Networking

Shawn McKee, Marian Babik
Co-chairs of WLCG Network and Transfer Metrics WG

ATLAS Sites Jamboree
27-29 January 2016

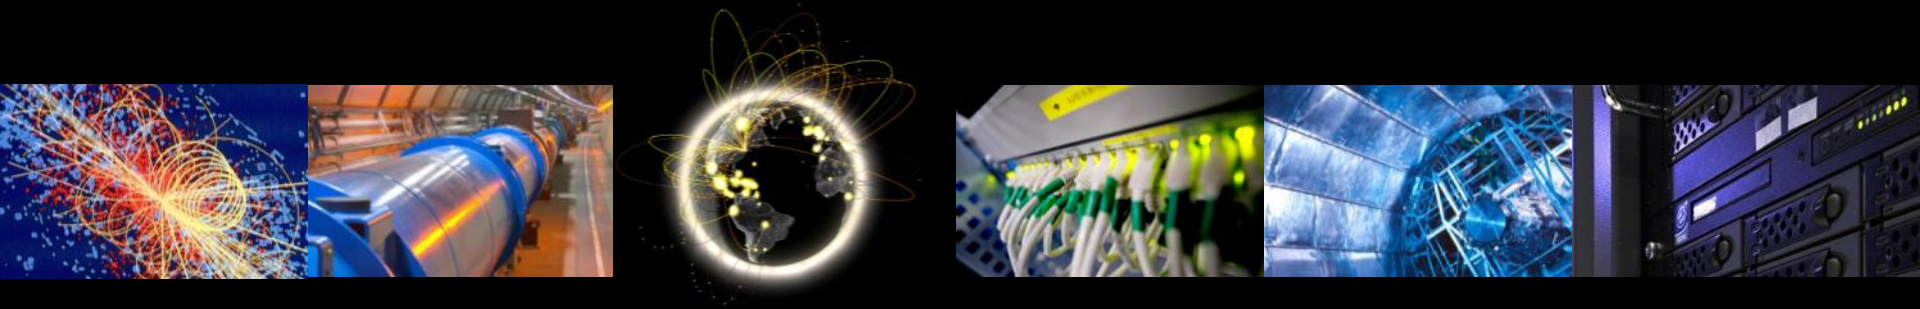

WLCG
Worldwide LHC Computing Grid

Open Science Grid

Overview of perfSONAR in WLCG/OSG

- Goals:
  - Find and isolate “network” problems; alerting in time
  - Characterize network use (base-lining)
  - Provide a source of network metrics for higher level services
- Choice of a standard open source tool: perfSONAR
  - Benefiting from the R&E community consensus
- Tasks achieved:
  - Finalized core deployment and commissioned perfSONAR network
  - Monitoring in place to create a baseline of the current situation between sites
  - Developed test coverage and made it possible to run “on-demand” tests to quickly isolate problems and identify problematic links 2

Importance of Measuring Our Networks

- End-to-end network issues are difficult to spot and localize
  - Network problems are multi-domain, complicating the process
  - Standardizing on specific tools and methods allows groups to focus resources more effectively and better self-support
  - Performance issues involving the network are complicated by the number of components involved end-to-end.
- perfSONAR provides a number of standard metrics we can use
- Latency measurements provide one-way delays and packet loss metrics
  - Packet loss is almost always very bad for performance
- Bandwidth tests measure achievable throughput and track TCP retries (using Iperf3)
  - Provides a baseline to watch for changes; identify bottlenecks
- Traceroute/Tracepath track network topology
  - All measurements are only useful when we know the exact path they are taking through the network.
  - Tracepath additionally measures MTU but is frequently blocked

