Introduction to

perfSONAR

Network Startup Resource Center

ESnet
ENERGY SCIENCES NETWORK

U.S. DEPARTMENT OF ENERGY
Office of Science

UNIVERSITY OF OREGON

NSRC
Network Startup Resource Center

Outline

- What is perfSONAR trying to solve?
  - Multi-domain, complex environments
  - Hard vs. software failures
- PerfSONAR as a solution
- Benefits
- Resources
- Collaborators

Problem Statement

The global Research & Education network ecosystem is compris of hundreds of international, national, regional and local-scale networks.

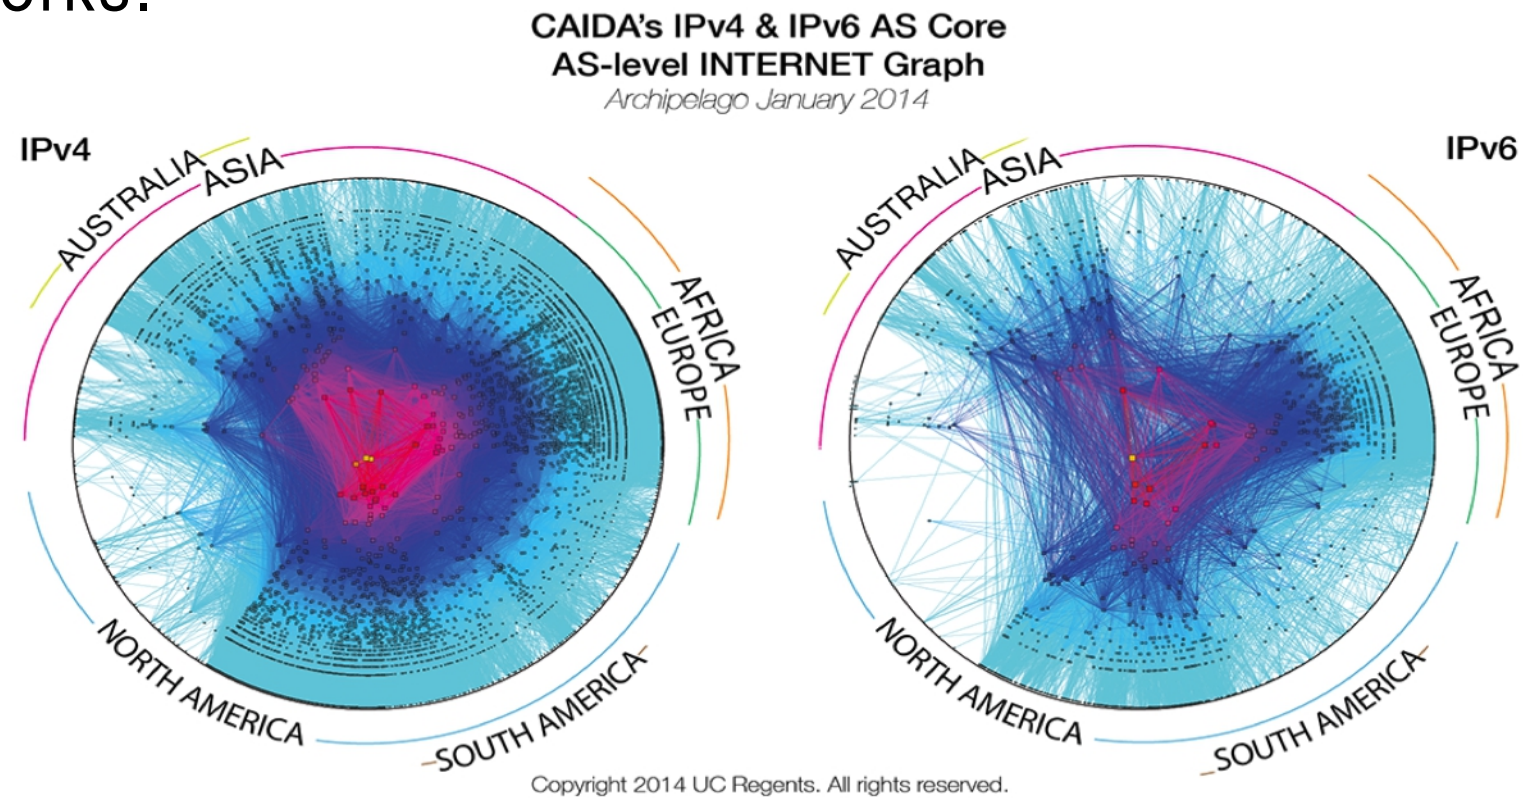

Problem Statement

While these networks all interconnect, each network is owned and operated b
separate organizations (called “domains”) with different policies, customers,
funding models, hardware, bandwidth and configurations.

