Scaling the PuNDIT Project for Wide Area Deployments

Jorge Batista, Gabriele Carcassi, Constantine Dovrolis,
Danny Lee and Shawn McKee (Presenting)

Overview

PuNDIT is a NSF SSI project which uses perfSONAR data to identify and localize network problems

PuNDIT is collaborating with ESnet, Internet2, OSG and WLCG

Meta goal is to meet an often expressed need to **automate** watching/analyzing perfSONAR metrics to inform users/site-admins when there are real network problems they should address.

- See further details at http://pundit.gatech.edu
- User GUI mock-up http://punditui.aglt2.org/

Top Level Architecture

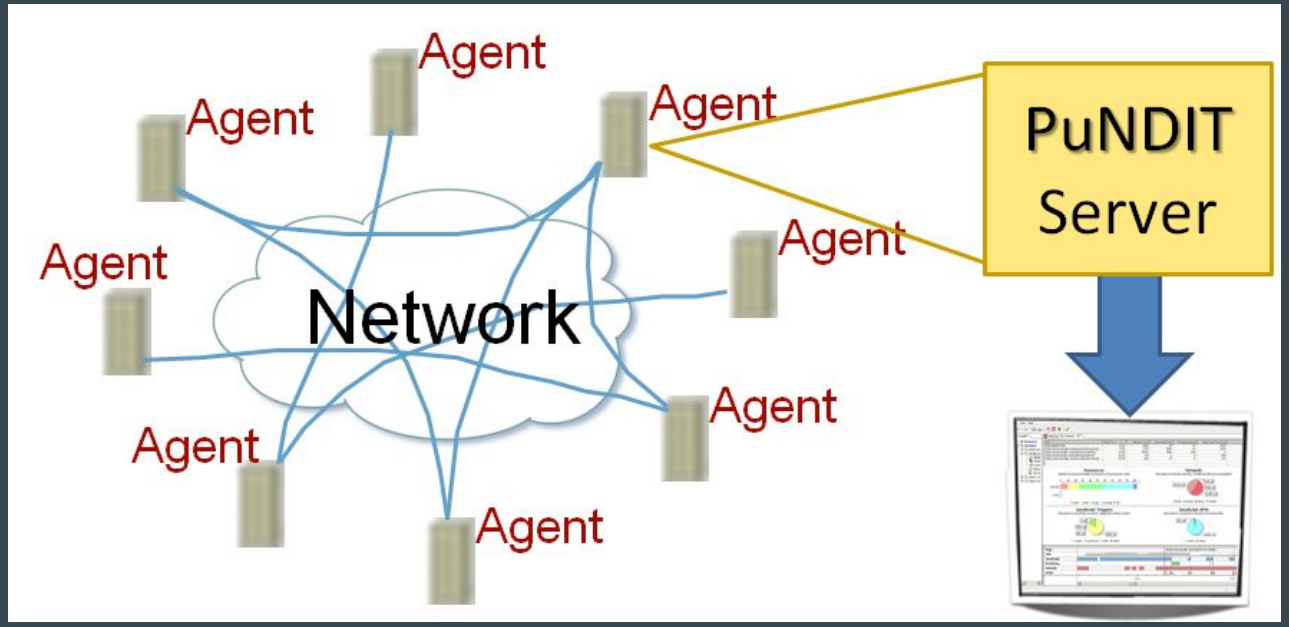

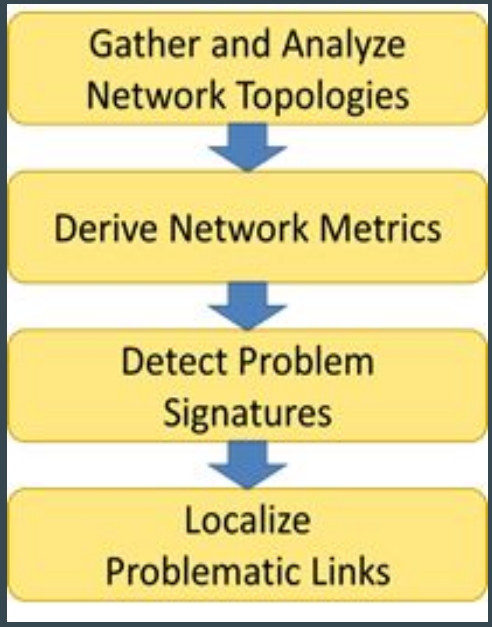

- **perfSONAR** provides the base measurement infrastructure
  - Collects network metrics like latency, loss and reordering
  - Collects topological information
- A lightweight PuNDIT process installed on each host performs detection (**PuNDIT** Agent)
- A central server holds event repository and runs localization algorithm (**PuNDIT** Server)

Distributed (Agent) Processing

- Agents gather perfSONAR measurements and perform statistical analysis
- Problem Detection is run, table on right shows types
- Summarized results are computed and continually sent to the central server
- Raw data retained on agent host if problem detected
    - This allows more detailed problem exploration for either improving PuNDIT or digging into specific problems that were found

Problems Detected

Excessive Delays, Losses and Reordering	When these parameters in a 5 second window exceed user-specified thresholds
Congestion	Delay increases correlated with losses, indicating possible congestion at an interface
Route Change	A change in route resulting in a sudden and significant change in delay measurements
Route Instability	Repeated route changes, which results in poor TCP performance
End-host Context Switch	Context switching of the active measurement applications on an end-host, which may cause measurement fluctuations to be erroneously detected as problems

Scalable Message Queuing with RabbitMQ

- One challenge for PuNDIT is reliably getting data shipped back to the central PuNDIT server for correlation analysis.
  - Must be robust and very low impact on the Toolkit node
- We chose to use an Advanced Message Queuing Protocol (AMQP) and selected RabbitMQ as our implementation with the following components/features
  - Exchange - A receiver/supplier of messages. Messages are always sent to exchanges, never directly to queues. Exchanges send messages to queues.
  - Binding - A relationship between an exchange and a queue. Queues are bound to exchanges by routing rules.
  - Queue - A buffer for storing messages (think of queues as named mailboxes).
  - Very robust and highly scalable

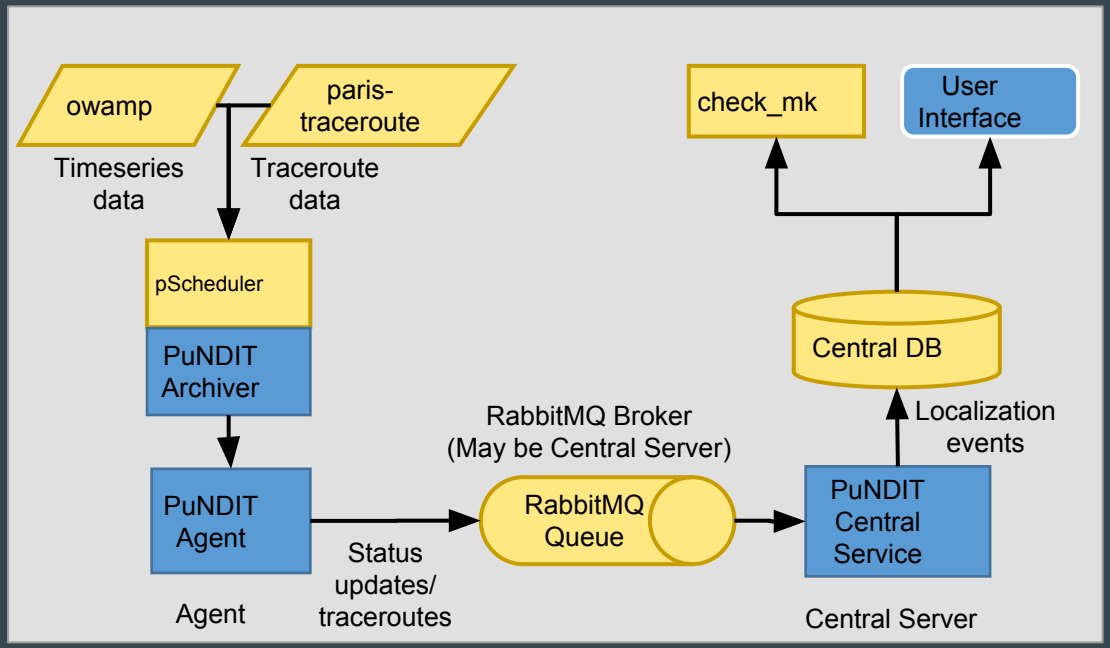

PuNDIT Agent CPU Load Analysis

- One concern that PuNDIT users may have is the impact of running the PuNDIT agent on their perfSONAR toolkits
  - We must prove the PuNDIT agent doesn't adversely impact the Toolkit's ability to make its measurements
- To analyze the impact we are using a ten node perfSONAR test-bed which is also used to test new releases of perfSONAR
  - Each node has the PuNDIT client (agent), Paris traceroute and RabbitMQ packages installed
  - All report to our PuNDIT server
  - Wide range of nodes in terms of memory, CPU and disk (very old to relatively new)
  - Running `atop` to monitor resource usage
  - Logging CPU, memory, disk, network loads

