


WLCG Network Throughput WG

Shawn McKee, Marian Babik for the Working Group

HEPiX Budapest
25 April 2017



Working Group

- WLCG Network Throughput WG formed in the fall of 2014 within the scope of WLCG operations with the following objectives:
 - Ensure sites and experiments can better understand and fix networking issues
 - Measure end-to-end network performance and use the measurements to single out on complex data transfer issues
 - Improve overall transfer efficiency and help us determine the current status of our networks
- Core activities include:
 - Deployment and operations of perfSONAR infrastructure
 - Gain visibility into how our networks operate and perform
 - Network analytics
 - Improve our ability to fully utilize the existing network capacity
 - Network performance incidents response team
 - Provide support to help debug and fix network performance issues

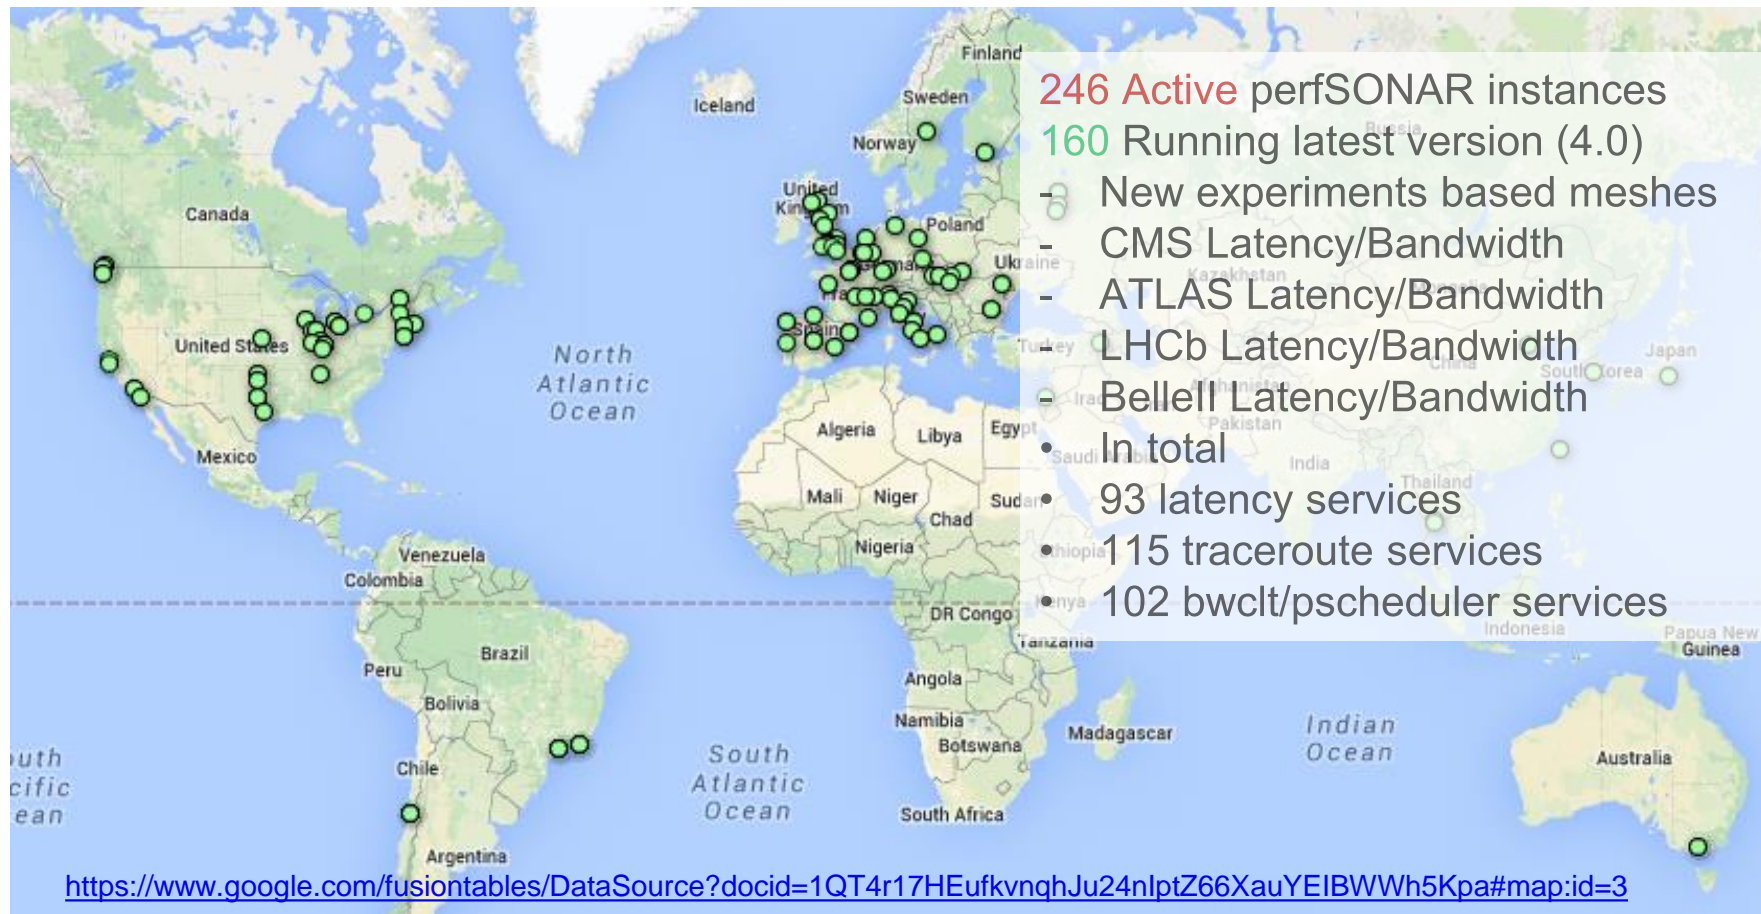
<https://twiki.cern.ch/twiki/bin/view/LCG/NetworkTransferMetrics>

