

WLCG/OSG Networking Update

Shawn McKee, Marian Babik

Spring 2018 HEPiX, Madison WI
14 May 2018



Overview

This is an update on the networking effort in OSG and WLCG

While there is a lot going on, I want to primarily cover two topics:

- In the **first part** I will be focusing on the changes in OSG and the resulting migration of services from the IU GOC
- In the **second part** I will discuss our and associated analytics

OSG Networking Components

- Network Monitoring via **perfSONAR**
 - Having perfSONAR fully deployed with a global dashboard is giving us powerful options for better management and use of our network
- A network **datastore** host all network metrics
- Tools to manage and maintain our infrastructure
 - A Modular dashboard (**MaDDash**); critical for quick “visibility” into networks. We can’t manage/fix/respond-to problems if we can’t “see” them.
 - **OMD/Check_mk** (used to monitor and verify the state of many globally distributed perfSONAR services); required to maintain the overall proper functioning of the monitoring infrastructure.
 - The development of the “**mesh-configuration**” and corresponding GUI interface; critical to creating a scalable, manageable deployment for WLCG/OSG
- Documentation --- Installation, debugging, How-tos
- Outreach and Support
 - With the network R&E community, VOs, software developers
 - OSG Support provides network ticket triage and routing

OSG Networking Evolution

The OSG network data pipeline has significantly evolved since its inception in 2012

- Initially we queried all **perfSONAR** toolkits centrally and stored the gathered metrics in ESmond
 - Users had a significant learning curve to be able to access and analyze the data; **it wasn't easy to do analytics**
- The system evolved to include parallel publication to an **ActiveMQ** bus at the end of 2015
 - Users could choose to subscribe to data of interest
 - Data could now easily be sent to Elasticsearch
- In 2017 we added a new destination to **RabbitMQ** (run by Nebraska and hosted in **AWS**)
 - New destination to Nebraska Elasticsearch
 - New 'long-term' repository for data
- In 2018 we are focusing on streamlining the system