PuNDIT Agent CPU Load Analysis Results

Test Hosts	Location	Processor	CPUs	Speed GHz	Memory GB	PuNDIT CPU Load
psum05.aglt2.org	Ann Arbor, MI	Xenon X5355	8	2.66	16	2%
psmsu05.aglt2.org	Lansing, MI	Opteron 270	4	1.99	4	~90%
psum09.cc.gt.atl.ga.us	Atlanta, GA	Xenon X5355	8	2.66	16	0%
netmon1.atlast-swt2.org	Arlington, TX	Opteron 2216	4	2.39	8	2%
iut2-net3.iu.edu	Bloomington, IN	Opteron 2216	4	2.61	8	2%
uct2-net3.mwt2.org	Chicago, IL	Opteron 275	4	2.19	8	2%
personar.unl.edu	Lincoln, NE	Xeno E5620	8	2.40	12	2%
ps3.ochep.ou.edu	Norman, OK	Pentium E2220	2	2.40	2	2%
nettest.lbl.gov	Berkeley, CA	Xeon E3-1275	8	3.50	32	1%
Ps-development.bnl.gov	Uptin, NY	Xeon X5650	24	2.67	24	2%

Most nodes show negligible load from running PuNDIT

However one node (psmsu05) showed significant load and we investigated further

Agent CPU Load Analysis Results

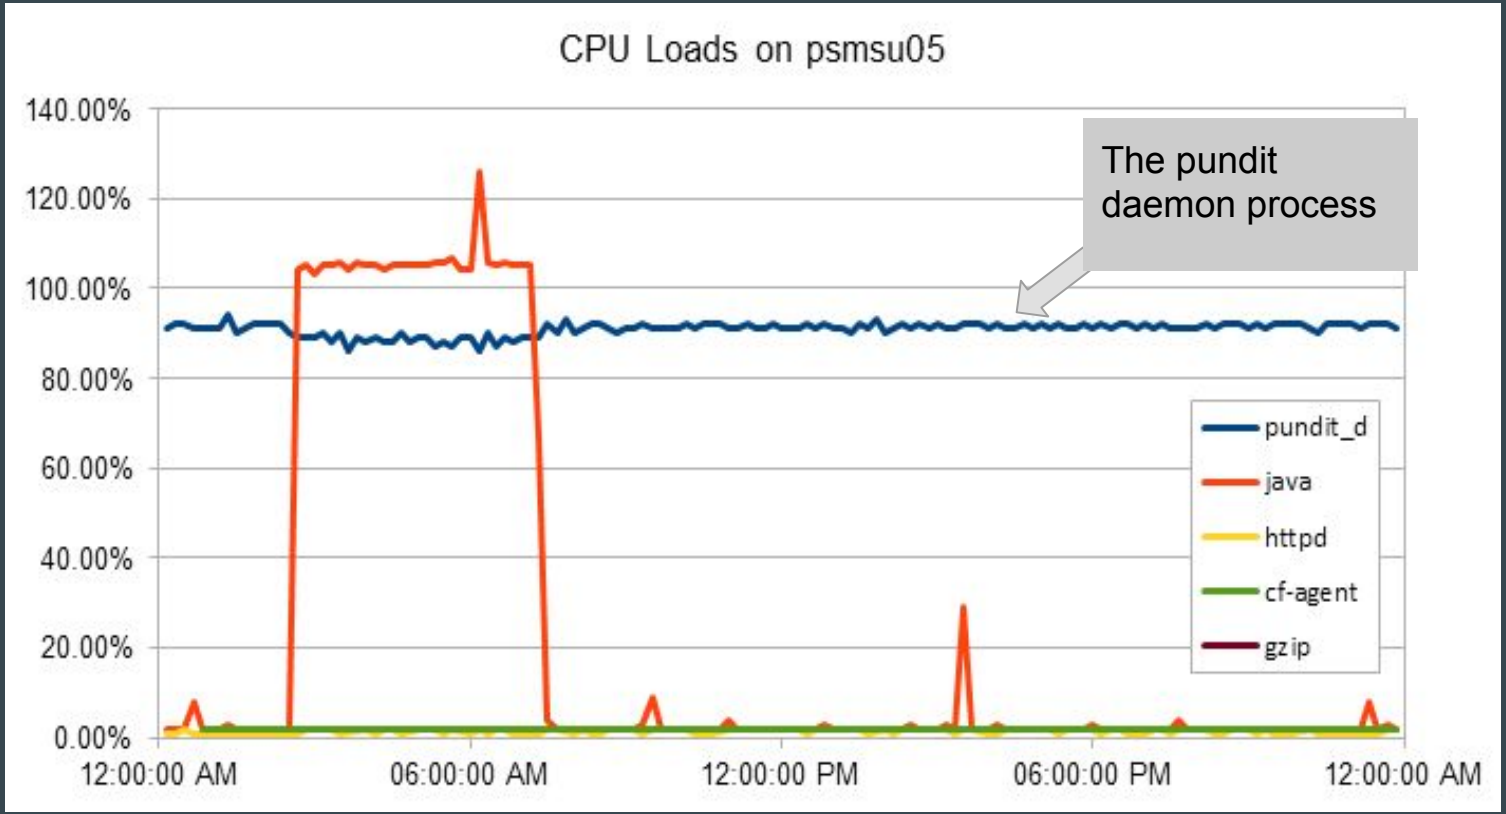

PuNDIT Agent CPU Load Analysis Results

- CPU load 2% or less for “normal” operating conditions
- The very high load on psmsu05 was caused by a file deletion failure
  - PuNDIT agent in perfSONAR 3.5 required patch to the Toolkit and had to clean-up (delete) files
  - After the delete problem was discovered and fixed psmsu05 also ran at ~2% CPU load
- New **perfSONAR 4.0** module (pScheduler) allows subscription to any data
  - Using pScheduler, PuNDIT no longer has to handle file deletions
- Disk, memory, and network loads also neglible for the scale of the full-mesh testing in the testbed (10-nodes)
  - Plans to test at larger scale during beta release sometime later this fall.

The PuNDIT Central Server: Centralized Algorithms

One of the most powerful features of **PuNDIT** is its ability to “see” many network paths and look for problem correlations by having a **central PuNDIT server**

- The central server is installed via RPM on a RHEL-compatible system
- Inputs: traceroute information and problem reports from all participating Toolkits
- Two kinds of tomography algorithms are used to localize problems
  - Boolean Tomography: Loss Problems
    - Returns True or False for possible interfaces
  - Range Tomography: Delay Problems
    - Returns an estimated range for delays in problematic interfaces
  - Algorithms are lightweight
- Multithreaded analysis implementation for parallel processing

