

# One Network Datastore to Track them All

Rob Quick <rquick@iu.edu>

Operations Officer - Open Science Grid

Manager High Throughput Computing – Indiana University

CHEP 2016 - San Francisco

Contributing Partners: Marian Babik (CERN), Shawn McKee (U of Michigan), Edgar Fajardo Hernandez (UCSD), and Soichi Hayashi, Tom Lee, Chris Pipes, and Scott Teige (Indiana University)

- As an international distributed high throughput computing facility the OSG and its stakeholders and peering collaborators are highly dependent on the underlying network infrastructure
- With this in mind we set out to provide an operational infrastructure service that would gather, organize and supply network metrics to the OSG users, collaborators and administrators to assure effective network usage
- In 2012 OSG added a Networking Area to track and support network issues affecting our stakeholders and peering infrastructures
- In late 2015 the internal Data Store, Custom Mesh, and Visualization Components and external perfSONAR clients operational infrastructure was made production by the OSG Network and OSG Operations Groups
- This is the story of that service...

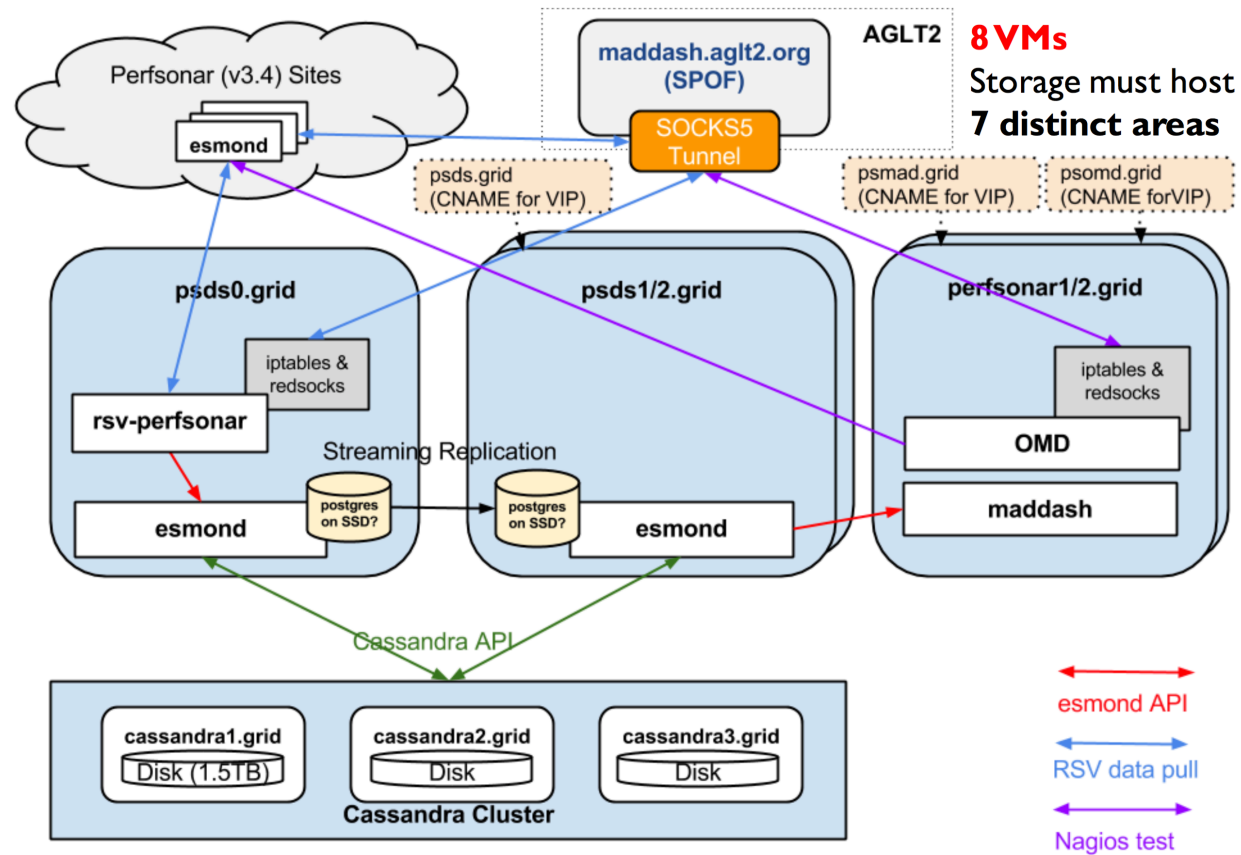
# Goals of the OSG Network Datastore

- To gather, in a single service, all network metrics from OSG and connected projects such as WLCG
- To provide reliable organization of these metrics via a publically accessible interface
- To allow custom visualization of metrics from network subsets (unique meshes)
- To publish the gathered metrics to interested collaborators for network support and research projects
- To provide alert and alarm mechanisms to allow quick info gathering and resolution of network related issues (leveraging PuNDIT)
- **NOTE**: While we heavily lean on a large number of perfSONAR clients to gather this information, this presentation will focus on the data collection, storage, and presentation, not on the perfSONAR clients themselves

