

# Sharing Data via perfSONAR

Andy Lake (ESnet)

2nd Telemetry and Data Workshop ▪ April 4, 2022

*perfSONAR is developed by a partnership of*



**ESnet**



**INDIANA UNIVERSITY**



# What is perfSONAR?

- An **open source 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



# What's the data look like?

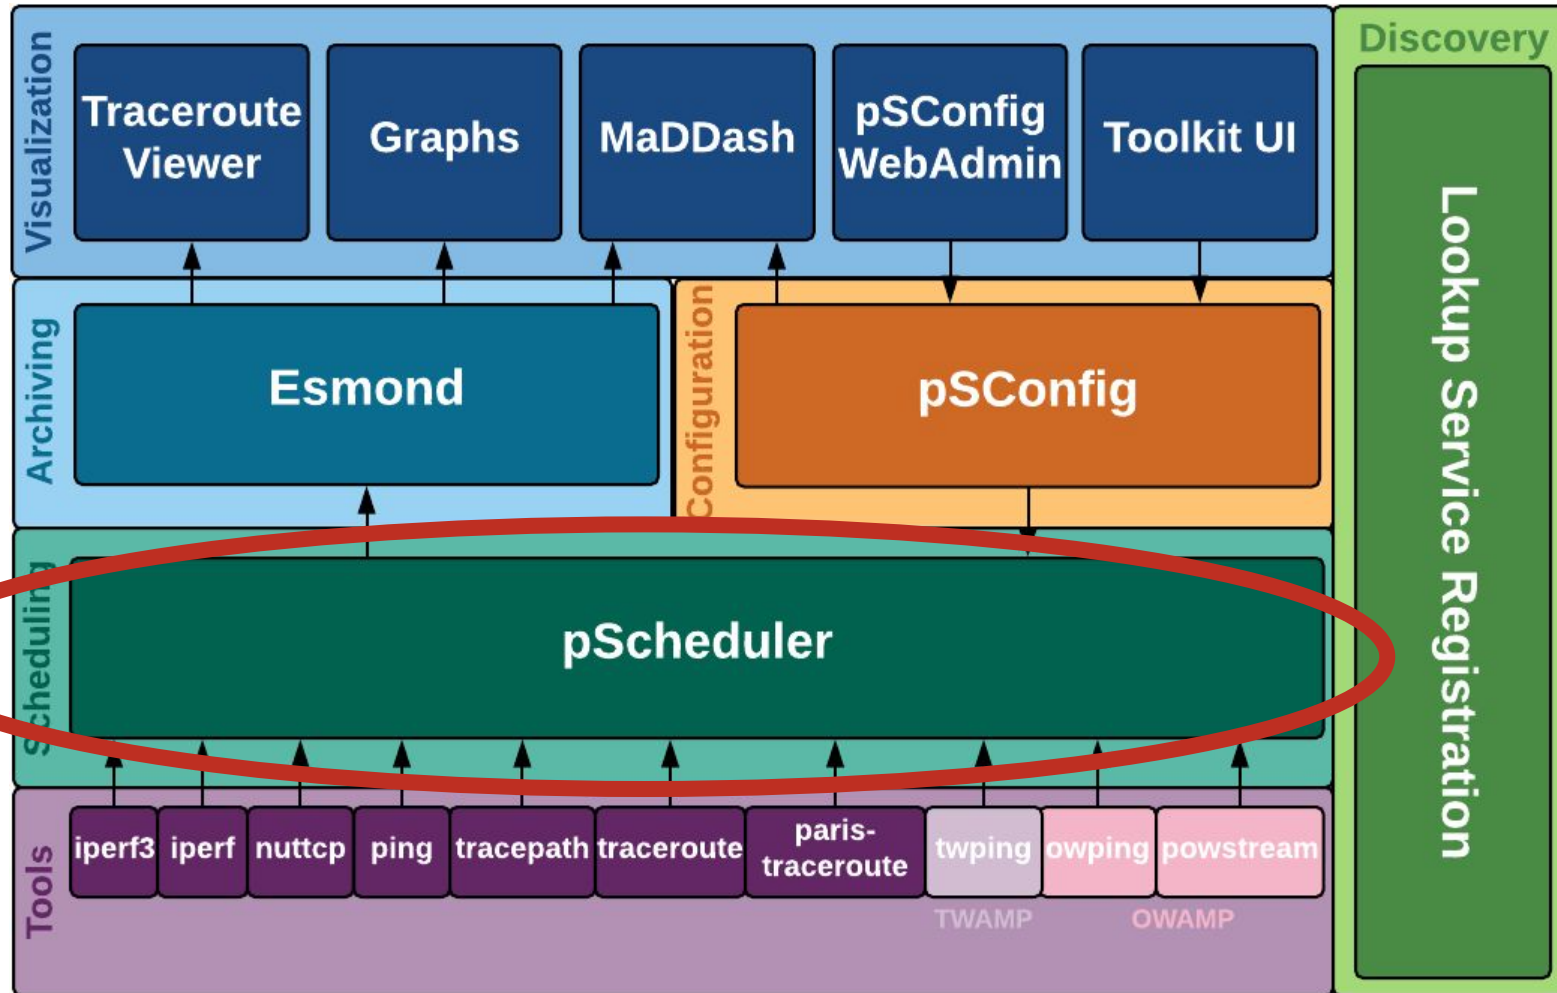
- Metadata that describes test parameters. Depending on type of test includes:
  - **IP addresses** - e.g. source and destination
  - **Scheduling information** - duration, testing interval, etc
  - **Test-specific info** - bandwidth, packet size, packet count, etc
- Values that describe the results of the test. Examples:
  - **Numeric values** - Throughout, packet loss, latency and derivatives
  - **Contextual information** - IPs returned by traceroute

