

Network Performance Monitoring with perfSONAR-PS

Shawn McKee

University of Michigan

Thanks to Philippe Laurens, Jason Zurawski,
Tom Wlodek and the perfSONAR-PS team



perfSONAR
powered



Outline

- Network monitoring and an overview of perfSONAR and perfSONAR-PS
- Experience in USATLAS
- Details about perfSONAR-PS as deployed
- Outlook and Plans

Network Monitoring for LHC

- Our **W**orldwide LHC **C**omputing **G**rid relies heavily (implicitly) on the underlying networks that interconnect our sites and resources
- Networks, because of their distributed nature, are typically difficult to debug when there are end-to-end issues impacting users/applications:
 - Multiple “owners” (administrative domains)
 - Lack of information to quickly isolate problem locations when “network” problems are suspected
 - Insufficient expertise/tools to diagnose networks

Why Network Monitoring?

- As you have seen in yesterday's [talk by Daniele](#) on LHCONE preparations, both ATLAS and CMS are working on verifying and monitoring our inter-site links and their end-to-end performance
- This is **critical** and **necessary** work to ensure a robust infrastructure for LHC...
- ...BUT **not** sufficient!
- We need additional information to differentiate **end-site** issues from **network** issues

perfSONAR and perfSONAR-PS

- perfSONAR (www.perfsonar.net) is:
 - An infrastructure for network performance monitoring
 - A services-oriented architecture with 3 views:
 - A **consortium** of organizations building and designing an interoperable set network monitoring middle-ware
 - A **protocol** based upon SOAP XML messages
 - Several interoperable **software packages**
- perfSONAR-PS is a perfSONAR development effort, by Internet2/ESnet and others, targeted at creating an easy-to-deploy and easy-to-use set of perfSONAR services

Robust Network is Key

- USATLAS started working closely with the perfSONAR-PS team more than 2.5 years ago.
- Goal was to instrument the network connections between US Tier 1 and all US Tier 2 sites in one uniform way.
- **Primary motives**
 - Aid in problem diagnosis and location identification
 - Differentiate end-site issues from network issues
 - Archive of standard regular measurements over time
- USATLAS Deployed perfSONAR-PS
 - Implemented two identical end points in each facility
 - Defined mesh of connection tests between all facilities

perfSONAR-PS Deployment for USATLAS

- Deployed the same inexpensive hardware at all sites
 - ~\$600 per KOI 1U system, 1Gb NIC
 - Now looking into Dell R410 for hardware refresh
- Used same perfSONAR_PS Toolkit linux live-CD
 - now most sites use the net-install, rpm distribution
- Dedicate one node for throughput and one node for latency at each site
 - Throughput tests are resource intensive and tend to bias latency tests
- Define a common set of test
 - Mesh of Throughput tests to/from all T1/T2 perfSONAR nodes
 - Mesh of Latency tests to/from all T1/T2 perfSONAR nodes
- Now augmented with summary via Dashboard


perfSONAR-PS Tools

- Web based GUI for admin to configure and for user to display measurement tests
 - After initial setup of local disk, IP, NTP
- Nodes may be declared being part of some “communities” (e.g. LHC or USATLAS) to help identification in a directory lookup service
- Two main test types
 - Throughput tests (bwctl) non-concurrent
 - Ping (PingER) and One-Way Latency tests (owamp) time stampedTests are scheduled and Measurement Archive manages results
- Also available
 - traceroute and ping (i.e. reverse route from remote PS host)
 - Network Diagnostic Tools (NDT,NPAD) on demand
 - Cacti installed

perfSONAR-PS: Web GUI

← → ↻ 🏠 psmsu02.aglt2.org/toolkit/

📁 Net 📁 Personar 📁 CRAC 📁 ATLAS 📁 AGLT2 📁 gDocs 📁 DZero 📁 Home



User Tools

- Local Performance Services
- Global Performance Services
- Java OWAMP Client [↗](#)
- Reverse Traceroute [↗](#)
- Reverse Ping [↗](#)

