



June 14th, LHCOPN/LHCONE Meeting

Jason Zurawski – Research Liaison

LHCONE Monitoring Thoughts

LHCONE Monitoring

- LHCONE Monitoring = Hard Problem(?)
- Current Use Case: USATLAS
- Possible Solutions



Hard Problem To Solve?

- Participants
 - Multiple Domains (e.g. Building, Campus, Regional, Backbone, Exchange Point)
 - Multiple Parties (e.g. VO management, Local/Regional/National IT staff)
- Technologies
 - Monitoring at all layers of the OSI stack (e.g. light levels all the way up to application performance)
- Governance
 - Conversations about this over last 2 days – who runs LHCONe? Can someone enforce monitoring rules?
 - Value add: installation of monitoring tools and someone to ensure they work
 - Some central facility to manage the tickets/process?

LHCONE Monitoring

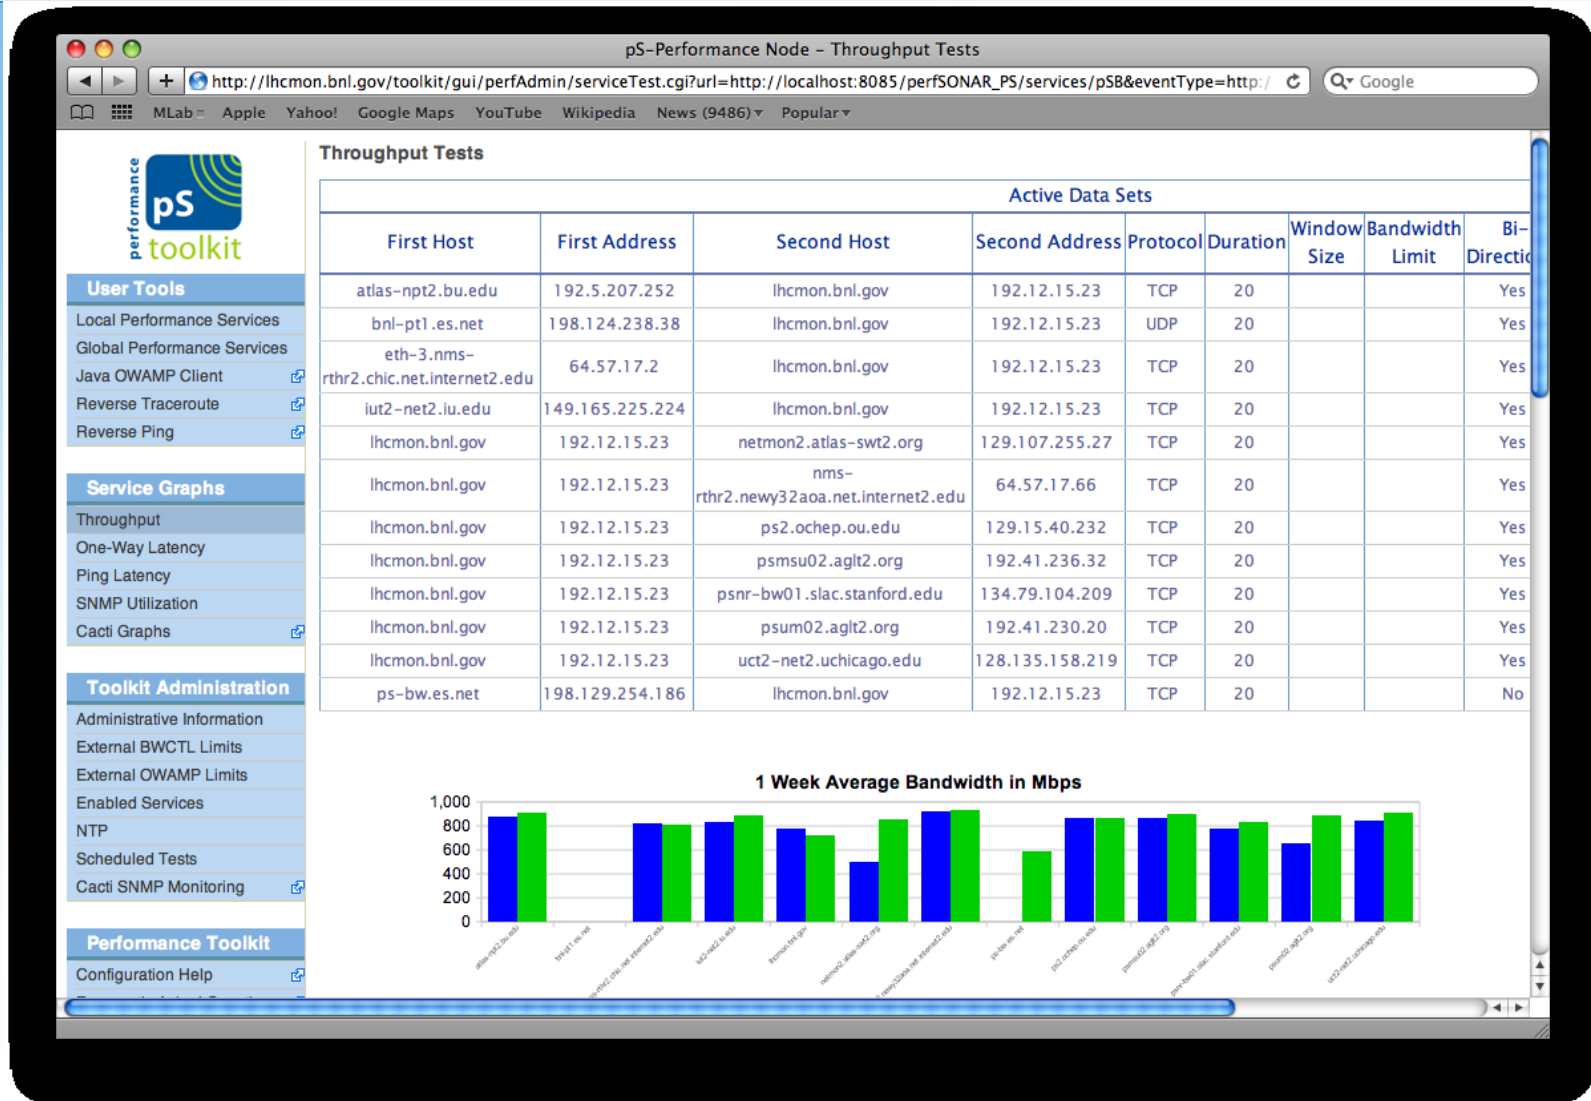
- LHCONE Monitoring = Hard Problem(?)
- **Current Use Case: USATLAS**
- Possible Solutions




