

LHCONE perfSONAR-PS Testing Plans and Status

⚠ NOTE: This page is now deprecated and this monitoring is superseded by the work being done in the WLCG perfSONAR-PS Task Force See <https://twiki.cern.ch/twiki/bin/view/LCG/PerfsonarDeployment> for installation instructions.

This page is intended to be the place where everyone can get information about setting up **perfSONAR-PS** for LHCONE testing. As noted it is now deprecated and this information is being kept for reference only.

NOTE: This bug has been fixed in the most recent perfSONAR-PS (v3.2.2)

⚠ A small perfSONAR-PS bug has been identified (thanks to Enzo Capone, Philippe Laurens and Andy Lake!) which causes DNS names with 20 or more characters before the first '.' to be interpreted as an IPv6 address. This currently impacts setting up tests to PIC. The work-around is to use PICs IP addresses when setting up the tests. They are:

- perfsonar-ps-latency.pic.es 193.109.172.189
- perfsonar-ps-bandwidth.pic.es 193.109.172.190

Please update your configurations for PIC using IP addresses for now! Thanks.

First a bit about the **Purpose** of all of this.

1. We want to be able to quickly characterize the **current** networking situation between those sites proposed to take part in testing LHCONE. The list of sites is shown in the table below.
2. After the all sites convert to using LHCONE we want to then measure the networking situation and compare with the previous measurements

The proposed tests below are **not** intended to be in place indefinitely. On the contrary, once we have completed the **before** and **after** measurements we should plan to remove the full mesh of measurements described on this page. Longer term we need to have a measurement infrastructure in place and we will need to discuss how best to do that. This page is **not** about a long-term measurement infrastructure.

Also a quick note on the physical location for the LHCONE perfSONAR-PS instances: **Our strong recommendation is to co-locate the two perfSONAR-PS nodes with the sites primary grid-storage.** The reason is that we want the perfSONAR-PS instances to measure as much of the **network** path as is possible, end-to-end. The perfSONAR-PS measurements are intended to represent what the network is doing end-to-end and can be used to differentiate network problems from end-host/storage/software problems.

All LHCONE site-networking details **should** be documented on <https://twiki.cern.ch/twiki/bin/view/LHCONE/LhcOneVRF>. We hope that it will help new sites to set up their router configurations, and provide help to those experiencing problems. In particular, sites should be able to check their BGP configurations and ensure that they are receiving the correct routes. Please make sure your site details are added there by either directly editing that Twiki (if you have access) or sending your details to Edoardo Martelli (edoardo.martelli@cernNOSPAMPLEASE.ch) so he can include it.

SiteList < LHCONE < TWiki

The following table documents the perfSONAR-PS sites involved in this initial LHCONE testing. Most columns are self-evident. There are 2 **Setup** columns, one after the latency instance and one after the bandwidth instance. For installed we put **Y** only if the specific instances is the latest version (currently 3.2.2) and if the corresponding services are configured. We put an **N** if the instance is not the latest, not running or not configured. We put a **?** if we haven't gotten information on a particular instance. The **LHCONE** column shows if the site has added the *LHCONE* community to their perfSONAR-PS install. The **MTU** column tracks what the MTU setting is on the **bandwidth** instance. The **Comments** lists specific concerns or notes about the site and its setup.