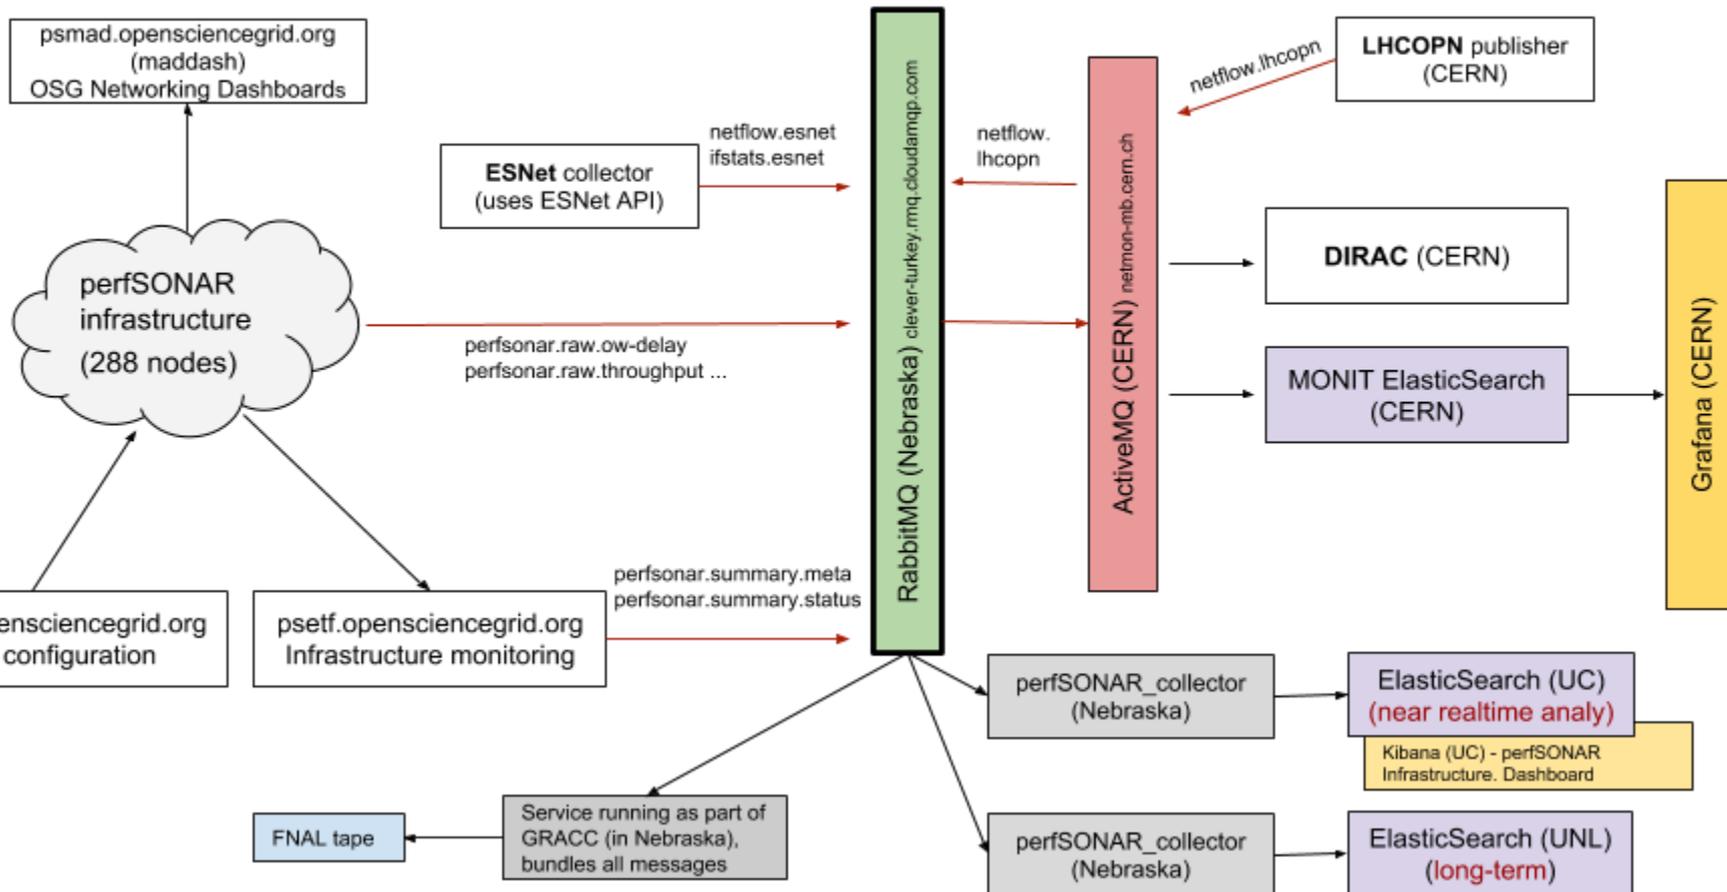
OSG Transition

- As many of you know the OSG is undergoing a transition.
 - Future funding will be coming as a combination of independently funded research grants and from the new S2I2 institute
- During this transition the OSG GOC at IU is being retired
 - All OSG services there are either being migrated or retired
- **The OSG networking services have already been migrated or shutdown**
 - AGLT2 (at MSU) is hosting the psetf, psconfig and psrv services (one production and one ITB of each)
 - Additional services are now in the cloud (RabbitMQ) or at University of Chicago or Nebraska

Target OSG Network Data Pipeline

Simpler, ALL data goes to RabbitMQ bus for fanning out to users and tools

May 12, 2018



Important Changes

- All URLs referring to *.grid.iu.edu are **deprecated**
 - Instead use *.opensciencegrid.org
- The central OSG network datastore is now hosted in Elasticsearch
 - The central OSG that used ESmond is turned off
- **Update your bookmarks!**
- Docs: <http://opensciencegrid.org/networking>
- MaDDash: <http://psmad.opensciencegrid.org/maddash-webui/>
- Meshconfig: <https://psconfig.opensciencegrid.org>

What you should know about perfSONAR v4.1

- Next release of perfSONAR is 4.1 and should be available “soon” (between now and July)
 - New snmp component will allow gathering local snmp data
- **perfSONAR 4.1 drops SL6 support which is the OS for most of our instances**
 - Our recommendation: **reinstall with CentOS7; don't worry about saving data**
- New Endpoint control capabilities
 - The mesh-config is being replaced by pSConfig. Lot's more options on what can be centrally managed and over-ridden on a per-instance basis
- pScheduler
 - Resource management – port pools
 - Pre-emptive scheduling support – improving client response time
 - **Requires port 443 to be open to all (potential) testing nodes**
- New plugins
 - Network traffic capture (via 'snmp')
 - Application-level (e.g. http response time)
- **Retirement of bwctl**
- **TWAMP support (two-way active measurement)**
 - ping alternative of owamp – routers/switches can participate in the tests
- Docker support



OSG Networking Data

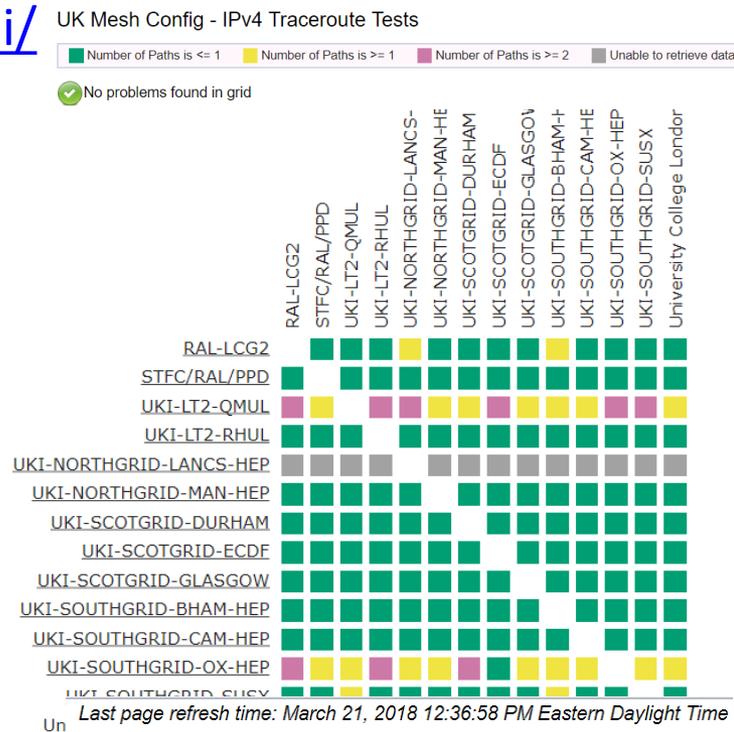
The OSG networking data pipeline is gathering many kinds of data

