

# How to make the most of your small node?

PMP Workshop 2019

**Antoine Delvaux**  
*perfSONAR Service Manager*

TNC19, Tallinn, Estonia, 16/06/2019

# Session Outline

- Using the perfSONAR GUI
  - Check your setup
  - Adding new regular measurements
  - Viewing results
- Add your node to another mesh
  - Using the CLI
  - Checking the schedule
- Troubleshoot an end to end path
  - pScheduler on the CLI
  - Finding useful perfSONAR nodes

## This is a hands-on session

- You will need some tools on your laptop
  - Web browser
  - SSH client
- You will need some details of your PMP perfSONAR node
  - Hostname / IP address
  - Login credentials to webadmin
  - Login credentials to ssh session (password or ssh key)
- If you don't have access to a perfSONAR device you can use: <https://psmall-poz2.man.poznan.pl>

# Using the perfSONAR GUI



## Accessing the perfSONAR GUI

- Go to the perfSONAR toolkit web GUI:
  - <https://your.perfSONAR.node/>
- Public page: for anyone to see
  - Check if everything looks correct
    - Services status
    - IP addresses (v4 and v6) and corresponding hostname
    - Do you see a table of test results at the bottom?
- Private configuration pages
  - Login with your webadmin username and password
    - pswwebadmin

## Check setup

- Check admin information
  - That is sent to the public Lookup Service / Service Directory
  - Better **not** to provide personal contact information
  - Metadata: advertise how others can use your node
  - Communities: tags, enable filtering
- Check NTP setup
  - Best to use **your own NTP servers** if you have
- Automatic updates should be enabled
  - Best way to keep everything up to date and in sync
  - Automatically set by our Ansible playbooks

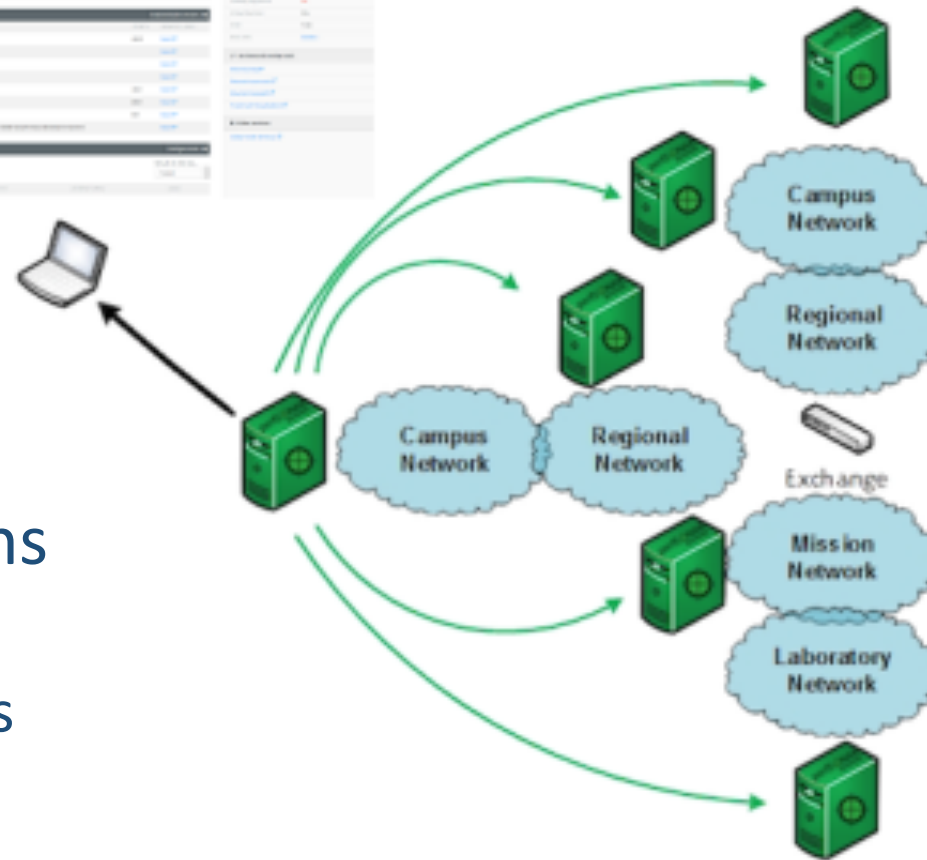


# Adding new regular measurements



Network	IP	Port	Protocol	Measurement
Campus Network	10.0.0.1	80	TCP	Web
Regional Network	10.0.0.2	80	TCP	Web
Mission Network	10.0.0.3	80	TCP	Web
Laboratory Network	10.0.0.4	80	TCP	Web

- We're combining **mesh**
- and **island** deployment patterns
  - PMP mesh
  - Your own ad-hoc measurements



# Adding new regular measurements

- Authentication required
- Navigating through the test setup page
  - Test type
  - Advanced parameters: change only if you know
  - Test members
    - Can browse Lookup Service, filter by community
- Adding new hosts to existing tests
- Adding new tests to existing hosts



## What and where to test?

- Know your users
- Know your big data flows
- This will help you identify destinations you want to test to
- A few regional destinations
- A few intercontinental destinations
- Test for
  - Packet drops and latency: always, inexpensive
  - Throughput: lightly and try to get in touch with remote admin
  - The perfSONAR toolkit will always add trace tests

# How to identify the perfSONAR targets?

- Browse the Service Directory (Global node directory link)
  - <http://stats.es.net/ServicesDirectory/>
  - Look at map
  - Filter by community
    - Project, Networks, Research groups
    - Interface speed
    - Any other community you have setup
- Get in touch with network admins on the other end
- When you have a target, check it's health
  - Toolkit public webpage
  - Remotely check schedule
- Try a few tests on the CLI first



# Participating to another mesh



## Another mesh, which mesh and why?

- You have other pS devices and your own mesh
- Your organisation is part of a group already using pS
- You just want to try it out for yourself



# Accessing your node through ssh

- ssh client
- Password authentication
- Better: key authentication
- Access policy and ssh configuration
  - ssh root access is always disabled for small nodes
  - fail2ban is running on the small node
    - Keep bad parties away
    - But can lock you out too if you don't remember password

## It's all pSconfig ... on the CLI

- Checking current setup

```
geantadmin@perfsonar:~$ psconfig remote list
=== pScheduler Agent ===
[
  {
    "url" : "https://pmp-central.geant.org/pscfig-psmall.json",
    "configure-archives" : true
  }
]
```

- Adding a new mesh

```
psconfig remote add URL.json
Psconfig remote add -configure-archives URL.json
```

## After adding

- Check status

```
psconfig pscheduler-stats
```

- Check log files
  - /var/log/perfsonar/psconfig-pscheduler-agent.log
  - Check for last run
  - Check for warnings and errors
- Make sure remote URL is reachable from the small node

```
curl -k URL.json | jq
```



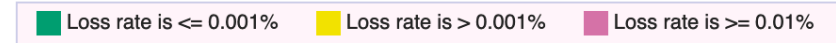
## Please note that...

