

July 22nd 2013, XSEDE Network Performance Tutorial

Jason Zurawski – Internet2/ESnet

Kathy Benninger - Pittsburgh Supercomputing Center

Interpretation of Measurements from Remote Hosts

Outline

- Setup Machines
- Overview of pSPT Functionality

Virtual Machines – Web Interface

- We will be using VMs:
 - perfsonar-ws-2.internet2.edu – perfsonar-ws-10.internet2.edu
 - **Note – Some of you have to share**
- Visit your VM in a web browser first, e.g.:
 - <http://perfsonar-ws-XX.internet2.edu> (where XX is your number)

The screenshot displays the pS-Performance Node web interface. The browser address bar shows the URL `workshop101.kanren.net/toolkit/`. The interface is organized into several sections:

- ps performance toolkit** logo and navigation links.
- 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.
- pS-Performance Node** section containing:
 - Host Information**: Organization Name, Host Location, Administrator Name, Administrator Email.
 - Communities This Host Participates In**: pS-NPToolkit-3.2.1.
 - Host Status**: Primary Address (workshop101.kanren.net), MTU (1500), NTP Status (Synced), Globally registered (No).
 - Services Offered**:
 - Bandwidth Test Controller (BWCTL)**: Running. Includes links for `tcp://[2001:49d0:3c00:2:250:56ff:fe99:2bbf]:4823` and `tcp://workshop101.kanren.net:4823`.
 - Lookup Service**: Running. Includes link for `http://[2001:49d0:3c00:2:250:56ff:fe99:2bbf]:9995/perfSONAR_PS/services/hLS`.

Virtual Machines – Logging In

- Click on “Enabled Services”
 - Note you may need to ‘ok’ a security warning
- Username: “root”
- Password: “psworkshop”

Virtual Machines – Enabling SSH

The screenshot shows a web browser window with the URL https://workshop101.kanren.net/toolkit/admin/enabled_services/. The page title is "Enabled Services Configuration Tool". On the left is a sidebar with the "performance pS toolkit" logo and three main sections: "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), and "Toolkit Administration" (Administrative Information, External BWCTL Limits, External OWAMP Limits, Enabled Services, NTP). The main content area has "Save" and "Reset" buttons at the top. Below is a table of services with checkboxes and descriptions. The "SSH" row is circled in red. At the bottom, there are two filter buttons: "Only Enable Bandwidth Services" and "Only Enable Latency Services", followed by "Save" and "Reset" buttons. The "Save" button at the bottom is also circled in red.

Services	Description
<input checked="" type="checkbox"/> PingER	Enables this host to perform scheduled ping tests. These tests will periodically ping configured hosts giving administrators a view of the latency from their site over time.
<input checked="" type="checkbox"/> perfSONAR-BUOY Throughput Testing	Enables this host to perform scheduled throughput tests. These tests will run periodically giving administrators a view of the throughput to and from their site over time.
<input checked="" type="checkbox"/> perfSONAR-BUOY Latency Testing	Enables this host to perform scheduled one-way latency tests. These tests will run periodically giving administrators a view of the latency from their site over time.
<input checked="" type="checkbox"/> perfSONAR-BUOY Measurement Archive	Makes available the data collected by the perfSONAR-BUOY Latency and Throughput tests.
<input checked="" type="checkbox"/> NDT	Allows clients at other sites to run NDT tests to this host.
<input checked="" type="checkbox"/> NPAD	Allows clients at other sites to run NPAD tests to this host.
<input checked="" type="checkbox"/> BWCTL	Allows clients at other sites to run Throughput tests to this host
<input checked="" type="checkbox"/> OWAMP	Allows clients at other sites to run One-Way Latency tests to this host
<input type="checkbox"/> SSH	Allows administrators to remotely connect to this host using SSH
<input checked="" type="checkbox"/> SNMP MA	Makes available SNMP statistics collected by Cacti (Note: you must configure cacti for this to work)
<input checked="" type="checkbox"/> Traceroute MA	Makes available results of data collected by scheduled traceroute tests
<input checked="" type="checkbox"/> Traceroute Scheduler	Enables this host to run scheduled traceroute tests.
<input checked="" type="checkbox"/> Lookup Service	Registers your services into the global set of perfSONAR services so that they can be discovered

