

Initial Network Monitoring for LHCONE

Shawn McKee/University of Michigan

LHCOPN/LHCONE Meeting - Berkeley

January 31st, 2012

Monitoring LHCONE: Goals/Purpose

- ❄ We need to understand how a transition to LHCONE will impact our infrastructure.
- ❄ **First step:** get monitoring in place to create a baseline of the current situation
- ❄ **Second step:** as sites transition to using LHCONE, characterize the impact based upon measurements
- ❄ To gather the before/after measurements for LHCONE we chose the **perfSONAR-PS** toolkit: already demonstrated in USATLAS and the capabilities of the modular dashboard.
- ❄ **perfSONAR's main purpose is to aid in network diagnosis** by quickly allowing users to isolate the location of problems. **In addition it can provide a standard measurement of various network performance related metrics over time as well as “on-demand” tests.**

Summary for LHCONE

- ❄ As noted, our specific goal in setting up perfSONAR-PS for LHCONE is to acquire **before** and **after** network measurements for the selected early adopter sites. This is **not** the long-term network monitoring setup for LHCONE...that is **TBD** (being discussed at this meeting)
- ❄ Details of which sites and how sites should setup the perfSONAR-PS installations is documented on the Twiki at: <https://twiki.cern.ch/twiki/bin/view/LHCONE/SiteList>
- ❄ Sites should carefully review the info on the Twiki!
- ❄ In the next few slides I will highlight some of the relevant details

LHCONE perfSONAR-PS

- ❄ We want to measure (to the extent possible) the entire network path between LHC resources. This means:
 - ❑ We want to locate perfSONAR-PS instances as close as possible to the storage resources associated with a site. The goal is to ensure we are measuring the same network path to/from the storage.
- ❄ There are two separate instances that should be deployed:
latency and bandwidth
 - ❑ The **latency instance** measures one-way delay by using an NTP synchronized clock and send 10 packets per second to target destinations
 - ❑ The **bandwidth instance** measures achievable bandwidth via a short test (20-60 seconds) per src-dst pair every 4 hour period

Network Impact of perfSONAR

- ❄ To provide an idea of the network impact of a typical deployment here are some numbers as configured in the US
 - ❑ **Latency tests** send 10Hz of small packets (20 bytes) for each testing location. USATLAS Tier-2's test to ~10 locations. Since headers account for 54 bytes each packet is 74 bytes or the rate for testing to 10 sites is **7.4 kbytes/sec**.
 - ❑ **Bandwidth tests** try to maximize the throughput. A 20 second test is run from each site in each direction once per 4 hour window. Each site runs tests in both directions. Typically the best result is around **925 Mbps on a 1Gbps link for a 20 second test**. That means we send $4 \times 925 \text{ Mbps} \times 20 \text{ sec}$ every 4 hours per testing pair (src-dst) or about 5 Mbps average.
 - ❑ Tests are configurable but the above settings are working fine.

perfSONAR-PS Issues Observed

- ❄ Getting working monitoring deployed is a **big** part of the battle
 - ❑ Focusing on a set of inter-site monitoring configuration raises awareness of the current shortcomings in our infrastructure
- ❄ Two primary problems we seem to have:
 - ❑ Traffic between Tier-2Ds and Tier-1s often **routed on congested GPN** links or is **passing thru a firewall, limiting performance**
- ❄ Issue with MTU setting. Suggestion for LHCONE is to use jumbo frames. We need to understand the impact on our measurements.
- ❄ Test durations: 1G vs 10G. 20 seconds OK for 1G, but what about 10G? 60 seconds seems more reasonable.
- ❄ Getting alerts running: Issues with false positives.
- ❄ Higher level alarms: when, how?
- ❄ Modular dashboard: intro, use, future, issues

Modular Dashboard

- ❄ Thanks to Tom Wlodek's work on developing a "modular dashboard" we have a very nice way to summarize the extensive information being collected for the near-term LHCONE network characterization.
- ❄ The dashboard provides a highly configurable interface to monitor a set of perfSONAR-PS instances via simple plug-in test modules. Users can be authorized based upon their grid credentials. Sites, clouds, services, tests, alarms and hosts can be quickly added and controlled.

