OSG Networking Area

Shawn McKee/University of Michigan

OSG Internal Review / FNAL

March 14th 2014

Open Science Grid

OSG Networking Area Mission

- OSG Networking was added at the beginning of OSG's second 5-year period in 2012
- The "Mission" is to have OSG become the network service data **source** for its constituents
  - Information about network performance, bottlenecks and problems should be easily available.
  - Should support our VOs, users and site-admins to find network problems and bottlenecks.
  - Provide network metrics to higher level services so they can make informed decisions about their use of the network *(Which sources, destinations for jobs or data are most effective?)*

Ingredients for the Mission

Open Science Grid

- The primary motivator for adding networking to OSG was the very successful deployment of the perfSONAR-PS Toolkit in, first, the USATLAS VO and then in USCMS
  - perfSONAR, developed by a global collaboration of Research and Education R&E) networks, provides a standardized set of network tools, data storage schema and methods to create an easy-to-install network measurement point.
  - The perfSONAR-PS Toolkit is the implementation created by ESnet, Internet2 and others and provides scheduled and on-demand tests between the toolkit and other relevant instances.
- USATLAS had also developed a "modular dashboard" which could gather, test and display perfSONAR network metrics to allow quick visualization of the state of the network.
  - perfSONAR+dashboard was(is) very effective in identifying hidden, longstanding network issues in USATLAS infrastructure.
- Initial versions of network related documentation (How-tos, install guides, FAQs, etc) had been created as part of HEPs efforts to-date. These were used to kick-start the OSG documentation.

OSG Networking Area Effort

- The smallest area in OSG. Currently 25% of me plus 15% of Mike Blodgett/Wisc.
  - **Also draws upon other OSG areas as appropriate** (Operations, Technology and Software)
  - However this area is leveraging effort in Internet2/ESnet (perfSONAR-PS development) and HEP/WLCG (perfSONAR-PS global deployment and efforts in ATLAS and CMS)
- Nice to leverage external effort BUT that makes us very dependent upon effort we don't control…

Year 1 Goals and Key Initiatives in Network Area

Open Science Grid

- Year 1 of OSG Networking was primarily focused on getting network monitoring in place
  - **Deploying perfSONAR-PS:** Instrumenting OSG sites with standardized tools to gather network metrics
  - **OSG Network Service:** Gathering OSG network metrics centrally and making them available for users and applications
  - **Network Documentation:** Creating documentation for OSG user and VO managers to guide them in understanding and diagnosing network issues

Achievement Summary Year 1

Open Science Grid

- Significant improvements in perfSONAR-PS Toolkit:
  - Much more robust and resilient
  - Default configuration well tuned to our deployment environment
  - New feature: mesh-configuration allows centralized and federated management of tests and meshes
- OSG Network Service deployed in MyOSG/OIM
  - New OIM registration schema setup for perfSONAR-PS instances
- Modular Dashboard project created and beta release out
  - License issues cleared (BSD) and source in GitHub
  - Much of the original functionality recreated in a scalable modular way
- New documentation for users and admins in place
  - Guidance for debugging network issues, using OSG network tools and deploying OSG network infrastructure (toolkit)

Year 2 OSG Networking Plans

- Two primary components:
  - Finishing Year 1 items
  - New efforts
- Finishing Year 1
  - Complete deployment of perfSONAR-PS once v3.3 is out
    - Had 18 sites…targeting ~100
  - Improving the modular dashboard
    - Complete replication of ALL original features
    - Optimize resource use
    - Extend “client” API to best support evolving use-cases
- Other areas for year 2
  - Optimize network metrics and collection
  - Enable alarming and problem analysis based upon network metrics
  - Improve tools and documentation from user perspective

OSG Networking Service

- OSG is building a centralized service for gathering, viewing and providing network information to users and applications.
- Basis of this service was the follow-on to the original BNL dashboard project: https://github.com/PerfModDash
  - Recoding to create a modularized, scalable system which maintains all functionality from the original service
  - **Thanks to Soichi/Rob/Operations, OSG has this service up: announcement date was June 11, 2013**
    - Matrix View http://tinyurl.com/bpzbjrw
    - Host Status http://tinyurl.com/adbhv6n
- **Goal**: OSG becomes the "source" for networking information for its constituents, aiding in finding/fixing problems and enabling applications and users to better take advantage of their networks