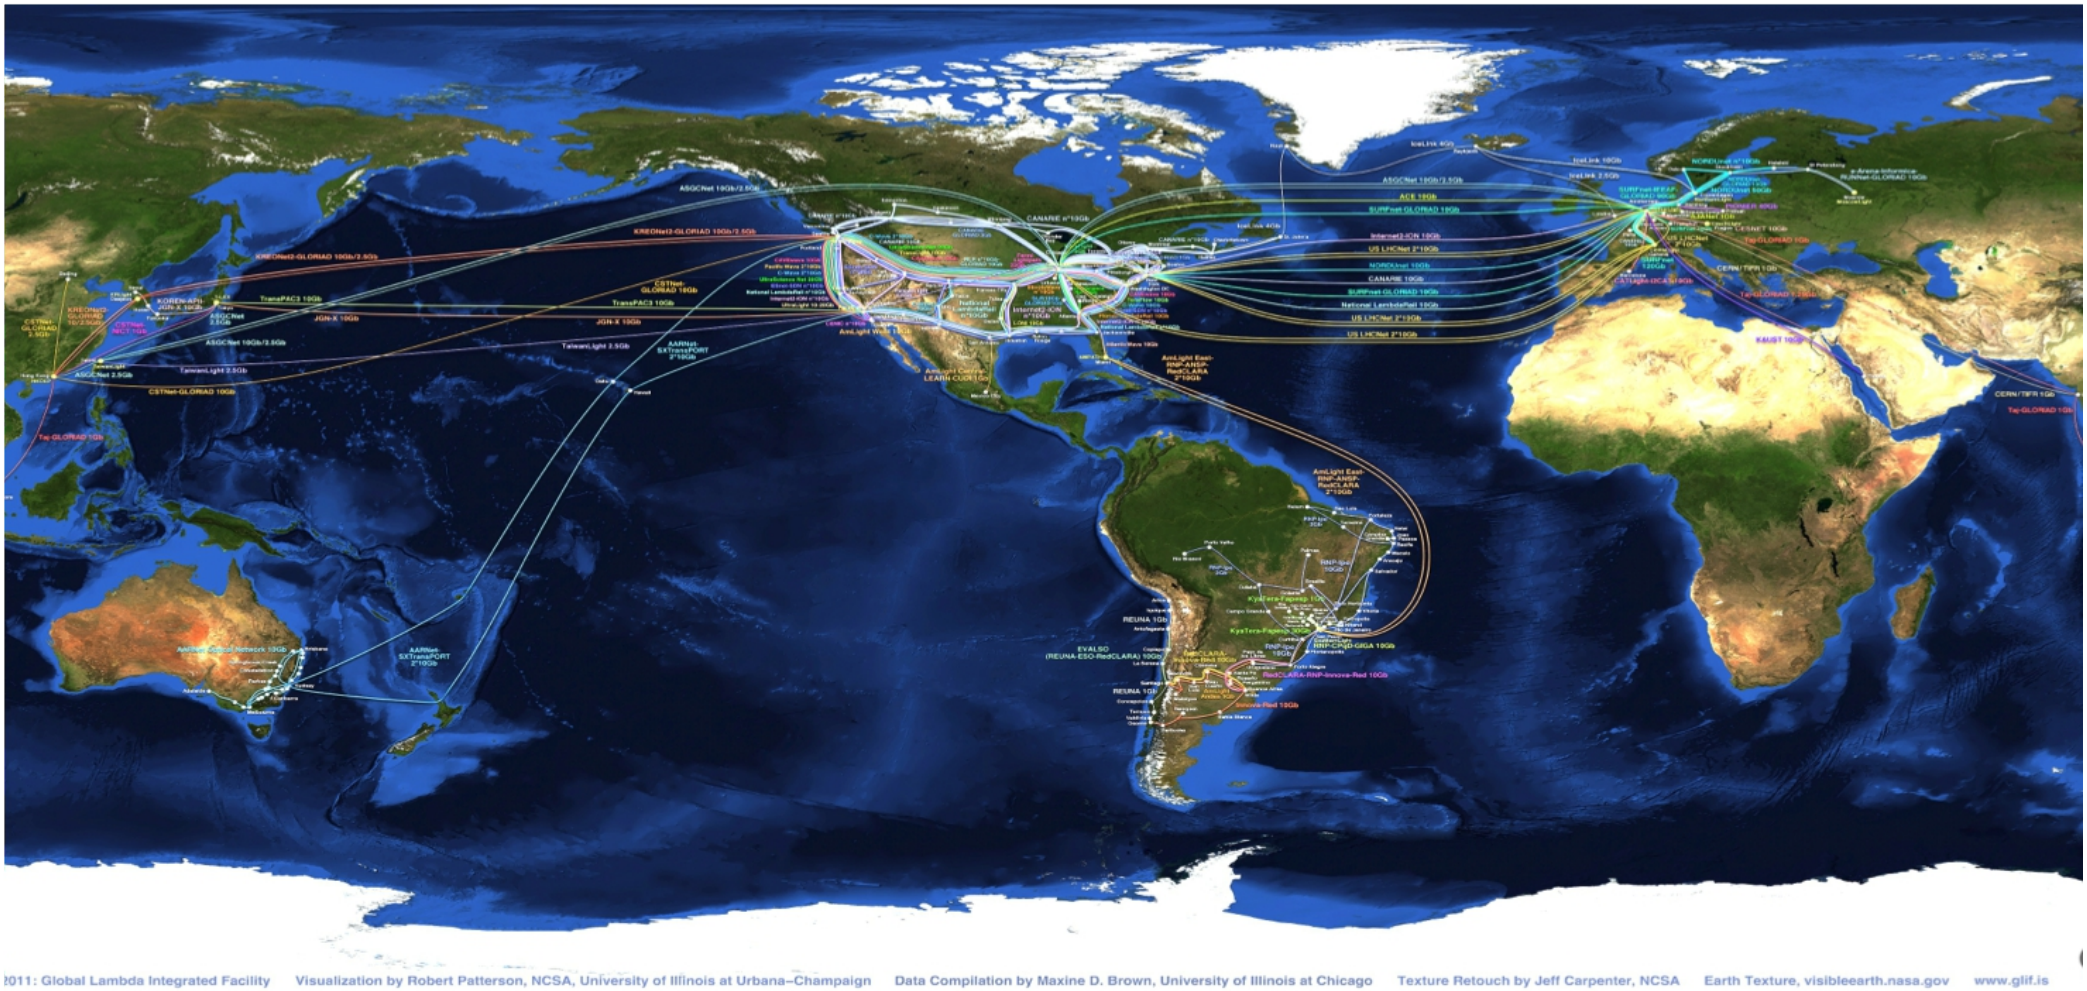

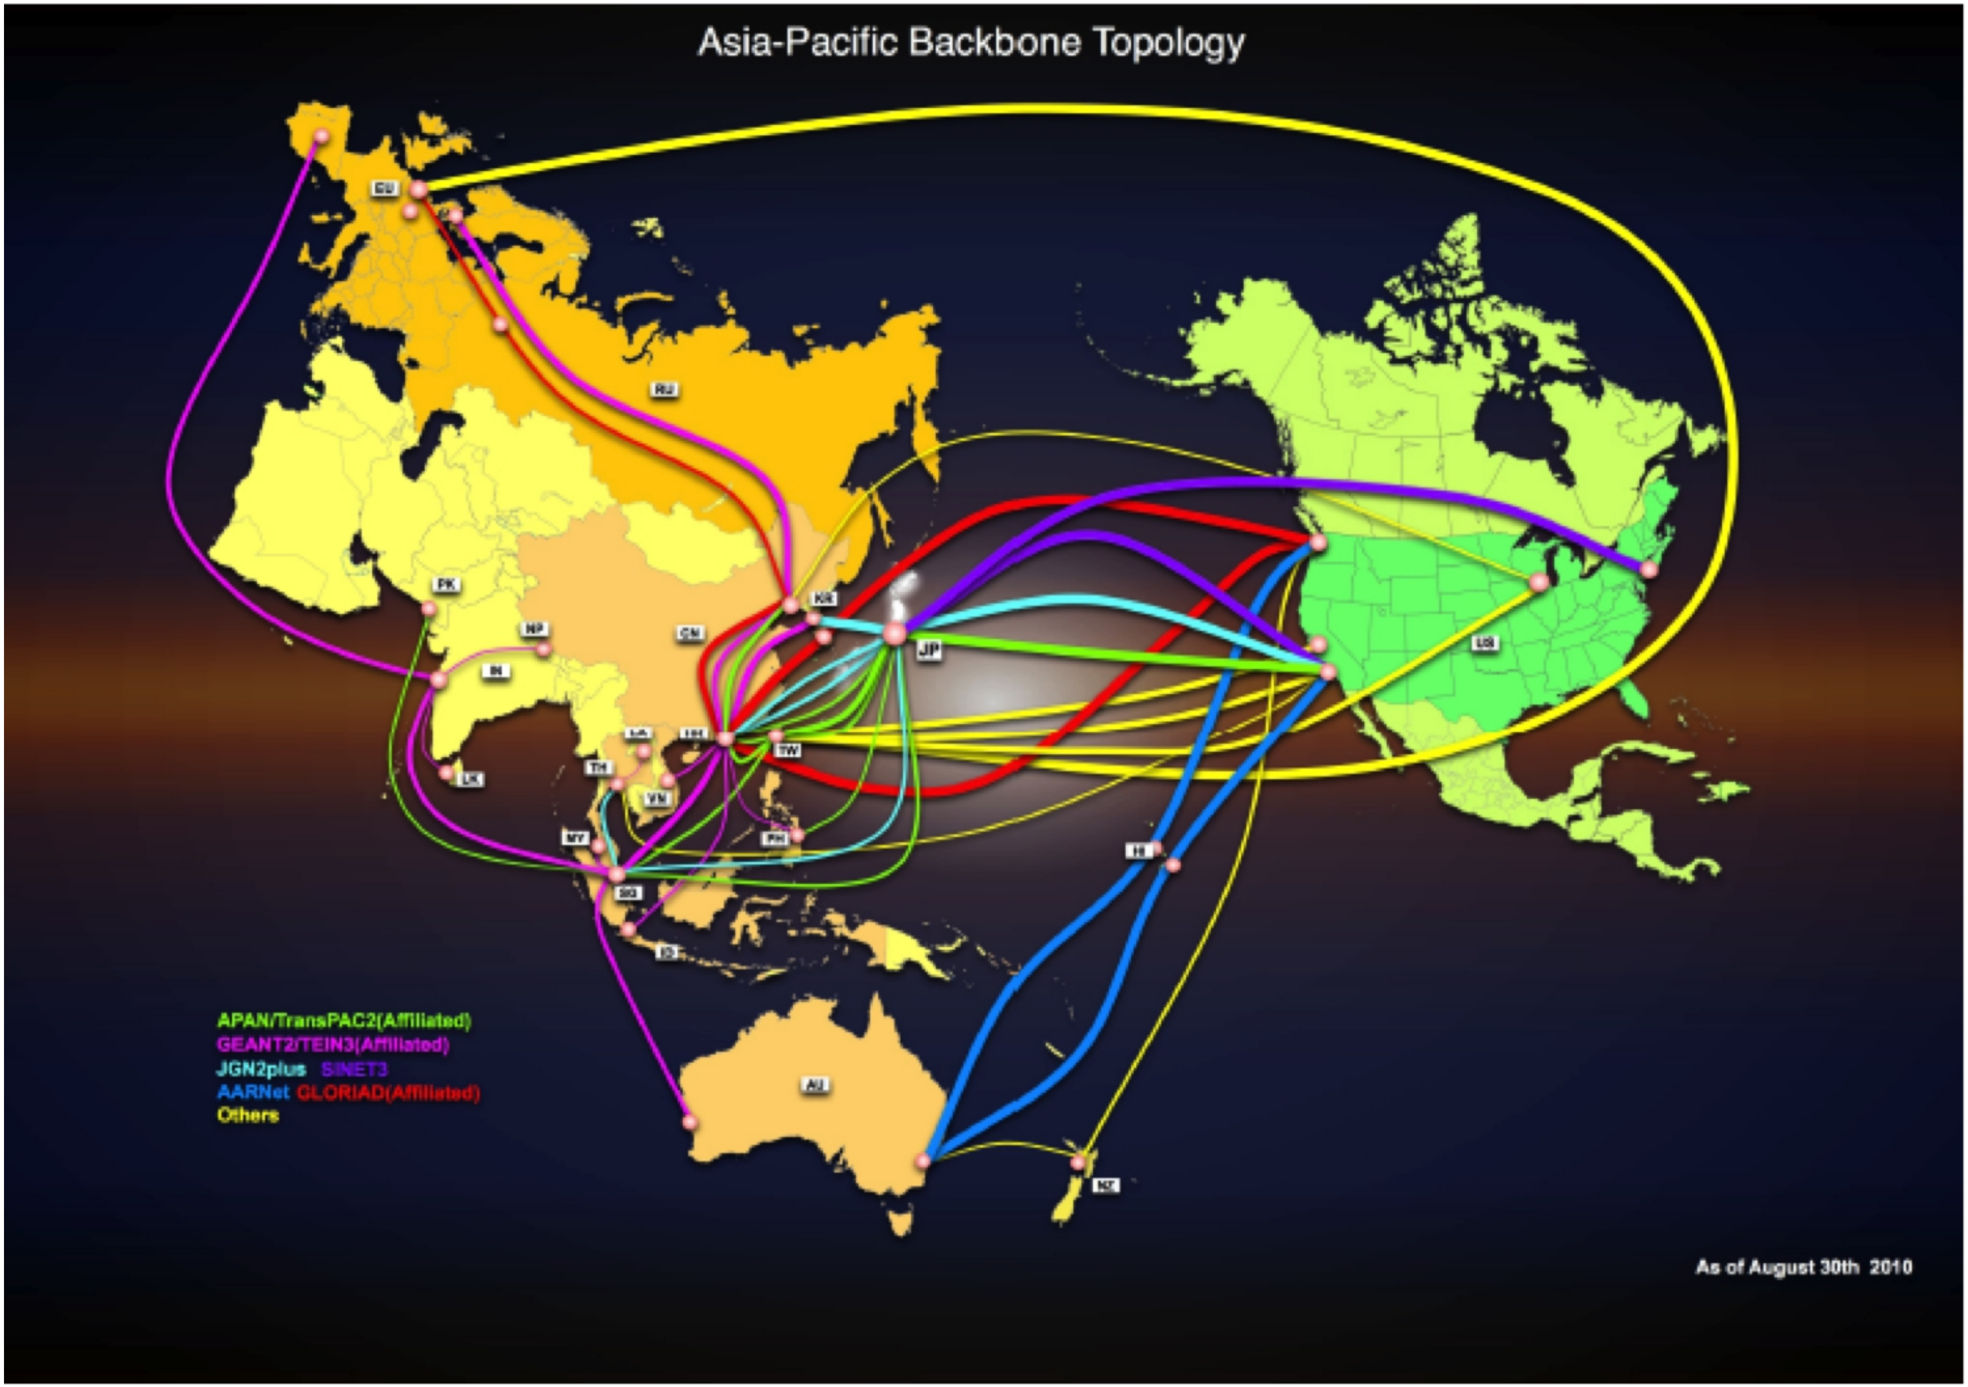

Asia-Pacific Backbone Topology
EU
RU
PK
NP
IN
CN
KR
JP
TW
TH
VN
LK
MY
SG
PH
ID
AU
NZ
HI
US
APAN/TransPAC2(Affiliated)
GEANT2/TEIN3(Affiliated)
JGN2plus SINET3
AARNet GLORIAD(Affiliated)
Others
As of August 30th 2010

GÉANT At the Heart of Global Research Networking

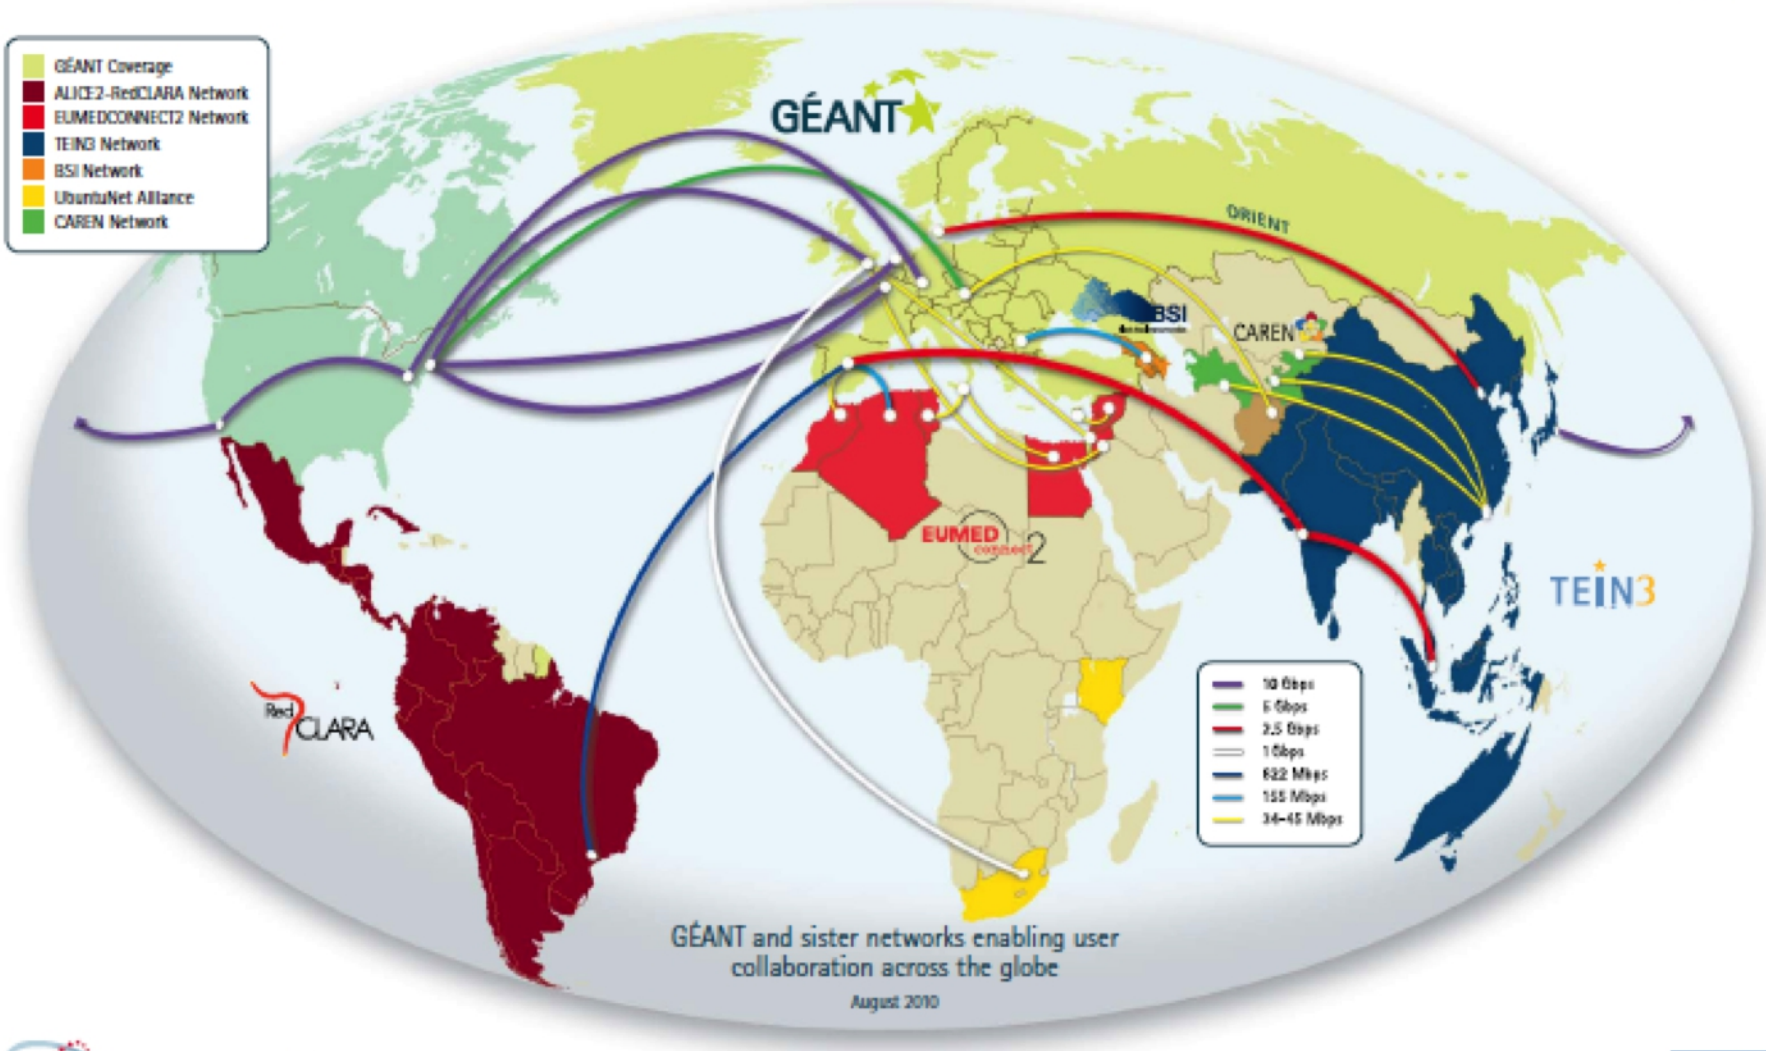

DANTE www.dante.net GÉANT www.geant.net

connect • communicate • collaborate

GÉANT is co-funded by the European Commission within its 7th RésD Framework Programme.
This document has been produced with the financial assistance of the European Union. The contents of this document are the sole responsibility of DANTE and can under no circumstances be regarded as reflecting the position of the European Union.