News

- perfSONAR 4.0 released on April 17
 - Major release bringing many new features
 - Default installations should auto-update
 - Some LHCONE/LHCOPN instances may require manual intervention.
- OSG Network Measurement Service updates coming along with perfSONAR 4.0
 - New Mesh-configuration Agent (MCA)
 - Now part of perfSONAR 4.0; available as new central component
 - Updated WLCG perfSONAR monitoring based on ETF
 - New analytics capabilities using ELK + Jupyter and new data available via messaging/streaming API

Current perfSONAR Deployment

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



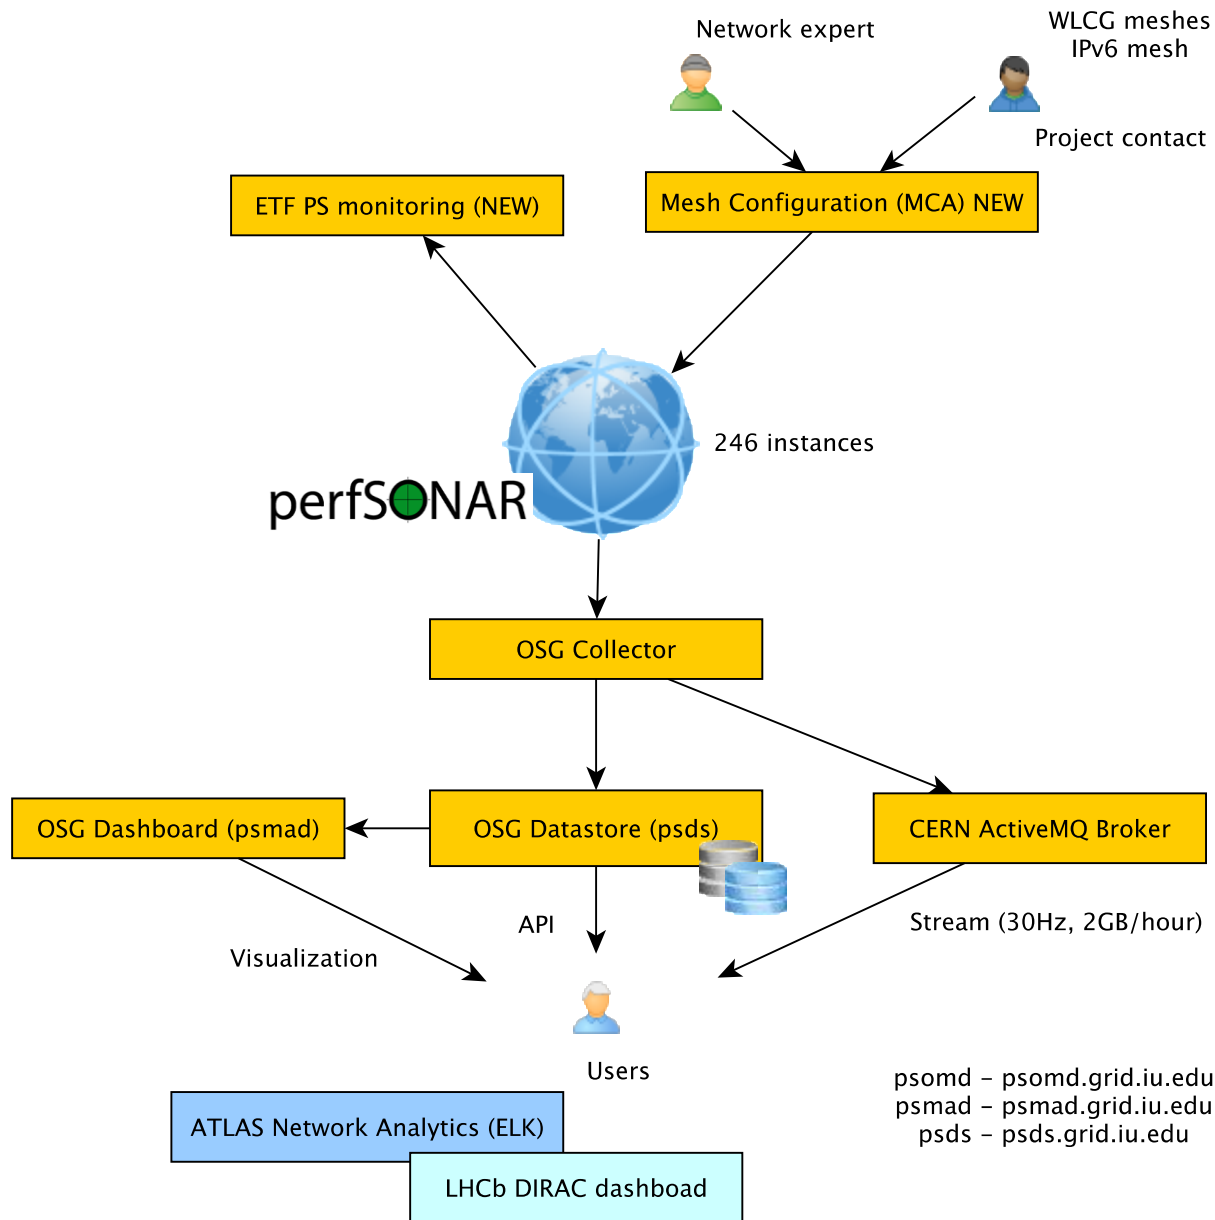
- Initial deployment coordinated by WLCG perfSONAR TF
- Commissioning of the network followed by WLCG Network and Transfer Metrics WG

