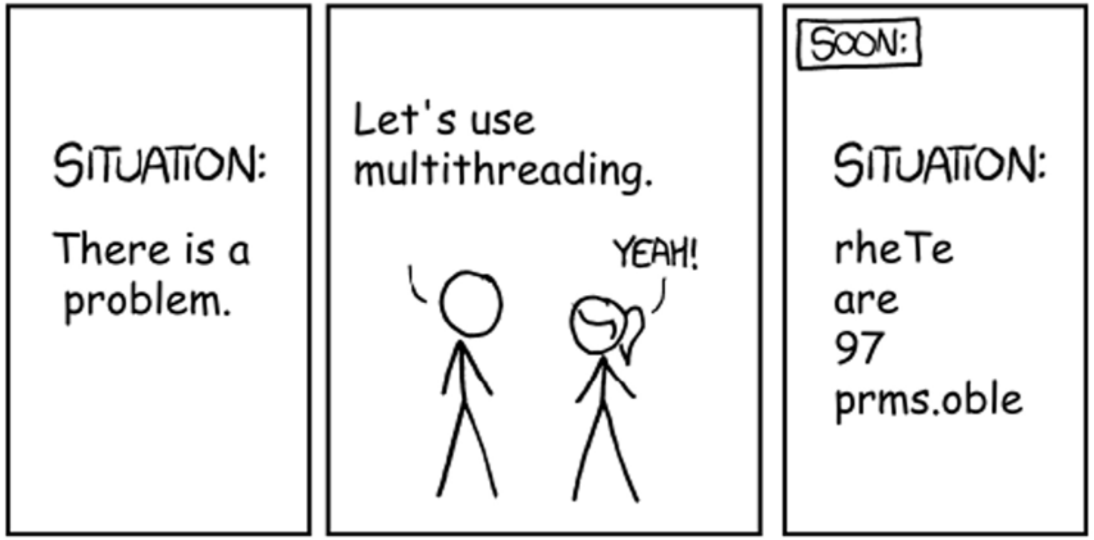

Performance Testing

- Test Setup:
  - Computers:
    - Back-to-back, no routers
    - 200 Gbps link
    - 16 cores, 2 packages, 32 cores total
    - More cores than parallel streams
    - Mellanox ConnectX6
    - ~3 GHz AMD CPU

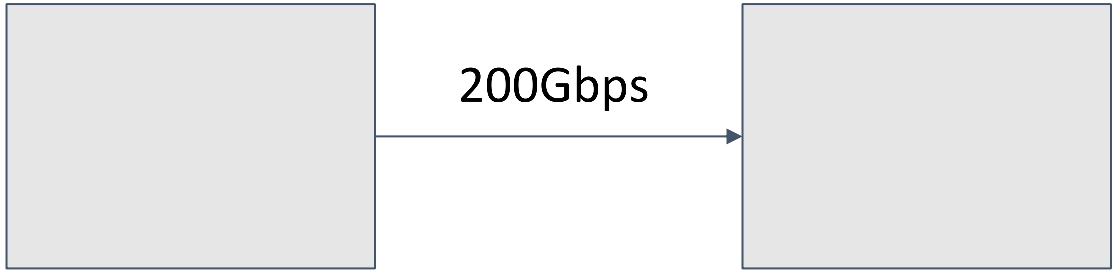

Effects on Performance

iperf3-mt has significants improvements in throughput performance over iperf3.

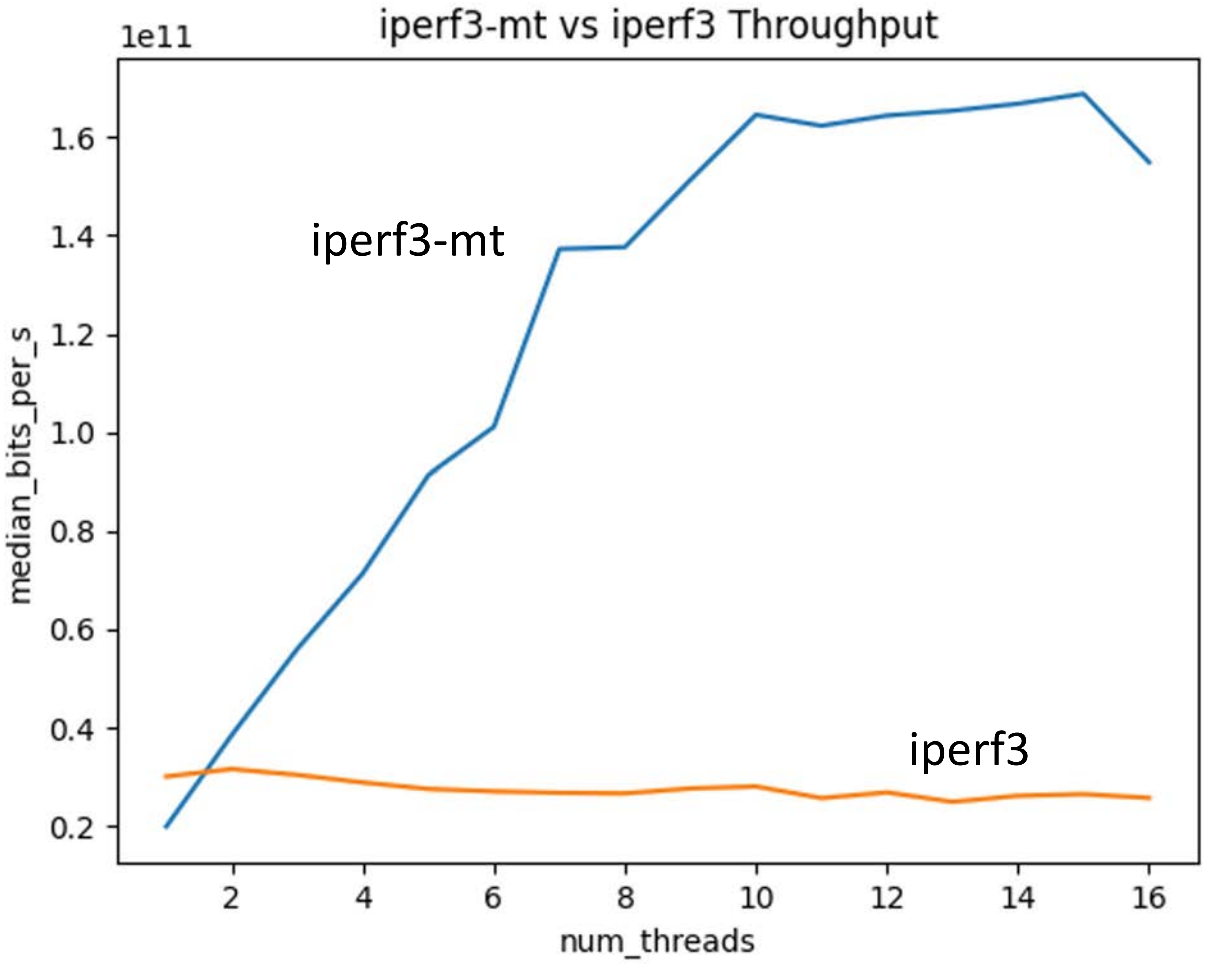

How does it compare to iperf2?

Multi-threaded iperf3 has similar throughput results to iperf2

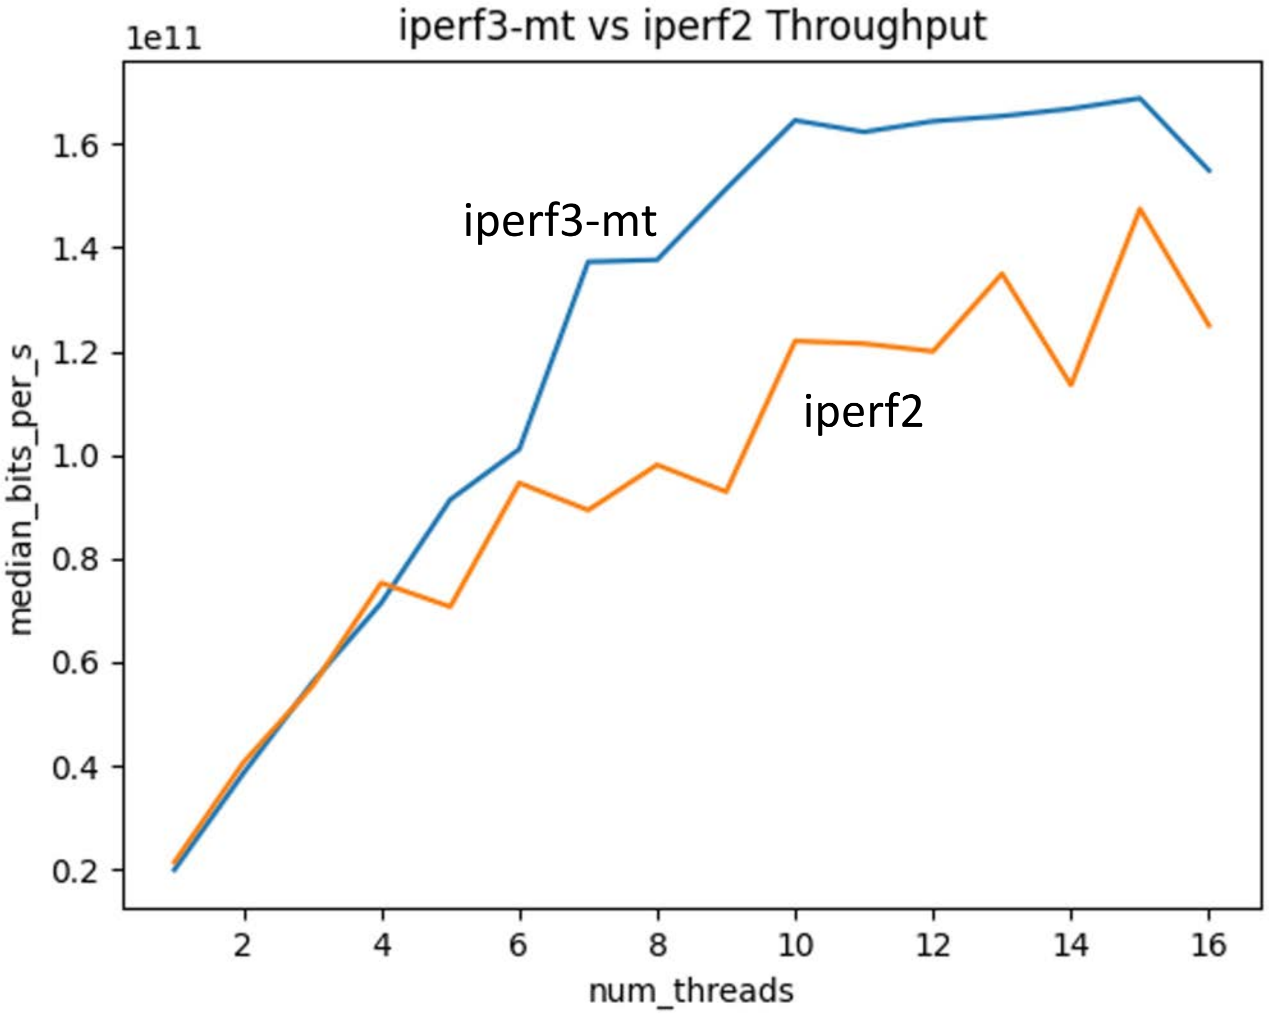

What’s Next?