Virtual Machines – Enabling SSH

- Click ‘SSH’ to enable the SSH service
- Click “Save”
 - A progress bar will appear
 - When done “Configuration Saved And Services Restarted” will appear
 - Note: If you are sharing, only one of you will need to make this change
- SSH is now available on your host

Virtual Machines – SSHing to Host

- Open a terminal
- SSH to root@perfsonar-ws-XX.internet2.edu (where XX) is your number):

```
zurawski — root@perfSONAR-ws-1:~ — ssh — ttys000 — 80x24
[zurawski@poteen ~]$ ssh root@perfsonar-ws-1.internet2.edu
The authenticity of host 'perfsonar-ws-1.internet2.edu (207.75.164.236)' can't be established.
RSA key fingerprint is 60:0d:fe:d5:be:72:0a:b3:24:4e:81:69:a0:60:47:10.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added 'perfsonar-ws-1.internet2.edu,207.75.164.236' (RSA) to the list of known hosts.
root@perfsonar-ws-1.internet2.edu's password:
Last login: Tue Mar 12 08:20:03 2013
Welcome to the Internet2 pS-Performance Toolkit v3.2.2

To start configuration, run:

sudo /opt/perfsonar_ps/toolkit/scripts/nptoolkit-configure.py

Once you set passwords, you can login to the web interface and finish configuration. The web interface should be available at:

https://perfsonar-ws-1.internet2.edu/toolkit
[root@perfSONAR-ws-1 ~]#
```


Virtual Machines – Done (For Now)

- Your Virtual Machine is “ready” to go.
 - Diagnostic services are ready to be used
 - Regular testing is scheduled
- If installing @ Home ...
 - Burn a LiveCD (<http://psps.perfsonar.net/toolkit>):
 - Put into bay, and simply boot. This **will not** affect the machine’s hard drive **unless** you configure storage
 - Burn a NetInstall CD (<http://psps.perfsonar.net/toolkit>):
 - Install like a standard Linux distribution.
 - When done you will have the same setup as the VMs.
 - Detailed configuration instructions are also available:
 - <http://code.google.com/p/perfsonar-ps/wiki/pSPerformanceToolkit33>

Outline

- Setup Machines
- Overview of pSPT Functionality

Home Page

The screenshot shows a web browser window titled "pS-Performance Node - pS-Performance Node" with the address bar displaying "http://npw.internet2.edu/toolkit/". The browser's search bar contains "Google". Below the address bar, a navigation bar lists links: MLab, Apple, Yahoo!, Google Maps, YouTube, Wikipedia, News (7937), and Popular. The main content area is titled "pS-Performance Node" and contains several sections:

- User Tools**: A sidebar menu with links to Local Performance Services, Global Performance Services, Java OWAMP Client, Reverse Traceroute, and Reverse Ping.
- Service Graphs**: A sidebar menu with links to Throughput, One-Way Latency, Head Ping Latency, Red PC Ping Latency, Green PC Ping Latency, Blue PC Ping Latency, SNMP Utilization, and Cacti Graphs.
- Performance Toolkit**: A sidebar menu with links to Configuration Help, Frequently Asked Questions, About, and Credits.
- Host Information**: A table with fields for Organization Name, Host Location, Administrator Name, and Administrator Email.
- Communities This Host Participates In**: A table with one row.
- Host Status**: A table with one row showing Primary Address as npw.internet2.edu.
- Services Offered**: A table with three rows:
 - Bandwidth Test Controller (BWCTL)**: Status "Running". Links: tcp://npw.internet2.edu:4823, tcp://ipv6-annarbor-ofc.internet2.edu:4823, tcp://192.168.0.1:4823.
 - Lookup Service**: Status "Disabled". Links: http://npw.internet2.edu:9995/perfSONAR_PS/services/hLS, http://ipv6-annarbor-ofc.internet2.edu:9995/perfSONAR_PS/services/hLS, http://192.168.0.1:9995/perfSONAR_PS/services/hLS.
 - Network Diagnostic Tester (NDT)**: Status "Running". Links: tcp://npw.internet2.edu:3001, http://npw.internet2.edu:7123, tcp://ipv6-annarbor-ofc.internet2.edu:3001, http://ipv6-annarbor-ofc.internet2.edu:7123, tcp://192.168.0.1:3001.