- perfSONAR metrics
 - **Latency** and **packet-loss** (10Hz of test packets per source-destination)
 - **Bandwidth** (one test per 4 or 6 hour window)
 - **Traceroute**: monitoring network path topology
- ESnet and LHCOPN interface and flow data
- FTS (File Transfer Service) metrics
 - Rate per file
 - Total traffic generated per site and link
 - Queue times

Monitoring Metrics

- We use MaDDash to view metric summaries
 - Provide quick view about how networks are working
- OSG hosts a production instance at:
<http://psmad.grid.iu.edu/maddash-webui/>

- Metrics are displayed via source-destination matrix
- Multiple dashboards (meshes) can be selected
- Custom menus link to relevant resources
- There is integrated analysis which detects common problems and provides advice



OMD/Check_MK Service Monitoring

We are using OMD & Check_MK to monitor our perfSONAR hosts and services. Provides useful overview of status/problems

https://psetf.opensciencegrid.org/etf/check_mk/

[Requires x509 in your browser; update for InCommon authentication planned]

The screenshot shows the Check_MK web interface. The top navigation bar includes the 'Check MK' logo, a 'Raw' data link, and a breadcrumb trail: 'All hosts' > '288 rows' > '/DC=ch/DC=cern/OU=Organic Units/OU=Users/CN=mckee/CN=500323/CN=Shawn Mc Kee (admin) 13:20'. The main content area is divided into two columns of host/service lists. The left column is titled 'Local site etf' and the right column is titled 'etf'. Each list has a header row with columns for 'state', 'Host', 'Icons', and a set of status metrics (OK, Wa, Un, Cr, Pd). The 'state' column uses color-coded indicators: green for 'UP', yellow for 'Warning', and red for 'Down'. The 'Icons' column shows small icons representing different services or components. The 'OK', 'Wa', 'Un', 'Cr', and 'Pd' columns contain numerical values representing the count of items in each state.

state	Host	Icons	OK	Wa	Un	Cr	Pd
UP	atogr009.nipne.ro		14	3	0	0	0
UP	btw-bw.t1.grid.kiae.ru		14	3	0	0	0
UP	btw-lat.t1.grid.kiae.ru		12	3	0	0	0
UP	cassandra2		52	0	0	0	0
UP	ccperfsonar2.in2p3.fr		12	3	0	0	0
UP	egperfer.ph.bham.ac.uk		18	3	0	0	0
UP	hcc-ps02.unl.edu		12	5	0	0	0
UP	heplnx130.pp.rl.ac.uk		12	5	0	0	0
UP	hepsonar2.ph.liv.ac.uk		12	3	0	0	0
UP	iut2-net2.iu.edu		14	3	0	0	0
UP	iut2-net4.iu.edu		14	3	0	0	0
UP	lapp-ps02.in2p3.fr		12	3	0	0	0
UP	lcg-lat.sfu.computecanada.ca		12	3	0	0	0
UP	lcgperfradar.dnp.fmph.uniba.sk		12	3	0	0	0
UP	lcgps01.gridpp.rl.ac.uk		12	3	0	0	0

state	Host	Icons	OK	Wa	Un	Cr	Pd
UP	btw-bw.grid.kiae.ru		14	3	0	0	0
UP	btw-lat.grid.kiae.ru		12	3	0	0	0
UP	cassandra1		52	0	0	0	0
UP	ccperfsonar1.in2p3.fr		12	5	0	0	0
UP	clrperf-owamp.in2p3.fr		12	3	0	0	0
UP	hcc-ps01.unl.edu		12	3	0	0	0
UP	hep-ps-bw-bp.pp.rl.ac.uk		16	1	0	0	0
UP	hepsonar1.ph.liv.ac.uk		12	5	0	0	0
UP	iut2-net1.iu.edu		12	3	0	0	0
UP	iut2-net3.iu.edu		13	2	0	0	0
UP	lapp-ps01.in2p3.fr		13	4	0	0	0
UP	lcg-bw.sfu.computecanada.ca		14	3	0	0	0
UP	lcg-sonar01.hep.ucl.ac.uk		18	3	0	0	0
UP	lcgperfsonar.dnp.fmph.uniba.sk		14	3	0	0	0
UP	lcgps02.gridpp.rl.ac.uk		12	5	0	0	0