Site Name	Country	Tier	Contact	LAT Nod
AGLT2 (MSU)	US	Tier-2D	Philippe Laurens laurens@paNOSPAMPLEASE.msu.edu	psmsu01.aglt2
AGLT2 (UM)	US	Tier-2D	Shawn McKee smckee@umichNOSPAMPLEASE.edu	psum01.aglt2
DESY-HH	DE	Tier-2D	Kars Ohrenberg Kars.Ohrenberg@desyNOSPAMPLEASE.de	perfsonar-ps-01.c
GRIF/IRFU	FR	Tier2	irfuGRID_ADMINISTRATION@ceaNOSPAMPLEASE.fr	perfsonar01.datagrid.c
GRIF/LAL	FR	Tier-2D	Michel Jouvin jouvin@lalNOSPAMPLEASE.in2p3.fr	psonar1.lal.in2p
GRIF/LPNHE	FR	Tier-2D	Victor Mendoza mendoza@lpnheNOSPAMPLEASE.in2p3.fr	lpnhe-psl.in2p
LRZ-LMU	DE	Tier-2D	Christoph Anton Mitterer christoph.anton.mitterer@lmuNOSPAMPLEASE.de	lcg-lrz-perfs1.gri
MWT2(UC)	US	Tier-2D	Rob Gardner rwg@hepNOSPAMPLEASE.uchicago.edu	uct2-net1.uchicag
Napoli	IT	Tier-2D	Enzo Capone ecapone@naNOSPAMPLEASE.infn.it	perfsonar2.na.in
Prague	CZ	Tier-2D	Petr Vokac petr.vokac@cernNOSPAMPLEASE.ch	ps01-l.farm.parti
Tokyo	JP	Tier-2D	Tomoaki Nakakura tomoaki@iceppNOSPAMPLEASE.s.u-tokyo.ac.jp	perfsonar1.icep
Toronto	CA	Tier-2D	Leslie Groer groer@physicsNOSPAMPLEASE.utoronto.ca	ps-latency.scinet.uto
ASGC	TW	Tier-1	Wenshui Chen chenws@twgridNOSPAMPLEASE.org	lhc-latency.twgr
BNL	US	Tier-1	John Bigrow big@bnlNOSPAMPLEASE.gov	lhcperfmon.bnl
CERN	CH	Tier-1	Virginie Longo neteng@cernNOSPAMPLEASE.ch	perfsonar-ps2.ce
PIC	ES	Tier-1	Fernando Lopez network@picNOSPAMPLEASE.es	psl01.pic.es

SARA	NL	Tier-1	Sander Boele sanderb@saraNOSPAMPLEASE.nl	ps.lhcopn-ps.sa
TRIUMF	CA	Tier-1	Vitaliy Kondratenko vitaliyk@triumfNOSPAMPLEASE.ca	ps-latency.lhcopn-mo
KIT	DE	Tier-1	Bruno Hoeft bruno.hoeft@kitNOSPAMPLEASE.edu	perfsonar2-de-kit.g

LHCONE perfSONAR-PS Test Configuration

For all sites in the above table, we want to configure a "full-mesh" of tests. We plan on having:

- Latency (OWAMP) tests
- Bandwidth (BWCTL) tests
- Traceroute tests

The proposal is to do things in two steps: 1) Get all sites configured and advertising membership in the **LHCONE** community (See above), and 2) Setup the tests above to each of the other LHCONE sites in the table.

Step 1 (Target Date ASAP)

Step one is to get the appropriate perfSONAR-PS services installed and participating in the LHCONE community. The plan is to have all sites finish step 1) **ASAP**.

The perfSONAR-PS release notes are visible at:

http://psps.perfsonar.net/toolkit/releasenotes/pspt-3_2_2.html

The quick-start Wiki is here: <http://code.google.com/p/perfsonar-ps/wiki/pSPPerformanceToolkit322>

Some additional information for LHCONE testing sites:

- You may want to install the `NetInstall` version which will install to the local system disk. The system can then use `yum` to update itself.
- After installing (either the `NetInstall` or `LiveCD` versions) you will need to setup the services running on each type of node. Our convention so far has been to make the first node (by name or IP) the `Latency` node and the second node the `Bandwidth` node. This is easy to configure by using the Web GUI and selecting `Enabled Services` on the left hand navigation panel under `Toolkit Administration`. You can select the button at the bottom for enabling only `Latency` or only `Bandwidth` services. On the `Bandwidth` node you should make sure to enable the two `Traceroute` services (the `MA` and `Scheduler`).
- Each site should fill out the appropriate `Administrative Information` (under `Toolkit Administration` on left of Web GUI). The `Communities` section (see <http://code.google.com/p/perfsonar-ps/wiki/pSPPerformanceToolkit322#Communities>) should have `LHCONE` added in addition to whatever other communities the site wants to list (`ATLAS`, `LHC`, etc.)
- The `NTP` servers need to be setup carefully for the `Latency` node. Ideally at least 4 `good` servers should be configured (add `local` or regional ones if they are not in the distributed list).
- Firewalls may be an issue (See comments in table above). If you suspect your site will block **ANY** of the sites listed above, can you update your firewalls to allow just the specific set of perfSONAR-PS

instances for LHCONE to connect to your instances?

- We you have finished installing your two instances (Latency and Bandwidth) please *update* the table above or send the information to Shawn McKee (smckee@umichNOSPAMPLEASE.edu) .j

Once all sites have done the above it should be easy to add the required tests. By using the LHCONE community it should be easy to find the appropriate sites when setting up the Scheduled Testing in step 2 below.

Step 2 (Target Date ASAP)

For step 2) we want to implement a full set of scheduled tests between the various LHCONE sites in the table above. There are 3 tests that we want to configure to every other LHCONE test-site:

- Latency tests (10 packets/sec via OWAMP)
- Bandwidth tests (4 hour testing window using TCP Iperf with a 30 second test)
- Traceroute tests (A traceroute every 10 minutes using the defaults for this test)

Once the other sites are visible in the Community Lookup service it is easy to add tests. See this section of the notes: http://code.google.com/p/perfsonar-ps/wiki/pSPerformanceToolkit322#Scheduled_Testing. **NOTE: When you go to configure your site, some other sites in the table above may not be advertising their participation in the "LHCONE" community. You can directly add sites in any of the above tests by typing in the needed DNS entries from the table above.**

Latency Test Details

On your site's Latency node's web GUI, login and click on Scheduled Tests under Toolkit Administration . Then click the Add New One-Way Delay Test button. Under Description use LHCONE Latency Test and leave the Packet Rate and Packet Size at the defaults of 10 and 20. You will be brought to a new screen showing No Members in Test under Test Members . You should be able to click on the LHCONE community under the Find Hosts To Test With area. For each of the Latency hosts in the list I will distribute you should click the Add To Test link after it. Once those are all added **and you click SAVE**, you have setup the Latency tests. You can find the current **checklist** of latency nodes here: LHCONE_perfSONAR-PS_latencynode.txt. Please make sure all of them have latency tests configured from your latency node.

Here is a reference latency configuration from **psum01.aglt2.org** (NOTE: the psum01.aglt2.org host is **not** listed since the test is running there. Other sites should be sure to include it in their configuration of course!):

Test Parameters	
Description	LHCONE Latency Test
Packet Rate (packets per second)	10
Packet Size (bytes)	20
Edit Test Parameters	

Test Members		
lhc-latency.twgrid.org	OWAMP Server at ASGC	Delete
perfsonar1.icepp.jp	OWAMP Server at The University of Tokyo in ICEPP	Delete
perfsonar-ps2.cern.ch	OWAMP Server at CERN	Delete
psmsu01.aglt2.org	OWAMP Server at AGLT2 at Michigan State University in East Lansing, MI	Delete
psonar1.lal.in2p3.fr	OWAMP Server at GRIF/LAL	Delete
lhcpfmon.bnl.gov	OWAMP Server at BNL	Delete
ps-latency.scinet.utoronto.ca	OWAMP Server at SciNet High Performance Computing Consortium in University of Toronto	Delete
perfsonar-ps-latency.pic.es	OWAMP Server at ES-PIC in Barcelona, Spain	Delete
ps.lhcopn-ps.sara.nl	OWAMP Server at SARA	Delete
perfsonar-ps-01.desy.de	OWAMP Server at DESY-HH	Delete
perfsonar.farm.particle.cz	OWAMP Server at Institute of Physics AS CR in Prague	Delete
ps-latency.lhcopn-mon.triumf.ca	OWAMP Server at TRIUMF in Vancouver	Delete
perfsonar2.na.infn.it	OWAMP Server at INFN Napoli - Atlas Tier2 in Naples - Italy	Delete
perfsonar-de-kit.gridka.de	OWAMP Server at KIT	Delete