Service Graphs


- Throughput
- One-Way Latency
- Ping Latency
- SNMP Utilization
- Cacti Graphs [↗](#)

Toolkit Administration

- Administrative Information
- External BWCTL Limits
- External OWAMP Limits
- Enabled Services
- NTP
- Scheduled Tests
- Cacti SNMP Monitoring [↗](#)

Performance Toolkit

- Configuration Help [↗](#)
- Frequently Asked Questions [↗](#)
- About [↗](#)
- Credits [↗](#)



pS-Performance Node For AGLT2 at Michigan State University In East Lansing, MI, USA


Host Information	
Organization Name	AGLT2 at Michigan State University
Host Location	East Lansing, MI, USA
Administrator Name	Philippe Laurens
Administrator Email	aglt2-admin@umich.edu

Communities This Host Participates In	
USATLAS LHC	

Host Status	
Primary Address	psmsu02.aglt2.org

Services Offered	
Bandwidth Test Controller (BWCTL) [1]	Running
• tcp://psmsu02.aglt2.org:4823	
Lookup Service [1]	Running
• http://psmsu02.aglt2.org:9995/perfSONAR_PS/services/hLS	
Network Diagnostic Tester (NDT) [1]	Running
• tcp://psmsu02.aglt2.org:3001	
• http://psmsu02.aglt2.org:7123 ↗	
Network Path and Application Diagnosis (NPAD) [1]	Running
• tcp://psmsu02.aglt2.org:8100	
• http://psmsu02.aglt2.org:8000 ↗	
One-Way Ping Service (OWAMP) [1]	Disabled
• tcp://psmsu02.aglt2.org:861	
perfSONAR-BUOY Regular Testing (Throughput) [1]	Running
perfSONAR-BUOY Measurement Archive [1]	Running
• http://psmsu02.aglt2.org:8085/perfSONAR_PS/services/pSB	
perfSONAR-BUOY Regular Testing (One-Way Latency) [1]	Disabled
PingER Measurement Archive and Regular Tester [1]	Disabled
• http://psmsu02.aglt2.org:8075/perfSONAR_PS/services/pinger/ma	
SNMP Measurement Archive [1]	Running
• http://psmsu02.aglt2.org:9990/perfSONAR_PS/services/SNMPMA	

perfSONAR-PS: Throughput Tests web page



User Tools

- Local Performance Services
- Global Performance Services
- Java OWAMP Client
- Reverse Traceroute
- Reverse Ping

Service Graphs

- Throughput
- One-Way Latency
- Ping Latency
- SNMP Utilization
- Cacti Graphs

Toolkit Administration

- Administrative Information
- External BWCTL Limits
- External OWAMP Limits
- Enabled Services
- NTP
- Scheduled Tests
- Cacti SNMP Monitoring


Performance Toolkit

- Configuration Help
- Frequently Asked Questions
- About


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
atlas-npt2.bu.edu	192.5.207.252	psmsu02.aglt2.org	192.41.236.32	TCP	20			Yes	-- Select --	-- Select --
iut2-net2.iu.edu	149.165.225.224	psmsu02.aglt2.org	192.41.236.32	TCP	20			Yes	-- Select --	-- Select --
lhcmn.bnl.gov	192.12.15.23	psmsu02.aglt2.org	192.41.236.32	TCP	20			Yes	-- Select --	-- Select --
netmon2.atlas-swt2.org	129.107.255.27	psmsu02.aglt2.org	192.41.236.32	TCP	20			Yes	-- Select --	-- Select --
osgx4.hep.uiuc.edu	192.17.18.41	psmsu02.aglt2.org	192.41.236.32	TCP	20			Yes	-- Select --	-- Select --
ps2.ochep.ou.edu	129.15.40.232	psmsu02.aglt2.org	192.41.236.32	TCP	20			Yes	-- Select --	-- Select --
psmsu02.aglt2.org	192.41.236.32	psnr-bw01.slac.stanford.edu	134.79.104.209	TCP	20			Yes	-- Select --	-- Select --
psmsu02.aglt2.org	192.41.236.32	psum02.aglt2.org	192.41.230.20	TCP	20			Yes	-- Select --	-- Select --
psmsu02.aglt2.org	192.41.236.32	uct2-net2.uchicago.edu	128.135.158.219	TCP	20			Yes	-- Select --	-- Select --

