LHCOPN/LHCONE perfSONAR Update

Shawn McKee / U Michigan

LHCONE/LHCOPN Meeting

https://indico.cern.ch/event/681168/

Cosener's House Abingdon, UK

March 6th, 2018

Overview of Talk

- LHCONE and LHCOPN infrastructure overview
  - Status and changes in our meshes
- Tools Status and ongoing projects
  - New perfSONAR release, network analytics, mesh-config
- Other Items

Acknowledgements: Marian Babik and Ilija Vukotic led much of the work on the new, interesting items I will cover.

LHCONE MaDDash – 04 Apr 2017

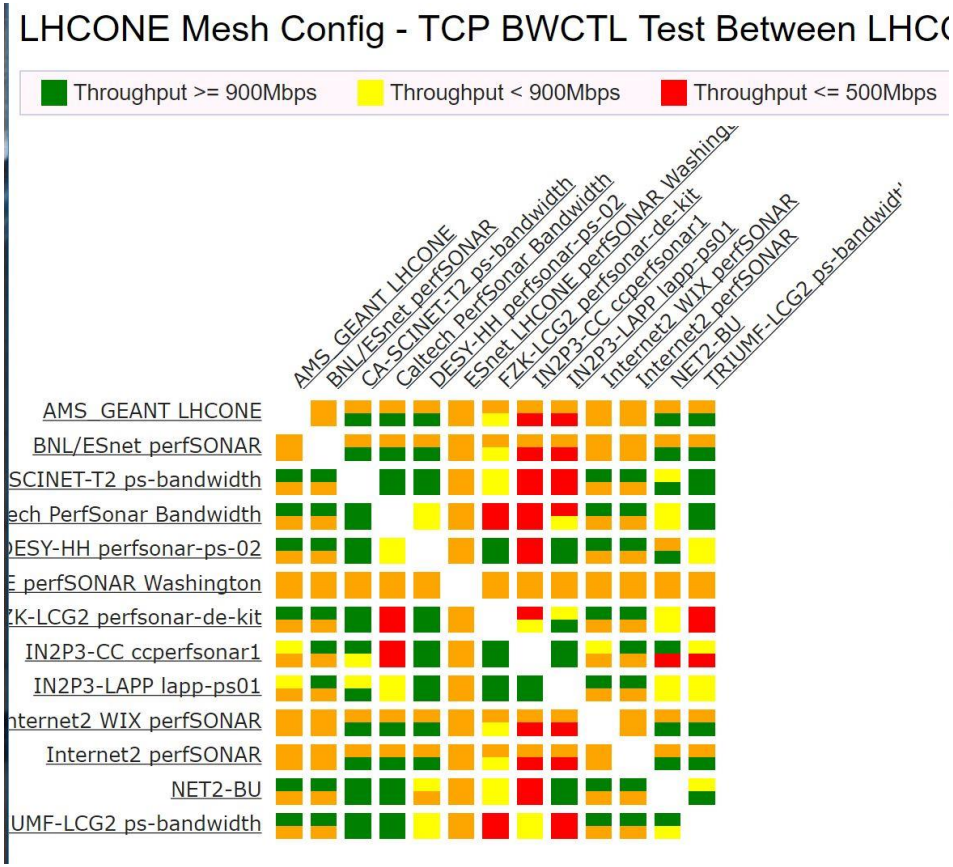

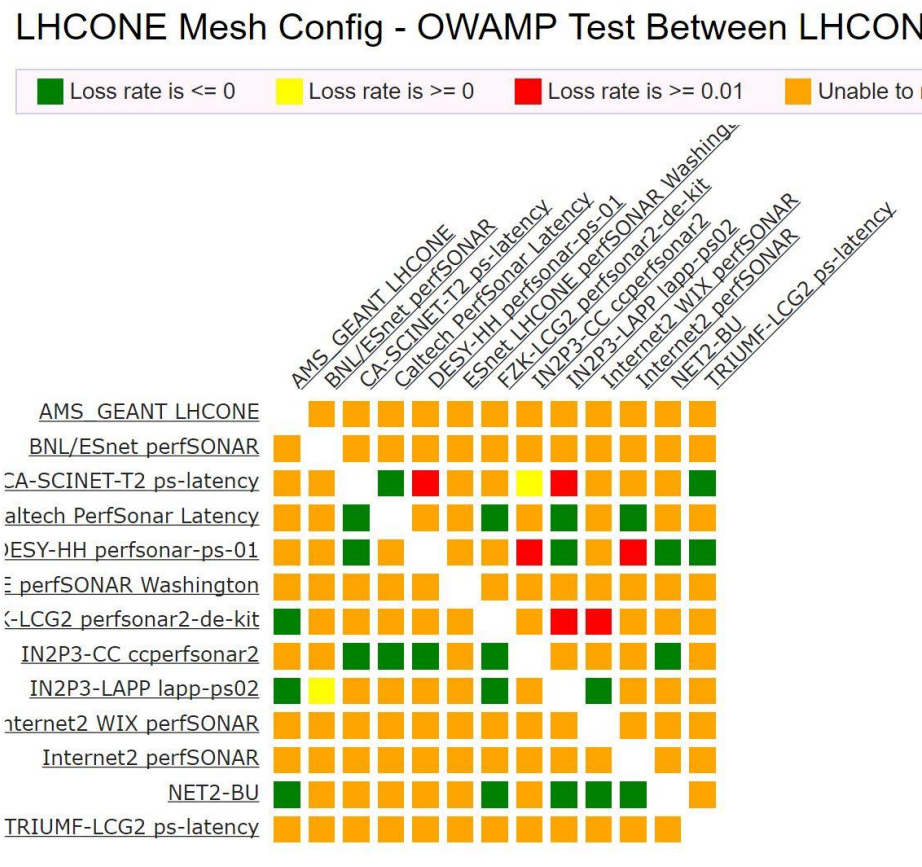

As you can see we have had lots of degradation in the SERVICE itself. Missing tests and results.