perfSONAR 4.0

- Focus is on “Control and Stability”
 - New standalone mesh-config (MCA)
 - New test scheduler (pScheduler) replacing BWCTL:
 - Shared by all tests and aware of the resources each uses
 - Containing finer grained controls about who can run tests and what tests they are allowed to run.
 - Increased visibility and control as to when tests will be run
- CentOS 7 and Debian 8 support
 - perfSONAR 4.1 plan drops SL6 support; **requires reinstall for SL6 hosts**
- New Endpoint selection capabilities
 - Better metadata, topology information will be available
- New Dashboard (MaDDash w/MadAlert) part of this release
- New minimum HW requirements
 - Dual-core (2GHz+) and 4GB RAM
 - Heavily used latency instances would benefit from more memory
 - But new deployment models planned for 4.1: small-factor, VMs, docker
- Toolkit now exposes info via JSON, used in our monitoring:
 - <http://psum06.aglt2.org/toolkit/?format=json>

Network Measurement Platform

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.



MCA

- Gathers and organizes information on hosts from a combination of sources
 - Imports from perfSONAR global lookup service entries; we can discover perfSONARs deployed anywhere if they join the right community
 - Able to gather information from grid registration databases like GOCDB / OIM
- Auto-completion / entry of values
- Context dependent user interface (see Testspecs)
- Can be easily installed outside of OSG
- Provides a RESTful interface to allow easy monitoring and software-controlled config
 - <https://meshconfig-itb.grid.iu.edu/apidoc/>
- Supports filtering and dynamic host groups
 - Can now build dynamic meshes, e.g., all CentOS6 hosts who are members of the ATLAS community
- Becomes part of perfSONAR software stack – major upstream contribution from OSG/WLCG

MCA

Config-view of **Mesh Config Administrator** shown below. We now have the ability to do much finer-grained control and can setup as many mesh-config 'admins' as needed.

New production instance being activated in OSG in the next week or two.

The screenshot displays the Mesh Config Administrator web interface. The left sidebar contains navigation options: Configs, Hosts, Host Groups, and Testspecs. The main content area shows a list of 18 registered mesh configs, including 'perfSONAR Testbed', 'test', 'Soich's Test Config', 'Soichi's Dynamic HostGroup Test', 'USATLAS Mesh Config', 'USCMS Mesh Config', 'UK Meshconfig', 'LHCONE Mesh Config', 'WLCG LHCb Latency Mesh', 'OPN Config', 'WLCG CMS Latency Mesh', 'WLCG LHCb Bandwidth Mesh', and 'WLCG CMS Bandwidth Mesh'. The 'OPN Config' is selected, and its details are shown in the right pane. The 'MeshConfig URL' is 'http://meshconfig-itb.grid.iu.edu/pub/config/opn-all'. The 'Name' is 'OPN Config'. The 'Admins' section lists three users: Soichi Hayashi, Shawn McKee, and Marian Babik. The 'Tests' section shows a test named 'TCP BWCTL Test Between OPN Bandwidth Hosts' with a service type of 'Bandwidth (bwctl)' and a mesh type of 'Mesh'. A world map is displayed below the test details, showing several red location pins across various continents.

ETF PS

- New monitoring based on WLCG [ETF](#)
 - We have an etf_ps version in docker: mbabik/etf:1.2.8p20
 - See https://etf.aglt2.org/etf/check_mk/ for a running example
- Goal is to replace the OSG psomd.grid.iu.edu and perfsonar-itb.grid.iu.edu check_mk instances with this

The screenshot displays the Check MK web interface. The top navigation bar includes the 'Check MK' logo, version 'Raw 1.2.8p18', and the title 'Services of Host psum05.aglt2.org'. On the right, it shows '15 rows /DC=ch/DC=cern/OU=Organic Units/OU=...'.

