Oregon Gigapop Traffic Characterization

Internet2/NLANR Joint Techs
May 16th, 2001, Lincoln NE

Joe St Sauver, Ph.D.
(joe@oregon.uoregon.edu)
Computing Center
University of Oregon

I. Introduction

How Does Oregon Connect to The World?

- The Oregon Gigapop (located at UO in Eugene) has two POS OC3's from Abilene: one to Denver, and one to Sunnyvale
- Commodity transit (UUNet & CWIX) is handled separately via OWEN/NERO (http://www.nero.net/)
- Peering (Verio, Globix, Akamai, etc.) is handled separately via the Oregon Internet Exchange (http://www.oregon-ix.net/)

The Abilene OC3s

Normal weekly traffic, Denver OC3

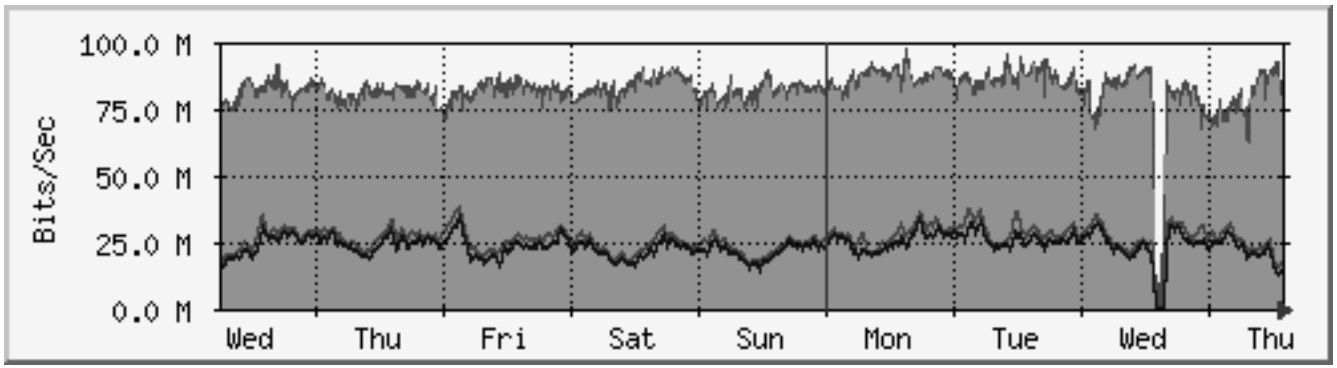

Normal weekly traffic, Sunnyvale OC3

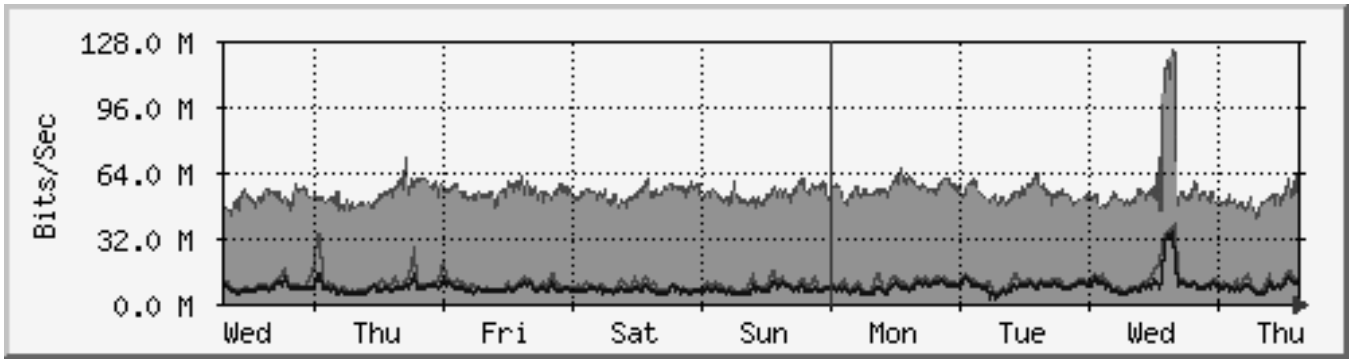

Who's At the Oregon Gigapop?

