



Monitoring the US ATLAS Network Infrastructure with perfSONAR-PS

Shawn McKee/University of Michigan

Andrew Lake/ESnet, Philippe Laurens/MSU, Horst Severini/OU,

Tomasz Wlodek/BNL, Stephen Wolff/I2 and Jason Zurawski/I2

on behalf of the ATLAS Collaboration

CHEP 2012, New York, New York

May 21st, 2012



❄ Summary of this talk:

- ❑ Motivations and perfSONAR overview
- ❑ perfSONAR-PS in USATLAS: goals and deployment details
- ❑ Modular dashboard
- ❑ Example use of the tools to find/fix problems
- ❑ Future work

Motivations for Monitoring the Network



- ❄ LHC collaborations are:
 - ❑ Data intensive
 - ❑ Globally distributed
 - ❑ and rely upon the network as a **critical** part of their infrastructure
- ❄ Yet finding and debugging LHC network problems can be **difficult** and, in some cases, take months.
- ❄ How can we quickly identify when problems are network problems and help isolate their locations?
- ❄ We don't want to have a network monitoring system per VO!
- ❄ The perfSONAR project was designed to help do this

What is perfSONAR?



❄ A collaboration

- ❑ Production network operators focused on designing and building tools that they will deploy and use on their networks to provide monitoring and diagnostic capabilities to themselves and their user communities.

❄ An architecture & set of communication protocols

- ❑ Web Services (WS) Architecture
- ❑ Protocols established in the Open Grid Forum
 - ⌘ Network Measurement Working Group ([NM-WG](#))
 - ⌘ Network Measurement Control Working Group ([NMC-WG](#))
 - ⌘ Network Markup Language Working Group ([NML-WG](#))

❄ Several interoperable software implementations

- ❑ [perfSONAR-MDM](#)
- ❑ [perfSONAR-PS](#)

❄ A Deployed Measurement infrastructure

perfSONAR Architecture Overview



- Interoperable network measurement middleware designed as a Service Oriented Architecture (SOA):
 - Each component is modular
 - All are Web Services (WS) based
 - The global *perfSONAR* framework as well as individual deployments are decentralized
 - All *perfSONAR* tools are Locally controlled
 - All *perfSONAR* tools are capable of federating locally and globally
- *perfSONAR* Integrates:
 - Network measurement tools and archives (e.g. stored measurement results)
 - Data manipulation
 - Information Services
 - Discovery
 - Topology
 - Authentication and authorization

perfSONAR in USATLAS



- ❄ Since the network is so fundamental to our work on the Large Hadron Collider (LHC) USATLAS targeted deployment of a **perfSONAR** instance at all US ATLAS primary facilities.
 - ❄ The perfSONAR-PS toolkit was selected because of our close working relationship with both ESnet/Internet2(<http://psps.perfsonar.net/toolkit>)
- ❄ **perfSONAR's main purpose is to aid in network problem diagnosis**
 - ❄ It can quickly allow users to isolate the location of problems.
 - ❄ **It provides a standard measurement of various network metrics over time via scheduled tests**
 - ❄ It can provide “on-demand” tests.
 - ❄ **It can be used to alert users to significant changes in the network**
- ❄ **Both USATLAS and USCMS** have expressed a strong interest in broadly deploying these tools in a consistent way

perfSONAR-PS Deployment



- ❄ USATLAS wanted a set of tools that:
 - ❑ Are easy to install
 - ❑ Measure the “network” behavior
 - ❑ Provide a baseline of network performance between end-sites
 - ❑ Are standardized and broadly deployed
 - ❑ Provide a history/archive for forensic reference
- ❄ Details of how sites should setup the perfSONAR-PS installations are documented on the Twiki at:
<https://twiki.cern.ch/twiki/bin/view/LHCONE/SiteList>
 - ❑ These instructions originated from USATLAS’s experience
 - ❑ Tests are configured to measure achievable bandwidth, latency, packet-loss and routing.
- ❄ Next I will highlight some of the relevant considerations

perfSONAR-PS Deployment Considerations



- ❄ We want to measure (to the extent possible) the entire network path between USATLAS resources. This means:
 - ❑ We want to locate perfSONAR-PS instances as close as possible to the storage resources associated with a site. The goal is to ensure we are measuring the same network path to/from the storage.

- ❄ There are two separate instances that should be deployed: latency & bandwidth (Two instances to prevent interference)
 - ❑ The **latency instance** measures one-way delay by using an NTP synchronized clock and send 10 packets per second to target destinations (Important metric is **packet-loss!**)
 - ❑ The **bandwidth instance** measures achievable bandwidth via a short test (20-60 seconds) per src-dst pair every 4 hour period

perfSONAR-PS Deployment Considerations



- ❄ Each “site” should have perfSONAR-PS instances in place.
 - ❑ If a Tier-2 has more than one “network” location, each should be instrumented and made part of scheduled testing.

