Section 1. Introduction

Section 1 Keypoints:

- This bandwidth audit has been prepared pursuant to Oregon Senate Bill 622 (1999).

- The areas of interest addressed by this report include:

    - a technical narrative description of the OWEN/NERO network
    - a comparison of usage by OWEN/NERO to that of comparable networks
    - consideration of a “bandwidth usage standard per user,” or a “bandwidth cost standard per user”
    - analysis of what OWEN/NERO’s bandwidth is being used for
    - identification of any opportunities which might result in improved efficiency for OWEN/NERO’s use of bandwidth
    - consideration of bandwidth demand growth, and strategies for meeting that demand, given limited opportunities for incremental legislative funding
    - the relationship between OWEN/NERO and the new State of Oregon Enterprise Network

- The approach taken to this report is objective (descriptive, normative and comparative), rather than subjective and proscriptive; technical recommendations for OWEN/NERO future directions are not included in this report.

- Additional limitations on the scope of this report include:

    - We’ve not considered LAN or campus infrastructure-related bandwidth issues.
    - We’ve not considered annecdotal network performance data, nor did we survey end users about network bandwidth sufficiency or lack thereof.
    - We’ve not considered Internet2/high performance research network connectivity, nor connectivity obtained via local exchange points (such as the Oregon Internet Exchange).
    - Because of the unlimited distribution of this report, we crafted our report 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 report was prepared with the expectation that its audience would have a non-technical background; consequently, sufficient technical background information has been provided to allow a non-technical reader to fully understand and properly interpret the data that has been provided.

Legislative Context for This Bandwidth Audit

Conference Committee No. 2 Amendments to the A-Engrossed version of Senate Bill 622 of the 1999 Legislative session requires that "The Department of Higher Education shall complete an audit of bandwidth utilization and report to the Joint Legislative Committee on Information Management and Technology during the Seventy-first Legislative Assembly in the manner provided in ORS 192.245."[1]

ORS 192.245 reads:

> 192.245 Form of report to legislature. Whenever a law of this state requires a written report be submitted to the Legislative Assembly, the requirement shall be met by distribution of an executive summary of no more than two pages sent to every member of the Legislative Assembly and one copy of the report to the office of the Speaker of the House of Representatives, one copy to the office of the President of the Senate and five copies to the Legislative Administration Committee. This requirement does not preclude providing a copy of any report to a specific legislative committee if required by law.

Generally Understood Areas of Audit Interest

Based on committee member comments during the January 2000 hearing of the Oregon Joint Legislative Committee on Information Management and Technology (JLCIMT) and based on earlier informal discussions, we were aware of some generally understood areas of audit interest including:

— What is this thing we call "OWEN/NERO"? Where does it go? Who does it serve? What does it cost? (a technical narrative description of the network being examined)

— Is OWEN/NERO's actual aggregate bandwidth consumption consistent with that of comparable network consortia elsewhere? (normative/comparative macroscopic bandwidth study)

— Has OWEN/NERO deployed a reasonable amount of bandwidth per user, or are inappropriately high levels of bandwidth being delivered to users? (e.g., endeavor to develop a bandwidth standard per user, or a cost standard per user for wide area bandwidth costs)

— What is OWEN/NERO's bandwidth being used for? Is it being used for purposes consistent with the missions of the attached agencies? Do agencies have acceptable use policies and are they being enforced? (study of usage controls)

---

1. http://www.leg.state.or.us/99reg/measures/sb0600.dir/sb0622.a2c.html

— What (if anything) can be done to use OWEN/NERO's bandwidth more efficiently? (e.g., will a process study identify opportunities for improved efficiency?) Is OWEN/NERO bandwidth demand going to continue to grow without limit? If so, how does OWEN/NERO's operators propose to meet that demand, recognizing that state support for networking is only one of many areas competing for legislative financial support? What about OWEN/NERO and its relationship to the new State of Oregon Enterprise Network (SOEN)?

The remainder of this report addresses those questions subject to the limits described below.

Descriptive/Normative/Comparative Emphasis

We should also note that in this study we've endeavored to objectively approach OWEN/NERO's bandwidth data from three particular angles:

— We have endeavored to describe the network usage that we see today,

— We have endeavored to develop norms or standards for current usage to serve as a benchmark for future OWEN/NERO bandwidth audits (if any)

— We have endeavored to compare OWEN/NERO's network traffic with other network traffic reports so as to provide context for observed OWEN/NERO performance

What Has Been Excluded From Analysis

We have intentionally avoided a (necessarily subjective) prospective approach, and hence we will not be offering opinions about substantive future directions for OWEN/NERO. Opinions about what should be done with the OWEN/NERO network in the future will necessarily vary from person to person, and are properly the subject of executive decision making, and hence are beyond the scope of this review.

We also wish to note some additional limits on the scope of this report:

— We did not considered LAN (local area network) or campus infrastructure-related bandwidth issues. That area has been defined by OUS as being outside the scope of this study, because the problems there are largely well understood by the responsible parties, and because remediating those particular issues is largely a matter of identifying sufficient funding.

— We did not consider any end user annecdotal network performance data, nor did we systematically survey end users about network bandwidth sufficiency or lack thereof. Again, that area was defined to be outside the scope of this study, and the large number of potentially confounding causes for observed poor performance (if any) to a given remote site makes it hard to effectively employ annecdotal or survey research data for troubleshooting and problem isolation.

— Nor did we consider Internet2/high performance research network connectivity in depth. That connectivity has a known, fixed cost determined as a condition of grant funding received from the National Science Foundation, its usage is constrained to a limited number of sites having a research and educational character, and its bandwidth appears to be sufficient to meet forseeable requirements.

— We also have not studied connectivity obtained via local exchange points (such as the Oregon Internet Exchange) which has no associated direct costs at any depth.

— We were also quite mindful of the potentially broad distribution of this report. A broadly circulated document must necessarily be crafted with some care so as to respect users privacy while accomplishing this document's underlying goal of providing the objective information that's needed for legislative oversight and informed public debate. The Oregon University System has requested and received a legal opinion that the information presented does not compromise individual privacy rights protected under FERPA, the Electronic Communications Privacy Act or other applicable statutes.

This Report's Anticipated Audience

In preparing this report, we did so with the expectation that its audience would be diverse and largely comprised of non-network engineers, including (but not limited to) Oregon legislators, senior administrators within the Oregon University System, interested members of the press, and members of the general public.

As a result, and because of the inescapably technical nature of the material being covered, we've included sufficient technical background information to allow an interested reader to fully understand (and properly interpret) the data that has been provided. Thus, in the next section, we begin with a broad discussion of network bandwidth. Where that information recapitulates material which a particular reader has already mastered, please accept our apologies — we would rather briefly bore some than leave others without the introductory foundation that later material requires.

Section 2. Understanding the Basics of Network Bandwidth and Its Measurement

Section 2 Keypoints

- Section 2 introduces basic network engineering background information the reader will need to know.

- Bandwidth nomenclature:

  - bandwidth is measured in bits per second (bps).
  - commonly encountered units of bandwidth include kilobits per second (1,000's of bits per second, or Kbps), and megabits per second (1,000,000's of bits per second, or Mbps)
  - eight bits make up an octet (or one byte). A page of text is about 2000 octets long.
  - data is shipped across the network in chunks called packets, not as individual bits or octets.
  - a wide range of data transmission speeds are in use today. Commonly encountered speeds include: 56Kbps modems, 1.544Mbps T1s, 10Mbps ethernet, 45Mbps DS3s, 100Mbps fast ethernet, and 1000Mbps gigabit ethernet.

- Circuit types: in addition to having different nominal speeds, wide area connectivity can be provisioned in a variety of different ways, including frame relay, point-to-point circuits and via colocation. These different provisioning methods have widely varying costs, and the method used to effect a connection can also affect how much actual bandwidth is available on a sustained basis from that connection.