Latency and packet loss matters

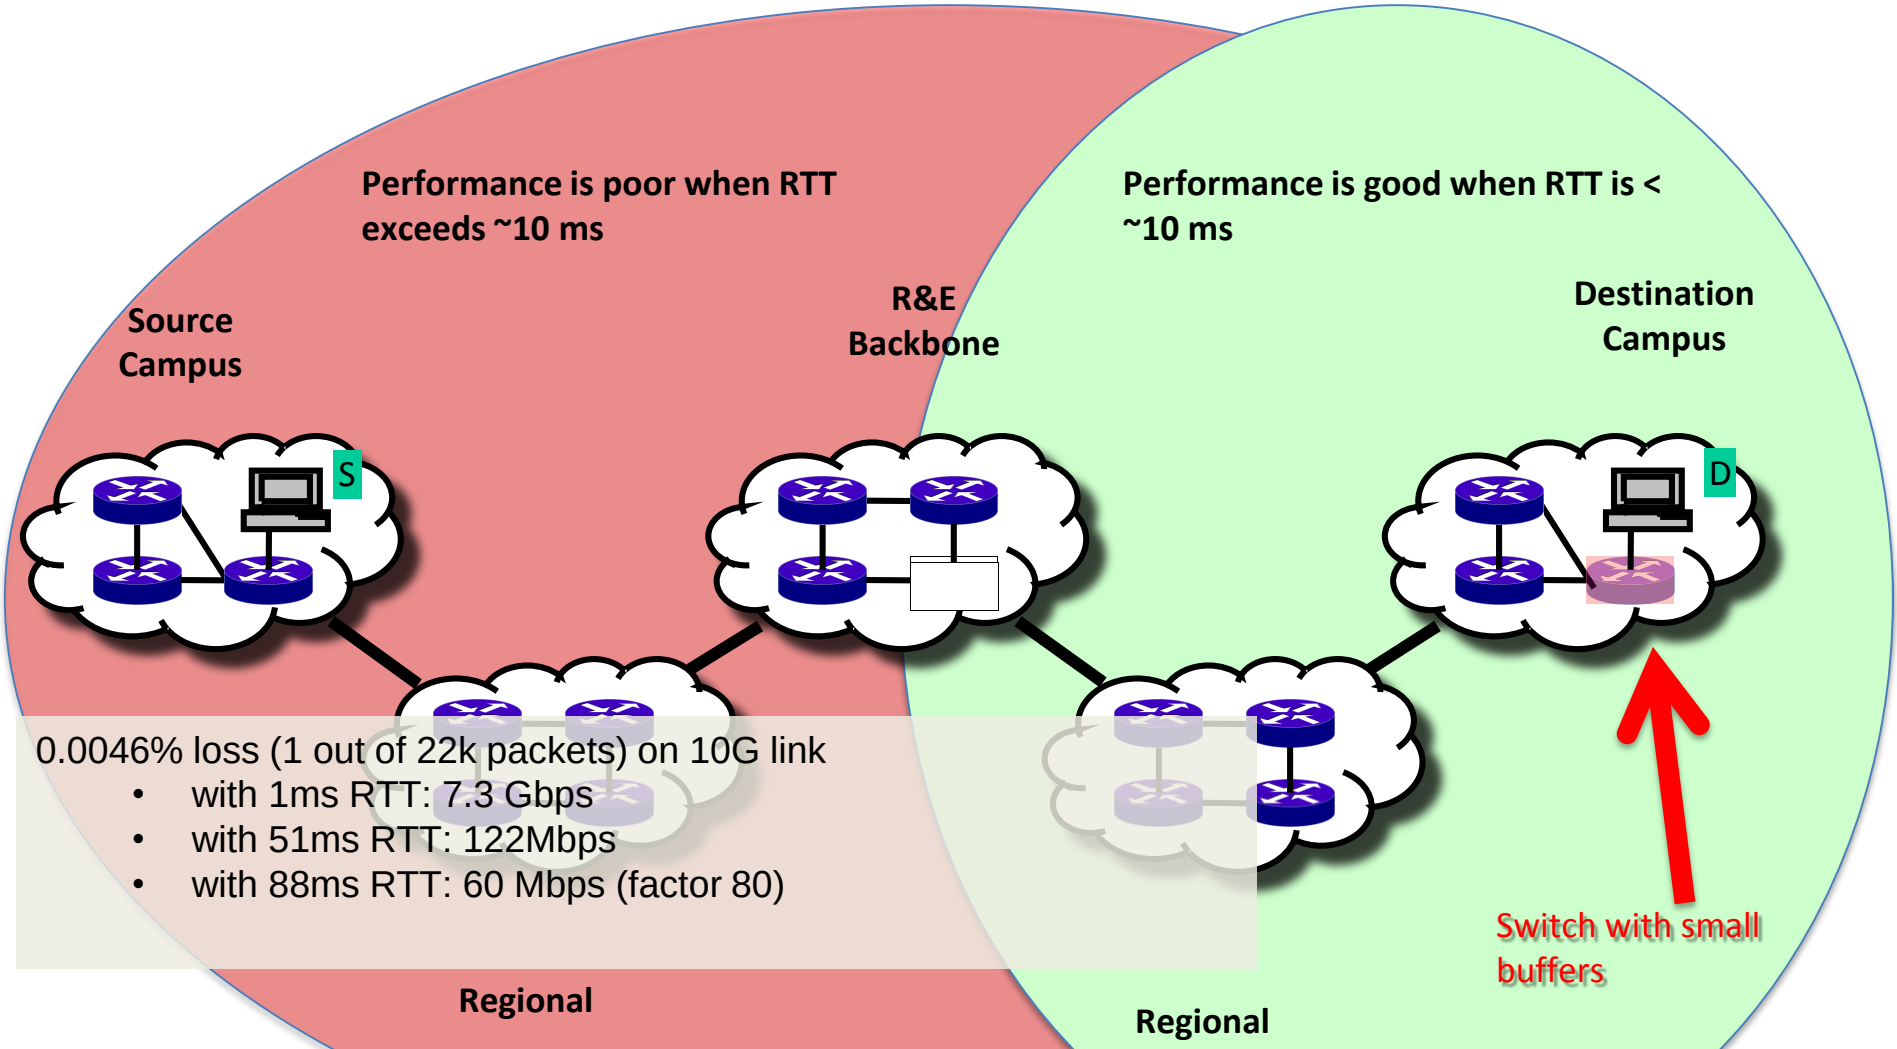

Credits Jason Zurawski, ESNet

Current perfSONAR Deployment