At the bottom left of the page, there is a logo for "perfSONAR powered" and a "perfsONAR toolkit" logo.


Reverse Traceroute

The screenshot shows a web browser window titled "pS-Performance Node - pS-Performance Node" with the URL <http://npw.internet2.edu/toolkit/>. The page features a sidebar with a "performance pS toolkit" logo and two main sections: "User Tools" and "Service Graphs". In the "User Tools" section, "Reverse Traceroute" is highlighted with a red circle. Other tools listed include "Local Performance Services", "Global Performance Services", "Java OWAMP Client", and "Reverse Ping". The "Service Graphs" section lists various performance metrics like "Throughput", "One-Way Latency", and "Head Ping Latency". The main content area is titled "pS-Performance Node" and contains several sections: "Host Information" (with fields for Organization Name, Host Location, Administrator Name, and Administrator Email), "Communities This Host Participates In", "Host Status" (showing Primary Address as npw.internet2.edu), "Services Offered" (listing Bandwidth Test Controller (BWCTL) as Running, Lookup Service as Disabled, and Network Diagnostic Tester (NDT) as Running), and "Performance Toolkit" (with links for Configuration Help, Frequently Asked Questions, About, and Credits). The footer of the page includes the "perfSONAR powered" logo and the "PSC PITTSBURGH SUPERCOMPUTING CENTER" logo.

Reverse Traceroute Use Case

- Use in conjunction with regular traceroute:
 - Workstation -> Server via your terminal
 - Server -> Workstation via the web
- Detect asymmetric routes
- Cautionary notes:
 - Don't rely on traceroute for latency times (e.g. icmp and routers/switches)
 - Asymmetric routing by itself is not harmful, it's a way of life. If one of the routes has a performance problem then **it is** harmful.

Reverse Traceroute Results



Stanford
Linear
Accelerator
Center

traceroute from 192.168.0.1 (npw.internet2.edu) to 164.113.130.1 (www.kanren.net) for 64.134.175.41

