Network and Transfer WG Metrics Area Meeting

Shawn McKee, Marian Babik

Network and Transfer Metrics Kick-off Meeting
26h November 2014

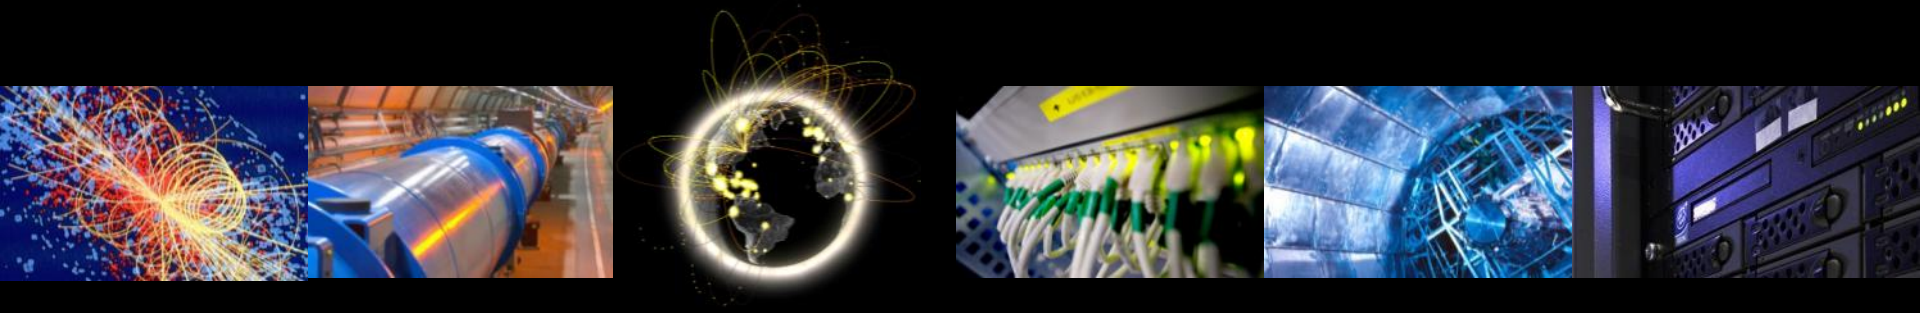

WLCG
Worldwide LHC Computing Grid

Outline

- Status and progress in perfSONAR
  - T2.1 Commissioning/Operations
  - T2.2 Storage
  - T2.3 Configuration
- Metrics area
  - T1.1: Gather requirements and use cases
  - T1.2: Review existing transfer and network metrics
  - T1.3: Determine current test coverage
  - T1.4: Topology mapping

Actions from last meeting

- ALL: Send comments and suggestions on the proposed list of topics/tasks and on the way WG will be organized
- ALL: Volunteer to lead tasks in the metrics area (T1s)
- Julia: Send a list of topics concerning xRootD tasks to the WG. To be discussed with WLCG OPS Coordination
  - Separate meeting held on the topic – agreed to follow up on status of GLED deployment and support
- Marian: Setup WG JIRA and report to WLCG OPS coordination every 2 weeks on the status of ongoing tasks.
  - JIRA at https://its.cern.ch/jira/browse/METRICS
  - WLCG OPS Reports at https://twiki.cern.ch/twiki/bin/view/LCG/NetworkTransferMetrics#Reports
- Marian, Shawn: Prepare abstract for CHEP2015 (done)


Network Monitoring Status perfSONAR

perfSONAR ops

- perfSONAR 3.4 released Oct 14th
- Restructuring support and operations
  - Introduced site-level support via GGUS
- Rewritten documentation
  - https://twiki.opensciencegrid.org/bin/view/Documentation/DeployperfSONAR
- Responded to ShellShock and Poodle
  - Sites advised to terminated their instances
  - Performed security audit and established security procedures
- Testing and validation of the new perfSONAR central configuration
- perfSONAR 3.4 update campaign
  - Includes migration to the new configuration system
  - Security considerations documented
  - Progressing well (111 sonars updated out of 214)
  - Deadline 8th January

perfSONAR config and store

- Deployed in OSG production
- Introduces central interface to reconfigure the entire network
  - All aspects – tests parameters, mesh participation
  - List of available sonars taken from GOCDB and OIM
  - Supports hierarchical support model (per mesh admins)
  - Web interface
  - Connected to OSG crawler and perfSONAR infrastructure monitoring
- Site reconfiguration needed to adopt
  - Run as part of 3.4 campaign
- perfSONAR data store
  - Deployed in OSG ITB – several major issues fixed
  - Scale tests on-going this week
  - Operationally ready for production

