

# **OSG Networking Area Planning and Discussion**

**Shawn McKee/University of Michigan**

**OSG Retreat / UCSD**

**May 11<sup>th</sup> 2016**



**Open Science Grid**

❄ 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?*)

# Review: OSG Networking Motivations



- ❄ Networks underlie our distributed computing model but are historically only **indirectly** visible. This led many to feel most problems with a WAN involved were network problems (and sometimes that was true).
- ❄ perfSONAR is part of an evolving infrastructure where the network plays an increasingly visible role. With perfSONAR we can monitor our networks, understand capacity, find bottlenecks and detect problems. It is NOT the only thing we need.
- ❄ With Software Defined Networking slowly creeping into our network hardware we will have more opportunity in the future to integrate the network we need into our end-to-end systems.
- ❄ **Our long-term goal in OSG is to make the network visible and controllable to improve our infrastructure, avoid congestion, work around failures and improve efficiency.**

# OSG Networking Area Effort



- ❄ The smallest area in OSG. Currently 40% of me plus:
  - ❑ Also draws upon other OSG areas as appropriate (Operations, Technology and Software)
    - ⌘ **Significant help** for Datastore development, bug fixes and optimization from **Soichi, Edgar and Brian**. **Rob, Thomas, Gabriele and Chander** also contributing to oversight and operations.
  - ❑ This area is leveraging effort in **Internet2/ESnet** (perfSONAR development) and **HEP/WLCG** (perfSONAR global deployment and efforts in ATLAS and CMS)
  - ❑ Coordinating with **PuNDIT** satellite project: targeting problem identification/localization
  
- ❄ **This last year has been very successful.**
  - ❑ We have raised awareness about the network, its status and performance. Our colleagues know about perfSONAR and trust its data
  - ❑ We used our OSG network services, infrastructure and metrics to quickly find and fix network issues for our science domains (ATLAS, CMS, Belle II)

# OSG Year 4 Goals Summary



## ❄ Three primary goals all completed

### ❑ Network Datastore into production

- ⌘ New hardware installed by August 1, 2015
- ⌘ Production version of data pipeline operational November 1, 2015

### ❑ USATLAS and USCMS sites using OSG data to find and fix problems

- ⌘ MaDDash meshes for USATLAS and USCMS ready July 15 2015
- ⌘ Site and network issues identified and acted on starting September 2015

### ❑ Data available via API

- ⌘ Initial version ready August 30 2015
- ⌘ v1.0 ready October 1, 2015 using ESmond schema, broadcast to ActiveMQ/CERN

❄ We have **instrumented** our networks, users and VOs are benefiting and data is available for use.

## ❄ Stretch goals

### ❑ Network cost matrix for LHC experiments

- ⌘ Initial prototype developed in January 2016 working with ANSE and WLCG. Needs hardening.

### ❑ Alerting and Alarming on network issues.

- ⌘ Have testing using Check\_mk rule based notification starting. MadAlert and PuNDIT

# Network Performance Issues

We have been using our OSG infrastructure and data to identify and fix numerous network issues.

The table on the right shows the recent (since Feb 2016) WLCG support unit issues. Using OSG perfSONAR data we even know the root cause of the current open issues.

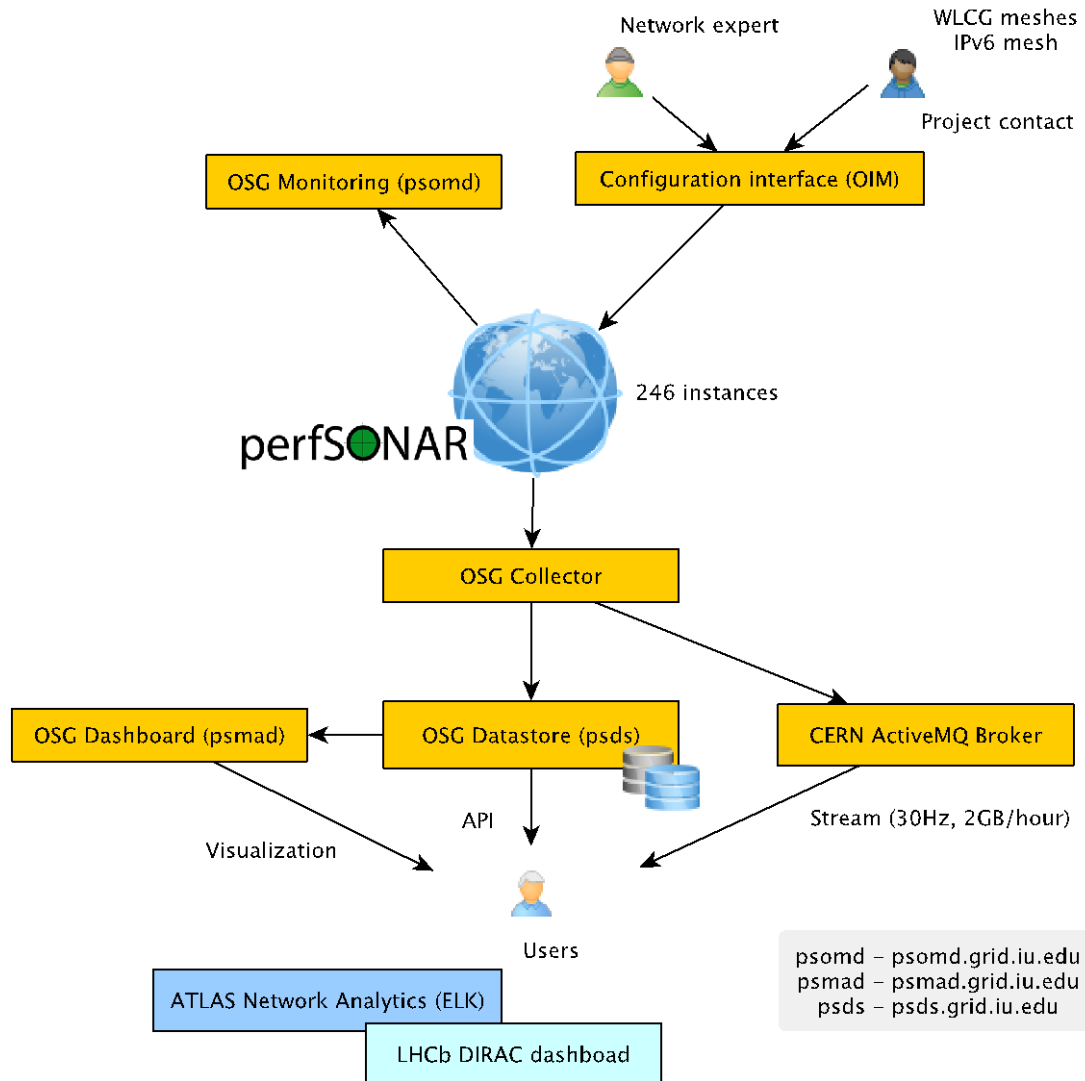
Within OSG we have additionally found and fixed issues involving Caltech, Boston University, BNL, FNAL, OU, Michigan, Michigan State, Langston, Indiana, University of Chicago, University of Illinois, Nebraska and Georgia Tech