1 Week Average Bandwidth in Mbps




perfSONAR-PS: Throughput graphs



perfSONAR-PS: Latency Tests web page

← → ↻ 🏠 psmsu01.aglt2.org/toolkit/gui/perfAdmin/serviceTest.cgi?url=http://localhost:8085/perfSONAR_PS/services/pSB&eventType=http://ggf.org ☆

📁 Net 📁 Perfsonar 📁 CRAC 📁 ATLAS 📁 AGLT2 📁 gDocs 📁 DZero 📁 Home



User Tools

- Local Performance Services
- Global Performance Services
- Java OWAMP Client
- Reverse Traceroute
- Reverse Ping

Service Graphs

- Throughput
- One-Way Latency
- Ping Latency
- SNMP Utilization
- Cacti Graphs


Toolkit Administration

- Administrative Information
- External BWCTL Limits
- External OWAMP Limits
- Enabled Services
- NTP
- Scheduled Tests
- Cacti SNMP Monitoring

One-Way Latency Tests

Active Data Sets					
First Host	First Address	Second Host	Second Address	Bi-Directional	Graph
atlas-npt1.bu.edu	192.5.207.251	psmsu01.aglt2.org	192.41.236.31	Yes	4 Hours ▾
iut2-net1.iu.edu	149.165.225.223	psmsu01.aglt2.org	192.41.236.31	Yes	4 Hours ▾
lhcpfmon.bnl.gov	192.12.15.26	psmsu01.aglt2.org	192.41.236.31	Yes	4 Hours ▾
netmon1.atlas-swt2.org	129.107.255.26	psmsu01.aglt2.org	192.41.236.31	Yes	4 Hours ▾
ps1.occhep.ou.edu	129.15.40.231	psmsu01.aglt2.org	192.41.236.31	Yes	4 Hours ▾
psmsu01.aglt2.org	192.41.236.31	psmsu01.aglt2.org	192.41.236.31	Yes	-- Select -- ▾
psmsu01.aglt2.org	192.41.236.31	psnr-lat01.slac.stanford.edu	134.79.104.208	Yes	4 Hours ▾
psmsu01.aglt2.org	192.41.236.31	psum01.aglt2.org	192.41.230.19	Yes	4 Hours ▾
psmsu01.aglt2.org	192.41.236.31	uct2-net1.uchicago.edu	128.135.158.216	Yes	4 Hours ▾


perfSONAR-PS: Latency Graph



perfSONAR-PS: Reverse traceroute

← → ↻ ⬆ iut2-net1.iu.edu/toolkit/gui/reverse_traceroute.cgi?target=psmsu01.aglt2.org ☆ ⚙ * 🎨 ⬆ 📄 🔍

Net Perfsonar CRAC ATLAS AGLT2 gDocs DZero Home Other bookmarks



Stanford
Linear
Accelerator
Center

traceroute from 149.165.225.223 (iut2-net1.iu.edu) to 192.41.236.31 (psmsu01.aglt2.org) for 35.9.71.18

CGI script maintainer: [Les Cottrell, SLAC](#). Script version 4.60, 8/9/09, Les Cottrell.
[Download perl source code.](#)

To perform a traceroute from iut2-net1.iu.edu, enter the desired target [host domain](#) (e.g. www.yahoo.com) or [Internet address](#) (e.g. 137.138.28.228) in the box below:

Enter target name or address: then push 'Enter' key.