perfSONAR metrics

Metrics and Their Use

- We gather a number of metrics:
  - Topology/path-information via traceroute
  - One-way delay via OWAMP
  - Packet-loss via OWAMP
  - Usable bandwidth via BWCTL
- WLCG perfSONAR coverage
  - http://grid-monitoring.cern.ch/perfsonar_coverage.txt
- ESnet has some nice pages on using perfSONAR to identify problems
  - http://fasterdata.es.net/performance-testing/evaluating-network-performance/
- Some examples from Jason Zurawski follow

Traceroute

- Traceroute is fundamental to any of the other metrics. Without it we don't know what path was being measured. In the toolkit:

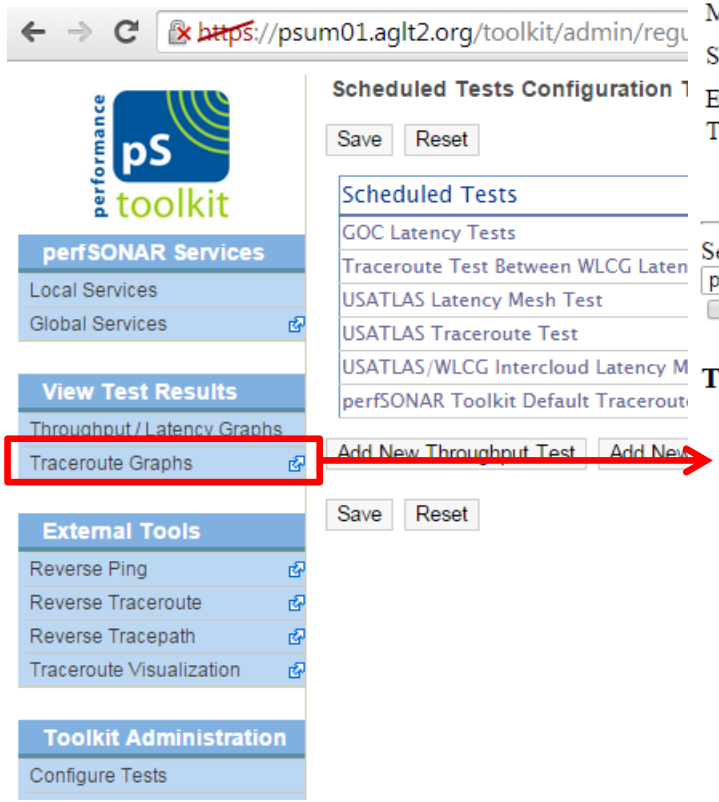