- ❄ Standardized hardware and software is a **good** idea
 - ❑ Measurements should represent what the **network** is doing and not differences in hardware/firmware/software.
 - ❑ USATLAS has identified and tested systems from Dell for perfSONAR-PS hardware. Two variants: R310 and R610.
 - ⌘ R310 cheaper (<\$900), can host 10G (Intel X520 NIC) but not supported by Dell (Most US ATLAS sites choose this)
 - ⌘ R610 officially supports X520 NIC (Canadian sites choose this)
 - ⌘ Orderable off the **Dell LHC portal**
 - ❑ We try to coordinate upgrades USATLAS-wide

Network Impact of perfSONAR-PS



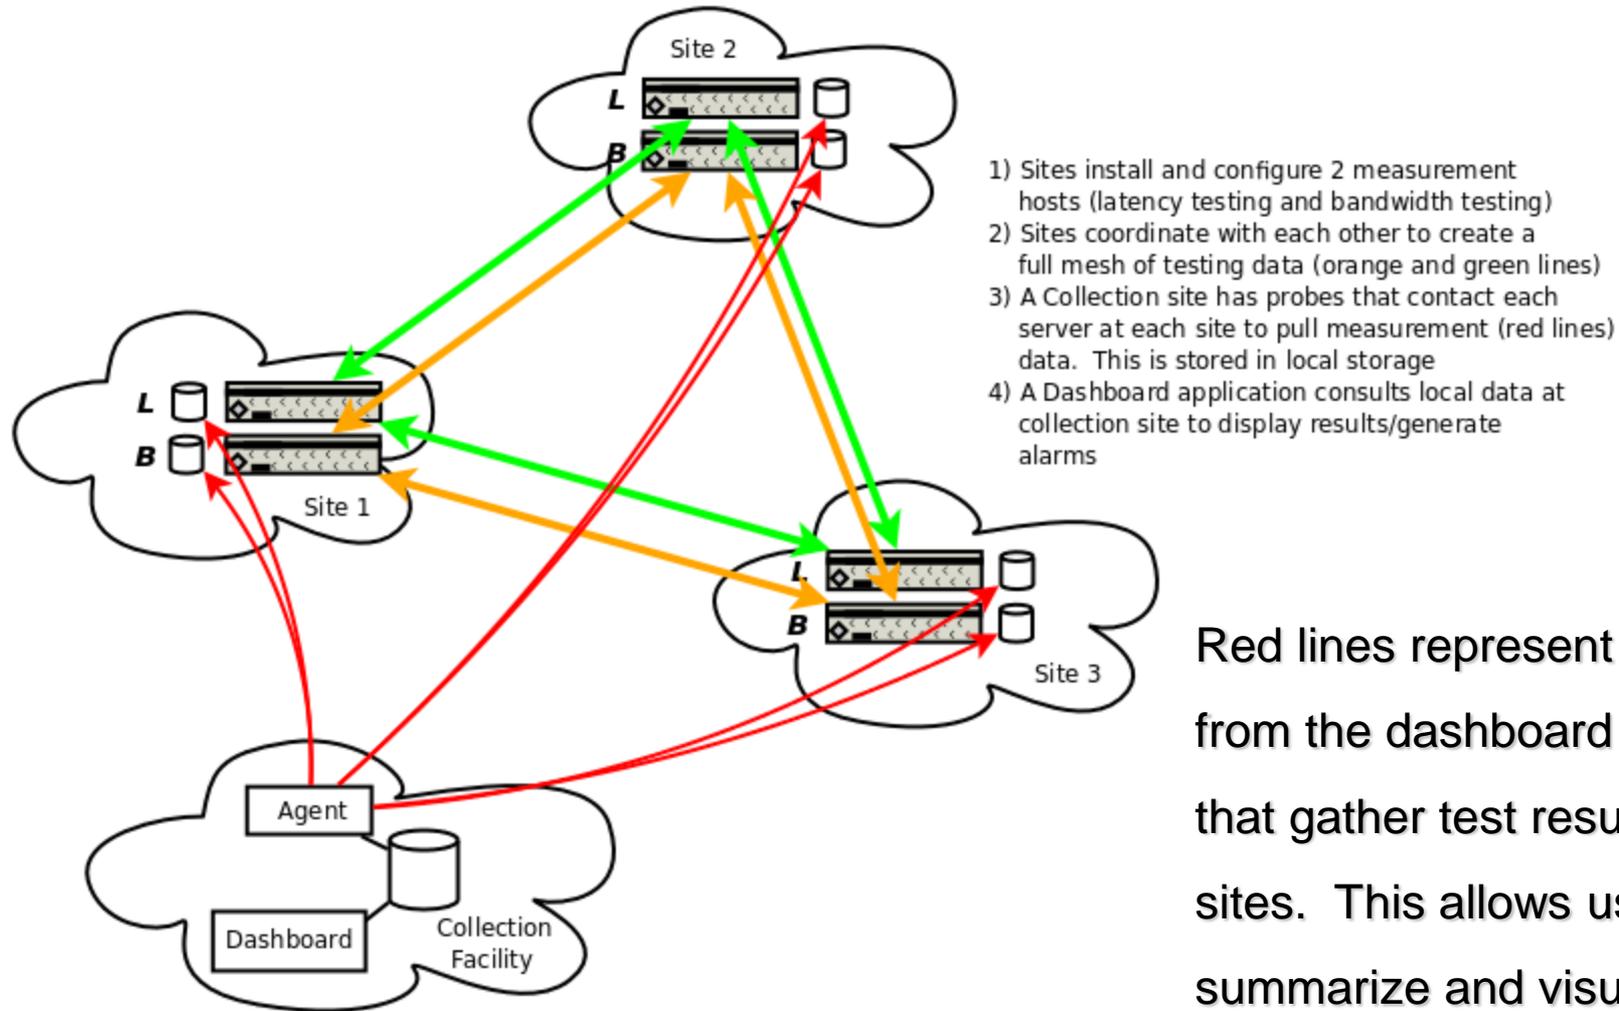
- ❄ To provide an idea of the network impact of a typical deployment here are some numbers as configured in the US
 - ❑ **Latency tests** send 10Hz of small packets (20 bytes) for each testing location. USATLAS Tier-2's test to ~9 locations. Since headers account for 54 bytes each packet is 74 bytes or the rate for testing to 9 sites is **6.7 kbytes/sec**.
 - ❑ **Bandwidth tests** try to maximize the throughput. A 20 second test is run from each site in each direction once per 4 hour window. Each site runs tests in both directions. Typically the best result is around **925 Mbps on a 1Gbps link for a 20 second test**. That means we send $4 \times 925 \text{ Mbps} \times 20 \text{ sec}$ every 4 hours per testing pair (src-dst) or about **46.25 Mbps** average for testing with 9 other sites.
 - ❑ **Traceroute tests** are negligible in terms of bandwidth used
 - ❑ Tests are configurable but the above settings are working fine.