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

To perform a traceroute from npw.internet2.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 npw.internet2.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, 164.113.130.1, 140)
traceroute to 164.113.130.1 (164.113.130.1), 30 hops max, 140 byte packets
 1  prodserv-non-rtr.mgmt.internet2.edu (10.165.5.1)  0.449 ms  0.489 ms  0.577 ms
 2  192.122.200.73 (192.122.200.73)  0.644 ms  0.687 ms  0.697 ms
 3  xe-0-1-0x56.aa3.mich.net (198.108.22.137)  0.281 ms  0.273 ms  0.279 ms
 4  xe-0-2-0x22.wsu5.mich.net (198.108.23.51)  1.813 ms  1.845 ms  1.870 ms
 5  v0x1004.rtr.wash.net.internet2.edu (192.122.183.10)  13.729 ms  13.737 ms  13.745 ms
 6  64.57.28.100 (64.57.28.100)  32.533 ms  32.308 ms  32.286 ms
 7  ge-6-2-0.0.rtr.kans.net.internet2.edu (64.57.28.36)  40.835 ms  40.842 ms  40.847 ms
 8  64.57.28.178 (64.57.28.178)  40.933 ms  40.907 ms  40.983 ms
 9  kr-ku-e1-5.kanren.net (164.113.192.54)  41.903 ms  41.964 ms  41.883 ms
10  kr-noc-vlan603.kanren.net (164.113.200.10)  42.226 ms  42.225 ms  42.279 ms
11  kanren.net (164.113.130.1)  41.949 ms !X  41.954 ms !X  41.980 ms !X
```

NDT Tests

The screenshot shows a web browser window titled "pS-Performance Node - pS-Performance Node" with the URL <http://npw.internet2.edu/toolkit/>. The page features a sidebar with navigation links and a main content area with several sections. A red circle highlights the "Network Diagnostic Tester (NDT)" section.

performance toolkit

User Tools

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

Service Graphs

- Throughput
- One-Way Latency
- Head Ping Latency
- Red PC Ping Latency
- Green PC Ping Latency
- Blue PC Ping Latency
- SNMP Utilization
- Cacti Graphs

Performance Toolkit

- Configuration Help
- Frequently Asked Questions
- About
- Credits

pS-Performance Node

Host Information

Organization Name
Host Location
Administrator Name
Administrator Email

Communities This Host Participates In

--

Host Status

Primary Address	npw.internet2.edu
-----------------	-------------------

Services Offered

Bandwidth Test Controller (BWCTL) [\[1\]](#) **Running**

- [tcp://npw.internet2.edu:4823](http://npw.internet2.edu:4823)
- [tcp://ipv6-annarbor-ofc.internet2.edu:4823](http://ipv6-annarbor-ofc.internet2.edu:4823)
- [tcp://192.168.0.1:4823](http://192.168.0.1:4823)

Lookup Service [\[1\]](#) **Disabled**

- http://npw.internet2.edu:9995/perfSONAR_PS/services/hLS
- http://ipv6-annarbor-ofc.internet2.edu:9995/perfSONAR_PS/services/hLS
- http://192.168.0.1:9995/perfSONAR_PS/services/hLS

Network Diagnostic Tester (NDT) [\[1\]](#) **Running**


- [tcp://npw.internet2.edu:3001](http://npw.internet2.edu:3001)
- <http://npw.internet2.edu:7123>
- [tcp://ipv6-annarbor-ofc.internet2.edu:3001](http://ipv6-annarbor-ofc.internet2.edu:3001)
- <http://ipv6-annarbor-ofc.internet2.edu:7123>
- [tcp://192.168.0.1:3001](http://192.168.0.1:3001)

perfSONAR powered

NDT = Network Diagnostic Tool