- A small node has limited resources
- It is intended as a first contact, for you, with perfSONAR
- We can help you design your own pS deployment plan if you wish to go further


# Use Cases: real examples

# Dashboard: a quick overview to spot issues

## PMP - IPv4 OWD - Loss



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

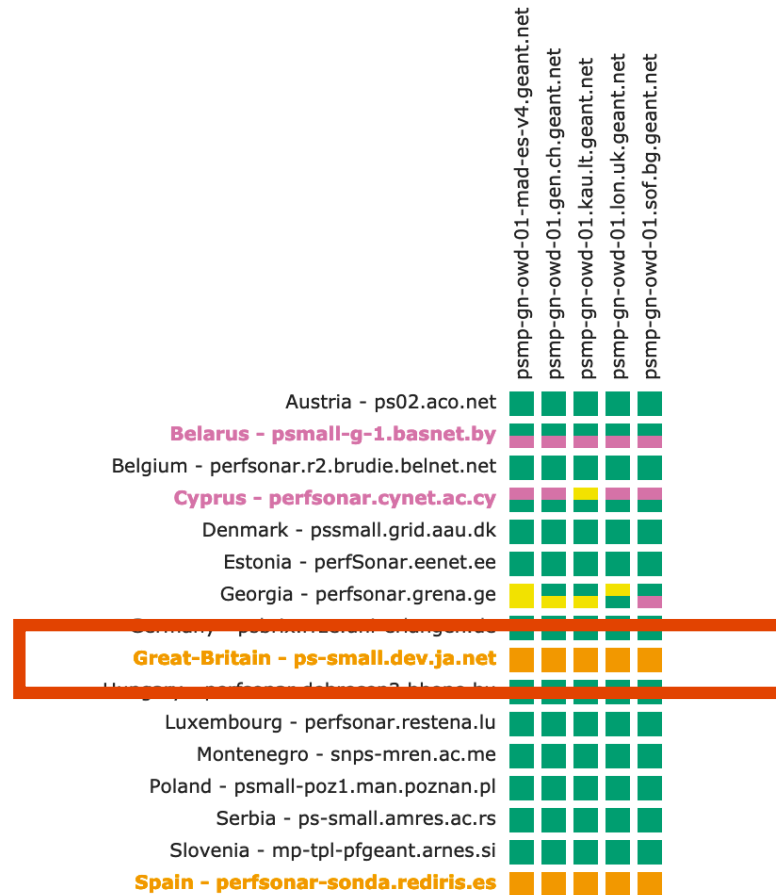


# A disconnected node

## PMP - IPv4 OWD - Loss

■ Loss rate is  $\leq 0.001\%$  ■ Loss rate is  $> 0.001\%$  ■ Loss rate is  $\geq 0.01\%$