On the left, there is a sidebar with a 'Tactical Overview' section showing 287 hosts and 14 problems, and 4548 services with 2066 problems. Below this is a 'Quicksearch' field containing 'psum05.aglt2.org' and a 'Views' menu with options like Overview, Hosts, Host Groups, Services, and Service Groups.

The main content area shows a table of monitoring services for 'psum05.aglt2.org'. The table has columns for State, Service, Icons, Status detail, Age, Checked, and Perf-O-Meter. Several services are in a 'WARN' state, including 'OSG datastore freshness: owamp' and 'perfSONAR esmond freshness: owamp'.

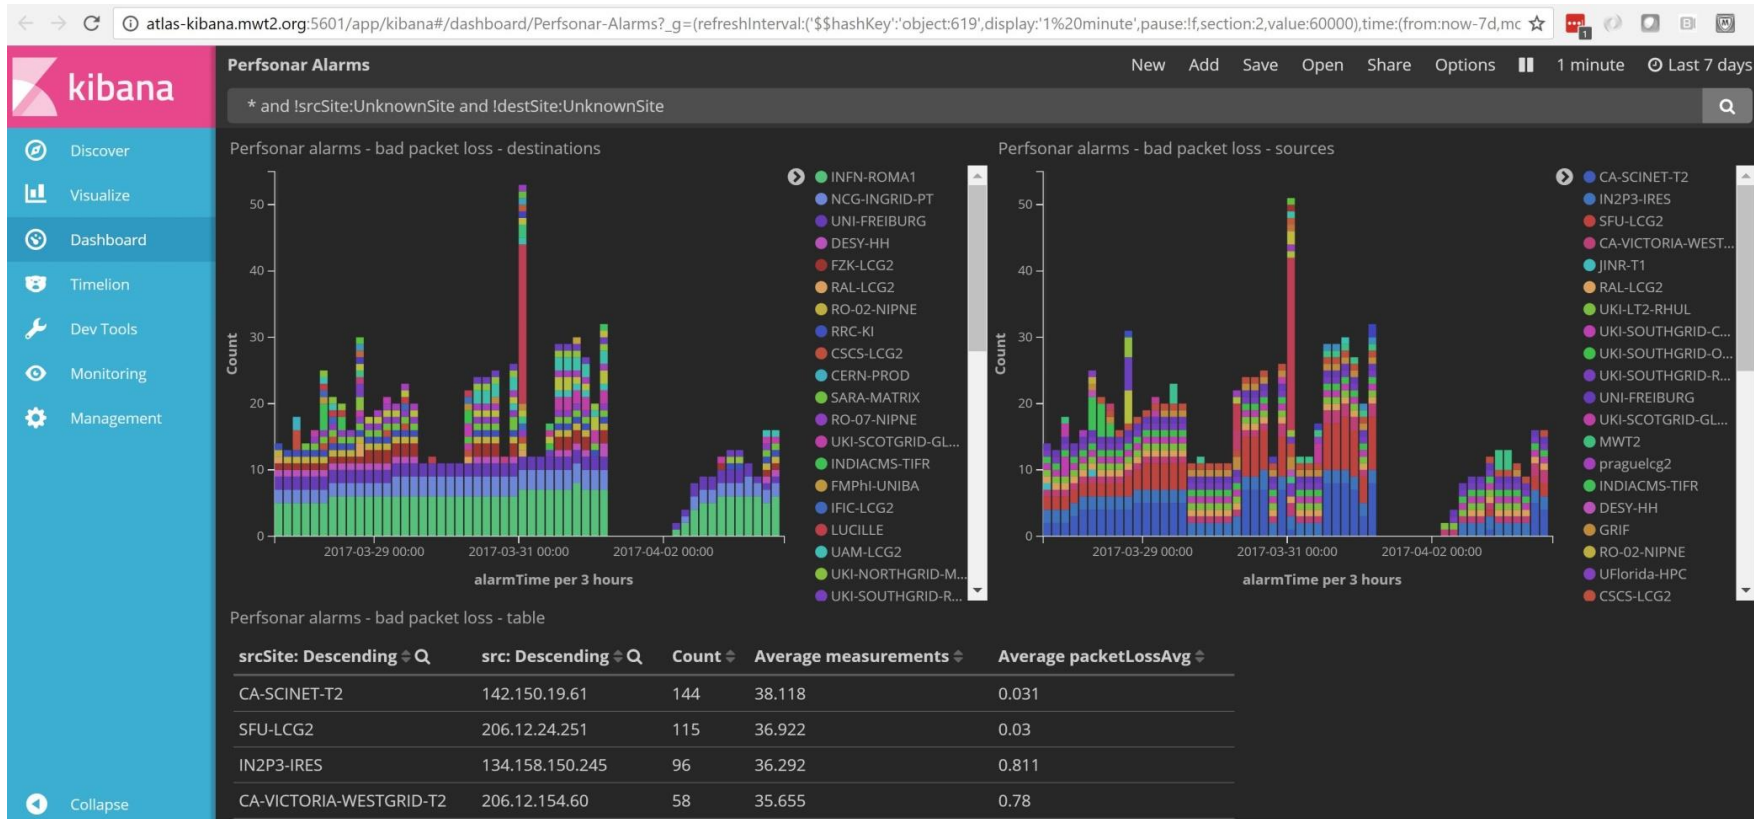
State	Service	Icons	Status detail	Age	Checked	Perf-O-Meter
WARN	OSG datastore freshness: owamp		WARNING - 9 stale hosts for event type: histogram-owdelay	4 hrs	42 min	
OK	perfSONAR configuration: contacts		OK - Contact and organization found	38 hrs	2 sec	
OK	perfSONAR configuration: location		OK - Location: -83.736719/42.276845	38 hrs	2 sec	
OK	perfSONAR configuration: meshes		OK - Meshconfig auto URL configured	38 hrs	2 sec	
WARN	perfSONAR esmond freshness: owamp		WARNING - 1 stale hosts for event type: histogram-owdelay	2017-04-06 12:27:07	36 min	
WARN	perfSONAR esmond freshness: owamp rev		WARNING - 1 stale hosts for event type: histogram-owdelay	2017-04-06 12:31:03	32 min	
OK	perfSONAR hardware check		OK - CPU:2/8cores/2659.935Mhz RAM:16GB NIC:1Gbps/9000MTU/IPv4 only	38 hrs	2 sec	
OK	perfSONAR json summary		OK - Toolkit metadata successfully retrieved	38 hrs	14 sec	
OK	perfSONAR services: http/https		OK - Toolkit homepage reachable	38 hrs	2 sec	
OK	perfSONAR services: ndt/npad disabled		OK - NDT/NPAD disabled		2 sec	
OK	perfSONAR services: ntp		OK - NTP synchronized	38 hrs	2 sec	
OK	perfSONAR services: owamp		TCP OK - 0.133 second response time on psum05.aglt2.org port 861	2017-04-11 04:04:08	6 min	
OK	perfSONAR services: pscheduler		OK - pScheduler stats retrieved	2017-04-06 12:11:03	58 min	
OK	perfSONAR services: regular testing/pscheduler		OK - pScheduler is enabled and running	38 hrs	2 sec	
OK	perfSONAR services: versions		OK - Toolkit version found 4.0.0.1-1.el6	38 hrs	2 sec	