USATLAS

- Hardware/Software
 - perfSONAR-PS Performance Toolkit (<http://psps.perfsonar.net/toolkit>)
 - 2 Dedicated Machines per T1 and T2 (Bandwidth and Latency Monitoring)
- Use Case
 - Regular full mesh testing (OWAMP/BWCTL/PingER)
 - Diagnostic tools on demand (NDT/NPAD)
 - Alarms built using NAGIOS
 - Throughput drops below threshold
 - Loss/Latency increase beyond threshold
 - Monitoring hosts/services become unreachable

USATLAS – Setting Up Tests



USATLAS – Simple Graph

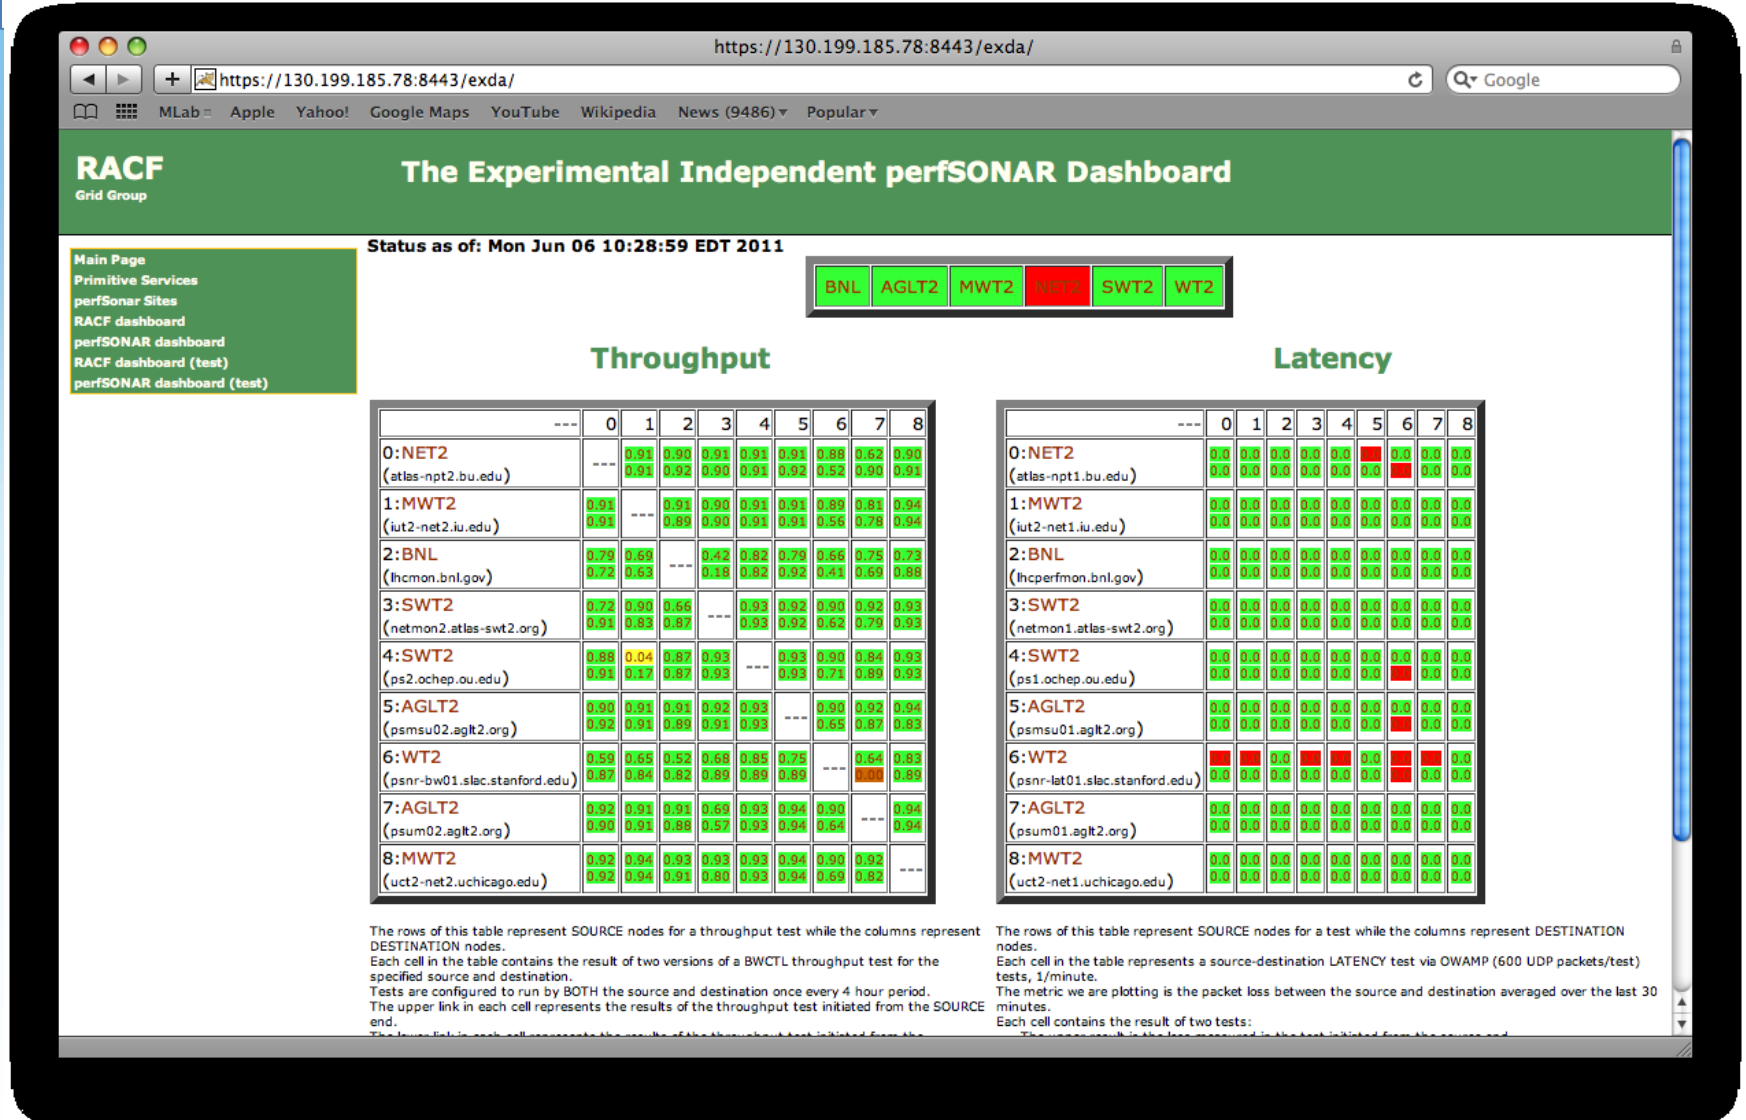


USATLAS

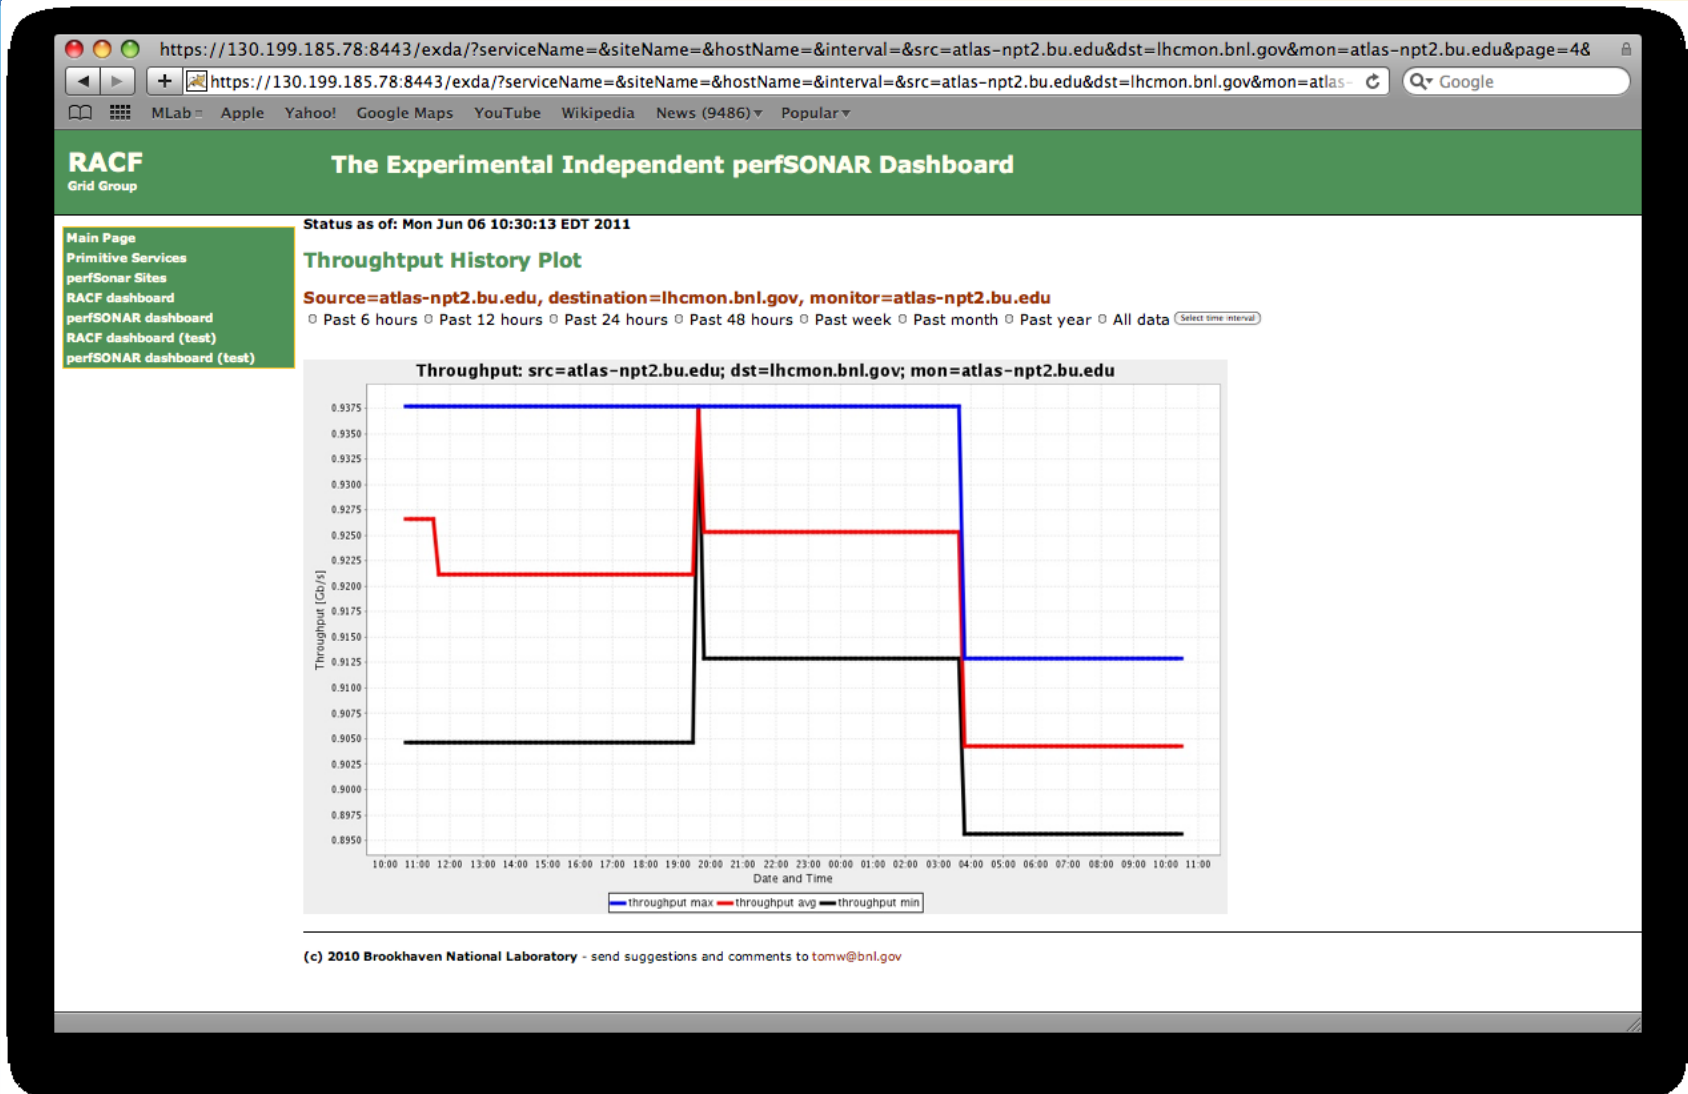
- Implementing other Components
 - Dashboard
 - Python based, integrated with perfSONAR-PS NAGIOS probes.
 - Web Service calls to remote instances to gain status info
 - Developed by BNL for USATLAS
 - Integration into data movement software
 - Still in pipe-dream phase – use perfSONAR-PS APIs to get data from monitoring hosts
 - Intelligent decisions about data movement (e.g. who to download from ala bit-torrent, or when to start a dynamic circuit vs use IP)