PuNDIT GUI: User Interface Example

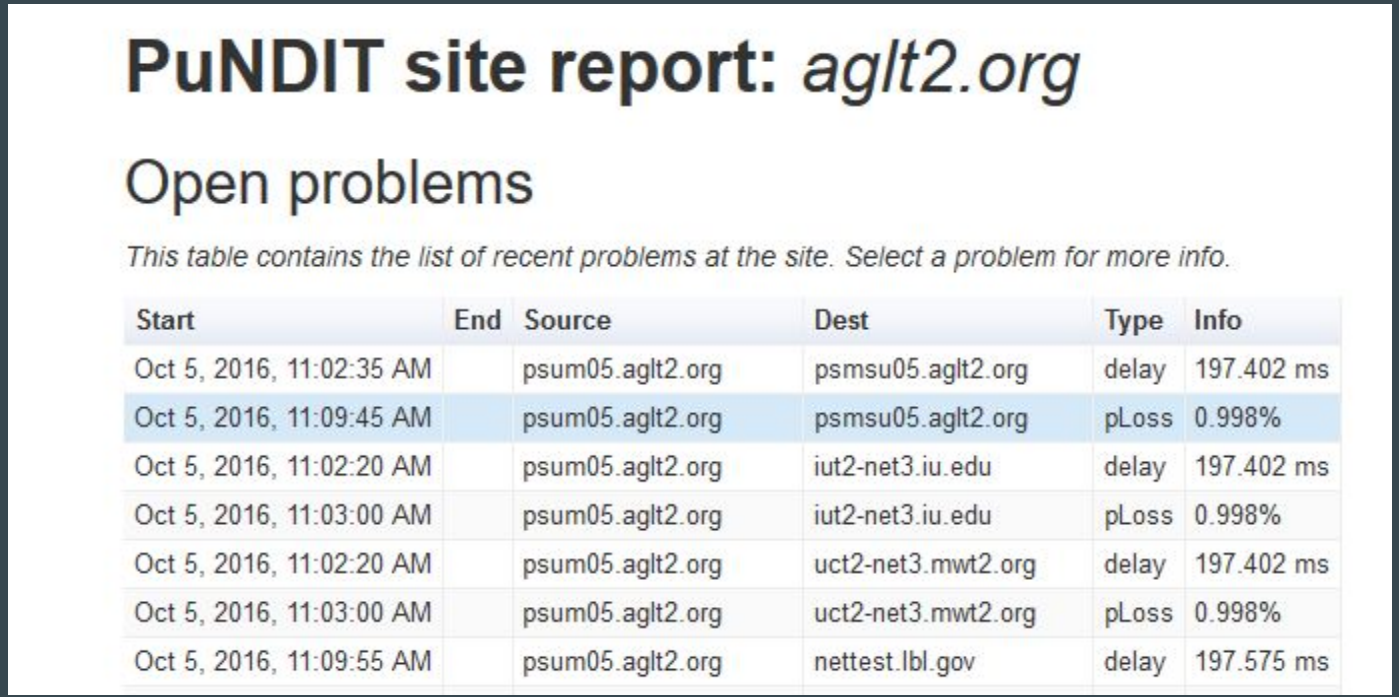

PuNDIT site report: *aglt2.org*

Open problems

This table contains the list of recent problems at the site. Select a problem for more info.

Start	End	Source	Dest	Type	Info
Oct 5, 2016, 11:02:35 AM		psum05.aglt2.org	psmsu05.aglt2.org	delay	197.402 ms
Oct 5, 2016, 11:09:45 AM		psum05.aglt2.org	psmsu05.aglt2.org	pLoss	0.998%
Oct 5, 2016, 11:02:20 AM		psum05.aglt2.org	iut2-net3.iu.edu	delay	197.402 ms
Oct 5, 2016, 11:03:00 AM		psum05.aglt2.org	iut2-net3.iu.edu	pLoss	0.998%
Oct 5, 2016, 11:02:20 AM		psum05.aglt2.org	uct2-net3.mwt2.org	delay	197.402 ms
Oct 5, 2016, 11:03:00 AM		psum05.aglt2.org	uct2-net3.mwt2.org	pLoss	0.998%
Oct 5, 2016, 11:09:55 AM		psum05.aglt2.org	nettest.lbl.gov	delay	197.575 ms