Lookup: [host name](#) | [mail domain](#) | [domain name](#) | [Locating a Host](#) | [visual traceroute](#) | [contacting someone](#)

Related web sites
[Traceroute servers.](#)
[Monitoring tutorial.](#)
[Internet monitoring](#)
[What is my IP address?](#)

Please note that traceroutes can appear similar to port scans. If you see a suspected port scan alert, for example from your firewall, with a series of ports in the range 33434 - 33465, coming from iut2-net1.iu.edu it is probably a reverse traceroute from our web based reverse traceroute server. Please do NOT report this to us, it will almost certainly be a waste of both of our times. For more on this see [Traceroute security issues](#).

```
Executing exec(traceroute, -m 30 -q 3, 192.41.236.31, 140)
traceroute to 192.41.236.31 (192.41.236.31), 30 hops max, 140 byte packets
 1 149.165.225.254 (149.165.225.254)  1.068 ms  1.033 ms  1.023 ms
 2 xe-0-2-0.2012.rtr.ictc.indiana.gigapop.net (149.165.254.249)  0.369 ms  0.367 ms  0.358 ms
 3 xe-0-0-0.1.rtr.ll.indiana.gigapop.net (149.165.254.26)  0.350 ms  0.341 ms  0.484 ms
 4 xe-0-0-1.110.rtr.atla.net.internet2.edu (149.165.254.21)  84.994 ms  84.976 ms  84.967 ms
 5 ge-0-2-0.0.rtr.chic.net.internet2.edu (64.57.28.103)  35.100 ms  35.093 ms  35.084 ms
 6 198.32.11.46 (198.32.11.46)  35.631 ms  35.497 ms  35.613 ms
 7 198.32.43.158 (198.32.43.158)  46.469 ms  46.456 ms  46.438 ms
 8 psmsu01.aglt2.org (192.41.236.31)  40.504 ms  40.739 ms  40.506 ms
```

Centralized Monitoring of the Distributed Monitoring: The BNL Dashboard

- We monitor 9 separate T1/T2 sites, i.e. 18 perfSONAR nodes
 - Total of **108** critical services, **72** throughput tests, **72** one-way latency tests
 - While perfSONAR is, by design, a decentralized architecture, we need a centralized Dashboard to keep track of the overall mesh
 - Developed by BNL (Tom Wlodek) for USATLAS (and now other clouds)
 - First within Nagios (but complex and hard to access)
 - Now rewritten as a standalone project accessible by all (and portable)
 - Use probes to monitor proper operation of critical services on each node
 - Alert emails sent to site admins on failing services
 - Use probes to retrieve the latest test results on pre-defined mesh of measurements (Throughput & Latency)
 - Both measurements about link A → B measured by BOTH A & B
 - Thresholds on results for label (OK, CRITICAL, etc) and color code
 - History and plot of service status and mesh measurements
- **Presents a compact overview of (perfSONAR nodes health and) all USATLAS inter-site network connections**

Dashboard in Nagios

← → ↻ ⬆ <https://nagios.racf.bnl.gov/nagios/cgi-bin/prod/perfSonar.php?page=123> ★ ⚙ ⚡ 📄 ⚙ ⚡ ⚡
 Net Perfsonar CRAC ATLAS AGLT2 gDocs DZero Home Other bookmarks

RACF
Grid Group

The BNL perfSONAR Dashboard

Main Page
 perfSonar Compact View
 perfSONAR Sites
 System Overview
 Subsystems:
 perfSONAR Hosts
 perfSONAR Services
 Throughput Matrix
 Latency Matrix
 Other monitoring links:
 RACF Nagios
 perfSONAR Nagios Table
 Dashboard instances:
 RACF dashboard
 perfSONAR dashboard (Nagios)
 perfSONAR dashboard (new)
 Gritta/RSV/perfSONAR (experimental)

perfSONAR Overview

All perfSONAR sites:


BNL WT2 AGLT2 MWT2 NET2 SWT2

Throughput Matrix										Latency Matrix									
Status as of: July 7, 2011, 9:30 am	0	1	2	3	4	5	6	7	8	Status as of: July 7, 2011, 9:30 am	0	1	2	3	4	5	6	7	8
0:psmsu02.aglt2.org (AGLT2)	-	OK OK	OK OK	OK OK	OK OK	OK OK	OK OK	OK OK	OK OK	0:psmsu01.aglt2.org (AGLT2)	OK OK	OK OK	OK OK	OK OK	OK OK	OK OK	OK OK	OK OK	OK OK
1:psum02.aglt2.org (AGLT2)	OK OK	-	OK OK	OK OK	OK OK	OK OK	OK OK	OK OK	OK OK	1:psum01.aglt2.org (AGLT2)	OK OK	OK OK	OK OK	OK OK	OK OK	OK OK	OK OK	OK OK	OK OK
2:lhcmn.bnl.gov (BNL)	OK OK	OK OK	-	OK OK	OK OK	OK OK	OK OK	OK OK	OK OK	2:lhcmn.bnl.gov (BNL)	OK OK	OK OK	OK OK	OK OK	OK OK	OK OK	OK OK	OK OK	OK OK
3:uct2-net2.uchicago.edu (MWT2)	OK OK	OK OK	OK OK	-	OK OK	OK OK	OK OK	OK OK	OK OK	3:uct2-net1.uchicago.edu (MWT2)	OK OK	OK OK	OK OK	OK OK	OK OK	OK OK	OK OK	OK OK	OK OK
4:iut2-net2.iu.edu (MWT2)	OK OK	OK OK	OK OK	OK OK	-	OK OK	OK OK	OK OK	OK OK	4:iut2-net1.iu.edu (MWT2)	OK OK	OK OK	OK OK	OK OK	OK OK	OK OK	OK OK	OK OK	OK OK
5:atlas-npt2.bu.edu (NET2)	OK OK	OK OK	OK OK	OK OK	OK OK	-	OK OK	OK OK	OK OK	5:atlas-npt1.bu.edu (NET2)	OK OK	OK OK	OK OK	OK OK	OK OK	OK OK	OK OK	OK OK	OK OK
6:netmon2.atlas-swt2.org (SWT2)	OK OK	OK OK	OK OK	OK OK	OK OK	OK OK	-	OK OK	OK OK	6:netmon1.atlas-swt2.org (SWT2)	OK OK	OK OK	OK OK	OK OK	OK OK	OK OK	OK OK	OK OK	OK OK
7:ps1.occhep.ou.edu (SWT2)	OK OK	OK OK	OK OK	OK OK	OK OK	OK OK	OK OK	-	OK OK	7:ps1.occhep.ou.edu (SWT2)	OK OK	OK OK	OK OK	OK OK	OK OK	OK OK	OK OK	OK OK	OK OK
8:psnr-bw01.slac.stanford.edu (WT2)	OK OK	OK OK	OK OK	OK OK	OK OK	OK OK	OK OK	OK OK	-	8:psnr-lat01.slac.stanford.edu (WT2)	OK OK	OK OK	OK OK	OK OK	OK OK	OK OK	OK OK	OK OK	OK OK

Status as of: July 7, 2011, 9:30 am

(c) 2010 Brookhaven National Laboratory - send suggestions and comments to tomw@bnl.gov

Standalone Dashboard



RACF
Grid Group

The Experimental Independent perfSONAR Dashboard

Status as of: Thu Jul 07 09:33:04 EDT 2011

Main Page

Clouds:

USATLAS

IT

Primitive Services

perfSONAR Sites


RACF dashboard

perfSONAR dashboard

RACF dashboard (test)

perfSONAR dashboard (test)