Example of Dashboard for LHCONE

Browser address bar: <https://130.199.185.78:8443/exda/?page=25&cloudName=LHCONE>

RACF The Experimental Independent perfSONAR Dashboard
Grid Group

Status as of: Sun Jan 29 09:51:37 EST 2012

Cloud LHCONE

Sites of LHCONE cloud

BNL	AGLT2	INFN Napoli	SARA	ASGC	PIC
KIT	TRIUMF	Toronto	Prague	Tokyo	DESY-HH
GRIF/LAL	LRZ-LMU				

- Main Page
- All Clouds
- Individual Clouds:
 - USATLAS
 - IT
 - LHCOPN
 - LHCONE
 - CA-ATLAS
- Inter Cloud Tests:
 - AGLT2-IT
 - FR-US
- Primitive Services
- perfSONAR Sites
- List of Hosts
- List of Matrices
- List of Alarms
- List of Clouds
- List of Sites
- List of Schedulers
- Administrator Page
- Manage Users
- Define of Edit Alarms
- RACF dashboard
- perfSONAR dashboard (old)
- RACF dashboard (test)

LHCONE Throughput Matrix

	---	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
0:BNL (lhcmom.bnl.gov)	---	0.00	0.32	0.00	0.46	0.00	---	0.01	---	---	0.00	0.00	0.58	0.00	0.42	0.00	
1:AGLT2 (psmsu01.aglt2.org)	---	0.91	0.34	0.00	0.00	0.00	---	0.10	---	---	0.33	0.00	0.00	0.00	0.31	0.00	
	0.92	---	0.94	0.19	0.00	0.00	---	0.00	---	---	0.00	0.00	0.00	0.00	0.40	0.93	
	0.91	---	0.94	0.09	0.00	0.00	---	0.00	---	---	0.00	0.00	0.00	0.00	0.41	0.00	

LHCONE Latency Matrix

	---	0
0:BNL (lhcmom.bnl.gov)	---	0.0
1:AGLT2 (psmsu01.aglt2.org)	---	0.0

See <https://130.199.185.78:8443/exda/?page=25&cloudName=LHCONE>

LHCONE Latency Matrix

RACF

Grid Group

The Experimental Independent perfSONAR Dashboard

Status as of: Sun Jan 29 09:54:05 EST 2012

LHCONE Latency Matrix

- Main Page
- All Clouds
- Individual Clouds:
 - USATLAS
 - IT
 - LHCOPN
 - LHCONE
 - CA-ATLAS
- Inter Cloud Tests:
 - AGLT2-IT
 - FR-US
- Primitive Services
- perfSonar Sites
 - List of Hosts
 - List of Matrices
 - List of Alarms
 - List of Clouds
 - List of Sites
 - List of Schedulers
- Administrator Page
- Manage Users
- Define of Edit Alarms
- RACF dashboard
- perfSONAR dashboard (old)
- RACF dashboard (test)
- perfSONAR dashboard (old,test)
- Dashboard documentation

	---	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
0:BNL (lhcperfmn.bnl.gov)		0.0 0.0	0.0 0.0	0.0 0.0	0.0 0.0	0.0 0.0	0.0 0.0	0.0 0.0	5.0 5.0	---	---	0.0 0.0	0.0 0.0	0.0 0.0	0.0 1.0	0.0 0.0	0.0 0.0
1:AGLT2 (psmsu01.aglt2.org)		0.0 0.0	0.0 0.0	0.0 0.0	0.0 0.0	0.0 0.0	0.0 0.0	0.0 0.0	5.0 3.0	---	---	0.0 1.0	0.0 0.0	0.0 0.0	9.0 0.0	0.0 0.0	0.0 0.0
2:AGLT2 (psum01.aglt2.org)		0.0 0.0	0.0 0.0	0.0 0.0	0.0 0.0	0.0 0.0	0.0 0.0	0.0 0.0	0.0 4.0	---	---	0.0 1.0	0.0 1.0	0.0 0.0	0.0 9.0	0.0 0.0	0.0 0.0
3:ASGC (lhc-latency.twgrid.org)		1.0 2.0	0.0 0.0	0.0 0.0	0.0 0.0	0.0 0.0	0.0 0.0	0.0 0.0	4.0 4.0	---	---	0.0 0.0	0.0 0.0	0.0 0.0	0.0 0.0	0.0 0.0	0.0 0.0
4:CERN (perfsonar-ps2.cern.ch)		0.0 0.0	0.0 0.0	0.0 0.0	0.0 0.0	0.0 0.0	0.0 0.0	0.0 0.0	0.0 0.0	---	---	0.0 0.0	0.0 0.0	0.0 0.0	0.0 0.0	0.0 0.0	0.0 0.0
5:DESY-HH (perfsonar-ps-01.desy.de)		0.0 0.0	0.0 0.0	0.0 0.0	0.0 0.0	0.0 0.0	0.0 0.0	0.0 0.0	0.0 0.0	---	---	0.0 0.0	0.0 0.0	0.0 0.0	0.0 0.0	0.0 0.0	0.0 0.0
6:GRIF/LAL (psonar1.lal.in2p3.fr)		0.0 0.0	0.0 0.0	0.0 0.0	0.0 0.0	0.0 0.0	0.0 0.0	0.0 0.0	0.0 0.0	---	---	0.0 0.0	0.0 0.0	0.0 0.0	0.0 0.0	0.0 0.0	0.0 0.0
7:INFN Napoli (perfsonar2.na.infn.it)		6.0 6.0	7.0 0.0	5.0 0.0	7.0 5.0	0.0 0.0	0.0 0.0	0.0 0.0	0.0 0.0	---	---	7.0 5.0	8.0 6.0	0.0 0.0	0.0 9.0	0.0 0.0	0.0 0.0
8:KIT (perfsonar-de-kit.gridka.de)		---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---
9:LRZ-LMU (lcg-lrz-perfs1.grid.lrz.de)		---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---
10:PIC (perfsonar-ps-latency.pic.es)		0.0 0.0	0.0 0.0	0.0 0.0	0.0 0.0	0.0 0.0	0.0 0.0	0.0 0.0	3.0 2.0	---	---	0.0 0.0	2.0 0.0	0.0 0.0	0.0 0.0	0.0 0.0	0.0 0.0