Network Analytics

- Ilija Vukotic/U Chicago has been leading an effort within ATLAS 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>
- Uses CERN IT streaming/messaging interface, which is publicly available:
 - Contains all perfSONAR measurements gathered by Network Measurement Platform including latencies, packet-loss, TTLs, throughputs, TCP retransmits, traceroutes/tracepaths
 - Now also includes some derived metrics, such as [simulated link utilization](#) for all LHCONE links
 - As well as LHCOPN and ESNet traffic/netflows

Network Analytics Example

- Example Kibana dashboard looking at sites with network paths having $>2\%$ packet-loss over a 1 hour period (site is dest(left) or src(right))



Alerts/Notifications

- OSG is targeting alerting users when sites they are interested in are experiencing problems
 - We must be careful; don't spam recipients for problems not relevant to them
 - Solution is to allow self-subscription for sites of interest
 - Example email:

Dear Shawn McKee,

This mail is to let you know that there are significant changes in the number of paths with large packet-loss detected by perfSONAR for sites you requested alerting about.

The site CA-SCINET-T2 (142.150.19.61)'s network paths have improved, the count of src-destination paths with packet-loss went from 7 to 0.

The site UKI-SOUTHGRID-OX-HEP (163.1.5.210)'s network paths have improved, the count of src-destination paths with packet-loss went from 21 to 1.

UKI-SOUTHGRID-OX-HEP (163.1.5.210) ---> INFN-ROMA1 (141.108.35.18)

To get more information about this alert message and its interpretation, please visit:

<http://twiki.opensciencegrid.org/bin/view/Documentation/NetworkingInOSG/PacketLossAlert>

If you suspect a network problem and wish to follow up on it please email the appropriate support list:

For OSG sites: goc@opensciencegrid.org using Subject: Possible network issue

For WLCG sites: wlcg-network-throughput@cern.ch using Subject: Possible network issue

Please include this alert email to help expedite your request for network debugging support.