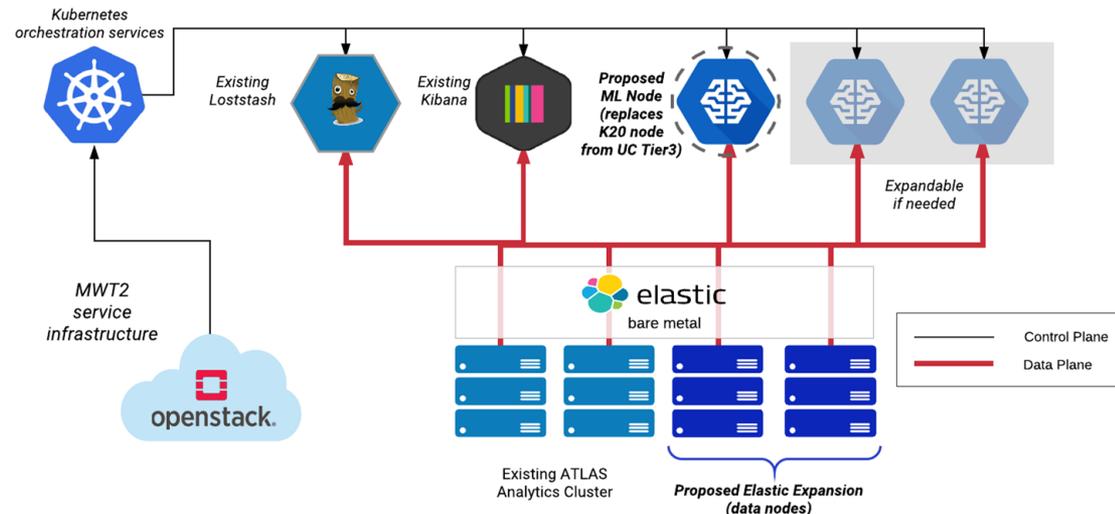


Analytics platform at UC

- Monitoring
 - Long retention
 - Fast & reliable
 - Has all data (ESnet, perfSONAR, FTS,...)
- Debugging issues
 - Has to have raw data
- Analysis
 - Open remote access to all the data
 - ML platform
 - Python, Jupyterl
 - Keras, Scikit-lear
 - GPU support
- Alarm & Alert

Recent Upgrade:

- Adding 5 ES nodes (doubling disk space)
- Adding dedicated ML node (8 GTX 1080 Ti)



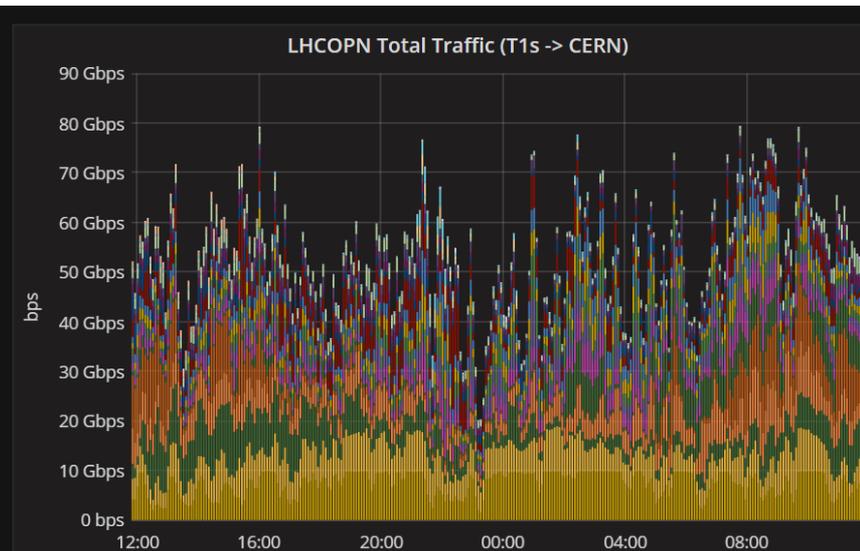
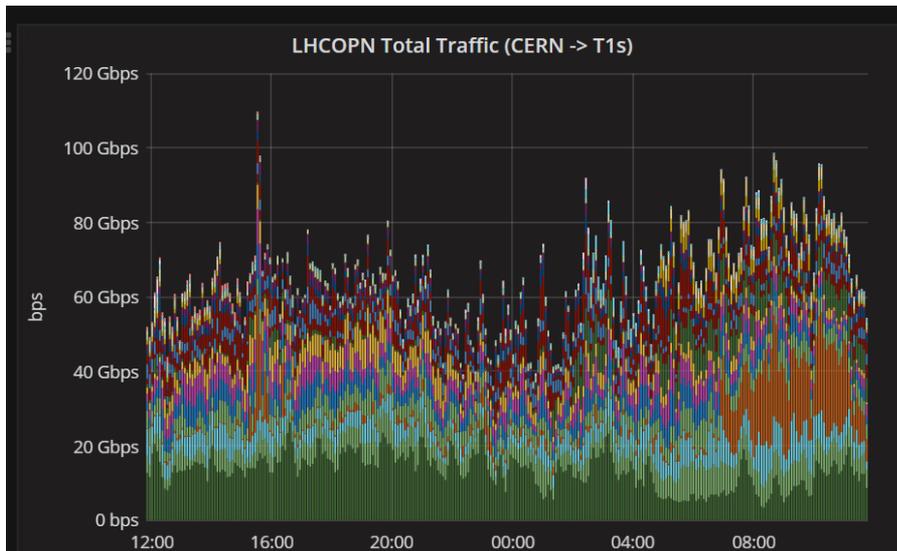
Example Grafana Dashboard

- Example dashboard showing side-by-side comparison of perfSONAR data, LHCOPN traffic and FTS transfers
- Also available with ESNet traffic data