http://punditui.aglt2.org/site-report.html?site=aglt2.org

PuNDIT GUI: User Interface Example

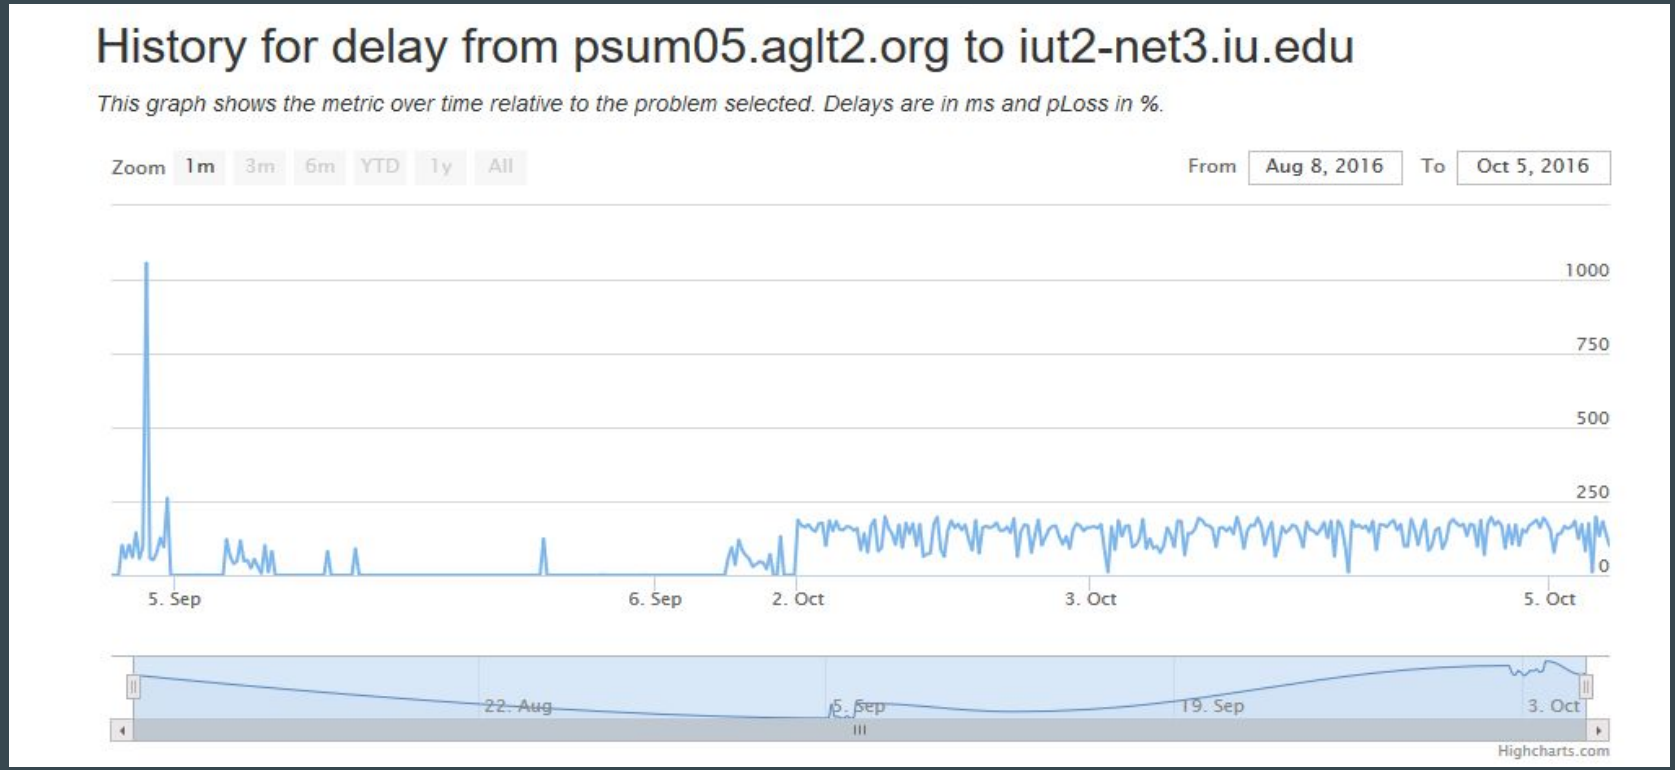

http://punditui.aglt2.org/site-report.html?site=aglt2.org

PuNDIT GUI: User Interface Example

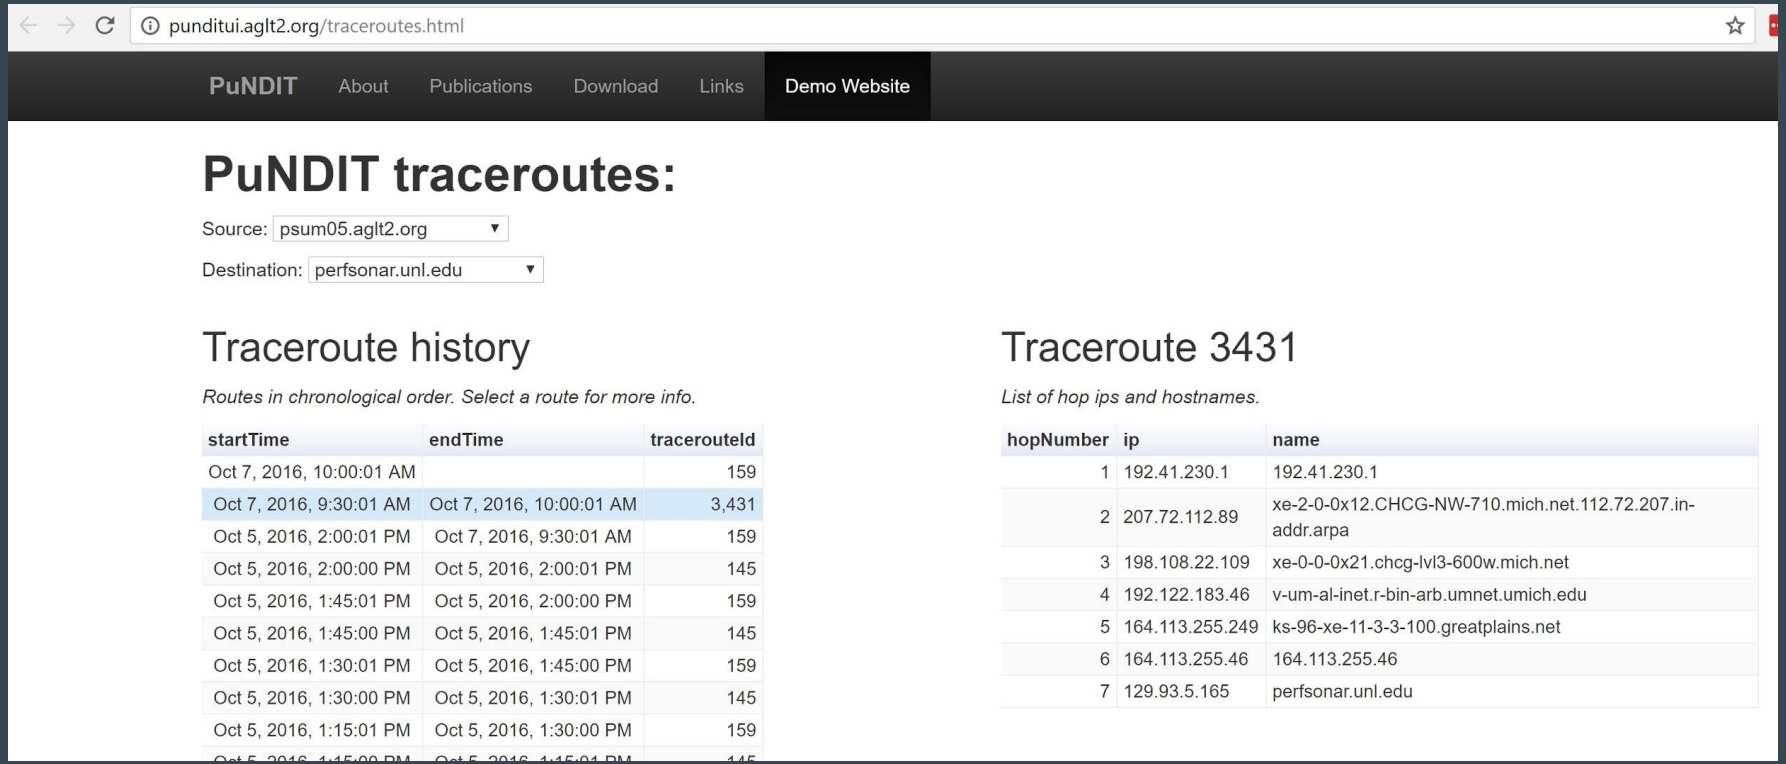

http://punditui.aglt2.org/traceroutes.html

Timeline and Plans

We are working on finalizing the first production release of the PuNDIT agent

- Integration with the new perfSONAR **pScheduler** has been a focus
  - Removes the need to patch the perfSONAR Toolkit; we just subscribe to the data we need
- Squashing a few bugs that our testing identifies is part of the process

In parallel we are working on the Central PuNDIT server

- Optimizing the schemas used to store the problems, net-topology and interface to the user GUI
- Tweaks and enhancements of the user GUI to more easily view and understand problems

Goal is to have two RPMS available shortly after the release of perfSONAR 4.0

- One RPM to be installed on each perfSONAR Toolkit, one for central services
- Packaging and testing are ongoing now...

Conclusion

- **We are building a distributed network problem identification and localization system with scalability in mind**
- Agents perform distributed (local) analysis and send summaries to a central server
  - Minimize amount of data being transferred
  - Low CPU and Network loads from PuNDIT Agents
- Uses scalable message queueing techniques to reliably transport data
- Central server receives all reports and topology and localizes problems
- Web GUI provides user visibility and interaction with the central PuNDIT server
- Targeting a release shortly after perfSONAR Toolkit 4.0 is out

Questions?

