

Update: Network Monitoring using perfSONAR-PS for LHC VOs

Shawn McKee/University of Michigan

WLCG Grid Deployment Board

September 12th, 2012

Introduction

- ❄ This is an update on the LHC network monitoring, primarily the perfSONAR-PS toolkits and associated modular dashboard, as well as some “network news”
- ❄ Outline:
 - ❑ A brief history and review
 - ❑ Current status and example issues
 - ❑ New since the last presentation
- ❄ Feel free to ask questions at anytime during the presentation...

Motivations for Common LHC Network Monitoring

- ❄ LHC collaborations rely upon the network as a critical part of their infrastructure, yet finding and debugging network problems can be difficult and, in some cases, take months.
- ❄ There is no differentiation of how the network is used amongst the LHC experiments. (Quantity may vary)
- ❄ We need a standardized way to monitor the network and locate problems quickly if they arise
- ❄ We don't want to have a network monitoring system per VO!

History of perfSONAR

- ❄ perfSONAR a joint effort of ESnet, Internet2, GEANT and RNP to standardize network monitoring protocols, schema and tools
- ❄ USATLAS adopted perfSONAR-PS toolkit starting in 2008. All Tier-2s and the Tier-1 instrumented by 2010.
- ❄ **Modular dashboard** developed by Tom Wlodek/BNL based upon USATLAS requirements to better understand deployed infrastructure
- ❄ LHCOPN choose to adopt in June 2011...mostly deployed within 3 months (by September 2011).

Network Monitoring for LHC: Goals/Purpose

❄ Goals:

- ❄ Find and isolate “network” problems; alerting in a timely way
 - ❄ Characterize network use (base-lining)
 - ❄ Provide a source of network metrics for higher level services
- ❄ **First step:** get monitoring in place to create a baseline of the current situation between sites
- ❄ **Next:** as sites evolve (join LHCONE, update their infrastructure) track the impact based upon measurements
- ❄ **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.**

LHC perfSONAR-PS Deployments

- ❄ 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

Modular Dashboard: Centralized Info

- ❄ Having a large number of perfSONAR deployments is great for instrumenting our networks, but all these instances are not easy to track, summarize or understand.
- ❄ The current modular dashboard is being used to track a large number of LHC perfSONAR-PS installations:

<https://perfsonar.racf.bnl.gov:8443/exda/>

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.

LHC “Clouds” Being Tracked

❄ ATLAS (includes some multi-VO sites)

- ❑ USATLAS (United States; 10 sites)
- ❑ IT (Italy; 4 sites)
- ❑ CA-ATLAS (Canada; 6 sites)
- ❑ UK (United Kingdom; **14 sites**)
- ❑ LHC-FR (France; 4 sites)

❄ CMS

- ❑ USCMS (United States; 10 sites)
- ❑ Some sites in Brazil and Europe are being added

❄ LHCOPN (all Tier-1 and CERN; 12 sites)

❄ LHCONE (Selected sites to track progress; 18 sites)

❄ Almost doubled the number of sites in ~9 months

Example of Dashboard (LHCONE)

RACF
Grid Group

The Production Instance of perfSONAR Dashboard

Status as of: Wed Sep 12 06:16:08 EDT 2012

Cloud LHCONE

Sites of LHCONE cloud

BNL	AGLT2	INFN Napoli	SARA	ASGC	PIC
KIT	TRIUMF	Toronto	Prague	Tokyo	LRZ-LMU
GRIF-LAL	DESY-HH	MWT2(UC)	GRIF/LPNHE		

- Main Page
- All Clouds
- Individual Clouds:
 - USATLAS
 - USCMS
 - IT
 - LHCOPN
 - LHCONE
 - CA-ATLAS
 - UK
 - LHC-FR
- 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
- Probes
- Manage Users
- Define or 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	16	17
0:BNL (lhcomon.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	0.00	0.00	0.00	0.00	0.00	0.00
1:AGLT2 (psmsu02.aglt2.org)	1.35	---	3.84	0.09	0.35	0.22	0.06	0.00	0.09	0.15	0.03	1.96	0.02	0.14	0.00	0.15	0.40	0.02	0.00
2:AGLT2 (psum02.aglt2.org)	0.55	2.73	---	0.05	0.16	0.20	0.04	0.00	0.09	0.14	0.05	2.62	0.01	0.16	0.00	0.11	0.26	0.05	0.00
3:ASGC	0.00	0.69	0.76	---	0.37	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.11	0.43	0.52	0.00

See <https://perfsonar.racf.bnl.gov:8443/exda/?page=25&cloudName=LHCONE>

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 anyone interested in network monitoring to use the dashboard to check the capabilities:
<https://perfsonar.racf.bnl.gov:8443/exda/?page=25&cloudName=LHCONE>
- ❄ Authorization for Mgmt via X509 supported.

perfSONAR-PS Issues Observed

- ❄ Getting working monitoring deployed was a **first main step**
 - ❑ Focusing on a set of inter-site monitoring configuration raises awareness of the current shortcomings in our infrastructure
- ❄ Two primary problems we noted:
 - ❑ Traffic between Tier-2Ds and Tier-1s is:
 - ⌘ **Often routed on congested GPN links**
 - ⌘ **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

Improving perfSONAR-PS Deployments

- ❄ Based upon the issues we have encountered we setup a Wiki to gather best practices and solutions to issues we have identified:

<http://www.usatlas.bnl.gov/twiki/bin/view/Projects/LHCperfSONAR>

- ❄ This page is being shared with the perfSONAR-PS developers and we expect many of the “fixes” will be incorporated into future releases.
- ❄ Please feel free to add to the Wiki (either directly or by emailing me updates/changes/additions).

