

Summary: OWEN/NERO Bandwidth Audit

Our Charge: In 1999, SB622 directed OUS to complete an audit of bandwidth utilization for the Oregon Legislature and the JLCIMT. Questions included (a) comparison of Oregon's bandwidth with that of other network consortia, (b) development of a bandwidth standard on a per-user basis, (c) examination of how current bandwidth is used, (d) identification of opportunities for improving the efficiency of current bandwidth usage, (e) strategies for satisfying future demand for bandwidth given funding limitations, and (g) a discussion of how the existing network, OWEN/NERO, relates to the new State of Oregon Enterprise Network (SOEN).

Character of the Report: Our full report was prepared with sufficient background material to allow a non-technical reader to understand and assess the data provided. Because that report would be a public record, we crafted it with some care so as to respect user privacy (as mandated by law) while providing the objective information needed for legislative oversight and informed public debate. This executive summary highlights some of the key findings of that two hundred and thirty page document.

Description of The OWEN/NERO Network: OWEN/NERO is Oregon's "Network of Networks," providing intrastate and Internet connectivity for all public (OUS) universities in the state, virtually all public elementary and secondary schools (via OPEN, the Oregon Public Education Network), and all state agencies (via the State of Oregon Department of Administrative Services). All in all, OWEN/NERO services in excess of 620,000 Oregonians (roughly 530,000 K12 students; 60,000 or so OUS students, faculty and staff; and 31,000 state agency employees), plus the members of the public who interact with those OWEN/NERO users. OWEN/NERO has three hub sites, located in Portland, Eugene and Corvallis. Those sites are connected via 45Mbps leased line circuits (with the Portland—Eugene circuit currently in the process of being upgraded to 155Mbps). Internet connectivity is provided via 121 Mbps worth of "commodity Internet transit" purchased from two major tier one service providers, UUNet, and Cable and Wireless. OWEN/NERO also has connectivity to Internet2 via the Oregon Gigapop, and also exchanges customer traffic (without paying any financial settlements) with a number of network providers who are present at the Oregon Internet Exchange (OIX) in Eugene.

Primary Focus of the Report — Peaking Inbound Commodity Internet Transit Bandwidth Usage: While all of OWEN/NERO's connectivity is valuable and important, we focused on *peaking inbound commodity Internet transit bandwidth*. Why? Quite simply *commodity transit bandwidth* is OWEN/NERO's largest expense, and an expense whose magnitude continues to grow (for OWEN/NERO, as for the Internet as a whole). We focused on *inbound* commodity transit bandwidth because inbound load dominates outbound load. Since OWEN/NERO cannot separately provision Internet bandwidth at different levels for inbound vs. outbound load, inbound load (because it is greatest) establishes the bandwidth required. Moreover, we homed in on *peaking* inbound load because load varies greatly during the course of a day, and, like a utility, OWEN/NERO needs to insure that it has

capacity sufficient to meet empirically observed peak demand (or OWEN/NERO's customers will become disgruntled and purchase services from an alternative provider). *Sizing transit circuits to meet inbound peaking traffic loads is thus clearly key.*

The \$1000/Mbps per Month Model: We found that OWEN/NERO has been doing an excellent job of sizing its commodity Internet transit capacity to meet observed demand, in part because OWEN/NERO partners specify and contract for a maximum level of inbound transit traffic, paying \$1000/Mbps per month for that contracted capacity. Sustained load in excess of that customer stipulated rate is dropped. Traffic which benefits all partners (e.g., a web caching traffic, the consortia's centralized news feeds, etc.) is excluded from that costing structure. Partner traffic profiles are available online at: http://www.nero.net/cgi-bin/rddcust.cgi/pritar=Traffic_Profile

On Establishing a Per User Bandwidth Usage or Bandwidth Cost Standard: We looked at the possibility of establishing a per-user bandwidth standard to insure that unrealistically high levels of throughput aren't being provisioned. Looking at current bandwidth usage, OWEN/NERO currently delivers 195 bits per second per user on average (121,000,000 bps/620,000 users) — for comparison, a typical dialup modem today has a nominal speed of 56,000 bits per second. It is hard to conceive of a saleable per user bandwidth standard which would be lower than the existing de facto standard.