Backup Slides

Agent CPU Load Analysis Results

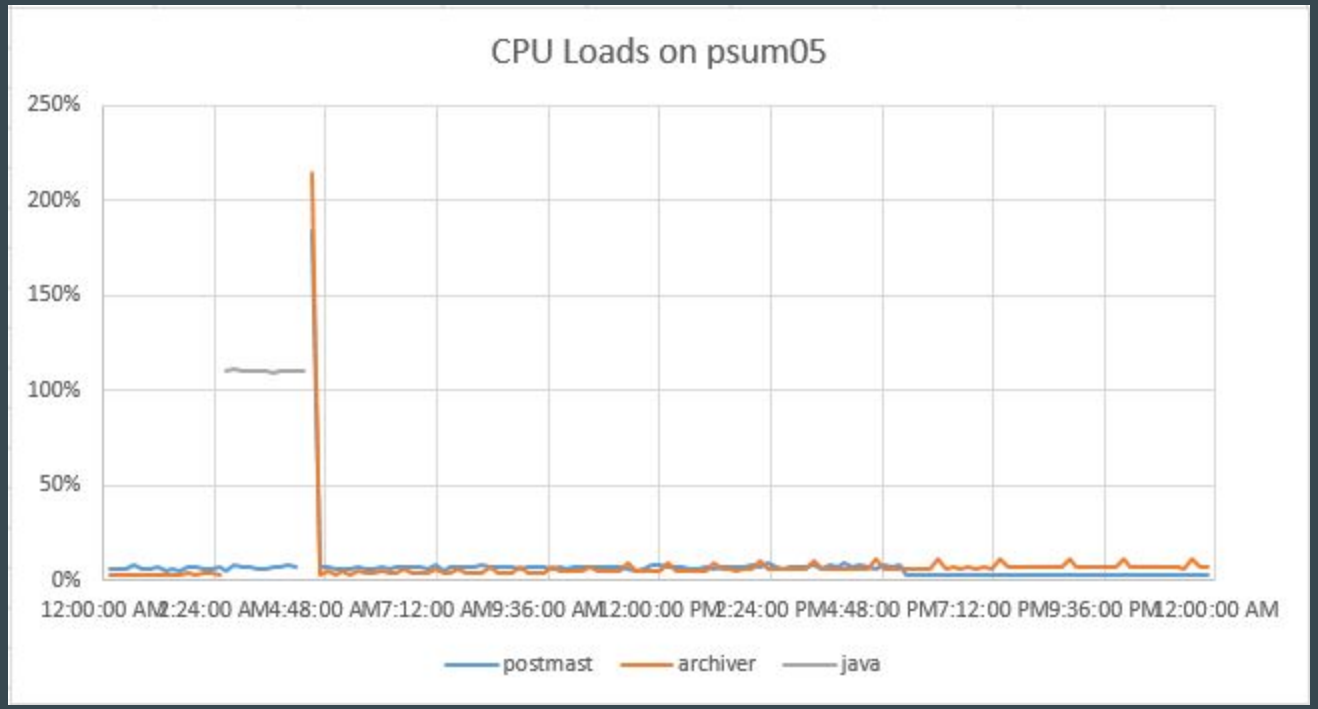

Central Server CPU Load Analysis Results

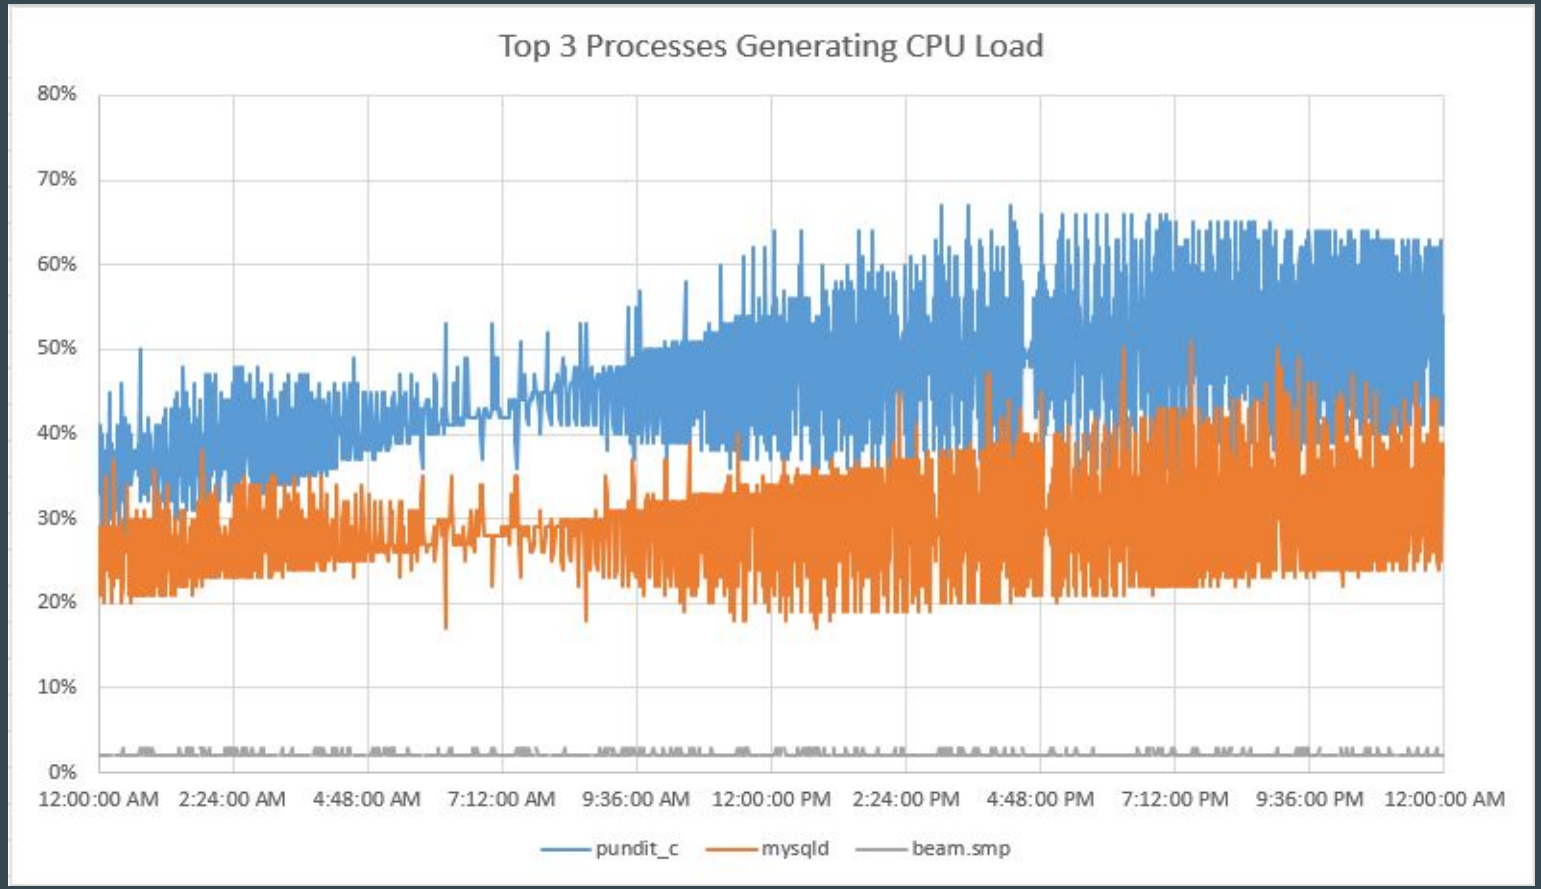