OSG Network Service in Year 2

- For year 2 we continue implementing and improving the OSG network service with our goal of OSG becoming the "source" of network information for OSG sites in mind
  - The OSG service should provide an integrated GUI component that quickly visualizes test metrics (matrices by metric type)
  - The datastore should have a client API that meets the needs of OSG VO users AND applications
  - We need to ensure the service maintains scalability for OSG use-cases (See discussion about WLCG at end of talk)
- The Operations team is critical in this effort
- **In December 2013, the primary dashboard developer left BNL.** That dashboard code is now curated by OSG operations but we need a supported replacement

Example of Original Dashboard

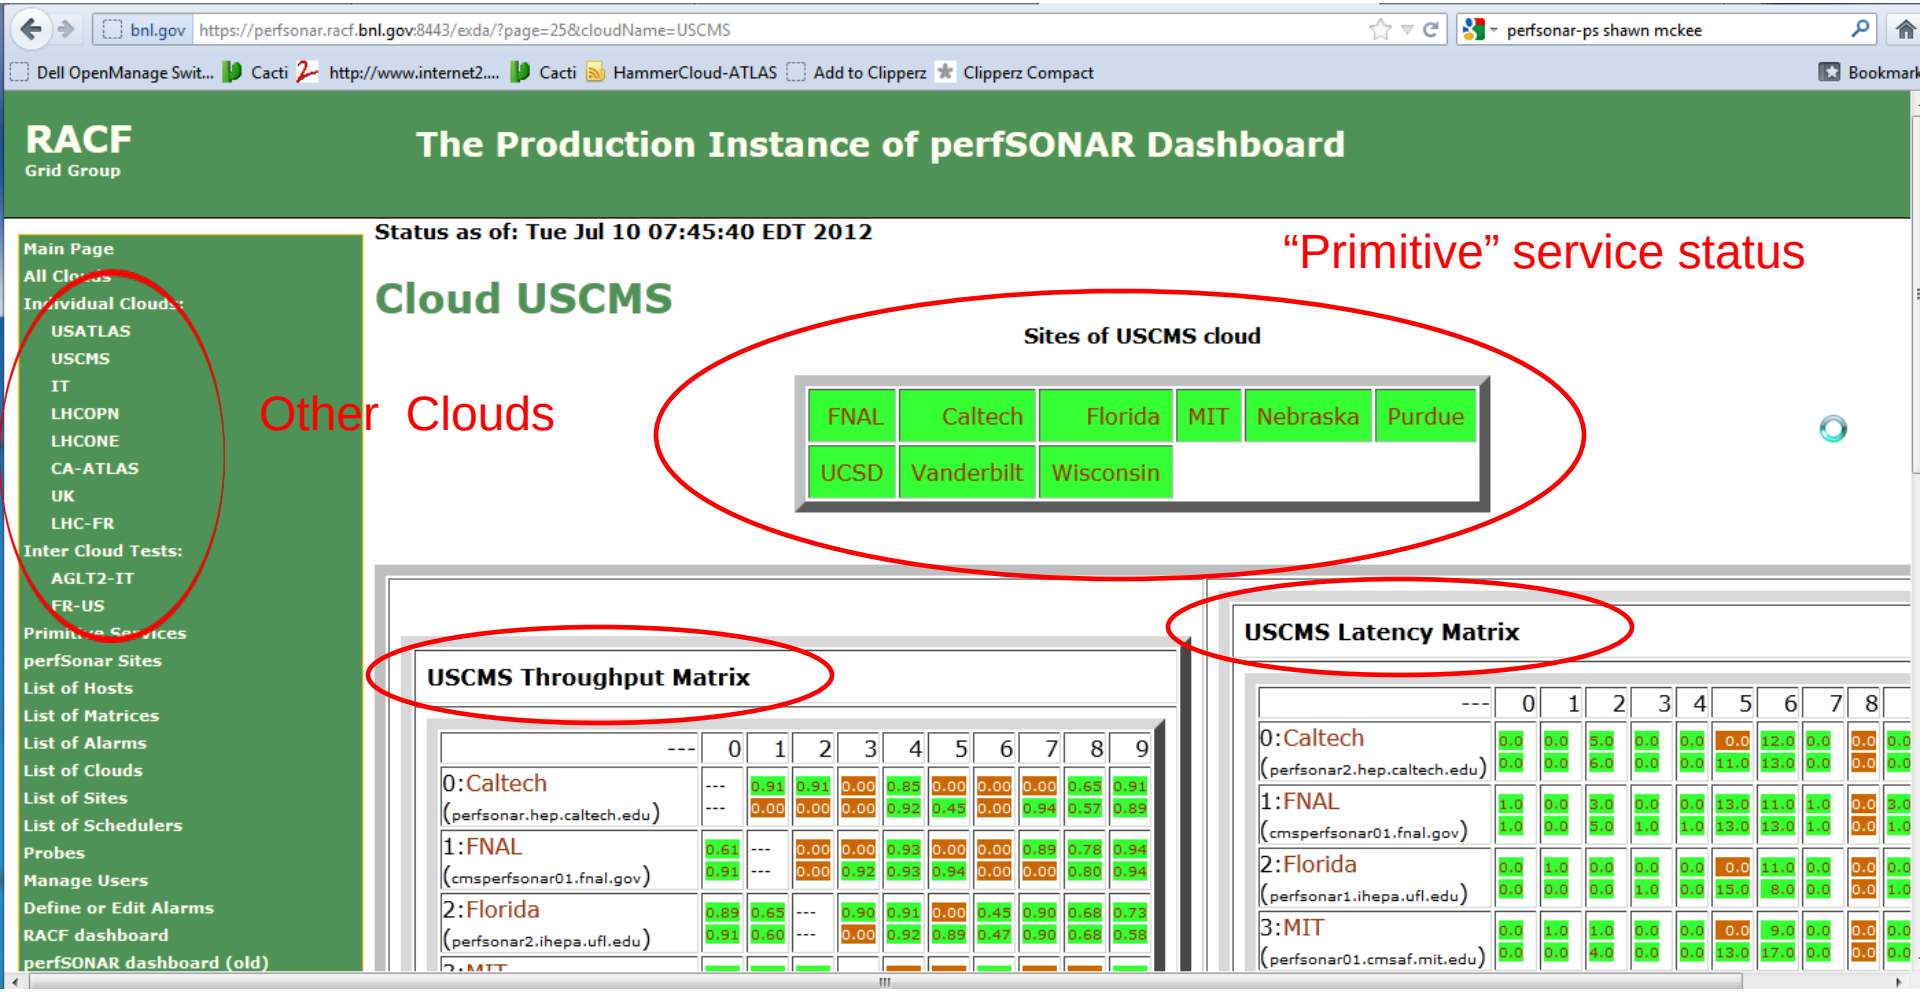

Orphaned since December 2013. Prototyping replacement (MaDDash/OMD)