Looking at this differently, if each OWEN/NERO user paid even a dollar/month for Internet commodity transit bandwidth, OWEN/NERO would receive \$14.8 million/biennium (vs. its current budget of less than \$2 million/biennium).

OWEN/NERO Transit Bandwidth Compared to Other Network Consortia: We also compared OWEN/NERO transit bandwidth usage to that of other network consortia for which data was available, including CALREN-2 (CA), Great Plains Network (covering universities in AR, KS, NE, ND, OK & SD), MichNet (Michigan), More.Net (Missouri), NCNE (serving five schools in PA & WV), NCREN (North Carolina), Net.Work.Virginia (Virginia), the Washington State K-20 Network, and WiscNet (Wisconsin). While none of those networks are a perfect comparator for OWEN/NERO, in all cases except that of the Great Plains Network, those consortia had commodity Internet transit bandwidth greater than that deployed by OWEN/NERO. For example, North Carolina's NCREN has 465 Mbps worth of commodity transit bandwidth and Washington State's K-20 network has 445 Mbps (vs. OWEN/NERO's 121 Mbps). Clearly, OWEN/NERO's transit bandwidth usage is quite modest relative to other statewide network consortia for which data was available.

The Flow Study: In order to characterize OWEN/NERO's existing peaking inbound commodity transit traffic, we sampled inbound network traffic flows from 2:00-2:15PM on April 6th and 7th, 2000. Samples were drawn without notice to avoid intentional or unintentional changes in user behavior. Those two fifteen minute samples from the UUNet and CWIX circuits resulted in a total of 1,761,170 flows for analysis. Statistics associated with those samples were consistent with those reported by Kevin Thompson for the MCI backbone as of 1997, and with

McCreary & Claffy's more recent 10 month CAIDA study of the Ames Internet Exchange ending in March 2000.

Flow Data Breakdown By Application: We looked at that flow data both on a per flow basis, and on an octet-weighted basis. On a per flow basis, nearly 3/4ths of all flows sampled were http (e.g., world wide web), which is very consistent with Thompson's reported value of 75%. Roughly 9% of all OWEN/NERO flows sampled were domain name system related; no other single application accounted for a significant number of flows (e.g., >=5% of all flows). On a per octet basis, 63.3% of all octets were http (e.g., world wide web); the only other single application accounting for more than 5% of inbound octets was nntp (e.g., Usenet News), at 10.0%. For comparison, McCreary and Claffy found that 58.9% of the NASA Ames Internet Exchange traffic (by octets) was http, and 11.7% was nntp. OWEN/NERO's traffic composition (on a per application basis) is thus very consistent with what the NSF's Cooperative Association for Internet Data Analysis reported having seen at NASA Ames.

Flow Data Breakdown by Source Autonomous System: On a per flow basis, 13.3% of incoming flows were from Exodus, 7% were from AOL, 6.4% were from GlobalCenter, and 5.6% were from Abovenet. No other single ASN accounted for at least 5% of all flows. On an octet weighted basis, only two ASNs had at least 5% or more of all inbound octets: Exodus (12.5%) and GlobalCenter (7.3%). (Abovenet just missed the 5% threshold at 4.9% of all octets; AOL had 4.3%). This data indicates that no single ASN accounted for a disproportionate amount of incoming traffic.

