NANOG 34 – Internet Exchange Operator Panel

Chris Malayter TDS

Network Engineering



Order of Operations

- Introductions
- Overview
- Speakers
- Questions Time Permitting
- Closing

Order of Speakers

- Mike Hughes LINX
- Dave Meyer OregonIX
- Tom Bechly MCI/MAE
- Troy Davis SIX
- Celeste Anderson Pacific Wave

Upgrading to 10G and beyond: Real World Experiences Mike Hughes, CTO mike@linx.net NANOG34 IXP Panel

Current Status

- 173 members from 27 countries
- Present at 6 Carrier hotels in London
 Not owned by LINX, independent
- Dual platform, multiple 10G ring design
 - Two separate public peering LANs, available across all sites
 - Foundry and Extreme switches
- Public switch traffic: 56 Gb/sec
 - With steady growth
- Managed Private Interconnects: ~25 Gb/sec

10G Deployment: Part 1

- Ran 10Gig Ethernet in backbone since 2001
- Recently completed an overhaul of the Foundry network
 - Deployed Foundry MG8
 - Wasn't completely problem free (linecard crashes)
- Now connecting up our first 10G participants
- Growth in link-aggregated connections
 - But some bugs also found in member routers
 - e.g. f00d.beef.f00d floods

We love pain! 10G: Part 2

- Two networks give us lots of extra redundancy and flexibility
 - Does mean we get to do things twice, though!
- This year, LINX will upgrade the Extreme platform to be of an equivalent spec
 - Both networks need to be roughly equal
- Currently in the selection process
 - So, I'm slightly limited what I can say...

Interesting Selection Datapoint

Really small packets in the majority



Packet Size Distribution at LINX

What shouldn't happen...



Sleeping on packaging = bad!

So, what's changed this time?

- No "big bang"
 - Hardware will be trickled in
 - Outweighs disadvantages of homogeneous network
- Test as much as possible, then test it again!
 - Can you be too thorough?
- Agreed acceptance criteria with vendor
 - Especially for the first system
 - Acceptance criteria for payment and future orders

Other Projects

- We're also currently working on projects in the following areas:
 - Multicast distribution
 - Using PIM snooping to control multicast behaviour
 - DDoS mitigation
 - Using a British-built cleaner
 - Trying to solve the VLAN issues
 - Provisioning, tag-space clash issues, renumbering, etc
 - Collecting sFlow data
 - For traffic management purposes

Oregon-IX

David Meyer

Agenda

- Brief History
- Current Status
- Futures

First, Where in the world is Oregon



Oregon is the kinda square state *between* California and Washington

Geography: A slightly closer look



Brief History of the (OIX) Universe

- Randy Bush and I started what became the OIX
 - Circa 1995
 - Was (and continues to be) a cooperative regional effort
- Peers
 - Rainnet, UO, NERO
- Switching
 - Well, 10mb/s small switch
- Fiber
 - Qwest
- "Colo Space"
 - Hijacked closet space in the UO Research Park

Current Status

- Peers
 - O(30)
 - Mostly regional participants
 - Downturn had substantial impact
- Switching
 - Dual HP 2800 (10/100/1000)
- Fiber
 - Several local, regional, and national networks
- Collocation
 - Telco-grade conditioned environment
 - UO Switchroom

Services

- Cross Connects
- Remote Hands
- "Consulting"

Futures

- Additional fiber facilities
- Switch fabric Upgrades (-> 10G)
- Services
 - Experimental VoIP peering
- Layer1/Layer2 extentions
- And of course we would like more traffic

Thanks

http://www.oregon-ix.net

MAE® Services NANOG 34 – IX Operators Panel

Tom Bechly May 17, 2005 tom@mae.net

MAE® Services and Remote Peering

- Peering requires delivering traffic from own network to peer network
 - -Build out national or global network
 - Lease private line from own network to peer network or exchange point
 - -Extend reach via service provider network
- Remote Peering may meet needs of cable networks, content providers, regional ISPs, and their peers
 - -Lower transit or reduce network costs
 - -Control over routing
 - –Improved customer experience (latency, throughput)

MAE Services Extended (MAE EXT)