- Flows: flows are unidirectional sequences of packets going between two points on the network; the core of this report's usage analyses are done on flows.

- Flows have numeric source and destination addresses associated with them; those addresses (such as 128.223.142.13) are called "dotted quads."

- Dotted quads can sometimes (but not always) be mapped to symbolic internet addresses (fully qualified domain names, or "FQDNs," such as www.uoregon.edu)

- Some addresses are assigned dynamically (via a protocol called DHCP) to different users at different times; these dynamic addresses are common and useful, but can complicate traffic analyses.

- Dotted quads are assigned in chunks called "network blocks." Network blocks vary widely in size (from dozens of addresses to millions of addresses). A given organization may be assigned multiple non-contiguous blocks. Other organizations may informally be permitted to use parts of blocks that aren't formally assigned to them. For these reasons, netblocks are generally a poor unit for analysis.

- ASNs (“autonomous system numbers”) are a more commonly used network aggregate, and represent a “connected group of IP networks that adhere to a single and clearly defined routing policy.” This report will include analysis of traffic sources on an ASN basis.

- Ports: Each flow, in addition to having a source address and a destination address, has a source port and a destination port. Each service that a system offers over the Internet (such as world wide web pages, or electronic mail) is offered via a specific (generally agreed upon) port. Port numbers are the primary way we have of inferring the type of traffic a particular flow represents, however port numbers are not always used consistently, some port numbers are available for dynamic assignment for a multiplicity of uses, and in some cases one application may masquerade itself by running on another application’s port, making application analysis by port number an imprecise art at best.

- In general, flows are independent of each other, however some emerging applications (such as Napster) require sequential analysis of multiple flows in order to identify traffic associated with a particular application.

- Flows can take place via either TCP or UDP protocols. A port operating at a given value using TCP can be offering an entirely different application than that same port using UDP. A programmer’s choice of TCP or UDP also affects how fast an application can go.

- In addition to the ports associated with network flows, there are also “ports” on network hardware into which network cables get inserted. Some types of network monitoring watch the traffic level coming out particular network hardware ports, and the proper analysis of that data requires knowledge of what cable has actually been plugged into each such port.

How is network bandwidth measured?

Bandwidth rates are normally measured in terms of “bits per second” or “bps.” A bit is a single binary digit (or the fundamental ability of a circuit to be on or off, and thereby convey information).

Each character included in a web page or email message normally requires the use of eight bits to represent a particular letter, number, or special symbol, and such a character is normally referred to as a “byte” or as an “octet.” The availability of eight bits per character means that a total of 256 (2^8) unique symbols (letters, numbers, etc.) can theoretically be represented by an eight-bit octet.[2] To obtain an expanded representational range, some non-roman character sets (e.g., Chinese and Japanese) are written using double byte character sets,[3] but we do not need to consider that issue further for our purposes here.

For reference and familiarization purposes, a typical double-spaced page of typewritten text is normally about 2000 octets or 16,000 bits long, a typical floppy disk holds 1.4MB, a typical CDROM holds 650MB, and desktop class 26GB IDE hard drives are now routinely available for less than $200 retail. 36GB and 72 GB SCSI drives for servers are now routinely commercially available, too.

For convenience, when referring to higher bandwidth rates, thousands of bits per second are normally called “kilobits per second” and are written in abbreviated form as ‘kbps.’ If we have a million bits per second, that’s normally called a “megabit per second” or “mbps.” A billion bits per second is referred to as “gigabit per second” or “gbps.”

If we’re referring to octets or bytes, we typically talk about kilobytes per second (kBps), megabytes per second (mBps), or gigabytes per second (gBps). Note that the “B” in each of those abbreviations is formally capitalized, representing bytes, rather than lowercased (which would represent bits).

We should also mention that some people strictly define “kilo” to mean 1024 (2^{10}) rather than 1000 as a multiplier, and “mega” to mean 1,048,576 (2^{20}) rather than 1,000,000, and “giga” to mean 1,073,741,824 (2^{30}) rather than 1,000,000. In general, we will use the decimal rather than binary definitions of those term, and indicate that usage by capitalizing the “K,” “M,” or “G” in abbreviations.

Another commonly seen network term is “packet.” When data is sent over the network, it doesn’t go as individual bits, or individual octets, rather it is sent in clumps called packets. A packet consists of a payload

---

2. For the purposes of this discussion, we also ignore the issue of parity bits and start and stop bits, which can effectively push the bits per character up above 8 bits, and we will also disregard 7 bit ASCII and EBCDIC encodings.
3. See, for example, CJKV Information Processing, Ken Lunde, O’Reilly, Sebastopol CA, 1999.

(one or more octets worth of actual data) and a header. The header contains information about where a packet's from, and where it's going, and represents "overhead" which potentially reduces the actual amount of information that can be conveyed. Packet sizes generally range from forty octets to around 1500 octets. We introduce the concept of packets here because some network studies you may see report traffic in terms of packets; we belive reporting traffic in terms of octets is a more readily comprehended unit of measure.

Finally, we'd like to note that bandwidth is a rate per unit time, not a measure of total bits transferred. Standard network business practice is to sell a circuit of a given capacity, the customer paying the same whether that circuit is completely quiescent or operating at one hundred percent of capacity around the clock. You pay for the size of the pipe, not how much flows through it. For that reason we do not focus on total traffic transferred, but rather the usage at peak times which determine the capacity which we need to provision.

What speeds are commonly used?

Network circuits are available in a variety of different speeds. The current, historical and now emerging commonly encountered speeds for data transmission are:

45 bps: 45 bps is about the slowest speed communication speed that has been in routine historical/current use. TDD's (Telecommunication Devices for the Deaf) communicate at this speed in the United States.[4] And yes, TDDs running at 45 bps are still in widespread/routine use today.

110 bps: The speed of old ASR 33 Teletypewriters (TTY's).[5] TTY's date from the late 1960's, and were commonly connected to phone lines using an acoustical coupler.[6] 110 bps speeds are not routinely in use today.

195 bps: **<u>Current OWEN/NERO bandwidth/user</u> (computed by dividing total commodity transit bandwidth by the number of OWEN/NERO users being serviced; see page 85, below.)**

4. http://tap.gallaudet.edu/TTY-basics.htm
5. http://www.telnet.hu/hamster/pdp-11/egyeb/asr33.jpg
6. Because at that time customer installed equipment was not allowed to interconnect electrically with the phone system, acoustical couplers with a pair of rubber "cups" were used to allow the microphone and speaker in the phone handset to make a transient acoustical connection instead.

300 bps: Bell 103 or V.21 modem standard. This is the speed of LA36 DECWriterII printing terminals (ca. 1975); the original Hayes 300 baud modems were introduced in 1981. While 300 bps was a big improvement over 110 bps TTY's, 300 bps was still horribly slow. 300 bps speeds aren't routinely used today, except for some high frequency packet radio applications.[7]

1200 bps: Bell 212A or V.22 modem standard. At 1200 bps, video terminals such as the Televideo 910 (circa 1982) became a popular choice. Obsolete today, except for some HF packet radio applications.

2400 bps: The V.22bis modem standard. Rarely seen in use today.

9600 bps: V.29 standard (Group III fax)/V.32 standard (modem). Still in very common use as a fax transmission standard today.

14.4kbps: V.32bis standard. The slowest speed modem still routinely encountered in actual routine use; at the end of its practical life.

28.8kbps: The original V.34 modem standard dating from the spring of 1995. Still quite common.

33.6kbps: The improved V.34 modem standard dating from the fall of 1996. Very common. The fastest speed supported by analog connections.