LHCONE MaDDash – 5 Mar 2018

LHCONE Mesh Config - TCP BWCTL Test Between LHCONE Bandwidth Hosts

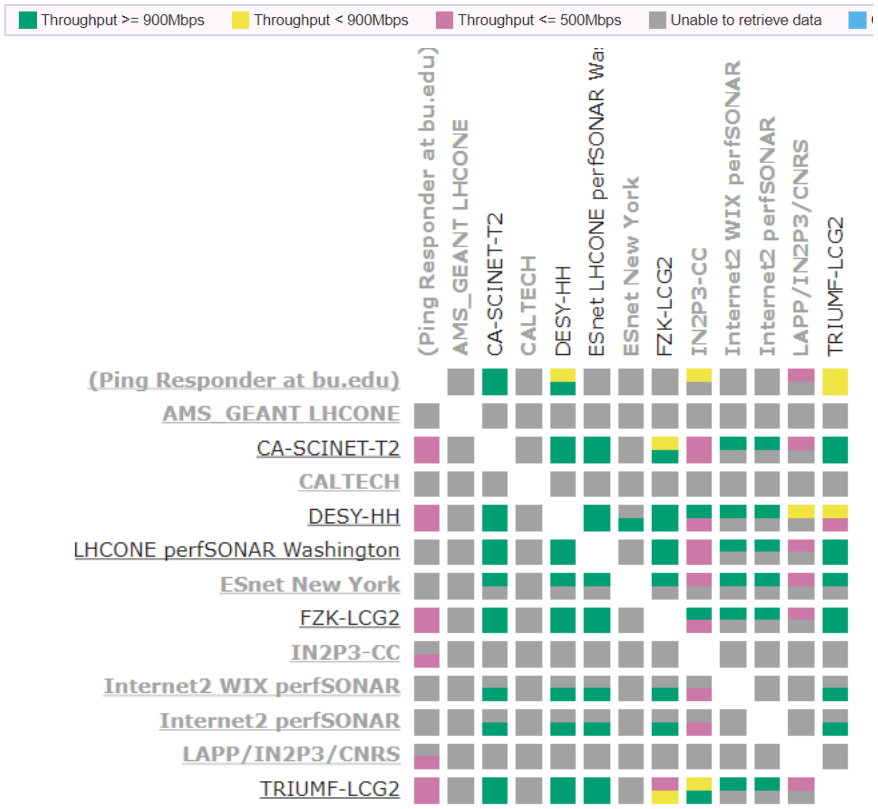

LHCONE Mesh Config - OWAMP Test Between LHCONE Latency Hosts

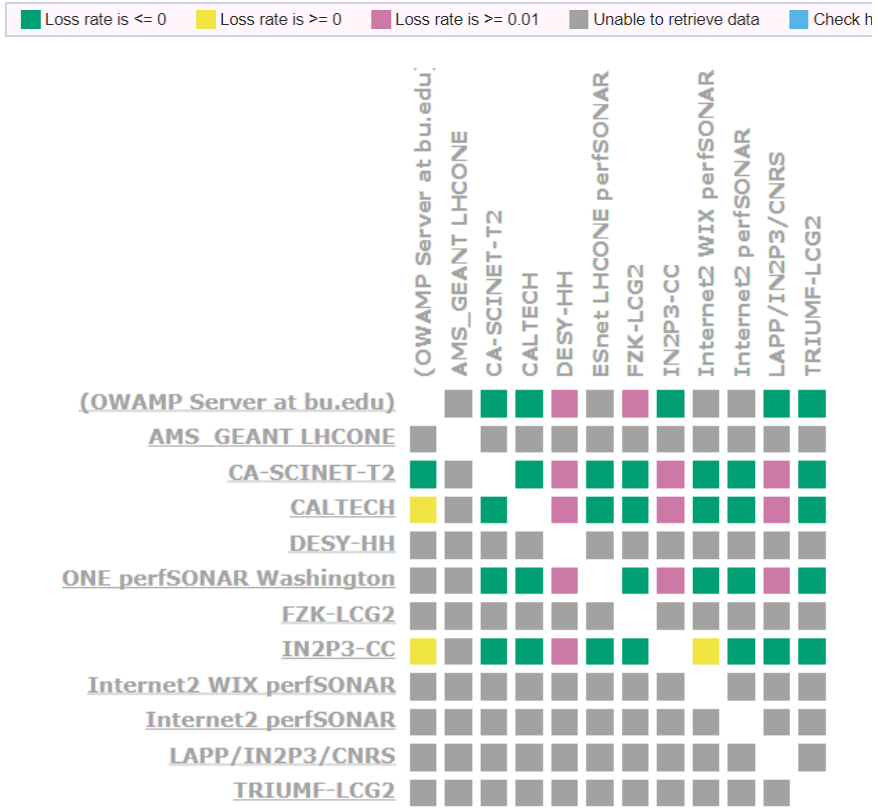

Compared to April 2017, things seem a bit better regarding available data

NOTE: New version of MaDDash and “unknown” data is now grey instead of orange

LHCOPN MaDDash – 18 Sep 2016

