


LHCOPN monitoring with perfSONAR-PS

John Shade /CERN IT-CS

September 2011 LHCOPN/LHCONE Meeting Amsterdam



Introduction

- Washington LHCOPN meeting in June agreed to use perfSONAR for LHCONE, and to deploy perfSONAR-PS in the LHCOPN
- Thanks to a sustained collaborative effort from all sites, hardware was found and the toolkit has been deployed in all eleven Tier1s, without exception, and at the Tier0!

Acknowledgements

Site Deployment

Magnus Bergroth - NDGF
John Bigrow – BNL
Sander Boele - SARA
Laurent Caillat-Vallet - CC-IN2P3
Wenshui Chen – ASGC
Ian Collier – RAL
Donato De Girolamo – CNAF
Vitaliy Kondratenko - TRIUMF
Virginie Longo - CERN
Fernando Lopez – PIC
Julia Rohlfig - KIT
Darryl Wohlt - FNAL


Technical Support

Jason Zurawski – Internet2

Dashboard Development

Tom Wlodek – BNL

Dashboard




RACF The Experimental Independent perfSONAR Dashboard

- Main Page
- All Clouds
- Individual Clouds:
 - USATLAS
 - IT
 - LHCOPN
- Primitive Services
- perfSONAR Sites
- List of Hosts
- List of Matrices
- List of Clouds
- List of Sites
- RACF dashboard
- perfSONAR dashboard
- RACF dashboard (test)
- perfSONAR dashboard (test)

Status as of: Sat Sep 24 08:07:22 EDT 2011

Cloud LHCOPN

Sites of LHCOPN cloud



LHCOPN Cloud Throughput Matrix

	---	0	1	2	3	4	5	6	7	8	9	10	11
0:BNL (ncom.bnl.gov)	---	0.20	0.20	0.20	0.20	0.20	0.20	0.20	0.20	0.20	0.20	0.20	0.20
1:ASGC (perfsone-as.bvgrid.org)	0.20	---	0.20	0.20	0.20	0.20	0.20	0.20	0.20	0.20	0.20	0.20	0.20
2:CC-IN2P3 (coperfsone-hcoopn.in2p3.fr)	0.20	0.20	---	0.20	0.20	0.20	0.20	0.20	0.20	0.20	0.20	0.20	0.20
3:CERN (perfsone-ps.cern.ch)	0.20	0.20	0.20	---	0.20	0.20	0.20	0.20	0.20	0.20	0.20	0.20	0.20
4:CNAF (perfsone-ps.cnafr.infn.it)	0.20	0.20	0.20	0.20	---	0.20	0.20	0.20	0.20	0.20	0.20	0.20	0.20
5:FNAL (psone1.fnal.gov)	0.20	0.20	0.20	0.20	0.20	---	0.20	0.20	0.20	0.20	0.20	0.20	0.20
6:KIT (perfsone-de-kit.griak.de)	0.20	0.20	0.20	0.20	0.20	0.20	---	0.20	0.20	0.20	0.20	0.20	0.20
7:NDGF (perfsone-ps2.ndgf.org)	0.20	0.20	0.20	0.20	0.20	0.20	0.20	---	0.20	0.20	0.20	0.20	0.20
8:PIC (perfsone-ps.pic.es)	0.20	0.20	0.20	0.20	0.20	0.20	0.20	0.20	---	0.20	0.20	0.20	0.20
9:RAL (perfsone-ps01.gridpp.rl.ac.uk)	0.20	0.20	0.20	0.20	0.20	0.20	0.20	0.20	0.20	---	0.20	0.20	0.20
10:SARA (ss.lhcops.sara.nl)	0.20	0.20	0.20	0.20	0.20	0.20	0.20	0.20	0.20	0.20	---	0.20	0.20
11:TRIUMF (perfsone-ps.lhcops-mon.triumf.ca)	0.20	0.20	0.20	0.20	0.20	0.20	0.20	0.20	0.20	0.20	0.20	---	0.20

The rows of this table represent SOURCE nodes for a throughput test while the columns represent DESTINATION nodes. Each cell in the table contains the result of two versions of a B/W/T, throughput test for the specified source and destination. Tests are configured to run by BOTH the source and destination once every 4 hour period. The user link in each cell represents the results of the throughput test initiated from the SOURCE and the over link in each cell represents the results of the throughput test initiated from the DESTINATION end. The user link in each cell represents the results of the throughput test initiated from the SOURCE and the over link in each cell represents the results of the throughput test initiated from the DESTINATION end. A cell is OK (green) if the measured bandwidth (averaged over all measurements in the last 24 hours) is >= 100 Mbits/sec. A cell is WARNING (yellow) if the measured bandwidth (averaged over all measurements in the last 24 hours) is >= 10 Mbits/sec and < 100 Mbits/sec. A cell is CRITICAL (red) if the measured bandwidth is not available (no test defined) or is < 10 Mbits/sec (averaged over all tests in the last 24 hours).