56kbps: The V.90 modem standard, the currently-prevailing best-available dialup modem service. Adopted in 1998 by the ITU. FCC power regulations cap throughput at 53kbps, and in practice, due to line quality issues, users may only see 44Kbps or less.

DS0: The slowest speed commonly provisioned leased line/frame circuit. A DS0 64kbps circuit delivers 56kbps of usable capacity, or roughly the same speed as a typical current generation dialup modem. A DS0 equals one basic copper phone line, just like the ones used for residential phones.

---

7. See, for example: http://www.mfjenterprises.com/packradio/mfj1278b.html

ISDN:	ISDN BRI (Basic Rate Interface) — ISDN 2B+D service delivers 128Kbps. An unpopular service that never really caught on in Oregon due to its tariffing.
xDSL:	As normally provisioned, xDSL typically delivers 256kbps (however asymmetric rates of up to 7Mbps may potentially be available to some customers for an additional fee).
NxDS0:	Multiple DS0 circuits can be combined, or "inverse multiplexed" together to form a circuit with larger than DS0 capacity (this is normally only done for 384Kbps video-conferencing circuits).
T1:	**1.544 Mbps, or 24 times a DS0. This is the speed at which many individual K12 schools or smaller colleges connect.**
NxT1:	multiple T1 circuits can be combined, or "inverse multiplexed" together to form a circuit with larger than T1 capacity (this is normally only done for two to six T1s at the most (e.g., for 3Mbps to 9Mbps speeds).
Ethernet:	**10Mbps (the speed of most desktop network connections).**
11Mbps:	The speed of 802.11b standard high speed wireless equipment[8] currently being rolled out at UO and many other sites.
Fractional DS3	A DS3 circuit can carry up to 45Mbps, however many times sites buy only part of the full capacity of that circuit; fractional DS3's rates generally go from 12Mbps up to 42Mbps in 3Mbps increments.
DS3:	**44.736Mbps ("45Mbps") or 28 times a T1. This is the speed of the OWEN/NERO CWIX connection at this time, and the speed of our intrastate backbone circuits. (The Eugene-Portland backbone circuit is being upgraded to OC3).**
NxDS3:	Just as multiple T1 circuits can be combined, so too can multiple DS3 circuits. For example, the OWEN/NERO UUNet connection is 76Mbps.

8. See, for example: http://www.wavelan.com/products/

Fast Ethernet:	**100Mbps, or 10 times the speed of "regular" ethernet. This is the speed at which many network servers and a growing number of desktop workstations connect to the network.[9]**
OC3:	**155Mbps, or 3 times a DS3. This is the speed of the Oregon Gigapop's Sacramento and Denver links to Internet2. (OC3 is the smallest capacity circuit available from Abilene).**
OC12:	622Mbps, or 4 times an OC3. (This is one of two other speed circuits available from Abilene).
Gigabit Ethernet:	**1000Mbps, or 10 times the speed of fast ethernet. This is the speed at which OWEN/NERO customers' fastest servers connect to the network today; for example, the University of Oregon's large shared hosts connect to UONet via gigabit ethernet.**
OC48:	2.4Gbps, or 24 times the speed of fast ethernet. (This is the speed of the Abilene national backbone, and the third type of end-site or gigapop connection available from Abilene today)
OC192:	10Gbps, or 100 times the speed of fast ethernet. (This is the speed of NTONC[10] in Portland) — OC192 speeds are still quite uncommon except for the very largest network service providers and some experimental networks such as NTONC).
10 Gigabit Ethernet	Emerging follow-on to the gigabit ethernet standard; however, final standards for 10 gigabit ethernet have not yet been produced, and 10 gigabit ethernet hardware isn't routinely available as production equipment yet.

9. Fast ethernet speed connections are becoming increasingly routine because fast ethernet cards are down to $20/each from some vendors, and most popular ethernet switches (such as the HP 4000M) come stock with 10/100 ports capable of delivering either 10Mbps or 100Mbps service, whichever is wanted.
10. http://www.ntonc.org/

What other choices need to be made when provisioning a circuit?

Besides deciding on the capacity of your connection, you also need to select the technology that you'll use to make the actual connection. The most common options are:

Frame Relay: In Oregon, commonly used for 56Kbps through DS3 class circuits that have intermittent or "bursty" traffic. CIR ("committed information rate") often runs only 50% of nominal rate, although clear channel (full bandwidth) frame relay circuits are also available. Customarily priced in a distance insensitive way. User connects to a telco provider's frame relay "cloud," which in turn connects to the Internet service provider's point of presence (POP). Frame relay is the technology used in the State of Oregon's Fast Packet contract.

Dedicated Point-to-Point Circuit: Commonly used for T1, DS3, OC3, and faster circuits. Conceptually, the user contacts the incumbent local exchange carrier (ILEC) or a competitive local exchange carrier (CLEC) and arranges to lease a dedicated circuit of a specified speed between two specified points (normally between the customer's premises and the Internet service provider's POP). The fee paid for a dedicated point-to-point circuit reflects both the speed of the circuit and the distance between the two points it connects. The user typically is guaranteed the availability of the full bandwidth associated with that circuit. Used for OWEN/NERO backbone circuits and elsewhere within OWEN/NERO.

Colocation: If the customer's premises and the service provider's POP are in the same physical facility (the facilities are "colocated"), connection may be effected by simply using Cat 5 twisted pair cable (for most 10Mbps or 100Mbps connections), or by using multimode or single mode fiber optic cable (for comparatively longer runs at 10Mbps or 100Mbps, or for most connections made at gigabit speeds). Obviously colocation is the cheapest and easiest solution, where it is available as an option.

Are there additional network concepts I need to be familiar with?

Yes. There are a number of additional network concepts you'll need to familiarize yourself with in order to understand the remainder of the bandwidth audit report.

Network Flows

For example, you should be familiar with the concept of a network traffic flow.

Think of a network traffic flow as consisting of a series of packets (chunks of data) going between two locations on the network. All the packets in a given flow may follow the same path, just like water in a well-worn streambed, or some packets may take one path while other packets may take a different route, like water flowing through a network of coastal tributaries.

There are a number of attributes associated with each network flow. First, just like a flow of water, each flow of network packets has a "direction." Packets flow from their source to their destination. Most network applications actually require network flows in both directions, however those flows can and should be conceptualized as two independent flows, one in each direction, rather than a single bidirectional pipe.

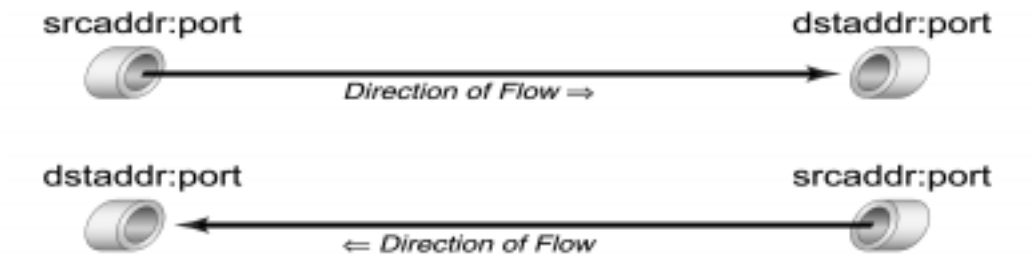

Network Addresses

The flow source and the flow destination each have an address, customarily abbreviated as the srcaddr and the dstaddr.

These addresses are normally represented as four integer values (with each integer value ranging from 0 to 255), separated by dots. For example: 128.223.32.18 is the address of a system at the University of Oregon. Addresses represented in this fashion are said to be written as "dotted quads."