- UO, Oregon State and Portland State
- Four sponsored participants: Eastern Oregon Univ, Oregon Inst of Technology, Southern Oregon Univ, Western Oregon Univ
- And now, as a SEGP, the Oregon Public Education Network (OPEN), servicing ~600,000 public K12 folks statewide

What's Not Included In This Flow Study

- Commodity transit and peerage traffic won't be seen in this HPC-only flow study
- Most intrastate educational traffic is handled externally from I2 (the exception being traffic to/from OHSU/OGI in Portland, which connects to I2 via the University of Washington gigapop)
- SEGP folks weren't live at the time this data was collected (but they are now)

Privacy Issues

- Traffic measurements were an integral part of what was originally proposed and funded by NSF in our Connections grant
- Still, we still recognize and appreciate the importance of privacy-related concerns.
- Consistent with that, we report data in this study only in aggregated form.

An Earlier Commodity Transit Bandwidth Study Is Available

- https://web-vms.uoregon.edu/~joe/bw2/owen/index.html

  Note that since the time of that study, a year ago, new applications have emerged, network traffic patterns may have changed, and a number of new providers have joined the Oregon IX, so you should interpret that historical data cautiously if you look at it.

Generalizability

- We make **no claim** that the Oregon Gigapop's traffic is "typical" of Abilene traffic in general (as if there is such a thing as "typical" Abilene traffic!)
- We firmly believe that every participant and every gigapop will differ, often materially
- Individual users matter greatly, and can cause dramatic changes in traffic patterns
- We'd encourage ALL I2 sites/gigapops to do their own traffic characterization study

II. The Flows

~1.45 Million Flows

- 1,454,544 flows from the OGIG Cisco 12008 via sampling netflow with a sampling factor=10 (ftp://ftpeng.cisco.com/isp/12.0S/Features/sampled_netflow.txt), resulting in a sample totalling 5.3716E10 octets worth of traffic.
- The data file itself, in the format as it was analyzed with SAS, is 181,463,583 octets or 25,033,364 octets when gzip -9'd

Times...

Fri 4/27/01 13:29-15:40 (131 minutes)
363,636 flows and 1.442E10 octets
Mon 4/30/01 12:54-14:47 (113 minutes)
363,636 flows and 1.302E10 octets
Tue 5/01/01 13:30-15:32 (122 minutes)
363,636 flows and 1.353E10 octets
Wed 5/02/01 13:21-15:12 (111 minutes)
363,636 flows and 1.275E10 octets

Representative and Comparable Times

- The days selected were typical days… (not vacation periods, nor on weekend days)
- Mid afternoon times are our traditional peak for our inbound commodity transit pipes; we also wanted times consistent with our earlier commodity transit bandwidth study
- HPC traffic patterns are relatively time invariant round the clock anyway...
- The data presented here was collected prior to our SEGP folks going live

Directionality

- **Inbound:** 832,093 flows and 1.27E10 octets (mode 40 octets/flow, median 115, mean 15.3K, min 28, max 17.82M, std dev 232K)
- **Outbound:** 622,451 flows and 4.1E10 octets (mode 40 octets/flow, median 108, mean 65.9K, min 28, max 18.37M, std dev 430K)
- That 3.2-to-1 outbound-to-inbound octet ratio is consistent with what MRTG shows

III. Protocols

Protocols Seen

TYPE	FLOW COUNT	OCTETS
TCP	1,038,519 (71.4%)	4.513E10 (84.0%)
UDP	249,162 (17.1%)	8.563E9 (15.9%)
ICMP	127,591 (8.8%)	1.693E7 (0.03%)
IPMP	34,382 (2.4%)	1.937E6 (0.00%)
PIM	4,180 (0.29%)	2.594E6 (0.00%)
IPv6	521 (0.04%)	56,816 (0.00%)
IGMP	189 (0.01%)	8,712 (0.00%)

For Comparison...

- Trends in Wide Area IP Traffic Patterns: A View from Ames Internet Exchange (www.caida.org/outreach/papers/AIX0005/) reported 91% TCP octets, 5.1% UDP octets, 2.7% GRE octets and 0.7% ICMP octets (plus some other miscellaneous protocols)

Non TCP/UDP/ICMP Protocols In The Oregon Flow Data...

- Proto=169 ==> IPMP (http://watt.nlanr.net/AMP/IPMP/)
- Proto=103 ==> PIM (all local router to local router traffic only)
- Proto=41 ==> IPv6 (associated with UO's 6bone-gw router and/or one of our partners tunneling to dnvr-v6.abilene.ucaid.edu)
- Proto=2 ==> IGMP

IV. ICMP

ICMP Traffic...