The bad news is that most of these cases were not initiated proactively but rather in response to data transfer issues

Incident	Ticket	Comments
BNL RAL CERN		Still open
BNL SARA CERN	<a href="https://ggus.eu/index.php?mode=ticket_info&amp;ticket_id=120957">https://ggus.eu/index.php?mode=ticket_info&amp;ticket_id=120957</a>	Resolved, issue with ESNet router at CERN
ASGC CERN IJS	<a href="https://ggus.eu/index.php?mode=ticket_info&amp;ticket_id=119820">https://ggus.eu/index.php?mode=ticket_info&amp;ticket_id=119820</a>	Still open
CBPF	<a href="https://ggus.eu/index.php?mode=ticket_info&amp;ticket_id=120081">https://ggus.eu/index.php?mode=ticket_info&amp;ticket_id=120081</a>	Resolved by RNP stopped publishing to ESNet CBPF IPs
FNAL CERN	<a href="https://ggus.eu/ws/ticket_info.php?ticket=119551">https://ggus.eu/ws/ticket_info.php?ticket=119551</a>	Didn't reach WLCG throughput SU, was assigned directly to CERN LHCOPN, was fixed by Esnet - faulty router interface in New York
PIC inbound	via mailing list <a href="https://twiki.cern.ch/twiki/bin/view/LCG/NetworkTransferMetrics#Network_Throughput_Support_Unit">https://twiki.cern.ch/twiki/bin/view/LCG/NetworkTransferMetrics#Network_Throughput_Support_Unit</a>	10 Gbps link WAN at PIC sharing LHCOPN, LHCONE was completely saturated causing input discards
BNL to PIC	via mailing list	LHCOPN link CERN-PIC was flapping a lot due to an issue with the Geant fibre to Spain
MAINZ CA	via mailing list	MAINZ uses a "commercial" network provider and Canadian sites only peer with R&E networks
OU inbound	via mailing list	Narrowed down to a faulty switch on site
CA EU	<a href="https://ggus.eu/index.php?mode=ticket_info&amp;ticket_id=118748">https://ggus.eu/index.php?mode=ticket_info&amp;ticket_id=118748</a> , <a href="https://ggus.eu/index.php?mode=ticket_info&amp;ticket_id=118730">https://ggus.eu/index.php?mode=ticket_info&amp;ticket_id=118730</a>	Trans-atlantic channel instability, resolved by Canarie by re-routing

# Overview of perfSONAR Pipeline

The diagram on the right provides a high-level view of how OSG is managing perfSONAR deployments, gathering metrics and making them available for use.



- ❄ Ilija Vukotic/U Chicago has been leading an effort to get network metrics into an analytics platform
- ❄ This analytics service indexes historical network related data while providing predictive capabilities for near term network throughput performance.
- ❄ Primary functions:
  - ❑ Aggregate, and index, network related data associated with WLCG “links”
  - ❑ Serve derived network analytics to ATLAS production, DDM & analysis clients
  - ❑ Provide a generalized network analytics platform for other communities in the OSG
- ❄ Part of ATLAS Analytics platform
  - ❑ <https://cds.cern.ch/record/2056257/files/ATL-SOFT-SLIDE-2015-752.pdf>