⚠ Found a total of 4 problems involving 4 hosts in the grid



# Worrying pattern on the dashboard

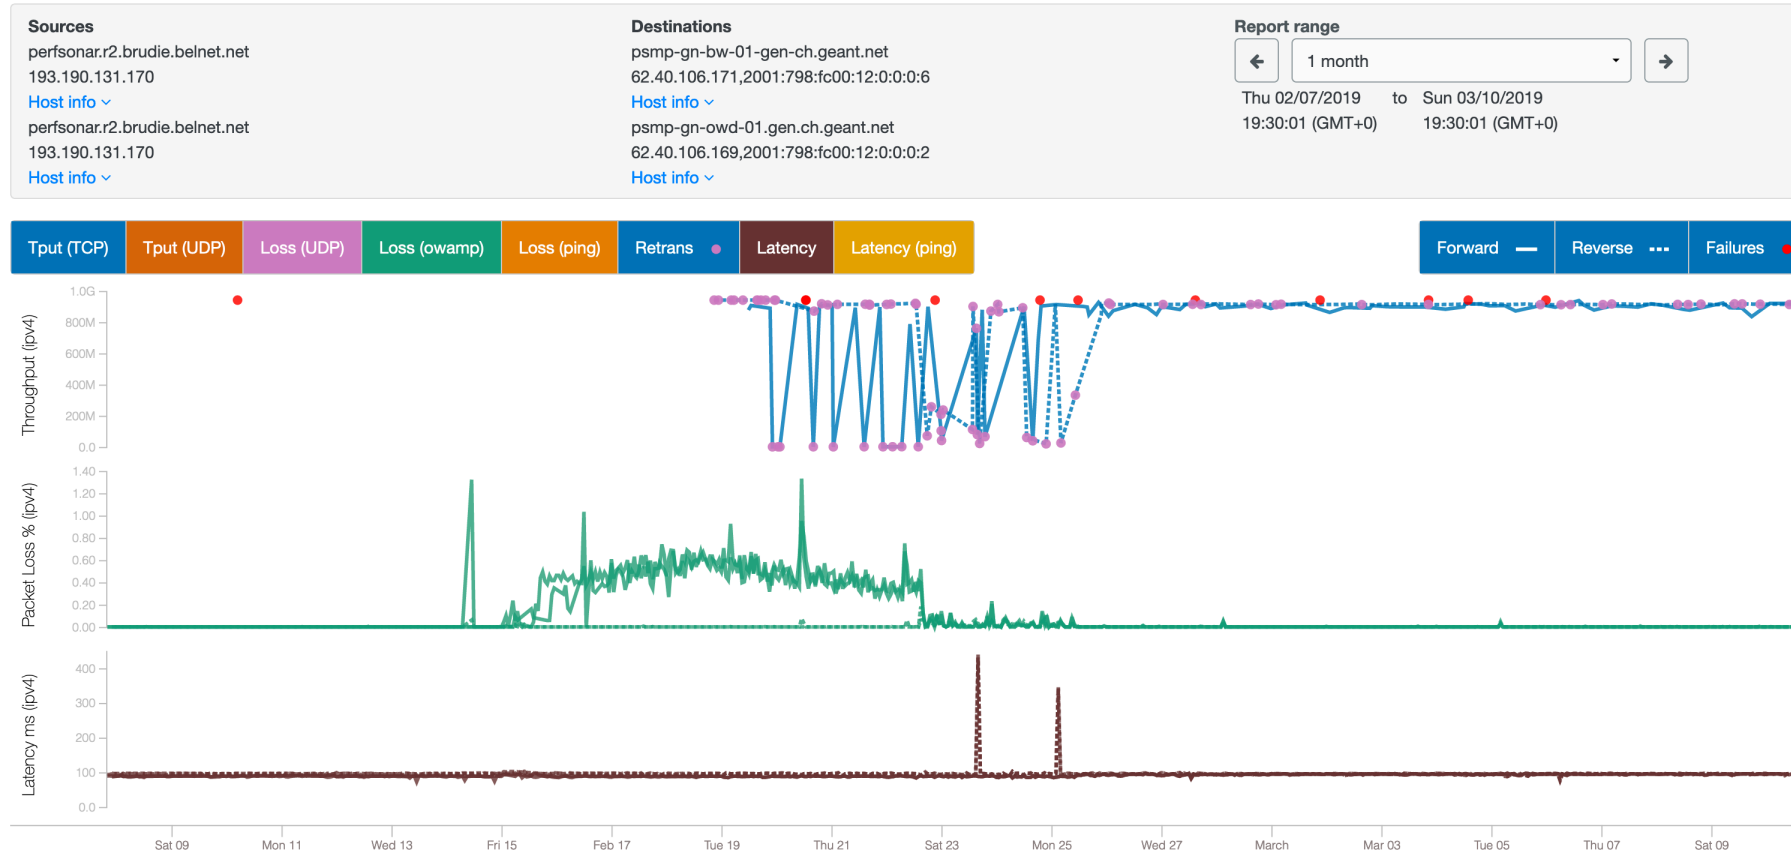
## PMP - IPv4 OWD - Loss

Loss rate is  $\leq 0.001\%$     Loss rate is  $> 0.001\%$     Loss rate is  $\geq 0.01\%$


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



# Historical graphs confirm the issue (to Geneva)



# To London and to all 5 GÉANT MP







# What happened?

## Contact with Network Operating Center

- They acknowledge a network issue on Feb 15<sup>th</sup>
  - Fiber cut or device / network card failure
  - That's the first peak in packet drops
- Traffic was rerouted through backup link, half capacity
- During the day, the main link was restored
  - But packet drop still present
  - During about 1 week
- Situation back to normal on Feb 23<sup>rd</sup>


# Confirmed by looking at link utilisation (main/backup)



## NREN access links


- Main link: LAG
  - 2 x 10 Gbps Link Aggregation
- Backup link: LAG
  - 2 x 10 Gbps ... or only 1 x 10 Gbps?
- Situation confirmed by looking at link utilisation and interface packet counters

# Traffic on each links of the main link (LAG)




■ Inbound Current: 2.45 G Average: 4.70 G Maximum: 10.91 G ■ Peak ■ 95th % In (8494.81 Mbps In)  
■ Outbound Current: 8.59 G Average: 9.24 G Maximum: 13.35 G ■ Peak ■ 95th % Out (12590.09 Mbps Out)

Full LAG



■ Inbound Current: 1.27 G Average: 3.12 G Maximum: 8.38 G ■ Peak  
■ Outbound Current: 4.38 G Average: 5.31 G Maximum: 9.71 G ■ Peak

Link 1







■ Inbound Current: 1.19 G Average: 1.59 G Maximum: 5.57 G ■ Peak  
■ Outbound Current: 4.21 G Average: 3.94 G Maximum: 6.66 G ■ Peak

Link 2

# Interface counters on main link: Link 1 and Link 2 (as seen from GÉANT side)

## Link 1




Packet discards out  
Saturation


Packet errors in  
Link/Card issue



# And on the backup link (discards, saturation)





Inbound	Current: 8.64 G	Average: 1.21 G	Maximum: 10.46 G	Peak	95th % In (6673.08 Mbps In)
Outbound	Current: 339.32	Average: 190.87 M	Maximum: 8.58 G	Peak	95th % Out (177.14 Mbps Out)