LHCOPN Cloud Latency Matrix

	---	0	1	2	3	4	5	6	7	8	9	10	11
0:BNL (ncom.bnl.gov)	---	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
1:ASGC (perfsone-as.bvgrid.org)	0.00	---	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
2:CC-IN2P3 (coperfsone-hcoopn.in2p3.fr)	0.00	0.00	---	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
3:CERN (perfsone-ps2.cern.ch)	0.00	0.00	0.00	---	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
4:CNAF (perfsone-ps.cnafr.infn.it)	0.00	0.00	0.00	0.00	---	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
5:FNAL (psone1.fnal.gov)	0.00	0.00	0.00	0.00	0.00	---	0.00	0.00	0.00	0.00	0.00	0.00	0.00
6:KIT (perfsone-de-kit.griak.de)	0.00	0.00	0.00	0.00	0.00	0.00	---	0.00	0.00	0.00	0.00	0.00	0.00
7:NDGF (perfsone-ps2.ndgf.org)	0.00	0.00	0.00	0.00	0.00	0.00	0.00	---	0.00	0.00	0.00	0.00	0.00
8:PIC (perfsone-ps.pic.es)	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	---	0.00	0.00	0.00	0.00
9:RAL (perfsone-ps01.gridpp.rl.ac.uk)	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	---	0.00	0.00	0.00
10:SARA (ss.lhcops.sara.nl)	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	---	0.00	0.00
11:TRIUMF (perfsone-ps.lhcops-mon.triumf.ca)	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	---	0.00

The rows of this table represent SOURCE nodes for a test while the columns represent DESTINATION nodes. Each cell in the table represents a source-destination LATENCY test via D/WAMP (800 UDP packets/test) tests, 1/minute. The metric we are plotting is the packet loss between the source and destination averaged over the last 30 minutes. Each cell contains the result of two tests. The upper result is the loss measured in the test initiated from the source and the lower result is the loss measured in the test initiated from the destination end. An OK (green) result is when the average packet loss is less than 2 out of 800 packets. A WARNING (orange) result is when the average packet loss >= 2 but < 10 out of 800 packets. A CRITICAL (red) result is when EITHER the test is not defined or the packet loss >= 10 out of 800 packets. An UNKNOWN (gray) result may indicate any other test outcome, including but not limited to: uncomprehensible test output, no response, test timed out etc.

LHCOPN Matrices

LHCOPN Cloud Throughput Matrix

	---	0	1	2	3	4	5	6	7	8	9	10	11
0:BNL (lhcmn.bnl.gov)	---	0.20	0.71	0.81	0.07	0.61	0.00	0.77	0.51	0.03	0.75	0.64	
	---	0.49	0.83	0.81	0.05	0.00	0.80	0.76	0.69	0.03	0.76	0.59	
1:ASGC (perfonar-ps.twgrid.org)	0.60	---	0.27	0.45	0.12	0.44	0.00	0.24	0.27	0.02	0.48	0.03	
	0.30	---	0.31	0.42	0.06	0.00	0.00	0.48	0.34	0.02	0.03	0.01	
2:CC-IN2P3 (ccperfonar-lhcopn.in2p3.fr)	0.29	0.03	---	0.94	0.03	0.05	0.93	0.63	0.39	0.08	0.82	0.19	
	0.37	0.04	---	0.94	0.03	0.00	0.93	0.73	0.49	0.08	0.78	0.08	
3:CERN (perfonar-ps.cern.ch)	0.82	0.49	0.94	---	0.45	0.65	0.94	0.93	0.70	0.40	0.93	0.37	
	0.81	0.50	0.94	---	0.57	0.00	0.00	0.93	0.71	0.39	0.93	0.16	
4:CNAF (perfonar-ps.cnaf.infn.it)	0.21	0.23	0.36	0.61	---	0.00	0.61	0.49	0.43	0.40	0.52	0.12	
	0.26	0.14	0.49	0.56	---	0.12	0.00	0.54	0.00	0.30	0.51	0.09	
5:FNAL (psonar1.fnal.gov)	0.41	0.00	0.00	0.00	0.09	---	0.00	0.82	0.00	0.00	0.00	0.10	
	0.80	0.42	0.00	0.00	0.00	---	0.62	0.00	0.36	0.00	0.84	0.10	
6:KIT (perfonar-de-kit.gridka.de)	0.75	0.50	0.94	0.94	0.05	0.83	---	0.92	0.91	0.35	0.93	0.53	
	0.74	0.00	0.94	0.94	0.05	0.00	---	0.92	0.00	0.32	0.93	0.10	
7:NDGF (perfonar-ps2.ndgf.org)	0.61	0.49	0.63	0.72	0.15	0.00	0.64	---	0.68	0.18	0.84	0.15	
	0.75	0.45	0.71	0.75	0.15	0.15	0.64	---	0.77	0.17	0.83	0.08	
8:PIC (perfonar-ps.pic.es)	0.22	0.12	0.84	0.87	0.26	0.20	0.20	0.65	---	0.15	0.62	0.19	
	0.22	0.12	0.61	0.77	0.31	0.11	0.90	0.70	---	0.23	0.75	0.05	
9:RAL (perfonar-ps01.gridpp.rl.ac.uk)	0.32	0.16	0.39	0.45	0.24	0.00	0.20	0.41	0.41	---	0.48	0.08	
	0.29	0.13	0.42	0.38	0.21	0.00	0.84	0.27	0.26	---	0.33	0.11	
10:SARA (ps.lhcopn-ps.sara.nl)	0.76	0.53	0.48	0.70	0.02	0.17	0.83	0.85	0.18	0.28	---	0.54	
	0.75	0.53	0.48	0.70	0.02	0.38	0.82	0.84	0.30	0.20	---	0.32	
11:TRIUMF (perfonar-ps.lhcopn-mon.triumf.ca)	0.08	0.01	0.12	0.05	0.04	0.24	0.11	0.24	0.08	0.01	0.00	---	
	0.24	0.01	0.12	0.06	0.03	0.40	0.04	0.00	0.11	0.01	0.38	---	