European Commission Information Society and Media

UNIVERSITY OF OREGON

NSRC Network Startup Resource Center

As at June 2014

NSRC
Network Startup Resource Center

Problem Statement

This complex, heterogeneous set of networks must operate seamlessly from “end to end” to support science and research collaborations that are distributed globally.

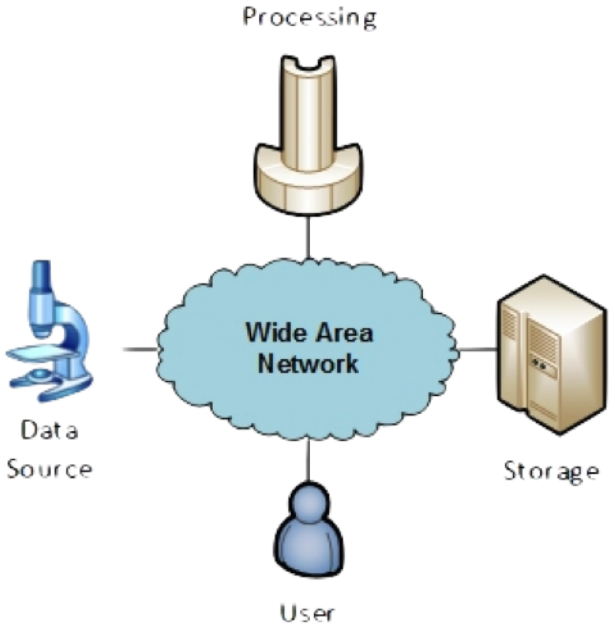

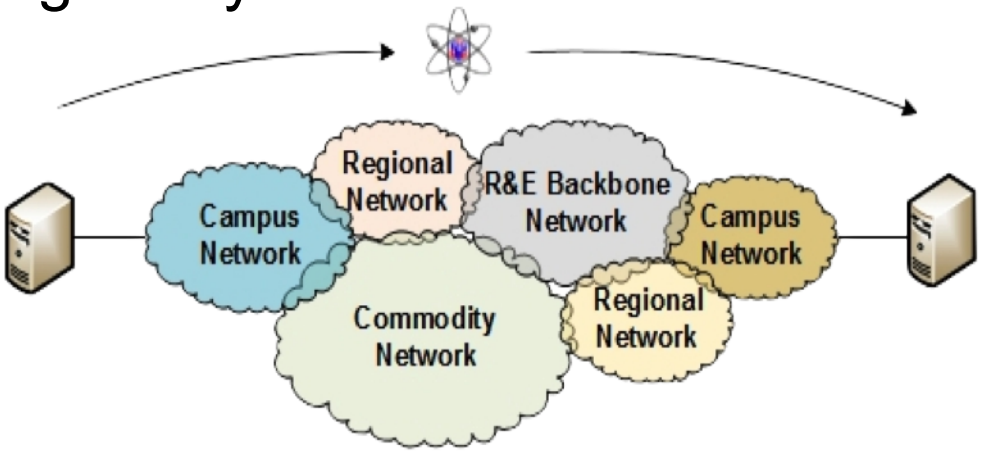

Problem Statement

- In practice, performance issues are prevalent and distributed.
- When a network is underperforming or errors occur, it is difficult to identify the source, as problems can happen anywhere, in any domain.
- Local-area network testing is not sufficient, as errors can occur between networks.

Problem Statement
Hard vs. Soft Failures

"Hard failures" are the kind of problems every organization understands

- Fiber cut
- Power failure takes down routers
- Hardware ceases to function

Classic monitoring systems are good at alerting hard failures

- i.e., NOC sees something turn red on their screen
- Engineers paged by monitoring systems

"Soft failures" are different and often go undetected

- Basic connectivity (ping, traceroute, web pages, email) works
- Performance is just poor

How much should we care about soft failures?

What's it take to transfer data...

Data set size	1 Minute	5 Minutes	20 Minutes	1 Hour
	Time to transfer			
10PB	1,333.33 Tbps	266.67 Tbps	66.67 Tbps	22.22 Tbps
1PB	133.33 Tbps	26.67 Tbps	6.67 Tbps	2.22 Tbps
100TB	13.33 Tbps	2.67 Tbps	666.67 Gbps	222.22 Gbps
10TB	1.33 Tbps	266.67 Gbps	66.67 Gbps	22.22 Gbps
1TB	133.33 Gbps	26.67 Gbps	6.67 Gbps	2.22 Gbps
100GB	13.33 Gbps	2.67 Gbps	666.67 Mbps	222.22 Mbps
10GB	1.33 Gbps	266.67 Mbps	66.67 Mbps	22.22 Mbps
1GB	133.33 Mbps	26.67 Mbps	6.67 Mbps	2.22 Mbps
100MB	13.33 Mbps	2.67 Mbps	0.67 Mbps	0.22 Mbps

But, just a little packet loss...