To change your alerts preferences please you the following link:

https://docs.google.com/forms/d/e/1FAIpQLSfwwtAvMrqp4Ot_LYfmNu75_v33dtAxiXg7ZvVdn1X5v7TEgg/viewform?edit2=2_ABaOnudW1aM_lw2Ub3gE1cMSCZJyqjIOaAvdg0pb68_3ViPF77dyEy62OK_BQ

Best regards,

ATLAS Networking Alert Service

Improving Transfer Efficiency

- One of the main goals of the WG to identify and help fix network issues in our infrastructure
- Network analytics, alerts/notifications are the primary drivers in this activity
 - Detection and analysis of the issues has significantly improved
- In addition, WG has created a support unit to help sites and experiments with network-related 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
- Two recent example cases: **ASGC** and **TRIUMF/SFU**

ASGC case

- A ticket was opened in March 2016
https://ggus.eu/index.php?mode=ticket_info&ticket_id=119820

“Hello,

ATLAS recently suffers from various network related problems with access to resources in ASGC. Problems reported via proper channels are solved but soon they reappear. Examples of recent issues: stratum-1 at ASGC (GGUS ticket 119557) ; NDGF - ASGC transfer problems (GGUS ticket 119276). WLCG weekly meeting suggest to open this general ticket and assign it to WLCG Network Monitoring unit, which could use perfsonar infrastructure to better understand these problems.

Jiri Chudoba as ATLAS Computing Run Coordinator”

- ASGC connected by dedicated 10Gbit via StarLight and then via AMS to CERN; perfSONAR coverage in Asia quite limited
- Testing to ASGC from many places showed no more than **500Mbit**; StarLight to ASGC was not more than **200Mbit**, so moved focus to ASGC
- Asked ASGC for network map and asked them to re-connect their bwctl node directly to border router while keeping latency node on Cisco N7K
 - Testing from StarLight directly to ASGC border router revealed Cisco N7K to be the problem...buffers too small for WAN transfers (72MB shared with 12 ASICs)
 - **Solution**: Move perfSONAR and DTNs off Cisco and onto border router

ASGC Result

- **Investigation took less than 30 days** (we posted recommendation to ASGC in May 2016 and there were some initial delays before we got going)
 - **perfSONAR** was critical to achieving this
 - We managed to gain factor 10, ASGC is at ~ **4.4Gbps** now
 - Can we get better ?
- Major achievement for **WLCG network throughput**, shows that at least in Asia (but we have other cases showing that also elsewhere) there are benefits to be gained from analyzing our metrics to locate problems.
- With respect to **LHCOPN/LHCONE** we're clearly missing guidance on high-latency (distant) sites network architectures and sizing the router buffers
 - Coordinating with ESnet/Fasterdata on getting this improved

TRIUMF/CA case

- The beta version of the alerting system shown earlier has been showing that all the Canadian WLCG sites are seeing periods of large packet loss.
- Rolf Seuster/UVic, the WLCG network contact for the Canadian cloud, was interested in trying to use our network metrics and associated co-located Jupyter instance to try to identify the location of the problem
 - He was given access to the system and pointed at a few existing examples of analyzing network metrics
 - Results from just a few days ago follows

Analyzing SFU to DESY Zeuthen

Using the network metrics available in ElasticSearch, Rolf created some Python code in Jupyter to analyze the data

- To make the problem manageable he focused on two specific sites: SFU and DESY_ZN
- For a **1h time window** when the packet loss was bad, he got all measurements for packet losses correlated with all traceroutes originating from SFU.
- Additionally, for the same time window he got all measurements going into DESY Zeuthen.
- Then just **count how many measurements showed packet-loss (bad) or not (good)** for each router between the sites
- This shows how often each router is part of a good or bad end-to-end packet-loss test.
- Only 1 router had no good measurements: **62.40.126.145** a.k.a. **canarie-lhcone-gw.mx1.lon.geant.net** has to be the faulty router.

Router(IP)	Good(CA)	Good(ZN)	Bad(CA)	Bad(ZN)
206.12.24.225	61	0	44	2
206.12.9.96	43	0	24	2
205.189.32.175	25	0	44	4
205.189.32.177	25	0	44	4
205.189.32.181	27	0	44	4
62.40.126.145	0	0	24	3
62.40.98.81	0	21	20	5
62.40.126.168	0	22	6	3
62.40.126.169	0	22	6	3
188.1.153.14	0	22	4	3
141.34.249.22	0	48	2	5
141.34.249.218	0	70	2	5
141.34.249.209	0	70	2	5
141.34.200.29	0	70	2	5

Plans

- perfSONAR deployment – targeting 4.1 release
 - New deployment models for OSG/WLCG
 - Upstream contributions to perfSONAR
 - Direct publishing to ActiveMQ/Kafka
 - SNMP pscheduler plugin to add site ingress/egress traffic from routers to perfSONAR MA
- OSG Network Measurement Platform
 - The goal is to provide an end-user friendly interface to our measurement and analytics data
 - Focus on supporting users in identifying network problems and bottlenecks through documentation and self-subscribed targeted alerting
- Analytics
 - Improve **Kibana** dashboards
 - Add infrastructure status dashboard
 - Test **Grafana** for better time-series visualization
 - Adding netflows to stream/messaging API
 - ESNet, GEANT, Internet2 and Canarie

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
- Work on new applications is underway
 - Notifications/alerting
 - Predictive capabilities
 - Current utilization and capacity planning
 - Evaluating network performance of commercial clouds
- We (OSG and WLCG) welcome feedback on how to further improve

Questions or Comments?