- ICMP traffic was somewhat higher than normal (8.8% of all flows vs. 1.5-2% of all flows in earlier commodity internet flow studies) due to OGIG partner participation in various network measurement projects, e.g.:
  -- OSU hosts an NLANR AMP box, and
  -- UO hosts an AMP box, a Surveyor, a NIMI box, a Skitter box, and a looking glass
- These 6 boxes were involved in 69.4% of all ICMP flows (and 36.8% of all ICMP octets)

But We Also Saw...

- … 326 outbound ICMP flows, a directed broadcast (‘smurf’) attack totaling 4.5MB (or 26.7% of the total ICMP traffic) directed against an IRC server in Europe (the vulnerable partner router has been corrected)
- …. 4,609 ICMP flows associated with a single dynamic address (which also saw 40 or 80 byte flows on 4661/TCP and related “large” flows on 4662/TCP)

An Aside Regarding the TCP/4641 and TCP/4642 Flows

- We believe those flows are associated with "eDonkey2000" -- see: www.bajapuntes.com/edonkey/faqes.html
- Their slogan (I kid you not): "Harness the power of 2000 electronic donkeys!"
- Obligatory pun about peer to peer networking programs and network *as*sets

We Also Saw Other P2P Apps Doing Network Latency Checks

- We saw ICMP traffic from some additional sources (and to some additional destinations) which would be consistent with interactive gaming-related network latency checks and/or peer-to-peer file sharing-related network latency checks
- We also saw some “normal” ICMP traffic.

V. UDP

UDP Traffic

Type	**Octets**	**Flows**
Multicast	8.01E9 (93.6%)	39,053
Real	2.55E8 (2.98%)	6,613
Half-Life	8.01E7 (0.94%)	29,952
ICQ	8.68E6 (0.10%)	3,956
DNS	7.36E6 (0.09%)	59,525
Remainder...	2.29%	44%

- Note: we are atypically multicast-intensive
- And note: the remaining 2.29% of UDP octets == ~44% of all UDP flows

About That Multicast Traffic

- Virtually all (99+%) of outbound multicast octets originated from UO sources
- One interesting multicast service/app which we became aware of in looking at the multicast flow data is described at http://emeetingportal.com (see also http://www.marratech.com)

VI. TCP

TCP Traffic

Type	Octets		Flows
nntp	3.66E10	(81.05%)	331,628
ftp	1.16E09	(2.56%)	16,047
gnutella	9.91E08	(2.20%)	36,897
napster	9.37E08	(2.08%)	10,797
http	6.63E08	(1.50%)	431,208
kazaa	4.28E08	(0.95%)	15,469
ldm	1.47E08	(0.33%)	1,683
hotline	6.88E07	(0.15%)	715

(continued...)

TCP Traffic (cont.)

qt/rtsp/real	4.02E07 (0.09%)	2,035
ssh	3.84E07 (0.09%)	1,676
smtp	3.60E07 (0.08%)	38,654
shoutcast	3.31E07 (0.07%)	1,204
aim	2.15E07 (0.05%)	154
icq	1.38E07 (0.03%)	1,515
dns	1.29E07 (0.03%)	606

- Note: the remaining 8.74% of all TCP octets == ~14.2% of all TCP flows

About Our NNTP Traffic

- OGIG partners have coordinated news feeds available from and managed by UO (although partners are free to seek supplemental feeds to fill in niches which may not be available from UO's feeds), thus it is not surprising the 98% of the inbound NNTP traffic flows to UO's news feed boxes, and 96+% of the outbound NNTP traffic is sourced from UO's news feed boxes

Our NNTP Traffic Level Looks Comparable to Other Reports

- In http://darkwing.uoregon.edu/~joe/how-to-go-fast.ppt at slide 39 I showed the relative fraction of octets represented by Usenet at three CANet3 sites (BCNet, MRNet, and RISQ)… NNTP traffic ranged from 72-77% of total octets (while OGIG's represented 81%)

How Much Traffic Does a “Full” Feed Represent?