- Perfect for the end user to run from their laptop
 - Convince your campus users to test to this, instead of 'speedtest.net, or FCC'. Install one on **every** campus
 - Runs client to server and reverse
 - Detects common problems
- Caveats
 - NDT heuristics are a little dated (e.g. may tell you your connection is a 'T1' or similar). These are not that important though
 - Detects common problems like queueing, loss, duplication, SACKs, duplicate ACKs. Won't tell you 'where' these are happening, so note it's end to end.
 - Uses the Web100 patched kernel on the server end. If this is disabled for any reason (e.g. initial release perfSONAR update or security update that installs a new kernel), NDT will not work.

NDT Home Screen



NPAD Home

The screenshot shows a web browser window titled "pS-Performance Node - pS-Performance Node" with the address bar showing "http://npw.internet2.edu/toolkit/gui/services/". The page features a sidebar on the left with a "Performance Toolkit" section containing links for "Configuration Help", "Frequently Asked Questions", "About", and "Credits". Below this is the "perfSONAR powered" logo. The main content area lists several services with their status and endpoints:

Service	Status	Endpoints
Head Ping Latency		
Red PC Ping Latency		
Green PC Ping Latency		
Blue PC Ping Latency		
SNMP Utilization		
Cacti Graphs		
Lookup Service	Disabled	<ul style="list-style-type: none">tcp://npw.internet2.edu:4823tcp://ipv6-annarbor-ofc.internet2.edu:4823tcp://192.168.0.1:4823
Network Diagnostic Tester (NDT)	Running	<ul style="list-style-type: none">tcp://npw.internet2.edu:3001http://npw.internet2.edu:7123tcp://ipv6-annarbor-ofc.internet2.edu:3001http://ipv6-annarbor-ofc.internet2.edu:7123tcp://192.168.0.1:3001http://192.168.0.1:7123
Network Path and Application Diagnosis (NPAD)	Running	<ul style="list-style-type: none">tcp://npw.internet2.edu:8100http://npw.internet2.edu:8000tcp://ipv6-annarbor-ofc.internet2.edu:8100http://ipv6-annarbor-ofc.internet2.edu:8000tcp://192.168.0.1:8100http://192.168.0.1:8000
One-Way Ping Service (OWAMP)	Running	<ul style="list-style-type: none">tcp://npw.internet2.edu:861tcp://ipv6-annarbor-ofc.internet2.edu:861tcp://192.168.0.1:861
perfSONAR-BUOY Regular Testing (Throughput)	Running	

NPAD = Network Path Diagnostics

- Similar to NDT, but spits out a lot more
 - Tries to “simulate” network use based on current conditions.
 - E.g. expectations are a latency and a speed, will tell you if this is or is not possible
- Also relies on Web100, same caveats apply
- “Mom Feature” ... you can retrieve the HTML results from past tests. Users can mail these in trouble ticket reports or network engineers can access the results directly on the server.

NPAD Start – Pick some Expectations