LHCOPN Cloud Latency Matrix

	---	0	1	2	3	4	5	6	7	8	9	10	11
0:BNL (lhcmn.bnl.gov)	0.0	0.0	0.0	0.0	1.0	0.0	0.0	0.0	0.0	5.0	20.0	0.0	56.0
	0.0	1.0	0.0	0.0	0.0	1.0	0.0	0.0	0.0	3.0	30.0	0.0	61.0
1:ASGC (perfonar-ps.twgrid.org)	7.0	0.0	1.0	0.0	0.0	1.0	0.0	1.0	5.0	30.0	0.0	70.0	
	8.0	0.0	0.0	1.0	0.0	0.0	0.0	1.0	4.0	23.0	0.0	65.0	
2:CC-IN2P3 (ccperfonar-lhcopn.in2p3.fr)	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	4.0	29.0	0.0	48.0	
	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	3.0	28.0	0.0	48.0	
3:CERN (perfonar-ps2.cern.ch)	0.0	0.0	1.0	0.0	0.0	1.0	0.0	0.0	3.0	29.0	0.0	63.0	
	0.0	0.0	1.0	0.0	1.0	0.0	0.0	1.0	2.0	24.0	0.0	66.0	
4:CNAF (perfonar-ow.cnaf.infn.it)	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	3.0	25.0	0.0	65.0	
	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	3.0	24.0	0.0	67.0	
5:FNAL (psonar2.fnal.gov)	1.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	4.0	24.0	0.0	111.0	
	1.0	1.0	600.0	600.0	0.0	0.0	0.0	0.0	4.0	28.0	0.0	106.0	
6:KIT (perfonar2-de-kit.gridka.de)	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	5.0	35.0	0.0	62.0	
	0.0	0.0	0.0	0.0	1.0	0.0	0.0	0.0	2.0	0.0	0.0	54.0	
7:NDGF (perfonar-ps.ndgf.org)	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	4.0	26.0	0.0	49.0	
	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	3.0	30.0	0.0	61.0	
8:PIC (perfonar-ps.pic.es)	0.0	1.0	0.0	1.0	0.0	0.0	0.0	0.0	2.0	0.0	20.0	0.0	47.0
	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	1.0	0.0	20.0	0.0	73.0
9:RAL (perfonar-ps02.gridpp.rl.ac.uk)	0.0	0.0	0.0	0.0	0.0	0.0	1.0	0.0	1.0	3.0	0.0	0.0	62.0
	0.0	1.0	1.0	2.0	0.0	1.0	2.0	1.0	5.0	0.0	0.0	0.0	0.0
10:SARA (ps.lhcopn-ps.sara.nl)	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	2.0	0.0	0.0	0.0	60.0
	0.0	0.0	0.0	0.0	1.0	0.0	0.0	1.0	3.0	0.0	0.0	0.0	56.0
11:TRIUMF (perfonar-ps.lhcopn-mon.triumf.ca)	1.0	2.0	1.0	1.0	0.0	16.0	0.0	0.0	3.0	0.0	1.0	0.0	
	0.0	2.0	1.0	0.0	21.0	10.0	0.0	2.0	3.0	27.0	2.0	0.0	