Bandwidth Test Details

On your site's Bandwidth node's web GUI, login and click on **Scheduled Tests** under **Toolkit Administration**. Then click the **Add New Throughput Test** button. Under **Description** use **LHCONE Bandwidth Test** and set the **Time Between Tests** to be **4 Hours**, make sure the **Test Duration** is **30 Seconds** and that the **Bandwidth Tests** is **Iperf**, the **Protocol** is **TCP** and the **Use Autotuning** box is checked. You will be brought to a new screen showing **No Members in Test** under **Test Members**. You should be able to click on the **LHCONE** community under the **Find Hosts To Test With** area. For each of the Bandwidth hosts in the list I will distribute you should click the **Add To Test** link after it. Once those are all added **and you click SAVE**, you have setup the Bandwidth tests. You can find the current **checklist** of bandwidth nodes here: `LHCONE_perfSONAR-PS_bandwidthnode.txt`. Please make sure all of them have traceroute tests configured from your node.

Here is a reference bandwidth configuration from **psum02.aglt2.org** (NOTE: the psum02.aglt2.org host is **not** listed since the test is running there. Other sites should be sure to include it in their configuration of course!):

Test Parameters	
Description	LHCONE Bandwidth Test
Test Duration (seconds)	30
Inter-Test Interval (seconds)	14400
Bandwidth Tester	Iperf
Protocol	TCP
Use Autotuning	yes
Edit Test Parameters	

Test Members		
ps-bandwidth.scinet.utoronto.ca	BWCTL Server at SciNet HPC Consortium in University of Toronto	Delete
lcg-lrz-perfs2.grid.lrz.de	BWCTL Server at LRZ-LMU	Delete
perfonar-ps.cern.ch	BWCTL Server at CERN	Delete
psmsu02.aglt2.org	BWCTL Server at AGLT2 at Michigan State University in East Lansing, MI, USA	Delete
lhc-bandwidth.twgrid.org	BWCTL Server at ASGC	Delete
perfonar-ps-bandwidth.pic.es	BWCTL Server at PIC	Delete
ps.lhcopn-ps.sara.nl	BWCTL Server at SARA	Delete
perfonar2.icepp.jp	BWCTL Server at The University of Tokyo in ICEPP	Delete
psonar2.lal.in2p3.fr	BWCTL Server at GRIF/LAL	Delete
perfonar.na.infn.it	BWCTL Server at INFN Napoli - Atlas Tier2 in Naples - Italy	Delete
perfonar-bw.farm.particle.cz	BWCTL Server at Prague	Delete
ps-bandwidth.lhcopn-mon.triumf.ca	BWCTL Server at TRIUMF in Vancouver, Canada	Delete
perfonar-ps-02.desy.de	BWCTL Server at DESY-HH	Delete
lhcomon.bnl.gov	BWCTL Server at BNL	Delete
perfonar2-de-kit.gridka.de	BWCTL Server at KIT	Delete

BWCTL Port Configuration

For BWCTL (Throughput) nodes you need to increase the number of ports available. The way that BWCTL works is that there is a connection that is done before `iperf` is run to synchronize the two testers, and then the connection for the `iperf` test itself. If you make a change through the GUI it splits the port range into two equal parts: the first range 'peer_port' is for the control connection; the second range 'iperf_port' is for the `iperf` connection. We recommend providing 500 ports for BWCTL's use: 5001-5500. If you want to edit the file manual it is `/etc/bwctld/bwctld.conf` on your throughput node. Change it to look something like:

```
group    bwctl
iperf_port 5251-5500
user     bwctl
peer_port 5001-5250
facility  local5
```