LHCONE Throughput Matrix

RACF

Grid Group

The Experimental Independent perfSONAR Dashboard

Status as of: Sun Jan 29 09:52:25 EST 2012

LHCONE Throughput Matrix

- Main Page
- All Clouds
- Individual Clouds:
 - USATLAS
 - IT
 - LHCOPN
 - LHCONE
 - CA-ATLAS
- Inter Cloud Tests:
 - AGLT2-IT
 - FR-US
- Primitive Services
- perfSonar Sites
- List of Hosts
- List of Matrices
- List of Alarms
- List of Clouds
- List of Sites
- List of Schedulers
- Administrator Page
- Manage Users
- Define of Edit Alarms
- RACF dashboard
- perfSONAR dashboard (old)
- RACF dashboard (test)
- perfSONAR dashboard (old,test)
- Dashboard documentation

	---	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
0:BNL (lhcomon.bnl.gov)	---	0.00	0.32	0.00	0.46	0.00	0.00	0.01	---	---	0.00	0.00	0.53	0.00	0.42	0.00	
	---	0.91	0.34	0.00	0.00	0.00	0.00	0.10	---	---	0.33	0.00	0.00	0.00	0.31	0.00	
1:AGLT2 (psmsu02.aglt2.org)	0.92	---	0.94	0.19	0.00	0.00	---	0.00	---	---	0.00	0.00	0.00	0.00	0.00	0.93	
	0.91	---	0.94	0.09	0.00	0.00	---	0.00	---	---	0.00	0.00	0.00	0.00	0.41	0.00	
2:AGLT2 (psum02.aglt2.org)	0.90	0.94	---	0.00	0.00	0.00	0.00	0.00	---	---	0.00	0.00	0.00	0.00	0.43	0.00	
	0.91	0.00	---	0.05	0.00	0.00	---	0.18	---	---	0.00	0.00	0.00	0.00	0.43	0.00	
3:ASGC (lhc-bandwidth.twgrid.org)	0.33	0.04	0.00	---	0.03	0.00	---	0.00	---	---	0.00	0.00	0.00	0.14	0.45	0.39	
	0.00	0.05	0.00	---	0.00	0.00	---	0.00	---	---	0.00	0.00	0.00	0.00	0.32	0.00	
4:CERN (perfsonar-ps.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	
	0.00	0.00	0.00	0.13	---	0.00	---	0.00	---	---	0.76	0.00	0.00	0.00	0.00	0.00	
5:DESY-HH (perfsonar-ps-02.desy.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	
	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:GRIF/LAL (psonar2.lal.in2p3.fr)	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	
	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	
7:INFN Napoli (perfsonar.na.infn.it)	0.01	0.01	0.00	0.00	0.00	0.00	---	---	---	---	0.00	0.00	0.00	0.00	0.00	0.00	
	0.00	0.00	0.00	0.00	0.00	0.00	---	---	---	---	0.01	0.00	0.00	0.00	0.00	0.00	
8:KIT (perfsonar2-de-kit.gridka.de)	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	
	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	
9:LRZ-LMU (lcz-lrz-perfs2.grid.lrz.de)	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	
	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	
10:PIC (perfsonar-bandwidth.pic.fr)	0.50	0.25	0.31	0.01	0.00	0.00	---	0.03	---	---	---	0.00	0.76	0.00	0.18	0.26	
	0.00	0.22	0.10	0.04	0.00	0.00	---	0.16	---	---	---	0.00	0.00	0.00	0.00	0.00	

Using the Dashboard

- ❄ The dashboard is very useful for all of us to use to get a quick picture of the status for a particular grouping (cloud)
- ❄ It is also **very useful for sites** to debug their configurations!
- ❄ Note that you can quickly drill down and get error details as well as history plots or tables.
- ❄ I strongly wish to encourage all the early adopter sites targeted for LHCONE to use the dashboard to check status
<https://130.199.185.78:8443/exda/?page=25&cloudName=LHCONE>
(quick demo of what you can do)

Discussion/Questions

- ❄ Any questions about this presentation?

- ❄ Next up: some slides from Jason, John and me, intended to seed discussion on monitoring. On to the next set of slides...