- MAE EXT is suite of services to support needs of ISPs
 - –Remote Peering, L2 VPN, Full and Partial transit all from one port
- MAE EXT remote Peering
 - -Connections can be established between any MAE Services POP
- MAE EXT Layer 2 VPNs
 - -Customer can extend their network reach thru VPN as an alternative to private line (i.e., OCn) connection
- MAE EXT Transit
 - –Full routes, a full view of the global Internet routing tables

MAE EXT - continued

- MAE EXT is a connection oriented service
 - -Gigabit Ethernet Access 802.1q tagged (SX and LX)
 - -POS Frame Relay access DLCI (OC3, OC12, and OC48 SMIR)
 - -ATM VPI/VCI (OC3 and OC12 SMIR)
- Platform performs IPv4/IPv6 interworking between Ethernet, Frame Relay, and ATM (Based on Draft Martini)
 –IPv6 GigE interworking with Frame Relay and ATM - 7/05
- Usage based billing
 - -Peering an L2VPN 95th percentile of traffic pushed to network
 - -Transit is billed 95th percentile greater of transmit or receive
- Transit service is currently only IPv4, IPv6 available 3Q05
- All Peering and L2VPN connections provisioned via PeerMaker

MAE® Services Locations

Current Locations

- Wash DC Vienna & Ashburn
- Dallas InfoMart
- San Jose Market Post Tower
- Miami NAP of the Americas
- Pending Location
- New York City (111 8th Ave)

- New York City Ave of the Americas
- Chicago Printers Square
- Los Angeles One Wilshire Bldg

MAE Services – Global Extended Peering 2005/06



seattleix.net

Seattle Internet Exchange your packets, our frames

NANOG 34 IX Operator Panel Troy Davis troy@nack.net

seattleix.net What's unique here?

- Regional exchange
 ... in a broom closet
- Non-profit
- All volunteer
- Cheap splay



Gentle hockey stick



seattleix.net Zero-cost ecosystem

- Petri dish
 - -Peer at any size, for any reason
 - -Small is temporary

Practical scope

Transparent



seattleix.net Seems valued

- 40% average annual growth
 - Despite simple, cheap private peering
 - -> 100% '04-'05 (P2P, NNTP)
- Quietly reliable
 - Physical & basic layer 2 restrictions
- "Critical mass" moving target

seattleix.net

Future

- SIX in a year
 - Forklift?
 - n*GE to (a few) peers
 - New locations?
- Spare broom closet?
 - Geography/topology on your side: build
 - No critical mass (or 7 years): extend
 - A means, not an end





Pacific Wave www.pacificwave.net

NANOG 34 Seattle, Washington May 15-17, 2005

Global Lambda Integrated Facility World Map – June 2005

Predicted international Research & Education Network bandwidth, to be made available for scheduled application and middleware research experiments by June 2005.



www.glif.is

Visualization courtesy of Bob Patterson, NCSA.





Features, Activities, Plans

- Separate VLANs for standard and jumbo frame size.
- Multicast in same unicast VLAN via PIM snooping.
- IPv6 coexists in same VLANs as IPv4 traffic so that participants operate a conventional dual-stack interface to the exchange.
- Also providing Lambda services as part of National Lamda Rail (NLR) for projects such as CaveWave, Ultralight, etc.
- Next year connections from the Translight and WHREN projects.
- Working on flexible interconnects for layer 1 and optical and how they would/could be managed (or not).
- Measurement!

PACIFIC



Contact Information

A Joint Project of the Corporation for Education Network Initiatives in California (CENIC) and the Pacific Northwest Gigapop (PNWGP) In collaboration with the University of Washington (UW) and the University of Southern California (USC) Web: www.pacificwave.net Email: info@pacificwave.net + 1 206 PAC WAVE (1 206 722 9283) + 1 888 PAC WAVE (1 888 722 9283) Canada & US

Other sites:

Global Lambda Integrated Facility (GLIF) <u>www.glif.is</u> National Lambda Rail (NLR) <u>www.nlr.org</u> Internet Educational Equal Access Foundation <u>www.ieeaf.org</u> Translight <u>www.startap.net/translight</u> Western-Hemisphere Research and Education Networks (WHREN) http://www.cenic.org/events/cenic2005/pres/whren.pdf

Questions?

Closing

Chris Malayter TDS Network Engineering chris.malayter@tdstelecom.com