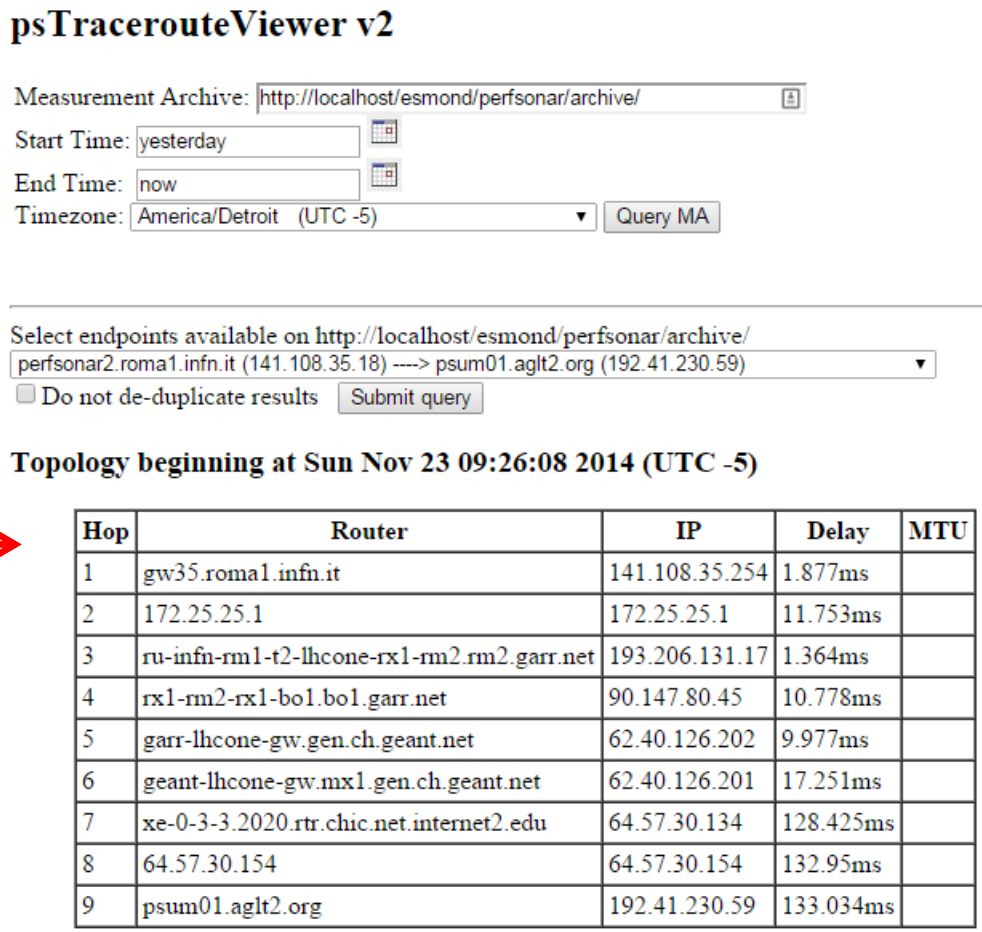

Hop	Router	IP	Delay	MTU
1	gw35.roma1.infn.it	141.108.35.254	1.877ms	
2	172.25.25.1	172.25.25.1	11.753ms	
3	ru-infn-rm1-t2-lhcone-rx1-rm2.rm2.garr.net	193.206.131.17	1.364ms	
4	rx1-rm2-rx1-bo1.bo1.garr.net	90.147.80.45	10.778ms	
5	garr-lhcone-gw.gen.ch.geant.net	62.40.126.202	9.977ms	
6	geant-lhcone-gw.mx1.gen.ch.geant.net	62.40.126.201	17.251ms	
7	xe-0-3-3.2020.rtr.chic.net.internet2.edu	64.57.30.134	128.425ms	
8	64.57.30.154	64.57.30.154	132.95ms	
9	psum01.aglt2.org	192.41.230.59	133.034ms	

OWAMP: Overloaded Link

Example of a campus with an overloaded uplink to the WAN. You can see daily overloads as campus gets busy each day