Modular Dashboard



- ❄ While the perfSONAR-PS toolkit is very nice, it was designed to be a distributed, federated installation.
 - ❑ Not easy to get an “overview” of a set of sites or their status
 - ❑ USATLAS needed some “summary interface”
- ❄ Thanks to Tom Wlodek’s work at BNL on developing a “modular dashboard” we have a very nice way to summarize the extensive information being collected for the near-term network characterization.
- ❄ The dashboard provides a highly configurable interface to monitor a set of perfSONAR-PS instances via simple plug-in test modules. Users can be authorized based upon their grid credentials. Sites, clouds, services, tests, alarms and hosts can be quickly added and controlled.

Modular Dashboard: Collecting Data



Red lines represent queries from the dashboard collector that gather test results from the sites. This allows us to summarize and visualize the USATLAS perfSONAR status

MODULAR Dashboard at BNL

Example of Dashboard for USATLAS



RACF
Grid Group

The Experimental Independent perfSONAR Dashboard

Status as of: Mon May 14 10:56:50 EDT 2012

“Primitive” service status

- Main Page
- All Clouds
- Individual Clouds:
 - USATLAS
 - USCMS
 - IT
 - LHCOPN
 - LHCONE
 - CA-ATLAS
 - UK
- Inter Cloud Tests:
 - AGLT2-IT
 - FR-US
- Primitive Services
- perfsonar Sites
- List of Hosts
- List of Matrices
- List of Alarms
- List of Clouds
- List of Sites
- List of Schedulers
- Administrator Page
- Manage Users
- Define or Edit Alarms
- RACF dashboard
- perfSONAR dashboard (old)
- RACF dashboard (test)
- perfSONAR dashboard (old,test)
- Dashboard documentation