Examples of Recent Issues

- ❄ Jason Zurawski/Internet2 has been central in helping to support and debug perfSONAR-PS instances and results
- ❄ Here are some recent typical issues (last few weeks):
 - ❑ SWT2/UTA: Site has bad performance to select destinations
 - ⌘ perfSONAR-PS data shows start of problem 3 months ago
 - ⌘ Under investigation
 - ❑ ASGC: Site is achieving poor throughput results
 - ⌘ Main issue is configuration related; large RTT requires either longer tests or increased network buffers
 - ❑ Univ of Oxford: Poor data transfer performance
 - ⌘ perfSONAR-PS data shows good network throughput
 - ⌘ Under investigation but may be application tuning issue
 - ❑ MWT2/UC/UIUC: Had transfer issues
 - ⌘ Used perfSONAR-PS traceroute to determine LHCONE routing issue

News and Updates (1/2)

- ❄ News: a number of new network-related items
- ❄ Open Science Grid has a new area in **networking**
 - ❑ Shawn McKee is area-coordinator
 - ❑ Focus on improving perfSONAR-PS installs, network tool/troubleshooting documentation and modular dashboard
 - ❑ Provide OSG service for collecting network metrics
- ❄ New development effort for Modular Dashboard
 - ❑ Needed to provide scalability and feature evolution
 - ❑ Joint effort of ESnet/Internet2/FNAL/BNL/OSG (and others)
 - ❑ See draft page at <https://www.opensciencegrid.org/bin/view/Documentation/NetworkingModularDashboard>

News and Updates (2/2)

- ❄ Internet2 is developing a perfSONAR-PS “Mesh Configuration” component
 - ❑ Draft info at <http://code.google.com/p/perfsonar-ps/wiki/MeshConfigurationInstallation>
 - ❑ Needs testing-of, and feedback-on, LHC use-cases
- ❄ In the US a newly funded proposal: ANSE
 - ❑ **A**dvanced **N**etwork **S**ervices for **E**xperiments (Caltech, Michigan, UTA, Vanderbilt)
 - ❑ Focus on incorporating network services into ATLAS/PaNDA and CMS/PhEDEx. Starts this Fall; funded for 2 years
- ❄ WLCG has setup a Network Working Group
 - ❑ Chair is Michael Ernst/BNL
 - ❑ <http://indico.cern.ch/getFile.py/access?contribId=7&resId=0&materialId=slides&confId=202090>

Challenges Ahead

- ❄ Getting hardware/software platform installed at all sites
- ❄ **Dashboard development:** Need additional effort to produce something suitable quickly. Ensure it meets our needs...
- ❄ Managing site and test configurations
 - ❑ **Testing and improving “centralized” (VO-based) configurations**
 - ❑ Determining the right level of scheduled tests for a site, e.g., Tier-2s test to other same-cloud Tier-2s (and Tier-1)?
 - ❑ Address 10G vs 1G tests that give misleading results
- ❄ **Alerting:** A high-priority need but complicated:
 - ❑ Alert who? Network issues could arise in any part of end-to-end path
 - ❑ Alert when? Defining criteria for alert threshold. Primitive services are easier. Network test results more complicated to decide
- ❄ Integration with VO infrastructures and applications

GDB Involvement

- ❄ The GDB could help to coordinate and motivate getting the needed deployments in place
- ❄ As noted: we want one set of network monitoring tools in place and have converged on perfSONAR
 - How to organize sites and VOs?
 - ⌘ perfSONAR-PS has communities to self-identify: recommendations?
 - ⌘ Scheduled tests: which ones, how often, what is missing?
 - Dashboard “infrastructure”. How best to configure and leverage it?
 - ⌘ Could use additional input in development process and coding
- ❄ **What ways are appropriate for the GDB to be involved?**

How to Make Progress?

- ❄ Using the LHCONE case as an example it seems possible to make significant progress in getting a standardized monitoring infrastructure in place quickly.
- ❄ Need to improve installs to be “set-it and forget-it”
- ❄ Integration with the experiments software stacks and DDM systems is now a high-priority (LS1 is an opportunity)
 - ❑ First network monitoring metrics
 - ❑ Next: SDN (Software Defined Networking)
- ❄ All VOs need to be aware of the need for network monitoring and the possibilities for sharing a common solution. Requires VO “pressure” to get sites to deploy
- ❄ VOs must assign effort to configure and gather VO view of network from shared perfSONAR measurement locations

Discussion/Questions

Questions or Comments?

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.