IN2P3 Throughput Tests

ps-Performance Toolkit - Throughput Tests - Monitoring

File Edit View History Bookmarks Tools Help

Hotel Casa 400 Amsterdam to Scienc... | ps-Performance Node - Throughput ... | GGUS: ID#73463 - Many transfer error... | LHCOPN and LHCONE joint meeting ...

ccperfsnar-lhcopn.in2p3.fr/toolkit/gui/perfAdmin/serviceTest.cgi?url=http://localhost:8085/perfSONAR_PS/services/psSB&eventType=http://ggf.org/ns/nmwg/tools/iperf/2.0

carouge run to run

Throughput Tests

Active Data Sets

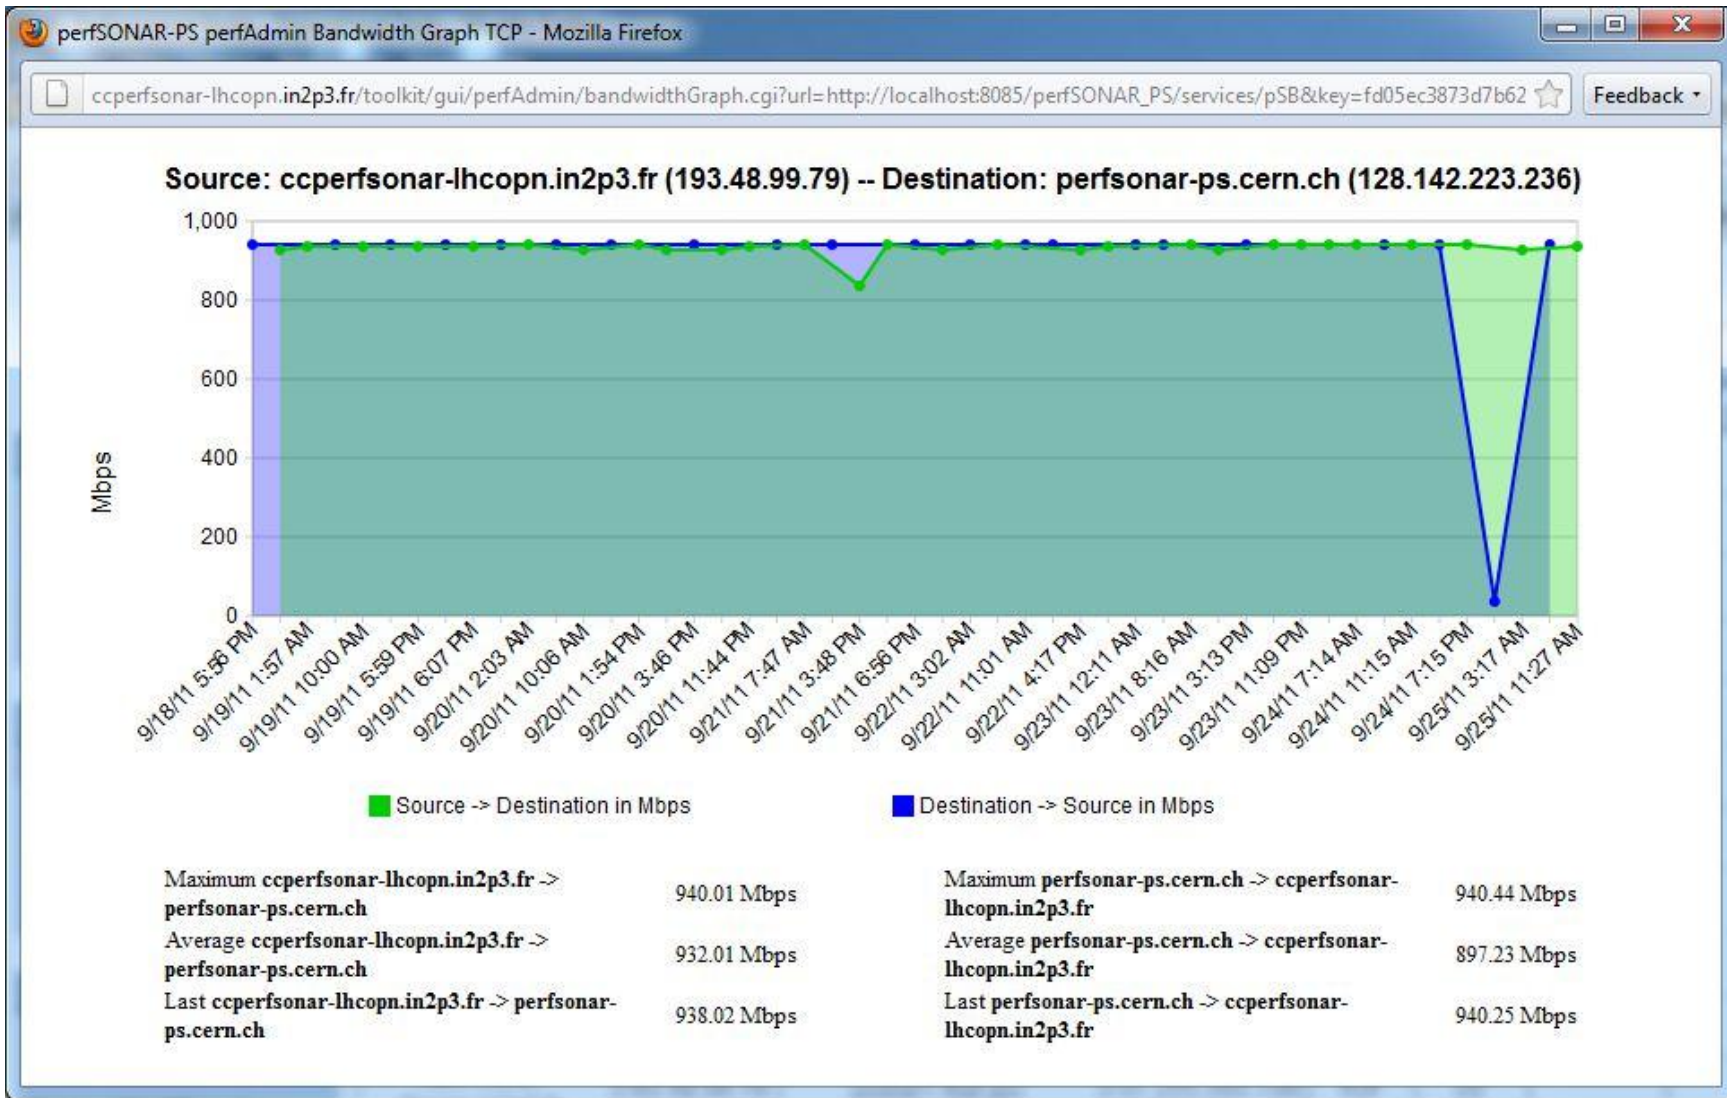
First Host	First Address	Second Host	Second Address	Protocol	Duration	Window Size	Bandwidth Limit	Bi-Directional	Line Graph	Scatter Graph
ccperfsnar-lhcopn.in2p3.fr	193.48.99.79	lhcomn.bnl.gov	192.12.15.23	TCP	20			Yes	-- Select --	-- Select --
ccperfsnar-lhcopn.in2p3.fr	193.48.99.79	perfsnar-de-kit.gridka.de	192.108.47.6	TCP	20			Yes	-- Select --	-- Select --
ccperfsnar-lhcopn.in2p3.fr	193.48.99.79	perfsnar-ps.cern.ch	128.142.223.236	TCP	20			Yes	1 Month	-- Select --