# Comparing to other common types of Network Data

01	perfSONAR	<ul style="list-style-type: none"><li>• IP addresses map to perfSONAR hosts that are usually publicly known</li><li>• Access controls for determining who is allowed to run tests</li><li>• General approach is if results are published at all, they are public for everyone. Otherwise, don't publish or use ACLs.</li></ul>
02	Interface Counters	<ul style="list-style-type: none"><li>• Sourced from protocols like SNMP or maybe streaming telemetry</li><li>• Data maps to interfaces at point in network. May map to traffic from single organization or network.</li><li>• Statistics aggregated so at least in backbone case can't directly isolate user traffic (but this depends on network profile)</li><li>• Many networks willing to share this to a degree</li></ul>
03	Flow	<ul style="list-style-type: none"><li>• Highly granular information in raw form. Able to determine endpoint IPs, application, more.</li><li>• Networks often very sensitive about sharing this data. International privacy laws apply.</li><li>• Possible to do aggregations to satisfy needs, but requires careful consideration</li></ul>

# What's perfSONAR general sharing policy?

- **perfSONAR strongly encourages users to share measurement results as widely as possible**
  - Acknowledgement that perfSONAR data generally less sensitive than other types of data, but also recognize policies vary to want to give users adequate knobs
- Most of access control surrounds limiting who can run tests
  - Much of the concern is about over-testing not over-sharing
  - Common sense defaults - we aren't trying to build a network of UDP cannons
  - Generally if you don't want someone to see a result, don't let them run the test in the first place
  - Can use firewalls and other mechanisms to limit access to stored results



Most enforcement here

# pScheduler Limits System

- *Know Your Limits* by Mark Feit -  
<https://www.youtube.com/watch?v=eGK02oo8JAM>
- Detailed language for controlling tests based on requester IP:
  - Automated lists of common ranges here:  
[http://stats.es.net/sample\\_configs/pscheduler/](http://stats.es.net/sample_configs/pscheduler/)
  - Relatively straight-forward directives for examining common fields
  - jq for more complicated rules
  - Can provide defaults for parameters like bandwidth, set scheduling priorities, etc
- **DON'T OVERCOMPLICATE YOUR RULES** - Just because you can do something doesn't mean you should. The community benefits the more places we can test to easily.

# What about the archive?

- Current *esmond* archive does not have any way to control who can see which results beyond ACLs
- Archiver plugins gives option to use something else with more controls
- perfSONAR 5.0 technically will change this with move to Elastic/OpenSearch
- perfSONAR plans to maintain its open access policy, but possibility there for brave souls





# What about the lookup service?

- perfSONAR servers have the ability to register their existence to a central lookup service (LS).
- Ability to opt-in or opt-out by enable/disable service...or based on reachability to LS
- Again, this is all or nothing approach, we do not enforce fine-grained access controls on LS
- Ability to run private LS for large private deployments




<http://stats.es.net/ServicesDirectory/>

# Private Deployments

- We love when data can be shared, but sometimes you want to test performance of an internal network or similar.
- ACLs and routing policy are your friends here, don't need to rely on perfSONAR
- Possible to even do things like setup a private lookup service
- Multiple examples of this in commercial and government sectors

# Summary


- perfSONAR encourages users to share data and allow testing as much as common sense and policy allow.
- perfSONAR provides a rich set of access controls at the test scheduling and execution layer to limit which tests can be run
- perfSONAR has been around a long-time and number of cases where users have combined application limits with traditional approaches using ACLs to meet needs.



# Questions and Answers

Question and answer icon by iconosphere from The Noun Project

# 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