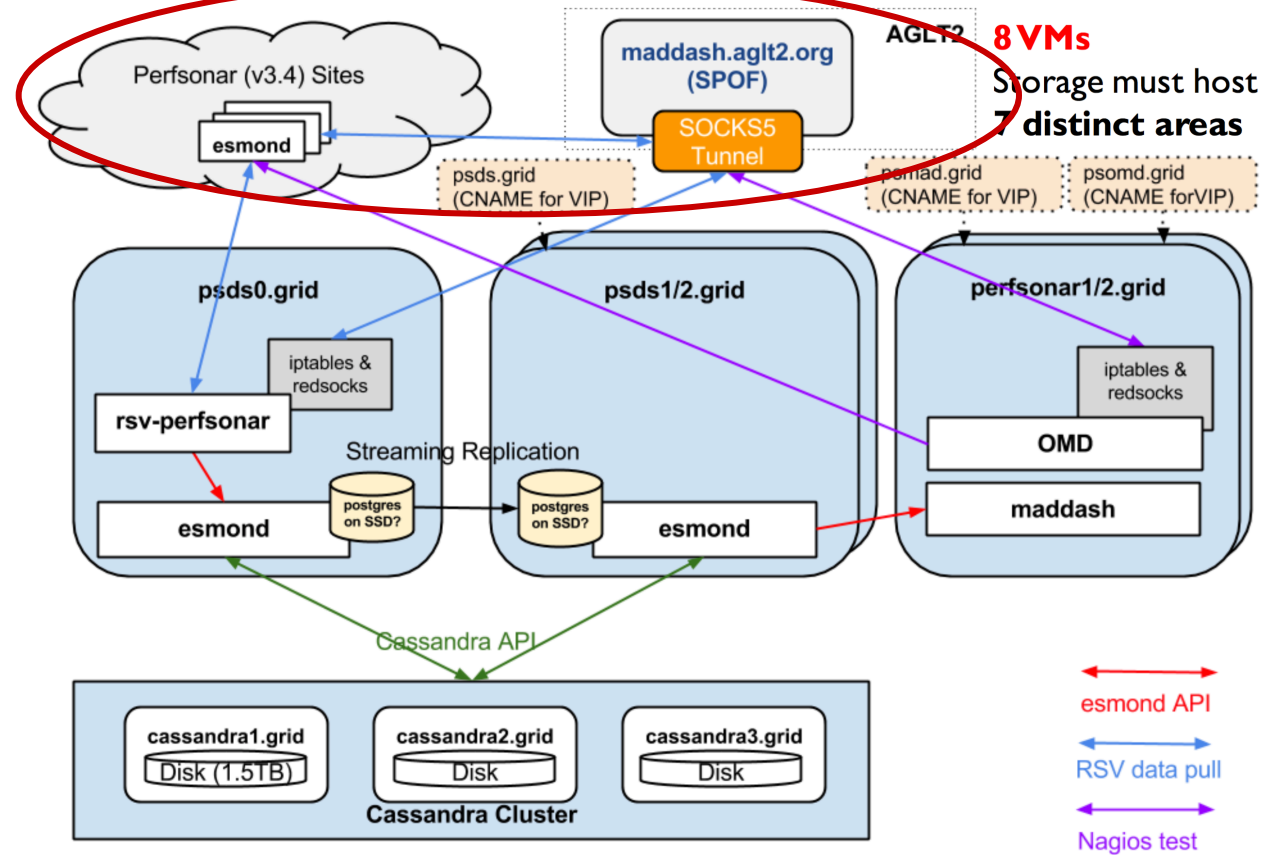
# Operational Infrastructure

- Development Environment and Prototype service at University of Michigan
- Integration Testing Environment hosted by Open Science Grid Operations in the OSG Integration Testbed at Indiana University
- Production Environment hosted by OSG Operations at Indiana University
- Metric publishing to outside collaborators including CERN and researchers at UNL and UCSD