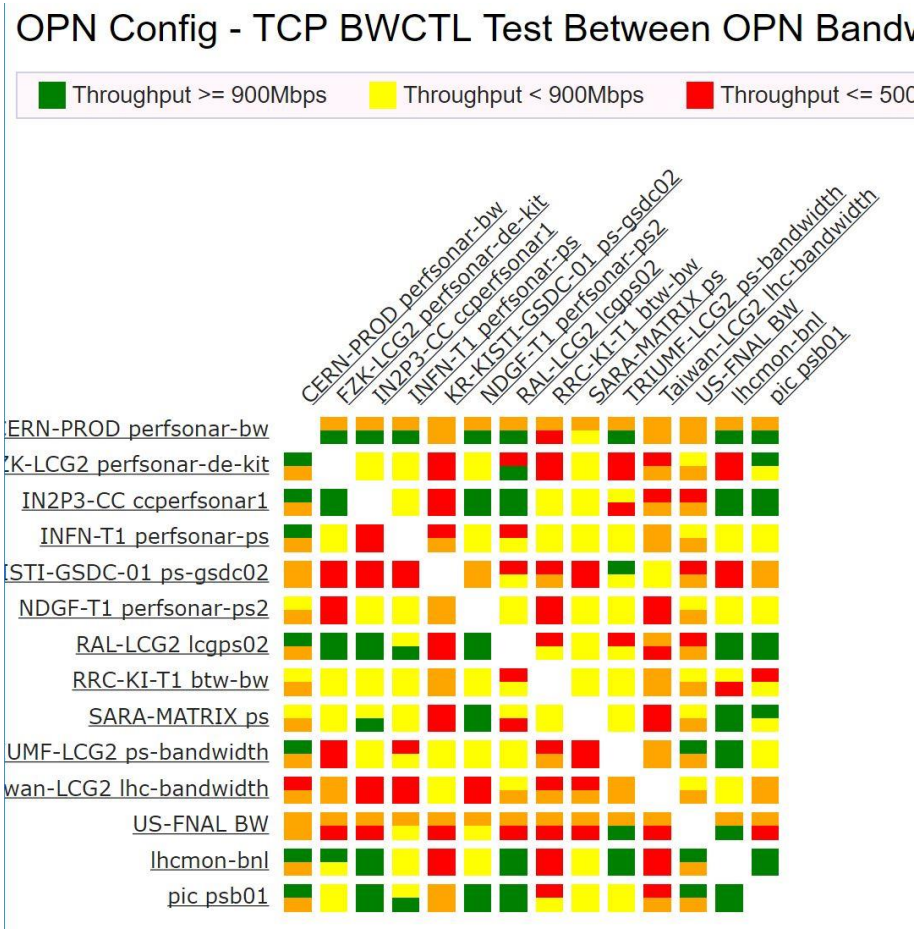

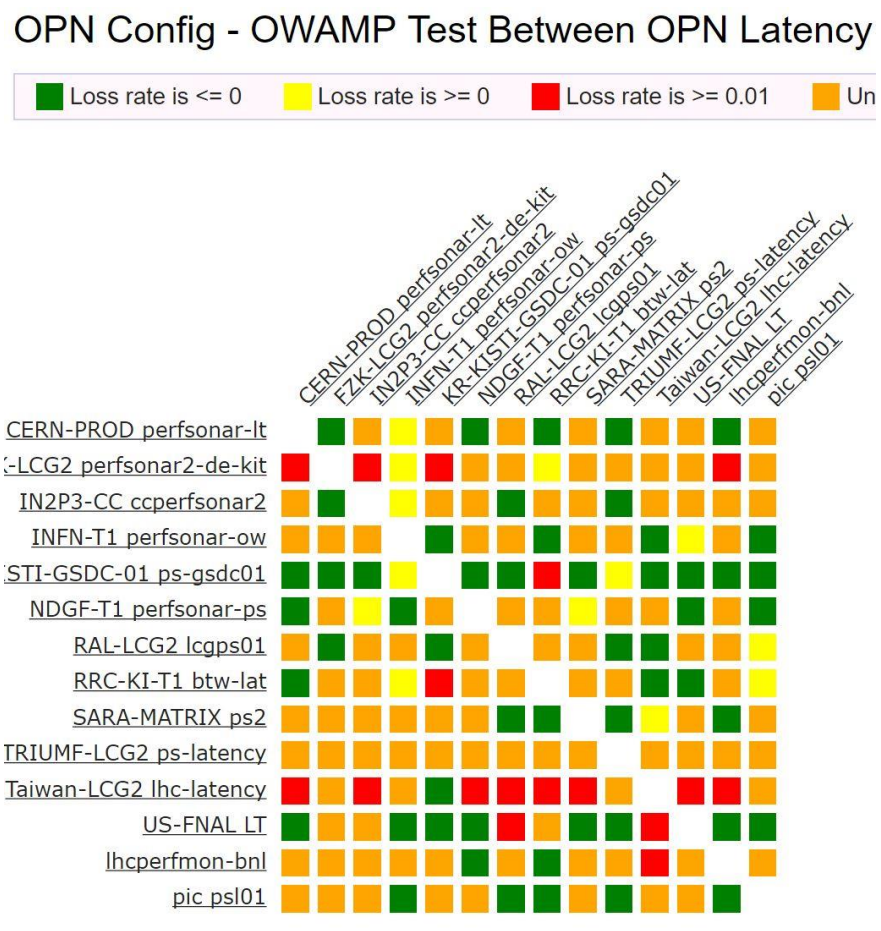

Again degradation of the SERVICE. **We should use the upcoming perfSONAR v4 release as an opportunity to get our installations fixed (and more resilient).**

LHCOPN MaDDash – 5 Mar 2018

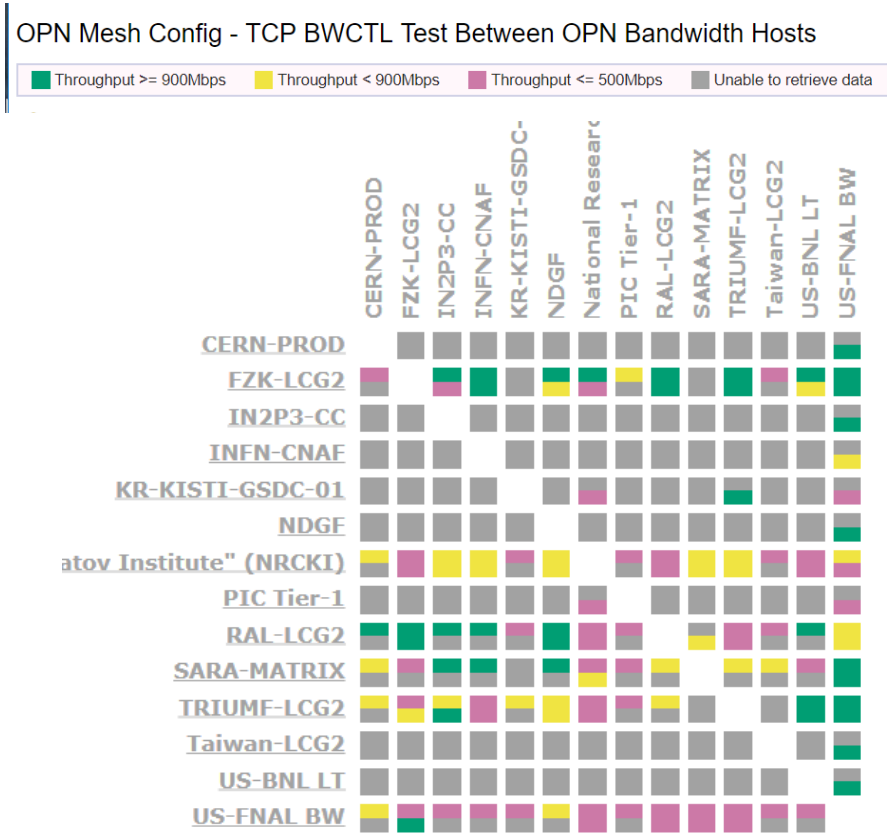

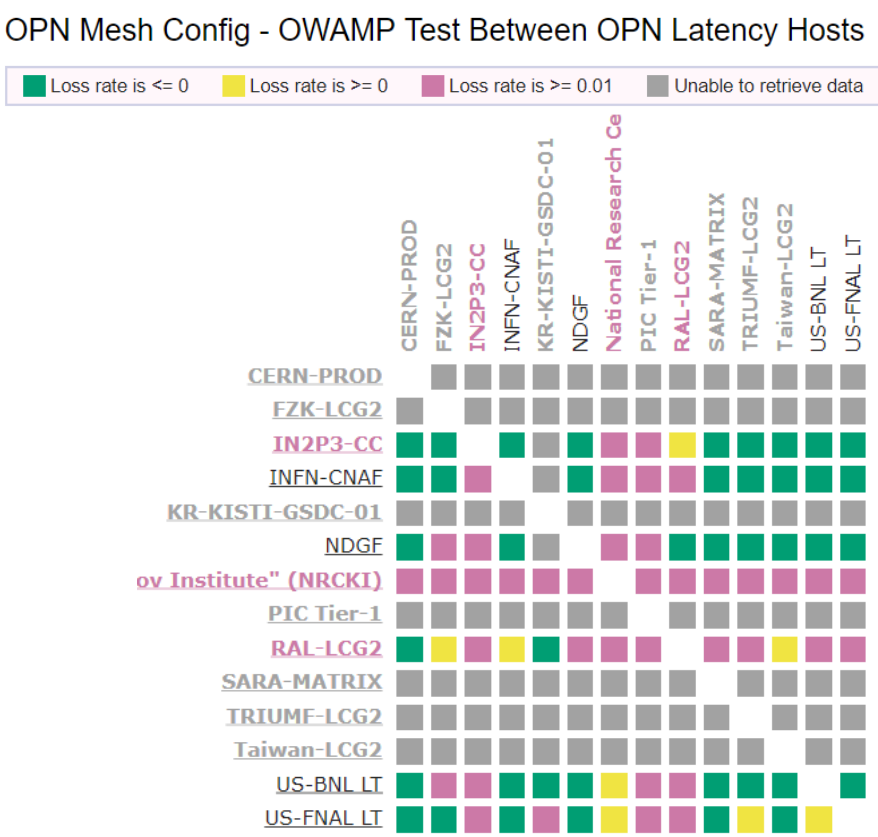

Less data since April 2017 A number of instances are not functioning in **originating measurements**

