Meeting Minutes from 12-July-2010 in Columbus

Performance Working Group at Joint Techs in Columbus, OH

July 12, 2010

Attending (partial list):

Jeff Boote
Chris Hawkinson
Carla Hunt
John Hicks
Greg Wickam
Matt Mathis
Tiziana Refice
Jose Augusto Monteiro
Daniela Bauer
Eric Pouyoul
Andy Lake
Aaron Brown
Brian Tiernery
Jason Lee
Yul Pyun
Seth Perry
Jon Dugan
Kavitha Kumar
Tom Throckmorton
Dan Pritts
Heorge Uhl
Inder Monga
Katsuhiro Sebayashi
Jason Zurawski

Internet2 Update   *(Jeff Boote)*

- **perfSONAR Software**
 pSPT 3.1.3 was released in April. 3.2RC1 released today. CentOS 5.5 based.
 Uses RPMs (all pSPS softare available as RPMs). Has the web100 kernel
 support on a stock CentOS kernel. RC, so we need testers. Also the LiveCD
 and install to disk option.

- New Employee(s)
 Welcome to Kavitha Kumar. Looking for more help, try to distribute this to
 as many people as you can: http://www.internet2.edu/about/staff/careers#netp

- NDT
 3.6.3 released.

- Internet2 ION Transition
 Complete. R&D staff helped move essential services such as lookup and
 topology. Web interface getting stats via perfSONAR for the junipers

- Standards
 NMC - Targeting a final draft of base doc by Oct 2010
 NM - Editing track for 'base' doc, deadline is Oct 2010
 NML - published a topology doc. Working on schema doc.

- JET Demo
 Interoperability between different networks. Bandwidth and Latency testing.
 NASA/NLR/Inernet2/ESnet/UEN/NOAA. Come to JET meeting for more.

Network Diagnostic Tool and the API from the Virginia Tech Perspective

- *Seth Peery, Virginia Tech Geospatial Information Sciences*

 Slides: http://serenity.gis.vt.edu/projects/ndt/presentations/jt/jtsummer2010.html

 Basic idea of talk - present the NDT JavaScript API. Talk about why it was
 needed and what it is good for. Also where the discussion on NDT needs to
 go.

Wanted to originially develop a framework for linking NDT with other web apps. Our use case was for mapping - e.g. Seth is in charge of GIS data at the university. Wanted to be able to map things like the performance measurements and geo-locate the sources. Useful for making things like a community broadband map - shows the VT example (google map of test results).

How did we get there? Initially added some simple hooks to NDT (its open source), set up a server and had people test to it. Couldn't start with a whole lot else (e.g. tools like ookla are closed source). Started slow, simple upload/download hooks. The changes were made to the applet portion, then applications can be developed against deployed servers. Others thought this was a good idea, then came to ask for more variables, and more formal treatment, e.g. MCNC, Broadband Census, Virginia Government.

Measurement lab approached as well - were doing work with the FCC. With the changes to the applet being targeted as a way to make new applications (that hid details of NDT). Hardened the changes to the API, made it a little nicer and documented it. Other MLab bits - deployed infrastructure, using tools like DONAR to help users test to a close server.

Getting other apps to use the API - simple code import (see web above). Creates new 'front ends' and at the same time each test will record tests for MLab.

Q: All the variables from web100 exposed?

A: Not all - popular ones that made sense. Can always add more.

Implications on 'what have we done' - NDT is a diagnostic tool that tells lots of things. Does adding it into a simple 'speed test' make sense? Does it do this well enough? Read recent MIT report and complaints when compared against something like Ookla. Basically bringing a research tool for one purpose into the 'mass market' speed test arena. Is this the right direction?

Audience feedback:

Chris H: NDT is the tool we turn too first for debugging things and getting a good baseline. Easy to run (for users), rich results to go though. After getting started, migrate to new tools for the most part. NDT is a 'tcp performance tester', wouldn't think it does well at raw speed tests.

Brian T: Research community will benefit from the data that mlab is collecting, even if it is 'flawed' as pointed out by MIT. Ways to show that the data is still useful, error estimations, etc. Also some thoughts on if TCP reno (used in NDT) is good for measurements compared to other implementations.

Some talk about what MLab is doing - e.g. Big Query (see talk later this week on it). Seth brings us back to the question - is NDT the right tool here?

Jeff B: This WG should be used to gather requirements. What should happen to NDT? It needs a lot of maintenance to keep it relevant. E.g. what happens in 100G testing? Heuristics need work, need to answer the congestion control question.

Matt M: Brings up some old IETF knowledge. E.g. IPPM wanted a simple way to test capacity. Knew it was hard. In NDT's case it has a repeatability problem. Works well at what it does (being able to tell very well what TCP is doing).

Seth: Wants to see someone from R&E (someone who will be more open) be the standards bearer for consumer broadband testing. Do need to make sure that NDT doesn't fall out of favor and not be able to go forward in the space. Notes that tools like ookla work well at what they do, we just don't know how they work...Would rather an imperfect tool that admits its flaws and people can see them vs an imperfect tool that doesn't admit failure and doesn't allow people to see what is going on.

Presta 10G Update

Come see demos at SC10

Discussion on pS/Performance Directions - (Jeff Boote)

We spent a long time making perfSONAR stable. Where do we go from here? What else does the community want? Perhaps: netflow? auth? closest MP

Do we need a task force from this group ? Will be summarizing the
results of the July 8-9, 2010 perfSONAR workshop soon. That could be a good indicator of what is
needed.

Some discussion on the NOC process - and how integrating the tools is still
a work in progress, but happening. Not enough representatives from an
end of line network (campus) in the room. Some thoughts from an exchange
point (latency is the enemy, little problems made worse).

GridFTP work - expose the results of tests. Limit active testing?
Web100 update later in the week, to see where that is going. Still valuable
for what it can expose. Tools use it.

A question on 'layer 2 visibility'. Also some clarification on what pS is
(a data model and schema) and what is isn't (re-making a lot of performance
tools).