- Testing on ESnet perfSONAR nodes and integration with pscheduler
- Further analysis with how threading impacts performance, virtual machines vs hardware
- Ensure correctness of threading prototype
- Experiments with tuning

When?

“Soon”

But seriously Q4 2023/Q1 2024

We’re going to test it on ESnet perfSONAR hosts before releasing the final version.

Any interest in testing or want the current working version: https://github.com/esnet/iperf/tree/mt

Summary

The new multi-threaded iperf3 can better test performance of faster paths.

GitHub: https://github.com/esnet/iperf

Multi-threaded: https://github.com/esnet/iperf/tree/mt

Contact email: iperf@es.net

Questions?

UP NEXT
MARK
FEIT
C L
PULL-ON/VOL
V.HOLD
COLOR
TINT
BRIGHT
PICTURE

2023 INTERNET2 TECHNOLOGY exchange

perfSONAR on the Internet2 Backbone

Mark Feit – Internet2 / perfSONAR Development Team
mfeit@internet2.edu

WE DID A THING

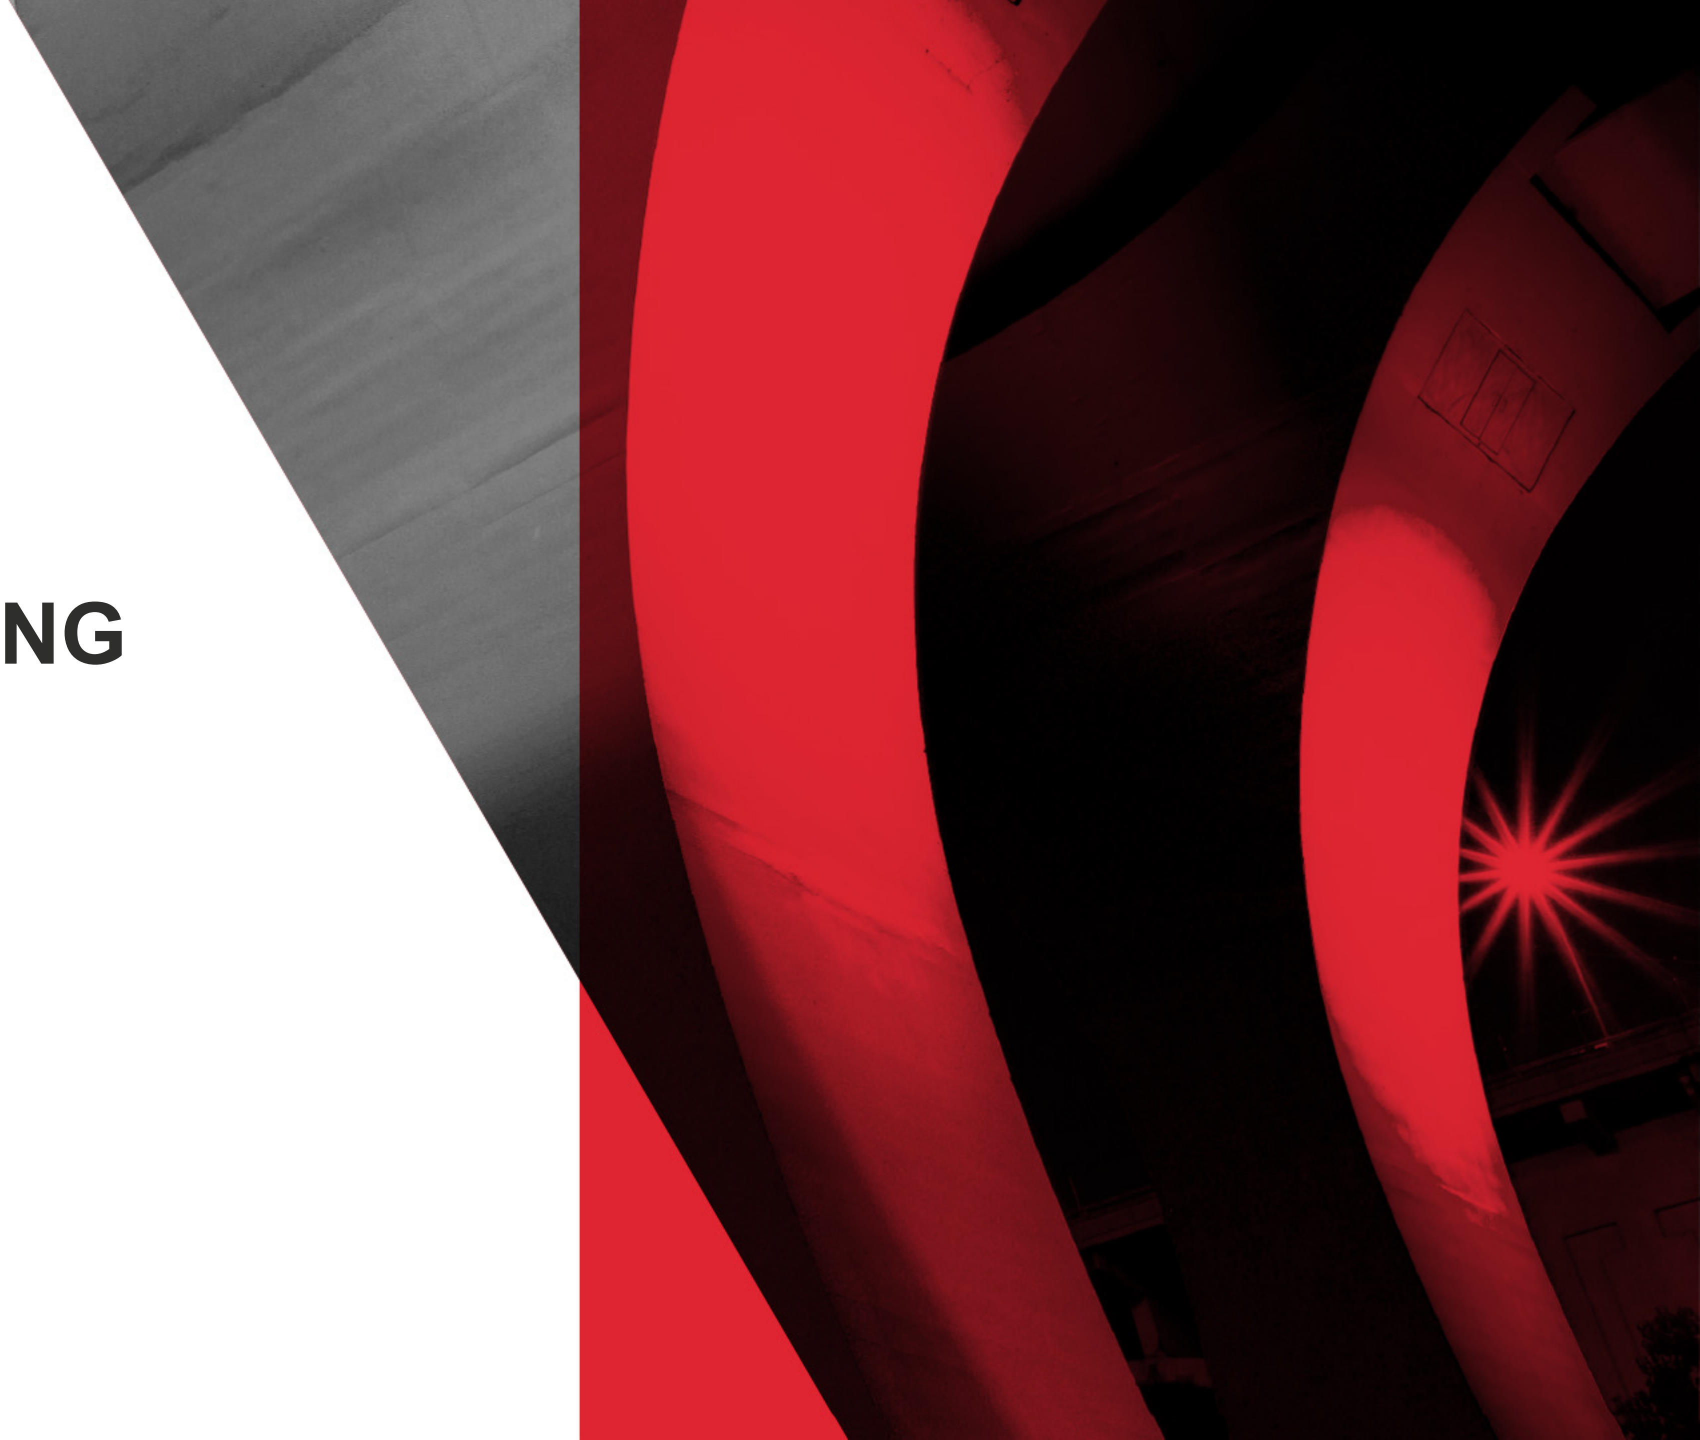

A Plethora of Public perfSONAR Points in PoPs

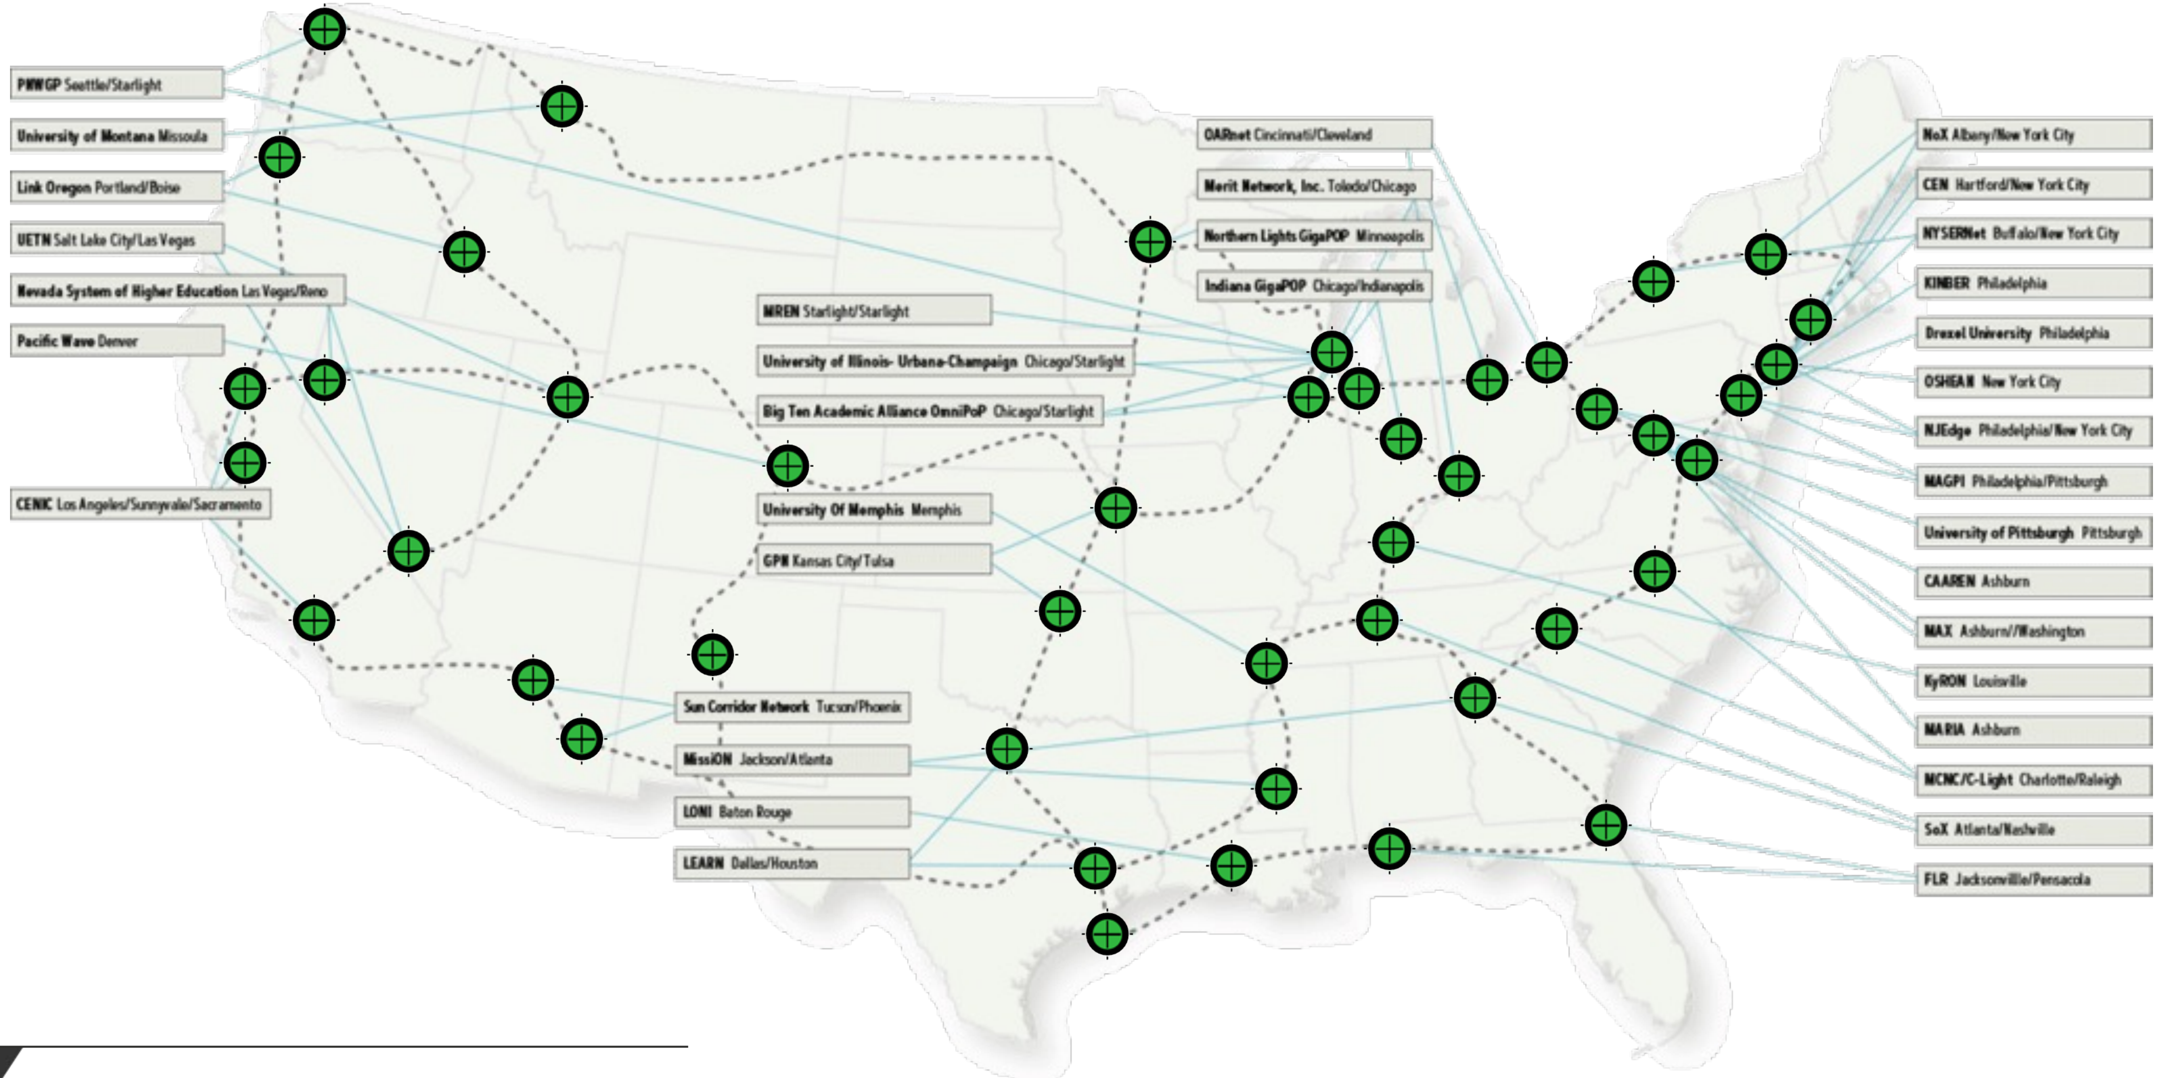
PNWGP Seattle/Starlight
University of Montana Missoula
Link Oregon Portland/Boise
UETN Salt Lake City/Las Vegas
Nevada System of Higher Education Las Vegas/Reno
Pacific Wave Denver
CENIC Los Angeles/Sunnyvale/Sacramento
MREN Starlight/Starlight
University of Illinois- Urbana-Champaign Chicago/Starlight
Big Ten Academic Alliance OmniPoP Chicago/Starlight
University Of Memphis Memphis
GPN Kansas City/Tulsa
Sun Corridor Network Tucson/Phoenix
MissiON Jackson/Atlanta
LONI Baton Rouge
LEARN Dallas/Houston
OARnet Cincinnati/Cleveland
Merit Network, Inc. Toledo/Chicago
Northern Lights GigaPOP Minneapolis
Indiana GigaPOP Chicago/Indianapolis
NoX Albany/New York City
CEN Hartford/New York City
NYSERNet Buffalo/New York City
KINBER Philadelphia
Drexel University Philadelphia
OSHEAN New York City
NJEdge Philadelphia/New York City
MAGPI Philadelphia/Pittsburgh
University of Pittsburgh Pittsburgh
CAAREN Ashburn
MAX Ashburn//Washington
KyRON Louisville
MARIA Ashburn
MCNC/C-Light Charlotte/Raleigh
SoX Atlanta/Nashville
FLR Jacksonville/Pensacola

Using the Internet2 Public perfSONAR Nodes