USATLAS – “Complete” View



USATLAS – Per-Pair Performance



USATLAS

- Support Structure
 - “Open Source” Software = “Open Source” support.
 - Community mailing lists, meetings with pS PS engineers for debugging/feature requests
 - Support on installation/upgrading (2 or so times a year) as required
 - No PERT – performance problems are handled by USATLAS with the help of Internet2/ESnet typically organizing testing and resource coordination with peer networks
- Difference vs MDM
 - No Help Desk
 - Machines are under local control only (we don’t maintain persistent login access)
 - Testing is up to the VO, we can help get things started.



LHCONE Monitoring

- LHCONE Monitoring = Hard Problem(?)
- Current Use Case: USATLAS
- Possible Solutions



Potential Solutions

- Similar to USATLAS Approach
 - Mandate direct participants purchase at least 1 (preferable 2) machines for monitoring purposes.
 - Stationed at the network core (near the storage/processing equipment)
 - Bonus deployment at the edge
 - Encourage Backbone/Regional/XP operators to do the same. Harder to enforce outside of the VO...
 - Compile list of desirable functionality (e.g. regular testing, on demand testing, complete OS vs packages, etc.).
 - Market based study of what is available vs what could be developed.
 - Form (or use this) WG to serve as community support
 - Installation
 - Configuration
 - Trouble shooting



Potential Solutions

- Possible Enhancements
 - New Software Development
 - LHC Community is not afraid to innovate – do the current solutions in the monitoring space scale? Is anything new needed? Anything need to be changed?
 - PERT/Help Desk
 - Home for the homeless w/ regards to trouble tickets.
 - Work with the networking partners to track progress
 - Handle issues with installation/configuration in the event that the open source model is not sufficient
 - Non-Local Control of monitoring
 - Central authority to own/maintain the infrastructure instead of allowing domains to manage this role



Discussion