See: http://newsfeed.mesh.ad.jp/flow/

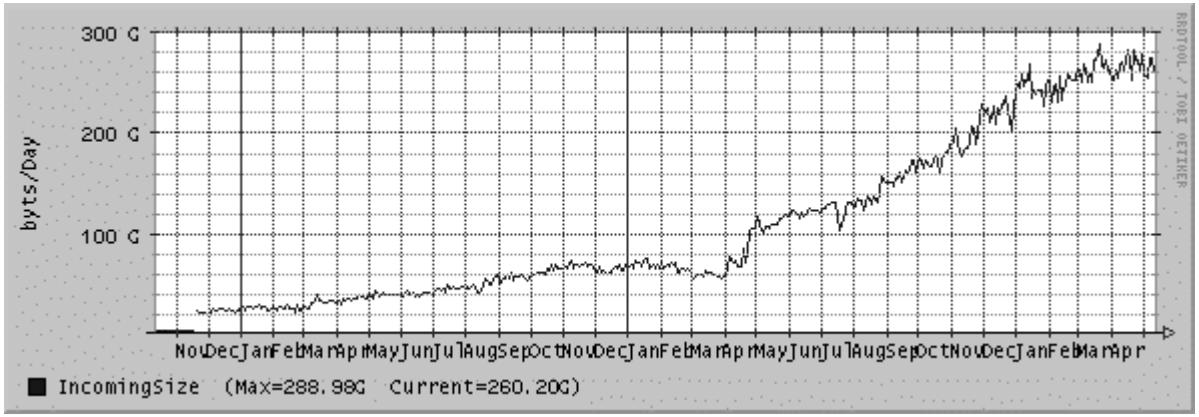

The recent leveling-off of traffic shown is may be associated with article size limits imposed by a major satellite Usenet firm

Traffic Volume vs. Article Size: “The 80/20 Rule” Holds True...

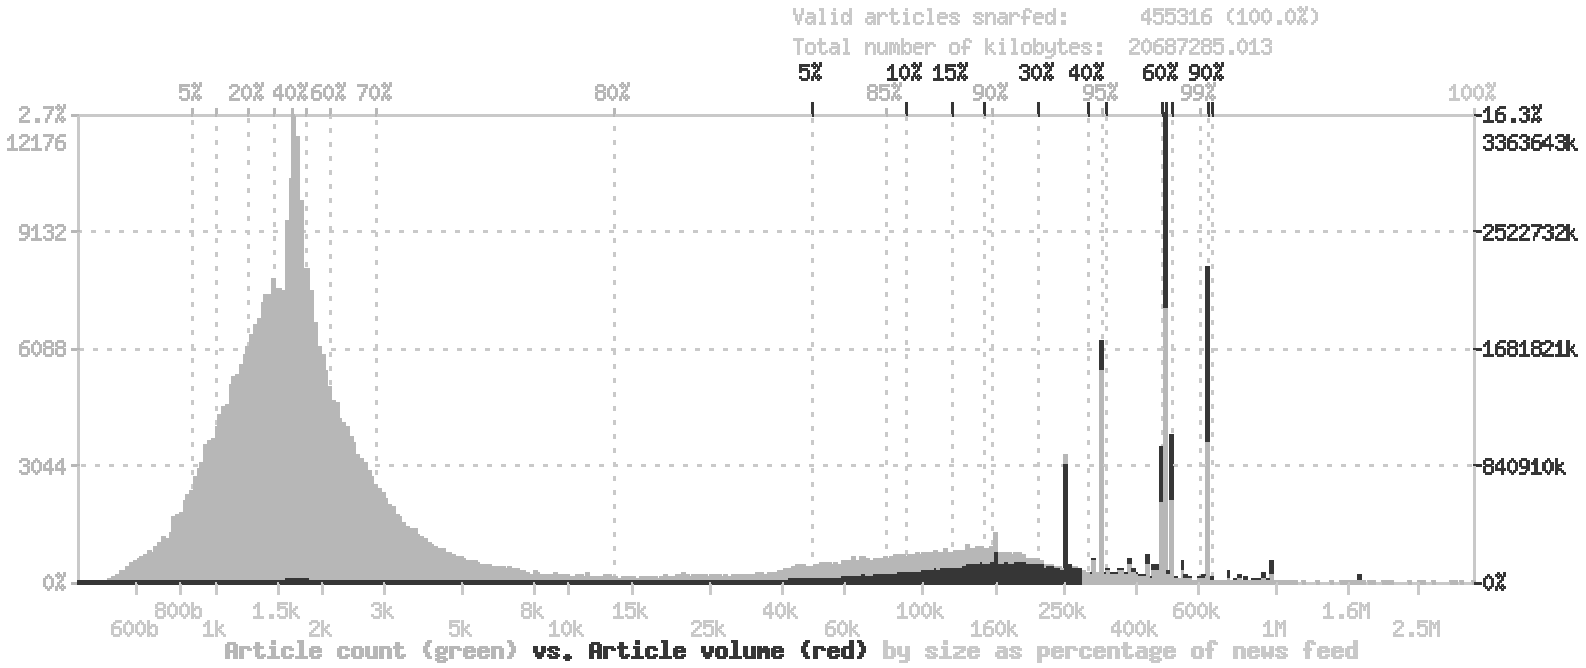

http://news.inet.tele.dk/innreport/news-notice.2001.05.10-04.30.26.html

More About That NNTP Traffic

- We peer with most other Internet2-connected sites which do Usenet, and we also peer with many non-I2 sites (but not over I2, of course)
- Having said that, we should note that we are **NOT** the busiest/highest ranked .edu Usenet site; see http://www.freenix.fr/top1000/ or http://news.anthologeek.net/

