

# Performance Area Group

Matt Zekauskas, [matt@internet2.edu](mailto:matt@internet2.edu)

Tuesday 22-Jan-2008

Hawaii Joint Techs



# What the heck is a PAG?

- The Performance Area Group is a repurposing of the Measurement SIG
- And subsuming some of the functions of the previous End-to-End TAG (e.g., act to provide some direction, too)
- For all of the Internet2 infrastructure

# What the heck is a PAG?

- A place to hash out ideas related to measurement and performance
- Primarily discussion, not presentation
- Provide input on these issues for the Internet2 Network to the NTAC
  - For example, Observatory functions; Layer 1 monitoring; things NOC should be doing
- Large, specific, issues can spawn a Working Group (need workers first!)

# This is a work in progress

- Discussion started December 2006 at the Fall Member meeting in Chicago
- Some discussion in the NTAC
- This is the third meeting, and (still) somewhat disorganized at that 😊
- Feedback welcome!

# Mailing list?

- Er, no. We could generate if interest
- There is also the old Measurement SIG list

# Pointers to other related sessions

- Tools tutorial, last Sunday
- perfSONAR WG yesterday
- Circuits stuff this morning
- Performance & perfSONAR update  
Today 4:30 to 5p (and surrounding sessions)

# General Topics

- Things we are doing
  - Observatory function recap
  - Ongoing monitoring
  - Resuming multi-vendor Ethernet circuit testing
- Things that could be done
  - On Internet2 Network
  - Internet2-wide
- Anyone with an interesting measurement or performance issue?
- Open floor

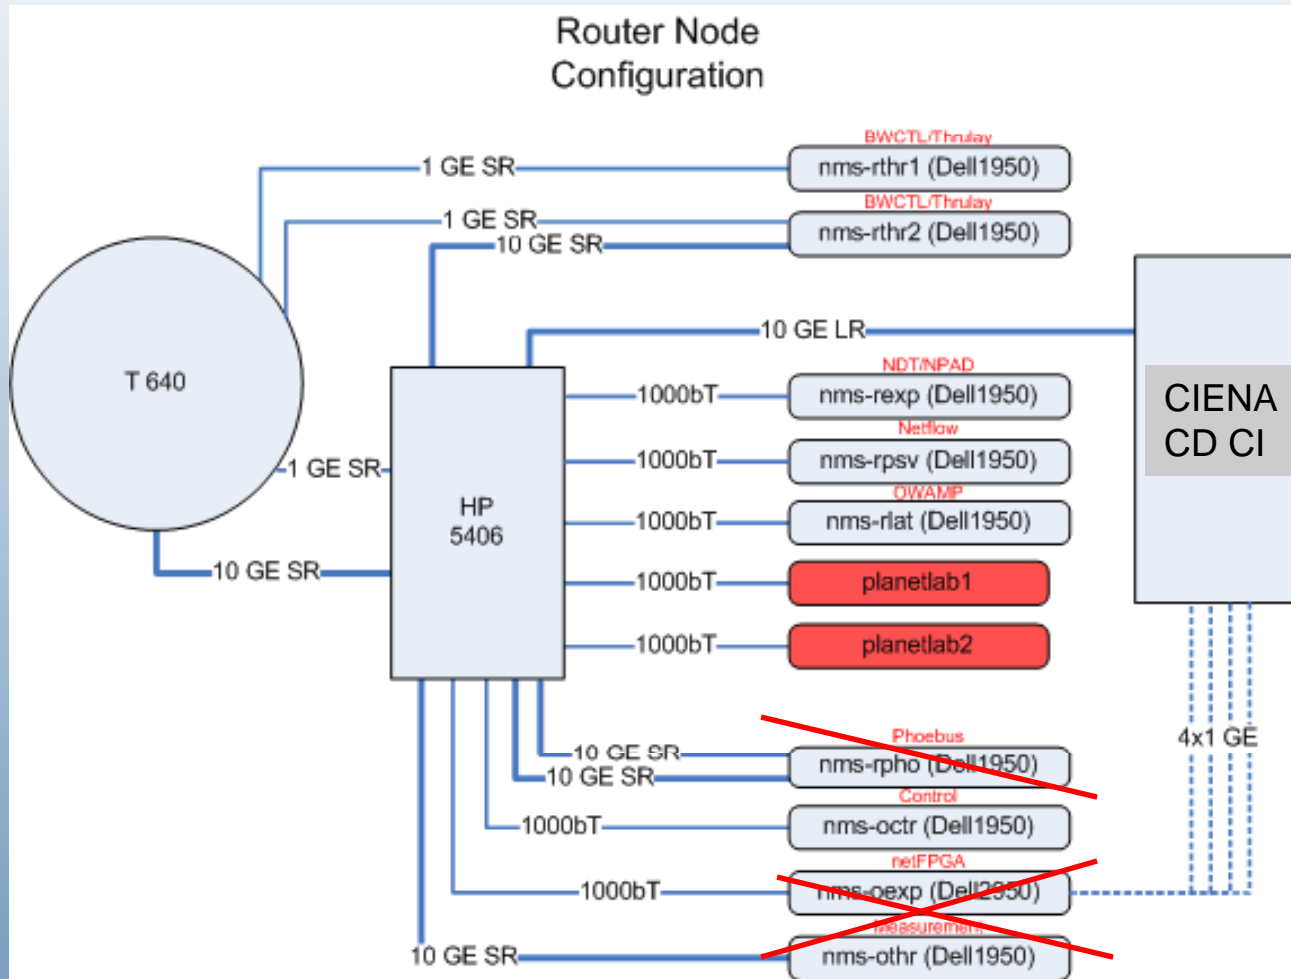
# Existing Observatory Capabilities

- One way latency, jitter, loss
  - IPv4 and IPv6 (“owamp”)
- *Regular* TCP throughput tests – ~1 Gbps
  - IPv4 and IPv6; On-demand available (“bwctl”)
  - ~10GE now also possible (Myricom and Dell 1950, must ask)
- SNMP
  - Octets, packets, errors; collected 1/min
- Flow data
  - Addresses anonymized by 0-ing the low order 11 bits
- Routing updates
  - *Both IGP* and BGP - Measurement device participates in both
- Router configuration
  - Visible Backbone – Collect 1/hr from all routers
- Dynamic updates
  - Syslog; also alarm generation (~nagios); polling via router proxy