***PoP*.ps.internet2.edu**

- IPv4 and IPv6 available

- Log into a perfSONAR system (yours, not ours) and run a task with pScheduler

```
pscheduler task throughput --dest PoP.ps.internet2.edu
```

PoP Directory

Albany	alba
Ashburn	ashb
Atlanta	atla
Boise*	bois
Boston**	bost
Charlotte	char
Chicago	chic
Chicago	eqch
Chicago	star
Cincinnati	cinc
Cleveland	clev
Dallas	dall
Dallas	dall3
Denver	denv
El Paso	elpa
Fargo**	farg
Hartford	hart2
Houston	houh
Houston	hous
Indianapolis	indi
Jackson	jcsn
Jacksonville	jack
Kansas City	kans
Las Vegas	lasv
Los Angeles	losa
Los Angeles	losa2
Lousville	loui
Minneapolis	minn
Missoula	miss2
Nashville	nash
New York	newy2
New York	newy32aoa
Pensacola*	pens
Philadelphia	phil
Phoenix	phoe
Pittsburgh	pitt
Portland	port
Raleigh	rale
Reno	reno
Sacramento	sacr
Salt Lake City	salt
San Jose	sanj
Seattle*	seat
Sunnyvale	sunn
Toledo	tole2
Tucson	tucs
Tulsa	tuls
Washington	wash

*Not yet in service **Future

Topology

PoP Type	Router Connection
Distributed	Core*
Multi-Degree	
Interconnect	
Split Interconnect	Aggregation*

*First where more than one is installed

Network Reachability

<u>Now</u>

R&E + I2PX

<u>Later</u>

Elsewhere

Initial Administrative Limitations

- Caps on **`throughput`** test bandwidth
  - R&E* 10 Gb/s
  - Elsewhere 1 Gb/s
  - Higher bandwidth considered on a case-by-case basis
- No, **`disk-to-disk`**, **`s3throughput`**, **`idleex`** or **`wifibssid`** tests

- Tests with a **`duration`** parameter are limited to one minute
- Repeating tests
  - Not more-frequently than once per hour
  - Must wrap up within 24 hours
  - **`repeat-until`** parameter not allowed

- Testing priorities
  - Internet2 Internal
  - R&E Networks
  - Everyone Else
- These limits will be refined periodically to make sure the community's needs are being met.

*Determined using ESnet's R&E network list: `http://stats.es.net/sample_configs/pscheduler/ren`

Beta Period

- Now through January
- Feature set close to production
- Feedback is appreciated

  `pas@internet2.edu`

β

Beta Will Be Beta

- Working through some teething with the NICs in some PoPs
  - Systems disappear from the network