Central Server CPU Load Analysis Results

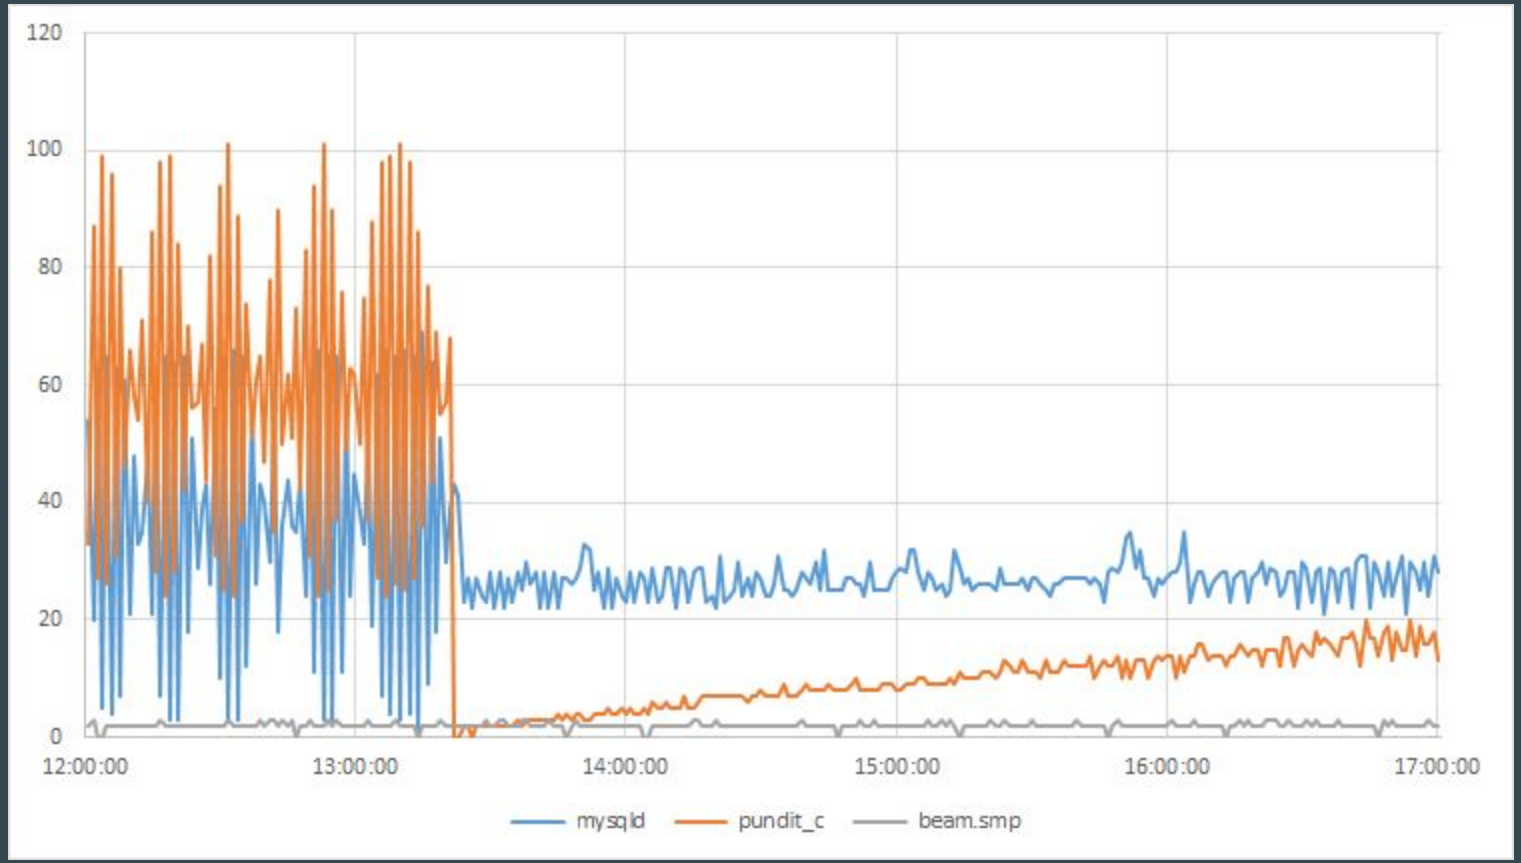

Central Server Disk Load

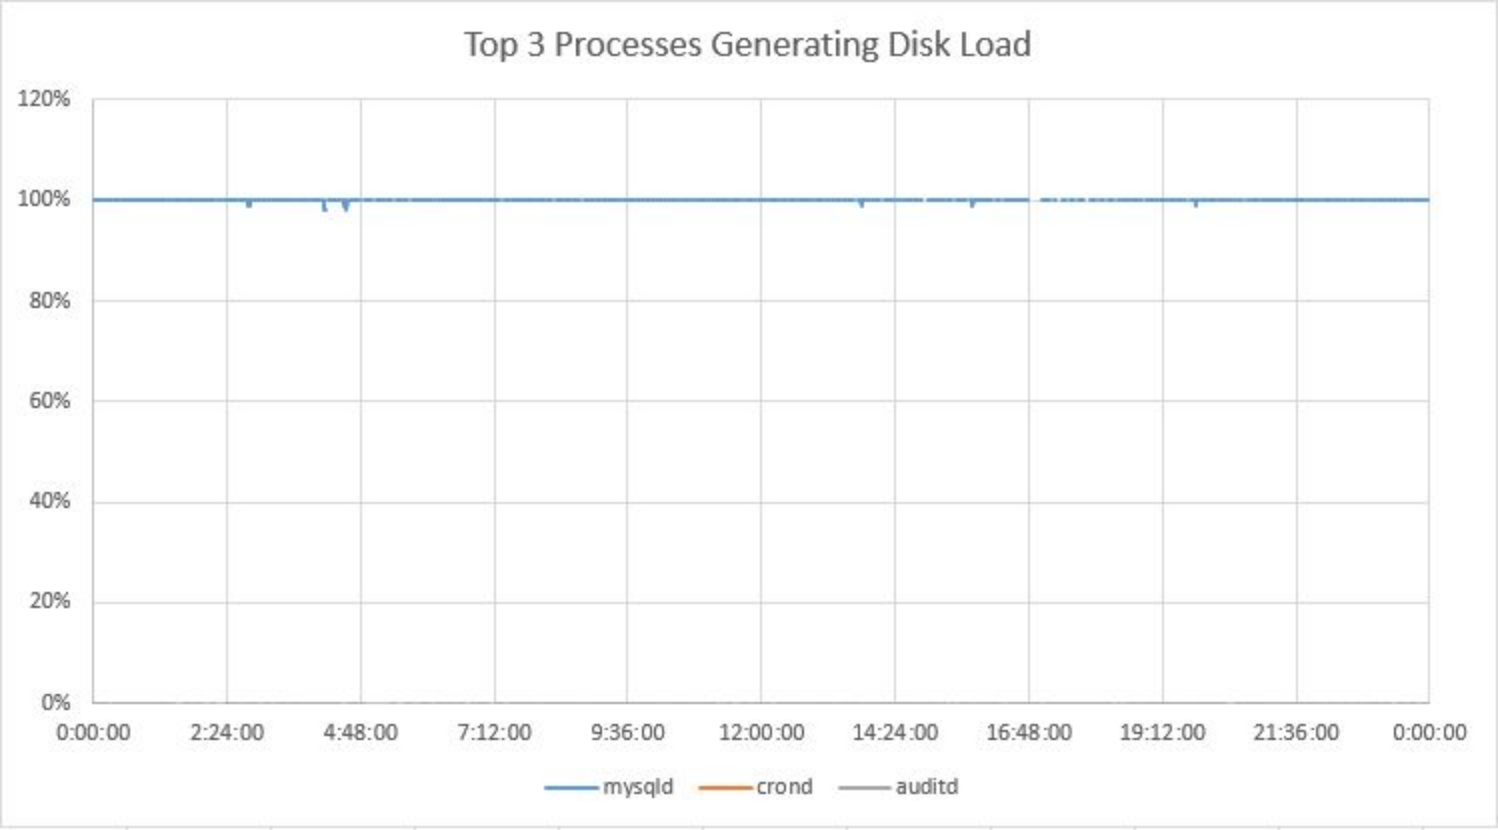

Central Server Memory Load