http://grid-monitoring.cern.ch/perfsonar_report.txt for stats

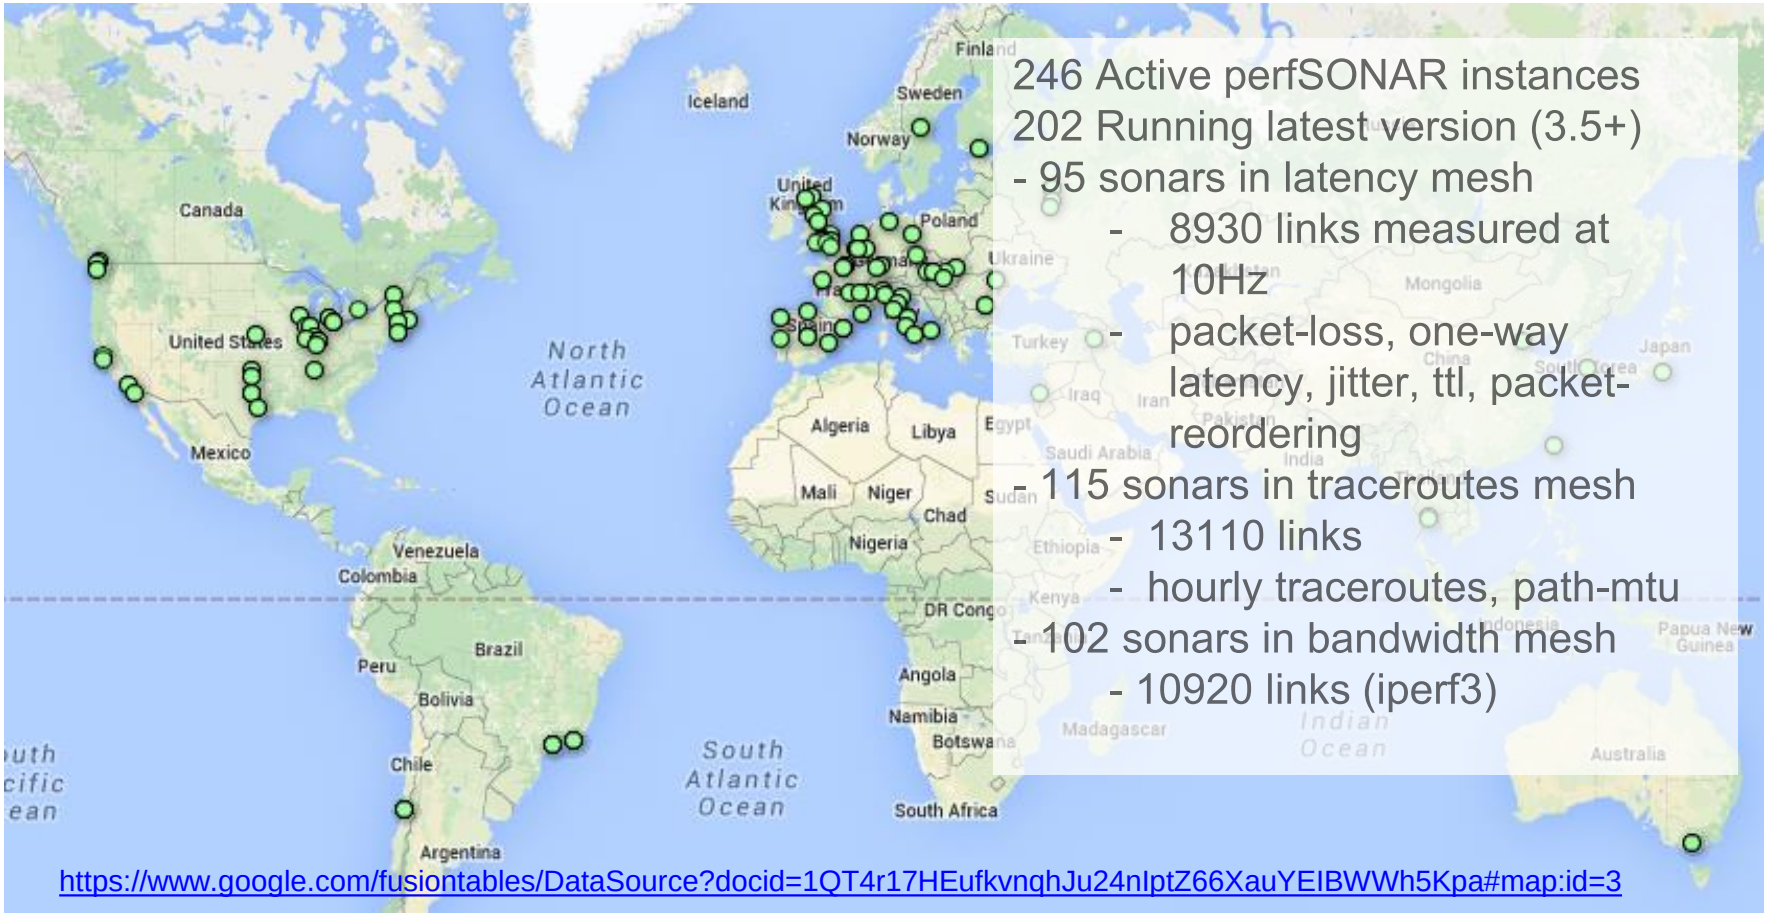

- Initial deployment coordinated by WLCG perfSONAR TF
- Commissioning of the network followed by WLCG Network and Transfer Metrics WG