- ❄ PuNDIT (OSG Satellite project) has ~8 months left. Goal is to provide near realtime alerting on some network problems
- ❄ At Michigan, Gabriele Carcassi and I developed MadAlert: a web-based tool to analyze OSG bandwidth and latency meshes and identify obvious problems.
  - ❑ <http://madalert.aglt2.org/madalert/index.html>
  - ❑ Will become part of ESnet's next MaDDash release
- ❄ Soichi has produced an almost final version of a standalone mesh-configuration utility
  - ❑ Allows campuses or VOs to setup and control their own perfSONAR meshes
  - ❑ Needs more feedback and testing
  - ❑ Some features still to add and enable
  - ❑ Will require packaging

# Exploring SDN in Production

- ❄ A group of people in the US from AGLT2, MWT2, SWT2 and NET2 are planning to explore SDN in ATLAS
  - ❑ Working with the LHCONE point-to-point effort as well
- ❄ The plan is to deploy Open vSwitch on ATLAS production systems at these sites (<http://openvswitch.org/>)
  - ❑ IP addresses will be move to virtual interfaces
  - ❑ No other changes; verify no performance impact
  - ❑ Traffic can be shaped accurately with little CPU cost
- ❄ The **advantage** is the our data sources/sinks become **visible** and **controllable** by OpenFlow controllers like OpenDaylight
- ❄ Follow tests can be initiated to provide experience with controlling networks in the context of ATLAS operations.

# WLCG Network and Transfer Metrics Working Group



<https://twiki.cern.ch/twiki/bin/view/LCG/NetworkTransferMetrics>

- ❄ The WLCG efforts at CERN are being reorganized; this is an opportunity to chart future directions for the working group
- ❄ There was a review of the working group April 28<sup>th</sup> during the WLCG Operations Coordination meeting.
  - <https://indico.cern.ch/event/514078/>
- ❄ We have a number of areas (projects; see next slide) we are considering and we need to understand where these efforts should be housed (Stay in WG, move to GDB, to LHCONE)
- ❄ OSG will continue to rely upon the WG to help organize global perfSONAR maintenance and support

# WLCG Network Projects

[Details at http://tinyurl.com/WLCG-R-D](http://tinyurl.com/WLCG-R-D)

- ❄ **Title:** LHCONE Traffic engineering
- ❄ **Areas:** LHCONE, routing, debugging, network orchestration
- ❄ **Title:** LHCONE L3VPN Looking Glass
- ❄ **Areas:** LHCONE, monitoring, debugging
- ❄ **Title:** Integration of network and transfer metrics to optimize experiments workflows
- ❄ **Areas:** FAX/Phedex, Rucio, perfSONAR, DIRAC
- ❄ **Title:** Advanced notifications/alerting for network incidents
- ❄ **Areas:** WAN, Advanced Notifications/Alerting, perfSONAR, Hadoop/Spark
- ❄ **Title:** Network performance of the commercial clouds
- ❄ **Areas:** Clouds, WAN connectivity, WAN performance (perfSONAR), establishing and testing network equipment at the cloud provider (VPN)
- ❄ **Title:** Software Defined Network Production Testbed
- ❄ **Areas:** WAN, SDN, LHCONE/LHCOPN, Storage/Data nodes

- ❄ Continue to do what we do now, and:
- ❄ Develop effective **Alarming and Alerting**
- ❄ Support higher-level network services
  - ❑ We have proto-typed a proximity service to find nearest SE given perfSONAR or to find the nearest perfSONAR give and SE
  - ❑ Create network cost prediction service to predict quality and capacity of source-destination paths for network decision support
- ❄ Improve the ability to manage and use network topology and network metrics: Analytics Platform?
- ❄ Prepare-for and integrate **Software Defined Networking**

# Continuing to Do What We Do...

## ❄ Basic things still need to happen in all years:

- ❑ Supporting users, VOs for net problem identification and repair
- ❑ Upgrades and bug-fixes to tools that gather, display and provide network metrics
- ❑ Tuning and optimizing existing testing
- ❑ Maintenance and creation of documentation
- ❑ Support for new ideas and feature requests.
- ❑ Exploring needs for new metrics to better meet researcher needs.

## ❄ **But what interesting possibilities should we focus on given OSG's unique position regarding our hosting of network metrics from all of OSG and WLCG?**

- ❄ Being able to "alarm" on real network problems is a good target: indicate (via monitoring) there is a network problem
- ❄ The next step is to actually "alert" on network problems.
  - The difference between an alarm and an alert is the target. An alarm can appear in some monitoring system for an operator to respond to while an alert is targeted at a person or list of persons (email, page, etc.).
  - To ***effectively*** alert requires that we first have a valid 'network' alarm AND that we be able to localize the problem more specifically than "along the end-to-end path". **Alerts should be only sent to those able to fix the problem.**