LHCOPN Total Traffic

<http://monit-grafana-open.cern.ch/dashboard/db/lhcopn?orgId=16>



- US-T1-BNL primary Max: 26.2 Gbps Avg: 13.3 Gbps
- DE-KIT: secondary Max: 9.4 Gbps Avg: 5.2 Gbps
- RRC-KI-T1 LHCOPN Tier1 - Primary Max: 12.8 Gbps Avg: 4.9 Gbps
- US-FNAL-CMS primary Max: 33.8 Gbps Avg: 4.6 Gbps
- DE-KIT primary Max: 8.5 Gbps Avg: 4.3 Gbps
- UK-T1-RAL first Max: 9.9 Gbps Avg: 4.3 Gbps
- UK-T1-RAL second Max: 9.8 Gbps Avg: 4.2 Gbps
- ES-PIC primary Max: 8.5 Gbps Avg: 3.7 Gbps
- JINR-T1 primary Max: 10.3 Gbps Avg: 3.0 Gbps
- FR-CCIN2P3 secondary Max: 9.5 Gbps Avg: 3.0 Gbps
- CA-TRIUMF primary 10G Max: 5.5 Gbps Avg: 2.8 Gbps
- FR-CCIN2P3 primary Max: 7.5 Gbps Avg: 2.1 Gbps
- NL-T1 primary Max: 6.7 Gbps Avg: 1.7 Gbps
- NL-T1 secondary Max: 5.2 Gbps Avg: 1.7 Gbps
- UK-T1-RAL third Max: 3.4 Gbps Avg: 1.7 Gbps
- IT-INFN-CNAF primary Max: 5.0 Gbps Avg: 1.1 Gbps
- NDGF primary Max: 3.1 Gbps Avg: 1.0 Gbps
- IT-INFN-CNAF secondary Max: 4.2 Gbps Avg: 984 Mbps
- NDGF secondary Max: 4.6 Gbps Avg: 738 Mbps
- TW-ASGC primary Max: 3.8 Gbps Avg: 277 Mbps
- KR-KISTI Max: 2.5 Gbps Avg: 244 Mbps
- RRC-KI-T1 secondary 2 Max: 3 kbps Avg: 3 kbps
- RRC-KI-T1 secondary 1 Max: 3 kbps Avg: 3 kbps
- RRC-KI-T1 LHCOPN Tier1 - Secondary 3 Max: 3 kbps Avg: 3 kbps
- US-T1-BNL tertiary Max: 1 kbps Avg: 1 kbps
- US-T1-BNL secondary Max: 1 kbps Avg: 1 kbps
- US-FNAL-CMS tertiary Max: 848 bps Avg: 826 bps
- US-FNAL-CMS secondary Max: 677 bps Avg: 665 bps
- JINR-T1 LHCOPN Tier1 - Secondary Max: 406 bps Avg: 358 bps
- TW-ASGC Starlight Max: 169 bps Avg: 158 bps

- DE-KIT: secondary Max: 9.9 Gbps Avg: 7.3 Gbps
- ES-PIC primary Max: 9.1 Gbps Avg: 5.6 Gbps
- US-T1-BNL primary Max: 15.1 Gbps Avg: 5.1 Gbps
- RRC-KI-T1 LHCOPN Tier1 - Primary Max: 12.8 Gbps Avg: 4.9 Gbps
- US-FNAL-CMS primary Max: 19.2 Gbps Avg: 3.9 Gbps
- JINR-T1 primary Max: 10.3 Gbps Avg: 3.0 Gbps
- UK-T1-RAL second Max: 8.1 Gbps Avg: 2.6 Gbps
- NL-T1 primary Max: 9.9 Gbps Avg: 2.5 Gbps
- NL-T1 secondary Max: 9.9 Gbps Avg: 2.6 Gbps
- UK-T1-RAL third Max: 7.7 Gbps Avg: 2.7 Gbps
- UK-T1-RAL first Max: 7.7 Gbps Avg: 2.6 Gbps
- FR-CCIN2P3 primary Max: 8.1 Gbps Avg: 1.8 Gbps
- FR-CCIN2P3 secondary Max: 9.3 Gbps Avg: 1.8 Gbps
- CA-TRIUMF primary 10G Max: 5.3 Gbps Avg: 1.5 Gbps
- IT-INFN-CNAF secondary Max: 3.8 Gbps Avg: 1.3 Gbps
- DE-KIT primary Max: 3.4 Gbps Avg: 1.3 Gbps
- IT-INFN-CNAF primary Max: 3.6 Gbps Avg: 1.2 Gbps
- NDGF secondary Max: 3.4 Gbps Avg: 569 Mbps
- NDGF primary Max: 3.5 Gbps Avg: 540 Mbps
- KR-KISTI Max: 2.5 Gbps Avg: 245 Mbps
- TW-ASGC primary Max: 1.1 Gbps Avg: 113 Mbps
- RRC-KI-T1 secondary 2 Max: 3 kbps Avg: 3 kbps
- RRC-KI-T1 secondary 1 Max: 3 kbps Avg: 3 kbps
- RRC-KI-T1 LHCOPN Tier1 - Secondary 3 Max: 3 kbps Avg: 3 kbps
- TW-ASGC Starlight Max: 2 kbps Avg: 2 kbps
- US-T1-BNL tertiary Max: 1 kbps Avg: 1 kbps
- US-T1-BNL secondary Max: 1 kbps Avg: 1 kbps
- US-FNAL-CMS tertiary Max: 836 bps Avg: 807 bps
- US-FNAL-CMS secondary Max: 640 bps Avg: 627 bps
- JINR-T1 LHCOPN Tier1 - Secondary Max: 406 bps Avg: 358 bps



Alarming: First steps

- Simple alarms for now based only on PerfSONAR*
 - To make sure we have all the data, we set up alarms on issues with the data transport chain.
 - An alarm on capital problems at individual sites
 - if packet loss from a site to more than 6 other sites is greater than 2%. While highly specific, this test is not sensitive enough.
 - Coming soon: alarms on large increases in latencies and significant changes in path.