- Ongoing experimentation on systems in **CLEV**, **PHOE** and **TUSC** PoPs

β

Performance Tuning

New OS

+ New Kernel

+ New NIC Driver

= New System Tunings

- Higher-speed throughput requires additional test parameters to run well

THE SYSTEMS

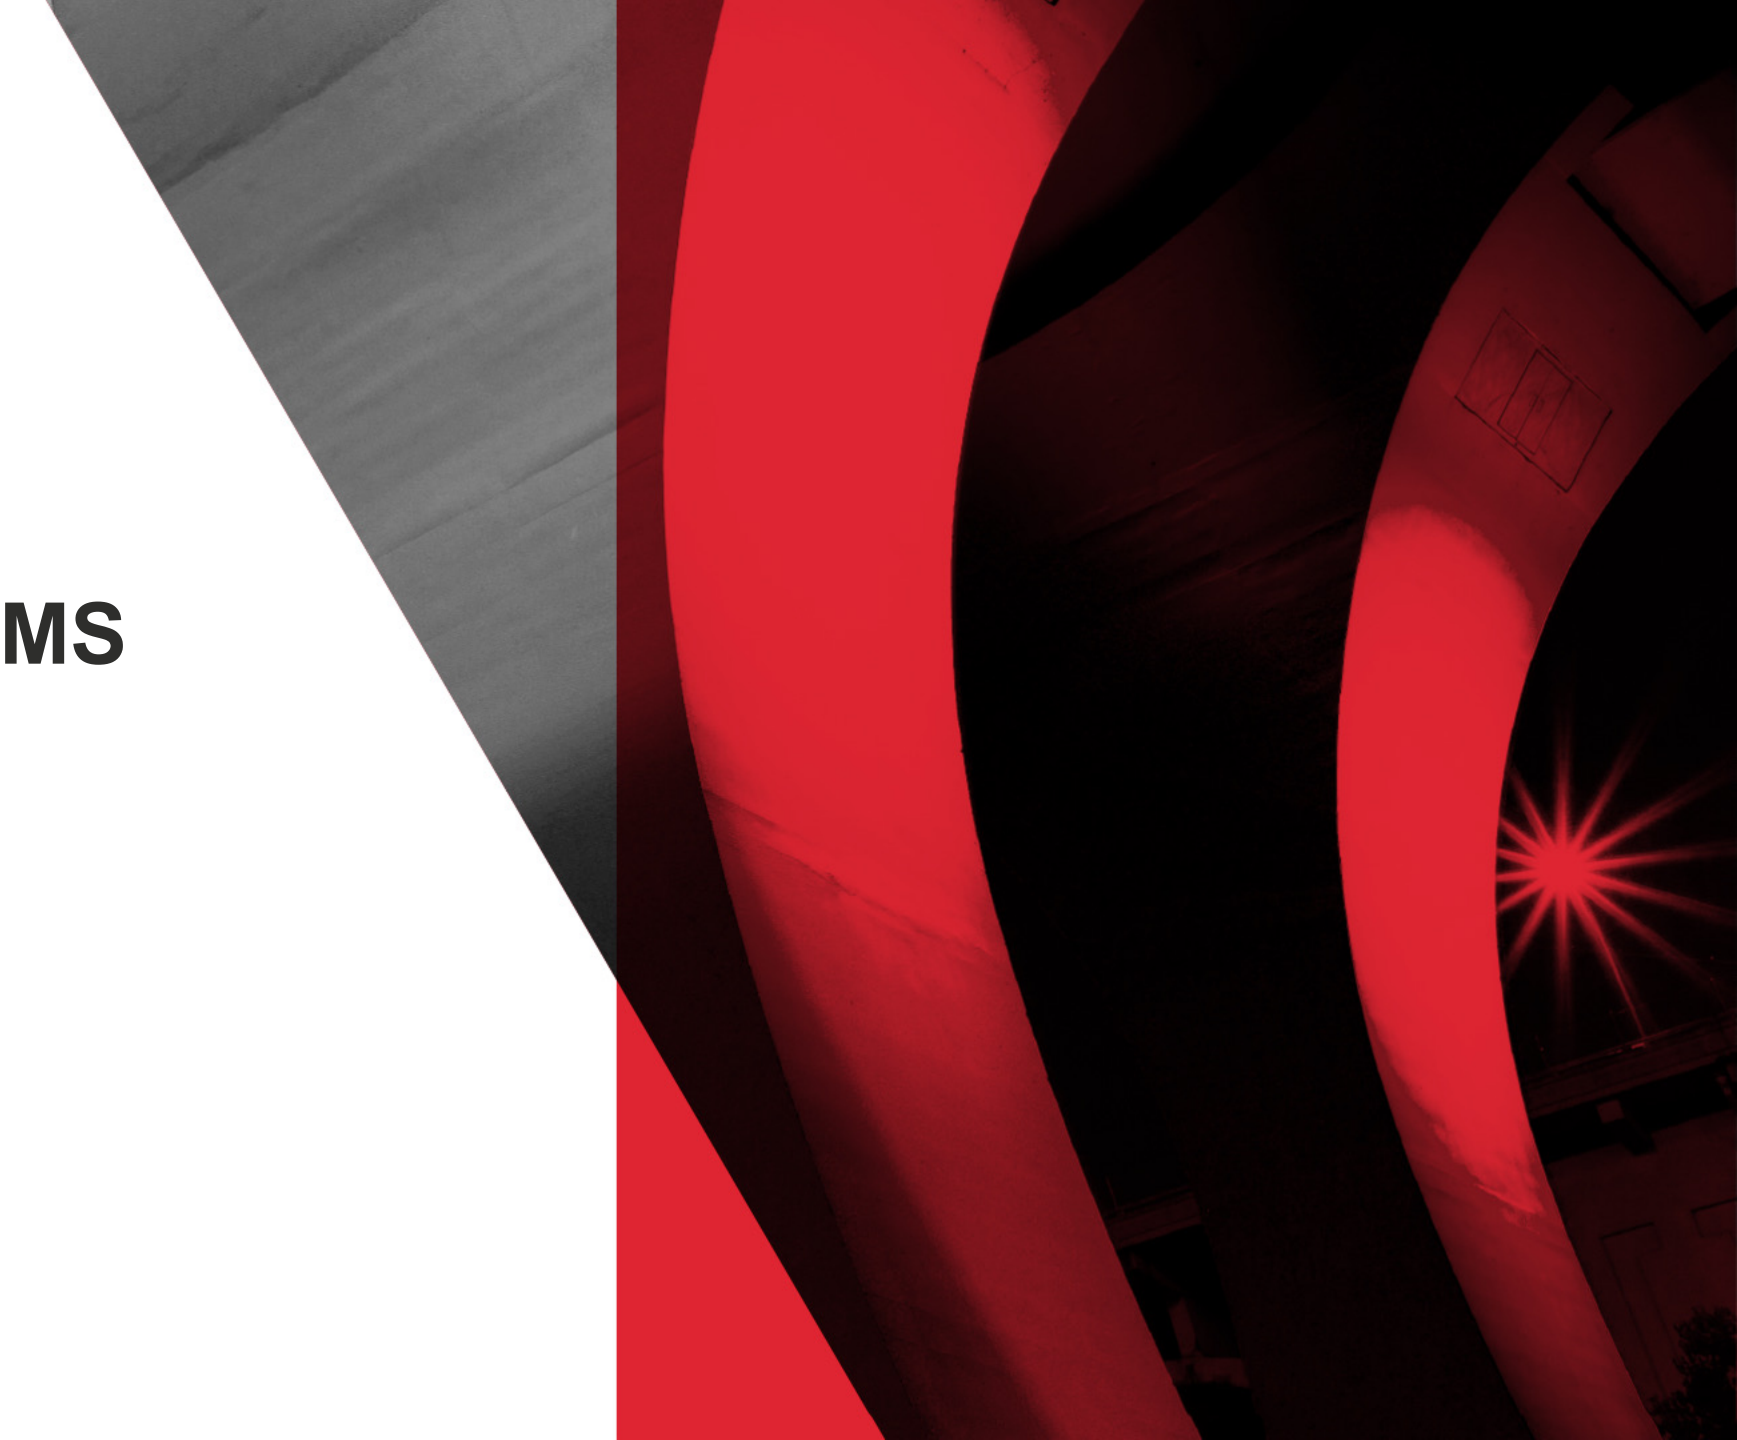

A ~~Chicken~~ System in Every ~~Pot~~ PoP

- Dell R6515 (NGI Buildout) and R6615 (Later)
- AMD EPYC 7402P 24 Cores / 48 Threads at 2.8 GHz
- 128 GiB RAM

- Broadcom 2x 10 GbE
- Mellanox Connect-X 5 2x 100 GbE

Connections to routers vary by PoP type.