# Example: BNL-SARA Issues

Source: **lhcmn.bnl.gov** - 192.12.15.23  
Capacity: 10G MTU: 9000

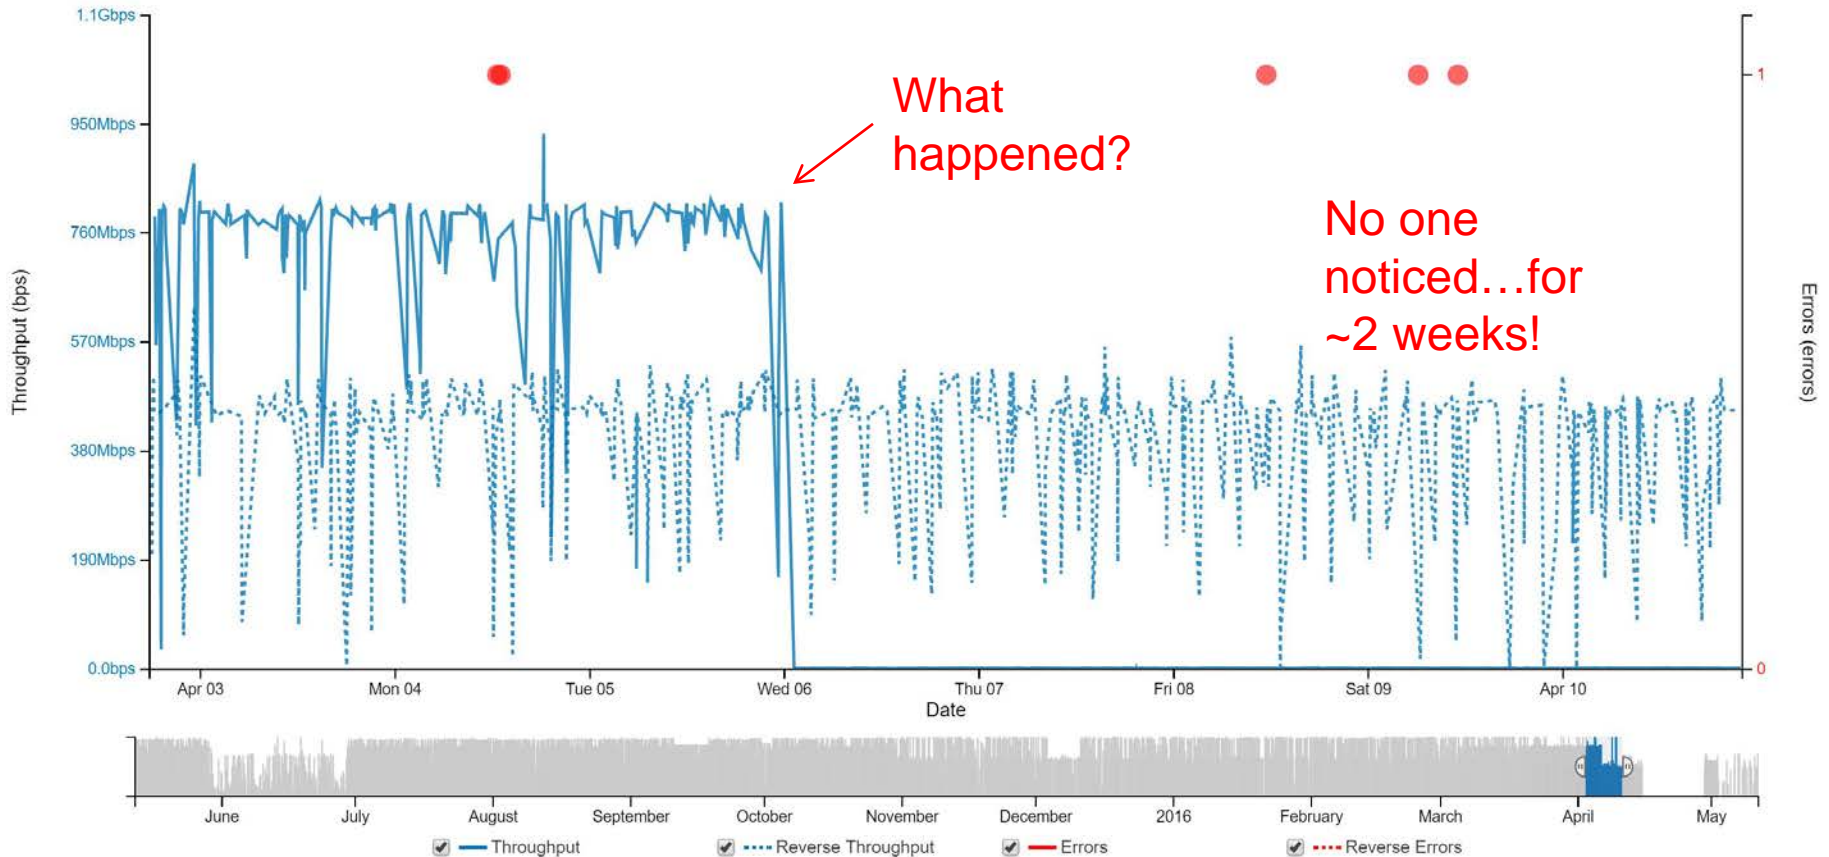
Destination: **ps.lhcopn-ps.sara.nl** - 145.100.17.9 [\[traceroute\]](#)  
Capacity: 1.0G MTU: 1500

[Link to this chart](#)

Zoom: 1d 3d 1w 1m 1y

Tue May 12 12:20:28 2015 -- Wed May 11 12:20:28 2016

[Previous 1y](#)





# Challenges for Alerting/Alarming

- ❄ We have ~125 sites which represent 15,625 paths
  - ❑ No way to have a person look at each one on a regular basis
  - ❑ Need to automate checks...simple first, add intelligence based upon experience
  
- ❄ We have multiple possible people (email lists) that could be notified
  - ❑ Need to gather and maintain contacts
  
- ❄ We must also maintain (and update) mappings between perfSONAR instances, experiment and WLCG site names