Info on MaDDash

- As you have noticed above, we are using a new (v2.0) version of MaDDash.
  - Names / labels are auto-scaled and vertical instead of slanted
  - We implemented “grey” as the default representing missing data
- There has been some confusion about where to go for data since we have two primary MaDDash instances
  - https://maddash.aglt2.org/maddash-webui This is the prototype instance and it **directly** gathers data from the perfSONAR toolkits
  - https://psmad.grid.iu.edu/maddash-webui This is the production instance and it gets its data from the OSG/WLCG central MA

perfSONAR v4.1 Update

- Next release of perfSONAR is 4.1 and should be available “soon”
  - **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 measuremnt)
  - ping alternative of owamp – routers/switches can participate in the tests
- Docker support

OSG/WLCG Net-Data Pipeline

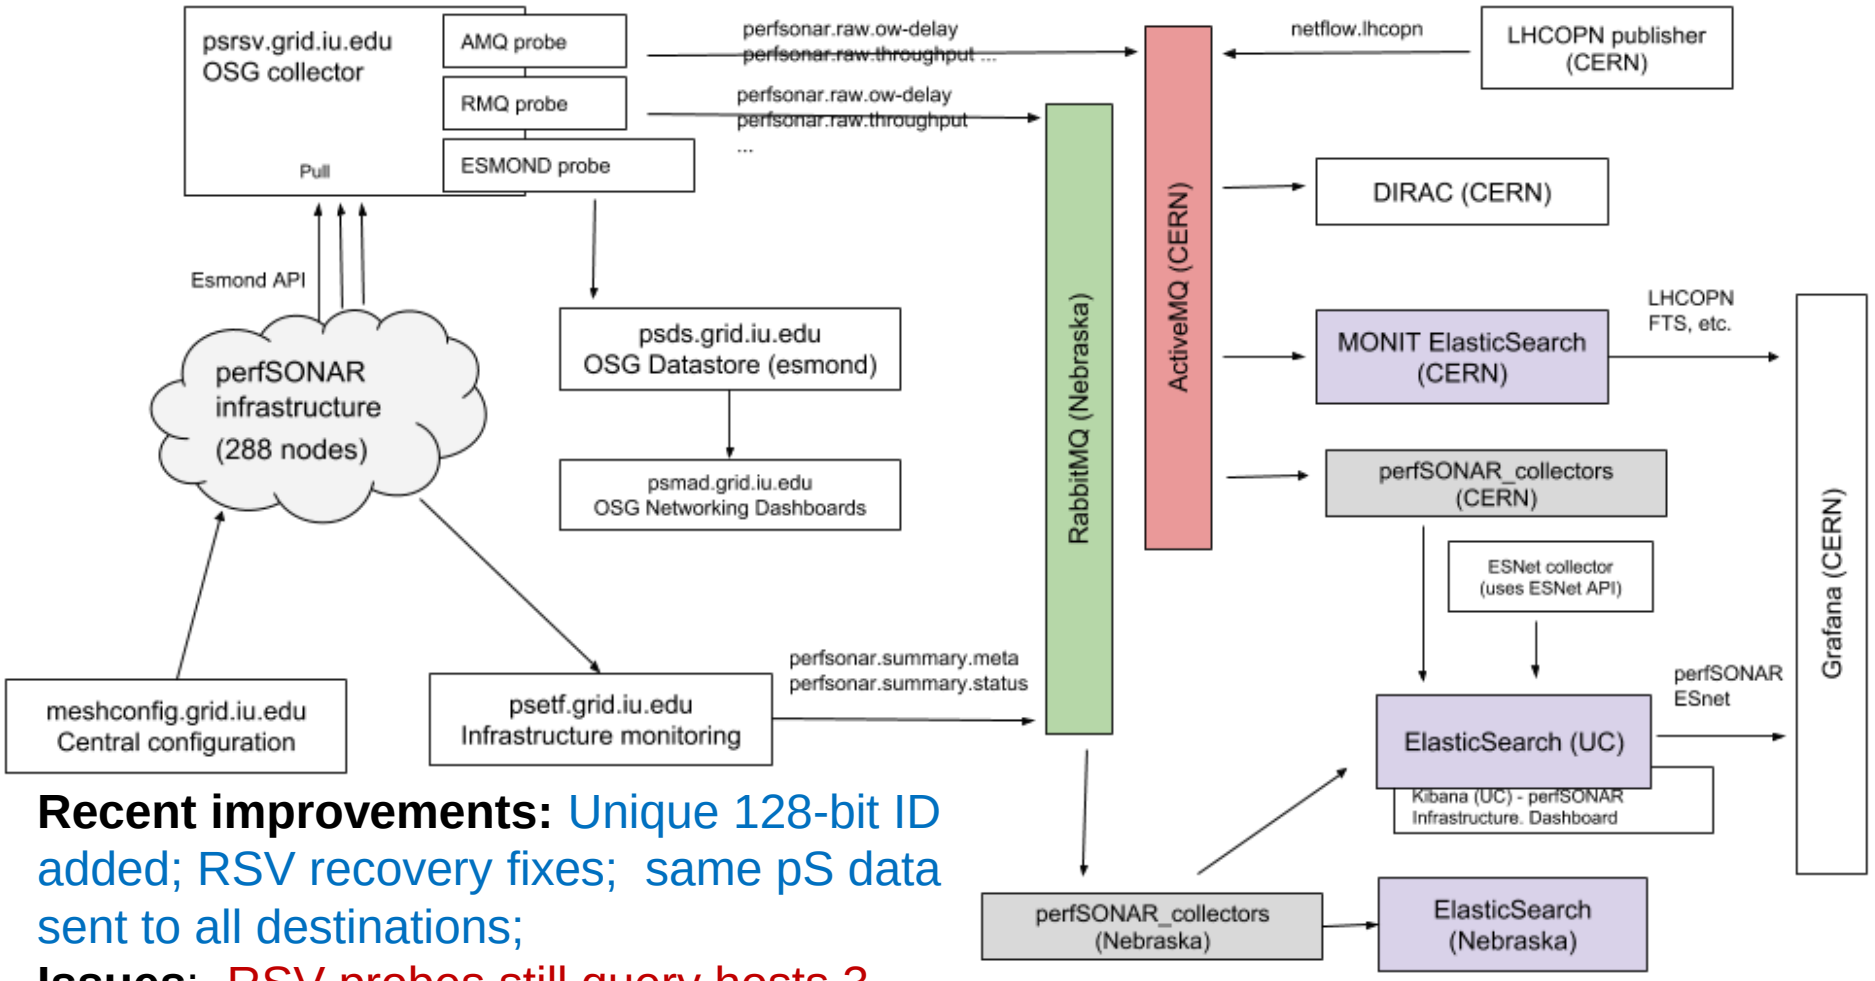

Recent improvements: Unique 128-bit ID added; RSV recovery fixes; same pS data sent to all destinations;