Cloud USATLAS

Sites of USATLAS cloud

BNL	AGLT2	MWT2	NET2	SWT2	WT2
MWT2-UIUC					

US cloud throughput measurement

	---	0	1	2	3	4	5	6	7	8	9
0:BNL (lhomon.bnl.gov)	---	0.00	0.00	0.78	0.81	0.77	0.00	0.82	0.88	0.78	
1:AGLT2 (psmsu02.aglt2.org)	0.00	---	0.00	0.59	0.54	0.54	0.00	0.53	0.51	0.57	
2:AGLT2 (psum02.aglt2.org)	0.52	0.44	---	0.38	1.40	0.45	0.17	0.20	0.11	0.33	
3:MWT2 (iut2-net2.iu.edu)	0.85	0.30	2.28	---	3.61	0.00	0.12	0.16	0.04	0.00	
4:MWT2 (uct2-net2.uchicago.edu)	0.00	0.44	0.39	0.94	---	0.00	0.14	0.13	0.10	0.00	
5:MWT2-UIUC (mwt2-ps02.campuscluster.illinois.edu)	0.90	0.94	0.93	0.94	0.94	---	0.92	0.93	0.91	0.90	
6:NET2 (atlas-npt1.bu.edu)	0.00	0.91	0.92	0.93	0.93	0.00	---	0.91	0.00	0.87	
7:SWT2 (ps2.ochep.pu.edu)	0.00	0.92	0.91	0.00	0.91	0.91	---	0.90	0.90	0.90	
8:SWT2 (netmon1.atlas-sw2.org)	0.88	0.00	0.91	0.91	0.91	0.00	0.90	0.93	---	0.81	
9:WT2 (psnr-bw01.slac.stanford.edu)	0.53	0.69	0.74	0.59	0.59	0.00	0.58	0.80	0.46	---	

The rows of this table represent SOURCE nodes for a throughput test while the columns represent DESTINATION nodes. Each cell in the table contains the result of two versions of a BWCTL throughput test for the specified source and destination. Tests are configured to run by BOTH the source and destination once every 4 hour period. The upper link in each cell represents the results of the throughput test initiated from the SOURCE end. The lower link in each cell represents the results of the throughput test initiated from the DESTINATION end.

US Cloud Owamp Packet Loss Measurement

		0	1	2	3	4	5	6	7	8	9
0:BNL (lhopermon.bnl.gov)		0.0	0.0	0.0	1.0	0.0	0.0	0.0	0.0	0.0	0.0
1:AGLT2 (psmsu01.aglt2.org)		0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
2:AGLT2 (psum01.aglt2.org)		0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
3:MWT2 (iut2-net1.iu.edu)		0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
4:MWT2 (uct2-net1.uchicago.edu)		0.0	0.0	1.0	0.0	0.0	1.0	0.0	2.0	0.0	0.0
5:MWT2-UIUC (mwt2-ps01.campuscluster.illinois.edu)		1.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	1.0	0.0
6:NET2 (atlas-npt1.bu.edu)		0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
7:SWT2 (ps1.ochep.pu.edu)		0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
8:SWT2 (netmon1.atlas-sw2.org)		0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
9:WT2 (psnr-lat01.slac.stanford.edu)		0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0

The rows of this table represent SOURCE nodes for a test while the columns represent DESTINATION nodes. Each cell in the table represents a source-destination LATENCY test via OWAMP (600 UDP packets/test) tests, 1/minute. The metric we are plotting is the packet loss between the source and destination averaged over the last 30 minutes. Each cell contains the result of two tests:

See <https://perfsonar.usatlas.bnl.gov:8443/exda/?page=25&cloudName=USATLAS>

perfSONAR Latency Monitoring



OWAMP (latency) tests send 600 pkts/minute to each site/node under test

For a given pair of nodes tests are run in **both** directions by **both** nodes.

Rows represent “source” site, columns are the “destination”. The numbers are the packet-loss (out of 600)

The top numbers in each row is the result of “source” tests while the bottom numbers is from “destination” tests.

Status based upon packet loss **average over 30 minutes**: **GREEN** is <2 pkt-loss out of 600, **YELLOW** is $2 \leq$ pkt-loss, **RED** is >10 pkt-loss OR test not defined OR error returned, **GREY** is timeout on query and **BROWN** is no response, unknown, garbled response, etc.