ccperfsnar-lhcopn.in2p3.fr	193.48.99.79	perfsnar-ps.cnaf.infn.it	131.154.254.11	TCP	20			Yes	-- Select --	-- Select --
ccperfsnar-lhcopn.in2p3.fr	193.48.99.79	perfsnar-ps.lhcopn-mon.triumf.ca	206.12.9.70	TCP	20			Yes	-- Select --	-- Select --
ccperfsnar-lhcopn.in2p3.fr	193.48.99.79	perfsnar-ps.pic.es	193.109.172.190	TCP	20			Yes	-- Select --	-- Select --
ccperfsnar-lhcopn.in2p3.fr	193.48.99.79	perfsnar-ps.twgrid.org	117.103.105.188	TCP	20			Yes	-- Select --	-- Select --
ccperfsnar-lhcopn.in2p3.fr	193.48.99.79	perfsnar-ps01.gridpp.rl.ac.uk	130.246.179.196	TCP	20			Yes	-- Select --	-- Select --
ccperfsnar-lhcopn.in2p3.fr	193.48.99.79	perfsnar-ps2.ndgf.org	109.105.124.88	TCP	20			Yes	-- Select --	-- Select --
ccperfsnar-lhcopn.in2p3.fr	193.48.99.79	ps.lhcopn-ps.sara.nl	145.100.17.9	TCP	20			Yes	-- Select --	-- Select --
ccperfsnar-lhcopn.in2p3.fr	193.48.99.79	psonar1.fnal.gov	131.225.205.139	TCP	20			No	-- Select --	-- Select --
ccperfsnar-lhcopn.in2p3.fr	193.48.99.79	psonar2.fnal.gov	131.225.205.141	TCP	20			No	-- Select --	-- Select --