perfSONAR v3.5 Toolkit

- perfSONAR v3.5 released on the 28th of September
- Features:
  - Support for low cost nodes ($200, LIVA)
  - Support for Debian, VMs, and other installation options
  - New GUIs
- Deployment on base OS via RPM bundles
  - perfSONAR Tools (just tools)
  - perfSONAR TestPoint (passive, no MA)
  - perfSONAR Core (+MA)
  - perfSONAR Complete (+Web and Toolkit Configuration)
- VMs - Still not recommended but possible, options
  - Full node VMs, VMs with dedicated physical NICs
  - OpenStack with pinning CPUs and SR-IOV
- Puppet and Docker support
- For support mail to wlcg-perfsonar-support or open GGUS ticket WLCG perfSONAR support

Existing Tools

- Number of tools available to help debug and understand network problems
  - iperf3, nutcp, bwctl, bwtraceroute, bwping, owping, scamper, etc.
- perfSONAR Toolkit installation gives you access to the entire network
  - tools can be used directly (run tests to/from a site)
  - but as well to run 3rd party tests (between two remote sites !)
- There are very good presentations on these tools in the training materials provided by ESNet/perfSONAR:
  - http://www.perfsonar.net/about/training-materials/
  - https://fasterdata.es.net/performance-testing/network-troubleshooting-tools/
  - https://fasterdata.es.net/performance-testing/troubleshooting/

Coordinating Network Issue Response

- The working group has created a support unit to coordinate responses to potential network issues
  - Tickets opened in the support group can be triaged to the right destination
  - Many issues are potentially resolvable within the working group
  - Real network issues can be identified and directed to the appropriate network support centers
- Documented at https://twiki.cern.ch/twiki/bin/view/LCG/NetworkTransferMetrics#Network_Performance_Incidents
- Most recent case CA<->EU GGUS-118730
  - resolved within hours of being reported
  - mainly due to our ability to narrow down using perfSONAR