US Cloud Owamp Packet Loss Measurement

	---	0	1	2	3	4	5	6	7	8	9
0:BNL (lhccperfmon.bnl.gov)		0.0 0.0	0.0 0.0	0.0 0.0	1.0 0.0	0.0 0.0	0.0 0.0	0.0 0.0	0.0 0.0	0.0 1.0	0.0 0.0
1:AGLT2 (psmsu01.aglt2.org)		0.0 0.0	0.0 0.0	0.0 0.0	0.0 1.0	0.0 0.0	0.0 0.0	0.0 0.0	0.0 0.0	0.0 0.0	0.0 0.0
2:AGLT2 (psum01.aglt2.org)		0.0 0.0	0.0 1.0	0.0 1.0	0.0 0.0						
3:MWT2 (iut2-net1.iu.edu)		0.0 2.0	0.0 0.0	0.0 0.0	0.0 0.0	0.0 0.0	0.0 0.0	0.0 0.0	1.0 0.0	0.0 0.0	1.0 0.0
4:MWT2 (uct2-net1.uchicago.edu)		0.0 0.0	0.0 0.0	0.0 0.0	0.0 1.0	0.0 0.0	1.0 0.0	0.0 0.0	2.0 1.0	0.0 1.0	0.0 1.0
5:MWT2-UIUC (mwt2-ps01.campuscluster.illinois.edu)		0.0 1.0	0.0 0.0	0.0 0.0	0.0 0.0	0.0 0.0	0.0 0.0	0.0 0.0	0.0 1.0	0.0 0.0	0.0 0.0
6:NET2 (atlas-npt1.bu.edu)		0.0 0.0	0.0 0.0	0.0 0.0	0.0 1.0	0.0 0.0	0.0 0.0	0.0 0.0	1.0 1.0	0.0 0.0	0.0 0.0
7:SWT2 (ps1.ochep.ou.edu)		0.0 1.0	0.0 0.0	0.0 0.0	3.0 3.0	0.0 0.0	0.0 0.0	0.0 0.0	0.0 0.0	0.0 0.0	0.0 0.0
8:SWT2 (netmon1.atlas-swt2.org)		0.0 0.0									
9:WT2 (psnr-lat01.slac.stanford.edu)		0.0 0.0	0.8 0.0	0.0 0.0	0.0 0.0	0.0 0.0	0.0 0.0	0.0 0.0	0.0 0.0	1.0 0.0	0.0 0.0

perfSONAR Throughput Monitoring



BWCTL (throughput) tests run iperf once per 4 hour window (non-concurrent)

For a given pair of nodes tests are run in **both** directions by **both** nodes.

Rows represent “source” site, columns are the “destination”. The numbers are the achievable bandwidth in a 20 sec test in units of Gbps

The top numbers in each row is the result of “**source**” tests while the bottom numbers is from “**destination**” tests.

Status based upon throughput **average over 24 hours**: **GREEN** is >100 Mbps, **YELLOW** is > 10 Mbps & < 100 Mbps, **RED** is <10 Mbps OR test not defined OR error returned, **GREY** is timeout on query and **BROWN** is no response, unknown, garbled response, etc.