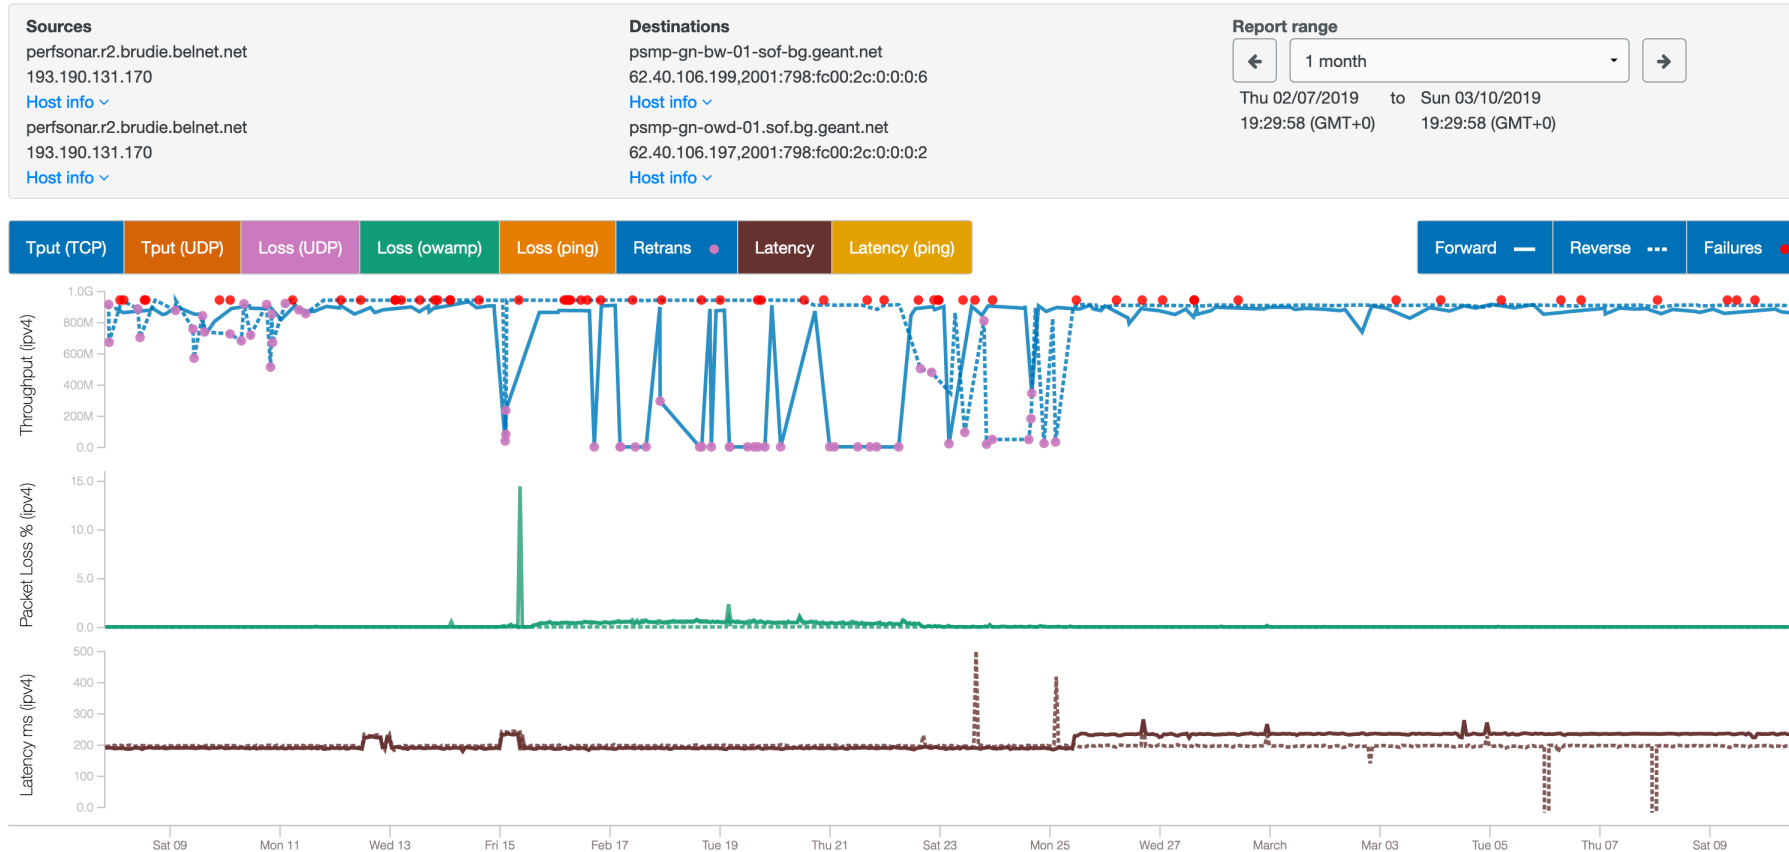
Discards In	Current: 0.00	Average: 0.00	Maximum: 0.00
Errors In	Current: 0.00	Average: 0.00	Maximum: 0.00
Discards Out	Current: 0.00	Average: 924.85 m	Maximum: 82.90
Errors Out	Current: 0.00	Average: 0.00	Maximum: 0.00

# After: traffic balanced on main and backup links





# But asymmetric route shown by perfSONAR



# Troubleshooting and end to end path

# Identifying measurement points along the path

- Start from a trace
- Identify the networks
  - Reverse DNS
  - ASN
- Look for pS measurement points in those networks
  - Use Service Directory
  - <http://stats.es.net/ServicesDirectory/>

# Divide and conquer!

- Divide the path in smaller chunks
  - If you cannot find a pS MP exactly on the path, try to find the closest one
- Test the smaller paths, one by one
  - For latency and **packet loss**
  - Then for throughput
- Testing segments only is usually not useful
  - Too small
  - Packet loss doesn't affect throughput that much

## Running some tests

- Latency (OWAMP)

```
pscheduler task latency --dest psmg-gn-bw-01-sof-bg.geant.net
```

- Throughput

```
pscheduler task throughput --dest psmg-gn-bw-01-sof-bg.geant.net
```

