

Meeting Minutes from SMM 2010 in Arlington

Performance Working Group Face-to-Face Meeting

at Internet2 Spring Member Meeting
April 26, 2010

Discussion

NSF/LSN/DOE sponsored perfSONAR workshop (Eric Boyd, Internet2)

<http://www.internet2.edu/workshops/perfSONAR/>

Internet2 Update (Jeff Boote, Internet2)

Key points:

- New hire in the R&D group starting in mid May.
- NSF Funding for two IRNC related projects (providing perfSONAR monitoring and dynamic circuit capability).
- Two releases of the pS Performance Toolkit (pSPT) -- bugfixes and security related releases
- Cisco AXP work (see Aaron Brown's presentation, below)
- ION service transition from Cienna to Juniper equipment. Conversion includes changing perfSONAR monitoring infrastructure
- Progress in standards areas (OGF)

ESnet Measurement Update (Joe Metzger, ESnet)

<http://www.internet2.edu/presentations/spring10/2010-04026-metzger-SMM.pdf>

- JET (joint engineering task force) consists of the us federal networks and others. Working on a perfSONAR demo (currently Ampath/NASA/Internet2 /ESnet/Internet2) to show the interoperability of the measurement tools. Policy is the hard thing they are trying to solve (particularly in respect to topology).
- ESCC is trying to develop a risk analysis template for installing perfSONAR tools at all DOE labs.
- DICE (ESnet/Geant/Canarie/Internet2/USLHCNET) is in the midst of a '2010 service challenge'. Try to use the same service defn, interdomain, at all levels (campus, regionals, backbones, countries) to deliver what users/operators want and need. Looking at a diagnostics, circuit monitoring, path characterization, and a 'portal' in a 12 month timeframe. Again this is a policy thing more than a development thing.
- E-Center is a project to develop a DOE centric measurement portal (Fermilab/ESnet/Bucknell). Web front end, access to analysis tools and the measurement infrastructure already deployed. ESnet will be providing measurements from the backbone as well as traceroute data. Traceroute data because most tools give you a point to point measurement, but not much info about the path itself. Related to the closest MP problem, etc.
- ESnet measurement infrastructure. About 60 machines (all 10g and 1g PoPs). 1 Gridftp server in response to common use case. Using GridFTP to move data between sites is common, need a common test point (in the middle) to assure performance
- Hiring new staff also - starting Mid may.

Update on the perfSONAR MA (Presta 10G card) -- (Kenji Shimizu , NTT Network Innovation Labs)

<http://www.internet2.edu/presentations/spring10/20100426-presta-shimizu.pdf>

- Originally talked about this at Joint Techs in Utah. Link to that presentation is at:
<http://www.internet2.edu/presentations/jt2010feb/20100202-shimizu.pdf>
- High-resolution traffic monitoring platform for 10G networks. 10G wire rate capture.
- Made some modifications to perfSONAR to make it capable of visualizing the new data format
- Slight modifications to XML messages
- New use case for this talk - Snow fiesta 2010 demo in Japan. Live video stream (compressed HDTV) over Japanese R&E network. Network is virtually sliced and shared.
- Dynamic vs. static assignment of underlying L1. Monitoring is used to ensure stable transmission and alert on problems (delay, jitter, burstiness).
- Fully featured multi-layer monitoring system
 - 'layer 7' - stream analyzing agent collects streaming headers
 - 'layer 3/4' - netflow probes

- 'layer 1' - high res pS MA introduced (ms resolutions)
- Visualizations to show things like traffic burstiness, delay distribution, one way delay, layer 7 traffic playback.
- Slides show what happens in the 'dynamic' configuration of the L1. System detects change, and recovers.
- Still ongoing development - integrating in pS.

perfSONAR on the Cisco AXP (Aaron Brown, Internet2)

<http://www.internet2.edu/presentations/spring10/20100426-perfWG-CiscoAXP.pdf>

- Embedded linux system (1.4ghz processor, 2G ram) that can be placed in the Cisco integrated services router. It has direct access to the backplane. Has an SDK
- Special thanks to Cisco and Steve Wolf for providing
- Benefits - server consolidation (power/space/cooling savings), one box to manage. Share the AXP across multiple instances.
- Does have direct access to the IOS (cisco operating system). Packet monitoring APIs are available, configuration management too. Event APIs.
- Development is done on a 'vhost' (single kernel, jailed linux for each development). RHEL4 compliant. Libraries to interact with the system for many languages (python/perl/c/java). Tools available to package things.
- Porting the perfSONAR PS tools was straightforward with some pitfalls. There is a perl environment available, doesn't include all the libs we use. Using PAT (perl archive toolkit) to make perl binaries instead. Two services: counters measurement archive (similar to SNMP MA, can use same GUIs to visualize data) and router measurement point (send router commands and get output).
- Future: other ps perfSONAR tools will be ported. bwctl/owamp. raw traffic monitoring.

XML Format for perfSONAR to Handle Hi-Res Measurement Data (Katsuhiro Sebayashi, NTT Network Innovation Labs)

<http://www.internet2.edu/presentations/spring10/20100426-NTT-sebayashi.pdf>

- HRA-MA (high res ma). captures and stores packets at 10g speed (100 NS resolution).
- HRA file formats are an extension of existing pS formats.
- Started with a pSPS SNMP MA. works with the presta 10G nic (seen earlier). Archive format uses something more specific than RRD files. File format takes this into account.
- Can use pSUI to view data
- Changed the XML message format slightly (didn't use the standard ISO 8601 formats).

Performance Service (Dale Finkelson, Internet2)

<http://www.internet2.edu/presentations/spring10/20100426-performance-finkelson.pdf>

- Performance service. Mostly from the point of view of a campus/user community. How can we derive some value from what is out there now?
- Goal should be to develop credible means of determining the capability on a connection. Also finding problems on a connection.
- Stuff is out there and deployed, good start. Need to get more.
- Identify problems that may exist (if they do exist), and where, and also provide a usable way to fix problems.
- Service definition - not every possible use of the network, but a couple of key ways to do so. Provide a useful interface to perfSONAR data (translate from 'operator' speak to English). Second level support in solving problems. Making the tools available.
- Use cases - data movement. What happens when the transfer fails, or is just poor, or has degraded over time? What about video - we care about pixelization, and bursts, and loss, and sound/video mismatch. Lastly - worry about clouds.
- The bottom line - moving bits. Just to need to help people do that effectively.