Replacement Prototype: MaDDash

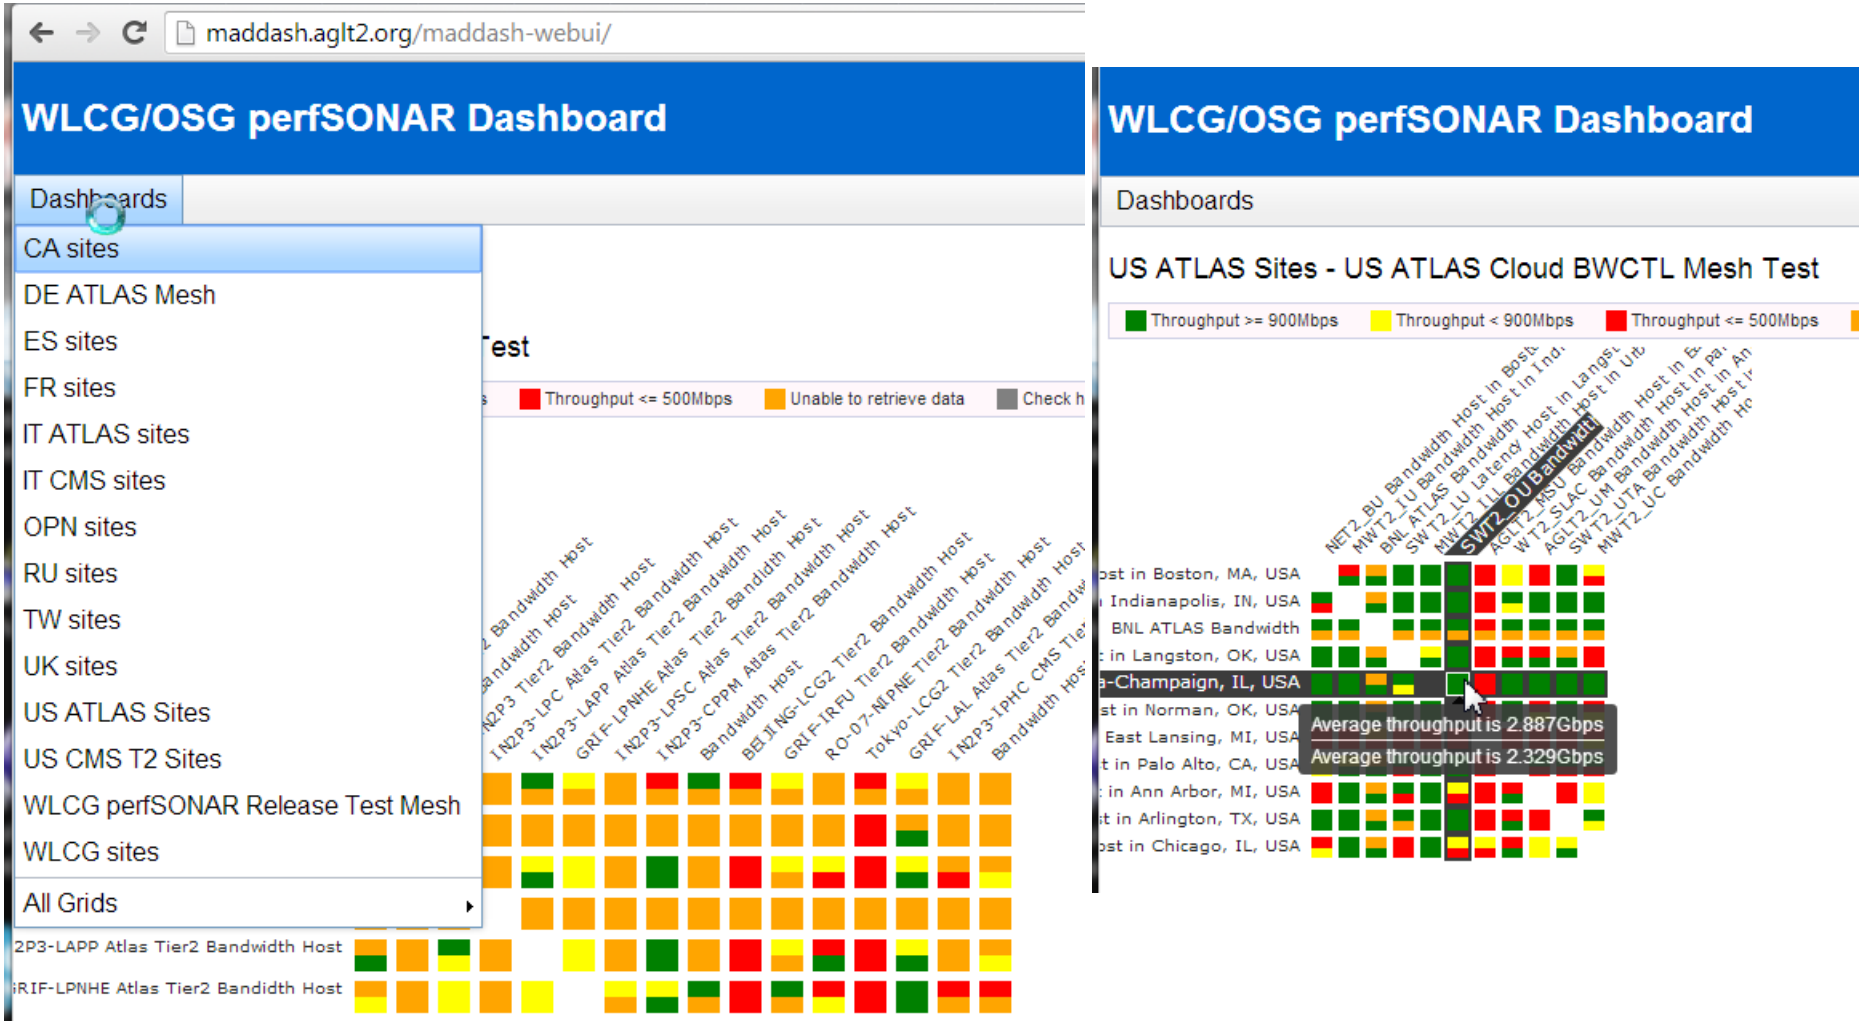

MaDDash (Monitoring and Debugging Dashboard) supported by ESnet

Prototype: Service Monitoring

Open Science Grid

OMD (Open Monitoring Distribution)