Services for USATLAS Cloud



Throughput

	---	0	1	2	3	4	5	6	7	8
0:BNL (lhcmn.bnl.gov)	---	0.87 0.92	0.63 0.92	0.91 0.92	0.92 0.91	0.92 0.90	0.48 0.57	0.87 0.90	0.82 0.48	
1:AGLT2 (psmsu02.aglt2.org)		0.91 0.91	---	0.72 0.73	0.91 0.91	0.94 0.80	0.87 0.92	0.90 0.92	0.88 0.64	
2:AGLT2 (psum02.aglt2.org)		0.82 0.91	0.82 0.92	---	0.91 0.91	0.92 0.83	0.74 0.92	0.48 0.67	0.91 0.92	0.68 0.64
3:MWT2 (iut2-net2.iu.edu)		0.85 0.81	0.75 0.91	0.56 0.74	---	0.94 0.93	0.91 0.91	0.90 0.90	0.91 0.91	0.89 0.56
4:MWT2 (uct2-net2.uchicago.edu)		0.92 0.92	0.94 0.93	0.75 0.90	0.93 0.94	---	0.93 0.93	0.93 0.93	0.90 0.69	
5:NET2 (atlas-npt2.bu.edu)		0.91 0.88	0.92 0.92	0.88 0.75	0.91 0.89	0.92 0.83	---	0.88 0.91	0.91 0.52	
6:SWT2 (netmon2.atlas-svt2.org)		0.90 0.71	0.92 0.92	0.77 0.90	0.88 0.90	0.93 0.93	0.89 0.91	---	0.93 0.62	
7:SWT2 (ps2.ochep.ou.edu)		0.82 0.67	0.93 0.93	0.75 0.88	0.11 0.17	0.93 0.93	0.88 0.90	0.93 0.93	---	0.90 0.71
8:WT2 (psnr-bw01.slac.stanford.edu)		0.54 0.86	0.65 0.88	0.63 0.75	0.63 0.82	0.83 0.85	0.59 0.85	0.72 0.88	0.89 0.90	---

Latency

	---	0	1	2	3	4	5	6	7	8
0:BNL (lhcmn.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	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	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 0.0	0.0 0.0
3:MWT2 (iut2-net1.iu.edu)		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
4:MWT2 (uct2-net1.uchicago.edu)		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:NET2 (atlas-npt1.bu.edu)		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:SWT2 (netmon1.atlas-svt2.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 0.0	0.0 0.0
7:SWT2 (ps1.ochep.ou.edu)		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
8:WT2 (psnr-lat01.slac.stanford.edu)		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

Dashboard: Primitive Services

RACF
Grid Group

The Experimental Independent perfSONAR Dashboard

Status as of: Thu Jul 07 09:35:47 EDT 2011

perfSONAR Services at BNL

Host: lhcmmon.bnl.gov

net.perfsonar.service.ma.perfSONAR_pSB lhcmmon.bnl.gov	net.perfsonar.service.ma.perfSONAR_hLS lhcmmon.bnl.gov	net.perfsonar.service.ma.NDT_port_7123 lhcmmon.bnl.gov
net.perfsonar.service.ma.bwctl_port_8570 lhcmmon.bnl.gov	net.perfsonar.service.ma.bwctl_port_4823 lhcmmon.bnl.gov	net.perfsonar.service.ma.NPAD_port_8001 lhcmmon.bnl.gov
net.perfsonar.service.ma.NDT_port_3001 lhcmmon.bnl.gov	net.perfsonar.service.ma.NPAD_port_8000 lhcmmon.bnl.gov	

Host: lhccperfmon.bnl.gov

net.perfsonar.service.ma.perfSONAR_pSB lhccperfmon.bnl.gov	net.perfsonar.service.ma.perfSONAR_hLS lhccperfmon.bnl.gov	net.perfsonar.service.ma.owamp_port_8569 lhccperfmon.bnl.gov
net.perfsonar.service.ma.owamp_port_861 lhccperfmon.bnl.gov		

(c) 2010 Brookhaven National Laboratory - send suggestions and comments to tomw@bnl.gov

Dashboard: Service History

RACF
Grid Group

The BNL perfSONAR Dashboard

Main Page

perfSonar Compact View

perfSONAR Sites

System Overview

Subsystems:

perfSONAR Hosts

perfSONAR Services

Throughput Matrix

Latency Matrix

Other monitoring links:

RACF Nagios

perfSONAR Nagios Table

Dashboard Instances:

RACF dashboard

perfSONAR dashboard (Nagios)

perfSONAR dashboard (new)

Gratia / RSV / perfSONAR (experimental)

History of Service CheckLookupService on host psmsu01.aglt2.org.