The screenshot shows a web browser window titled "NPAD Diagnostics" with the URL <http://npw.internet2.edu:8000/>. The browser's address bar and tabs are visible. The main content area contains instructions for using the NPAD Diagnostics tool. It includes a list of bullet points explaining the test process, such as having an end-to-end application performance goal, entering parameters, and clicking "Start Test". Below the instructions, there is a section titled "Test from server npw.internet2.edu to this machine" with input fields for "Round Trip Time (msec)" (set to 30) and "Target Rate (Mbps)" (set to 5). A "Start Test" button is next to these fields. Below the input fields is a "Log:" section with a text area for the test results. At the bottom of the page, there is a section titled "Command line client" with a note about using the command line diagnostic client if the Java applet shows errors.

NPAD Diagnostics

<http://npw.internet2.edu:8000/>

MLab Apple Yahoo! Google Maps YouTube Wikipedia News (7937) Popular

pS-Performance Node – pS-Perfo... NDT server NPAD Diagnostics

The test results are most accurate over a short network path. If this server is not near you, look for a closer server from the list of [Current NPAD Diagnostic Servers](#).

- Have an end-to-end application performance goal ([target round-trip time](#) and [target data rate](#)) in mind. Enter the parameters on the form below and click **Start Test**. Messages will appear in the log window as the test runs, followed by a diagnostic report.
- In the diagnostic report, failed tests (in red) indicate problems that will prevent the application from meeting the end-to-end performance goal. For each message, a question-mark link (?) leads to additional detailed information about the results.
- Every test is fully logged (including your IP address) and test results are [public](#). We use the logs and results to further refine the software.

For more information, see the [NPAD Documentation](#), especially the sections:

- [NPAD Diagnostic Procedure](#) - the full instructions.
- [Theory and Method](#) - why the the tests work.
- [Outcomes](#) - what to do next in the broader debugging context.

Note that tests take 2-5 minutes, depending on the parameters that you provide and the network path. If there is a queue, waiting times might be long.

Test from server npw.internet2.edu to this machine

Round Trip Time (msec): [Start Test](#)

Target Rate (Mbps):

Log:

Command line client

If the Java applet above exhibits errors or the form is blank, try the command line diagnostic client.

Applet DiagClient started

NPAD Results

Test Results

http://npw.internet2.edu:8000/ServerData/Reports-2011-04/64.134.175.41:2011-04-05-12:46:01.html

MLab = Apple Yahoo! Google Maps YouTube Wikipedia News (7937) Popular

pS-Performance Node - pS-Perfo... NDT server Test Results

Test conditions

Tester: npw.internet2.edu (207.75.164.109) [?]
Target: 64.134.175.41 (64.134.175.41) [?]
Logfile base name: 64.134.175.41:2011-04-05-12:46:01 [?]
This report is based on a 5 Mb/s target application data rate [?]
This report is based on a 30 ms Round-Trip-Time (RTT) to the target application [?]