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>
- Measurement Archive (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>
 - 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>

Backup slides

Latency and packet loss matters




Performance is poor when RTT exceeds ~10 ms

Performance is good when RTT is < ~10 ms

Source Campus

R&E Backbone

Destination Campus



0.0046% loss (1 out of 22k packets) on 10G link

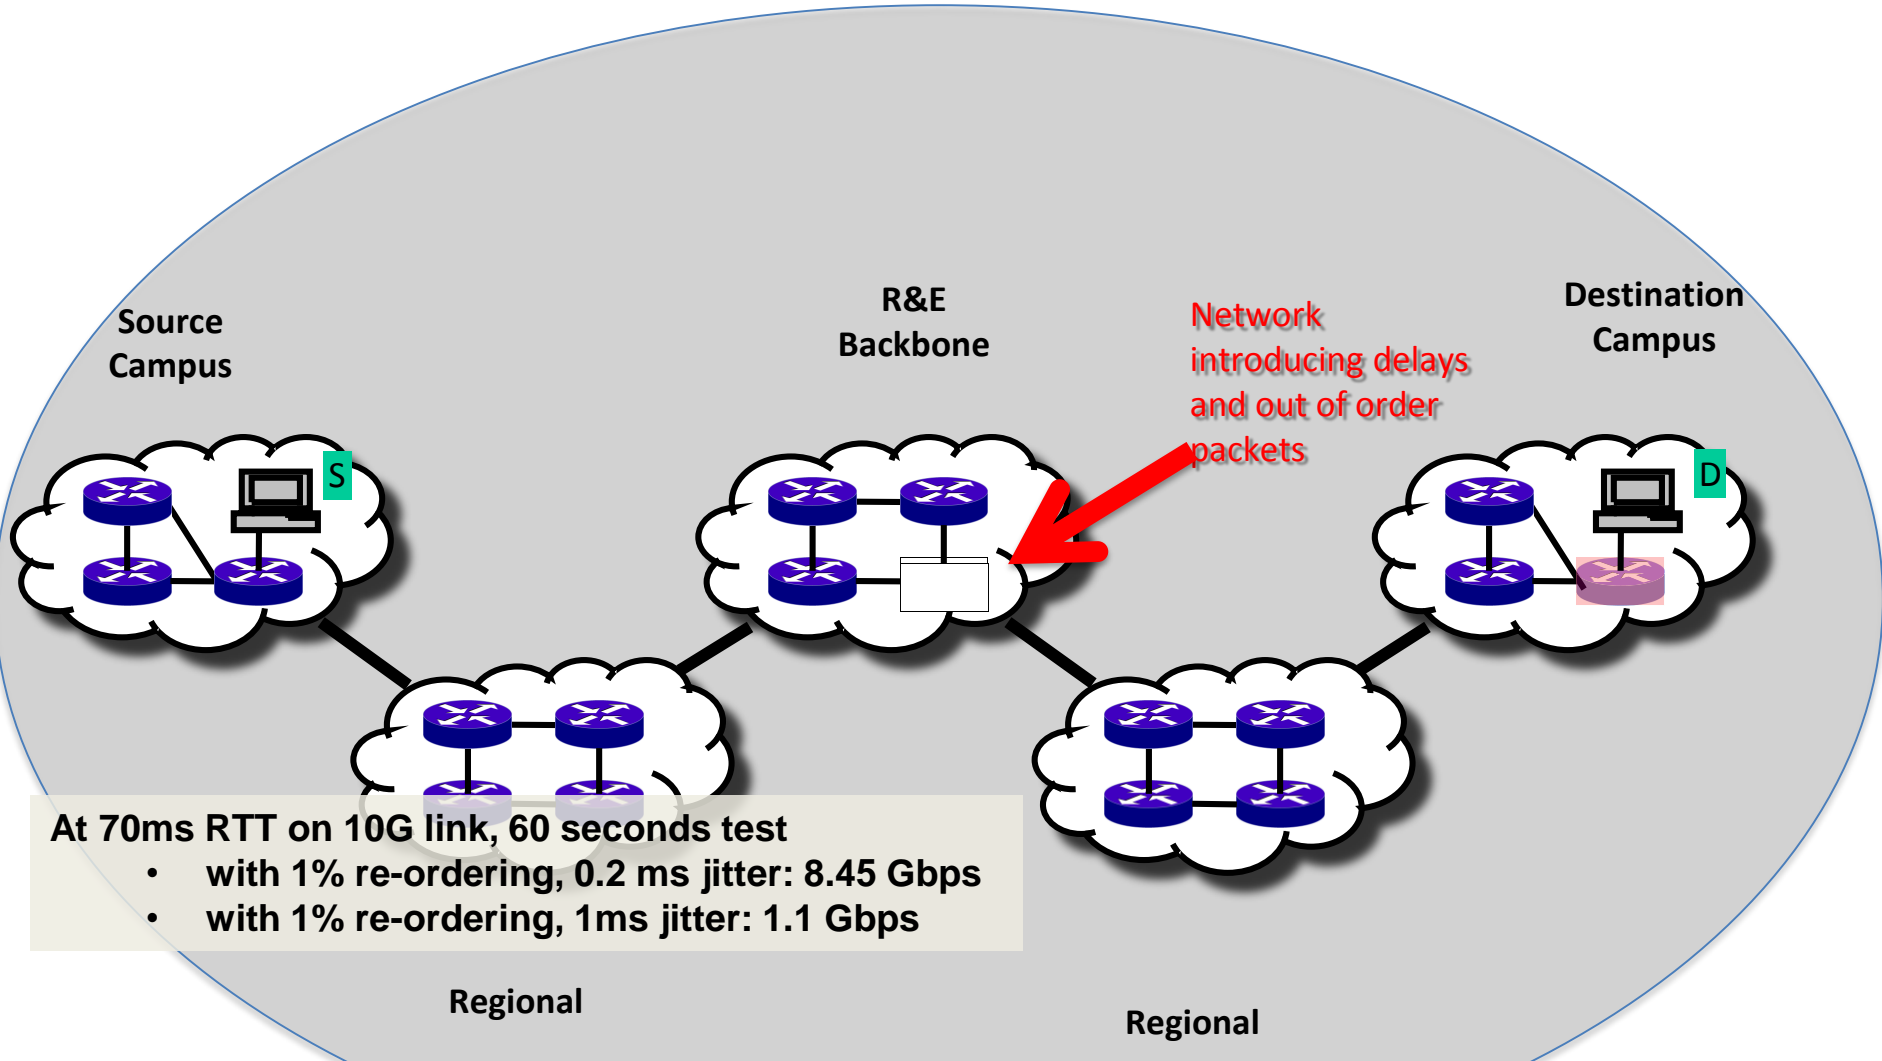
- with 1ms RTT: 7.3 Gbps
- with 51ms RTT: 122Mbps
- with 88ms RTT: 60 Mbps (factor 80)