Issues: RSV probes still query hosts 3 times for the same data; two main MQ have unique content: need inter-bus connector;

OSG/WLCG Recent Work

- We have been working to redo the data pipeline so that the complete end-to-end path was under OSG control/mgmt.
  - Data still stored in OSG's central (ESMond) measurement archive
    - psmad (MaDDash uses this); short-term storage (1-2 months)
  - Data flows to two Elasticsearch instances: 1) U Chicago for near-realtime and two Nebraska for long-term analysis (plots)
  - Data also sent to FNAL for archiving on tape
- Lots of work on the perfSONAR RSV probes that gather the data
  - Probes had some bugs and configuration errors now fixed
  - Challenges on retrieving data backlogs when a system has been down or unreachable; probes refactored to grab data in time-interval chunks to not overload the toolkits (like asking for 30 days of latency)
  - Still problems/challenges with firewalls and overloaded end-systems
  - **Longer-term awaiting the ability to directly have toolkits send to message bus destinations**

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

TWiki DASHBOARDS CONTACT

Networking alert

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.

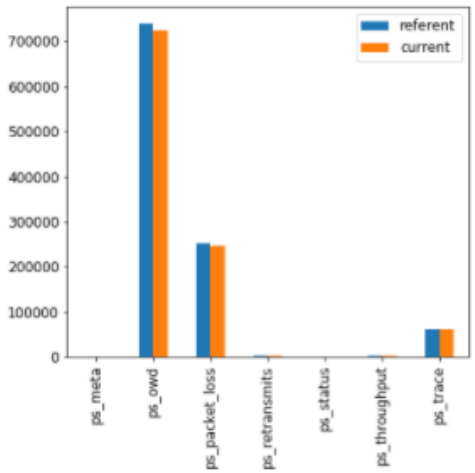

Network Data Gathering Stats - 30 Days

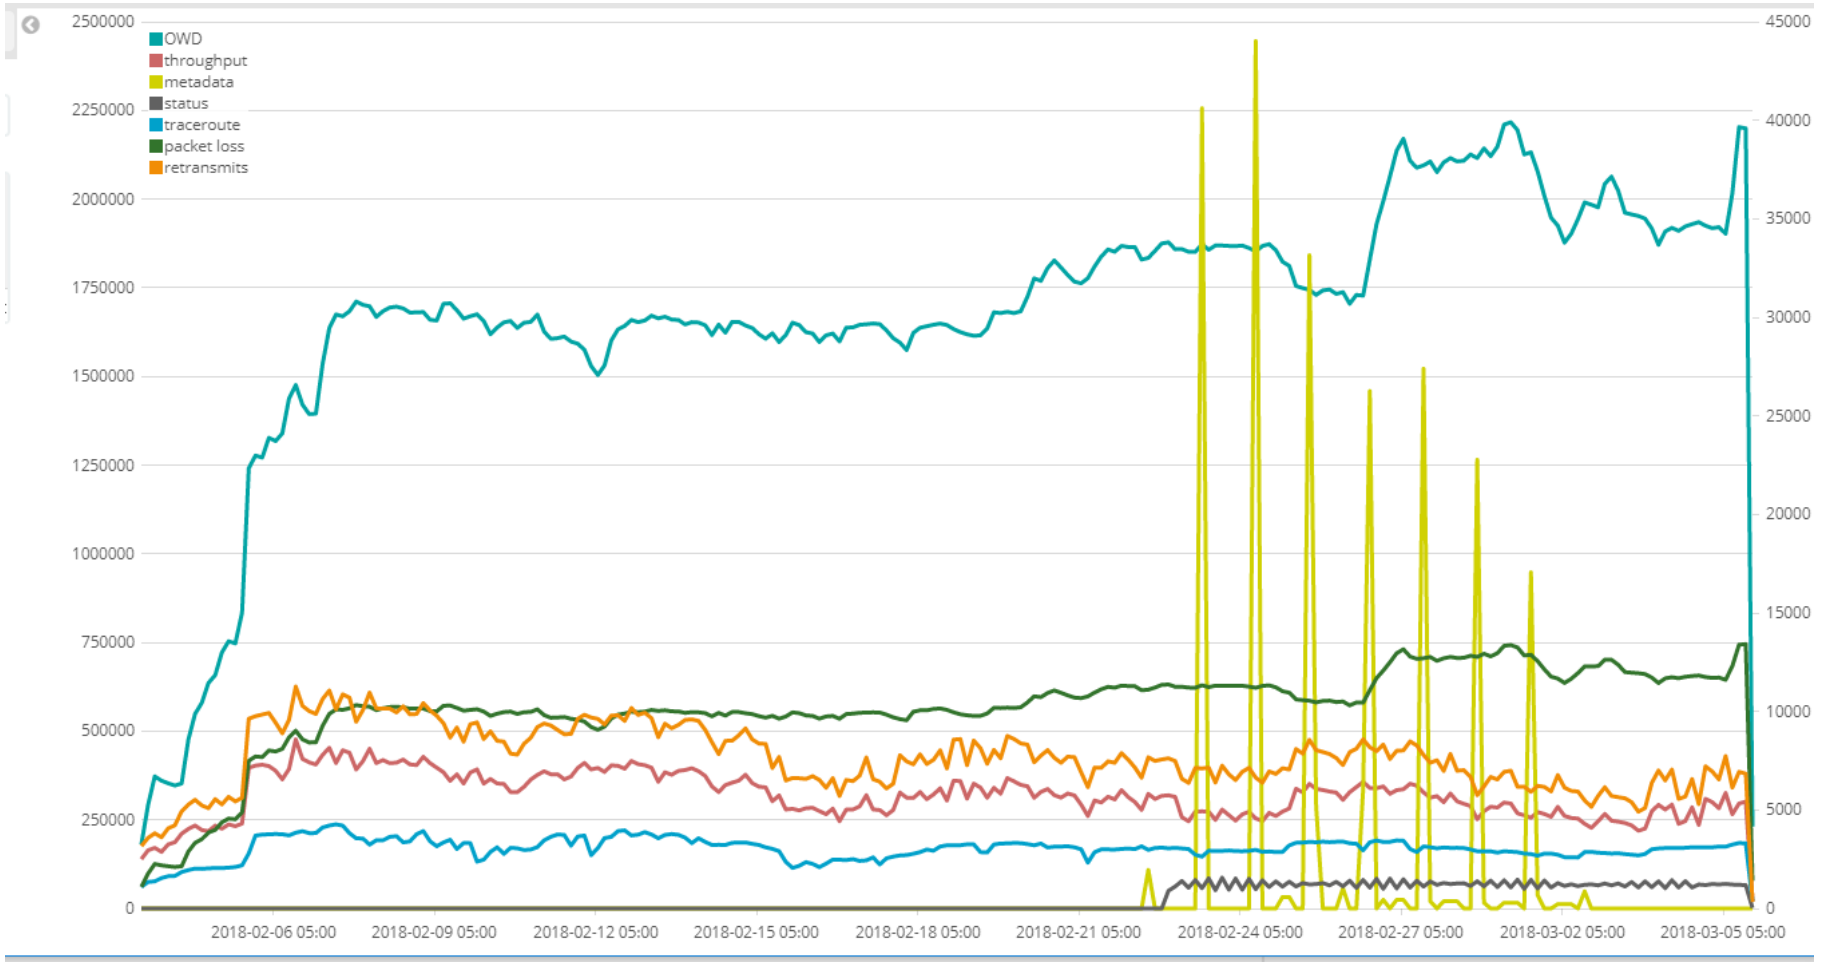

LHCOPN Total Traffic

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

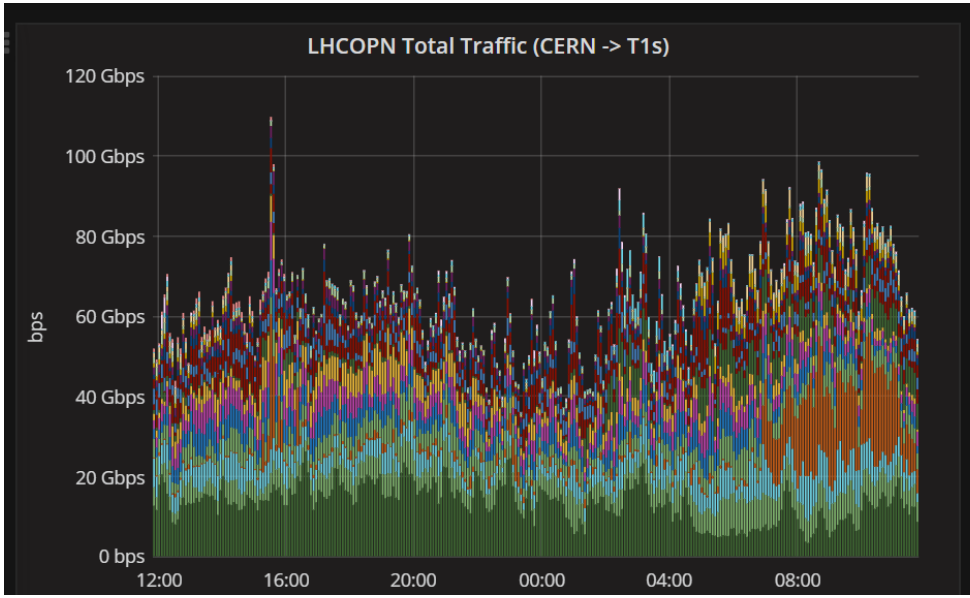

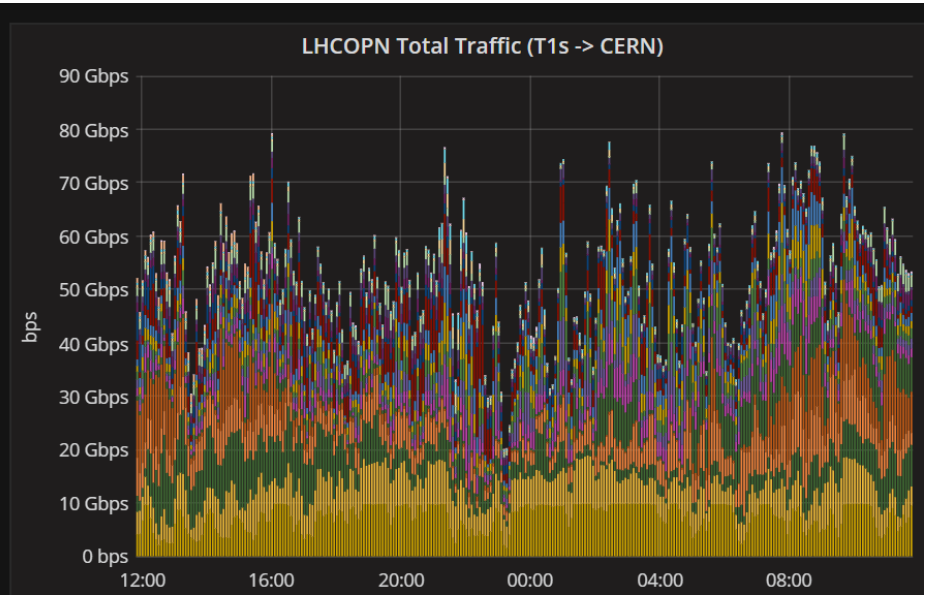

- US-T1-BNL primary Max: 26.2 Gbps Avg: 13.3 Gbps
- DE-KIT: secondary Max: 9.4 Gbps Avg: 5.2 Gbps
- RRC-KI-T1 LHCOPN Tier1 - Primary Max: 12.8 Gbps Avg: 4.9 Gbps
- US-FNAL-CMS primary Max: 33.8 Gbps Avg: 4.6 Gbps
- DE-KIT primary Max: 8.5 Gbps Avg: 4.3 Gbps
- UK-T1-RAL first Max: 9.9 Gbps Avg: 4.3 Gbps
- 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

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
DE-KIT: secondary	5.23 Gbps	9.41 Gbps	6.53 Gbps
	2.80 bps		4.03 Gbps
	3.78 bps		5.30 Gbps
	.48 bps		4.40 Gbps

LHCOPN Inbound Traffic Top 5

Metric	Avg	Max	Current
DE-KIT: secondary	7.36 Gbps	9.86 Gbps	8.37 Gbps
ES-PIC primary	5.64 Gbps	9.06 Gbps	3.69 Gbps
US-T1-BNL primary	5.19 G	15.15	8.94
RRC-KI-T1 LHCOPN Tier1 - Primary	4. G		
US-FNAL-CMS primary	3. G		

Karlsruhe Institute of Technology (DE-KIT)