1 Week Average Bandwidth in Mbps

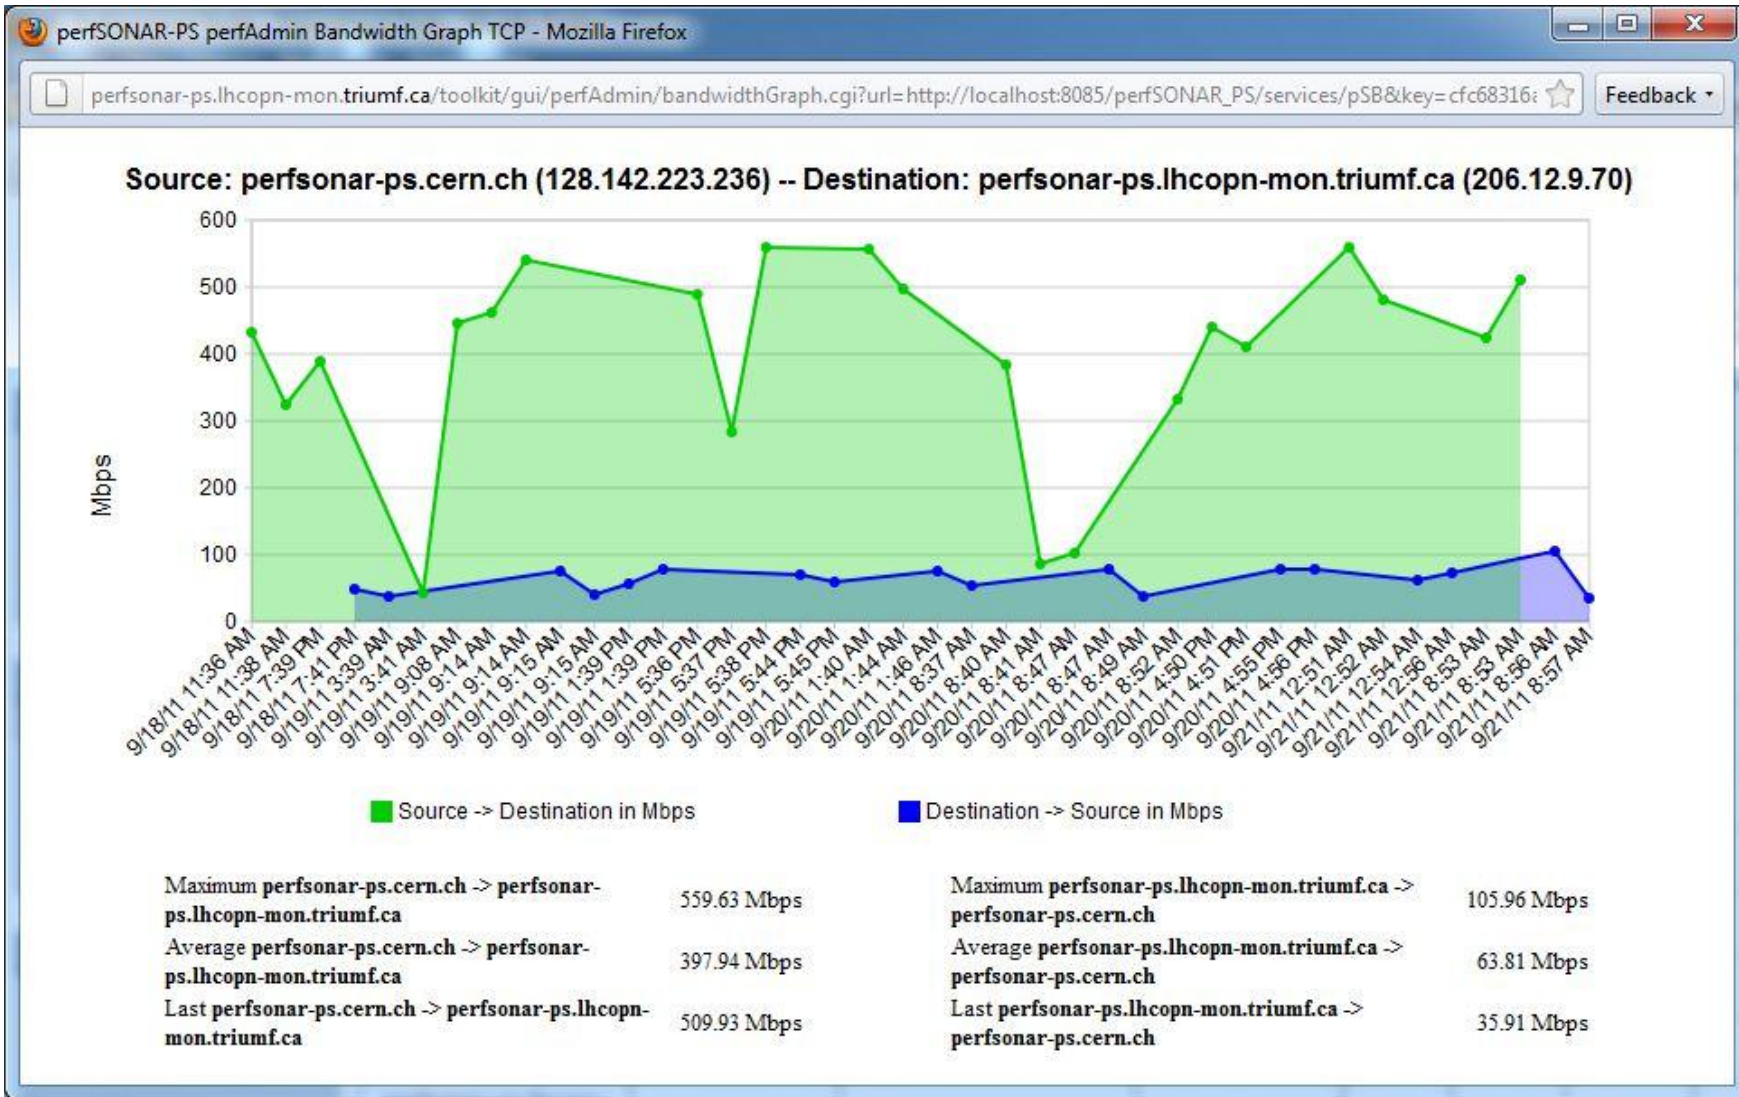
WLCG Worldwide LHC Computing Grid

perfSONAR

Throughput IN2P3-CERN



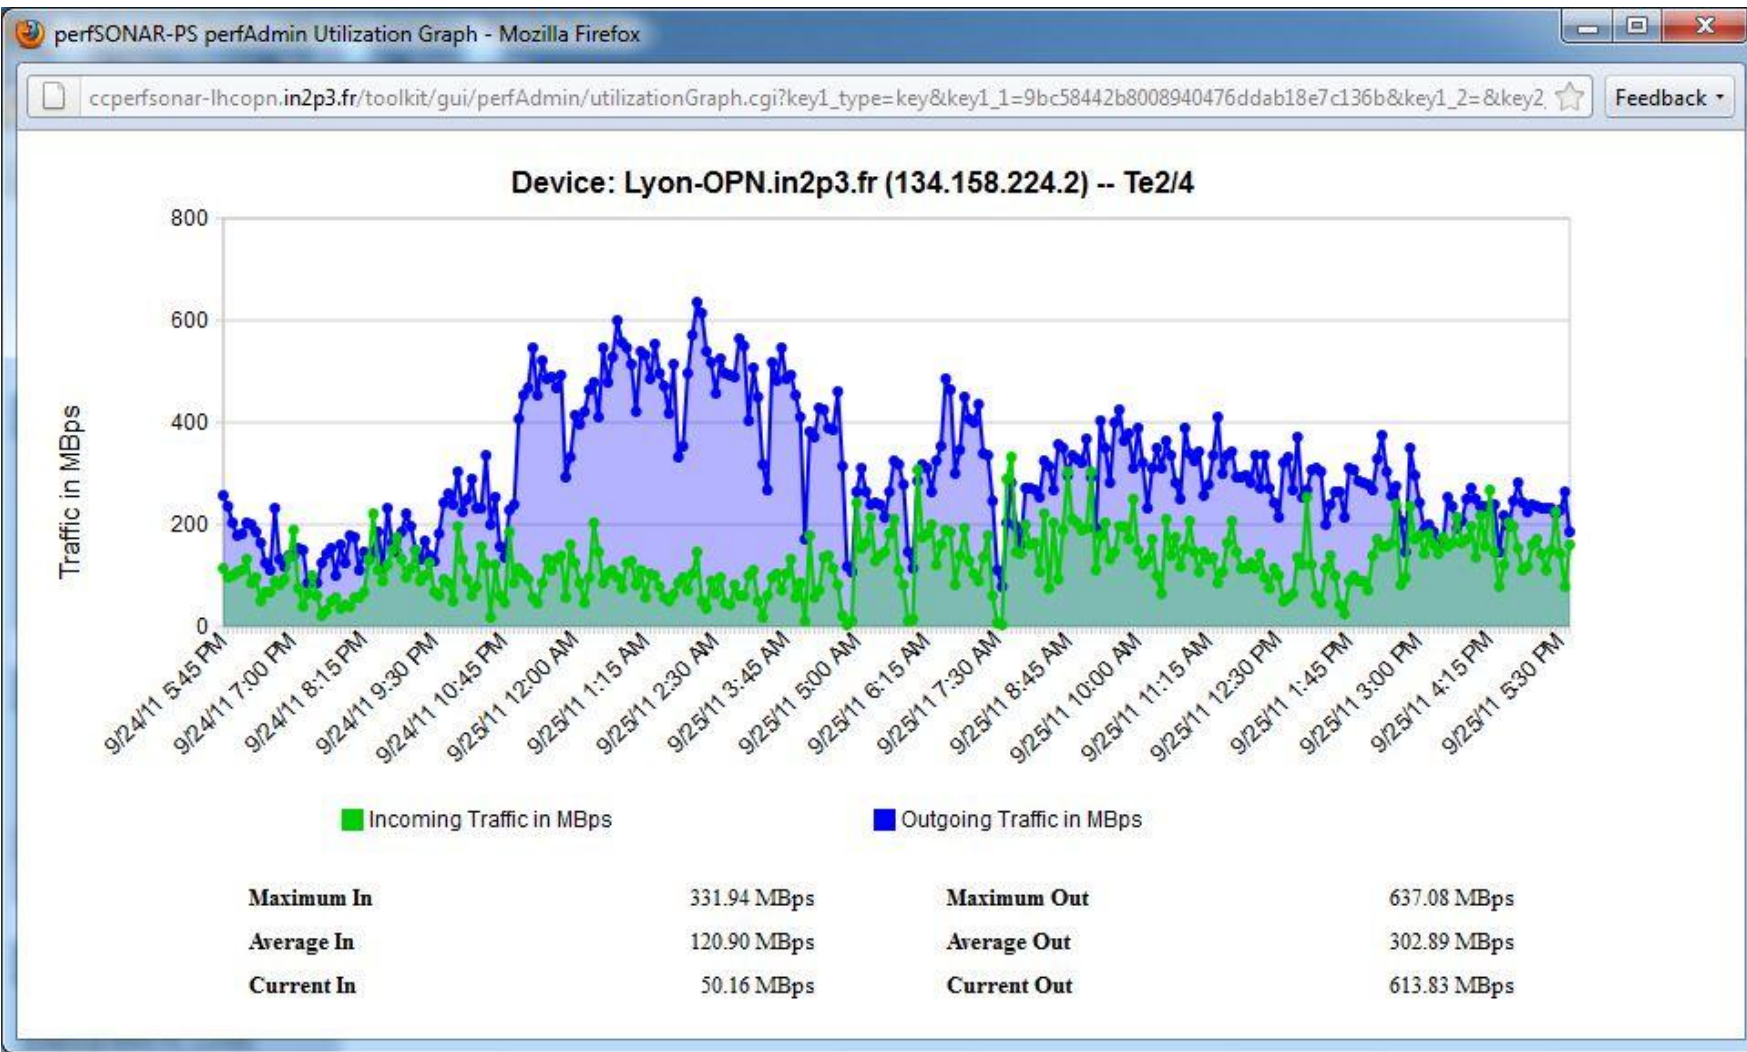
Throughput CERN-TRIUMF



IN2P3-Gridka Latency Tests



Link Utilisation



Observations

- The previous graphs were obtained from my hotel room (public WiFi). Do we want this?
- Need to harden perfSONAR-PS web-sites
 - CERN security team is testing our two servers and we will feed the results back to Jason and the developers.
 - Should we insist on valid certificates to access sites?
- perfSONAR-PS is mainly useful is the data is readily available!

Questions for the discussion

- How “open” do we make the perfSONAR-PS installations?
- How do we ensure stable operation of the perfSONAR-PS infrastructure?
- What happens to perfSONAR-MDM within the LHCOPN?
- What do we do with Sander’s dashboard?
- ... ?

Conclusion

- Great job done by all to get initial implementation working 😊 !!
- Now we need to fine-tune and achieve steady-state operations
- Few resources available, so things need to be “set-and-forget”
- Security is a concern!
- Current CERN-AGLT2 problem is already showing usefulness of perfSONAR-PS
(see GGUS: https://ggus.eu/ws/ticket_info.php?ticket=73463)

References

- Twiki:
<https://twiki.cern.ch/twiki/bin/view/LHCOPN/PerfsonarPS>
- perfSONAR Toolkit Page:
<http://psps.perfsonar.net/toolkit/>
- Tom's Dashboard:
<https://130.199.185.78:8443/exda/?page=25&cloudName=LHCOPN>
- DESY WLCG workshop presentations:
<https://indico.desy.de/conferenceOtherViews.py?view=standard&confId=4019>