Integrated package over Nagios

Checks/verifies primitive services are functional

Ensures we get good network metrics

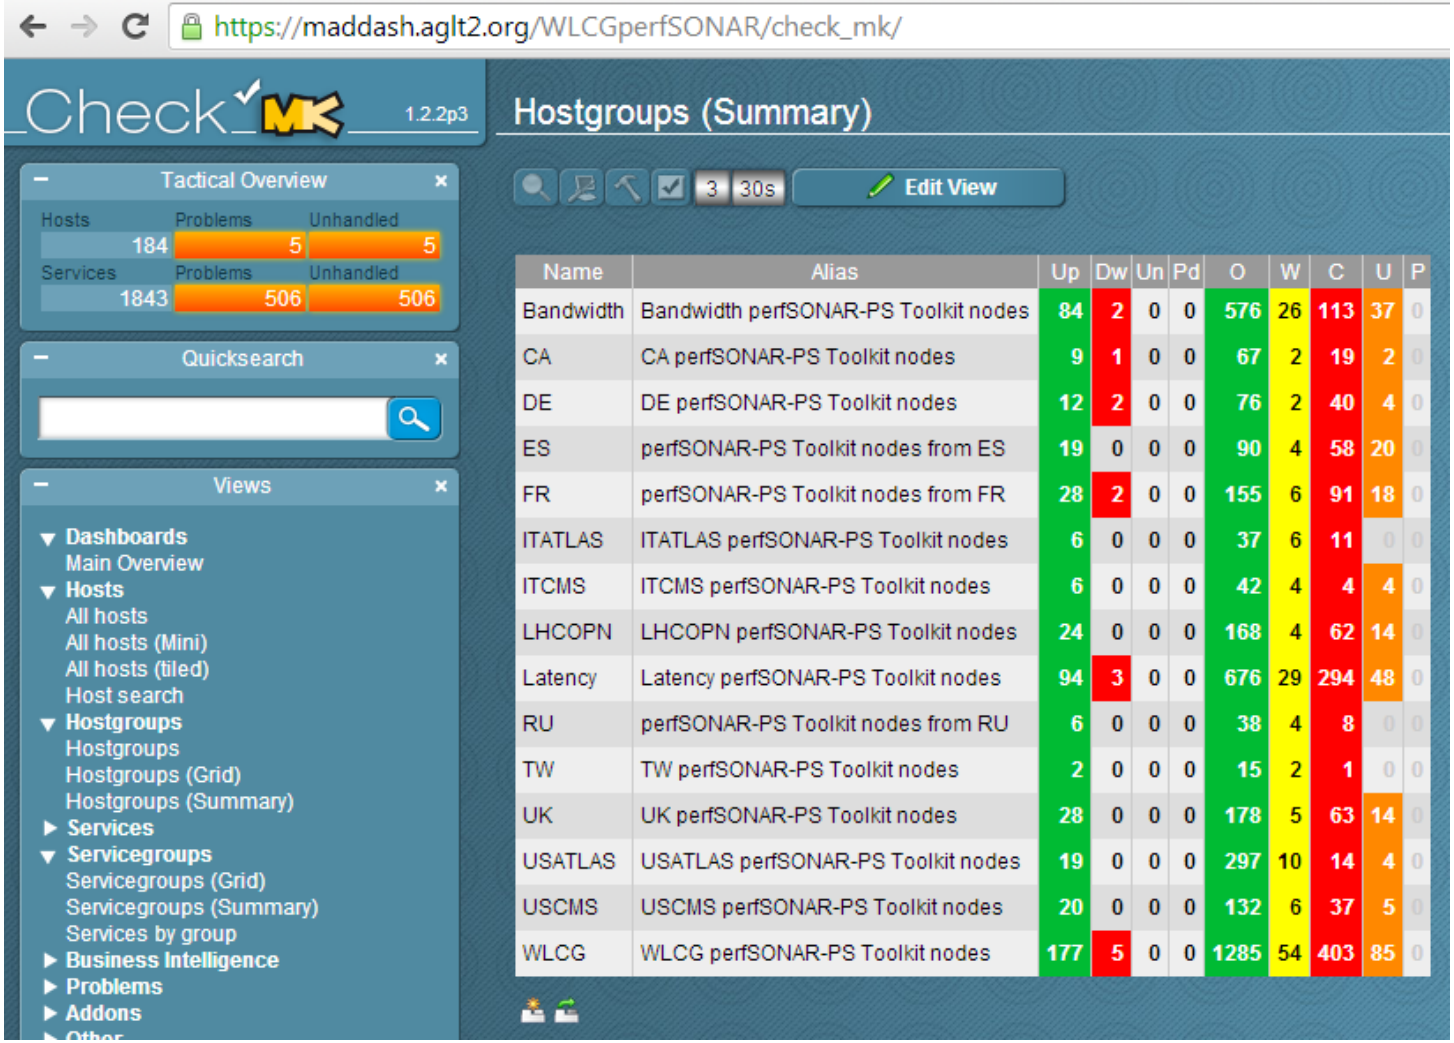

Name	Alias	Up	Dw	Un	Pd	O	W	C	U	P
Bandwidth	Bandwidth perfSONAR-PS Toolkit nodes	84	2	0	0	576	26	113	37	0
CA	CA perfSONAR-PS Toolkit nodes	9	1	0	0	67	2	19	2	0
DE	DE perfSONAR-PS Toolkit nodes	12	2	0	0	76	2	40	4	0
ES	perfSONAR-PS Toolkit nodes from ES	19	0	0	0	90	4	58	20	0
FR	perfSONAR-PS Toolkit nodes from FR	28	2	0	0	155	6	91	18	0
ITATLAS	ITATLAS perfSONAR-PS Toolkit nodes	6	0	0	0	37	6	11	0	0
ITCMS	ITCMS perfSONAR-PS Toolkit nodes	6	0	0	0	42	4	4	4	0
LHCOPN	LHCOPN perfSONAR-PS Toolkit nodes	24	0	0	0	168	4	62	14	0
Latency	Latency perfSONAR-PS Toolkit nodes	94	3	0	0	676	29	294	48	0
RU	perfSONAR-PS Toolkit nodes from RU	6	0	0	0	38	4	8	0	0
TW	TW perfSONAR-PS Toolkit nodes	2	0	0	0	15	2	1	0	0
UK	UK perfSONAR-PS Toolkit nodes	28	0	0	0	178	5	63	14	0
USATLAS	USATLAS perfSONAR-PS Toolkit nodes	19	0	0	0	297	10	14	4	0
USCMS	USCMS perfSONAR-PS Toolkit nodes	20	0	0	0	132	6	37	5	0
WLCG	WLCG perfSONAR-PS Toolkit nodes	177	5	0	0	1285	54	403	85	0

Alerting/Alarming for Network Issues

- **What most sites want is a tool that lets them know if there is a network problem (and ideally WHERE it is)**
- In year 2 we started to develop this capability for OSG sites
  - Primitive OSG perfSONAR-PS service monitoring is easy and we have Nagios-type plugins that check services
  - Much harder is deciding when network metrics gathered by perfSONAR-PS require an alert or alarm:
    - Is the change in metrics due to "normal" (heavy) network use or is there a new problem?
    - If there is a real problem, where is it located? This is critical because we should only alert someone if the problem is one they can fix
- Interesting project at Georgia Tech called Pythia (see Terena presentation https://tnc2013.terena.org/core/presentation/40)
  - Submitting (this week) to NSF SI2-SSE "PuNDIT" (Pythia Network Diagnosis Infrastructure) which targets OSG/WLCG
  - Goal is to provide this needed alerting/alarming component