.EDU's in the Top 50 of the Freenix 1000

Path Entry	Rank	Weight
news.maxwell.syr.edu	1	17.90
news-spur1.maxwell.syr.edu	11	8.70
newsfeed.stanford.edu	32	4.10
logbridge.uoregon.edu	48	3.49

Freenix data from 2/2001, the most recent data available (that data is consistent with the news.anthologeek.net data for 4/2001)

NNTP Traffic Levels and Traffic Exclusions

- We carry most (but not all) Usenet traffic, excluding, for example, warez (and other obviously problematic traffic), and traffic from chronically open/abused News servers
- Traffic volume is VERY sensitive to the inclusion/exclusion/poisoning of even individual newsgroups or individual servers
- See: http://www.newsadmin.com/

FTP Traffic

- Outbound FTP octets (ports 20/21 TCP) originated 61.6% at UO, and 37% at OSU.
- The #1 outbound ftp source (~16.9% of total ftp octets) was a UO anonymous ftp server which provides access to various software distributions (including Linux),
  #2 outbound (15%) was a major ftp server at OSU which also does software distribution

And The #3 FTP Server...

- … was an etree server (see www.etree.org)
- Etree is an interesting twist on traditional peer to peer music sharing applications…
- It uses good ’ol ftp
- It permits distribution ONLY of freely redistributable/uncopywritten music
- It uses a lossless shn format (disparaging the ‘low’ quality of lossy formats such as mp3’s)

A Couple of Issues To Consider About Etree

- Commodity transit bandwidth…
  etree shn format files tend to be huge, and demand for etree content substantial
- Infrastructural/policy issues… you may not want a passworded etree server popping up w/o permission on a subnet not configured for steady heavy loads, using a DHCP'd address and non-organizational DNS entries

And If Etree ftp'ers Don't Interest/Concern You...

- The number one audio client download at CNET for the week ending May 6th was AudioGalaxy Satellite (by a huge margin), and it uses… drumroll please… port 21 plus high number ports above 41000 (see: http://www.audiogalaxy.com/info/faq/satellite.php)

What About Other Peer-to-Peer Application Traffic?

- Napster and Gnutella were seen, but at relatively low levels (~2% of total TCP octets each).
- The comparatively low level of traffic seen for those applications of course raises the question: WHY are the levels of traffic seen from those applications so much low(er) at Oregon than at some other sites?

Napster/Gnutella May Be Comparatively Low Because...

- We have had a number of high profile online law enforcement incidents here at UO (e.g., the first federal felony conviction under the No Electronic Theft Act), see: http://www.cybercrime.gov/netconv.htm
- … and local busts haven't always been for copyright infringement done over the network; c.f.: http://www.dailyemerald.com/vnews/display.v/ART/2001/02/05/3a7eca642

Other Factors

- Some OGIG partners block/have tried blocking/are attempting to block Napster and Gnutella outright, or are experimenting with technical solutions to manage the traffic associated with them
- And finally, users have gone to other applications -- Napster and Gnutella are only two of many possible options if you have progressive/experimental students.

Kazaa

- Besides AudioGalaxy Satellite (mentioned on slide 40), at least based on what we've seen so far, Kazaa (see: www.kazaa.com) looks like it may be the "heir apparent" to Napster and Gnutella…
- Port 1214...
- You can try a sample search/download directly from their web site...

Hotline

- Another peer-to-peer application you may run into is Hotline (www.bigredh.com) (see also Carracho, www.carracho.com)
- Port 5500 (and others)
- Servers are listed on “trackers”
- Trackers are tracked by “tracker-trackers” including http://www.tracker-tracker.com/
- You can search for Hotline servers living in your netblock by doing a tracker-tracker server advanced search on your netblock

VII. Selected ASN Data

Why Look at ASN Data?

- Traditionally, ASN matricies have been useful in formulating peering policies and understanding traffic sources and sinks
- In this case, we simply wanted some insight into our network partners -- who were our top talking ASNs? Were they folks on the west coast? East coast? Overseas?