# Observatory Physical Design Review

- For lots of detail, see the presentation at Chicago member meeting.  
<http://events.internet2.edu/2006/fall-mm/sessionDetails.cfm?session=2995&event=258>
- Some equipment deployment delayed (2/3 PCs in “optical nodes”, 3/9 PCs in (most) “router nodes”)
- Deployment (modulo Seattle) complete
- Feedback always welcome

# Router Nodes (Original) Plan



# Observatory Functions

Device	Function	Details
nms-rthr1	Measurement	BWCTL scheduled 1 Gbps router throughput, Thrulay
nms-rthr2	Measurement	BWCTL on-demand 1 (and 10) Gbps router throughput, Thrulay: <a href="#">bwctl.{router}.net...</a>
nms-rexp	Experimental	NDT/NPAD: <a href="#">ndt.{router}.net.internet2.edu</a>
nms-rpsv	Measurement	Netflow collector
nms-rlat	Measurement	OWAMP with locally attached CDMA/GPS timing: <a href="#">owamp.{router}.net.internet2.edu</a>
<i>nms-rpho</i>	Experimental	Phoebus 2 x 10GE to Multiservice Switch
nms-octr	Management	Controls Multiservice Switch
<i>nms-oexp</i>	Experimental	NetFPGA: 4 sites to be funded by 100x100
<i>nms-othr</i>	Measurement	On-demand Multiservice Switch 10 Gbps throughput

# Observatory software

- Netflow, SNMP, BGP, config: collection OK
- ISIS collection tickled Juniper bug, awaiting fix
- OWAMP working where we have good clocks
- **bwctl awaiting new distribution, ad-hoc measurements working**
- NDT should be working everywhere

# Observatory web pages

- Working on converting new pages to new layout, remove Abilene references

# General Topics

- Things we are doing
  - Observatory function recap
  - **Ongoing monitoring**
  - Resuming multi-vendor Ethernet circuit testing
- Things that could be done
  - On Internet2 Network
  - Internet2-wide
- Anyone with an interesting measurement or performance issue?
- Open floor

## Other monitoring...

- perfSONARBOUY, just mentioned
- perfSONAR export of utilization, topology
- perfSONAR export of OSCARS info (current dynamic circuits)
- In next 6 mos or so: perfSONAR export of Ciena performance data

## In addition

- NOC is working on monitoring the Ciena and Infinera gear independently, and putting that in their alerting system.
- This should give us Infinera status we can publish too.

# General Topics

- Things we are doing
  - Observatory function recap
  - Ongoing monitoring
  - **Resuming multi-vendor Ethernet circuit testing**
- Things that could be done
  - On Internet2 Network
  - Internet2-wide
- Anyone with an interesting measurement or performance issue?
- Open floor

# Circuit testing

- Would like Ethernet links originating on the Dynamic Circuit Network to be able to terminate outside the Internet2 Network.
- For example, at a European lab or university.
- A bunch of work being done by DANTE, ESnet and Internet2 to ensure interoperability

# Sample Path to Europe

- Dynamic Circuit Network
  - Mainly ESLM cards in the Ciena
- MANLAN – HDXc switch (not dynamic)
  - Possible termination on OME6500
- OC192 under the Atlantic ocean
- Alcatel equipment in Europe

## Many cases

- 1GE: Ciena through to Alcatel
- 10GE: Ciena through to Alcatel
- 1GE mediated through Nortel OME6500 in MANLAN
  - (so, two half-circuits rather than a complete end to end circuit, with another Ethernet in the middle)
  - Also potential for other customers of MANLAN than Internet2's DCN
  - Works other way too, as onramp to DCN
- 10GE mediated through OME6500

# 1GE to Alcatel 1GE

- Have a Spirent 600 courtesy John Moore and the NC-ITEC
- Ciena to Alcatel
  - Have found losses at certain packet sizes near maximum rates, and in one direction (not max or min packet sizes, though)
  - We are working to understand this with the vendors
- Nortel to Alcatel: so far, so good
- Nortel to Ciena

# 10GE to Alcatel

- Only some Ciena-Alcatel testing so far
- 10GE PC's on either side
  - Hardware not identical
- Running code by Richard Hughes-Jones of Manchester
- ...and some single-stream Iperfs
- “OK”, so far, but see some PC config limitations that are working to eradicate first.

# General Topics

- Things we are doing
  - Observatory function recap
  - Ongoing monitoring
  - Resuming multi-vendor Ethernet circuit testing
- **Things that could be done**
  - On Internet2 Network
  - Internet2-wide
- Anyone with an interesting measurement or performance issue?
- Open floor

## Joe St Sauver, a while ago:

- “The traffic mix is changing due to things like the commodity peering initiative, etc. I'd be interested in any sort of pre-post analysis if the data to do that is available, although I understand that peering terms may limit visibility into that traffic.”
- ... there is a lot that could be done with information already available from the observatory. Any interest?

# The general question

- What information would you like to see from network monitoring?
- Consider two points of view:
  - Network Operations/Network Engineers
  - End Users
- Have recently heard would like Aggregate demand graph back

[www.internet2.edu](http://www.internet2.edu)