Many (but by no means all) numeric addresses also have symbolic names associated with them. For example, 128.223.32.18 has the symbolic name of "oregon.uoregon.edu" Symbolic names of this sort are normally called "fully qualified domain names" or FQDNs. In this example, for the host 128.223.32.18, oregon is the name of the system, uoregon.edu is the 2nd level domain name, and .edu is the top level domain name (or TLDN).

Traditional major TLDNs have been:

```
.com        Commercial organizations
.edu        Colleges and universities
.gov        US government
.mil        US military
.net        Networks
.org        Miscellaneous organizations
```

You will also see two letter ISO geographical domain names, e.g., us for the United States, de for Deutschland (Germany), es for Espanol (Spain), etc.

Addresses of K12 institutions are often of the form <something>.k12.<statename>.us (for example www.pps.k12.or.us). Addresses of city governments often use names of the format <something>.ci.<cityname>.<statename>.us (for example: www.ci.eugene.or.us) Addresses of state agencies often take the form <something>.state.<statename>.us (for example www.state.or.us).

Mapping Dotted Quads to FQDNs and Vice Versa

A given dotted quad may be associated with more than one FQDN, and a FQDN may be associated with more than one dotted quad. Similarly, a single given system may have multiple dotted quads associated with it and/or multiple fully qualified domain names.

On an interactive basis, dotted quads are normally translated to FQDNs using a program called nslookup. For example:

```
% nslookup 128.223.32.18
Server:     phloem.uoregon.edu
Address:     128.223.32.35

Name:       oregon.uoregon.edu
Address:     128.223.32.18
```

When one is converting a list of dotted quads, it is easy to write Perl code that will do that task on an automated (non-interactive) basis with output suitable for incorporation into later analyses.

We do not want to leave the impression, however, that converting dotted quads to FQDNs (or vice versa) is always a cut-and-dried routine matter, because it is not.

For instance, in some cases, a dotted quad may have NO fully qualified domain name associated with it — those types of dotted quads are normally called “unresolvable” address. There is nothing inherently improper about this — there is no Internet requirement per se that a given dotted quad also have a FQDN associated with it.

In other cases, FQDNs and dotted quads may map in ways you might not normally expect. Consider the following examples:

— the FQDN alpha.uoregon.edu resolves to **five** dotted quads (128.223.142.112, 128.223.142.106, 128.223.142.113, 128.223.142.109, and 128.223.142.111), and each of those dotted quads is a physically distinct machine (the five machines work as part of a computational cluster, and the determination was made that, in general, users connecting to the system should be distributed in a round-robin fashion among the five machines)

— the dotted quad 128.223.142.17 resolves to the primary fully qualified domain name of waterfall.uoregon.edu, but that machine is also known as search.uoregon.edu (This system became the University’s search engine after it was already running with the name “waterfall.uoregon.edu.” Because the name “waterfall” had no intrinsic association with web searching, we decided to add an extra system name (a “CNAME” in domain name system “lingo”), search.uoregon.edu, that would be more meaningful for users.

— the system darkwing.uoregon.edu has the primary address of 128.223.142.13, but that system also has additional IP addresses bound to it:

```
128.223.142.11 (lists.uoregon.edu)
128.223.142.12 (math.uoregon.edu)
128.223.142.21 (pdx.uoregon.edu)
128.223.142.22 (micro.uoregon.edu)
128.223.142.23 (cc.uoregon.edu)
128.223.142.24 (alumni.uoregon.edu)
128.223.142.25 (cas.uoregon.edu)
128.223.142.26 (bachfest.uoregon.edu)
128.223.142.27 (virtual-www.uoregon.edu)
128.223.142.28 (president.uoregon.edu)
128.223.142.29 (kwaxfm.uoregon.edu)
128.223.142.30 (dailyemerald.uoregon.edu)
```

— multiple virtual web servers may also “live on” a single IP address; for example:

> aaup.uoregon.edu, adaptive-tech.uoregon.edu, admissions.uoregon.edu, assembly.uoregon.edu, chem.uoregon.edu, chtl.uoregon.edu, clubsports.uoregon.edu, comm.uoregon.edu, continuo.uoregon.edu, craftcenter.uoregon.edu, culturalforum.uoregon.edu, deptcomp.uoregon.edu, directory.uoregon.edu, diversity.uoregon.edu, duckhunt.uoregon.edu, economics.uoregon.edu, emuchildcare.uoregon.edu, emufoods.uoregon.edu, financialaid.uoregon.edu, gensci.uoregon.edu, geology.uoregon.edu, giving.uoregon.edu, govt-aff.uoregon.edu, greeklife.uoregon.edu, healthcenter.uoregon.edu, uoig.uoregon.edu, lifesci.uoregon.edu, materialscience.uoregon.edu, microlab.uoregon.edu, natural-history.uoregon.edu, oracrao.uoregon.edu, osrp.uoregon.edu, outdoorprogram.uoregon.edu, philosophy.uoregon.edu, registrar.uoregon.edu, researchpark.uoregon.edu, scheduling.uoregon.edu, senate.uoregon.edu, staroffice.uoregon.edu, students.uoregon.edu, studentunion.uoregon.edu, studyabroad.uoregon.edu, telecom.uoregon.edu, uocard.uoregon.edu, uoexp.uoregon.edu, uoma.uoregon.edu, uosummer.uoregon.edu, uofamily.uoregon.edu, wfrn.uoregon.edu, and y2k.uoregon.edu

all are serviced from the address 128.223.142.27, which resolves to virtual-www.uoregon.edu (which is in turn an alias for darkwing.uoregon.edu).

To briefly recap, some important points to understand about network addresses include:

— while every network flow has a source and a destination address, those dotted quad addresses don’t necessarily resolve to a symbolic FQDN -- many will, but some will not

— even if a dotted quad does resolve, it may not be possible to tell which of several addresses is the “right” one to associate with a given flow -- if you use a web browser to access 128.223.142.27 using the symbolic name of admissions.uoregon.edu you get a far different set of web pages than you do if you access it using the symbolic name of uoma.uoregon.edu

— a given system may be associated with multiple addresses, and it can be very difficult if not impossible to determine which addresses should be aggregated to represent the total flow associated with a particular system

Nonetheless, because we know that most casual users think of Internet sites in terms of their domain names, in spite of all these shortcomings, we have provided summaries by domain names as part of this study. We do urge, however, that you recognize the limitations inherent in trying to associate symbolic addresses with dotted quads for the purpose of analyzing bandwidth usage.

A Special Case: DHCP and Dynamic Addresses

We should also mention one particular problem associated with mapping dotted quads or FQDNs to particular users or particular systems, and that's the problem of dynamic addresses assigned by DHCP (Dynamic Host Configuration Protocol).[11]

DHCP is designed to solve a particular problem, that of configuring hosts that need a network address temporarily, but don't necessarily need the same one from run to run, or any particular one. That is, assume you are connecting a desktop workstation to the network. You never plan to connect to the desktop workstation, you only plan to use it as a place to connect from when accessing content on the Internet.

You could assign that workstation an invariant IP address (a so-called "static IP address"), just as you would a server that's always up and constantly seeing service requests from all over the world, but that requires configuring the workstation to use that particular address, a sometimes tedious task that is error-prone and somewhat challenging for non-technical users. Moreover, in some situations (such as users dialing in to large blocks of modems), users might need to change their configuration every time they dialed in! Ugh!

DHCP solves that problem by automatically assigning users a temporary address, a "dynamic" IP address. This is very convenient for the user, and also is quite efficient from the point of view of the network adminisitrator.

The problem that dynamic addresses bring to network traffic analysis is that while dynamic address "foo" might have been assigned to user Sam Smith at one point in time, that same dynamic address might subsequently have been reassigned to user Jane Jones shortly thereafter. Thus, if it becomes necessary to determine who is associated with a dynamic address "foo," as in the case of reports of network abuse, one needs to cross- reference additional records to resolve a particular dynamic address to a particular user at a particular point in time.