The Round Trip Time for this path section is 24.000000 ms.
The Maximum Segment Size for this path section is 1344 Bytes. [?]

Target host TCP configuration test: Pass! [?]

TCP negotiated appropriate options: WSCALE=3, SACKok, and Timestamps. [?]
The target passed all tests! See tester caveats: [?]

Path measurements [?]

The path to the tester is too long for accurate measurements.
> *Test a shorter path section or reduce the target data rate and/or RTT.* [?]

Data rate test: Fail! [?]

The maximum data rate was 1.729524 Mb/s. [?]
This is below the target rate (5.000000 Mb/s). [?]
This test did not complete due to other problems with the path, target or tester.
> *Correct other problems first, and then rerun this test.* [?]

Loss rate test: Pass! [?]

Pass: zero losses in 2847 packets, loss rate less than 0.035125%. [?]
FYI: To get 5 Mb/s with a 1344 byte MSS on a 30 ms path the total end-to-end loss budget is 0.251889% (397 packets between losses). [?]

Suggestions for alternate tests

FYI: This path may pass with a less strenuous application: [?]
Try rate=1 Mb/s, rtt=232 ms
Or if you can raise the MTU: [?]
Try rate=1 Mb/s, rtt=1554 ms, mtu=9000 bytes

Network buffering test: Pass! [?]

Cacti

pS-Performance Node – pS-Performance Node For KanREN In 1102 Grand, Kansas City

https://ps-data.kanren.net/toolkit/gui/services/

performance **ps** toolkit

User Tools

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

Service Graphs

- Throughput
- One-Way Latency
- WSU Ping Latency
- KSU Ping Latency
- KU Ping Latency
- 1102 Grand Ping Latency
- OWAMP Jitter
- Cacti Graphs**

Toolkit Administration

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

pS-Performance Node For KanREN In 1102 Grand, Kansas City

Host Information

Organization Name	KanREN
Host Location	1102 Grand, Kansas City
Administrator Name	KanREN NOC
Administrator Email	noc@kanren.net

Communities This Host Participates In

KanREN

Host Status

Primary Address	ps-data.kanren.net
-----------------	--------------------

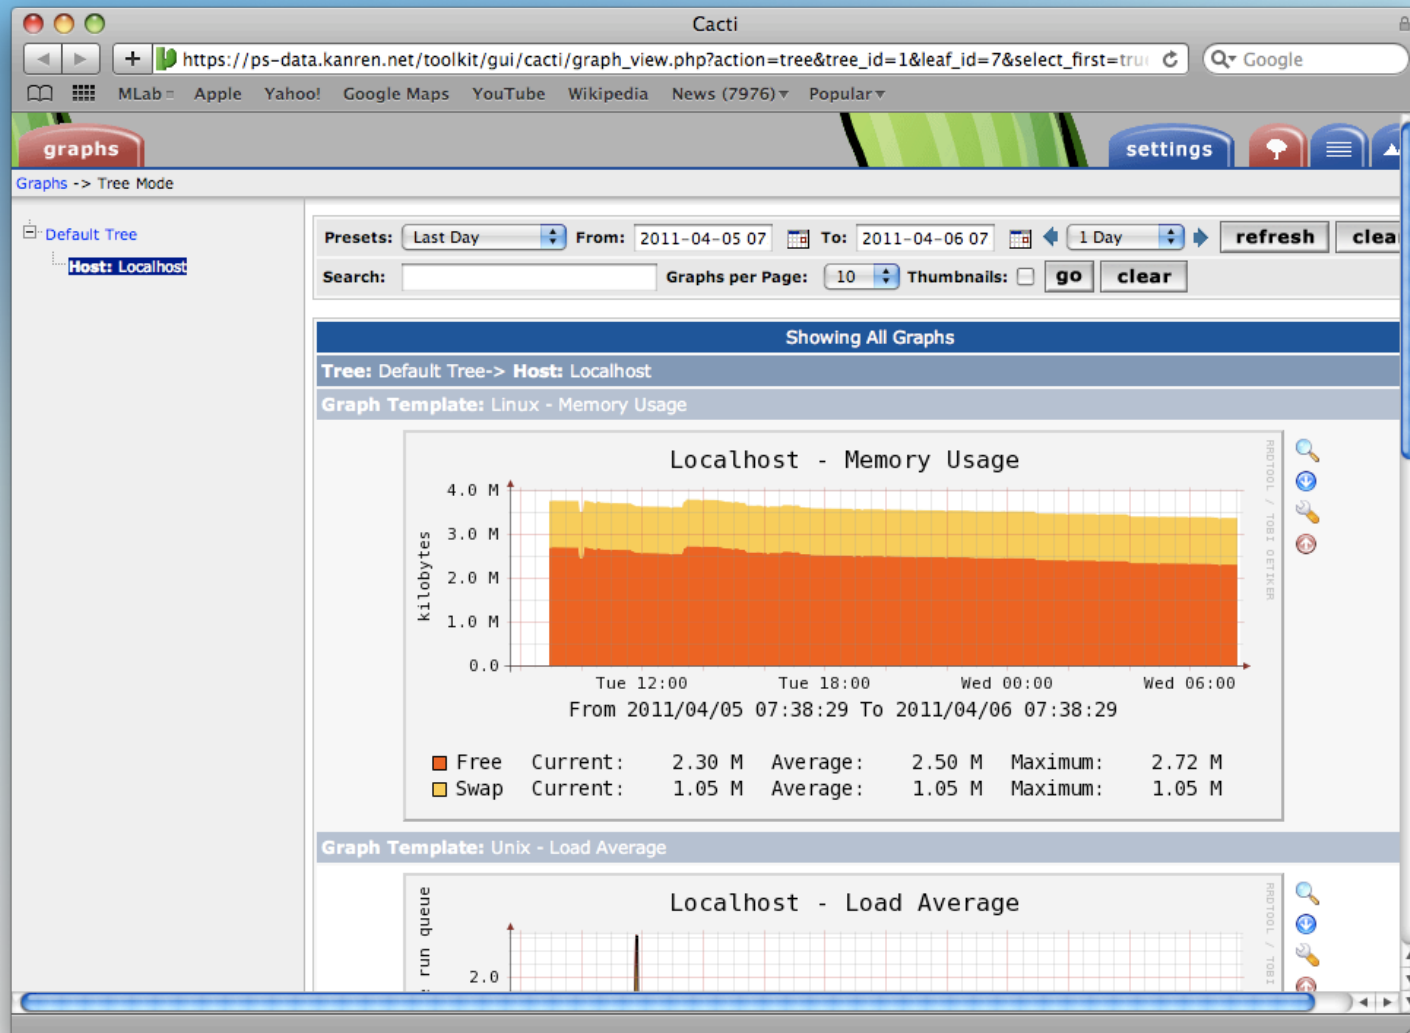
Services Offered


Bandwidth Test Controller (BWCTL)	Disabled
• tcp://ps-data.kanren.net:4823	
Lookup Service	Running
• http://ps-data.kanren.net:9995/perfSONAR_PS/services/hLS	
Network Diagnostic Tester (NDT)	Disabled
• tcp://ps-data.kanren.net:3001 • http://ps-data.kanren.net:7123	
Network Path and Application Diagnosis (NPAD)	Disabled
• tcp://ps-data.kanren.net:8100	

Using Cacti

- Currently only monitoring the localhost
 - Memory/cpu things
 - Plans to add NTP monitoring as well as some more advanced localhost things (process breakdown, etc.)
- Can be used for SNMP monitoring if desired
 - Note the poll is defaulted to about 5 minutes

Cacti Display





Interpretation of Measurements from Remote Hosts

July 22nd 2013, XSEDE Network Performance Tutorial

Jason Zurawski – Internet2/ESnet

Kathy Benninger - Pittsburgh Supercomputing Center

For more information, visit <http://www.internet2.edu/workshops/npw>