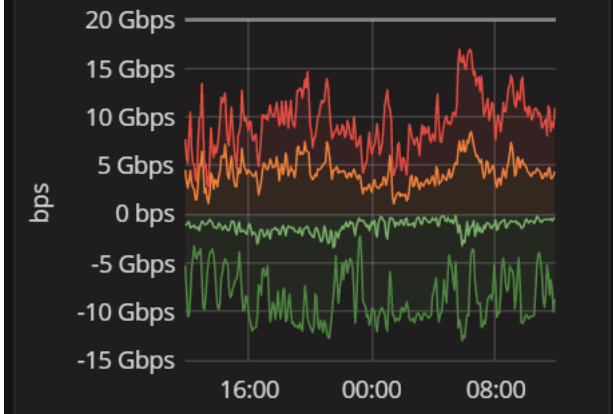

CERN to KIT primary Max: 8.48 Gbps Avg: 4.30 Gbps
CERN to KIT secondary Max: 9.41 Gbps Avg: 5.23 Gbps
KIT to CERN primary Max: 3.42 Gbps Avg: 1.26 Gbps
KIT to CERN secondary Max: 9.86 Gbps Avg: 7.36 Gbps

Fermi National Accelerator Laboratory (US-FNAL)

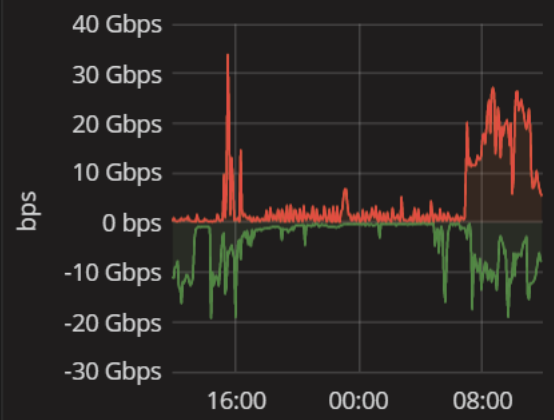

CERN to FNAL primary Max: 33.8 Gbps Avg: 4.6 Gbps
CERN to FNAL secondary Max: 677 bps Avg: 667 bps
FNAL to CERN primary Max: 19.2 Gbps Avg: 3.9 Gbps

pS Inter-regional Dashboard

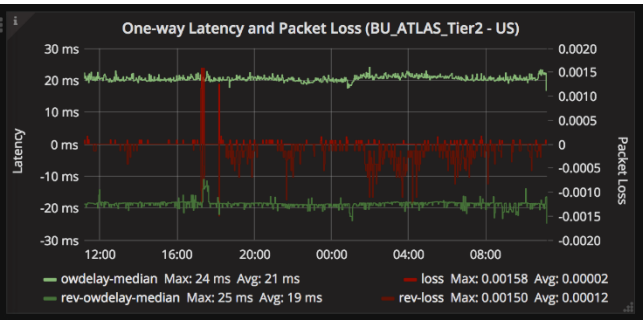

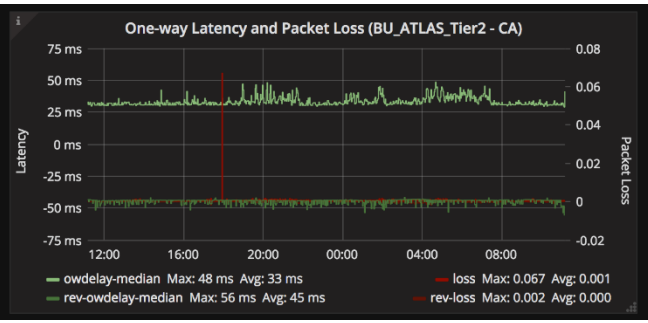

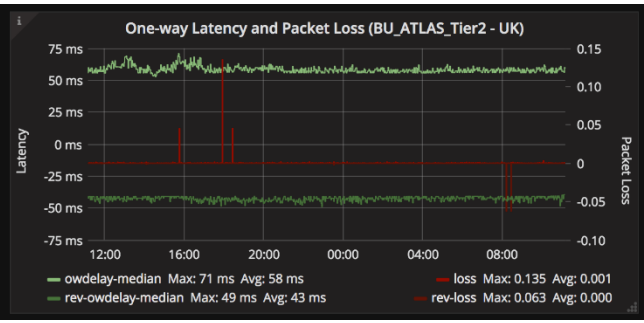

- Provides an overview of how the site performs in different regions
- Uses domain name wildcards to match regions

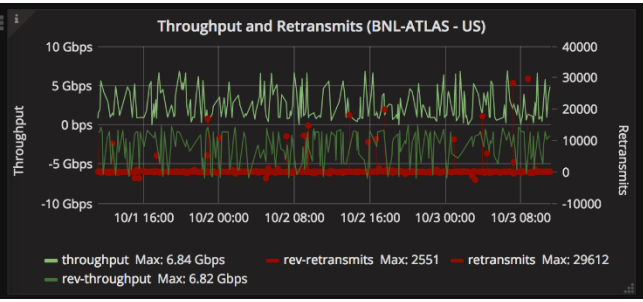

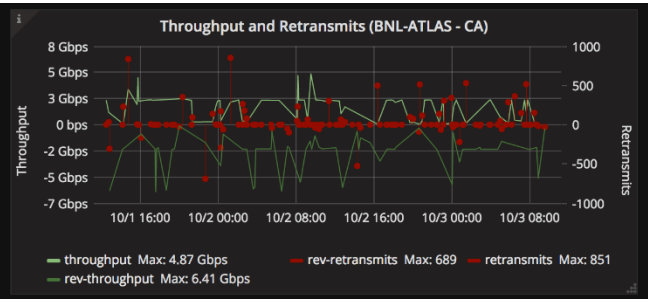

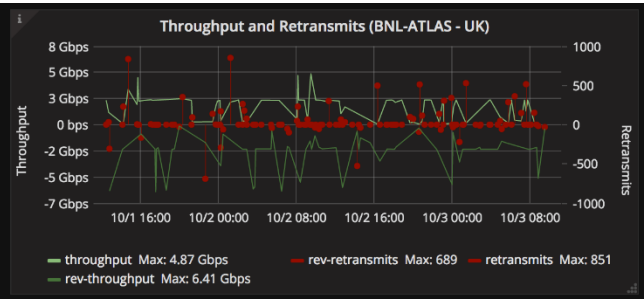

- ❄ Throughput and re-transmissions are done in the same way
- ❄ Granularity of testing is determined by meshes, which are experiment-based

Recruiting for Exploiting our Data

- With the current systems we have in place there are many opportunities to “extract value” from the data we collect
- We would like to encourage anyone with an interest in using the data or trying out new ways of exploiting the data to get involved.
- There a many possible areas to work in:
  - Network topology: visualization, analysis, user interfaces
  - Site analysis: identifying bottlenecks, misconfiguration, problems
  - Network issues: robust problem identification, localization, alarming
  - Notification systems: low-noise alerting, customization, user interfaces
- Let me know if you (or someone you know) is interested!

