

Meeting Minutes Joint Techs July 2009 in Indianapolis

Performance Working Group

Joint Techs, Indianapolis, IN

July 20, 2009, 12:30 - 1:30 p.m.

I. Community Updates

A. ESnet Updates

- Been installing perfSONAR at government labs: 28 bandwidth testers, 26 latency testers

At every site so far have found a performance issues

- Number of openings for developers
- Russ: What have been problems?
- Brian: Generally misconfigured routers, Joe M to update more

B. Internet2 Updates

- Tomorrow new NPToolkit will be available
- Latency and throughput systems at every single router locations on the backbone
- Will be making latency and throughput data available via web services very soon
- Throughput data is available, OWAMP (latency) data available by Internet2 Fall Member Meeting in October

II. Netflow - Joe St Sauver

<http://www.uoregon.edu/~joe/ipv6-mask.html>

- Working on developing IPV6 netFlow anonymization policy
- Tried different mask
- Sounds like 8 bit mask will be needed, that will primarily will hide site
- Some addresses need bigger masks
- Need study to determine if less restrictive mask would be warranted
- Some people want less but need more research
- Eric question on section 5: Mentioned members of Internet2 measurement staff to conduct study. Should there also be external people?
- Joe: External people can participate, must be an NDA
- Joe: Need to make sure that flows can be consistently tracked over time
- Eric: Show this to members of IPV6 WG? Also should we show it to RAC?
- Joe: Going next door to present to Ipv6, talking to RAC
- Joe: Some fields netflow v9 not otherwise available. This includes MAC addresses
- Joe: Would like you to take a look at the document and give him feedback. Linked of Wiki.
- Eric: What are next steps? Is the intenet to bring to NTAC?
- Joe: Show it to NTAC, SALSA and various other groups. Still long process to get it approved.
- Eric: I am happy to help.

III. Circuit Monitoring - Aaron Brown

[view slides](#)

- Internet2 starting pilot service for their circuit network
- Users interested in seeing activity on circuit
- Internet2's network a series of Cien cordirectors
- Using OSCARS and DRAGON software to control the network
- Goal to create a simple web page that users can visit and see some information.
- Statistics collection required to be monitored. If using SNMP router have a slew of tools available. Cienas speak TL1 which is less available.
- Perl script that polls the Cienas, gets variables of interest and stores in RRD files
- PerfSONAR service that reads RRD and makes available multiple components:
 - Agent - converts OSCARS IDC view of the network into measurable topology. Caches the topology
- Carla: Do you have to have an IDC?
- Aaron: Currently its tied to IDC to get circuit list
- Carla: What format is the topology?
- Aaron: XML, in UNIS format
 - Web Client: Retrieves topology from the cache and circuits from the cache

- On the Cienas there is Ethernet on the edges and SONET in the middle
- There are also VCGs and sub network connections that get created but aren't expressed in topologies
- Dell 5224 is generic Ethernet switch. It has Ethernet interfaces and VLANs created are new interfaces.
- Diagram in slides explains the process
- Future: Enable other domains to share information. Facilitates better debugging and could create better visuals tools
- Can leverage more of perfSONAR in future
- Aaron displayed Google maps of circuit
- Can filter list of circuits and only show you circuit and show utilization. Other data being collected but currently only utilizations
- Tom: Which TL1 toolkit are you using?
- Aaron: Wrote my own. I ran into issue with TI1 toolkit. They didn't support CoreDirector and didn't have the commands I needed. I was working in parallel with the U NOC and would have collaborated with them more had I know. URLs will be available later this afternoon

IV. Measurement Scheduling - Joe Metzger

- Wants guidelines for scheduling test. Primarily test with your peers but some things related to own backbone.
- How long should test run?
- Joe: it shouldn't be 10% of utilization and 10% seems high
- There is not good way to know when people are testing, but BWTCL does have a way to control how much testing is done to you - one of main purposes
- Jeff: During early Internet2 days decided on fix % based on capacity. It was 10% but at time were very overprovisioned.
- Jeff and Chris Small: In early days of Internet2, BWCTL was sometimes in 30%-50% range of utilization.
- Chris Small: Currently most of IPV6 traffic is measurement
- Joe: Anyone have issue with that?
- Tom: Traffic cyclical so doing % over a period of time can be tricky given this fact.
- Jeff: If you're over-provisioned then doesn't matter, but once you start hitting congestion you want to back off tests
- Martin: Can do heavyweight that tell user you can get X performance, and do smaller ones to track heavy one without problem.
- Group: Is one week enough for heavyweight and 4-6 times for lightweight?
- Tom: Interval needs to relate to what users is doing.
- Jeff: Once a week does not seem like enough because config changes happen more than that.
- Martin: Might depend on your network
- Vangelis: Do want to test to everyone (i.e. n-squared)?
- Jeff: No
- Jeff: You start running out of timeslots between 15 and 20 hosts if you're doing hourly because tests take time.
- Carla: would it make sense to make to create a document that has guidelines?
- General consensus was yes
- Carla: should share with NTAC for awareness
- Group: Version of NPToolkit released tomorrow supports that
- Martin: Can do it manually now but show also automate it in the future. We need more thought into how we will do that.
- Joe: How often you run tests, should that be related to who you are sending them to? i.e. send them more to people who are closer since it could have bigger impact.
- Jeff: You still have to run them a fair amount of time
- Brian: Too close is not useful either though
- Joe: I would guess that people will test with 6-12 sites
- Joe: Anyone who sets-up should setup DNS and preferably set up DNS with location records. Helps with visualization.
- Joe: QOS is set on ESnet as scavenger so if "real" traffic starts then measurement traffic is the first to go
- Brian: Should heavyweight have no QOS since it needs to tell you what you can get.
- Jeff: Depends on what you're after.
- Joe: In terms of scheduling tests Eli Dart has pushed the idea that we should define what a reasonable rate for the standard user is. Use that number for the tests.
- Rich: What about science community that needs more?
- Joe: They invest more in making performance better. That's minority of sites though and we want to tell the 95% of sites that don't have CMS or ATLAS data
- Rich: Do you want to set the bar higher so they can reach for it?
- Joe: We want to let them know there is a bar.
- Danno: What's the difference between butting up a web page is this is what you can get vs a page on how to tune system
- Jeff: Wondering how you are going to limit your tests to a certain amount.
- Jeff: It would be good to be able to show that 90% of sites get better performance than you, there must be something wrong
- Brian: It would be useful for NPToolkit to generate email alerts if BWTCL tests don't reach a threshold
- Joe: What should be X% capacity and Y% utilization for measurement tests be?
- Jeff: Depends on a lot of stuff. Three cases:
 1. What you do on your own network?
 2. What you do with your peers?
 3. What your users do?
- Carla: Let's keep brainstorming on mailing list

V. Deploying a Statewide Perf Measurement Network - Tom Throckmorton

[view slides](#)

- Big driver was connecting all public schools (115 school districts) to NCREN.
- Another 240 customers so emphasis on performance.
- Goal to be non-obtrusive with network provide more transparency
- What to increase active measurement and measure out to edge
- Connected 115 new systems but only had 120 prior to that so effectively doubles measurement points
- Challenge to provide meaningful and relevant output
- Challenge achieving centralized collection of data
- Few tools available to collect data
- Keeping everything simple yet maintainable.
- POPs are typically DC powered so finding hardware that was supported in environment was a challenge
- Trying to solve connection issues for participants such as when a connection seems slow. Challenges:
 - Customer has varying expertise and connection
 - Not very much documentation