## Logical Diagram of Datastore



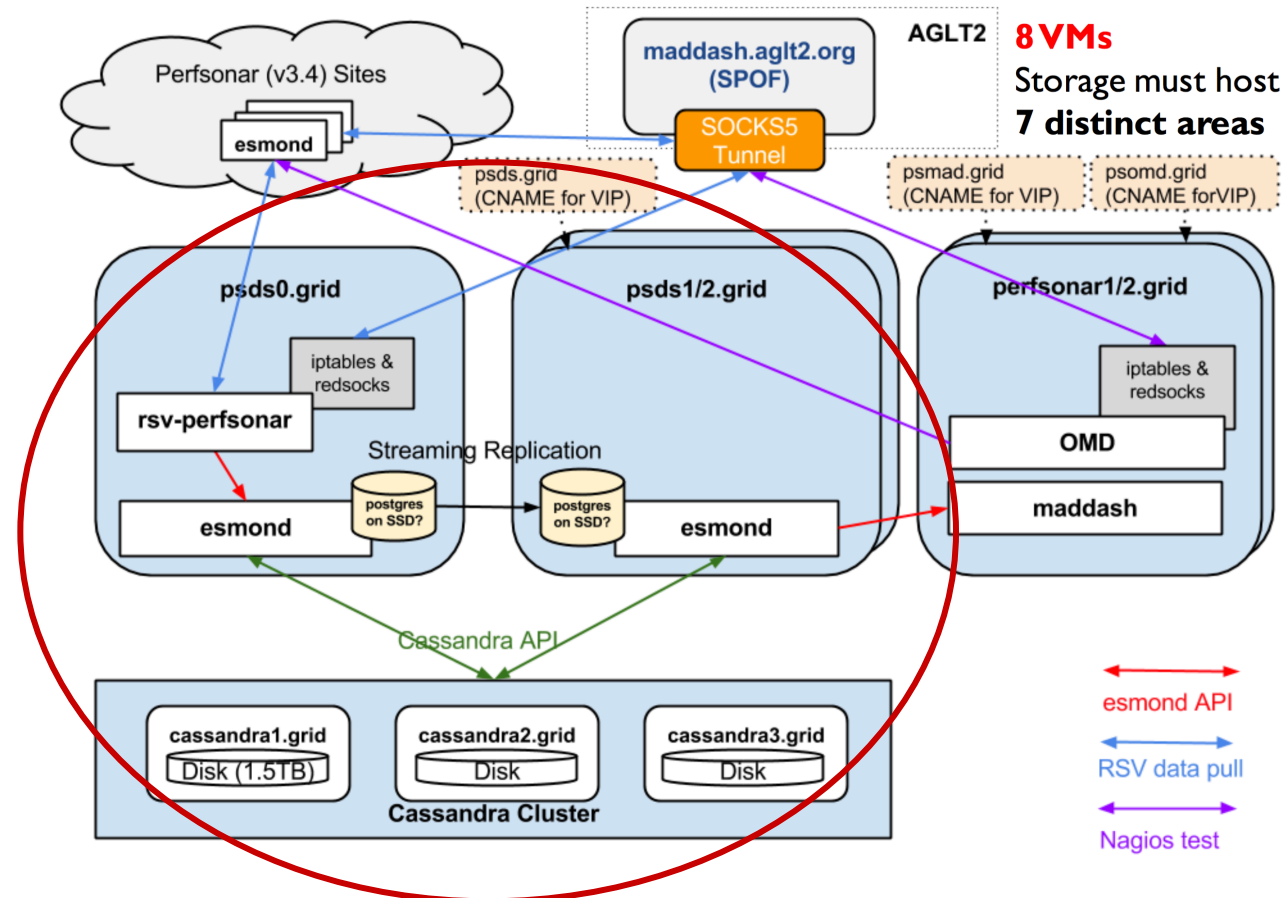
## Logical Diagram of Datastore



# perfSONAR deployment

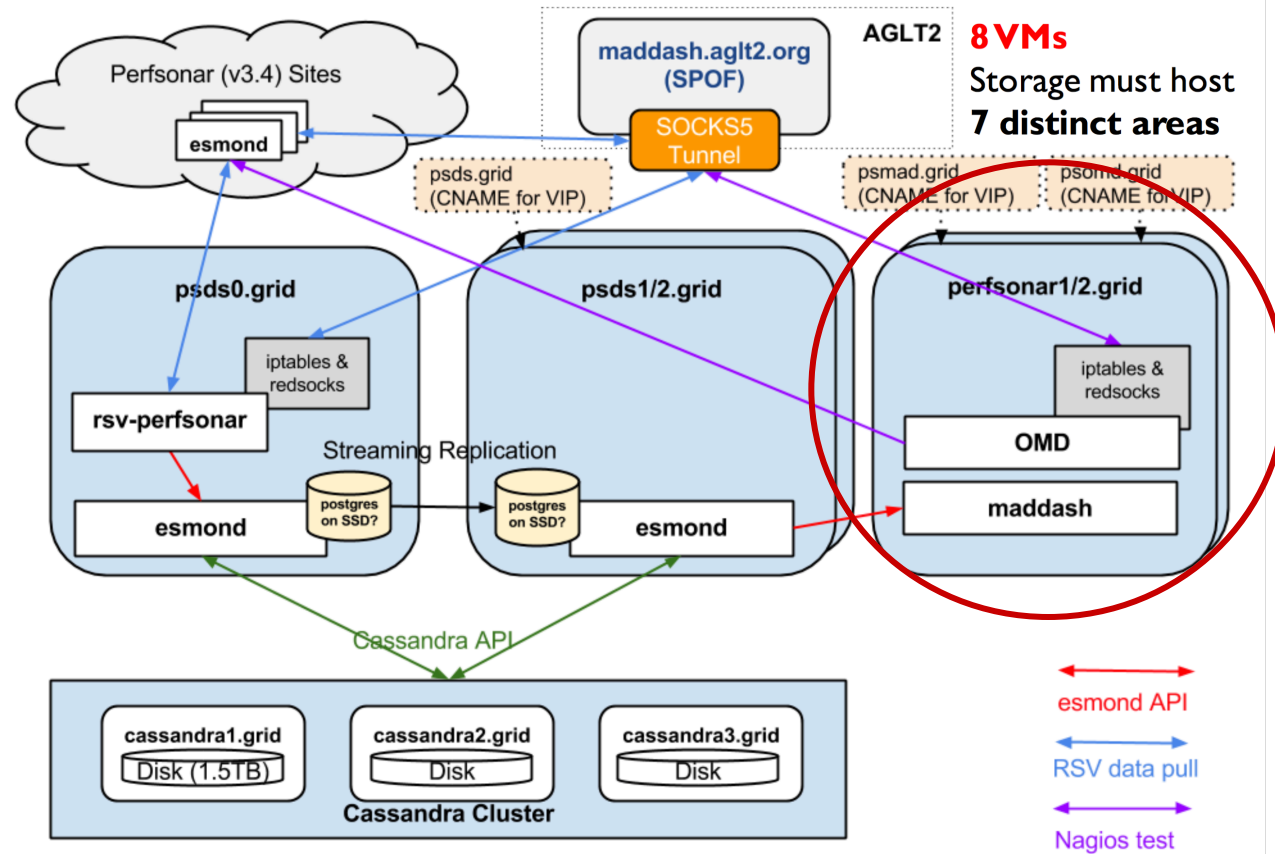
- Began in 2012 with WLCG Tier 1 & 2 resources
- Initial Rollout took several months
- Included several updates to clients to scale to the level needed by OSG
- Metrics published every ~5 minutes, caching built in
- Communicates metric results to the OSG Datastore via RSV probes
- As of 3 October 2016 271 clients registered and monitored for activity
- Due to firewall policy a small number of sites publish initially to AGLT2 before being sent to the OSG Datastore