Enabling Alarming (Work-in-progress)

- OSG has a goal of alerting and alarming on network issues and LHCONE/LHCOPN may want to piggy-back on this effort.
- Milestone completed: technical design of a suitable analysis system based upon existing time-series technologies
- Current operating implementation gathers all perfSONAR data OSG sends to CERN and puts it in ElasticSearch.
- Jupyter instance regularly runs cron tasks to analyze data
  - Near-term goal: anyone can subscribe to simple alert-emails.
  - Not “production” yet; needs further interface tweaking to make it easy for users to use
  - Still need to determine where this service (or its equivalent) should be homed (ATLAS specific resources, WLCG?, OSG?)
  - **Notify Shawn if you want to test this**

REMINDER: WLCG Support Unit

- **Reminder**: We have a GGUS support unit (WLCG Network Throughput; https://wiki.egi.eu/wiki/GGUS:WLCG_Network_Throughput) used to report incidents (mailing list: wlcg-network-throughput at cern.ch)
- Experiments can report potential network performance incidents.
  - WLCG perfSONAR support investigates and confirms if this is network related issue.
  - Once confirmed, it will notify relevant sites and will try to assist in narrowing down the problem to particular link(s). Tracking of ongoing incidents will be via the WG page.
- **Most common issues are**: MTU, MTU+LB, routing (mainly remote sites), site equipment/design, firewall, workloads causing high network usage
- Sites observing a network performance problem should follow their standard procedure, i.e. report to their network team and if necessary escalate to their network provider.
  - If confirmed to be WAN related, WLCG perfSONAR support unit can assist in further debugging. For the policy issues, sites escalate to the WLCG operations coordination.
  - https://opensciencegrid.github.io/networking/network-troubleshooting/ https://twiki.cern.ch/twiki/bin/view/LCG/NetworkTransferMetrics#Network_Throughput_Support_Unit
- LHCOPN/LHCONE experts are very important in this coordinated activity.

Other Items

Update on SDN in ATLAS

- Ben Mack-Crane and Galen Mack-Crane provided us documentation on OVS options and tested non-disruptive production node migration paths.
  - *Testbed in-place with SL6.8 and CentOS 7.3 physical nodes and bonded 10G links*
  - *Results posted here: https://www.aglt2.org/wiki/bin/view/Main/Open_vSwitch/InstallOpenvSwitch*
- AGLT2 deployed Open vSwitch on ATLAS production storage (dCache) systems (http://openvswitch.org/)
  - IP addresses was moved to virtual interfaces with no downtime required (brief 30-40 sec outage)
  - No other changes; verified no performance impact
  - Traffic could now be shaped accurately using 'tc' or OpenFlow with little CPU cost
- The **advantage** is the our data sources/sinks become **visible** and **controllable** by OpenFlow controllers like OpenDaylight
  - **BENEFIT**: Traffic shaping can result in significantly improved use of the WAN for some paths
- Challenge now is to get other ATLAS sites to follow.
  - Interest from MWT2, UVic, KIT and others. Requires concurrence from both networking and storage responsibles at each site
  - Shaping capabilities via OVS with Ryu being testing in the OSiRIS project (Ezra Kissel / IU leading)

HEPiX Network Function Virtualization WG

- A new HEPiX Working Group on Network Function Virtualization (NFV) has been formed.
  - Was proposed and approved at the Fall 2017 HEPiX meeting: https://indico.cern.ch/event/637013/contributions/2739266/
- To join the mailing list go to this URL: https://listserv.in2p3.fr/cgi-bin/wa?SUBED1=hepix-nfv-wg
- First meeting https://indico.cern.ch/event/699548/ with notes available at: https://docs.google.com/document/d/1zpcnAmqmlg_wULIP5ZtqAtqYO4mSH74owsZwTUbOTVc/edit
- Our next meeting will be Wednesday, March 14 https://indico.cern.ch/event/705126/
  - Hope you can join if you are interested!

Revisit our LHCOPN/LHCONE Toolkits?

- We have had lots of issues with the perfSONAR toolkits in LHCOPN/LHCONE since we initially deployed years ago
  - Many were not working for a long time or not in useful locations
  - Finding the right person who could fix/configure took a long time
- We should review the current list for usefulness
- **GOAL:** Find ~20 perfSONAR toolkit locations appropriate for LHCOPN and LHCONE (separately)
  - Need to identify the right contact for each; someone who can install/configure/repair if there are issues.
- (Re)install using CentOS7 (and 4.1 when available)
  - Systems should have at least 8 GB of RAM and 4 cores
  - Two NICs required to run both latency and bandwidth on the same system

Example: LHCONE perfSONAR BW Toolkits

- http://atlas-npt2.bu.edu/toolkit/
- http://ps01-nl.geant.net/toolkit (Dead!)
- http://ps-bandwidth.scinet.utoronto.ca/toolkit
- http://perfsonar.ultralight.org/toolkit (Dead!)
- http://lhcone-wash-opt1.es.net/toolkit
- http://lhcone-newy-opt1.es.net/toolkit (Dead!)
- http://perfsonar-de-kit.gridka.de/toolkit
- http://ccperfsonar1.in2p3.fr/toolkit
- http://lhcone.test.wix.internet2.edu/toolkit
- http://lhcone.test.manlan.internet2.edu/toolkit
- http://lapp-ps01.in2p3.fr/toolkit
- http://ps-bandwidth.lhcmon.triumf.ca/toolkit

Nothing in Chicago
Nothing in Asia
Only 3 in Europe

Reminder: About perfSONAR Toolkits

- WLCG is requiring all Tier-2 and above to install two perfSONAR instances: **bandwidth** and **latency**
  - As of version 4.0, this can be done with one server having 2 NICs
- For Tier-n sites, these instances should be connected on the same network location as the sites storage
  - We are trying to measure the end-to-end data network path!
- Our recommendation is the NIC for bandwidth measurement would ideally be the same as the storage
  - Latency/packet-loss NIC can be 1G to save cost
- To make the whole infrastructure more valuable, we would like as many perfSONAR toolkits along the R&E paths as possible!
- **Servers running perfSONAR toolkit should NOT be running any other applications or services!!**

Summary and Action Items

- **We need to plan for a campaign to clear up remaining LHCONE/LHCOPN problems and get v4.1 in place**
  - **Need more instances in Asia in the regional R&E networks!!**
  - **New alarming service should help focus on network problems**
  - **Everyone MUST open port 443 for things to work**
- **New monitoring, management and data analytics capabilities are coming soon.**
  - Need to setup LHCONE/LHCOPN to use these capabilities.
  - Need R&E network information from more network providers
- **As we fix known issues and get perfSONAR instances reliably operating, we free up time to pursue possible issues in the network itself, rather than the framework**
  - **As noted last-time: Looking for volunteers to get involved in network analytics, especially on LHCONE/LHCOPN issues**

Discussion/Questions/Comments?

References

- Network Documentation

  https://opensciencegrid.github.io/networking/

- Modular Dashboard Prototype
  - http://maddash.aglt2.org/maddash-webui
- OSG Production instances for OMD, MaDDash and Datastore
  - https://psetf.grid.iu.edu/etf/check_mk/
  - http://psmad.grid.iu.edu/maddash-webui/
  - http://psds.grid.iu.edu/esmond/perfsonar/archive/?format=json
- Mesh-config in OSG https://meshconfig.grid.iu.edu/