11. http://www.isc.org/products/DHCP/

Netblocks

Another concept you need to understand is that of netblocks. Network addresses are actually assigned, or "delegated" to organizations in "chunks" called network address blocks, not on a one-by-one basis.

For example, the University of Oregon has the block of addresses that range from 128.223.0.0 through 128.223.255.255, while Oregon State University has the block of network addresses that range from 128.193.0.0 through 128.193.255.255.

Each such network block has an associated netname, for example 128.223.0.0 through 128.223.255.255 is known as UONet, while 128.193.0.0-128.193.255.255 is known as ORST.

While both UONet and ORST are the same size netblocks (traditional class B netblocks), in general network address space is scarce, and new allocations are tightly rationed and require extensive justification. As a result, netblocks can (and do) vary dramatically in size, and may not be adjacent to other netblocks already assigned to an organization.

The basic size netblocks that you will commonly run into are as follows:

Netblock Sizes

Size	Class	# of Nets	# of Addrs	# of Class C's
/8	A	126	16,777,214	
		*		
		*		
		*		
/16	B	16,382	65,534	256
/17				128
/18				64
/19				32
/20				16
/21				8
/22				4
/23				2
/24	C	2,097,150	253 usable	1
/25			125 usable	1/2
/26			62 usable	1/4
/27			29 usable	1/8
/28				1/16
/29			6 usable	1/32
/30			1 usable	1/64

Address blocks used by the OWEN/NERO consortia are:

— University of Oregon

```
128.223.0.0/16
198.32.162.0/24
```

— Oregon State

```
128.193.0.0/16
199.201.139.0/24
```

— Portland State University

```
131.252.0.0/16
```

— OWEN/NERO (EOU, OCATE, OIT, SOU, WOU, and OUS administration)

```
140.211.0.0/16
207.98.64.0/18
```

— OPEN North

```
159.191.0.0/16          198.176.186.0/23
167.135.0.0/16          198.236.0.0/15
198.153.201.0/24        198.245.128.0/22
198.176.185.0/24        198.245.132.0/23
```

— OPEN South

```
157.246.0.0/16          198.140.208.0/24
163.41.0.0/16           198.237.0.0/19
167.128.0.0/16          199.79.32.0/20
192.220.64.0/18         204.214.97.0/24
198.68.17.0/24          204.214.98.0/24
198.74.32.0/21          204.214.99.0/24
198.74.40.0/23          206.99.0.0/19
198.98.8.0/22           207.98.0.0/18
```

— Eugene 4J School District

```
158.165.0.0/16
```

— State of Oregon DAS

```
159.121.0.0/16          199.2.160.0/19
167.131.0.0/16          199.48.32.0/20
170.104.0.0/16          199.195.16.0/20
192.133.23.0/24         204.27.190.0/24
192.152.7.0/24          204.89.128.0/24
198.68.186.0/24         204.94.0.0/19
198.176.0.0/22          205.143.224.0/21
198.176.0.0/21          205.167.4.0/23
198.176.4.0/23          205.167.156.0/23
198.176.229.0/24
```

With the exception of locally known address blocks, mapping individual addresses to netblocks typically requires either using the IPW[12] (IP whois) command, e.g.:

```
% ipw -a 128.223.32.18
128.223.0.0-128.223.255.255
```

or using the whois server running on the relevant regional registry — e.g., ARIN (covering the Americas), RIPE (Europe), APNIC (Asia), or NIPRNET (US military):

```
% whois -h whois.arin.net <address>
% whois -h whois.ripe.net <address>
% whois -h whois.apnic.net <address>
% whois -h whois.nic.mil <address>
```

Our experience has been that not all dotted quads will successfully map to assigned netblocks when using IPW, and that when IPW does successfully map a dotted quad to a netblock, the symbolic netblock name is often not particularly self-explanatory (nor even necessarily accurate, since a side effect of address scarcity and rationing is that many informal sub-delegations of network address space tend to occur).

There is also the issue that some organizations may use multiple non-adjoining netblocks. Use of multiple non-adjoining netblocks makes correct aggregation across those netblocks, but within a given organization, quite difficult to accomplish. (See, for example, the 19 distinct blocks comprising the DAS allocation, listed above). We should also note that netblocks are very poor (uneven) units of analysis for statistical purposes since some netblocks cover only a few dozen addresses, while others encompass millions. For all of these reasons, we have not done statistical analyses of network traffic by netblock, although in some cases we have used the netblock assignment information when nslookup failed to provide any guidance as to ownership of a given dotted quad of interest.

12. http://mjhb.marina-del-rey.ca.us/ipw/

ASNs

We do, however, provide summaries by ASN. ASNs are “autonomous system numbers” and represent “a connected group of IP networks that adhere to a single and clearly defined routing policy.”

Because each ASN represents an aggregation of multiple netblocks, and because you generally only can get an ASN if you are a major network that is multihomed (connected to multiple network service providers running BGP4), and because there are fewer ASNs than netblocks, ASNs a more appropriate level of network aggregation than network blocks for traffic reporting purposes. Examples of ASNs include:

```
3582        University of Oregon
4201        Oregon State University
6366        Portland State University
3701        NERO (includes EOU, OCATE, OIT, SOU, WOU, and OUS)
4222        OPEN South
6377        4J School District
7396        OPEN North
1798        State of Oregon DAS
```

A few examples of non-OWEN/NERO ASNs (out of the tens of thousands assigned) are:

```
3           MIT
68          Los Alamos
701         UUNet
1682        AOL
3356        Level3
13362       PC World Online
```

Any individual autonomous system number can be looked up by saying:

```
% whois -h whois.arin.net <ASN number>
% whois -h whois.ripe.net AS<ASN number>
% whois -h whois.apnic.net AS<ASN number>
% whois -h whois.nic.mil <ASN number>
```

In addition to ASNs being required[13] for sites that want to do BGP4 routing,[14] ASNs are also important because most peering-related decisions are based on traffic analyses performed at the ASN level of granularity. It is true that like netblocks, ASNs can vary widely in terms of their underlying size, but since they are such a common type of aggregator we felt we’d be remiss if we didn’t include an analysis of NERO/OWEN traffic by ASN.

13. http://www.arin.net/regserv/asnguide.htm
14. See http://www.ciscopress.com/catalog/titles/6522.html

Ports

Another important flow-related network concept is the concept of network “ports.” Each network service is offered to clients on a specific port. For example, most web servers listen on port 80 and most inbound mail is transfered to SMTP servers listening on port 25. If you think of a network address as telling you “where you’re going” you should think of a network port as specifying “what you’re going to do” once you get there.

More than one remote user may connect to the same incoming port number. For example, a given system may have a web server running on port 80, and it may handle tens or hundreds or thousands of users, all of whom are connecting to the “same” port number.

Port numbers can be divided into two ranges: privileged ports (numbered less than 1024), and general ports (running from 1024 up). On most systems, users are not permitted to create servers which listen/talk on privileged ports; only the system administrator (“root”) can install software to listen and respond to ports in that special range. In some cases, ports numbers are “well known”[15] and most (but not all) systems will use those port assignments. A brief summary of some of the more commonly used port assignments looks like:

```
20, 21          ftp (File Transfer Protocol)
22              ssh (secure shell)
23              telnet (remote login)
25              smtp (mail transfer between hosts)
53              dns (domain name service)
80              http (world wide web)
110             pop3 email
113             identd
119             nntp (Usenet News)
137             samba
139             netbios
143             imap email
161             snmp (simple network management protocol)
179             bgp
389             ldap/dropchute
443             https (secure world wide web)
554             qt4/rtsp/realaudio (real time streaming protocol)
563             secure nntp
1080            socks proxy
1723            pptp (point to point tunnelling protocol)
2049            nfs (network file system)
3128            squid (http proxy)
8080            http proxy
```

15. ftp://ftp.isi.edu/in-notes/rfc1700.txt

Additional ports are associated with particular interactive network games, and with various hacker/cracker programs which use the network while running *sub rosa* on compromised systems.