* For now FTS data doesn't tell us much about network performance (limited by the FTS itself).



Anomaly detection

Key system requirements:

- **Fast enough** - alarm to be raised in <3h.
- **Scalable** - with so many sites and links to monitor only minutes are available to evaluate (unless we go for a distributed system)
- **Sensitive** - down to single links.
- **Super Specific** - in a mesh that big, there are always some anomalies. Goal is to figure out which ones are significant enough.
- **Informative** - method should give us a debugging clue on:
 - **What happened.**
 - **Where it happened.**

Anomaly detection - cont.

Still some things to test:

- Data Smashing*
- Reinforcement Learning (Bootstrapped Q-learner)
- Create an OpenAI environment, see what others can come up with

In addition to alerts we could provide services predicting

* <http://rsif.royalsocietypublishing.org/content/royinterface/11/101/20140826.full.pdf>

Containerized analytics

A lot of people interested in trying out new methods, or debugging their cloud/site network.

To make it easy to start we created a container with all the tools (data access, neo4j for topology, jupyterlab, ML packages), codes, examples, documentation. **Have a look at:**

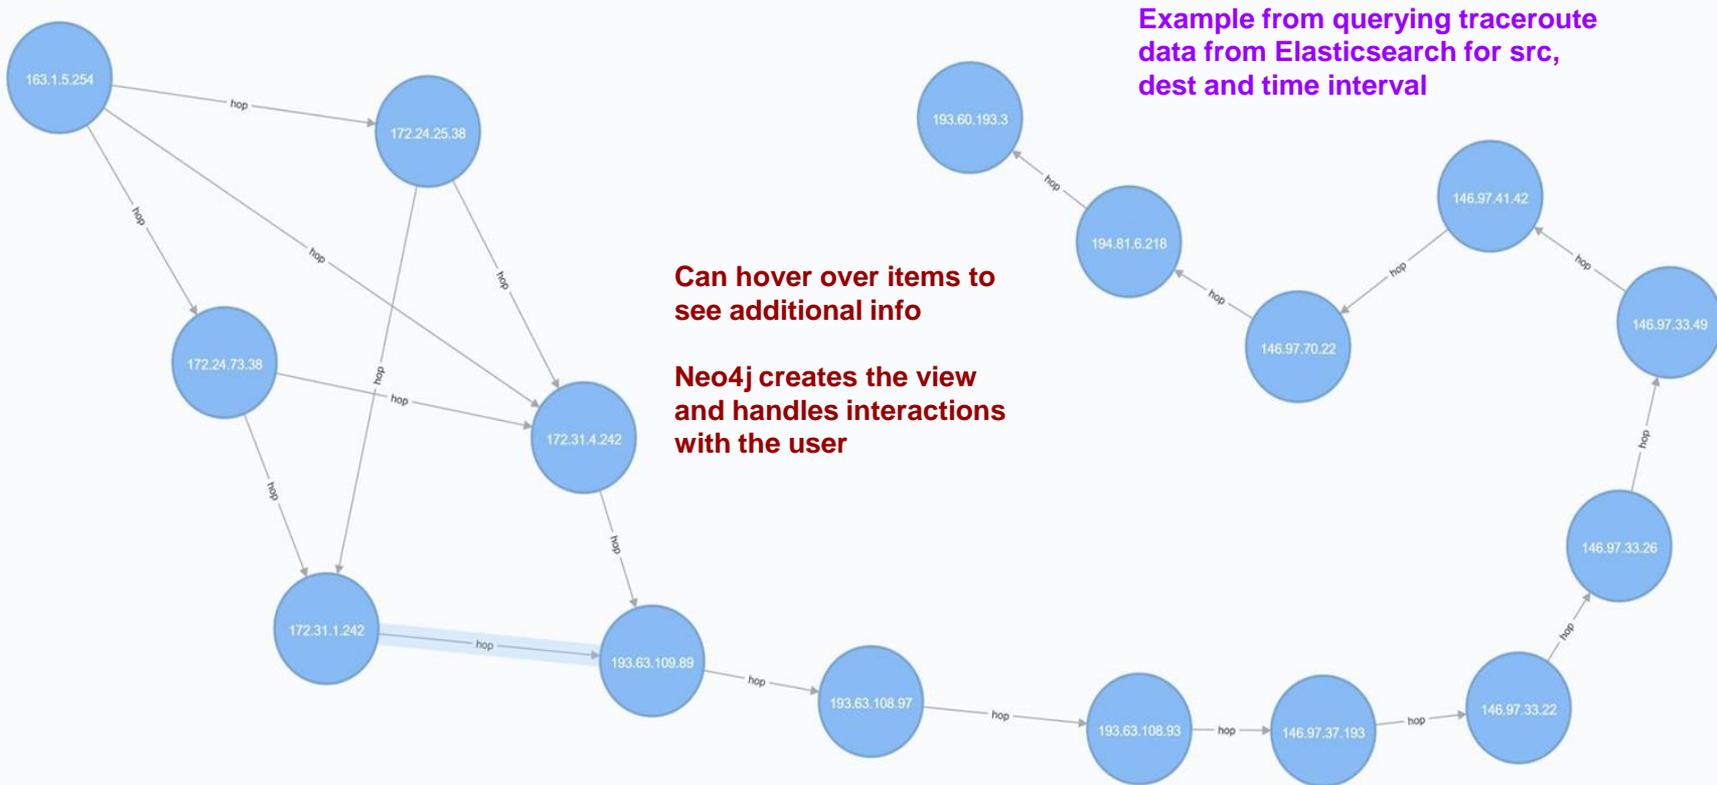
[opensciencegrid/network_analytics](https://opensciencegrid.org/network_analytics)