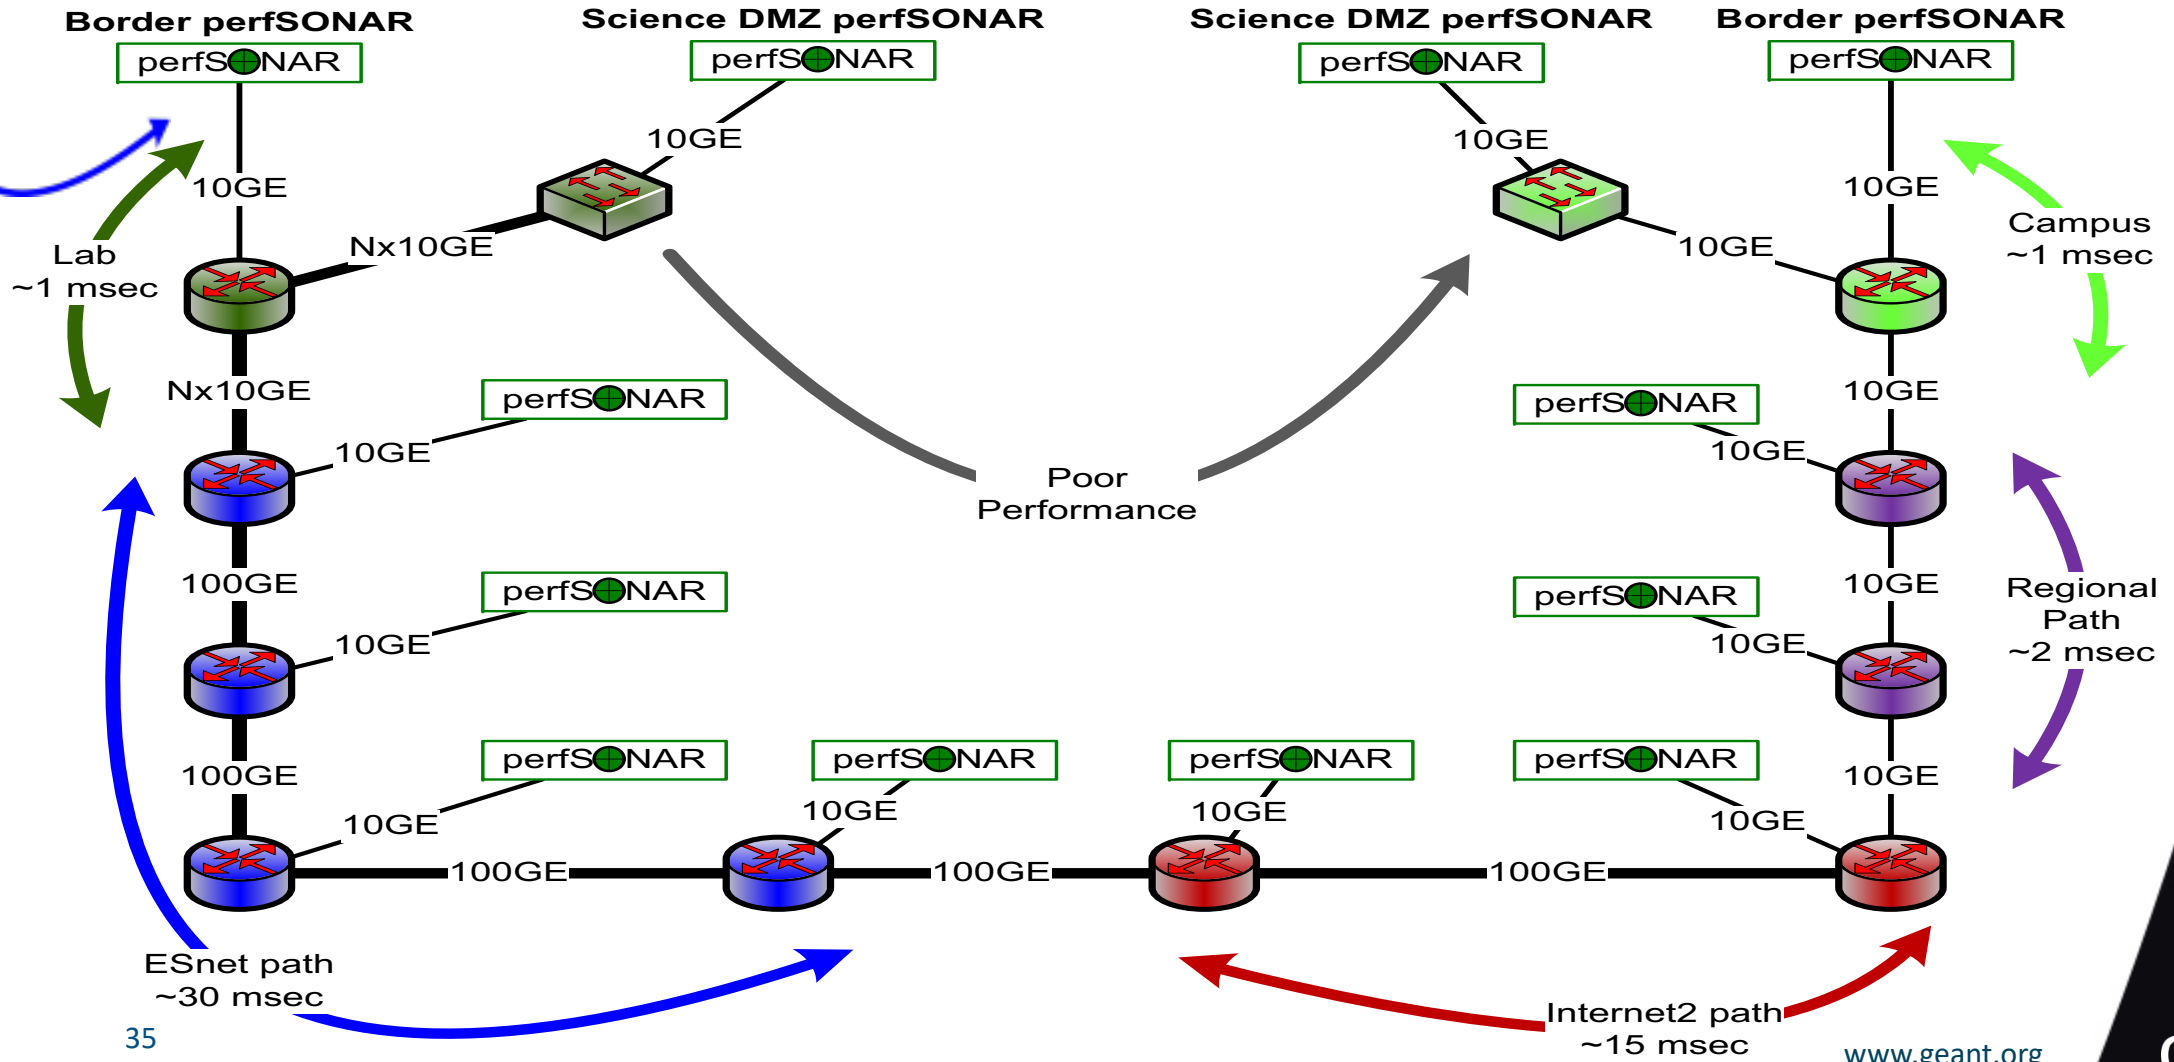
- Trace

```
pscheduler task trace --dest psmg-gn-bw-01-sof-bg.geant.net
```