☐ Past 6 hours ☐ Past 12 hours ☐ Past 24 hours ☐ Past 48 hours ☐ Past week ☐ Past month ☐ Past year ☐ All Data

Status	Time	Nagios Message
OK	2011-07-07 08:59:43	PS_LS_CHECK OK - Service functioning normally.
OK	2011-07-07 07:59:40	PS_LS_CHECK OK - Service functioning normally.
OK	2011-07-07 06:59:38	PS_LS_CHECK OK - Service functioning normally.
OK	2011-07-07 05:59:34	PS_LS_CHECK OK - Service functioning normally.
OK	2011-07-07 04:59:36	PS_LS_CHECK OK - Service functioning normally.
OK	2011-07-07 03:59:38	PS_LS_CHECK OK - Service functioning normally.
WARNING	2011-07-07 02:59:38	PS_LS_CHECK WARNING - Service returned unexpected response.
WARNING	2011-07-07 01:59:38	PS_LS_CHECK WARNING - Service returned unexpected response.
WARNING	2011-07-07 01:49:40	PS_LS_CHECK WARNING - Service returned unexpected response.
WARNING	2011-07-07 01:39:35	PS_LS_CHECK WARNING - Service returned unexpected response.
OK	2011-07-07 00:39:34	PS_LS_CHECK OK - Service functioning normally.
OK	2011-07-06 23:39:40	PS_LS_CHECK OK - Service functioning normally.
OK	2011-07-06 22:39:38	PS_LS_CHECK OK - Service functioning normally.


Status as of: July 7, 2011, 9:40 am

(c) 2010 Brookhaven National Laboratory - send suggestions and comments to tomw@bnl.gov

Dashboard: Throughput Measurement plot



Dashboard: Latency Measurement plot



Dashboard: other clouds

RACF
Grid Group

The Experimental Independent perfSONAR Dashboard

Status as of: Wed Jul 06 16:48:45 EDT 2011

Main Page

Clouds:

USATLAS

IT

Primitive Services

perfSonar Sites


RACF dashboard

perfSONAR dashboard

RACF dashboard (test)

perfSONAR dashboard (test)

Services for IT Cloud



Throughput

	---	0	1
0:Napoli (perfsonar.na.infn.it)	---	0.47 0.51	
1:Roma (perfsonar.roma1.infn.it)	0.44 0.51		---

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 BWCTL 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 upper link in each cell represents the results of the throughput test initiated from the SOURCE end.

The lower 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)

Latency

	---	0	1
0:Napoli (perfsonar.na.infn.it)	0.0 0.0	0.0 0.0	
1:Roma (perfsonar.roma1.infn.it)	0.0 0.0	0.0 0.0	---

The rows of this table represent SOURCE nodes for a test while the columns represent nodes.

Each cell in the table represents a source-destination LATENCY test via OWAMP (600 packets/test) tests, 1/minute.

The metric we are plotting is the packet loss between the source and destination averaged 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 end.

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 600 packets.

A 'WARNING' (orange) result is when the average packet loss is ≥ 2 but < 10 out of 600 packets.

A 'CRITICAL' (red) result is when EITHER the test is not defined or the packet loss is ≥ 10 out of 600 packets.

An 'UNKNOWN' (brown) result may indicate any other test outcome, including but not limited to: test timed out, no response, test timed out etc.