## Logical Diagram of Datastore

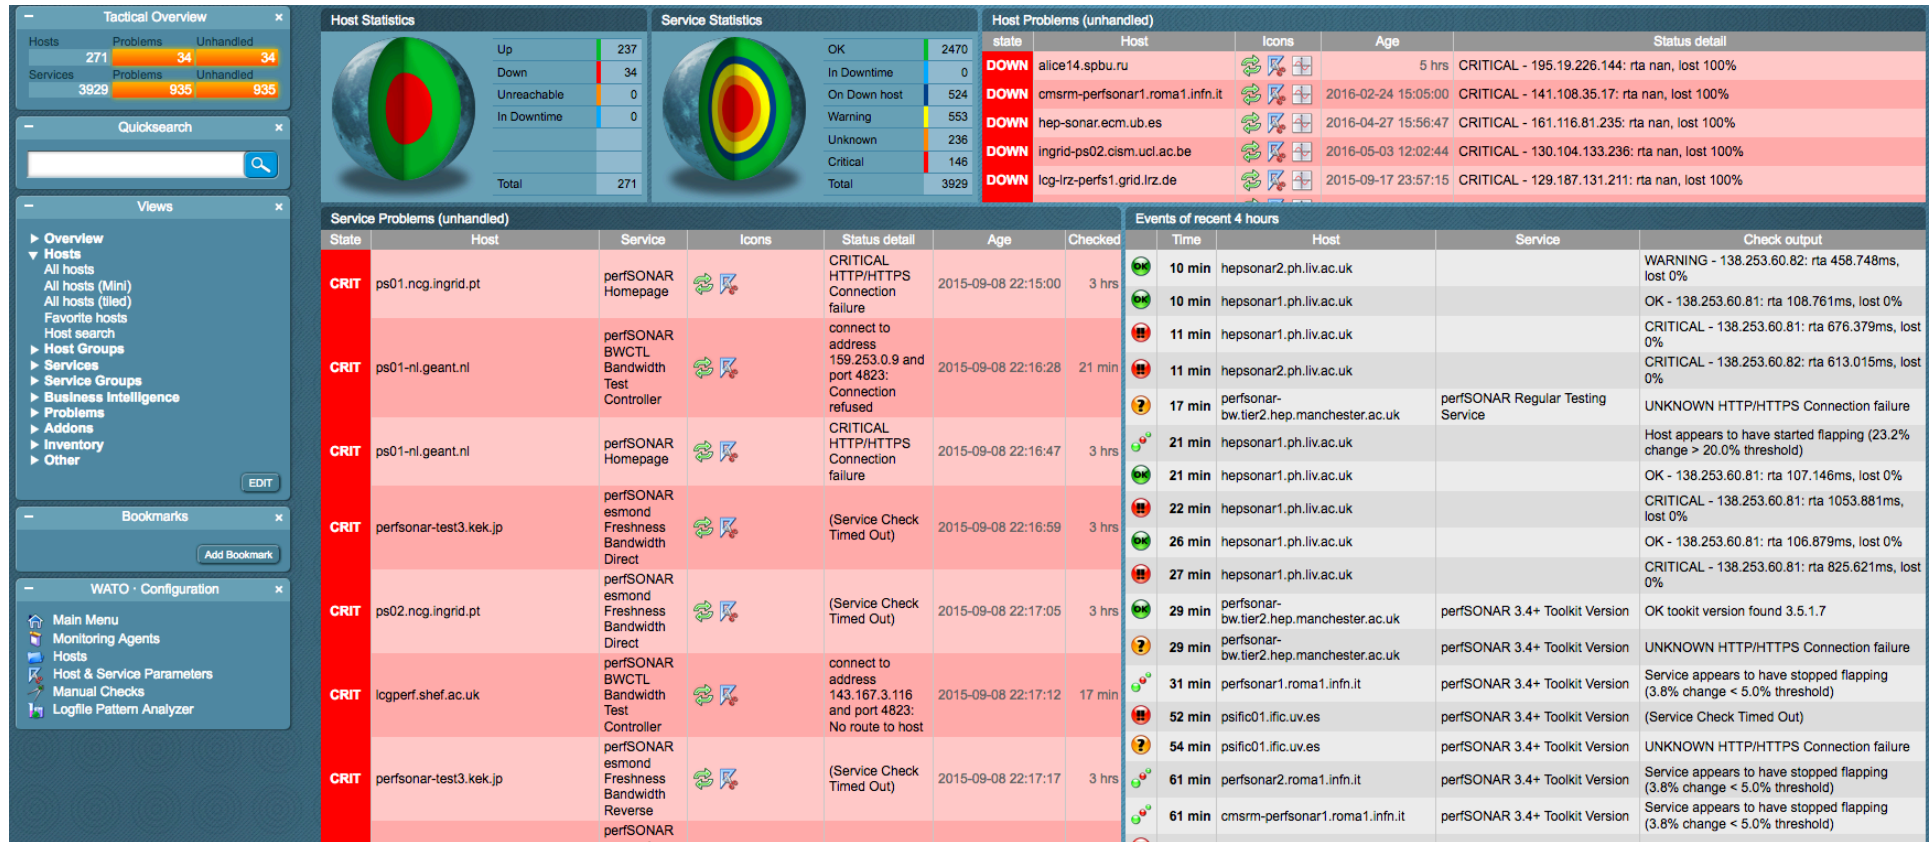


- Existing software used for communication of metrics from site level to datastore
  - RSV
  - Nagios
- RSV also used to publish to CERN Messaging Queue
  - Access to the published metrics is open
- Esmond - ESNM Monitoring Daemon
  - Esmond API used to publish Maddash mesh visualizations
  - Uses PostgreSQL on SSD for high activity
- Cassandra used for long term storage and archiving

## Logical Diagram of Datastore



- Open Monitoring Distribution – Nagios based monitoring tools allowing us to determine site metric publishing status




The screenshot displays the perfSONAR monitoring interface. On the left, there is a sidebar with navigation options like 'Overview', 'Hosts', 'Services', and 'WATO - Configuration'. The main area is divided into several sections:

- Tactical Overview:** Shows a summary of hosts (271) and services (3929) with counts for problems and unhandled issues.
- Host Statistics:** A globe icon and a bar chart showing host status: Up (237), Down (34), Unreachable (0), and In Downtime (0).
- Service Statistics:** Another globe icon and a bar chart showing service status: OK (2470), In Downtime (0), On Down host (524), Warning (553), Unknown (236), and Critical (146).
- Host Problems (unhandled):** A table listing critical issues for various hosts, such as 'alice14.spbu.ru' and 'cmsrm-perfsonar1.roma1.infn.it'.
- Service Problems (unhandled):** A detailed table of service issues, including 'ps01.ncg.ingrid.pt' and 'ps01-nl.geant.nl', with columns for state, host, service, status detail, age, and checked status.
- Events of recent 4 hours:** A log of recent events with columns for time, host, service, and check output.