Traceroute Test Details

To setup the Traceroute test, on your site's **Latency** node's web GUI, login and click on `Scheduled Tests` under `Toolkit Administration`. Then click the `Add New Traceroute Test` button. Under `Description` use `LHCONE Traceroute Test` and set the `Time Between Tests` to be 10 Minutes. The rest of the values can be left at the defaults. You will be brought to a new screen showing `No Members in Test` under `Test Members`. You should be able to click on the `LHCONE` community under the `Find Hosts To Test With` area. For each of the **Latency** hosts in the table above you should click the `Add To Test` link after it. Once those are all added **and you click SAVE**, you have setup the Traceroute tests. You can find the current **checklist** of latency nodes here: `LHCONE_perfSONAR-PS_latencynode.txt`. Please make sure all of them have traceroute tests configured from your latency node.

Here is a reference traceroute configuration from **psum01.aglt2.org** (NOTE: the psum01.aglt2.org host is **not** listed since the test is running there. Other sites should be sure to include it in their configuration of course!):

Description	LHCONE Traceroute Test
Time between tests (seconds)	600
Packet Size (bytes)	40
Time between packets (seconds)	0
Timeout for test (seconds)	30
Timeout for individual hops (seconds)	5
First hop to report	0
Maximum Number of Hops	64
Protocol	udp
Edit Test Parameters	

Test Members		
192.41.236.31	Traceroute Responder at AGLT2 at Michigan State University in East Lansing, MI	Delete
142.150.188.96	Traceroute Responder at SciNet HPC Consortium in University of Toronto	Delete
perfsonar-de-kit.gridka.de	Traceroute Responder at KIT	Delete
lhc-latency.twgrid.org	Traceroute Responder at ASGC	Delete
147.231.25.173	Traceroute Responder at Institute of Physics AS CR in Prague	Delete
lcg-lrz-perfs1.grid.lrz.de	Traceroute Responder at LRZ-LMU	Delete
157.82.112.68	Traceroute Responder at The University of Tokyo in ICEPP	Delete
90.147.67.252	Traceroute Responder at INFN Napoli - Atlas Tier2 in Naples - Italy	Delete
193.109.172.189	Traceroute Responder at ES-PIC in Barcelona, Spain	Delete
psonar1.lal.in2p3.fr	Traceroute Responder at GRIF/LAL	Delete
lhcperfmon.bnl.gov	Traceroute Responder at BNL	Delete
206.12.9.71	Traceroute Responder at TRIUMF in Vancouver	Delete
ps.lhcopn-ps.sara.nl	Traceroute Responder at SARA	Delete
perfsonar-ps2.cern.ch	Traceroute Responder at CERN	Delete
perfsonar-ps-01.desy.de	Traceroute Responder at DESY-HH	Delete
ps-latency.scinet.utoronto.ca	Traceroute Responder at Toronto	Delete

perfSONAR-PS Maintenance and Troubleshooting

Jason Zurawski has provided a PDF file which documents some basic maintenance, troubleshooting and repair steps to address some issues in perfSONAR-PS. Have a look at 20120204-USATLAS-pSPT.pdf.

NOTE: All LHCONE testing sites need to make sure they have provided a sufficient number of ports for testing...see section 6 in the PDF file.