- AlmaLinux 9

Host Architecture

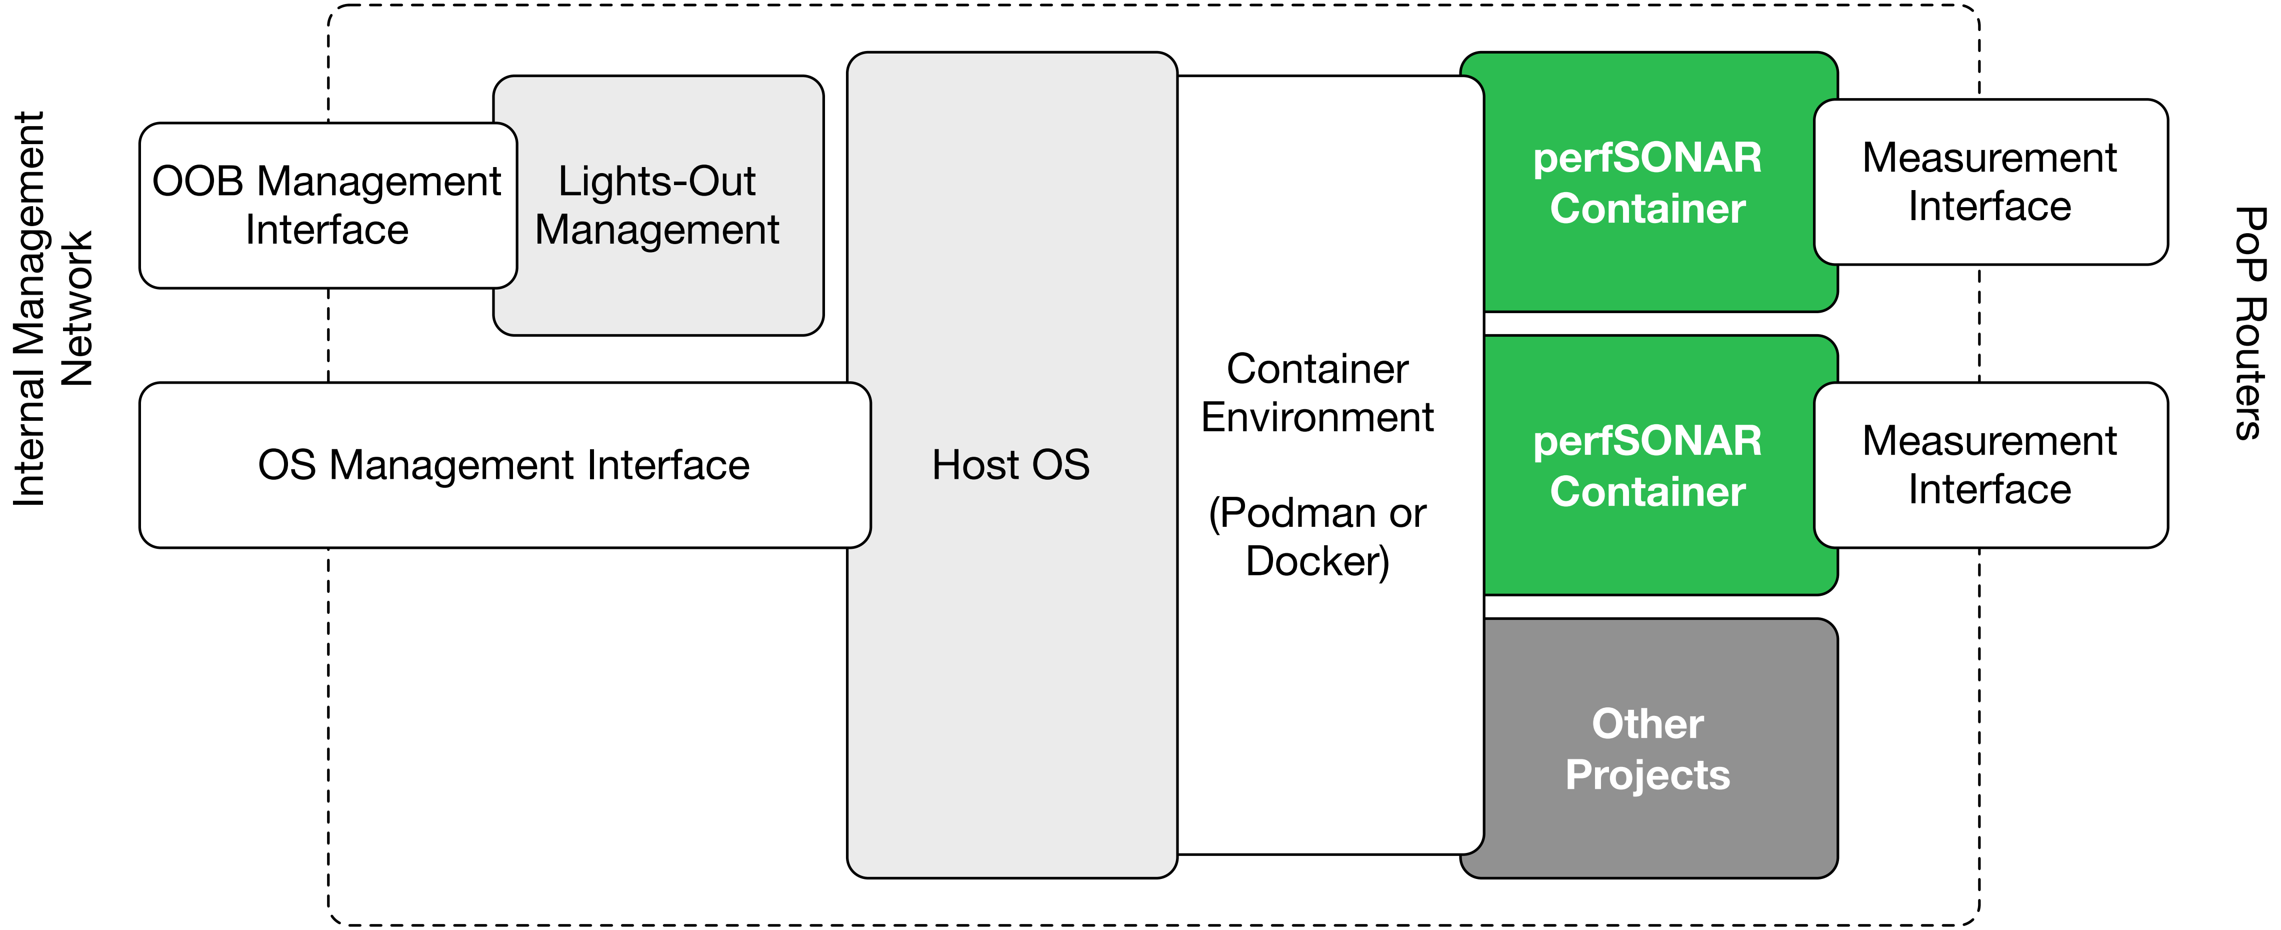

~~Special~~ Secret Sauce: The `macvlan` Network Driver

- Binds a host interface directly into a container
- Bypasses additional container networking code
- Negligible performance difference vs. bare metal
- No address assigned on the host
  - Prevents external access to the OS

Big MACVLAN

Resources for Performance

- 100 GbE Interface Containers
  - 12 Dedicated CPU Cores (Threads)
  - 32 GiB Dedicated RAM
  - IRQ Affinity (Tuning)
  - 93+ Gb/s

- 10 GbE Interface Containers
  - Shared CPU Cores
  - Shared Memory
  - No special performance tuning

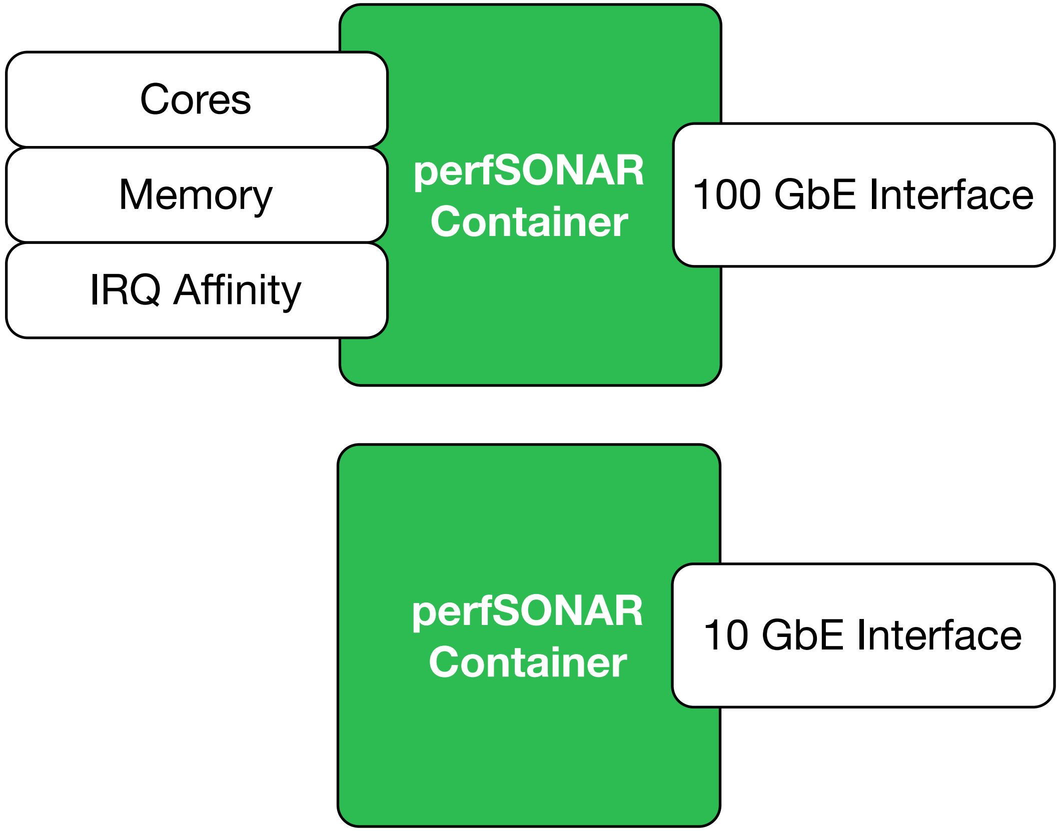

Deployment Technology Stack(s)

- Single data set with detailed information on PoPs, systems, interfaces and networks
  - System kickstart files
  - Ansible
  - Salt
  - Assorted shell scripts
  - DNS records
  - Internal proxy configuration and ACLs

- Ultimate goal is to use Salt

Why so many different technologies?

- The tools are buggy.

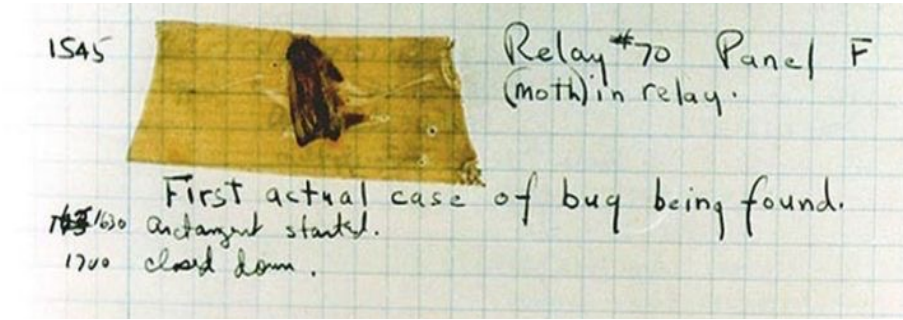

- IPv4 /31s used for point-to-point connections
  - Halves address use compared to /30s
  - See RFC 3021

- Docker couldn’t handle those at all
  - Patch submitted in 2021, released in Docker 23.0.0 (2022).

- Podman is fine with them but its web API isn’t.

The Long and Winding Road

- Began with Ansible for expediency
  - Problems with networks being re-created with each run
  - Destroyed/rebuilt the container. *No bueno.*
- Tried Salt
  - Stymied by the /31 problem
- Ended up with a set of shell scripts
  - Tied together for single-command provisioning of the entire system
  - Small bites that can be easily converted to Ansible or Salt

Future Development: Deployment Kit

- Will be derived from deployments at Internet2 and elsewhere
  - Based around Ansible
  - Minimal manual perfSONAR host configuration
  - Driven by data: Configure and go
  - Container-per-interface model (Plain or VLANs)

- Initial version targets Debian
  - EL to follow

Thanks!

Public perfSONAR

***PoP*.ps.internet2.edu**

Feedback

pas@internet2.edu

UP NEXT

LÆTITIA DELVAUX

perfSONAR in GÉANT

2023 edition

Lætitia A Delvaux, PSNC / GÉANT Project

GN5-1 WP6T3 Task Leader

TechEX23, Minneapolis, MN, USA

20 September 2023

Public (PU)

GN5-1

Multiple perfSONAR activites in GÉANT

- Lookup Service dashboards
- perfSONAR deployments
- Microdep integration
- On-demand perfSONAR Graphical User Interface (psGUI)

Lookup Service dashboards

- Display, filter and search the content of the Lookup Service
- Grafana 8 based
  - Filtering on text, domains, communities
  - Stats on hosts and services, maps
- https://stats.perfsonar.net
  - Replaces ESnet Service Directo

- Next steps:
  - Port to Grafana 9
  - Filter on multiple values

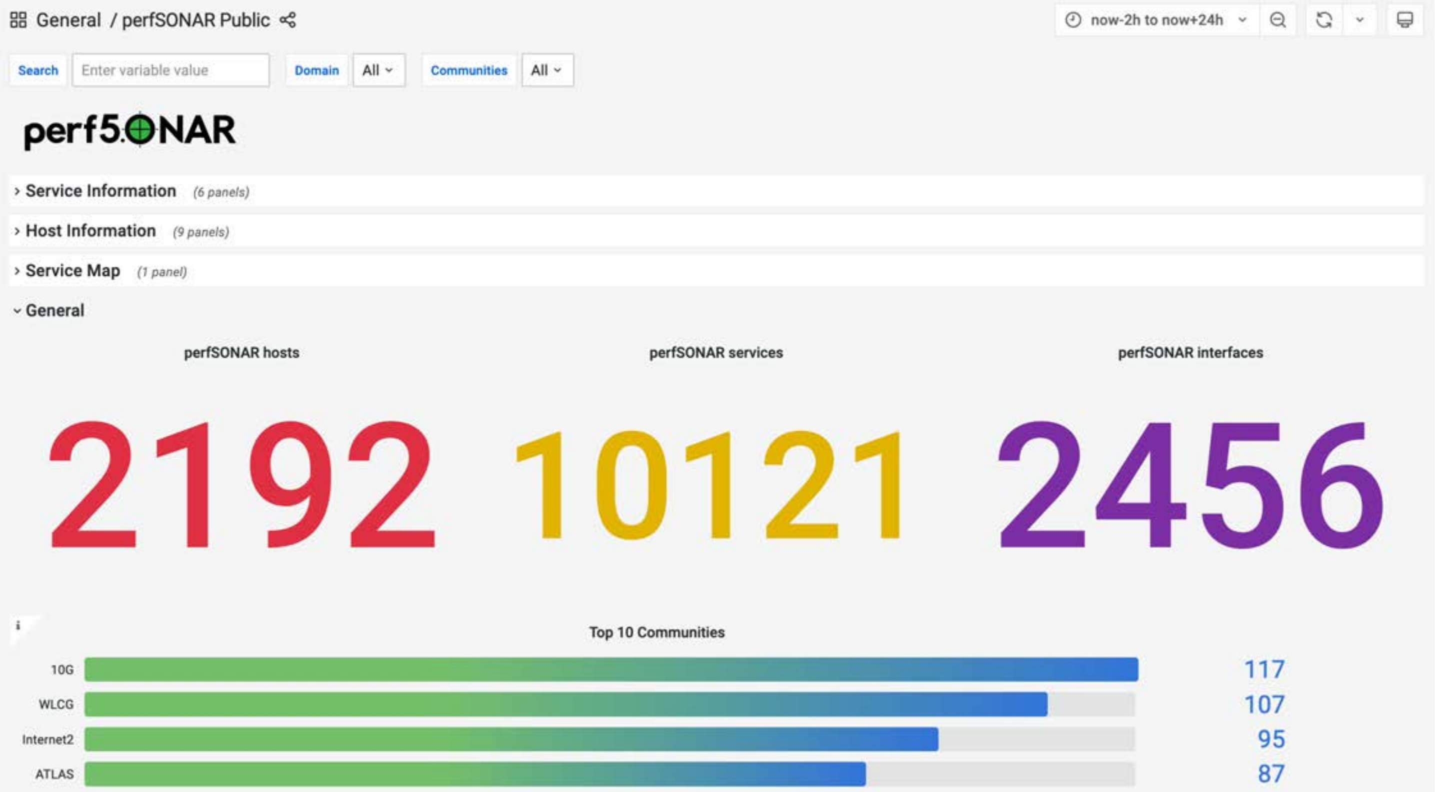

General / perfSONAR Public

now-2h to now+24h

Search: Enter variable value | Domain: All | Communities: All

perfSONAR

› Service Information (6 panels)