Overview of perfSONAR Pipeline

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

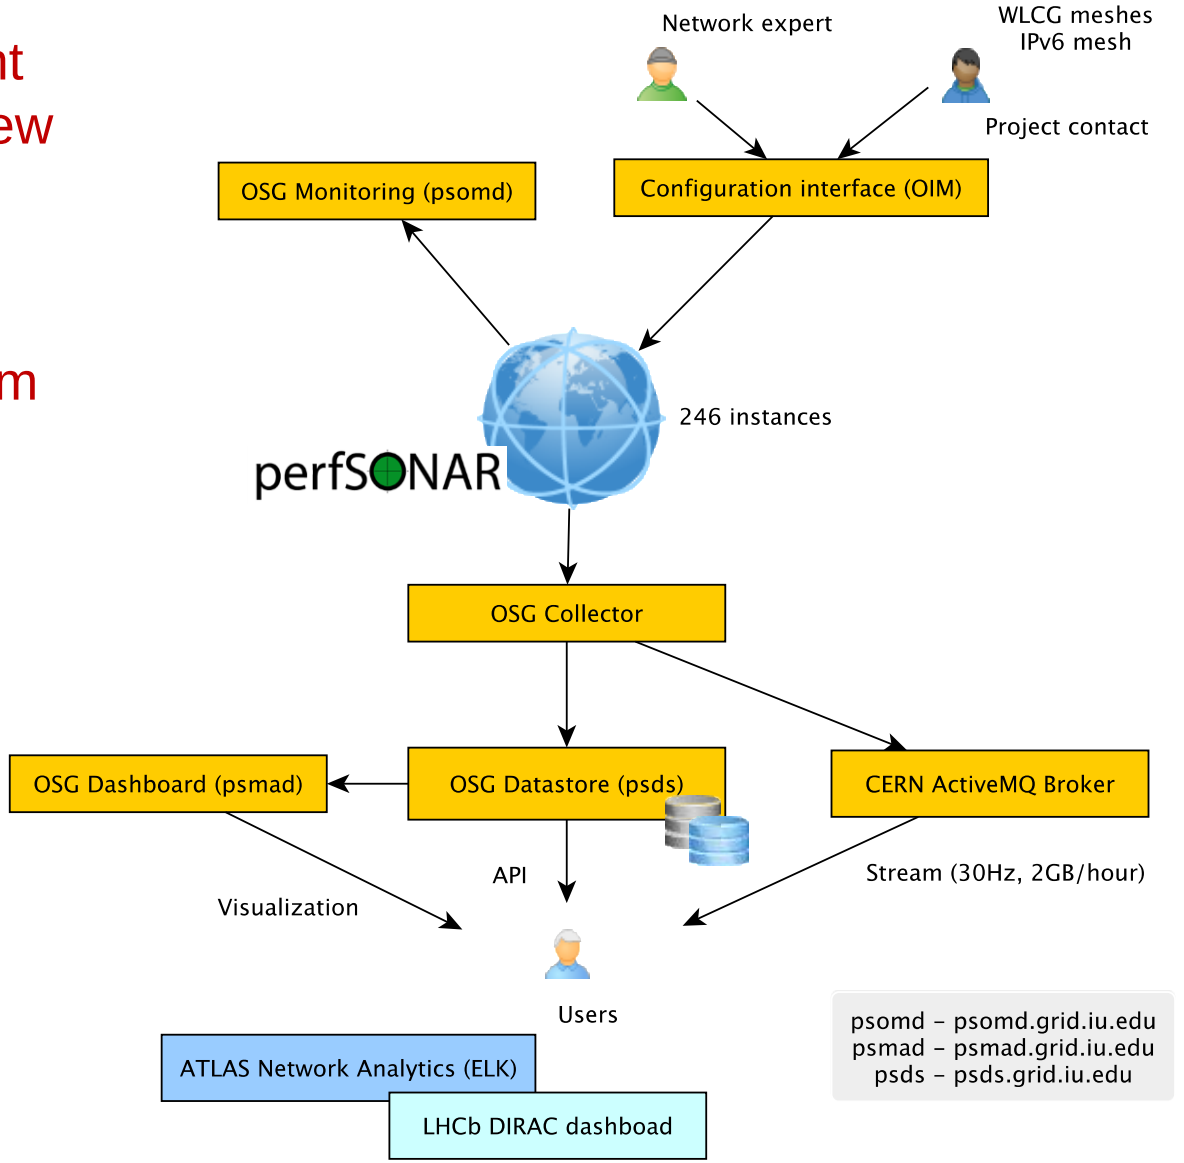

ATLAS Network Analytics

- Ilija Vukotic has been leading an effort to get network metrics into an analytics platform
- Diagram shows the flow