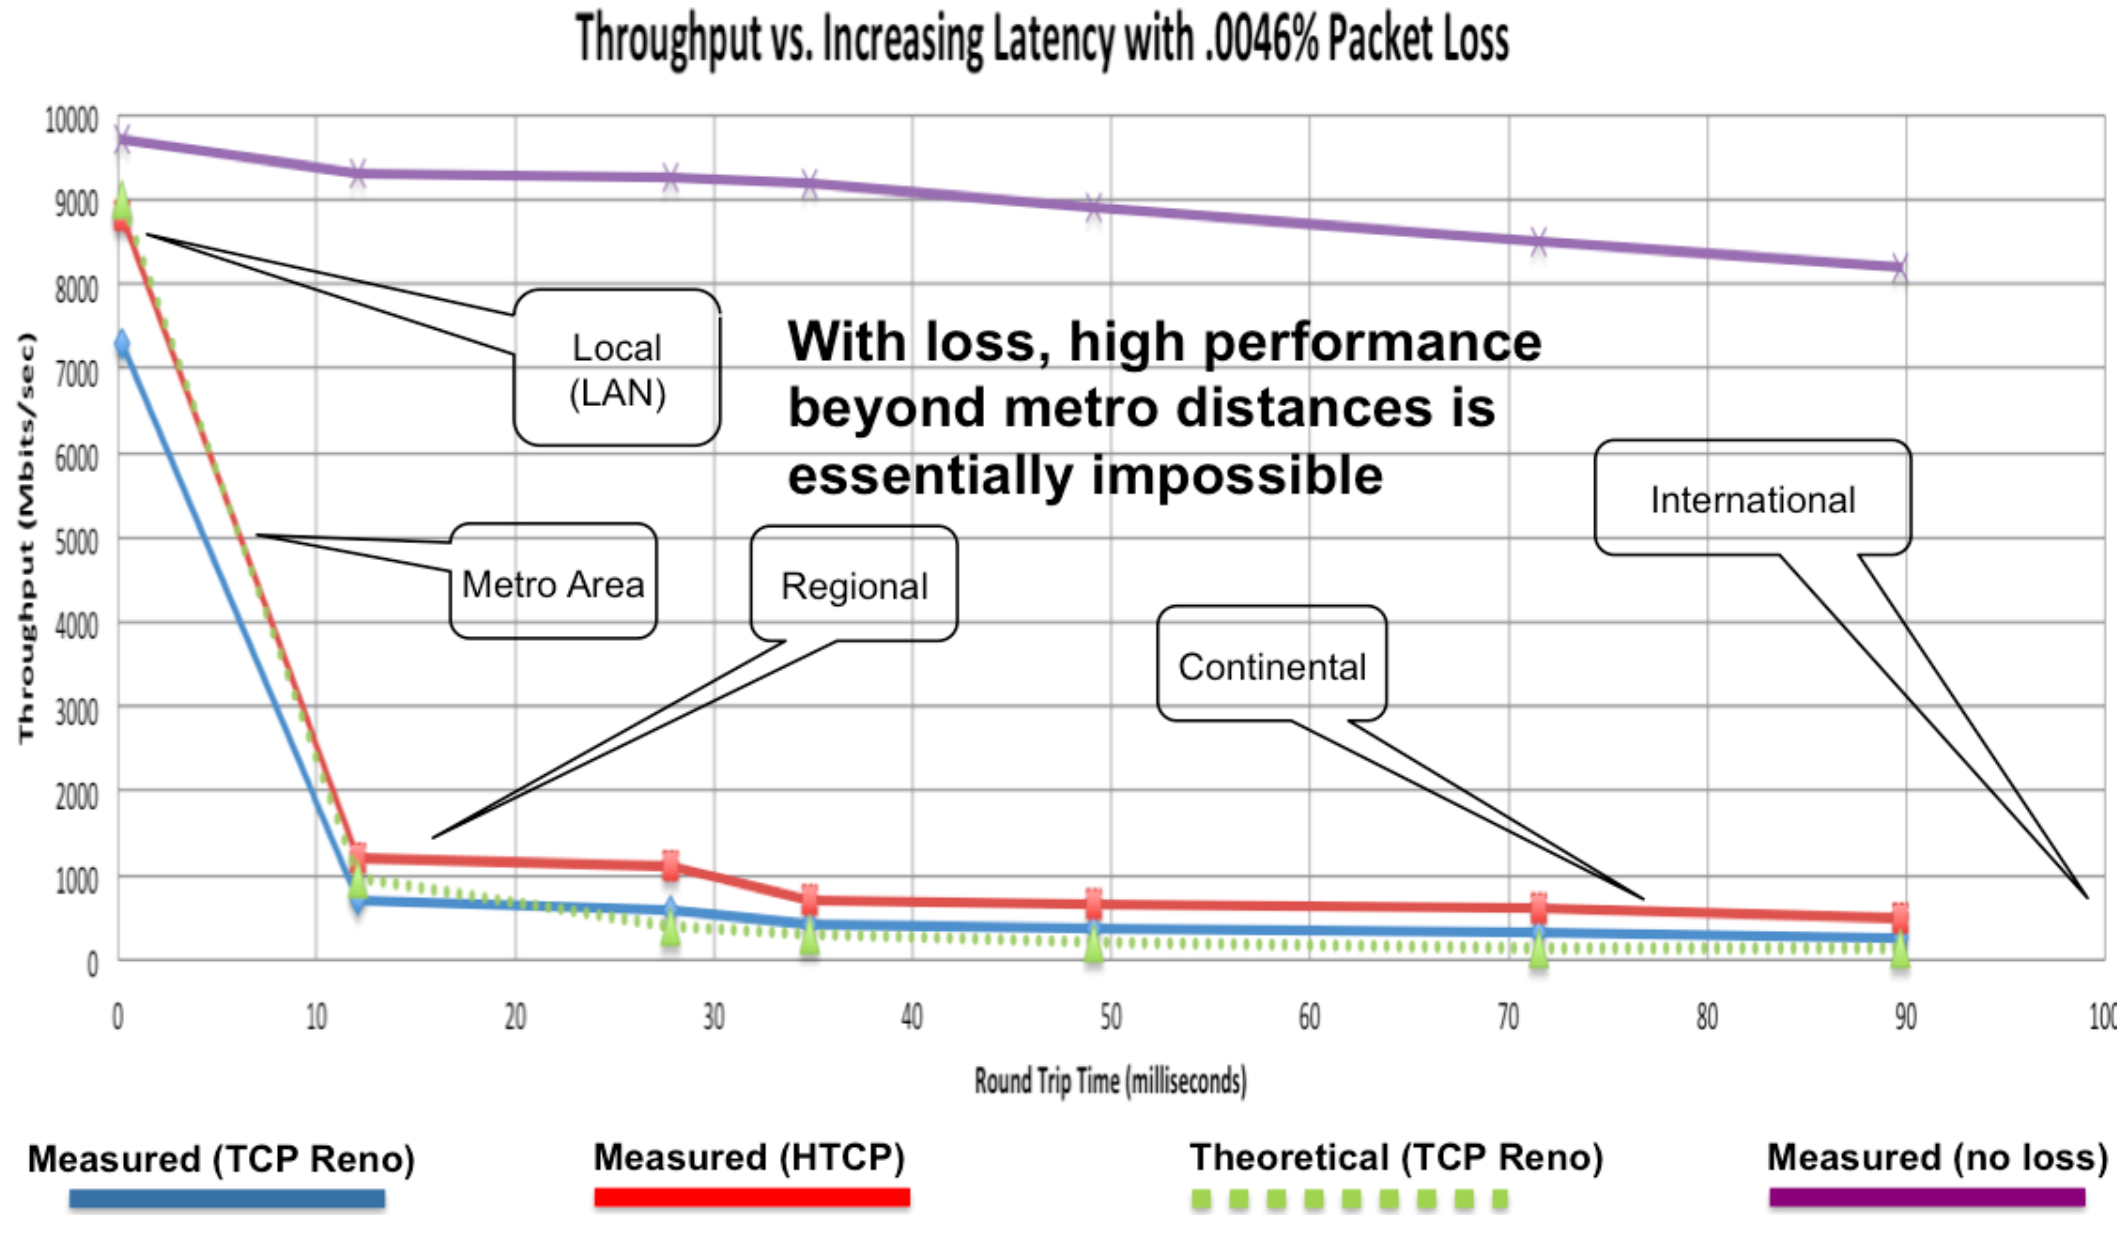

A Solution
perfSONAR and the perfSONAR Toolkit

A widely-deployed test and measurement infrastructure used by local, regional and national research networks and science facilities to actively track & troubleshoot network performance issues.

Open-source, community-developed software suite that:

- Provides consistent set of proven tools to pinpoint and resolve network performance issues including soft failures across complex, multi-domain data paths.
- Creates a standard way to visualize, publish, and archive network metrics and data for future analysis, and aids other networks in debugging issues.

PerfSONAR Toolkit

What do we care about measuring?

- **Latency:** Ping, OWAMP, NDT, NPAD (don't trust Traceroute)
- **Jitter:** OWAMP, UDP Iperf, UDP NuQcp, NDT, NPAD
- **Packet Loss, Duplication, unordered data:** OWAMP, UDP Iperf, UDP NuQcp, NDT, NPAD
- **InterfaceUtiliization/Discards/Errors:** – SNMP
- **Achievable Bandwidth:** Iperf, NuQcp, NDT, NPAD
- **Traveled Route:** Traceroute, Tracepath
- **MTU Feedback:** Tracepath, Ping, NDT, NPAD

Benefit: Multi-domain Insight

Over 1000 public perfSONAR nodes are deployed globally

- This allows you to test performance between your perfSONAR node and any other publicly available perfSONAR node. Very cool.
- Makes it possible to identify and work to debug issues even if they are outside your network domain.