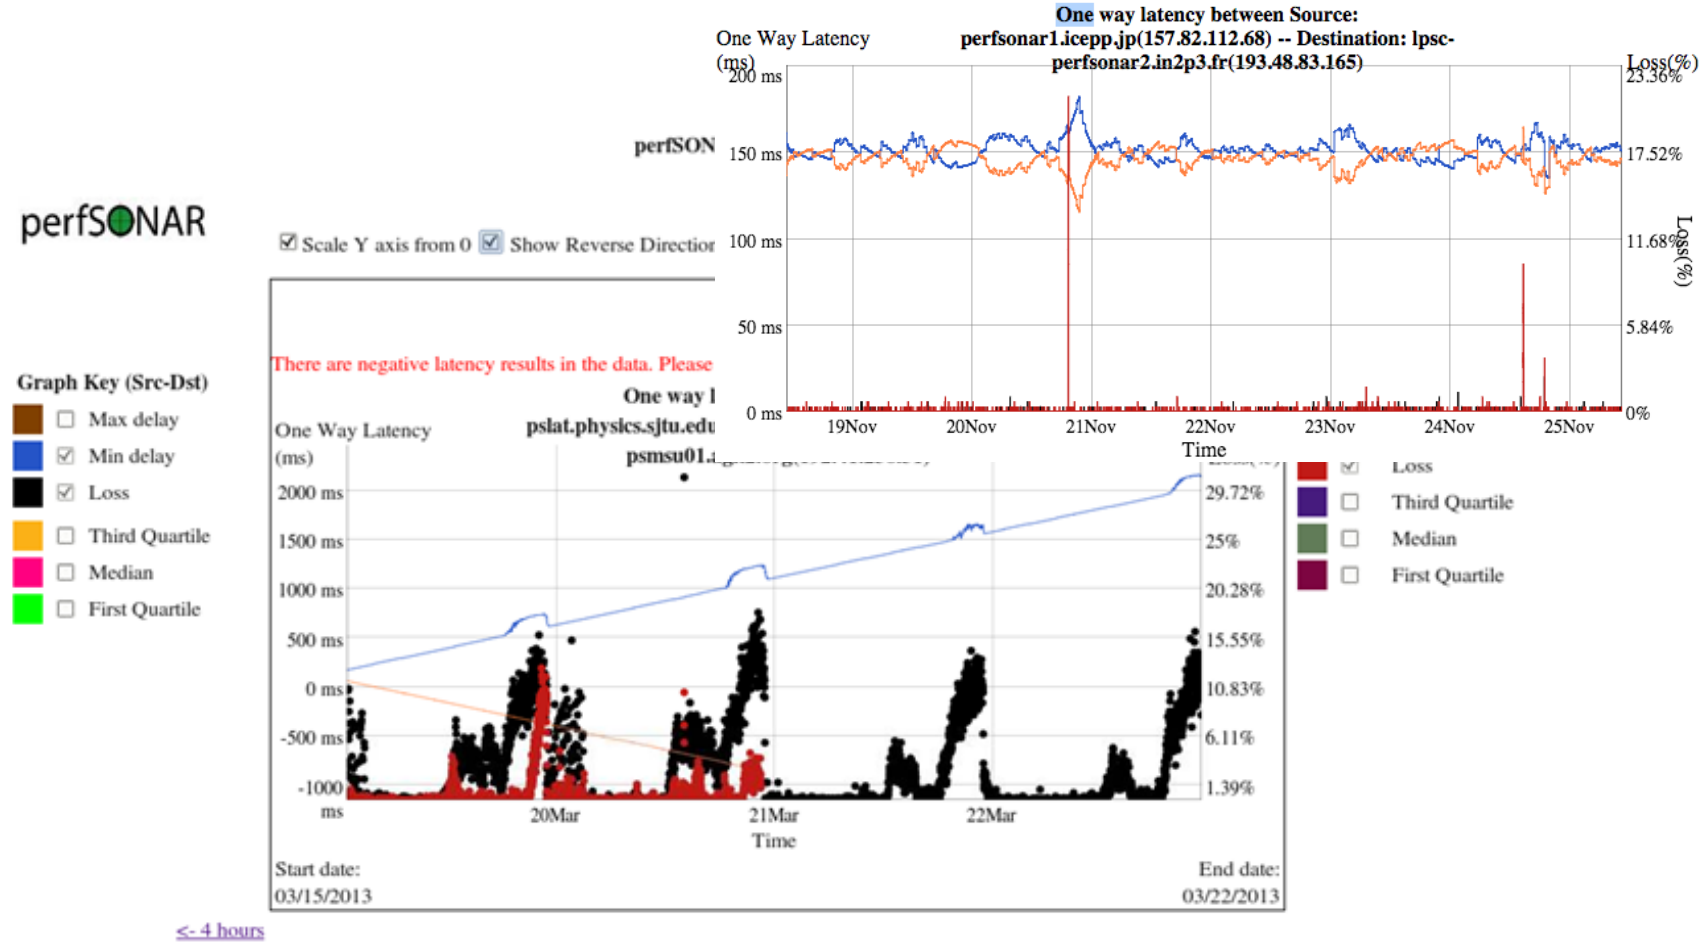

Bandwidth (iperf)

perfSONAR measures usable bandwidth
Bandwidth changes as MTU and window setting adjusted