US Cloud Throughput Measurement

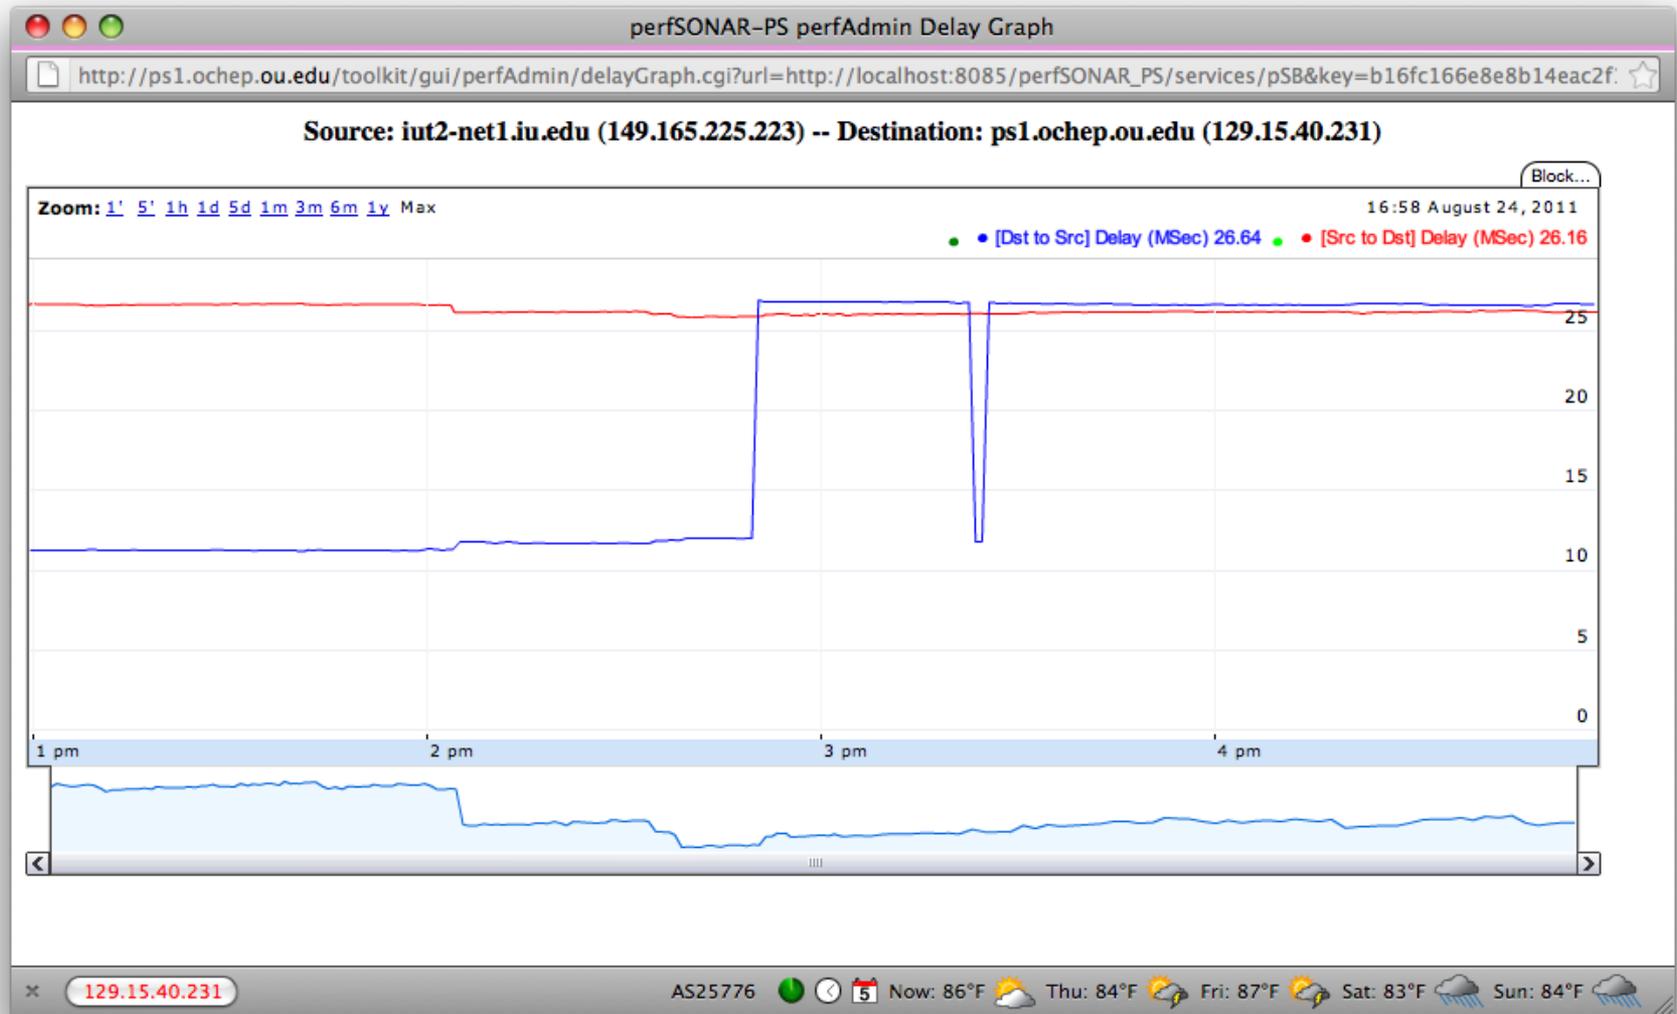
	---	0	1	2	3	4	5	6	7	8	9
0:BNL (lhcmn.bnl.gov)	---	0.90 0.00	0.41 0.76	0.00 0.59	0.44 0.31	0.00 0.00	0.86 0.77	0.00 0.76	0.53 0.45	0.78 0.49	
1:AGLT2 (psmsu02.aglt2.org)	0.85 0.90	---	0.00 0.00	0.92 0.66	0.91 0.70	0.94 0.94	0.92 0.92	0.93 0.92	0.92 0.91	0.84 0.64	
2:AGLT2 (psum02.aglt2.org)	0.68 0.32	0.94 0.93	---	1.60 1.58	0.29 0.00	0.59 0.58	0.21 0.00	0.19 0.20	0.09 0.11	0.30 0.28	
3:MWT2 (iut2-net2.iu.edu)	0.44 0.69	0.31 0.30	0.54 1.31	---	5.45 5.20	0.83 0.00	0.16 0.11	0.22 0.21	0.07 0.07	0.00 0.26	
4:MWT2 (uct2-net2.uchicago.edu)	0.24 0.00	0.00 0.66	0.40 0.00	0.80 2.26	---	0.64 0.87	0.10 0.12	0.07 0.15	0.07 0.05	0.17 0.25	
5:MWT2-UIUC (mwt2-ps02.campuscluster.illinois.edu)	0.84 0.00	0.94 0.00	0.93 0.93	0.91 0.00	0.99 0.98	---	0.00 0.00	0.00 0.93	0.92 0.00	0.90 0.00	
6:NET2 (atlas-npt2.bu.edu)	0.94 0.91	0.92 0.92	0.00 0.89	0.00 0.89	0.22 0.00	0.00 0.90	---	0.91 0.83	0.86 0.89	0.00 0.52	
7:SWT2 (ps2.oceph.ou.edu)	0.00 0.00	0.00 0.92	0.92 0.00	0.00 0.00	0.59 0.52	0.92 0.00	0.91 0.91	---	0.88 0.88	0.89 0.73	
8:SWT2 (netmon2.atlas-swt2.org)	0.86 0.88	0.92 0.92	0.92 0.92	0.00 0.63	0.35 0.24	0.00 0.92	0.91 0.91	0.91 0.93	---	0.89 0.61	
9:WT2 (psnr-bw01.slac.stanford.edu)	0.56 0.84	0.71 0.80	0.66 0.88	0.19 0.00	0.23 0.19	0.00 0.85	0.50 0.87	0.74 0.86	0.64 0.00	---	