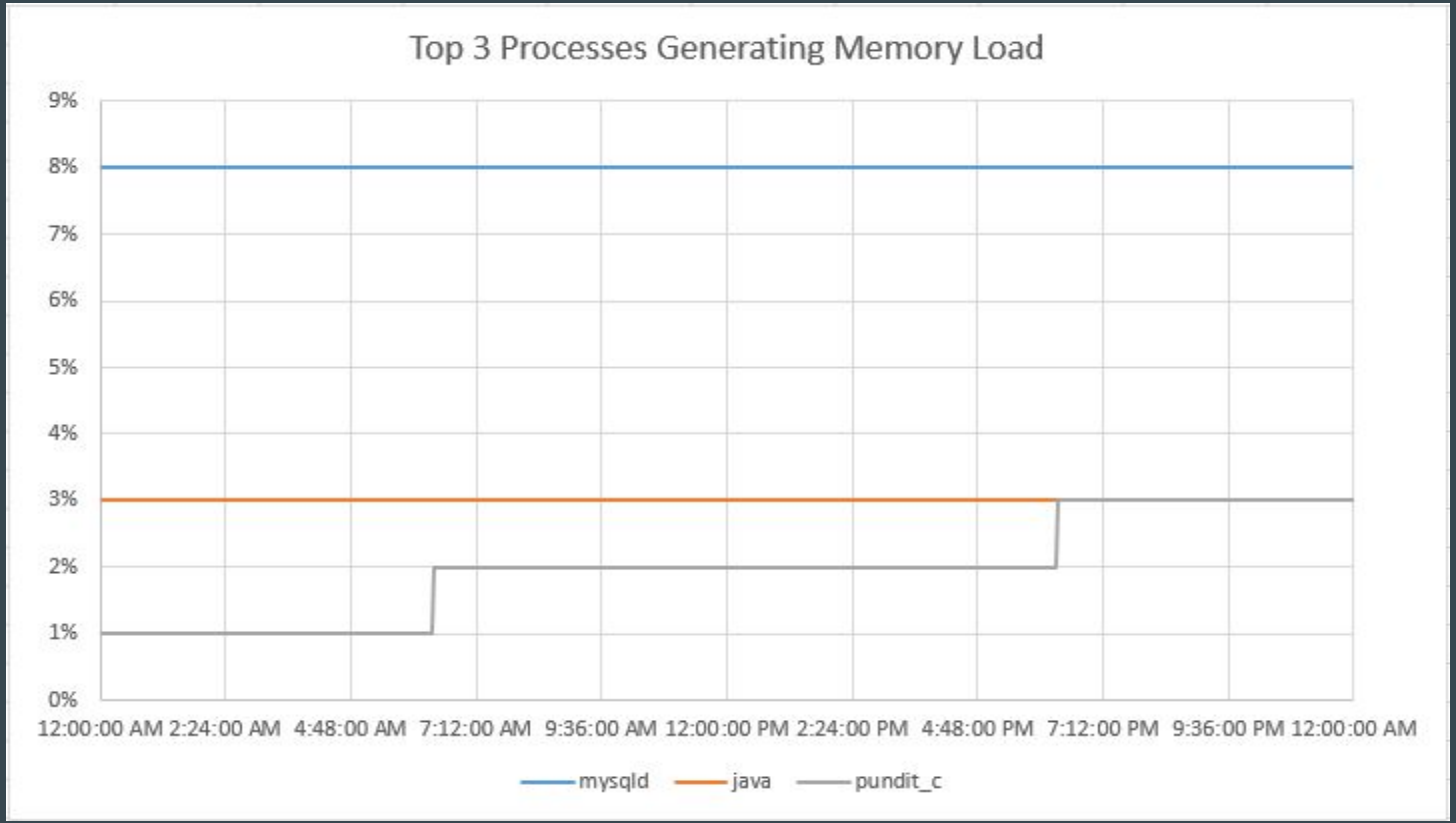

Central Server Network Load

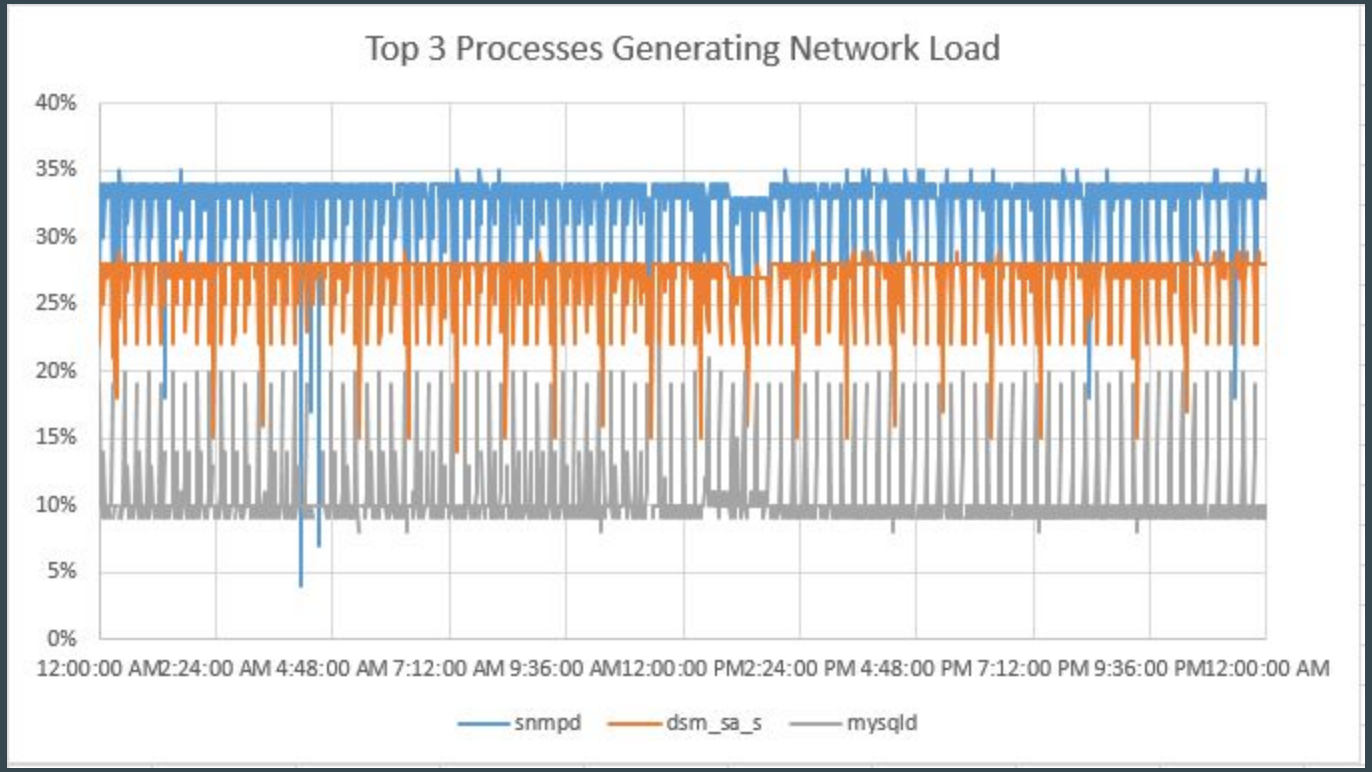

CPU Impact of varying number of problems

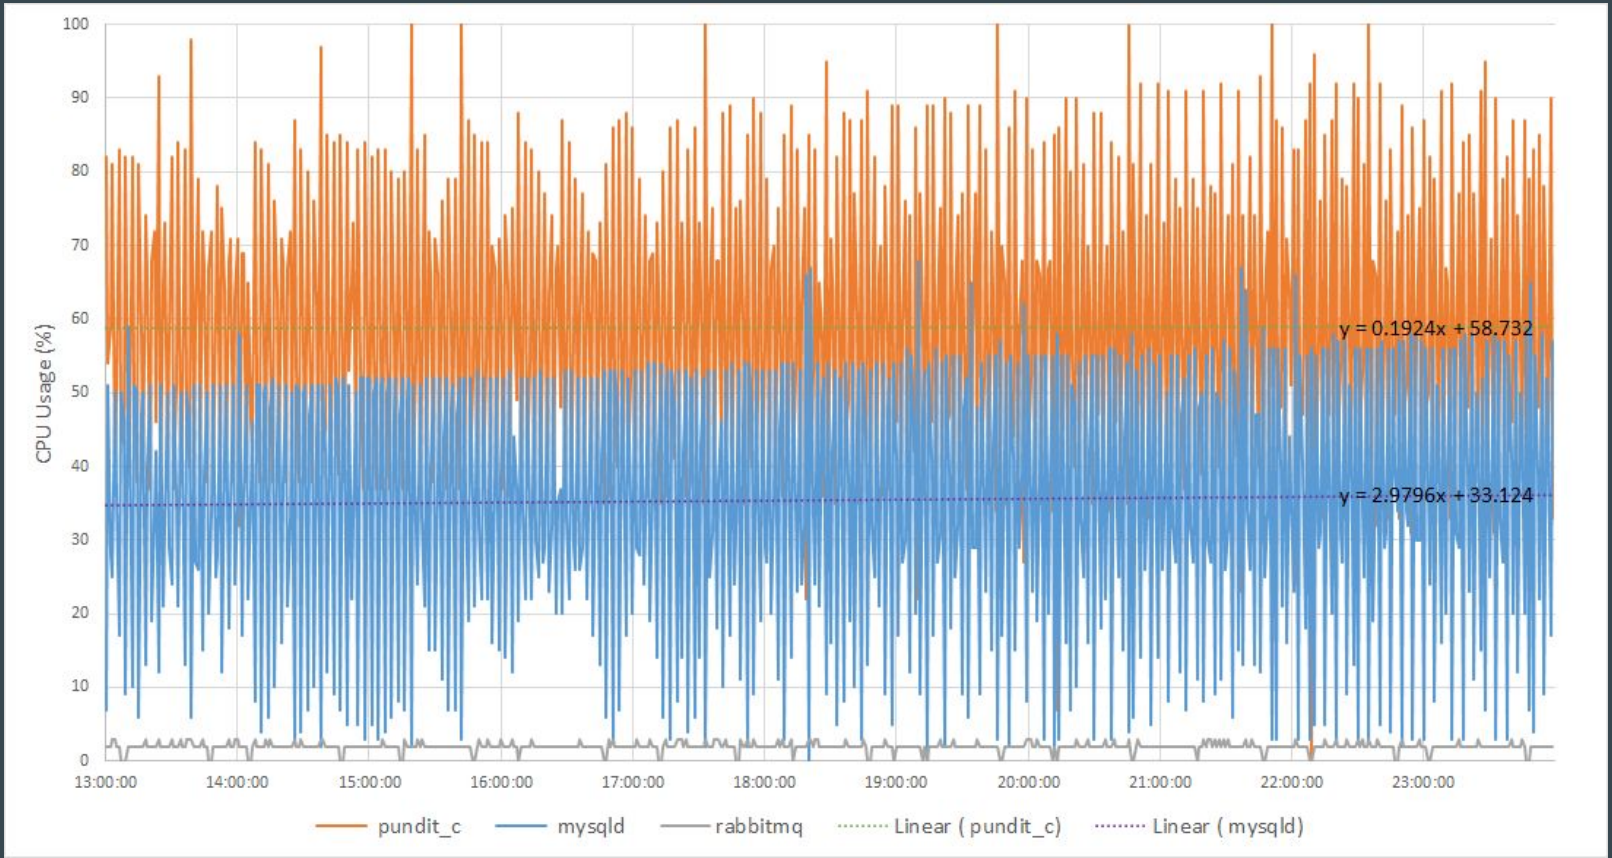