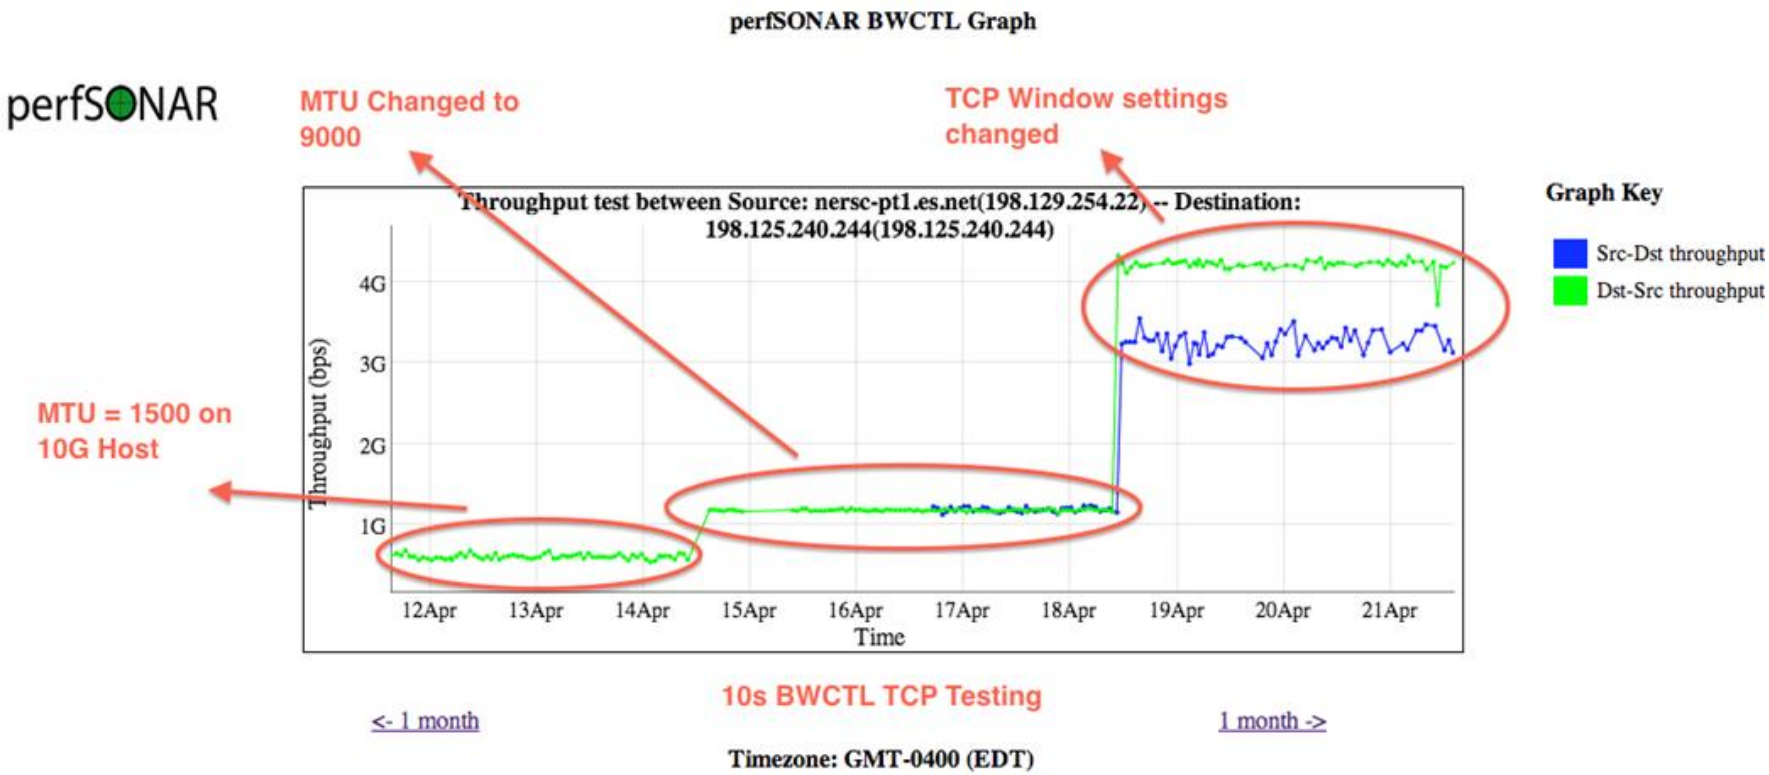

Bandwidth: Bad vs Good Routing

perfSONAR

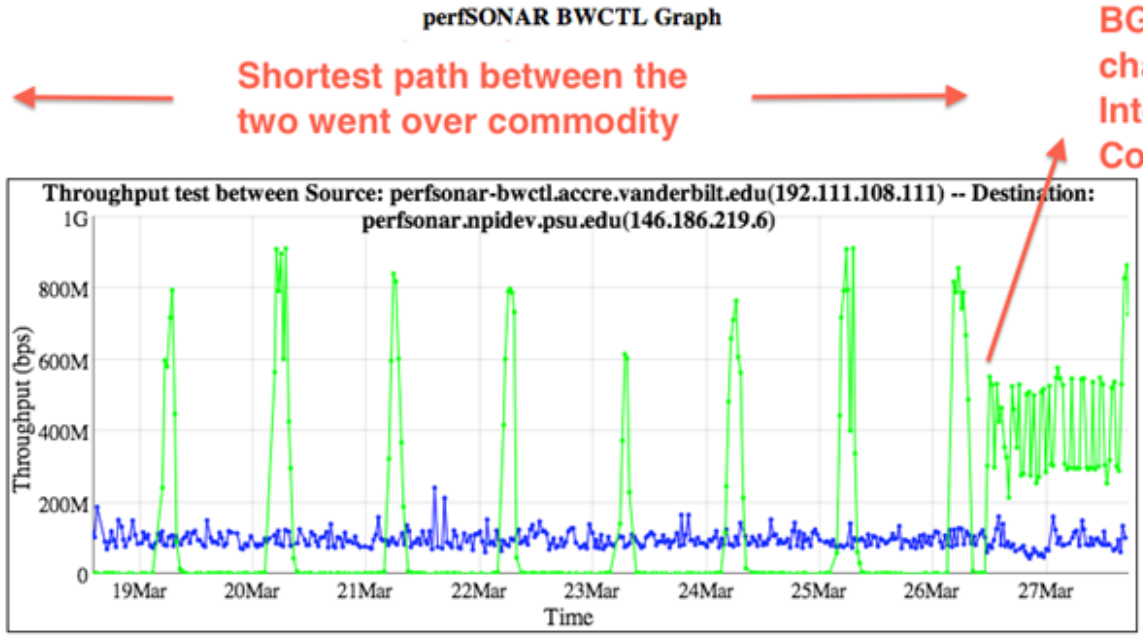

<- 1 month

1 month ->

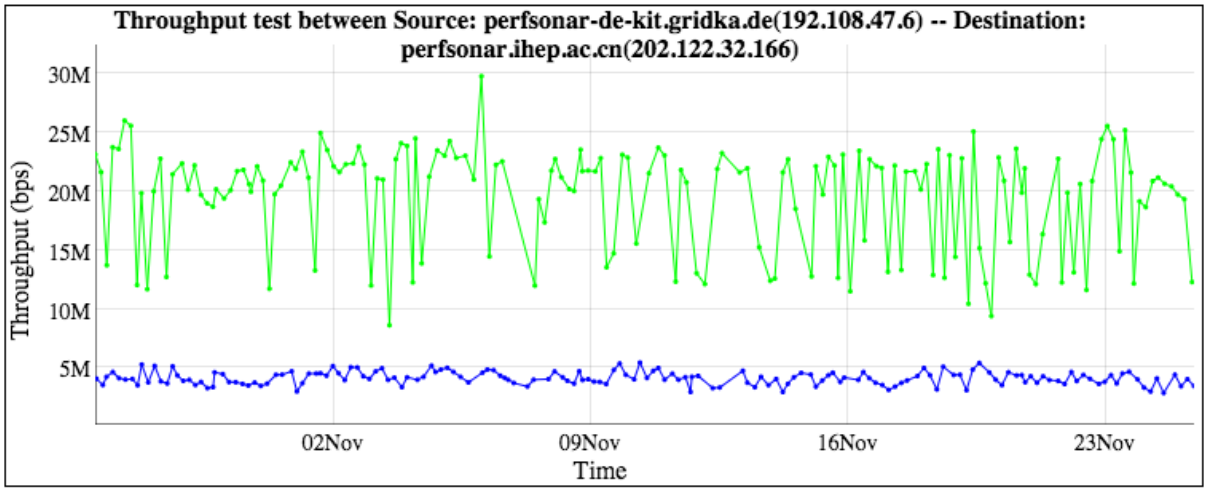

Drastic BW Change

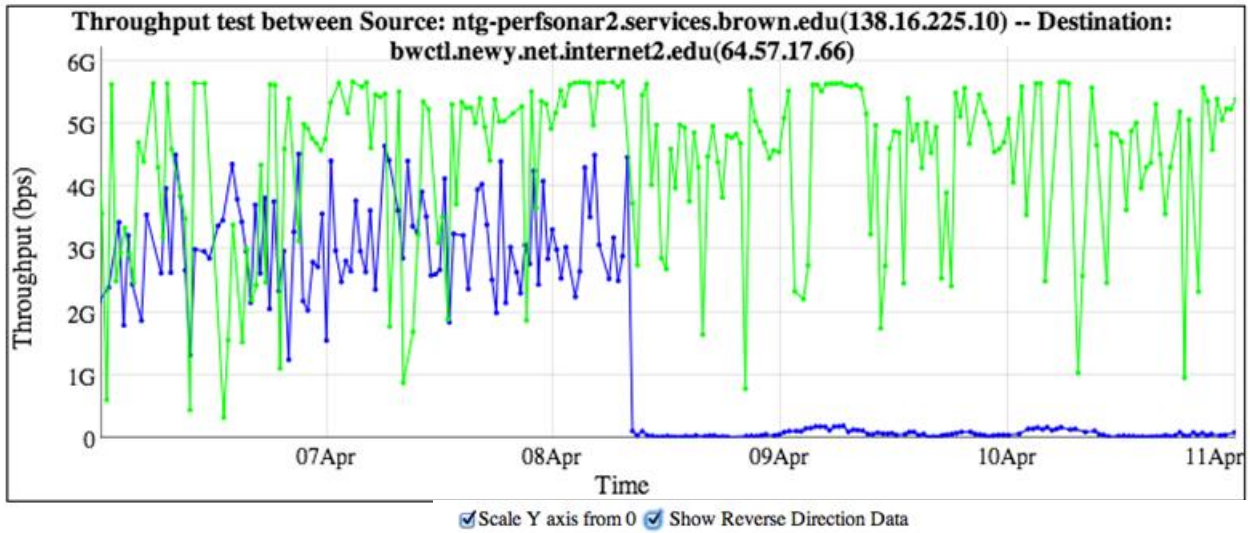

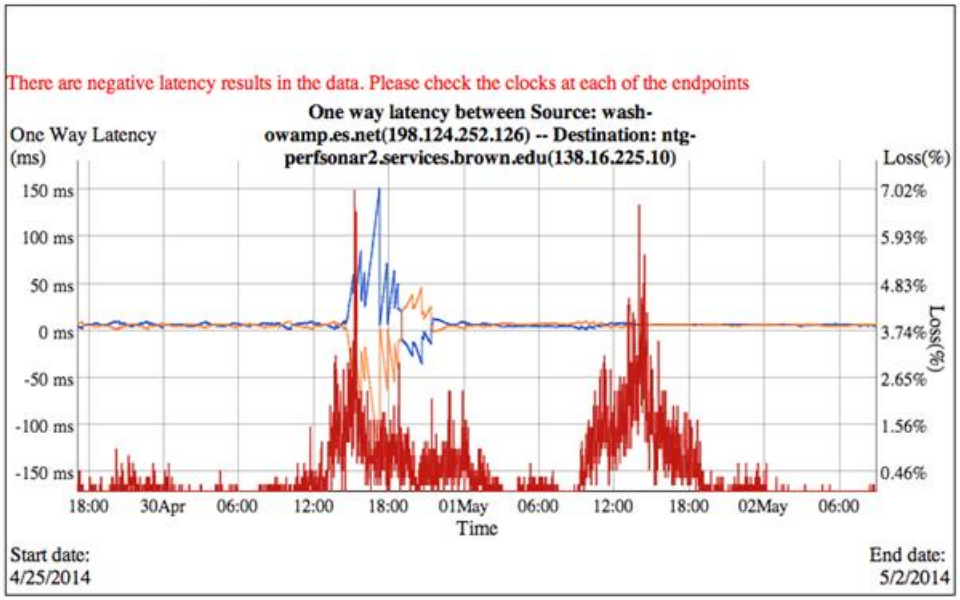

- Spikes of packet loss, almost always during business hours
- Function of the load on the line/time of day
- This was traced to **regional network**

Key Messages about Metrics

- The monitoring infrastructure is sensitive for a reason – so that it finds the problems in all layers of the OSI stack.
- End-to-end data transmission (or just about any other use case) suffers because of a problems that may be unseen or not understood.
- Understanding comes from learning to use the tools, learning to trust them, and having universal availability.
- Comprehensive solutions will save time in the end

Metrics Area

Questionnaire

- Available as Google document
  - https://docs.google.com/document/d/1ceiNlTUJCwSuOuvbEHZnZp0XkWkwdkPQTQic0VbH1mc/edit?usp=sharing
- Asking for your input
  - FTS, FAX, PhEDEx, Rucio, PanDA by Dec 5
  - Experiments by Dec 12
- Next year
  - Strawman – important to receive initial input on coverage, test characteristics, etc.
  - Regular meetings next year
    - 28 Jan, 18 Feb, 18 March, 8 Apr (all at 4pm CEST)

AOB

- Shawn at CERN next week
  - perfSONAR office on Thursday (4th Dec)