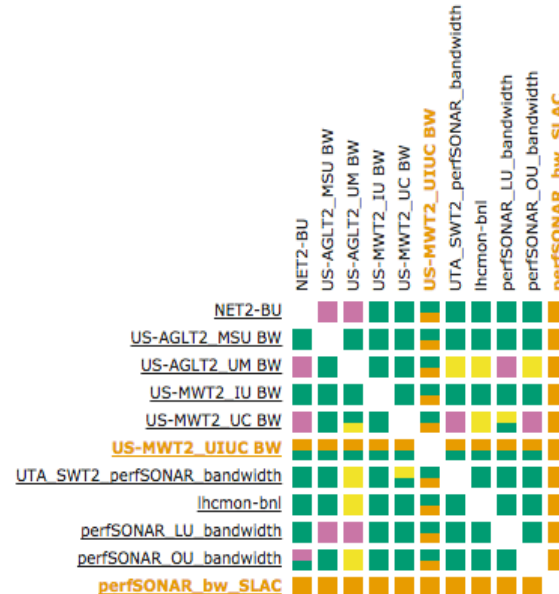
# Mesh Configuration and Visualization of Stats

- Custom Mesh Configuration Tool integrated with OSG Information Management Topology Database
- Mesh Visualization utilizing Maddash

USATLAS Mesh Config - USATLAS Bandwidth Mesh Test



⚠ Found a total of 2 problems involving 2 hosts in the grid



# So how is it working out?

- **January 2016 – Degraded performance from Canada to some European resources.**
  - Tracked to errors on the Washington-London link
  - A backup was provisioned over the NY to Amsterdam path
  - Troubleshooting with perfSONAR helped identify the problem
- **February 2016 – Various ongoing network issues at ASGC in Taiwan**
  - Recommendation of using routers with enough buffers to support long path transfers
  - perfSONAR again used to help identify location of issue and led network engineers to resolve the issue
- **March 2016 – Packet loss between CERN and Brazilian sites**
  - Problem tracked to ESnet – RNP stopped publishing to Esnet resolving the issue
  - perfSONAR tools used to track location of network issue and engineers suggested a workaround

# Some final thoughts

- We are still learning the behavior of this new service
  - A recent outage led to several operational process tweaks
  - Several new technologies adopted
- The value of the perfSONAR clients installations is obvious
  - Give the ability to locate source of problems
  - Provide a historical view of performance
- The datastore itself is available for researchers that want a long term view of international research network performance
- Alerting and alarming procedures will need continued tuning as the service matures
  - Both on the technical and support side
- Long term archiving has not been fully worked out
- What would you like to do with this data?

- Questions? Comments?
  - Feel free to ask us questions now or later:  
Rob Quick [rquick@iu.edu](mailto:rquick@iu.edu)  
Marian Babik [Marian.Babik@cern.ch](mailto:Marian.Babik@cern.ch)  
Shawn McKee [smckee@umich.edu](mailto:smckee@umich.edu)  
Edgar Fajardo Hernandez [emfajard@ucsd.edu](mailto:emfajard@ucsd.edu)  
Soichi Hayashi [hayashis@iu.edu](mailto:hayashis@iu.edu)  
Tom Lee [thomlee@iu.edu](mailto:thomlee@iu.edu)  
Chris Pipes [cpipes@iu.edu](mailto:cpipes@iu.edu)  
Scott Teige [steige@iu.edu](mailto:steige@iu.edu)