› Host Information (9 panels)

› Service Map (1 panel)

⌄ General

perfSONAR hosts	perfSONAR services	perfSONAR interfaces
2192	10121	2456

Top 10 Communities

10G	117
WLCG	107
Internet2	95
ATLAS	87

perfSONAR deployments in the GÉANT network (1/2)

- 10 public deployments on the core network: https://network.geant.org/perfsonar/

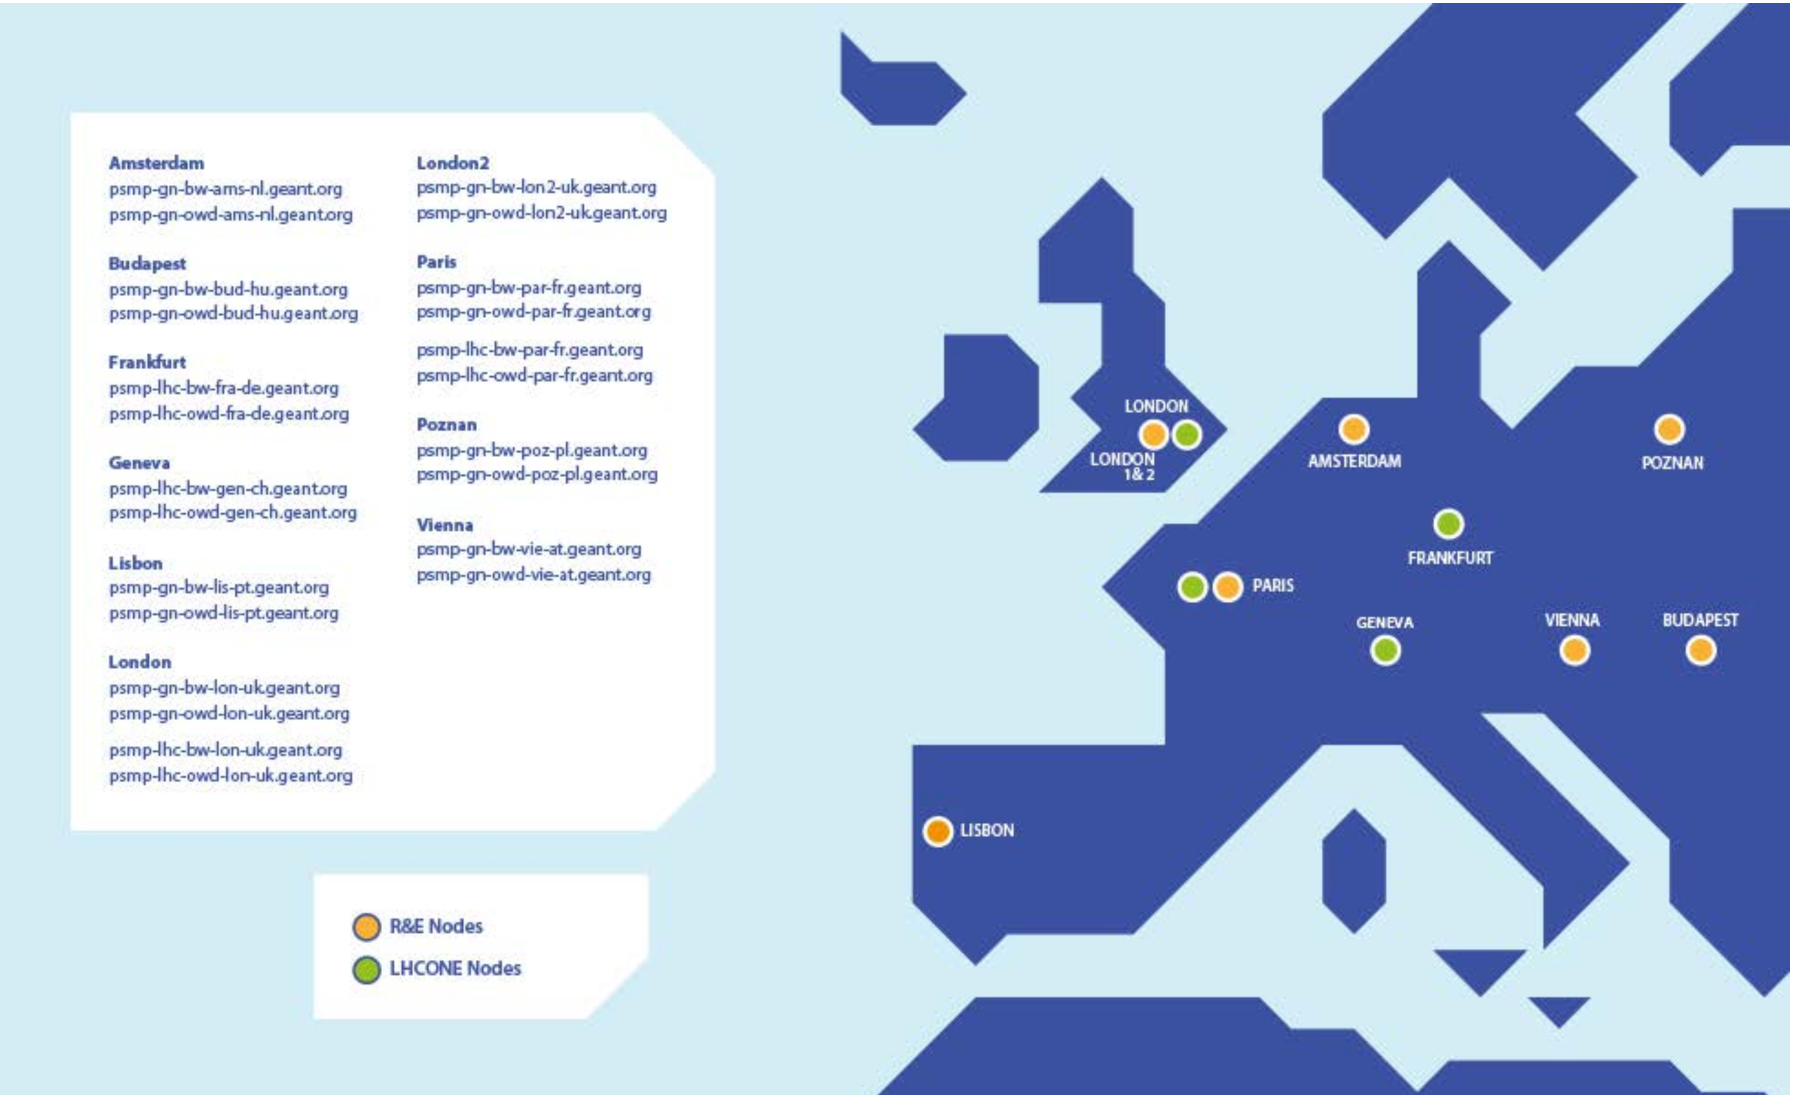

Amsterdam
psmp-gn-bw-ams-nl.geant.org
psmp-gn-owd-ams-nl.geant.org

Budapest
psmp-gn-bw-bud-hu.geant.org
psmp-gn-owd-bud-hu.geant.org

Frankfurt
psmp-lhc-bw-fra-de.geant.org
psmp-lhc-owd-fra-de.geant.org

Geneva
psmp-lhc-bw-gen-ch.geant.org
psmp-lhc-owd-gen-ch.geant.org

Lisbon
psmp-gn-bw-lis-pt.geant.org
psmp-gn-owd-lis-pt.geant.org

London
psmp-gn-bw-lon-uk.geant.org
psmp-gn-owd-lon-uk.geant.org

psmp-lhc-bw-lon-uk.geant.org
psmp-lhc-owd-lon-uk.geant.org

London2
psmp-gn-bw-lon2-uk.geant.org
psmp-gn-owd-lon2-uk.geant.org

Paris
psmp-gn-bw-par-fr.geant.org
psmp-gn-owd-par-fr.geant.org

psmp-lhc-bw-par-fr.geant.org
psmp-lhc-owd-par-fr.geant.org

Poznan
psmp-gn-bw-poz-pl.geant.org
psmp-gn-owd-poz-pl.geant.org

Vienna
psmp-gn-bw-vie-at.geant.org
psmp-gn-owd-vie-at.geant.org

perfSONAR deployments in the GÉANT network (2/2)

- Performance Measurement Platform (PMP)
  - Small nodes (Intel NUC) and VM
  - Deployed in GÉANT partners organisations
- Measurements
  - Diverse set of measurements
  - Regularly to GÉANT core network
  - Verify GÉANT access links
  - International connections (ESnet, Internet2, RNP, …)
  - 2nd tiers: University networks
- https://pmp-central.geant.org