CPU Impact of varying number of problems

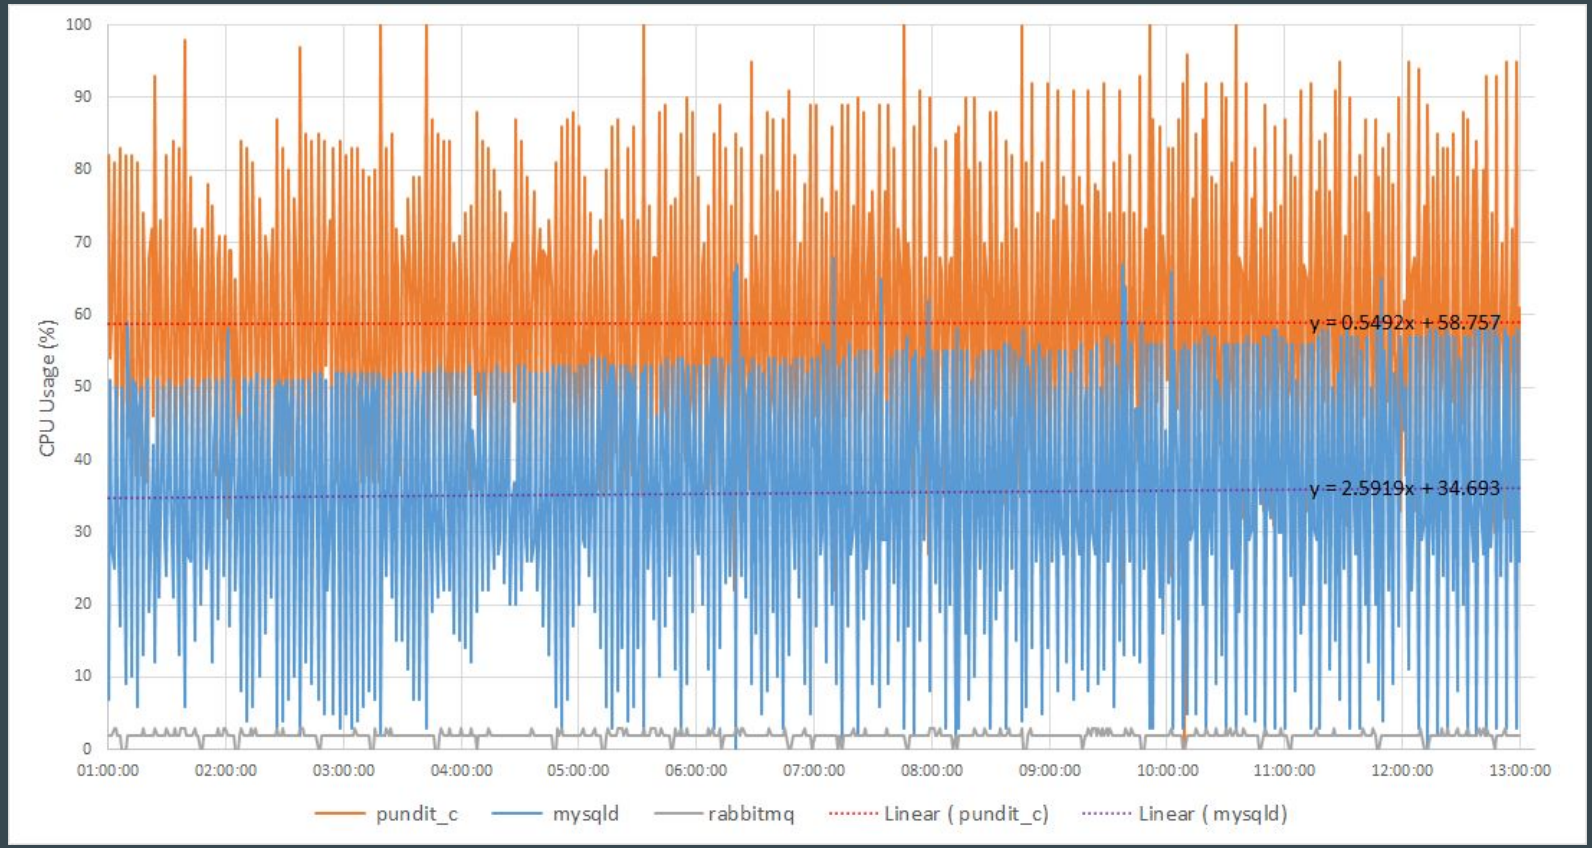

PuNDIT GUI: User Interface Example

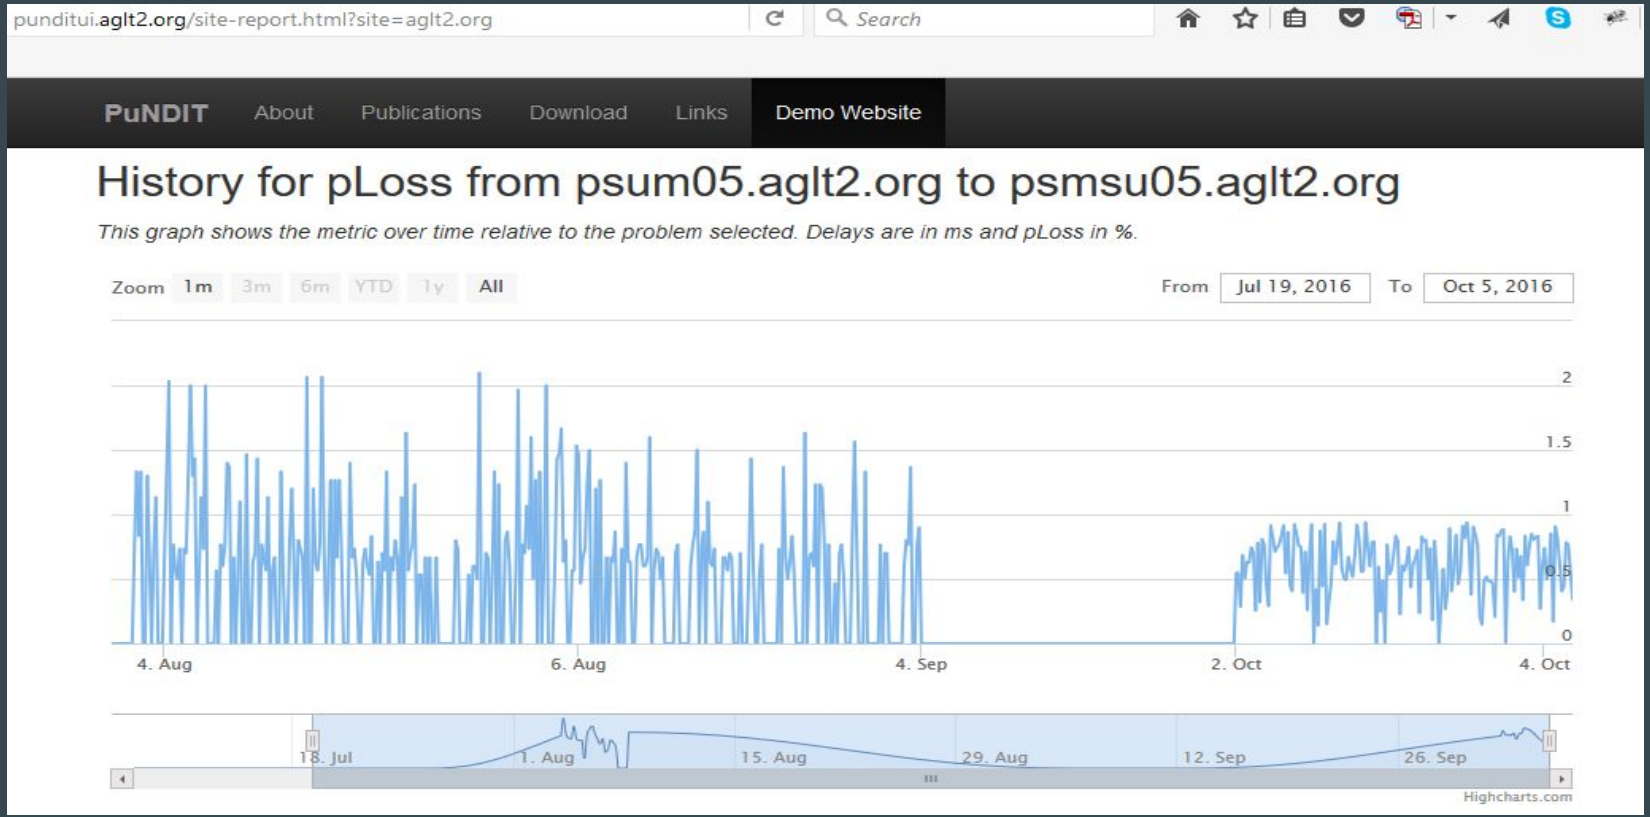

PuNDIT GUI: User Interface Example