There have been a few issues noticed when we utilize perfSONAR-PS at a scale that is larger than it was tested at. One example is the amount of local disk that is allowed to keep current test results. For latency tests with a mesh of about 10 sites we can exceed the default storage of 1GB of test results within a day. If your limit within perfSONAR-PS is set a 1GB, new tests will fail once you reach 1GB. There are automatic cleaning scripts which will repair this every day but it can cause testing failures during the day. The recommendation is to increase the allowed storage space to 3GB (assuming you are not pressed for local disk space). You should do this on your **latency** nodes:

- Login via the gui https://your_latency_node/toolkit/admin/owamp/ (or click "External OWAMP Limits" from the left-side of your latency node web interface)
- For the "Unprivileged Clients" box, click the "Edit Group Limits" URL

- Set 3GB (or something larger than 1GB) in the pop-up box:

The screenshot shows a dialog box titled "Edit Parameters Of Group Unprivileged Clients". It contains three input fields: "Maximum Bandwidth" with the value "1000" and a unit dropdown set to "Kbps"; "Maximum Disk Space" with the value "3" and a unit dropdown set to "Gigabytes"; and a checked checkbox for "Remove Client Results". At the bottom, there are two buttons: "Change" and "Cancel".

- Click "Save" at the bottom of the screen

Note we are trying to maintain a list of tips, maintenance items and troubleshooting at <https://www.usatlas.bnl.gov/twiki/bin/view/Projects/LHCperfSONAR> so please check there for new items.

Notes

When you go to configure your site, some other sites in the table above may not be advertising their participation in the "LHCONE" community. You can directly add sites in any of the above tests by typing in the needed DNS entries from the table above.

Tom Wlodek has been developing a Modular Dashboard to summarize perfSONAR test results. This is being used for ATLAS Tier-1 Clouds (Currently the US, UK, Italy and Canada) as well as the LHCOPN. You can see the dashboards here:

- **LHCOPN** <https://1perfsonar.racf.bnl.gov:8443/exda/?page=25&cloudName=LHCOPN>
- **Canada** <https://perfsonar.racf.bnl.gov:8443/exda/?page=25&cloudName=CA>
- **Italy** <https://perfsonar.racf.bnl.gov:8443/exda/?page=25&cloudName=IT>
- **US** <https://perfsonar.racf.bnl.gov:8443/exda/?page=25&cloudName=USATLAS>
- **UK** <https://perfsonar.racf.bnl.gov:8443/exda/?page=25&cloudName=UK>

For this LHCONE test phase we have implemented a monitoring page:

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

I recommend all sites in the "Updated/Ready" mode to implement the tests in Step 2). Note that some of the other sites **may** have changes to what is currently shown in the table above. If you do configure tests now, you may need to update them if the information for particular sites change.

I hope that sites can quickly setup the scheduled mesh of network tests required once all sites have completed step 1). I would like to have a goal of getting the mesh tests setup **ASAP**. Once all sites have tests configured and running we can start taking **baseline** data. It would be useful for sites to "capture" status on occasion by making screen shots of monitoring results or logging typical measurement values observed.

There is an open question about what kind of DDM tests are also planned between the proposed LHCONE Early Adopters.

Please send along any comments or suggestions about this information and planning. Also you can directly edit the Twiki but please send Shawn McKee (smckee@umichNOSPAMPLEASE.edu) a brief note when you do so I can keep everyone informed.

-- ShawnMcKee - 12-Dec-2011
 -- JohnShade - 08-Dec-2011

- LHCONE_perfSONAR-PS_latencynode.txt: List of sites and corresponding latency nodes (Used for latency and traceroute configuration). Note SARA instance likely needs updating.

- LHCONE_perfSONAR-bandwidthnode.txt: List of sites and corresponding bandwidth nodes (Used for bandwidth test configuration). Note SARA instance likely needs updating.
 - 20120204-USATLAS-pSPT.pdf: Updated version (Feb 4) Jason Zurawski's perfSONAR-PS maintenance/troubleshooting document
-

This topic: LHCONE > SiteList

Topic revision: r51 - 2018-06-10 - PetrVokac



Copyright &© 2008-2020 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.

Ideas, requests, problems regarding TWiki? Send feedback