- ❄ OSG needs to be able to support "higher-level services" that require network metrics to make decisions regarding data transfers and higher-level workflow optimizations involving the network.
  - ❑ Proximity service prototyped identify "close" SEs and perfSONARs e.g: <http://proximity.cern.ch/api/0.1/trace/tracepath?src=atlas-npt2.bu.edu&dst=heplnx130.pp.rl.ac.uk>
- ❄ What metrics and with what timeliness are best for meeting this need?
  - ❑ This can be very complicated to answer in practice.
  - ❑ We will need to work closely with the developers of such services and iteratively adapt what is provided to make this as effective as possible.
  - ❑ First need is a network "cost-matrix" of source-destination pairs
  - ❑ Interaction with users will point the way to missing components

# OSG Analytics Platform?

- ❄ Should OSG provide some level of analytics capability for our network metrics?
  - ❑ “Leverage” ElasticSearch ?
  - ❑ Incorporate other metrics (end-to-end test results) to support “interesting” questions?
- ❄ Concern: Another service (and hardware) to develop and operate
- ❄ Discussion?

- ❄ Can we create tools to **manipulate, visualize, compare** and **analyze** network topologies from the OSG network datastore contents?
- ❄ Can we build upon these tools to create a set of next-generation network diagnostic tools to make debugging network problems easier, quicker and more accurate?
- ❄ Even without requiring the ability to perform complicated data analysis and correlation, basic tools developed in the area of network topology-based metric visualization would be very helpful in letting users and network engineers better understand what is happening in our networks.

# Graph Databases: Neo4j

← → ↻ <https://topo.aglt2.org:7473/browser/>

**Neo4j 2.2.1**

**Node labels**

- Node
- perfSONAR\_Bandwidth
- perfSONAR\_Latency **SE**
- Trace

**Relationship types**

- CONTAINS
- LINK
- NETLINK

**Property keys**

- abuse\_emails accuracy
- address cidr city country
- country\_code country\_name
- created description dst
- error\_message handle hash
- host hostname id ip lat
- long misc\_emails mtu name
- postal\_code query range rtt
- service\_type src state
- success tech\_emails
- time\_zone ttl updated

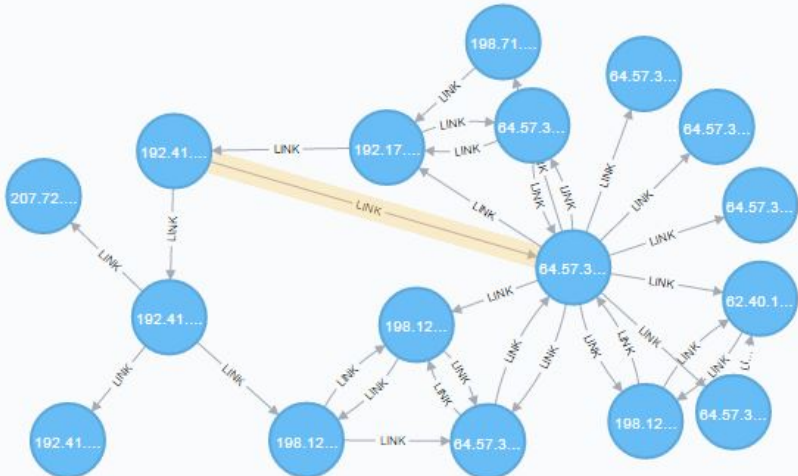
```
$ match (f:Node {ip:'192.41.230.59'})-[:LINK*1..2]->(x:Node) RETURN f, x
```

Graph

- \*(17) Node(17)
- perfSONAR\_Bandwidth(2)
- perfSONAR\_Latency(2)

\*(31) LINK(31)

Rows



LINK No properties

```
$ match (:Node) return count(*)
```

# Gather, Organize and Export Net Diagnostic “Work”



- ❄ In general, diagnosing and localizing network issues is difficult, even for experts.
- ❄ OSG should plan on making this process as straightforward as possible:
  - ❑ Collecting and organizing relevant information, automating as much of the process as possible.
    - ⌘ Mimic what net engineers do in gathering data and identifying issues.
  - ❑ Providing tools and tips to help describe, localize and characterize the problem
  - ❑ Package all the diagnostic information gathered to make it easy to hand off any debugging effort already worked on to other experts.

- ❄ Within the next few years evolving technology in the area of **Software Defined Networking(SDN)** may be able to provide researchers with the ability to construct their own Wide-Area networks with specified characteristics.
- ❄ What will OSG be able to do to integrate this type of capability with the rest of the OSG infrastructure?
- ❄ We need to plan for how best to enable evolving capabilities in the network for OSG users and admins
  - ❑ What is the impact on the OSG software stack?
  - ❑ What strategic modifications/additions are useful?



- ❄ There are a number of options to explore listed above. How much effort are we willing to devote for other possibilities?
  - ❑ Some argument to be made for adjusting as we move forward
- ❄ Effort to-date in OSG networking a combination of me, Operations ,Technology and Software/Release
  - ❑ As we add sites there will be a support load (me+Operations)
  - ❑ Exploring new capabilities may require help from Technology
  - ❑ Producing new packages may require Software/Release help
  - ❑ We leverage other efforts (WLCG, perfSONAR-PS developers, ESnet/MaDDash, satellite proposals)
  - ❑ We must ensure we continue to reliably gather and provide network metrics moving forward.