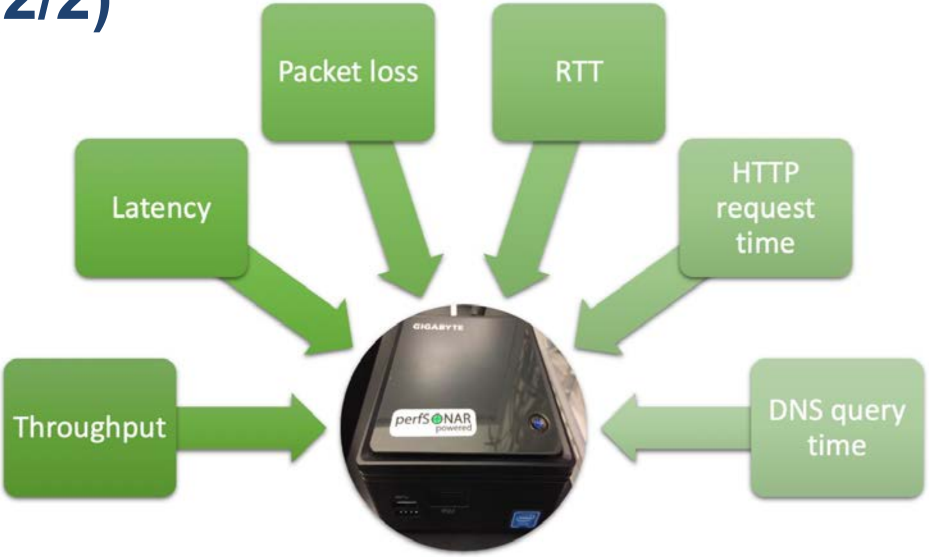

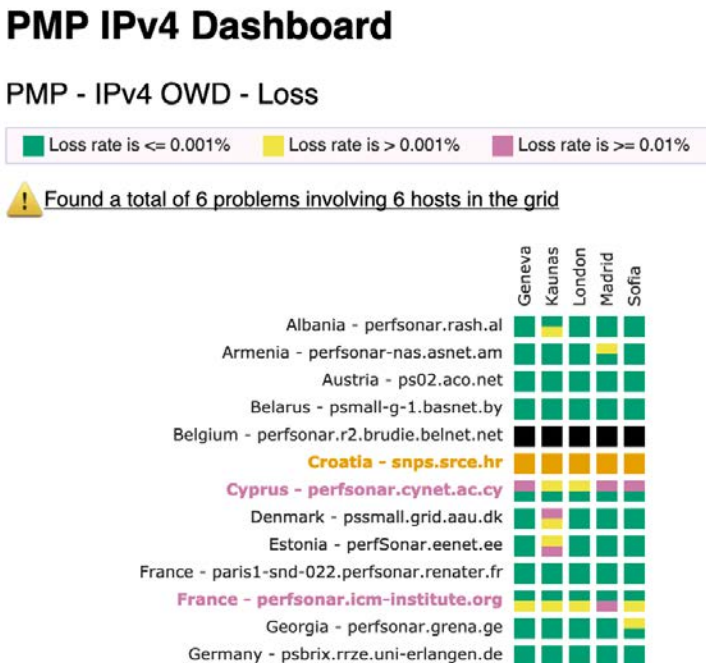

PMP IPv4 Dashboard

PMP - IPv4 OWD - Loss

🟩 Loss rate is <= 0.001% 🟨 Loss rate is > 0.001% 🟪 Loss rate is >= 0.01%

⚠ Found a total of 6 problems involving 6 hosts in the grid

	Geneva	Kaunas	London	Madrid	Sofia
Albania - perfsonar.rash.al	🟩	🟩🟨	🟩	🟩	🟩
Armenia - perfsonar-nas.asnet.am	🟩	🟩	🟩	🟨🟩	🟩
Austria - ps02.aco.net	🟩	🟩	🟩	🟩	🟩
Belarus - psmall-g-1.basnet.by	🟩	🟩	🟩	🟩	🟩
Belgium - perfsonar.r2.brudie.belnet.net	⬛	⬛	⬛	⬛	⬛
Croatia - snps.srce.hr	🟧	🟧	🟧	🟧	🟧
Cyprus - perfsonar.cynet.ac.cy	🟪🟩	🟨🟩	🟨🟩	🟪🟩	🟪🟩
Denmark - pssmall.grid.aau.dk	🟩	🟪🟨	🟩	🟩	🟩
Estonia - perfSonar.eenet.ee	🟩	🟨🟪	🟩	🟩	🟩
France - paris1-snd-022.perfsonar.renater.fr	🟩	🟩	🟩	🟩	🟩
France - perfsonar.icm-institute.org	🟩🟨	🟩🟨	🟩🟨	🟩🟪	🟩🟨
Georgia - perfsonar.grena.ge	🟩	🟩	🟩	🟩	🟨🟩
Germany - psbrix.rrze.uni-erlangen.de	🟩	🟩	🟩	🟩	🟩

Microdep integration with perfSONAR (1/2)

- Microdep is a packet loss analysis and visualisation tool
  - Spotting packet gaps, micro failures, ~10 packets loss
  - Using 100 packet/sec probes
  - Traceroutes and ICMP response monitoring
- Realtime event analysis:
  - Packet-loss (gaps)
  - Queues (jitter)
  - Route failures and changes (traceroute)
  - Joint event anomality and alarms (ELK and ML)

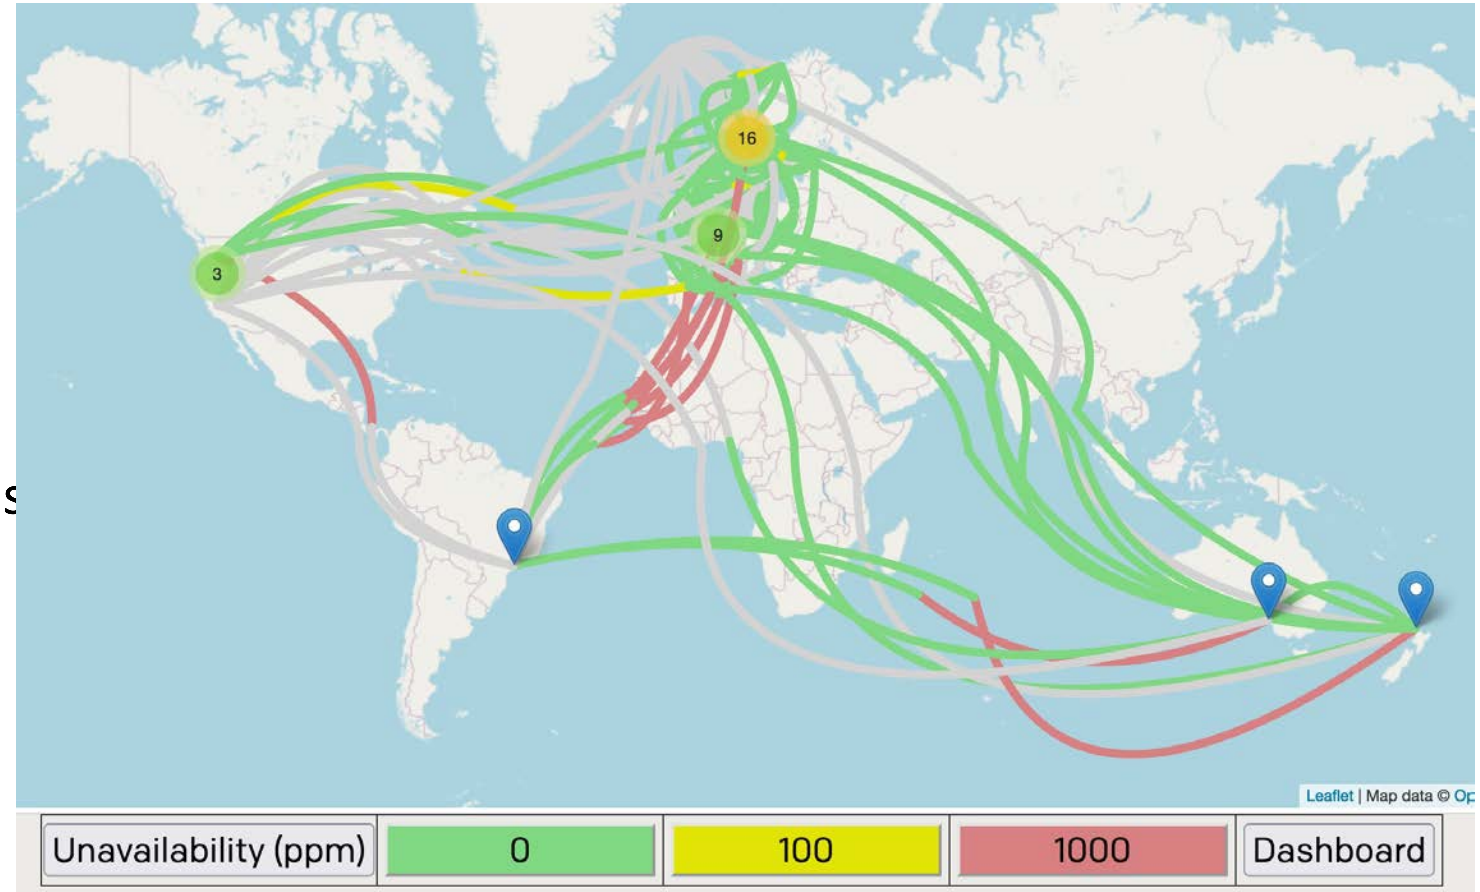

Microdep integration with perfSONAR (2/2)

- Using perfSONAR to generate probes
  - OWAMP for paced packets
  - Traceroute
  - Rely on 2000+ public perfSONAR hosts
  - Use pSConfig and pScheduler
- Adding a data pipeline to
  - Analyse packet gaps
  - Store history for further anal
- Next steps:
  - Package and bundle with pS

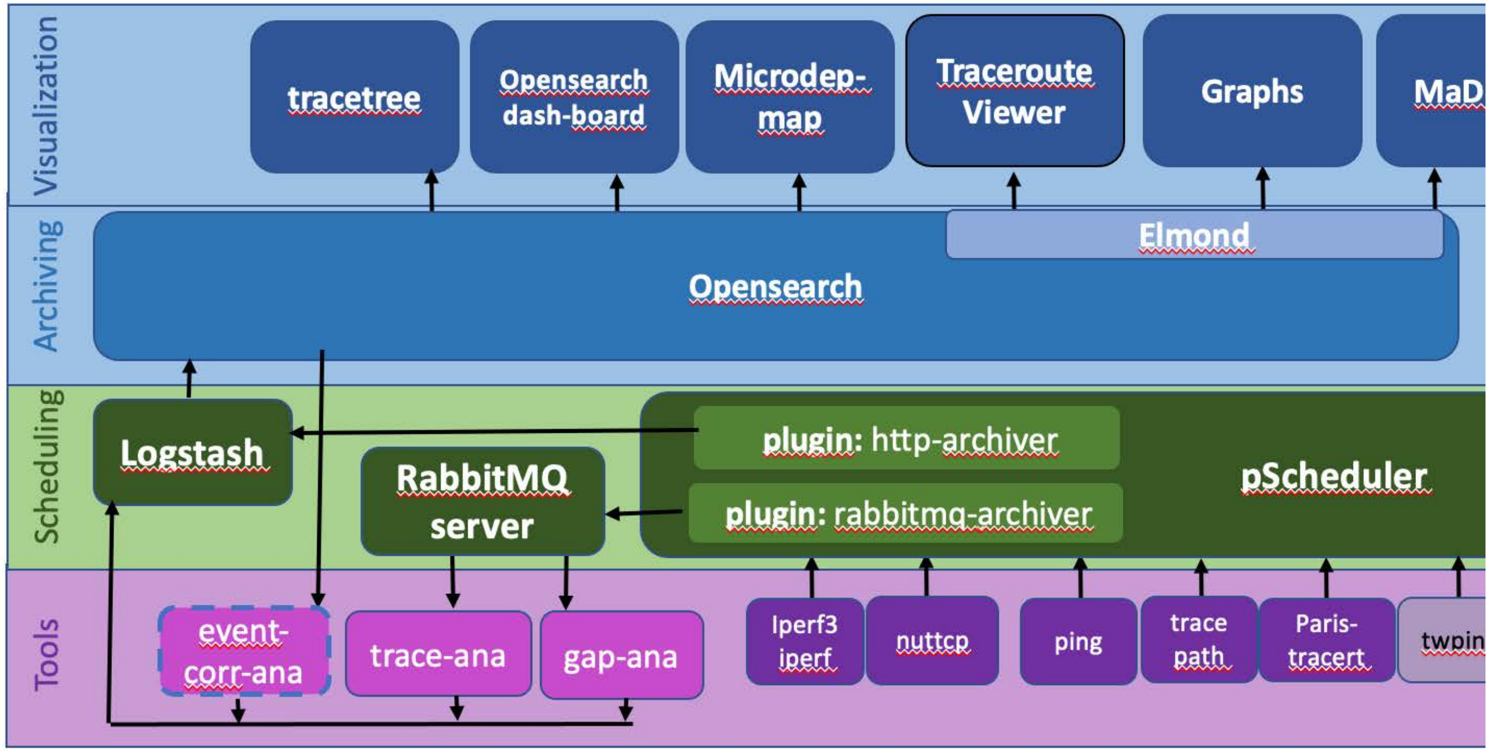

On-demand perfSONAR Graphical User Interface (psGUI) (1/2)