(c) 2010 Brookhaven National Laboratory - send suggestions and comments to tomw@bnl.gov

Diagnostics Capabilities

- **Throughput:** Notice problems and debug network, also help differentiate server problems from path problems
- **Latency:** Notice route changes, asymmetric routes
- Watch for excessive Packet Loss
- On-demand tests and NPAD/NDT diagnostics via web
- **Optionally:** Install additional perfSONAR nodes inside local network, and/or at periphery
 - Characterize local performance and internal packet loss
 - Separate WAN performance from internal performance
- Daily Dashboard check of own site, and peers

Example of diagnostics


Source: iut2-net1.iu.edu (149.165.225.223) -- Destination: psmsu01.aglt2.org (192.41.236.31)



- U.** Loss Observed (DST to SRC) Lost 7 packets out of 600
2011-5-10
- T.** Loss Observed (SRC to DST) Lost 1 packets out of 600
2011-5-10
- S.** Loss Observed (DST to SRC) Lost 1 packets out of 600
2011-5-10
- R.** Loss Observed (DST to SRC) Lost 342 packets out of 600
2011-5-10

psmsu02.aglt2.org/toolkit/gui/perfAdmin/bandwidthGraph.cgi?url=http://localhost:8085/perfSONAR_PS/services/pSB&key=453892e8c4088495c4700d3b7badcbcd&key2=29b

Source: iut2-net2.iu.edu (149.165.225.224) -- Destination: psmsu02.aglt2.org (192.41.236.32)




Maximum iut2-net2.iu.edu -> psmsu02.aglt2.org 924.88 Mbps
Average iut2-net2.iu.edu -> psmsu02.aglt2.org 915.93 Mbps
Last iut2-net2.iu.edu -> psmsu02.aglt2.org 912.80 Mbps

Maximum psmsu02.aglt2.org -> iut2-net2.iu.edu 912.82 Mbps
Average psmsu02.aglt2.org -> iut2-net2.iu.edu 239.58 Mbps
Last psmsu02.aglt2.org -> iut2-net2.iu.edu 912.42 Mbps

Asymmetric
throughput between
peer sites IU and
AGLT2 resolved

Examples of diagnostics...

- Most recently, after routing upgrade work, noticed small 0.7ms latency increase
- Traceroute found an unintended route change (packets destined to MSU were going through UM) → routing prefs quickly fixed



Outlook and Plans

- perfSONAR-PS being deployed in other ATLAS clouds
 - Italy started, Canada also in process
 - BNL Dashboard already monitoring IT cloud (at least for now)
 - Dashboard code will be packaged & distributed
- perfSONAR-PS being deployed at LHC T1 sites
 - LHCOPN voted in the June 2011 meeting to deploy it
- Connections between T2 of different clouds also need to be monitored to support file transfer - **Inter-cloud monitoring**
 - **Challenge is to identify appropriate subset; Cannot support full-mesh**
- perfSONAR is open source, with releases ~twice a year
- The more test points along the path, the better
 - Integrating information from backbone, routing points
 - Allows a “divide-and-conquer” approach to problem isolation
- **Has proven extremely useful for USATLAS to-date!**

Thank you! – Questions?

- perfSONAR
 - <http://www.perfsonar.net/>
 - <http://psps.perfsonar.net/toolkit>
 - Jason Zurawski zurawski@internet2.edu (I2/LHC contact)
- USATLAS perfSONAR Dashboard
 - Nagios <https://nagios.racf.bnl.gov/nagios/cgi-bin/prod/perfSonar.php> needs BNL login
 - Standalone <http://130.199.185.78:28080/exda/>
 - Tom Wlodek tomw@bnl.gov

perfSONAR-PS IPv6 Status

(Green check means activity is done. Yellow means in progress. Red means no firm plans to address.)

Tool	IPV6 Development Completed	Limited Deployment	Production Deployment Across Multiple Domains
BWCTL	✓	✓	✓
OWAMP	✓	✓	✓
pSB MA	✓	✓	
Lookup Services	✓		
Topology Service	✓		
SNMP MA	✓	✓	
PinGER	✓	✓	
NDT	✓		
NPAD	✗		
Toolkit Configuration Tools & GUI	✓		

Slide: Joe Metzger/ESnet