## ❄ **Priority 1:** Develop alarming / alerting for network problems

### ❄ **What:**

- ❑ Develop the ability to recognize network problems and **set alarm**
- ❑ When network problems can be localized, generate an **alert to the right contacts**
- ❑ Timescale to deploy prototype system: Fall 2016

### ❄ **How:**

- ❑ Initially GOC will scan for obvious problems and issue tickets
- ❑ Utilize “custom notification” features of Check\_MK (in OMD) to define specific criteria to alert when there are obvious problems
- ❑ **Development:** Find “level” shifts in path data...use ElasticSearch?
- ❑ Work with PuNDIT to enable for OSG
  - ⌘ Test ability to localize and alert and if successful, enable alerts
- ❑ Integrate with MyOSG/OIM to leverage existing capabilities

## ❄️ **Priority 2: Support higher-level services. Involves?**

- ❑ This is critical and primarily means providing an API the users and applications need to access the network information they need.
- ❑ Will require optimization to effectively support almost “real-time” data gathering that some services may require to steer their workloads or data transfer decisions.
- ❑ Must support end-site, VO based, test based and time-based queries
- ❑ Must function with OSG operational realities (access, resources)

## ❄️ **How?**

- ❑ Build upon the datastore and proximity service work
- ❑ Leverage Analytics Platform
- ❑ Use ANSE results and WLCG as initial clients to create/tune the API
- ❑ Will require Operations effort; may require Technology input?

# OSG Networking Work Plan



- ❄ **Continuing:** Develop tools to better support user experience in understanding, fixing and using the network
  - ❑ Work on automating perfSONAR configs for non-WLCG OSG sites
  - ❑ Improve visibility and usefulness of metric visualization
  - ❑ Expose topology information for diagnosis and network visualization
  - ❑ Create a way to “package” diagnostic information to hand-off initial problem troubleshooting to experts.
  
- ❄ Documentation will need to be augmented as we add capabilities for users
  
- ❄ The auto-mesh GUI should be provided as a standalone package to allow VOs and campuses to manage their perfSONAR deployments
  - ❑ Work with Soichi and OSG Software to finalize

## ❄ Other options should we focus on?

- ❑ How much effort?
- ❑ Adapt as we move forward (feedback from users, higher-level services; technology input for what is feasible/available)
- ❑ SDN may only start to be “real” after year 5
- ❑ Continuing need to do better at simplifying the process for finding and fixing problems in the network. Can we make a real change in how this is handled?

## ❄ Further discussion?

- Network Documentation  
<https://www.opensciencegrid.org/bin/view/Documentation/NetworkingInOSG>
- Deployment documentation for OSG and WLCG hosted in OSG  
<https://twiki.opensciencegrid.org/bin/view/Documentation/DeployperfSONAR>
- Measurement Archive (MA) guide  
[http://software.es.net/esmond/perfsonar\\_client\\_rest.html](http://software.es.net/esmond/perfsonar_client_rest.html)
- Modular Dashboard and OMD *Prototypes*
  - <http://maddash.aglt2.org/maddash-webui> [https://maddash.aglt2.org/WLCGperfSONAR/check\\_mk](https://maddash.aglt2.org/WLCGperfSONAR/check_mk)
- **OSG Production instances for OMD, MaDDash and Datastore**
  - <http://psmad.grid.iu.edu/maddash-webui/>
  - [https://psomd.grid.iu.edu/WLCGperfSONAR/check\\_mk/](https://psomd.grid.iu.edu/WLCGperfSONAR/check_mk/)
  - <http://psds.grid.iu.edu/esmond/perfsonar/archive/?format=json>
- Mesh-config in OSG <https://oim.grid.iu.edu/oim/meshconfig>
  - Being updated to a new standalone mesh-config application (ready for v3.6?)
- Use-cases document for experiments and middleware  
<https://docs.google.com/document/d/1ceiNITUJCwSuOuvbEHZnZp0XkWkwdkPQTQic0VbH1mc/edit>
- R&D projects <http://tinyurl.com/WLCG-R-D>