Backup

iperf INFN PIC

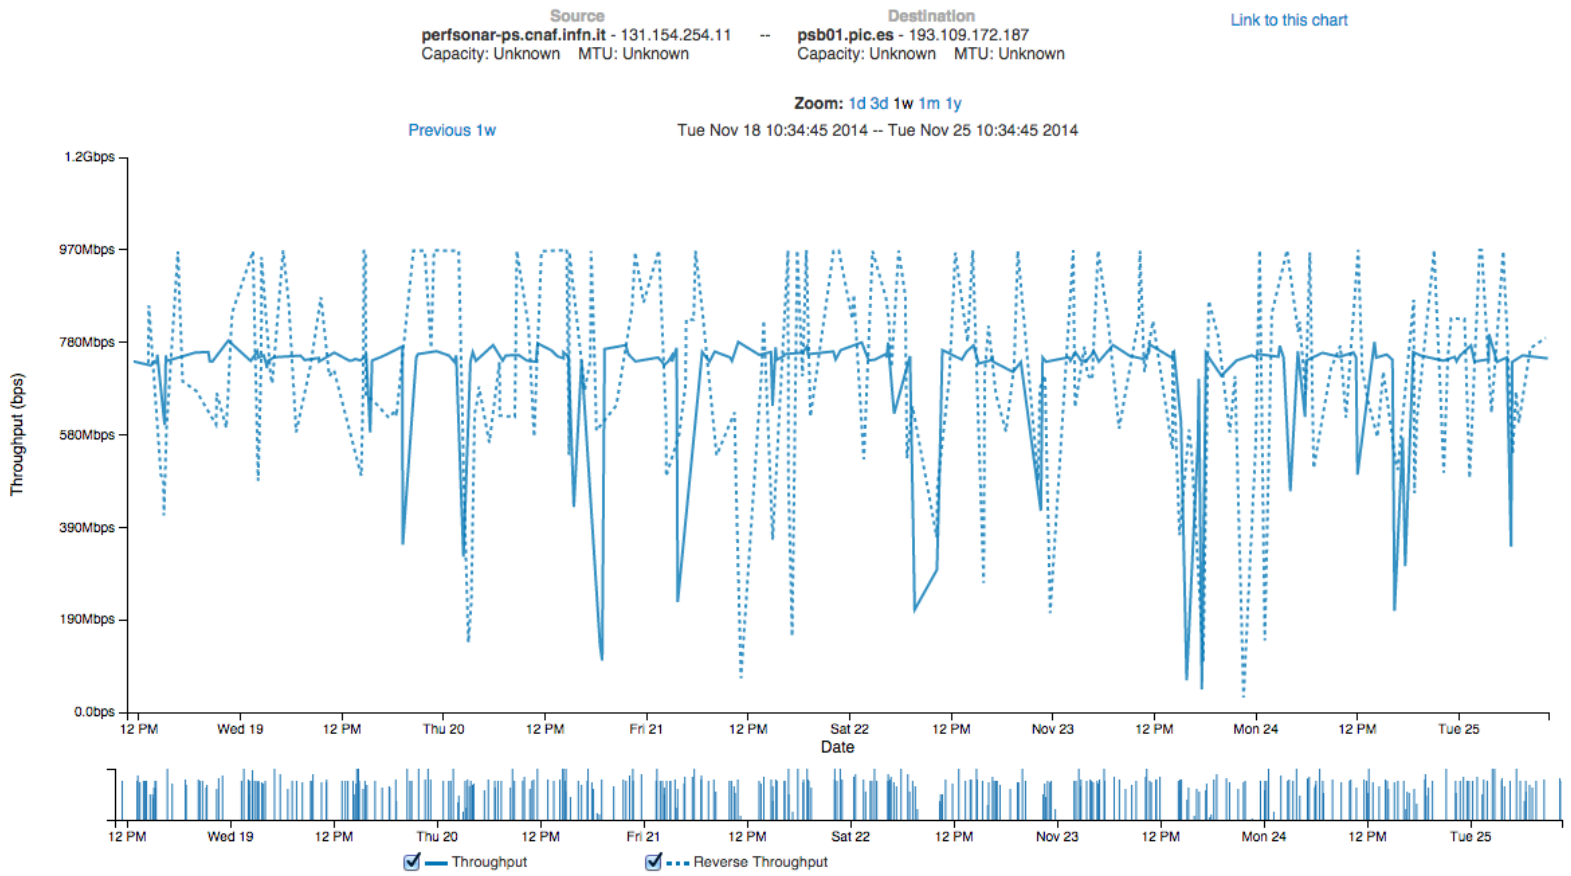

owamp+iperf INFN PIC