Top Inbound ASN Pairs

Source ==> Destination	Octets
UNI-C (1835) ==> UO (3582)	1,939,720,539
UCLA (52) ==> UO	1,575,405,342
Sunet Lulea (2831) ==> UO	1,040,216,294
UC Riverside (6106) ==> UO	604,472,700
Sunet KTH (1653) ==> UO	368,668,278
5Colleges MA (1249) ==> UO	246,658,860
Berkeley (25) ==> UO	240,757,571
VA Tech (1312) ==> UO	198,335,118
U Minn (217) ==> OSU (4201)	188,225,190
Utah Ed Net (210) ==> PSU (6366)	171,940,607

Top Outbound ASN Pairs

Source ==> Destination	Octets
UO (3582) ==> ASN 0	8,010,032,385
UO ==> MoreNet (2572)	1,810,147,864
UO ==> MichNet (237)	1,406,677,275
UO ==> U Alaska (7774)	1,064,602,948
UO ==> Cal State (2150)	1,050,989,542
UO ==> UT Austin (18)	1,043,281,555
UO ==> U Wash (101)	1,000,113,313
UO ==> U Hawaii (6360)	937,166,554

Top Outbound ASN Pairs (cont)

Source ==> Destination	Octets
UO ==> Purdue (17)	937,166,554
UO ==> GARR, Italy (137)	848,855,046
UO ==> SDSMT (11602)	838,399,339
UO ==> UAL Huntsville (10364)	812,746,479
UO ==> U Wash (73)	798,550,299
UO ==> Indiana U (87)	768,213,804
UO ==> 5Colleges MA (1249)	762,682,149
UO ==> Georgia Tech (2637)	753,894,764

A Surprise From the ASN Data

- In looking at the ASN data, while we focused (as we normally would) on traffic volume in octets, we also looked at ASNxASN cell flow counts
- We were surprised to see EBSCO (6932), a commercial online database vendor, in the #2 spot overall… Apparently they are connected via the vBNS+

VIII. Conclusion

What Seems To Be Working Well?

- The applications that we said would probably work well over I2 (see the bottom of dast.nlanr.net/Guides/writingapps.html, a piece adapted from an article I originally wrote for our CC newsletter), have in fact proven to work well on Internet2, and have traffic which is strongly represented in our observed traffic sample (Usenet News, IP multicast, measurement traffic, “ftp”).

“ftp”

- In http://darkwing.uoregon.edu/~joe/how-to-go-fast.ppt (slides 42-46) we had hypothesized that “ftp” traffic, the number two TCP app as reported in earlier CANet3 traffic studies, might not be “typical” ftp. We had hypothesized it might be server-to-server mirroring, but we now know that it is at least in part a variety of P2P audio-related apps generating “ftp”-ish traffic

Sharing Connections With Sponsored Participants Has Been Okay

- When sponsored/secondary participants were granted access to Abilene, there was considerable angst in some quarters that their traffic might prove problematic; based on this traffic study, we see no evidence of problems from sponsored participant traffic (and we believe that SEGP traffic will be similiarly a non-issue)

You Should Pay Attention to Peer-to-Peer Applications

- P2P applications are potentially hugely important on an operational level, albeit more for their impact on commodity transit than on I2 traffic levels per se.
- Attempts to control P2P traffic judicially or via technical means will not be wholly successful, so you may want to consider tracking travel levels on a per-connection basis via routine use of cflow or snmp

Commercial Partners Are Clearly Coming...

- EBSCO was noted this time (via vBNS+)
- For the sake of doing a thought experiment:
  s/EBSCO/Google/
  s/EBSCO/Yahoo/
  s/EBSCO/<your favorite content site>/
- Or how about
  s/EBSCO/@Home/
  s/EBSCO/<your favorite access provider>/

What About Your Flow Data?

- As we noted right up front, every site's data will definitely be different…. YOU should be looking at YOUR flow data, if you aren't already
- We should also note that we'd love to have the chance to work with some suitably sanitized flow data collected from other I2 sites. If you're interested in exploring this, please send me email.

A Plug for a New Measurement And Analysis Opportunity

- Dave Meyer and the Advanced Network Technology Center at the UO Computing Center operate a route viewer at the Oregon IX with views from many locations around the Internet (including an Abilene view)
- We've begun archiving that route view data at http://rv-archive.uoregon.edu/ and you may find it an interesting source of data for analysis

Thanks and Questions

- Thank you for the chance to talk today.
- Questions?