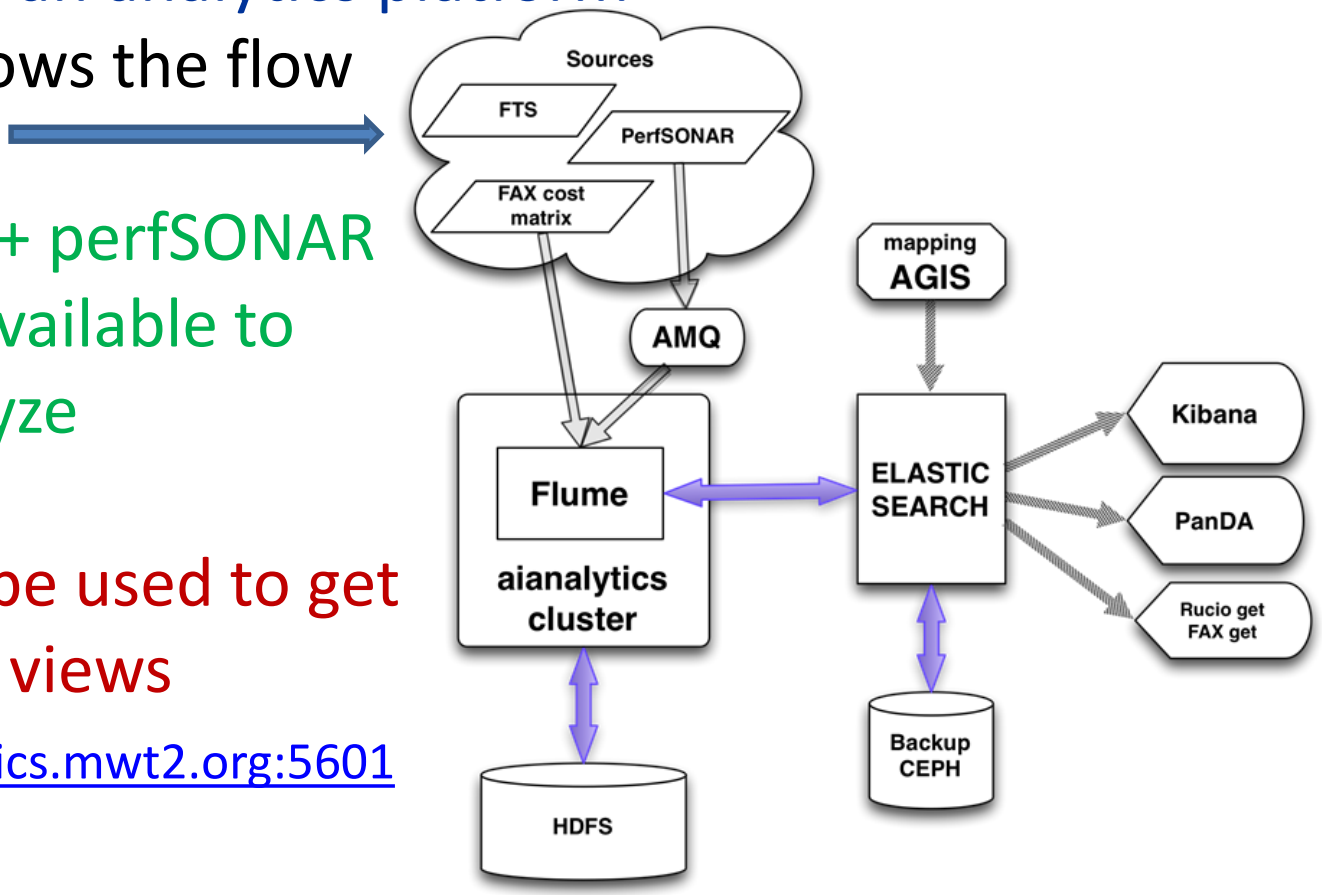

- End-to-end + perfSONAR data both available to jointly analyze
- Kibana can be used to get customized views

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

- See details in yesterday's Throughput meeting slides: http://tinyurl.com/gt92zwb

Playing with SDN in ATLAS

- 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.
- *For more details talk to Rob Gardner or Shawn McKee*

Summary

- We have a working infrastructure in place to monitor and measure our networks
- perfSONAR provides lots of capabilities to understand and debug our networks
  - New 3.5 provides new resiliency and install options
- Work on new applications is underway to make it more easier to find and fix problems
- We (OSG and WLCG) welcome feedback on how to further improve

References

- 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
- New 3.4 MA guide 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
- OSG Production instances for OMD, MaDDash and Datastore
  - http://psmad.grid.iu.edu/maddash-webui/
  - 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
- Use-cases document for experiments and middleware https://docs.google.com/document/d/1ceiNlTUJCwSuOuvbEHZnZp0XkWkwdkPQTQic0VbH1mc/edit