Most Popular Web Sites for OWEN/NERO Customers: Next, we narrowed in on just web-related traffic (since it accounted for roughly 3/4ths of all inbound flows, and nearly 2/3rds of all inbound octets). Getting a definitive understanding of what web sites OWEN/NERO customers visit is difficult for a variety of reasons, including the fact that (a) flow data files record raw numeric network addresses for each flow, but raw numeric network addresses may be associated with 10's or 100's of "virtual web hosts" (with no indication in the flow data of which host might actually have been visited); (b) we made a decision (in consultation with OUS) to limit analysis of individual web-related network addresses to those accounting for at least 0.1% worth of total web traffic, (c) the fact that many OWEN/NERO partners are using web caching (which means that a single page retrieval may satisfy many end users' subsequent requests for a popular page), and (d) the fact that material amounts of web traffic may be flowing via Internet2 or the Oregon Internet Exchange (for example, Verio (now part of NTT) is the world's largest web hosting company, but connects to OWEN/NERO via the OIX rather than via our commodity transit links, and hence that traffic will not be reflected in our analysis).

Most Popular Web Site Categories on a Per Flow Basis Notwithstanding those caveats, on a per flow basis, the web content category with the largest number of flows was the web advertising/online profiling category, comprising over 13% of all web flows. (Yes, there really are a ton of banner advertisements and a lot of web cookies being shoved at users by web sites out there!) The 2nd largest category on a per flow basis, at nearly 9% of all web flows, was distributed content delivery, virtually

entirely associated with Akamai (see: <http://www.akamai.com/>). Akamai will soon be at the Oregon IX, so soon their web pages will no longer factor into OWEN/NERO's transit bandwidth usage. The only other categories accounting for >=5% of web flows were the "mega Internet Service Providers" or default web portal sites, e.g., AOL/Netscape and Microsoft (collectively totalling 7.8%), and the search/web directory sites such as Yahoo and Excite (collectively totalling 7.3%). All other categories accounted for <5% of web flows. On a per flow basis, web flows associated with identifiable adult web sites amounted to less than 1% and web flows associated with identifiable hacker/cracker web sites amounted to only 0.1%.

Most Popular Web Sites Categories on an Octet Weighted Basis: On a per octet basis, the largest category of web traffic was distributed content delivery (e.g., Akamai) at just over 7%. The only other categories of web traffic accounting for at least 5% of web traffic octets were: web file sharing sites (5.4%), mega ISP/default portal sites (5.2%), and search/web directory sites (5%). On a per octet basis, web traffic associated with identifiable adult sites amounted to less than 1%, and web traffic associated with identifiable hacker/cracker sites amounted to only 0.1%

What about Napster? Despite media reports that Napster has been consuming disproportionate amounts of bandwidth at some sites, OWEN/NERO's identifiable Napster traffic totalled only 3.2% of total octets, a value consistent with what McCreary and Claffy saw at Ames. OWEN/NERO Napster usage may be lower than what some other sites have experience because some OWEN/NERO partners have banned Napster outright, while others have established a reputation for zero tolerance of copyright infringement (e.g., UO cooperated in the first successful federal felony criminal prosecution under the No Electronic Theft Act).

What about Online Games and Chat? Only 1.4% of all OWEN/NERO traffic (measured in octets) was associated with games/chat, a generally negligible amount of traffic, we believe.

Active network measurement programs: The Legislature should also know that OWEN/NERO partners actively participate in a number of ongoing network measurement programs, including the NLNR Active Measurement project, the Advanced.Org Surveyor project, the LBL NIMI projects, etc. Those project consistently show Oregon has excellent network performance.

Managing The Growth Challenge: Success in attracting network peers to the Oregon IX allows Oregon a way to add capacity without adding cost, and the \$1000/Mbps per month charge encourages customer bandwidth conservation and efficiency.

OWEN/NERO vis-a-vis SOEN: OWEN/NERO has been designated by DAS as the likely SOEN network service provider.

Efficiency Strategies Potentially Worth Exploring: Finally, peak shaving (shifting demand from peak times to off peak times via deployment of modems or use of flex time), promoting voluntary use of web caching by partner sites, and use of satellite services for web cache preloading and Usenet news distribution all have some potential for increasing OWEN/NERO's efficiency, although the network is already very, very, efficiently engineered.