- GUI to drive perfSONAR / pScheduler
- Use case:
  - MaDDash setup, grids, regular measurements
  - Want to do a one off, on-demand additional test
  - List of pS nodes coming from pSConfig file, MaDDash grids
- GUI list all possible t

  and main parameter

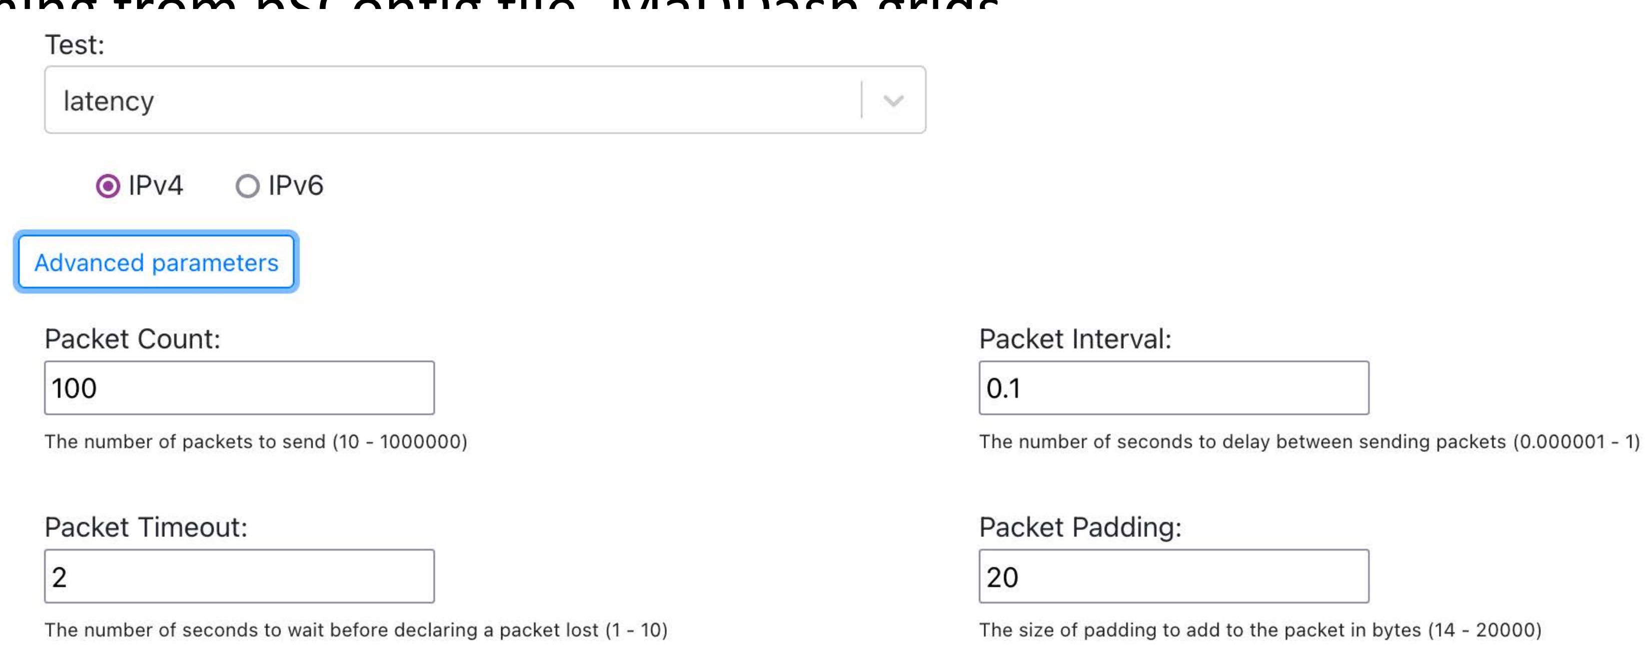

Test: latency

● IPv4 ○ IPv6

Packet Count: 100

The number of packets to send (10 - 1000000)

Packet Interval: 0.1

The number of seconds to delay between sending packets (0.000001 - 1)

Packet Timeout: 2

The number of seconds to wait before declaring a packet lost (1 - 10)

Packet Padding: 20

The size of padding to add to the packet in bytes (14 - 20000)

On-demand perfSONAR Graphical User Interface (psGUI) (2/2)

- Results:

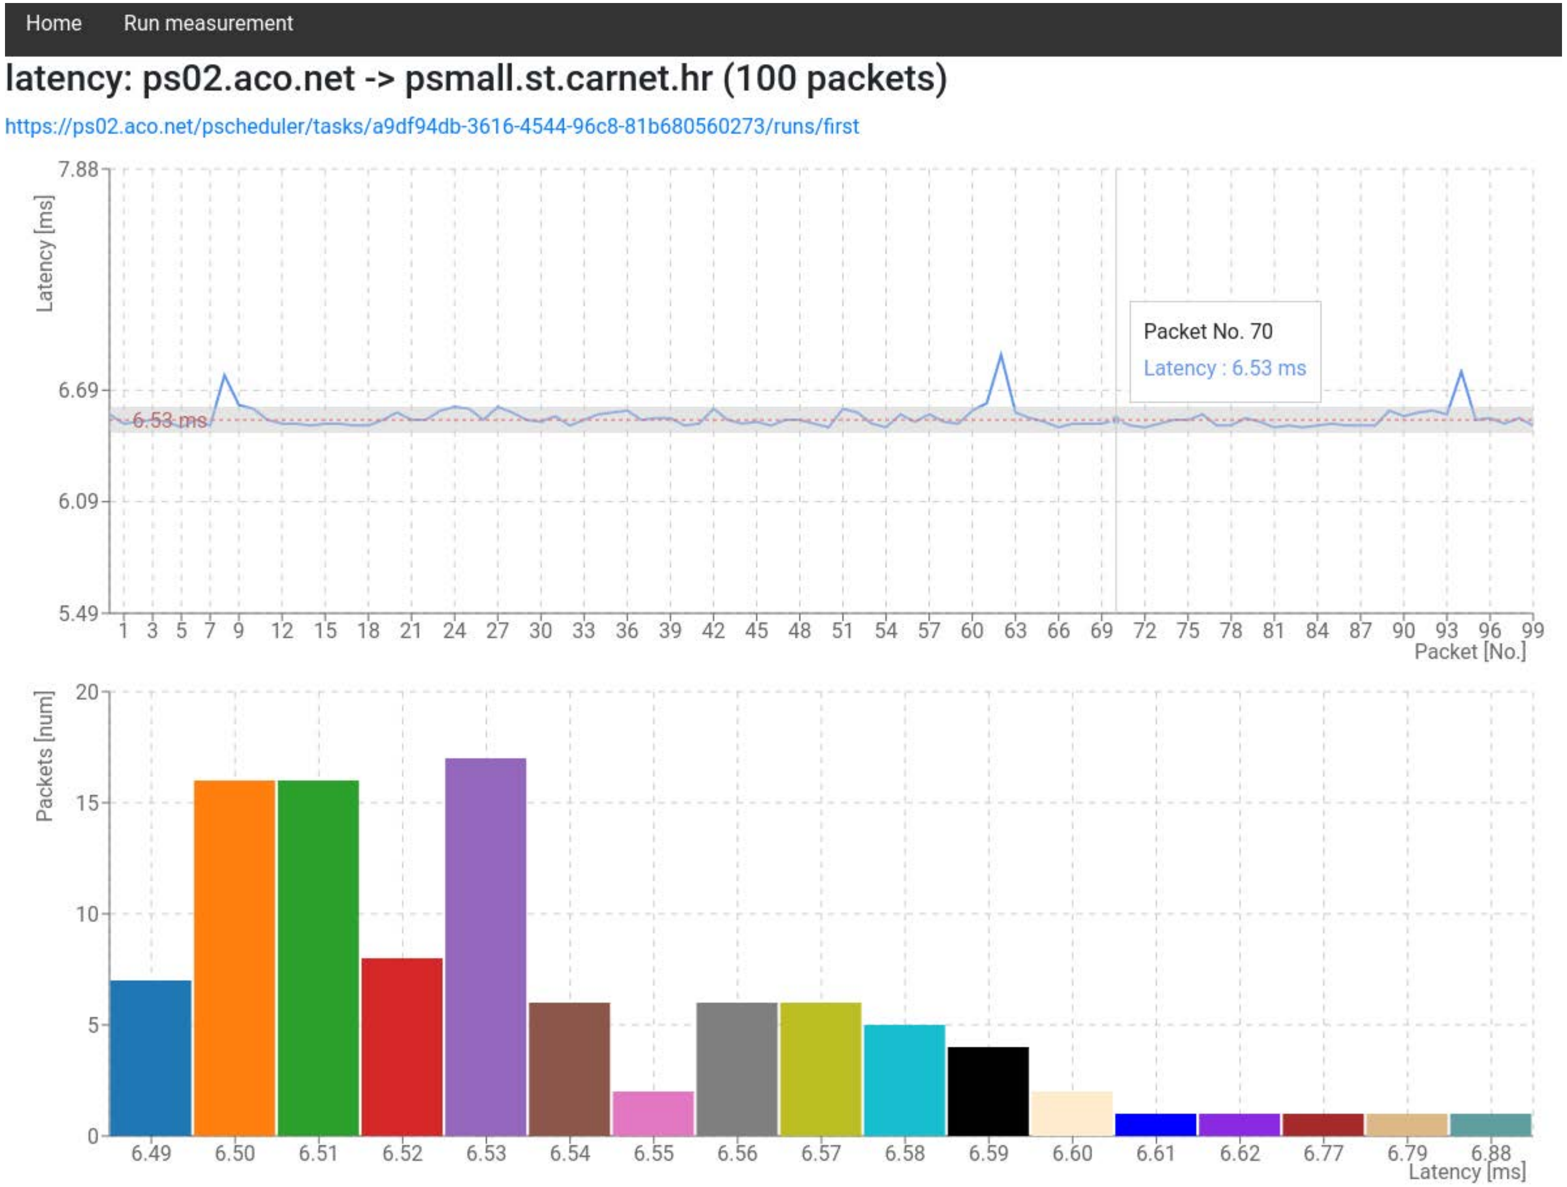

latency: ps02.aco.net -> psmall.st.carnet.hr (100 packets)

https://ps02.aco.net/pscheduler/tasks/a9df94db-3616-4544-96c8-81b680560273/runs/first

- Packaged as a Docker Image to be built
- https://github.com/perfsonar/psgui/

GÉANT

Thank You

Contact: perfsonar@lists.geant.org

www.geant.org

Co-funded by
the European Union

Cartoon text courtesy of textstudio.com