Contact Ilija Vukotic / UC if you have questions



Example Traceroute Visualization via Neo4J

We have tools to query the UC Elasticsearch for network path info and can send it to Neo4J



hop <id>: 2483 std_hop: 12.7830580754 hop: -4.73074792244 flow: 0.479097544791

Near-term Plans and Future Work

- **Updates to our network services**
 - perfSONAR v4.1 (will no longer support CentOS6)
 - MaDDash updates
 - pSConfig (replacing Meshconfig)
- **Updating our data pipeline to simplify and increase its robustness**
- Continue to iteratively debug, tune and optimize our set of network related alarms
- **Update/Upgrade Campaign for perfSONAR v4.1**
- **Enable more selective user-subscriptions to alerting based upon additional criteria**
- **Development of new ways to use our networking metrics**
 - For identification of network problems
 - To support problem diagnosis and localization
 - Improving user interfaces for network data exploration and use
- **Exploration of SDN (Software Defined Networking) capabilities as they become production ready**
- **Integration of network data with higher level services via our new analytics capabilities**

Summary

- We have made significant strides in making the network “visible” and easier to diagnose
 - ~250 OSG/WLCG perfSONAR deployments globally
 - All monitored, managed and orchestrated by OSG
 - Tools to manage/maintain our infrastructure are in place
 - We have production data pipeline providing both long term storage of all network metrics and near-realtime analytics capability
- Network dashboards can be used to get quick, specific details for subset of the network
- We have opportunities with our analytics platform to:
 - improve our network and storage resource use
 - enable identification of soft failures,
 - localize end-to-end problems
 - identify and remove bottlenecks

References / Links

- OSG networking documentation
 - <https://opensciencegrid.org/networking/>
- OSG Production instances for OMD, MaDDash and Datastore
 - https://psetf.opensciencegrid.org/etf/check_mk/
 - <http://psmad.opensciencegrid.org/maddash-webui/>
- OSG networking dashboards using the analytics platform:
 - <http://atlas-kibana.mwt2.org:5601/goto/7a9e388f4685965bcec6ced7705f7d9a>
 - <http://monit-grafana-open.cern.ch/dashboard/db/home?orgId=16>
 - <http://atlas-kibana.mwt2.org/app/kibana#/dashboard/Perfsonar-Alarms>
- Mesh-config in OSG <https://psconfig.opensciencegrid.org/>
- perfSONAR CentOS6->CentOS7 instructions (**do it now**): https://docs.perfsonar.net/install_migrate_centos7.html



Extra Slides



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?*)
- ***The GOAL: to make the most out of the networks we have!***

Vision: Networks Supporting Science

- While some of us are interested in (or worried about!) networks it is fair to say most scientists would rather not have to think about them.
 - Ideally networks are “transparent” and always do the right thing, allowing data to move as fast as possible, from anywhere to anywhere at anytime 😊
- The challenge is twofold:
 - Networks underlie all our distributed infrastructures and must work well for us to use our grids, clouds and HPC resources.
 - Problems in the network can be very hard to identify, isolate and fix
- OSG is working to better monitor, manage and diagnose our networks for all our benefit.

Status of Alerting and Alarming

- If you remember last years network presentation or saw the network analytics presentation yesterday, you know we have a prototype alerting system in place in OSG
- The alarming and alerting activity has been progressing but there have been some disruptions
 - Upgrades from the evolution of the data pipeline have, in some cases, broken existing alarms and dashboards
 - There are still “Unknown” instances that are not correctly mapped to their host sites by the collectors
 - We need to fix this, perhaps by adding an updateable lookup table
- Users are able to sign of for customized alerts via a Google spreadsheet.
 - Still fairly “raw”; alarms the will need to be improved to provide more user friendly (non-expert) information

Managing perfSONAR Deployments

- OSG originally developed a “meshconfig” GUI built within the OIM/MyOSG framework
 - We provided a GUI (MCA) to define and organize the regularly scheduled tests between specific sets of perfSONAR instances.
 - The mesh-config was a **huge** benefit; no longer need to use email to hundreds of system admins to make changes to network tests and their organization. The GUI made changes easy and consistent.
- **Problem:** not able to be made easily available to others within or outside OSG.
 - Campuses deploying many perfSONARs
 - Science VOs wanting to organize/customize their perfSONARs
- **The perfSONAR development team has taken over responsibility for this and will be providing updates and enhancements going forward**
 - Instructions are available for campuses and VOs to deploy their own MCA from containers hosted on GitHub
 - Next release of perfSONAR (v4.1) will change meshconfig to pSConfig (including and updated MCA).