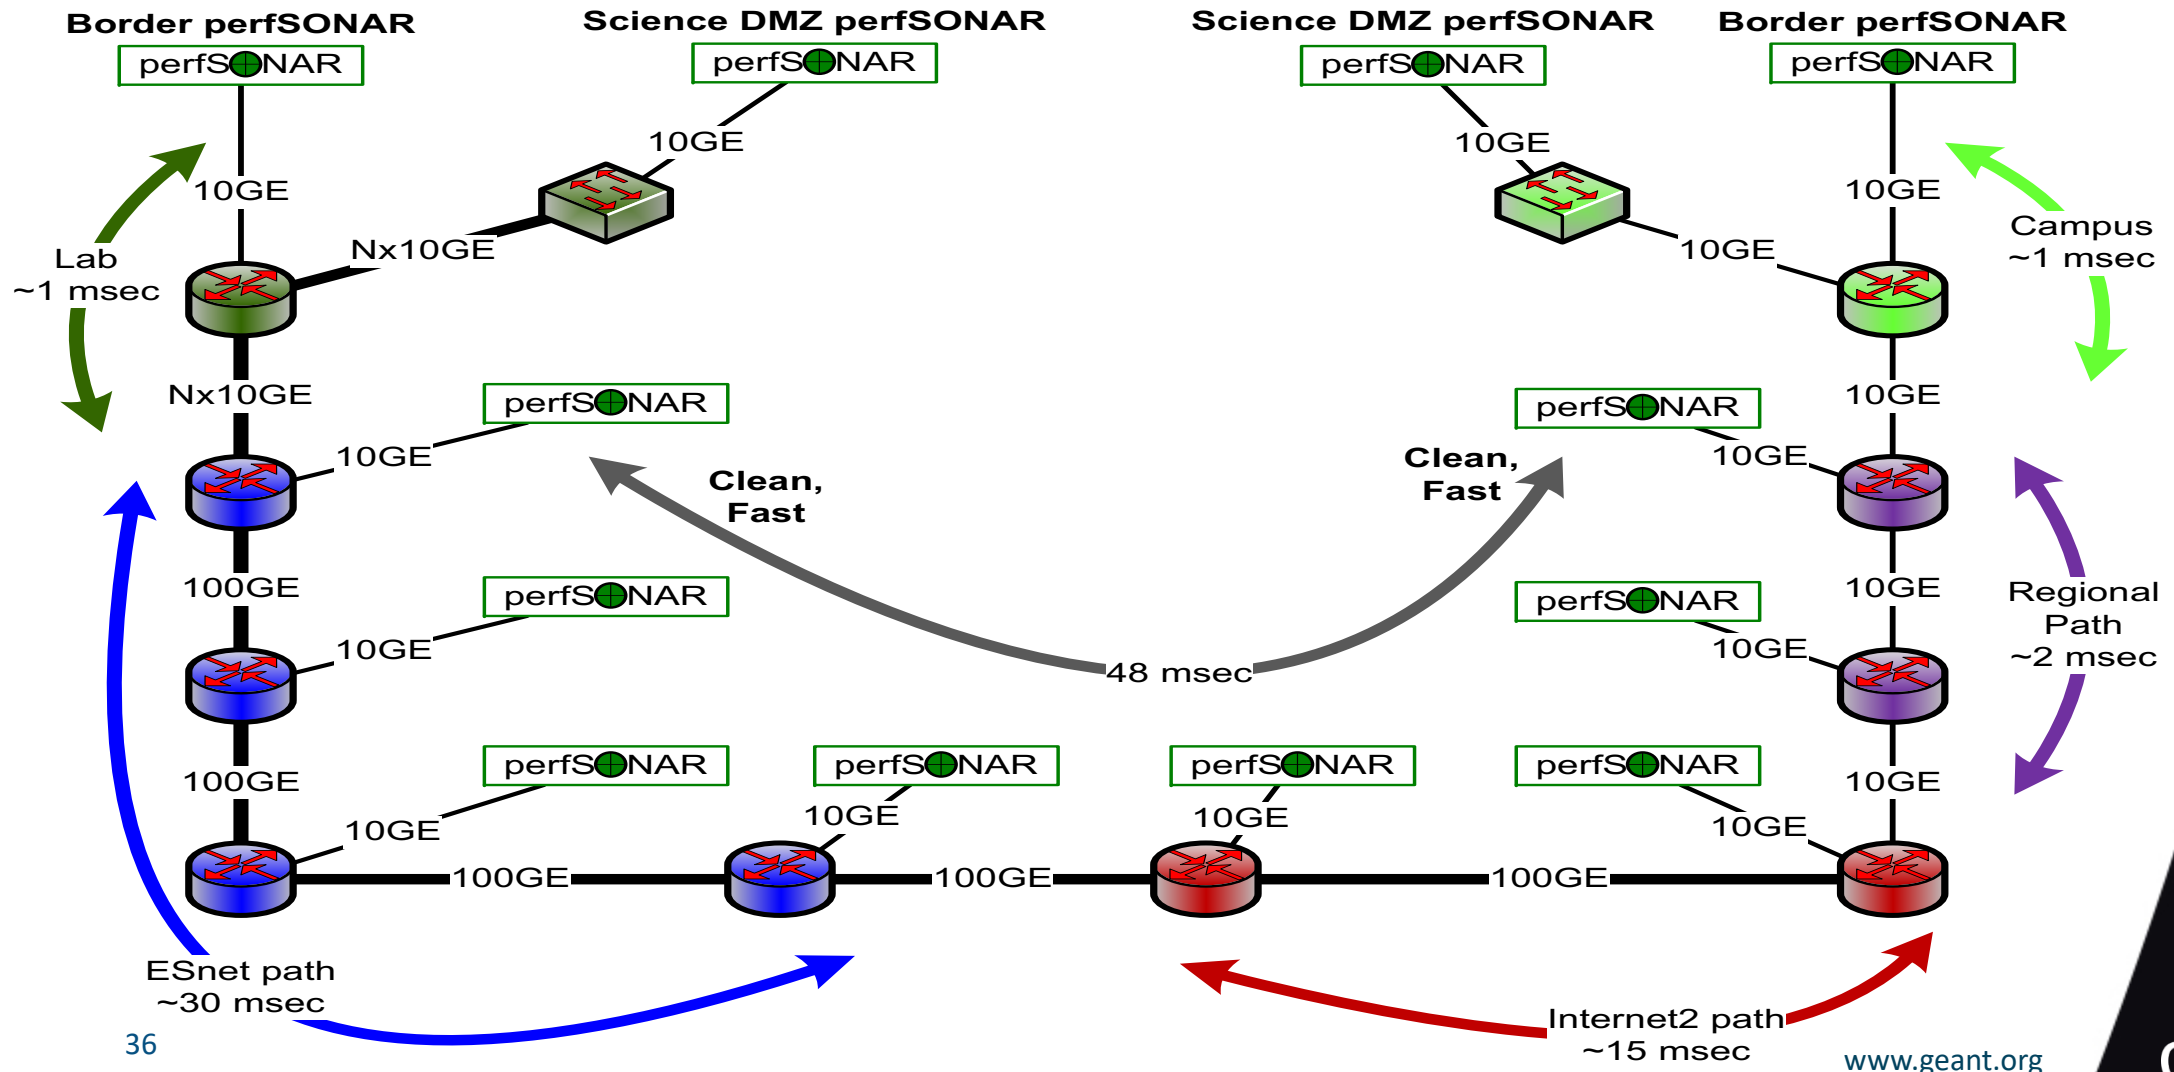
- Getting help

```
pscheduler --help  
pscheduler task latency --help  
pscheduler --troubleshoot psmg-gn-bw-01-sof-bg.geant.net
```