Switch with small buffers

Regional

Regional

Packet ordering and jitter



Throughput predictions

- 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
 - $\text{Rate} < (\text{MSS}/\text{RTT}) * (1 / \text{sqrt}(p))$
 - MSS – segment size
 - RTT – round trip time
 - p – packet loss
- Packet (re)ordering and jitter to be added as well

ATLAS Network Analytics

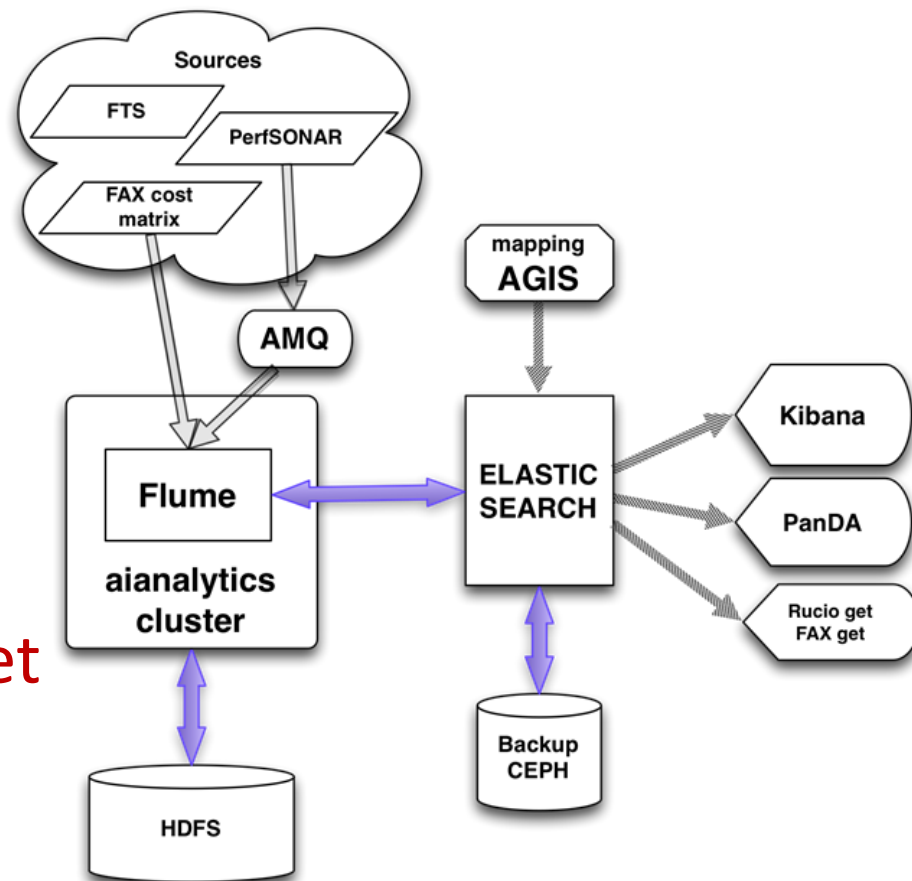
- Diagram shows the flow

- 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>



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*

Future Directions

- The WLCG efforts at CERN are being reorganized and this is an opportunity to chart future directions for the working group
- 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)
- There is a review of the working group scheduled for April 28th during the WLCG Operations Coordination meeting.

Possible Future Project Areas

- **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