Some ports may be used by a variety of different programs in a dynamic fashion that makes it virtually impossible to even guess what they are being used for (for example, ports 1024 and immediately upward are commonly used on a dynamic basis by a wide variety of applications, and there is no effective way of identifying what particular application is using port 1024, 1025, 1026, etc. at any particular time).

It is also worthwhile noting that some applications may intentionally masquerade on a port that is normally used for a different purpose in an effort to avoid detection or simplify passage through firewall rulesets. For example, the file sharing program "dropchute" remaps the default telnet port, port 23, to exchange files. In other cases, applications may open multiple ports, or try sequential ports until they find a port that is available (remember that only a single server may run on any given port).

What does all this mean? Well, categorization of the more obscure types of network traffic based on the port number the traffic may be using needs to be done with the clear recognition that it is not a cut-and-dried, absolutely reliable process. Categorization of traffic by port is an imperfect approximation at best.[16]

Sequential Flow Analysis

In general, when classifying flows, we look at flows on a flow by flow basis. That is, we do not look at, nor do we need to look at, flows which may have preceded or which may follow any particular individual flow.

Categorizing flows associated with Napster (an MP3-format music sharing application which has recently received much press coverage) is an exception to that rule. Unlike most applications, the bulk of flows associated with Napster do not go over a well-defined port, nor does the bulk of Napster related traffic go to/from a single well-defined central server.[17]

The way Napster works is that a Napster user runs the Napster application on his or her local PC. When he or she wants to retrieve a particular song, they use the Napster application to ask the central Napster server for a "pointer" (or referral) to where that song might be found from among other users who are running Napster. (That pointer is trivially small in terms of the amount of network traffic associated with it, and if we only summed up the Napster referral traffic, we'd be missing the bulk of the real activity associated with Napster.)

16. A nice summary of common ports is at http://www.robertgraham.com/pubs/firewall-seen.html
17. For details about the Napster protocol, see http://david.weekly.org/code/napster.php3

Having received the pointer to a site or sites with the desired song, the user's Napster application then connects to one of the servers identified by the Napster server, and proceeds to download the desired song or songs. Note that the server which actually provides the song is virtually never a server owned by the Napster company, it is a server that may be running at a university or college, on a user's home machine that happens to be dialed in, on a company desktop somewhere, etc. (It is this download-related flow that is typically quite material in size in terms of network traffic.)

Dave Plonka of the University of Wisconsin has pioneered and advocated a new method[18] when it comes to categorizing traffic as being Napster-related, taking advantage of that characteristic sequential pattern of activity, e.g.,

— User interacts with a referral server at Napster.Com,

— User downloads MP3s from suggested server

The trick is associating the second step (the downloading of the files) with the first step (the characteristic "tagging activity" of interacting with the Napster.Com server). In our analysis we implemented Plonka's approach in a conservative way, tagging flows as being Napster-related if one of two conditions holds:

— We tag flows as being Napster-related with relatively high confidence if the flows have srcaddr or dstaddr typically associated with Napster traffic (e.g., ports 6699, 8875, 4444, 5555, 6666, 7777, or 8888), or the srcaddr or dstaddr is in a network address block known to have been assigned to Napster for its use. As you will see below, that resulted in allocation of approximately 2.0% of all inbound octets as being Napster traffic.

— after doing all other flow type categorizations we can, we then make a second pass back through what is left over as uncategorizable, and if the destination (internal to OWEN/NERO) address is known to have had one or more Napster-related flows, we then assign remaining otherwise uncategorizable flows associated with that address as being presumptively Napster-related. Doing this results in the identification of an additional 1.2% of all inbound octets as being presumptively Napster-related. We will discuss those results further below.

At this time, no type of traffic other than Napster is amenable to this particular type of sequential analysis strategy.

---

18. http://net.doit.wisc.edu/data/Napster/

TCP vs. UDP

To understand network flow reports, you also should understand that there are two common types of network traffic — TCP traffic and UDP traffic. TCP (Transmission Control Protocol) traffic is:

— “reliable,” meaning that packets are guaranteed to be delivered once and only once, uncorrupted and in the correct order

— “rate adaptive,” meaning that TCP-based applications will slow down in the face of network congestion or if the remote peer cannot keep up, and should theoretically also be able to speed up, if appropriate, and if the flow is of sufficiently long duration

— “connection-oriented/stateful,” meaning that the status of the other end of a TCP flow knows about the status of the other end of the connection — for example, a TCP session can detect if the remote server crashes or becomes unreachable, and react accordingly.

TCP is generally used for most comparatively low bandwidth local-area and wide-area Internet services, including such mainstays as telnet, ftp, smtp, http, and nntp.

UDP (User Datagram Protocol) traffic is pretty much the exact opposite of TCP. In the case of this protocol, it is characterized by...

— “unreliable/best effort delivery,” packets may be lost, duplicated, delayed or delivered out of order. (This sort of non-acknowledge-delivery scenario may sound unsettling/unacceptable, until you recognize that in many ways UDP traffic mimics what happens when you drop a postcard in the post office box — you do not get confirmation of US Mail’s successful delivery unless you pay extra to purchase return receipt service, and if you put two postcards in the same post office box, one in the morning and one in the afternoon, there is nopredicting which postcard will be delivered first (or if they’ll both be delivered at the same time, or if only one will be delivered, and of course, on rare occasions neither of them will make it!)

— “non-connection-oriented/stateless,” this means that the server and client are freed of the need to try to keep track of what’s happening on the other end of the flow

— “non-rate adaptive,” packets get launched at a rate determined by the programmer and the hardware the program is running on — if the application is to survive network congestion or overly-loaded hosts, it is incumbent upon the application to include an means of doing so.

It is worth noting that UDP is the ONLY option if an application requires doing broadcasts (transmissions to all hosts on a given subnet) or multicast (one-to-many transmissions). TCP or UDP can be used for unicast applications (conventional one-to-one transmissions). Most often, UDP is used for local-area network applications such as:

— NFS (network file system), a file sharing protocol popular under Unix

— multimedia applications such as unicast (streaming) audio and video, and for IP multicast audio and video

— selected other miscellaneous applications such as NTP (network time protocol), AOL Instant Messenger, among others.

We mention the two types of traffic here because a given port (listening via TCP) can be used for a completely different purpose than that same port listening for UDP connections, and because it is generally harder for a TCP-based application to “go fast” than it is for a UDP based application. That is, while a system may be connected to the network via a 100Mbps fast ethernet connection, it would be unusual for that system to even achieve anything approaching 100Mbps throughput when running a TCP/IP connection,[19] due to TCP/IP protocol dynamics.