Example Performance Debugging

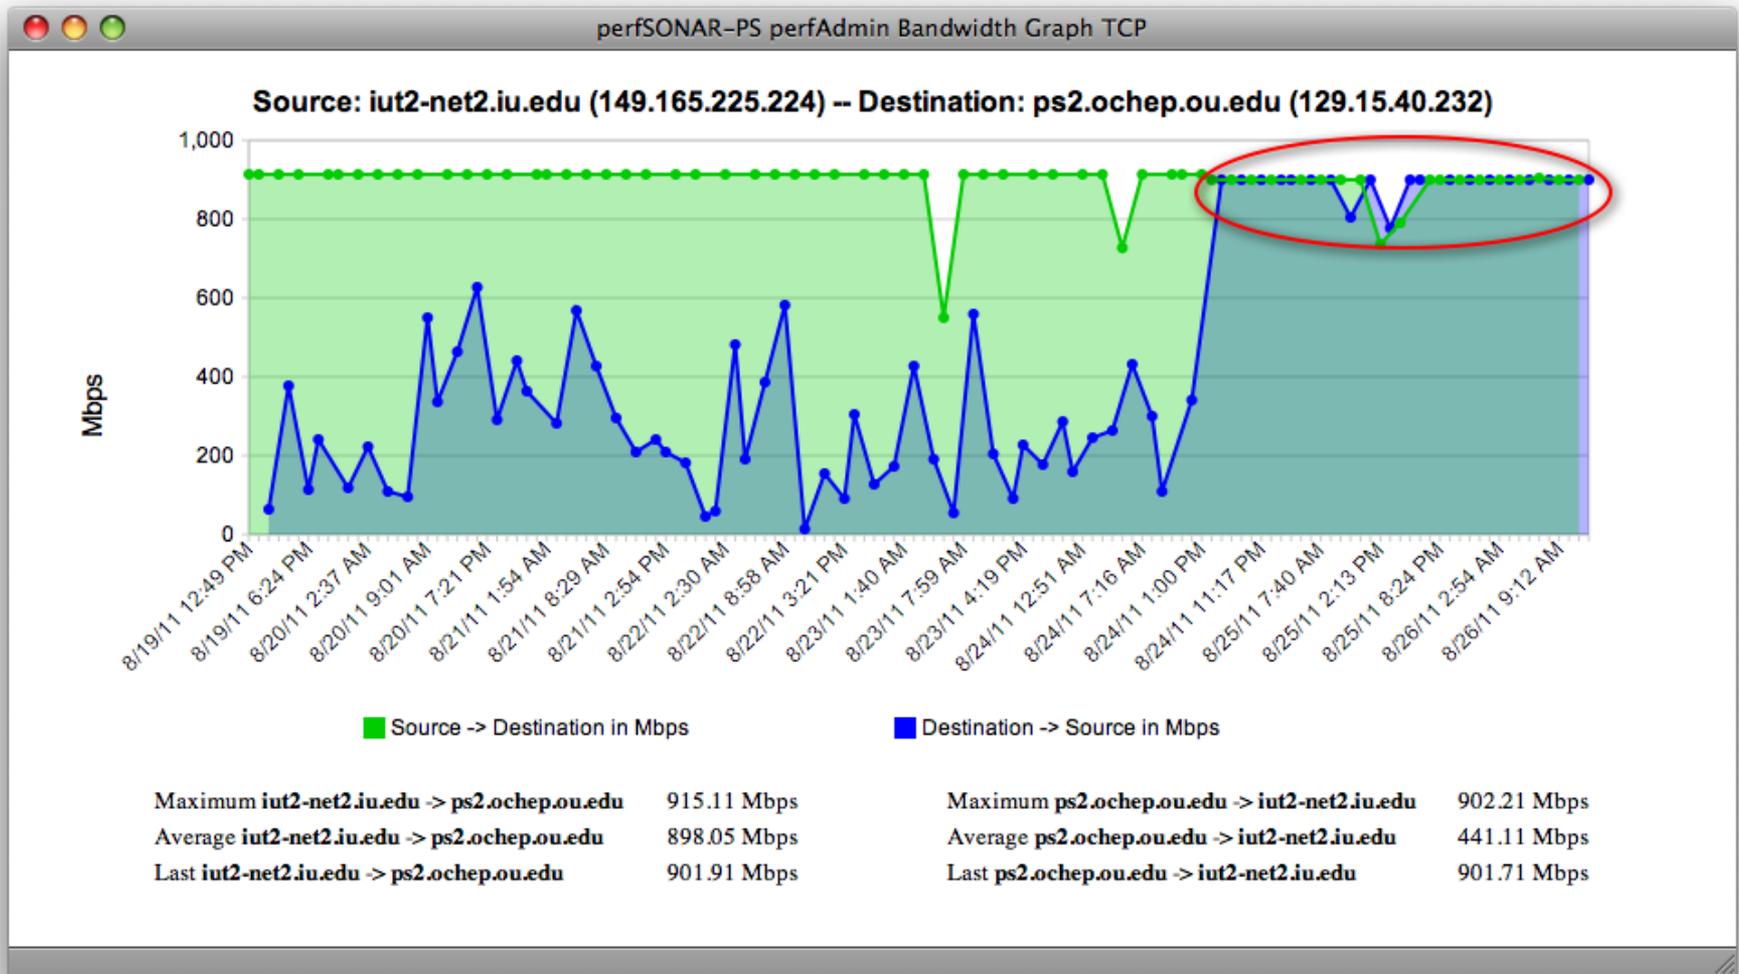


- ❄ The combination of lightweight latency tools with the heavier bandwidth based measurement (i.e. a simulation of a well tuned data movement application) has resulted in the discovery of several serious performance abnormalities between members of the USATLAS collaboration.
- ❄ In the next two slides I will show the observed latency measurements between two USATLAS sites (the University of Oklahoma and Indiana University) and a snapshot of bandwidth observations between the same two facilities.
- ❄ Path was initially asymmetric and showed a problem. Routing was fixed and the performance returns to a symmetric behavior after the routing change, thus implying that one of the paths contained a problem in need of further investigation.

OU/IU – Latency After Routing Change



OU/IU – BW One Week Later



Challenges Ahead



- ❄ Getting hardware/software platform installed at all sites
- ❄ **Dashboard development:** Currently USATLAS/BNL and soon OSG, Canada (ATLAS, HEPnet) and USCMS. More ?
- ❄ **Managing site and test configurations**
 - ❑ Determining the right level of scheduled tests for a site, e.g., Tier-2s test to other same-cloud Tier-2s (and Tier-1)?
 - ❑ Improving the management of the configurations for VOs/Clouds
 - ❑ Tools to allow “central” configuration
- ❄ **Alerting: A high-priority need but complicated:**
 - ❑ Alert who? Network issues could arise in any part of end-to-end path
 - ❑ Alert when? Defining criteria for alert threshold. Primitive services are easier. Network test results more complicated to decide
- ❄ **Integration with VO infrastructures.**

Summary and Conclusions



- ❄ To support the growing scale of global scientific projects, network operators and VOs alike must be cognizant of network performance considerations to assure proper operation.
- ❄ **Frameworks**, such as the pS Performance Toolkit, are capable of monitoring internal and external network performance metrics when properly deployed and managed.
- ❄ **Presentation layers**, such as the USATLAS dashboard, can deliver the raw results of this performance assurance in an easy to use and interpret format.
- ❄ This holistic approach to network measurement has resulted in the correction of numerous performance abnormalities, and saved the time and resources of strained operations staff.



Questions or Comments?