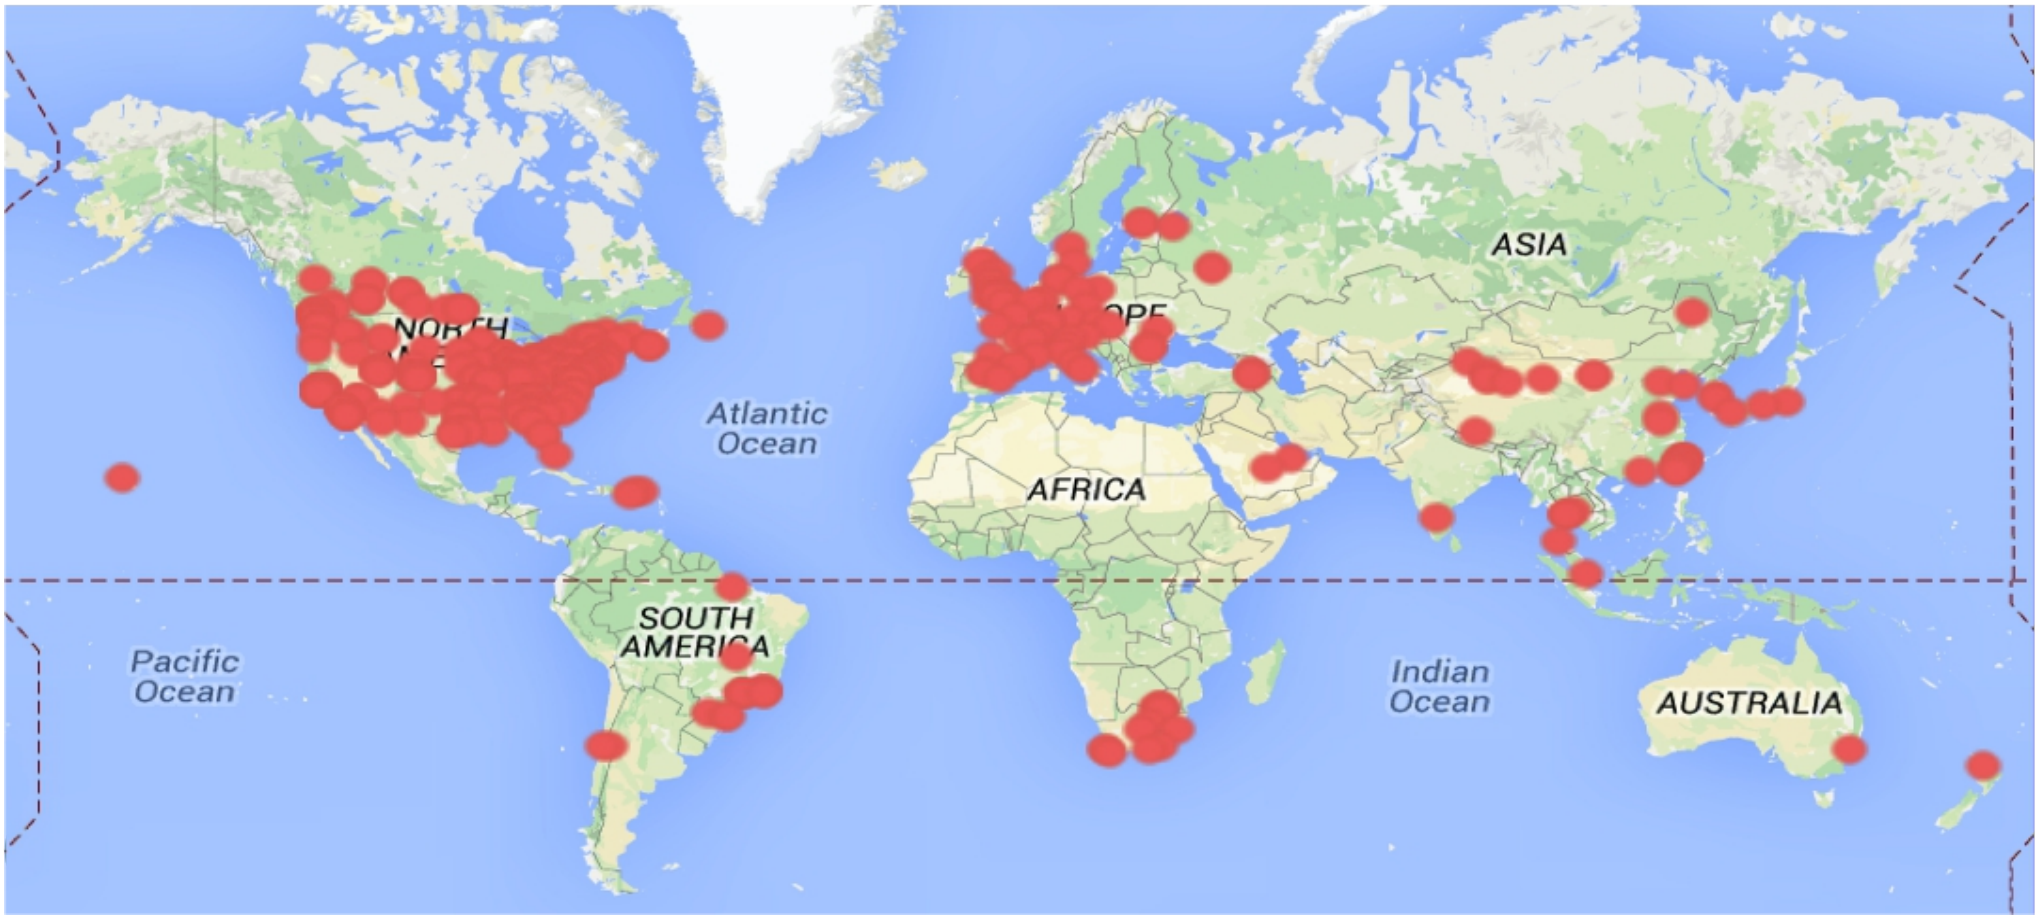

perfSONAR Community Benefits

Many of these services exists from other projects, but perfSONAR brings the concepts together.

Testing from client A between points B and C...

- **Latency:** Ping, OWAMP, NDT, NPAD (don't trust Traceroute)
- **Jitter:** OWAMP, UDP Iperf, UDP NuQcp, NDT, NPAD
- **Packet Loss, Duplication, unordered data:** OWAMP, UDP Iperf, UDP NuQcp, NDT, NPAD
- **InterfaceUtiliization/Discards/Errors:** – SNMP
- **Achievable Bandwidth:** Iperf, NuQcp, NDT, NPAD
- **Traveled Route:** Traceroute, Tracepath
- **MTU Feedback:** Tracepath, Ping, NDT, NPAD

perfSONAR Community Benefits

i.e., Testing from client A between points B and C...

```
[root@perfsonar100 ~]# bwctl -fM -t 30 -s 202.28.229.115
bwctl: Using tool: iperf
bwctl: 37 seconds until test results available

RECEIVER START
------------------------------------------------------------
Server listening on TCP port 5204
Binding to local address 203.159.61.100
TCP window size: 0.08 MByte (default)
------------------------------------------------------------
[ 15] local 203.159.61.100 port 5204 connected with 202.28.229.115 port 38953
[ ID] Interval       Transfer     Bandwidth
[ 15]  0.0-30.1 sec   210 MBytes  6.96 MBytes/sec
[ 15] MSS size 1448 bytes (MTU 1500 bytes, ethernet)

RECEIVER END
```

Then automate this within web GUI and build historical trends...