19. This may be subject to change, however, as more attention is spent on optimizing network IP stacks. See, for example: http://www.scd.ucar.edu/nets/projects/web100/

Interfaces, Ports, VLANs and Router Configuration

A final network-related concept we need to introduce and explain is that of router interfaces and network switch ports (these hardware “ports” are not the same as the “ports” associated with network flows). You should think of a router interface or a network switch port as being a physical socket on that network device, to which a network cable can be connected. For example:

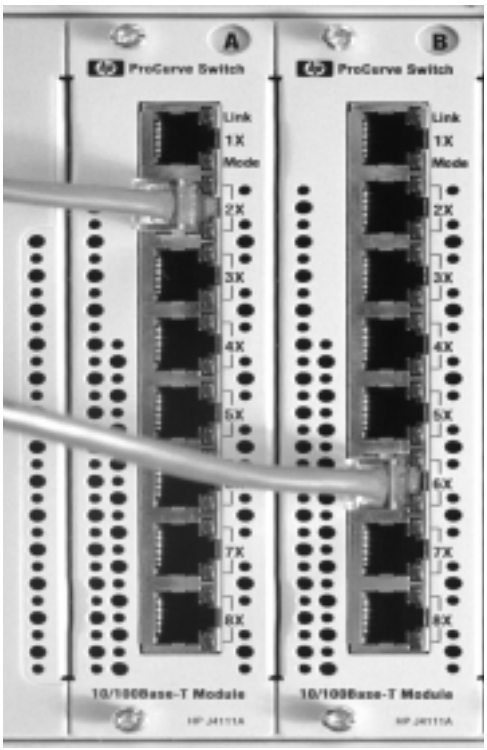

A typical network switch might have 8, 16, 40, or even 80 such ports (the above picture shows 16), while a router might have only three or four interfaces, or possibly far more.

Why do we bother to mention the concept of network interfaces and ports? Well, an important category of network monitoring, namely network monitoring via SNMP, which we’ll talk about later, is based on looking at counters associated with particular physical ports or interfaces.

Interpretation of data obtained via SNMP requires knowledge of how physical ports/interfaces are mapped to (or associated with) workstations, servers, or other network devices of interest. In a utopian world, network devices, like close friends, would be instantly and automatically recognized. In real life, however, the association of network devices to particular interfaces or ports is something that happens manually, as a bookkeeping matter, based on physical records created at the time the interconnections are made.

Pretty simple, right? Well, yes and no. The problem is that new stuff gets plugged in, old stuff gets unplugged, and existing stuff gets moved around, sometimes once or twice a year, sometimes daily. What is involved in deleting, adding or moving connections on a given network device?

On a network switch or hub, in the simplest of scenarios, all ports are fungible, and a given server or workstation can be unplugged from one and plugged back into another port having the same characteristics (speed, duplex setting, etc.) without problem... except that network management and monitoring software configurations often must now be manually updated to know that the system that used to be connected via port X is now on port Y.[20]

In a more complicated scenario, a single physical switch may use VLANs ("virtual lans") to break up a single physical switch into two or more virtual switches. That is, conceptually, ports 1-5 might be configured into one virtual network, ports 6-14 might be configured into a second virtual network, and ports 15-32 might be a third virtual network, just as if three physically separate switches were being used. Obviously, in that sort of a scenario, moving a cable from port 5 to port 20 would result in a non-trivial change, although a change from port 17 to port 24 (in our particular hypothetical scenario) would typically be transparent. (Also, VLAN configurations can be readily changed by a network administrator, which can further muddy the waters if record keeping isn't scrupulously maintained.)

In the case of a router, its interfaces are routinely configured on an interface-by-interface basis in a way that precludes just unplugging a connection from one interface and plugging that connection into another interface. However, a network engineer can change the configuration of the router at the same time the cables are getting rearranged, and then, just as in the other scenarios described above, a server can end up migrating from one interface to another. A common example of this arises when a network is moved from a regular 10Mbps ethernet interface to a fast 100Mbps ethernet interface located on a different physical router blade.[21]

---

20. Some network management software is "smart enough" to notice that the mac address, the unique hardware layer physical address assigned by the factory, has moved to a new port, and can either flag that change or take other action automatically.
21. Router "blades" are computer cards that plug into the router's chassis. Each blade might provide eight ethernet interfaces, or four fast ethernet interfaces, for example. By mixing and matching blades, a router can be configured to have whatever mix of ports a given network requires. The problem that large networks run into is that routers are typically quite expensive, and have chassis that can only hold a relatively modest number of blades (typically 5-7 blades). If you are operating large networks, there is thus great interest in routers that offer a "high port density" (e.g., having routers with chassis which can hold "lots" of blades, or blades that have a large number of interfaces (e.g., eight or more interfaces per blade rather than one or two interfaces) per blade. For example, Juniper Networks and Foundry Networks have been successfully eroding Cisco's market share for Internet core and campus routers, respectively, in large measure by offering port densities in excess of that routinely available in current Cisco products.

Section 3. Measuring Network Performance

Section 3 Keypoints

✔ Section 3 introduces key network measurement concepts.

✔ We can measure network traffic several different ways, including via SNMP, via flow analysis, via packet level passive monitoring, and via active measurement programs.

✔ SNMP (Simple Network Management Protocol) is a lightweight protocol that reports the value of simple counters associated with interfaces on network equipment. Interpretation of those counters requires knowledge of the underlying circuit's pupose. There can be an overwhelming number of SNMP variables; monitoring SNMP counters successfully is largely a measure of knowing which counters to pay attention to, and what constitutes "unusual" rates of change for those values.

✔ Flow analysis is another network measurement technique, and the one we rely on for the bulk of the data reported in this document. On a large production network the size of OWEN/NERO, there may be over three million flows in each direction per hour during peak usage times; hence, we only do flow analysis when circumstances require collection of flow data.

✔ Packet level passive monitoring can be used in some specialized circumstances when finer graularity is required than is available from flow based analyses, but can truely generate phenomenal levels of data and also poses privacy issues and security risks. Packet level passive monitoring wasn't done for this report.

✔ Active measurement programs, rather than watching traffic that happens to come by a particular sampling point, uses active probes to monitor network performance to remote sites of interest. OWEN/NERO partners participate in a variety of active measurement programs, however for the most part those measurements aren't applicable to the focus of this study (since they tend to involve non-commodity network connectivity).

✔ Most network measurement campaigns employ a combination of methods.

✔ Network measurement is still a very young discipline, with the first passive and active measurement workshop having been held just this spring at the University of Waikato in Hamilton, New Zealand.

So how do we measure network flows?

We can measure network performance several different ways.

SNMP (Simple Network Management Protocol)

SNMP is the Simple Network Management Protocol, and is defined in a number of RFCs.[22] It queries counters that are part of a MIB (Management Information Base).

SNMP is a lightweight standardized interface that allows most remotely manageable network devices to be queried on a manual or automated basis, reporting the state of the device as of that point in time. For example, connecting via SNMP to a router will allow us to identify the number of octets inbound and outbound per interface.

A MIB for a typical router or switch might have the following SNMP varaiables available:

```
mgmt/mib-2/interfaces/ifTable/ifEntry (OID: 1.3.6.1.2.1.2.2.1):

1   ifIndex.
2   ifDescr.
3   ifType.
4   ifMtu.
5   ifSpeed.
6   ifPhysAddress.
7   ifAdminStatus.
8   ifOperStatus.
9   ifLastChange.
10  ifInOctets.    [inbound octets, from the interface's point of view]
11  ifInUcastPkts.
12  ifInNUcastPkts.
13  ifInDiscards.
14  ifInErrors.
15  ifInUnknownProtos.
16  ifOutOctets.   [outbound octets, from the interface's point of view]
17  ifOutUcastPkts.
18  ifOutNUcastPkts.
19  ifOutDiscards.
20  ifOutErrors.
21  ifOutQLen.
22  ifSpecific.
```

22. RFCs are "Requests for Comments," somewhat misleadingly named documents promulgating Internet standards. A nice summary of relevant RFCs for SNMP can be found at http://www.hio.hen.nl/rfc/snmp/

Examination of SNMP ifInOctets and ifOutOctets counters can tell us, in aggregate, how much traffic is flowing over a given interface, and if we repeatedly poll the same counters, we can even get a picture of how that traffic varies over time.