LHOPN Detailed

<http://monit-grafana-open.cern.ch/dashboard/db/lhcopn-detailed?orgId=16>

LHCOPN Outbound Traffic Top 5

Metric	Avg	Max	Current
US-T1-BNL primary	13.32 Gbps	26.19 Gbps	8.37 Gbps

Karlsruhe Institute of Technology (DE-KIT)



- CERN to KIT primary Max: 8.48 Gbps Avg: 4.30 Gbps
- CERN to KIT secondary Max: 9.41 Gbps Avg: 5.23 Gbps
- KIT to CERN primary Max: 3.42 Gbps Avg: 1.26 Gbps
- KIT to CERN secondary Max: 9.86 Gbps Avg: 7.36 Gbps

6.53 Gbps
4.03 Gbps
5.30 Gbps
4.40 Gbps

LHCOPN Inbound Traffic Top 5

Metric	Avg	Max	Current
DE-KIT: secondary	7.36 Gbps	9.86 Gbps	8.37 Gbps

Fermi National Accelerator Laboratory (US-FNAL)



- CERN to FNAL primary Max: 33.8 Gbps Avg: 4.6 Gbps
- CERN to FNAL secondary Max: 677 bps Avg: 667 bps
- FNAL to CERN primary Max: 19.2 Gbps Avg: 3.9 Gbps

ES-PIC p
US-T1-E
RRC-KI-
Tier1 - F
US-FNA
primary



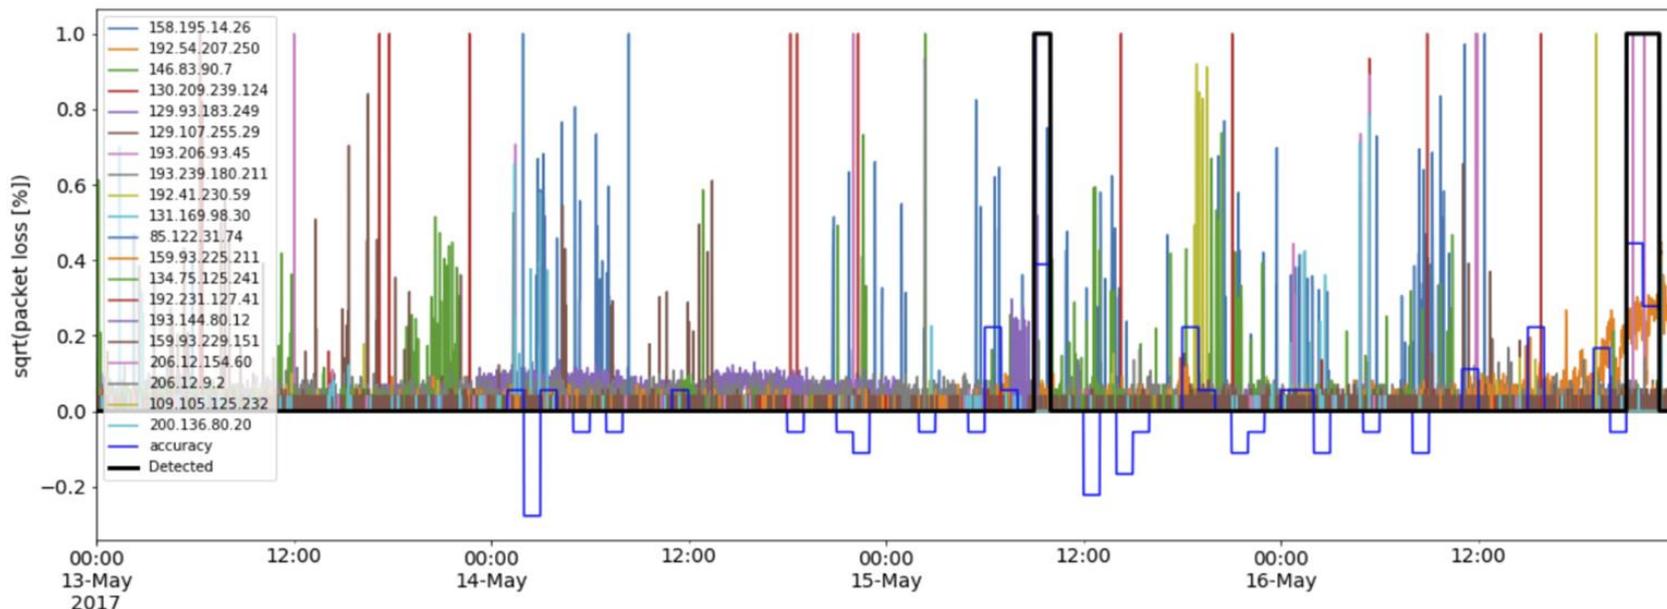
Anomaly detection - cont.

Tested few approaches:

Single class SVM - dimensionality problem

BASIC - Bayesian inference of simultaneous change-points in multiple timeseries - slow.

BDT and ANN - works well but needs validation



[*https://github.com/ATLAS-Analytics/AnomalyDetection/raw/f5851dfddd74fd6d053b7a5de64ac416c89f8f58/PerfSONAR/Paper/AD.pdf](https://github.com/ATLAS-Analytics/AnomalyDetection/raw/f5851dfddd74fd6d053b7a5de64ac416c89f8f58/PerfSONAR/Paper/AD.pdf)

Infrastructure Alerting

We have monitoring of the network dataflow into the Elasticsearch instances

Emails are sent to the core OSG/WLCG team if there are issues seen (drastic changes in information flow)

The email plot is clickable and can take us to a [Timelion](#) visualization accessible via Kibana



Alert on Elastic indexing rate [PerfSonar]

Dear Shawn McKee, this mail is to let you know that there is an issue in indexing Perfsonar data in UC Elasticsearch.