# ADDITIONAL SLIDES

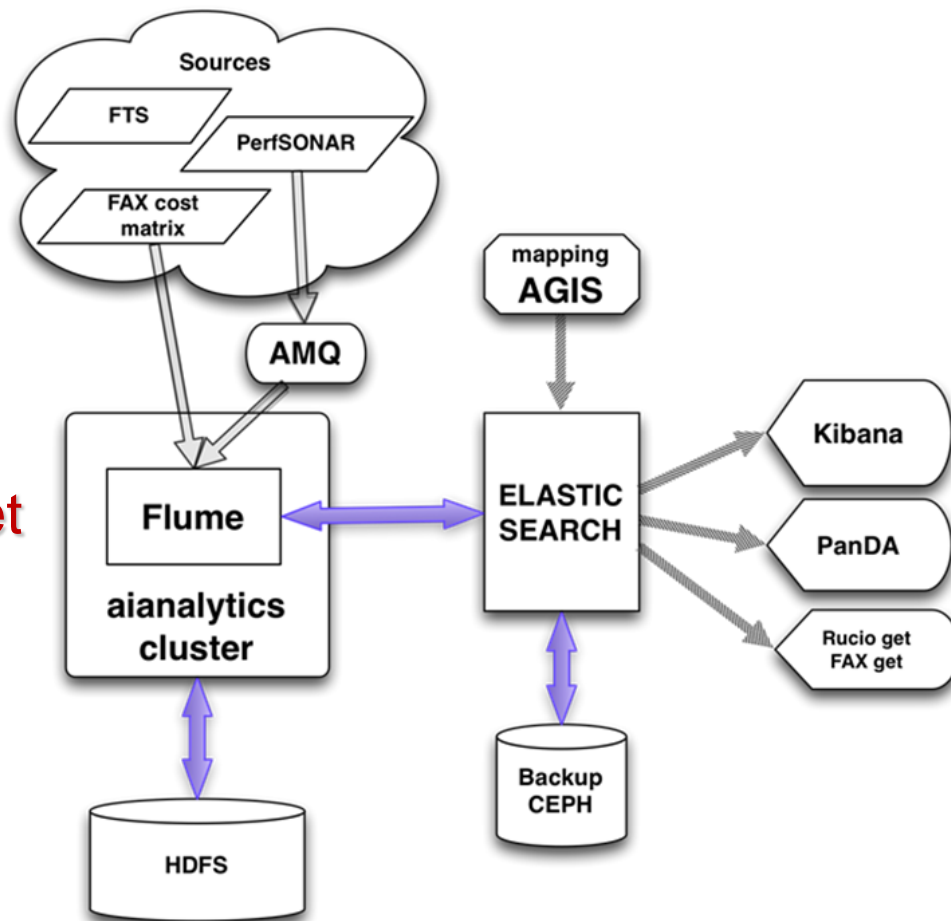
# ATLAS Network Analytics

❄ Diagram shows the flow

❄ End-to-end+perfSONAR  
data both available to  
jointly analyze

❄ Kibana can be used to get  
customized views

<http://cl-analytics.mwt2.org:5601>



❄ More details at: <http://tinyurl.com/gt92zwb>

## ❄ Core is our work with perfSONAR

- ❑ OSG provides lots of very useful feedback to the perfSONAR project and works close with Andy Lake/ESnet (as well as Jason Zurawski and Brian Tierney)
  - ⌘ Important because perfSONAR developers are targeting 100K deployments as the scale to support
- ❑ Main user (outside of ESnet) of MaDDash
- ❑ ESnet would really like Soichi's mesh-mgmt GUI as a package
  - ⌘ Sharing goes both ways...

## ❄ Working with ESnet closely on LHCONE and point-to-point network testbed

- ❑ Initial tests of “software-defined network” capabilities
- ❑ ESnet now serving US LHC Tier-2 sites



- ❄ PuNDIT was funded by the NSF (SSE-SI2) program:  
“PuNDIT will build upon the de-facto standard perfSONAR network measurement infrastructure to gather and analyze complex real-world network topologies coupled with their corresponding network metrics to identify possible signatures of network problems from a set of symptoms.”  
Website at <http://pundit.gatech.edu/>
- ❄ PuNDIT is currently using a number of OSG sites as a testbed.
- ❄ Project has ~8 montsh. Targeting initial deployment as part of perfSONAR v3.6/v4.0 (~Summer/Fall 2016)
- ❄ Goal is to enable PuNDIT for OSG/WLCG to find some kinds of near-real-time network issues
- ❄ Meeting TODAY to discuss how to get to this goal

# Using OWAMP Loss Metrics to Characterize Paths

- ❄ Throughput measurements are expensive so done at low frequency. Delays and packet loss rate are cheap.
- ❄ Idea is to use delays and packet loss rate to predict maximum possible throughput.
- ❄ Mathis formula is used to model impact of packet loss and latency on throughput
  - $\text{Rate} < (\text{MSS}/\text{RTT}) * (1 / \sqrt{p})$ 
    - ⌘ MSS – segment size
    - ⌘ RTT – round trip time
    - ⌘ p – packet loss
- ❄ Packet (re)ordering and jitter to be added as well



# Using Our Data

Host A is getting poor performance to Host B and seeing **3% packet loss**  
Normally we would start to investigate partial paths to isolate the problem



However we also see Host D to Host C is having problems and **2% packet loss**:

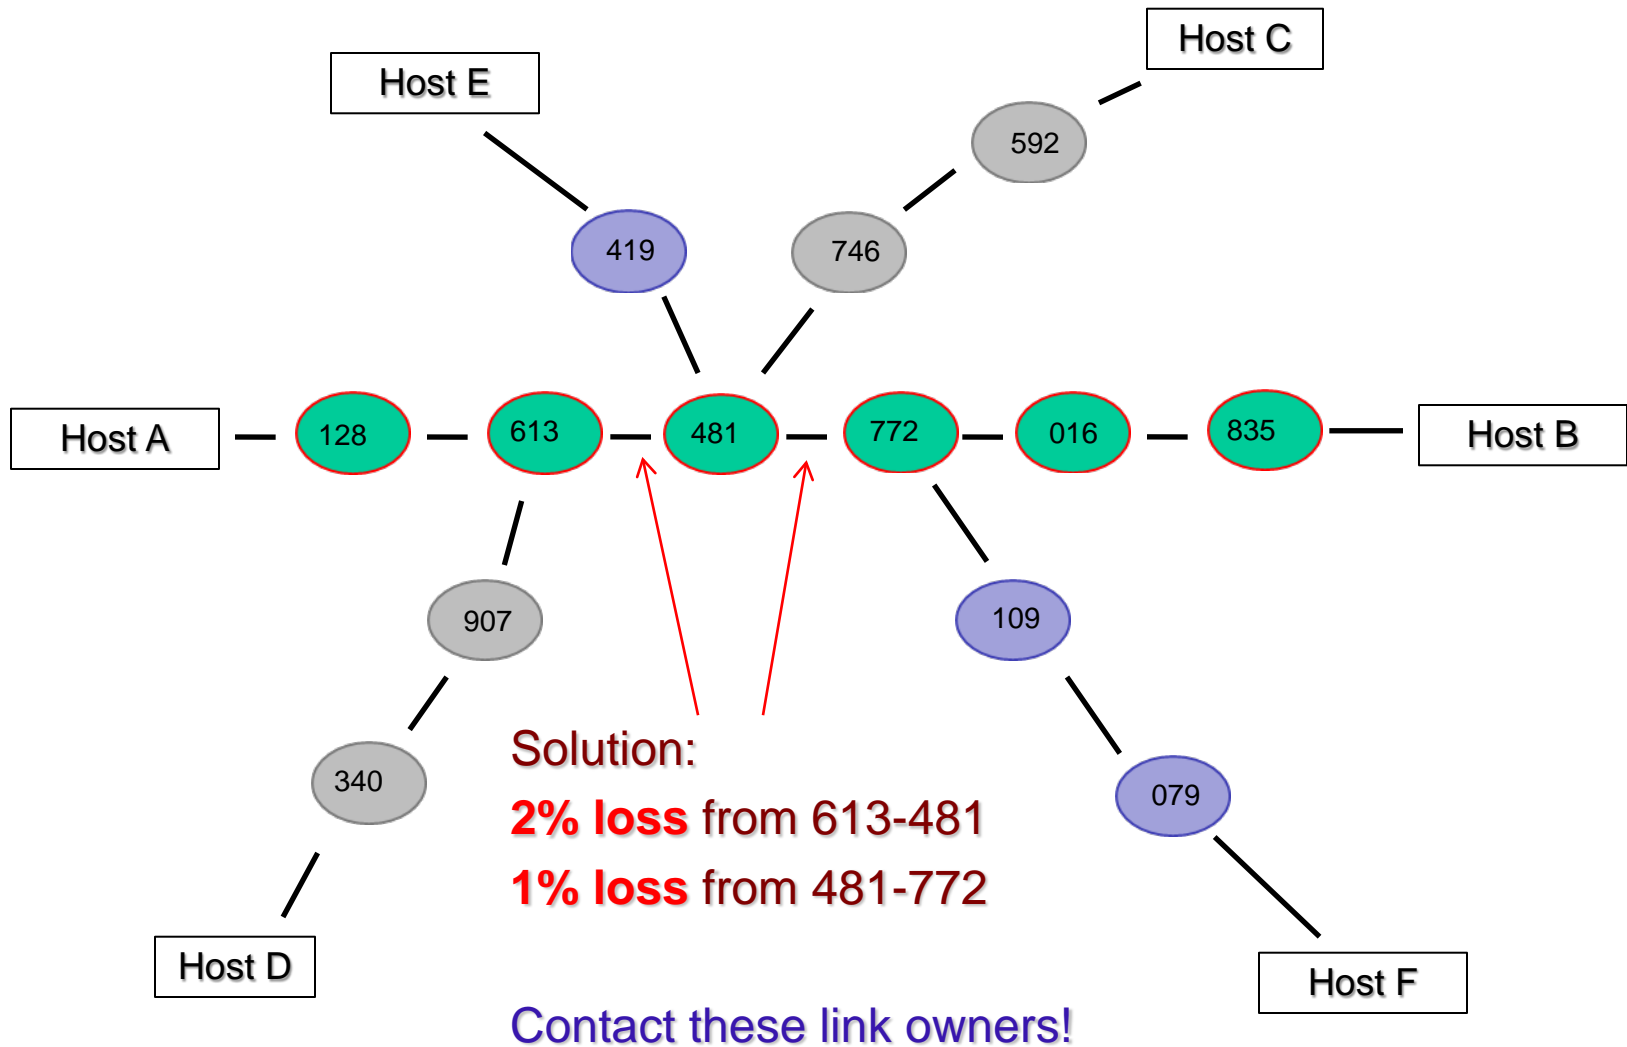


And there is a third pair (Hosts E and F) having **1% packet loss**:



Let's correlate these paths

# Topology View



# Longer Term OSG Networking Plans



- ❄ Years 4-5 need to build upon the first 3 years.
- ❄ What kinds of capabilities can we enable given a rich datastore of historical and current network metrics?
  - ❑ Users want "someone" to tell them when there is a network problem involving their site or their workflow.
  - ❑ Can we create a framework to identify when network problems occur and locate them? (\*Must\* minimize the false-positives).
- ❄ Issues that seem like "network issues" can often be due to problems at the ends (on the servers, in the software, in the configuration) or at least not WAN problems but LAN problems.

- ❄ Some important and interesting possibilities for what OSG might provide in the future include the creation of tools and visualization systems which manage network topologies (which are time-dependent)
  - ❑ Combining topology and metrics is powerful for identifying and localizing network problems; **currently a very manual process.**
- ❄ Using these tools users can look for correlations with the metrics measured across those topologies.
  - ❑ This type of tool can be used to help localize problems.
- ❄ **Note it is only by using the complete set of OSG network metrics that this becomes possible.**