# End to End path poor performance



# But middle path clean



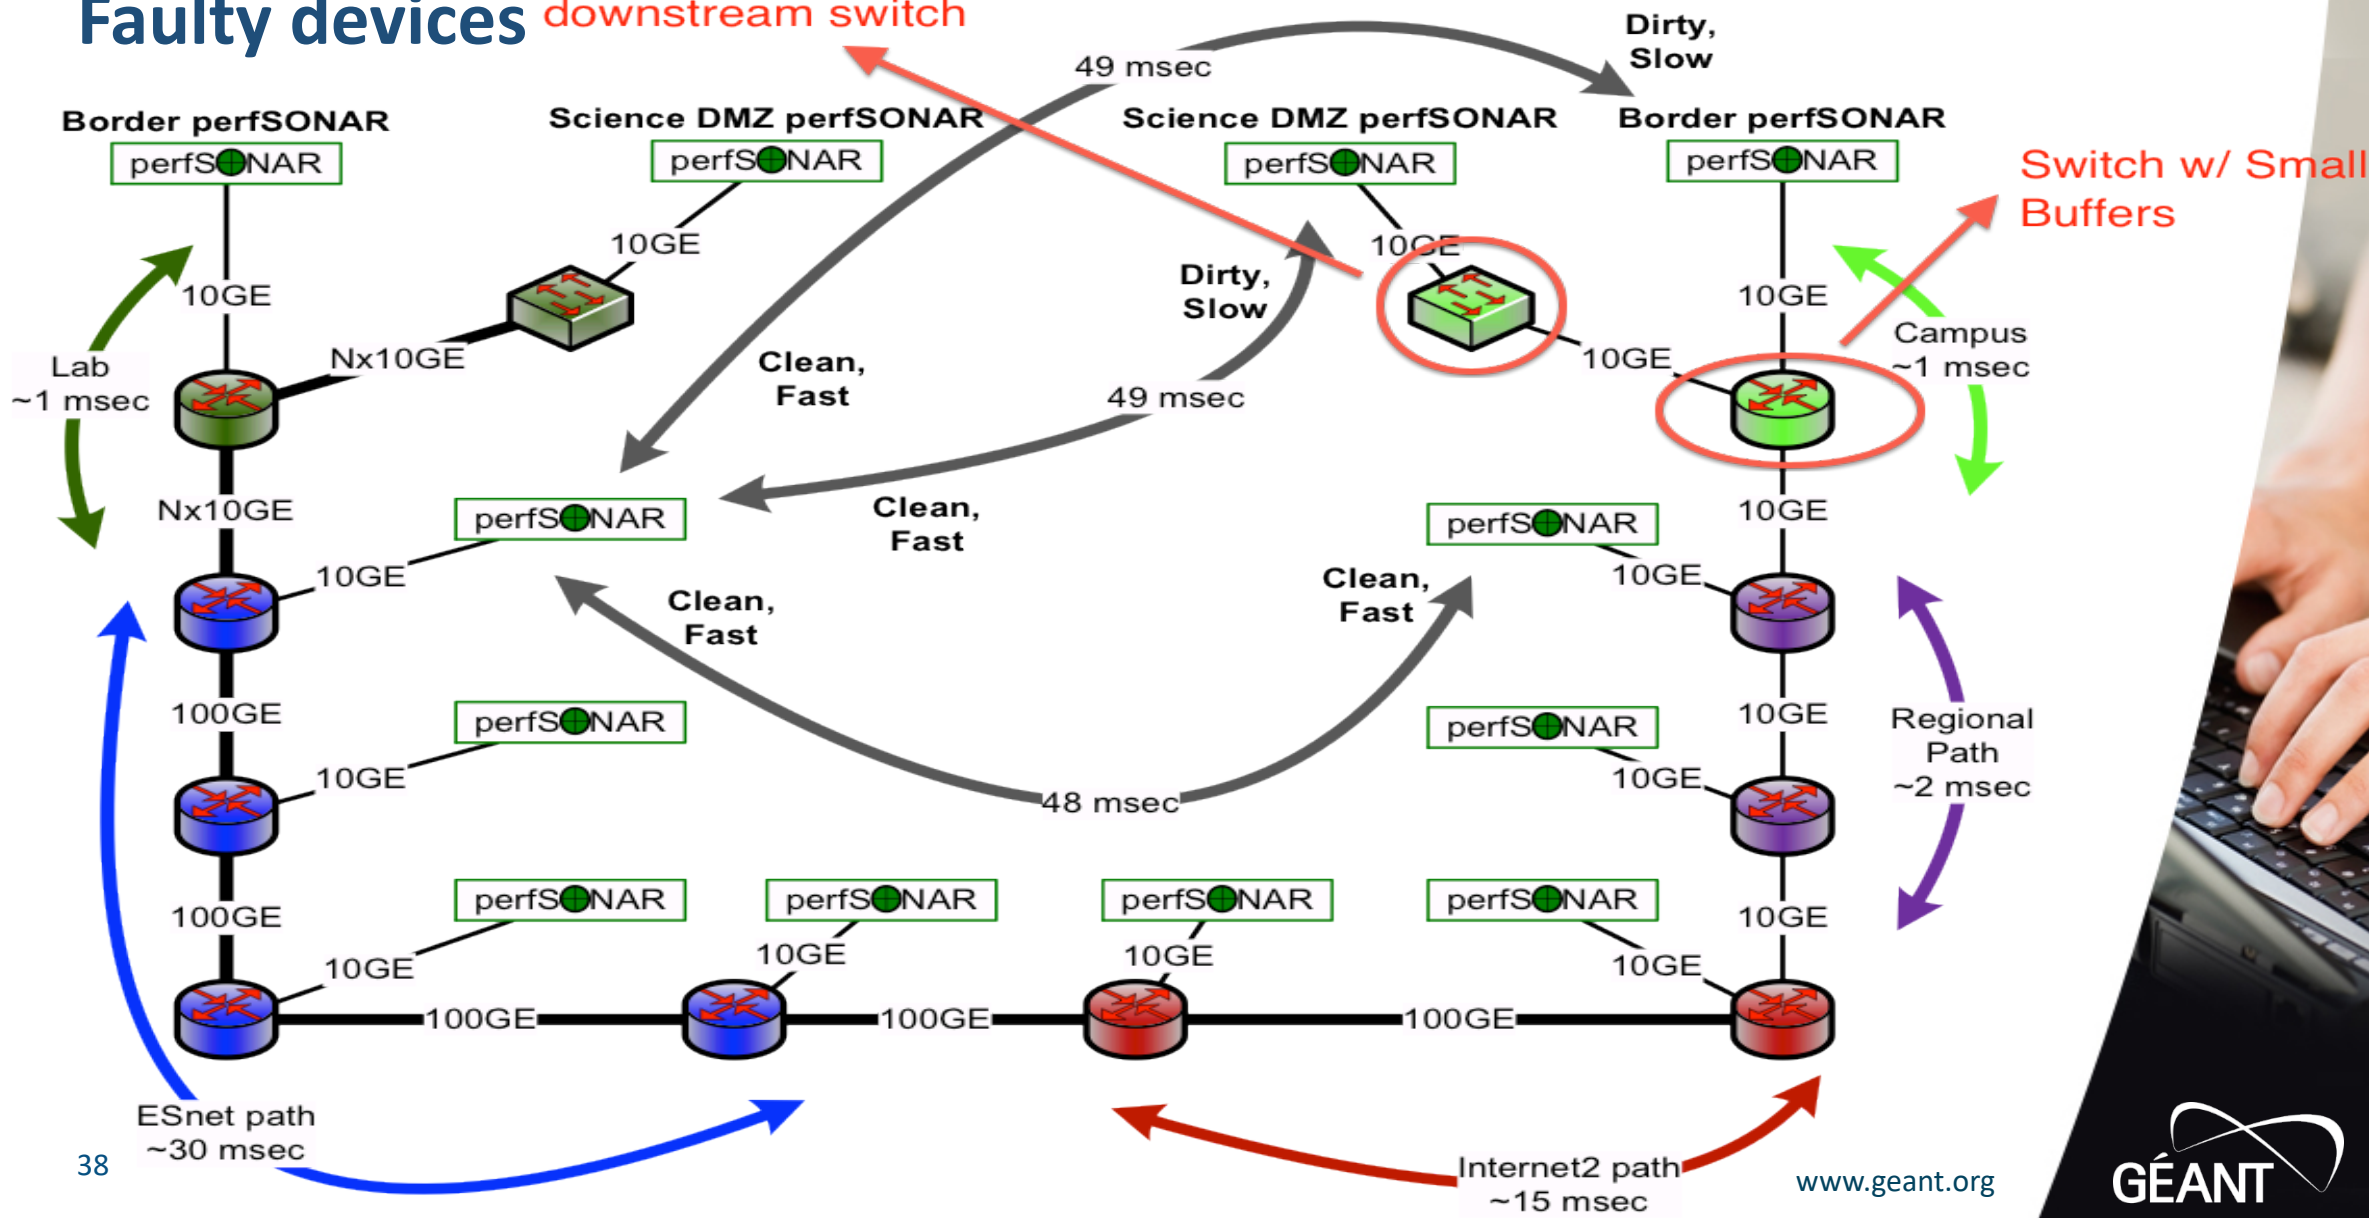
# Longer path slow





# Faulty devices

Impacted by downstream switch



## Back to the future...

- perfSONAR team currently working on a new service
- Called **pShooter**
  - Use of DNS records to advertise the closest perfSONAR MP next to a network device
  - Need each network to make useful advertisement
  - Will be provided as a webservice, to be called through an API
- Come again next year...

# To Conclude



## Maintaining your node

- Check your node is up and reachable
  - Add alerting for that is you want to (Nagios or else)
- Auto-updates will not upgrade the Linux kernel
  - Manual reboot required from time to time
- Check the PMP dashboard and your node color
- **Orange**: unreachable!
  - Node need to be verified
- **Yellow, Purple**: performance degradation
  - Might be worth investigating the network path towards GÉANT

## perfSONAR current development roadmap

- One release per year, looking to have more
- **4.2:** beta soon, final probably during summer
  - Disk-to-disk (GridFTP) plug-in
  - Measurement pre-emption (priorities)
  - Additional pSConfig utilities
  - Lookup Service improvements
- **4.3:** by end of 2019
  - Move to Python 3
- **4.4:** early 2020
  - New visualisation options (most probably, Grafana based)

## More resources...

- PMP dashboard:
  - <https://pmp-central.geant.org/maddash-webui/>
- Documentation
  - <https://docs.perfsonar.net>
- Youtube channel
  - <https://www.youtube.com/perfSONARProject/>
- Users lists
  - PMP list:
    - <https://lists.geant.org/sympa/info/perfsonar-smallnodes>
  - Global users list:
    - <https://lists.internet2.edu/sympa/info/perfsonar-user>
- GÉANT Consultancy and Expertise service
  - Contact us at [perfsonar@lists.geant.org](mailto:perfsonar@lists.geant.org)



# Thank you

Any questions?

[perfsnar-smallnodes@lists.geant.org](mailto:perfsnar-smallnodes@lists.geant.org)

[www.geant.org](http://www.geant.org)