Throughput Historical Trends

bwctl (throughput) testing over time with saved graphs

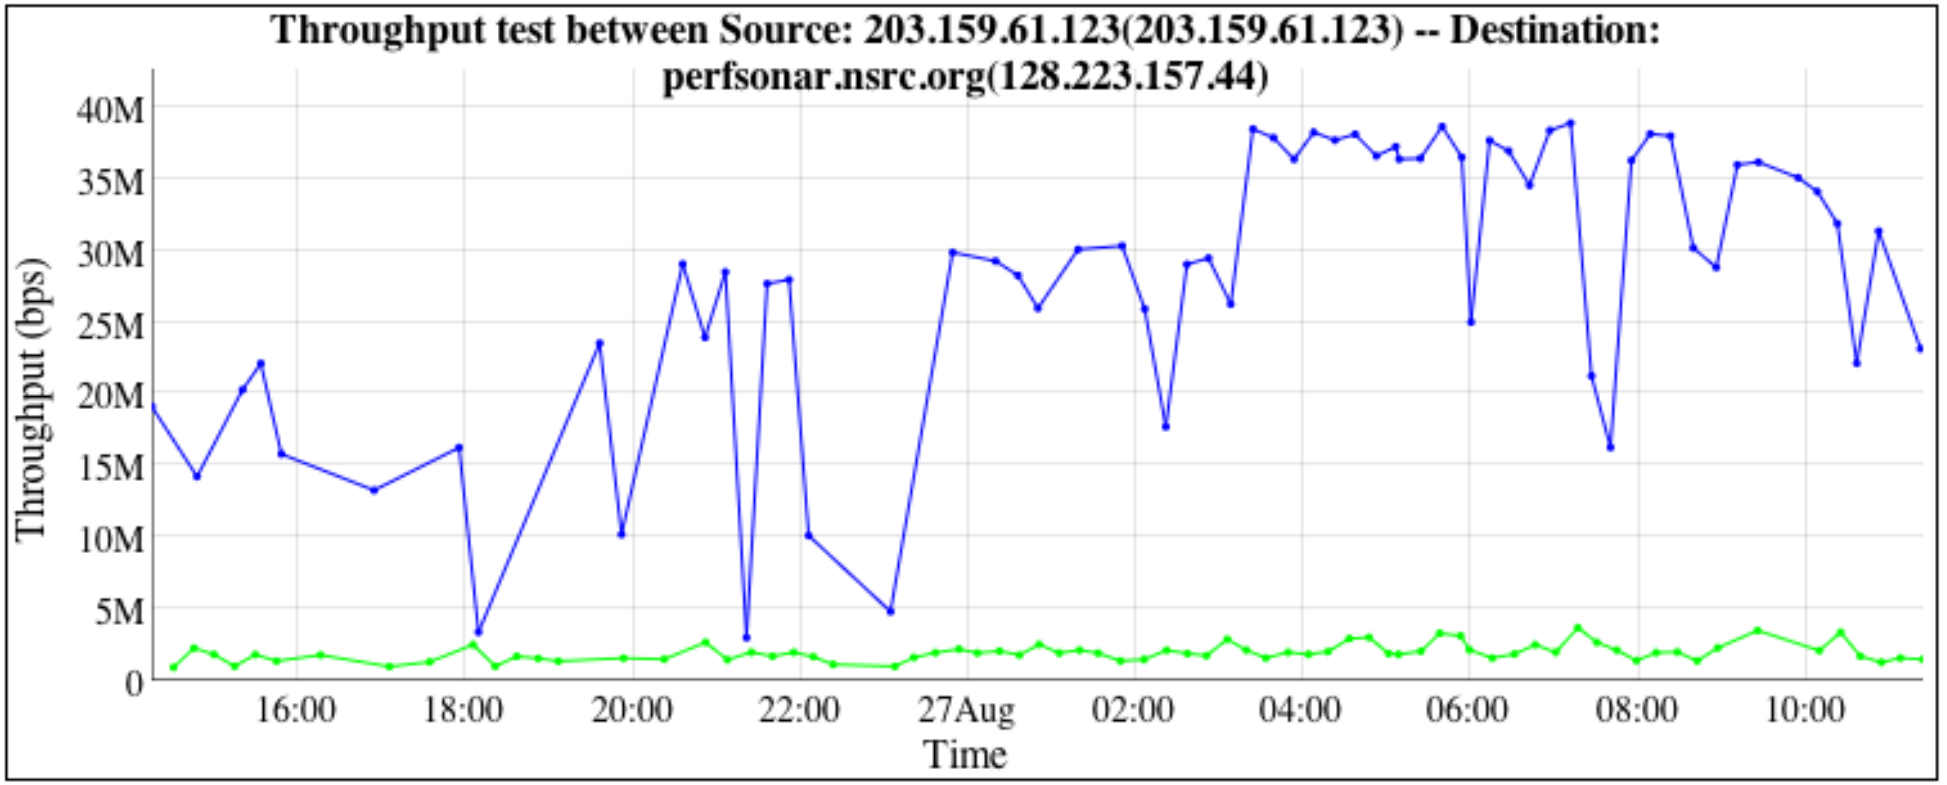

Benefits: Finding hidden problems

Above all, perfSONAR allows you to maintain a healthy, high-performing network because it helps identify the "soft failures" in the network path.

Classical monitoring systems have limitations

- ✓ Performance problems are often only visible at the ends
- Individual network components (e.g. routers) have no knowledge of end host state

- perfSONAR tests the network in ways that classical monitoring systems do not
- More perfSONAR distributions equal better network visibility.

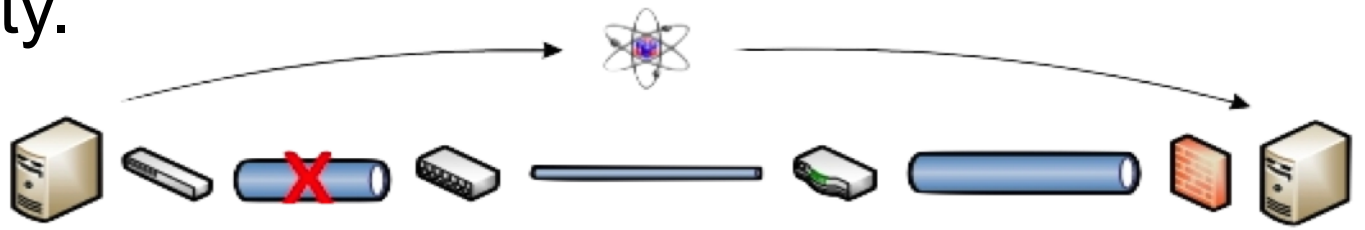

Benefits: Finding hidden problems

perfSONAR is designed to pinpoint and identify soft failures to accelerate resolution.

Example: Find and replace failing optics:

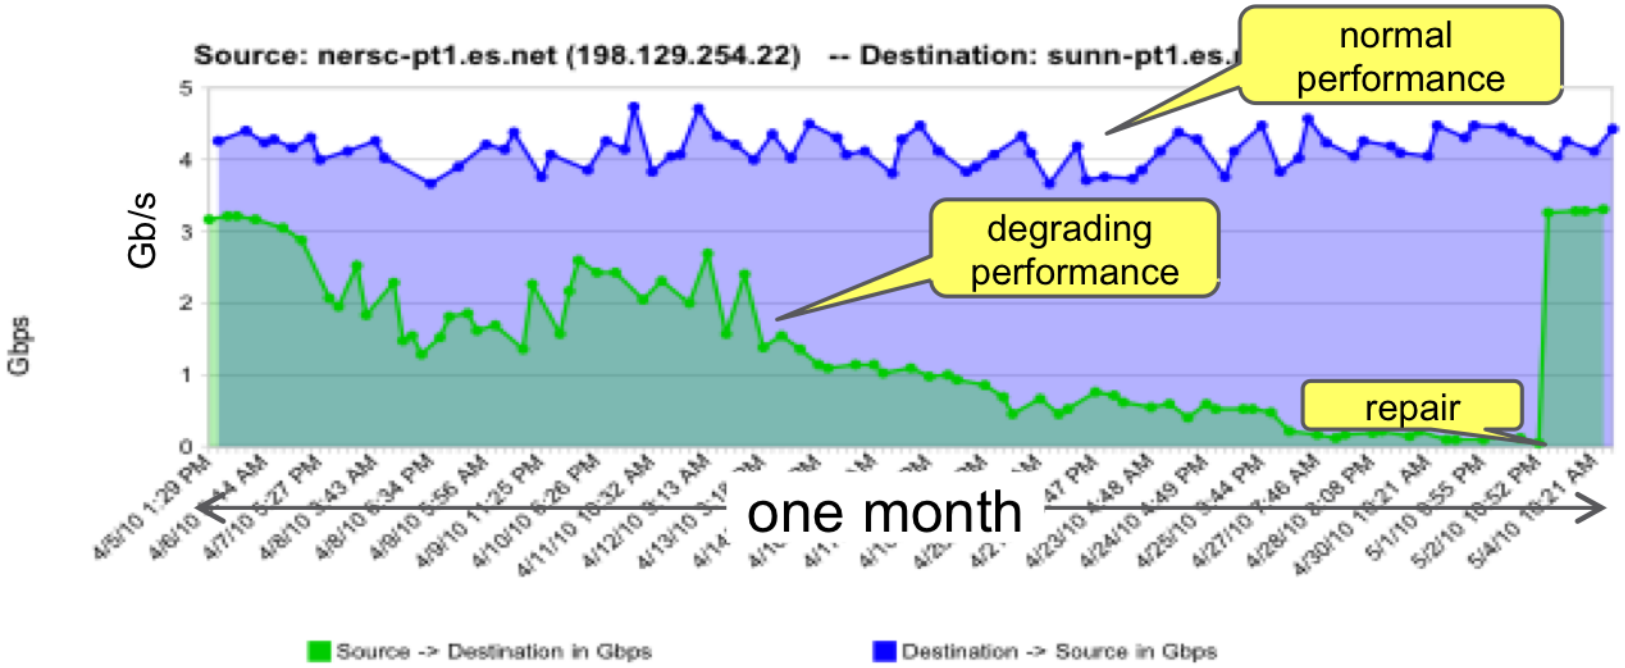

Benefits: Finding hidden problems

perfSONAR is designed to pinpoint and identify soft failures to accelerate resolution.

Example: Host Tuning:

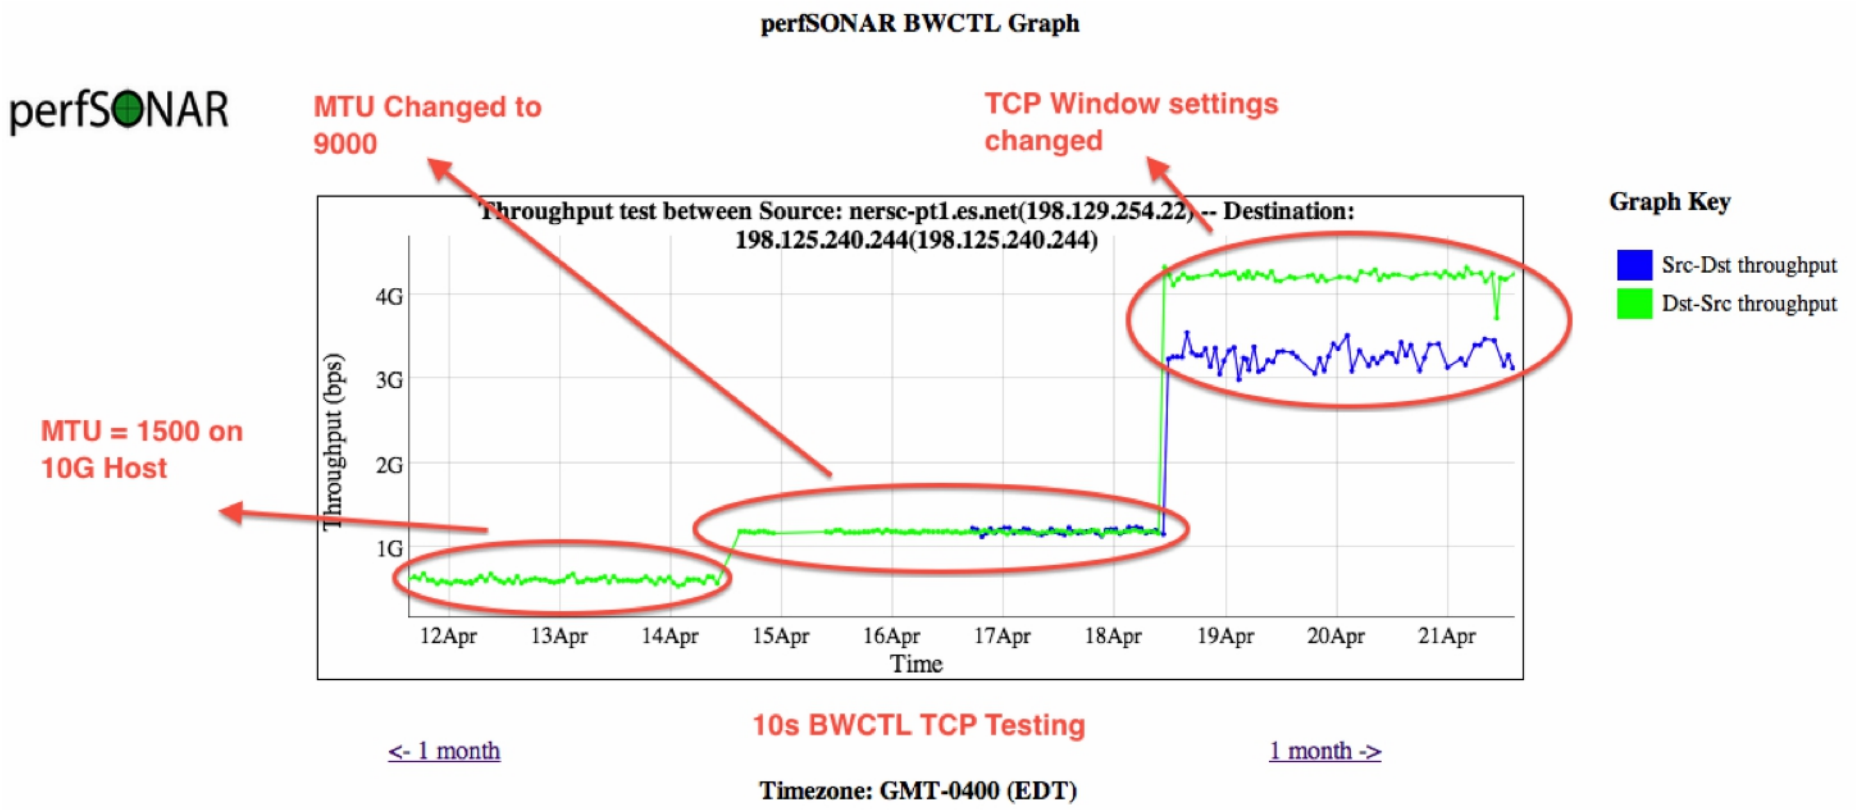

Benefits: Finding hidden problems

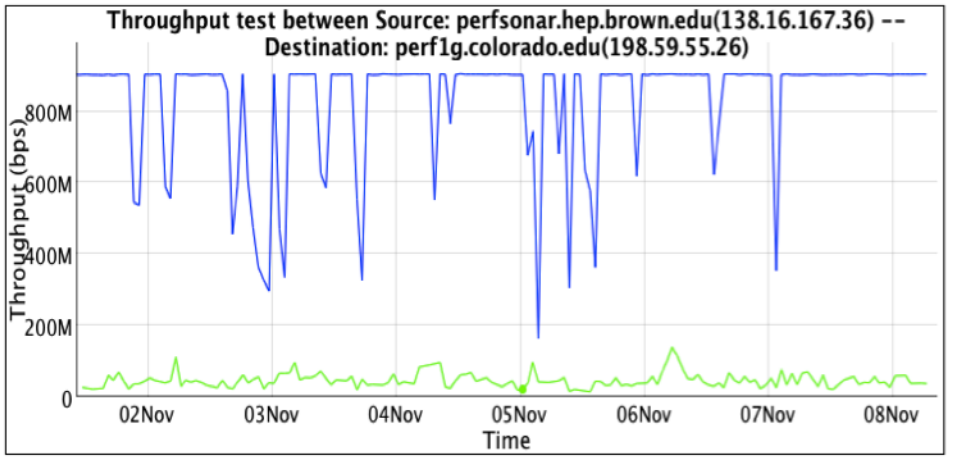

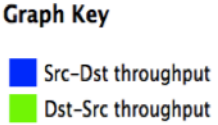

Inside the firewall

- One direction severely impacted by firewall
- Not useful for science data

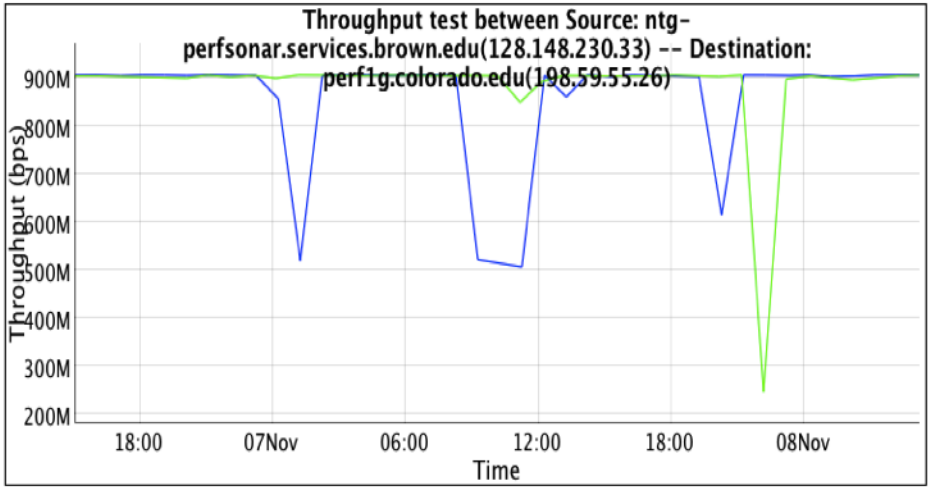

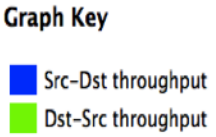

Outside the firewall

- Good performance in both directions

Benefits: Investment Planning, Verification

Test data and metrics verifies performance of current investments

- Can you get full use out of the carrier circuit you're paying for?
- If performance suddenly changes, what caused it?

Data informs future investments

- Put resources in places with demonstrated need
- Example: if your firewall is limiting performance, buying more bandwidth won't solve the problem

Make the case for architectural improvements

- How might a Science DMZ or alternate architecture change performance?
- How to evaluate potential purchases?
  - Put perfSONAR hosts on them and run tests
  - Prototype new model

Resources

perfSONAR Website
http://www.perfsonar.net/
perfSONAR mailing lists
http://www.personar.net/about/getting-help/
perfSONAR Directory of Users
http://stats.es.net/Services
FasterDate Knowledgebase
http://fasterdata.es.net/
perfSONAR Checklist
https://code.google.com/p/perfsonar-ps/wiki/PerformanceToolkitTestingChecklist

Resources (perfSONAR ecosystem)

Detailed perfSONAR Install, Configuration and Use Help
https://code.google.com/p/perfsonar-ps/wiki/pSPerformanceToolkit33

Network Troubleshooting Tools
http://fasterdata.es.net/performance-testing/network-troubleshooting-tools/

perfSONAR HowTos
http://fasterdata.es.net/performance-testing/perfsonar/ps-howto/

PerfSONAR-PS Toolkit Main Page
http://psps.perfsonar.net/

GEANT perfSONAR Pages (perfSONAR-MDM)
http://services.geant.net/PerfSONAR/Pages/Home.aspx

GEANT perfSONAR-MDM Product Brief
http://services.geant.net/PerfSONAR/Resources/Documents/PUB_12-055_perfsonar-product_brief_05.12.pdf

NDT/NPAD Overview
https://fasterdata.es.net/performance-testing/network-troubleshooting-tools/ndt-npad/

Network Diagnostic Tool (NDT)
http://software.internet2.edu/ndt/

Resources (perfSONAR ecosystem)

perfSONAR Lookup Services (LS)
http://psps.perfsonar.net/lookup/
Science DMZ
http://fasterdata.es.net/science-dmz/
Esnet Data Transfer Nodes (DTNs)
http://fasterdata.es.net/performance-testing/DTNs/
Globus
http://fasterdata.es.net/data-transfer-tools/globus/
Tracepath
http://linux.about.com/library/cmd/blcmdl8_tracepath.htm
bwctl command line
http://software.internet2.edu/bwctl/bwctl.man.html

Primary perfSONAR Development

For a full list of development collaborators and history of perfSONAR, please visit:

http://www.perfsonar.net/about/who-is-involved/