OSG Network Documentation

- Understanding, finding and fixing network problems or optimizing network usage can be complex. Documentation is **critical** to supporting OSG users and administrators
- Set of OSG network pages in place at https://www.opensciencegrid.org/bin/view/Documentation/NetworkingInOSG
- Network "triage" document for OSG users at: https://twiki.opensciencegrid.org/twiki/pub/Documentation/NetworkingTroubleShooting/20130204-OSG-Debug-SPM.docx
  - A "living" document, evolving as tools update.
- OSG network client tools updated / tested for client RPMs:

https://www.opensciencegrid.org/bin/view/Documentation/Release3/NetworkPerformanceToolkit

- Installation guide for perfSONAR-PS Toolkit in place at https://www.opensciencegrid.org/bin/view/Documentation/PerfSONARToolKit and https://www.opensciencegrid.org/bin/view/Documentation/RegisterPSinOIM
  - Continual tuning and updating as we gather experience and feedback

Network Info and Documentation

- As we make progress in improving perfSONAR-PS, the dashboard and OSG network services we need to maintain and expand our documentation.
- For year 2 we additionally have Mike Blodgett/UWisc working on “Network Monitoring Capabilities Design”
  - Test and evaluate the network monitoring capabilities OSG provides, especially from the perspective of users and network technical support perspectives.
  - Provide feedback for the developers of the perfSONAR-PS toolkit, the PerfModDash GitHub project and the OSG network services
  - Suggest new features as appropriate to better serve users and technical network support personnel.
  - Add/update documentation, targeted at users and technical network support personnel, on utilizing the OSG network monitoring to identify, localize and expedite the repair of network problems.“
- Document examples of successful use of OSG network tools

WLCG Request Discussion

- WLCG is deploying perfSONAR-PS Toolkits at ALL WLCG Tier-1/Tier-2 sites worldwide; deadline April 1, 2014
  - They need a centralized service to collect, organize and make available perfSONAR-PS network metrics
  - Originally they used the **BNL dashboard** (LHCOPN, LHCONE, various Tier-1 clouds (CA,IT,UK, FR) but now is orphaned
- **WLCG has requested that OSG (as a WLCG member) provide a network service including a perfSONAR-PS dashboard for WLCG**
- OSG is evaluating the impact of supporting the WLCG sites. Assuming the load (and hardware/people) is not too large, the plan is to support both OSG and WLCG

Lessons Learned

- Network issues can be “hard” to find/understand/fix, but most OSG users don’t care about networking until/unless they hit “network” problems or limitations.
- Relying upon “outside” effort (and timelines) can be ‘expensive’ for our scheduling and planning ☹
- Users drive improvements in community projects. We benefited from significant improvements in perfSONAR-PS from broad-scale testing and use.
- Users don’t always contribute to community projects. Need to find ways to engage more interest/effort in ours.
- Hardening/improving broadly deployed software is time-consuming

Challenges

- Providing a network service for both OSG and WLCG (scaling, user/application API and metric visualization)
  - Integrating the prototype MaDDash and OMD components into OSG Operations and the current setup (OIM/MyOSG)
  - Delivering an effective alerting/alarming component for network problem identification and localization
  - Creating a suitable datastore/API for clients of the metrics
- Extending perfSONAR-PS Toolkit deployments to all of OSG. (Most instances associated with WLCG/OSG sites)
- Automating creation and management of centralized perfSONAR-PS configurations from OIM/GOCDB

Discussion/Questions

Open Science Grid

Questions or Discussion?

Some References

- OSG Networking
  https://www.opensciencegrid.org/bin/view/Documentation/NetworkingInOSG
- perfSONAR-PS site http://psps.perfsonar.net/
- OSG perfSONAR-PS Toolkit install guide
  https://www.opensciencegrid.org/bin/view/Documentation/PerfSONARToolKit
- I2 Install/configuration guide: http://code.google.com/p/perfsonar-ps/wiki/pSPerformanceToolkit33
- Modular Dashboard: http://maddash.aglt2.org/maddash-webui
- OMD monitoring for perfSONAR:
  https://maddash.aglt2.org/WLCGperfSONAR/omd
- Tools, tips and maintenance:
  http://www.usatlas.bnl.gov/twiki/bin/view/Projects/LHCperfSONAR
- LHCOPN perfSONAR Deployment Details:
  https://twiki.cern.ch/twiki/bin/view/LCG/PerfsonarDeployment