Access to SNMP counters is typically limited to particular address ranges associated with network monitoring workstations, as well as password protected by a secret "community" string (which functions as a password, even though it isn't called one). In other cases, a default community string (often "public") may allow routine non-destructive (read only) access to SNMP data: For example, let's use SNMX[23] to access the SNMP values on a particular network device:

```
% snmx

SNMX> connect <address of device being monitored>

SNMX> cd /mgmt/mib-2/interfaces/ifTable/ifEntry/ifInOctets

SNMX> ls
Directory: /mgmt/mib-2/interfaces/ifTable/ifEntry/ifInOctets
OID: 1.3.6.1.2.1.2.2.1.10
-------------------------------------------------------------
Extension | -R-  Counter
-------------------------------------------------------------
.1        | 0
.2        | 0
.3        | 4195681094
.4        | 3706377452
[etc]

SNMX> ls
Directory: /mgmt/mib-2/interfaces/ifTable/ifEntry/ifInOctets
OID: 1.3.6.1.2.1.2.2.1.10
-------------------------------------------------------------
Extension | -R-  Counter
-------------------------------------------------------------
.1        | 0
.2        | 0
.3        | 175655722
.4        | 3726581991
[etc.]

SNMX> quit
```

If you compare the counter for interface 4 in the above example, you can see that it has changed by 3726581991-3706377452=20204539 octets between the first time that counter was queried and the second time that counter was queried.

23. http://www.ddri.com/Products/ace-snmx.html

Now look at interface 3. The alert observer will notice that the value for interface 3 at the time of the second polling actually is LOWER than the initial value. Is this because we've somehow magically "given some octets back?" No! The problem you are seeing is an example of SNMP "counter rollover" effects. That is, all SNMP counters have a maximum value[24] they can numerically represent, and when that value is exceeded the device routinely resets the counter to zero and resumes counting. Most SNMP management software explicitly handles that sort of rollover condition; we just mention it here as an illustration of the fact that SNMP is truly a "simple"/low level management protocol. To handle rollover and similar problems and to make SNMP-collected counter values "pretty" and more readily interpretable, most people graph those values using a product such as MRTG[25] or RRDtool.[26]

SNMP cannot, however, tell us very much about what traffic statistics mean.

SNMP lets us answer the "how much" question, but it doesn't — it can't — tell us very much about the "what" question or the "where from" and "where to" questions — at least not beyond the level of traffic statistics about what's going to physical ports on a given switch or interfaces on a router. (This is the port/ interface mapping problem we previously described in Section 2).

For those statistics to be changed from raw data to useful information, data about how networks or devices map to switch and router ports needs to be known, and you also need to understand the sort of load that is "normal" for a given port or interface.

Since a given switch or router may have eighty or more interfaces, obviously it isn't possible for a network technician to manually monitor all SNMP MIB (Management Information Base) variables for all interfaces at all times. The key to using SNMP successfully for network is to determine what you need to pay attention to, that is, what interfaces and variables need to be monitored more or less constantly, and what other interfaces and variables can be routinely (and safely) disregarded.

Sites that want an SNMP management station with a sophisticated graphical user interface and a full suite of features often purchase HP's commercial OpenView[27] product and run it on a dedicated network management workstation. Commercial network management software can be quite expensive.

---

24. In the case of counters represented as 32 bit values, that value is relatively small, only 2^{32} or only 4,294,967,296. Some network devices are now moving to using 64 bit values for SNMP counters, in which case their range is expanded to 2^{64} or 18,446,744,073,709,551,616.
25. http://ee-staff.ethz.ch/~oetiker/webtools/mrtg/mrtg.html
26. http://ee-staff.ethz.ch/~oetiker/webtools/rrdtool/
27. http://www.openview.hp.com/

Flow Analysis Using cflowd

For a finer level of granularity, network engineers can turn on flow accounting on selected router interfaces, and collect one record per active network flow using CAIDA's cflowd.[28] That is the method we rely on for the bulk of the statistical data described in this report. When we collect flow data, we generally get information about:

— flow starting time (in Unix "ticks," e.g., seconds since the start of the epoch)

— flow ending time

— source autonomous system number

— destination autonomous system number

— source network address (as a dotted quad)

— destination network address

— source port

— destination port

— network protocol (TCP, UDP, etc.)

— number of packets transferred, and

— number of octets transferred

Some may wonder why we don't collect this sort of information on an ongoing basis. The answer is simple: collecting flow data generates a phenomenal amount of data which must be stored and eventually processed, and collecting flow data also tends to delay packet forwarding on the router.[29] For example, a total of less than 30 minutes worth of flows from our commodity transit pipes translated to 1,761,170 flows and nearly 300MB worth of data to analyze!

28. http://www.caida.org/Tools/Cflowd/
29. [This is an interesting example of a Heisenbergian effect -- in measuring a phenomenon, by that sheer act of measuring, we change the phenomenon we're trying to assess.]

Note that collecting data at the flow level also raises privacy and security concerns since it may expose underlying network topologies and service information that may be valuable to cracker/hackers, and may reveal sensitive user level information (trivial examples include an employee visiting a web site for recovering alcoholics or suicide prevention, or an employee visiting a “head hunter” site to search for a new position, or a person looking for information about an incurable disease).

Packet Level Passive Monitoring

For still finer granularity, it is possible to attach a passive packet-level monitor to the network, after which one can then “sniff” (observe, see, eavesdrop upon) and record all traffic (or all traffic headers) flowing over that link. Obviously, since flows tend to consist of multiple packets/flow, doing this type of traffic analysis tends to yield far more data than monitoring on a per-flow basis, and can generate truely huge amounts of data when applied to busy connections.

Packet level monitoring also raises potentially very significant privacy and security concerns given that traffic being sniffed may include unencrypted usernames and passwords, the text of confidential email messages, credit card numbers being used for online commerce, etc. Traffic that is encrypted (for example, web transactions done with a secure server, or ssh remote login connections) obviously would not be vulnerable to eavesdropping in this way, but we believe that encryption of network transmissions is still the exception rather than the rule. The vast majority of connections (web sessions, email messages, telnet sessions, ftp sessions, etc.) all yield sensitive traffic whose contents are vulnerable to being sniffed.

No packet level passive monitoring has been relied on for the purpose of preparing this report.

Active One Way (and Round Trip) Ping Time/Packet Loss Measurements

A final way of measuring the sufficiency of network bandwidth consists of doing active one way or round trip ping time/packet loss measurement studies.

In this type of study, measurement sites are deployed around the network at a variety of sites of interest, and ping packets are then periodically sent from each of the measurement sites to each of the other measurement sites. The time it takes for the ping packets to get to each of the remote sites (in the case of one way measurements), or the time it takes for the ping packet to make a round trip (to the remote site and back) are then measured, recorded and summarized.

The “manual version” of this process looks something like:

```
% ping -s www.altavista.com
PING altavista.com: 56 data bytes
64 bytes from www.altavista.com (204.152.190.16): icmp_seq=0. time=33. ms
64 bytes from www.altavista.com (204.152.190.16): icmp_seq=1. time=30. ms
64 bytes from www.altavista.com (204.152.190.16): icmp_seq=2. time=28. ms
64 bytes from www.altavista.com (204.152.190.16): icmp_seq=3. time=29. ms
64 bytes from www.altavista.com (204.152.190.16): icmp_seq=4. time=30. ms
64 bytes from www.altavista.com (204.152.190.16): icmp_seq=5. time=29. ms
64 bytes from www.altavista.com (204.152.190.16): icmp_seq=6. time=29. ms
64 bytes from www.altavista.com (204.152.190.16): icmp_seq=7. time=31. ms
64 bytes from www.altavista.com (204.152.190.16): icmp_seq=8. time=33. ms
64 bytes from www.altavista.com (204.152.190.16): icmp_seq=9. time=30. ms
64 bytes from www.altavista.com (204.152.190.16): icmp_seq=10. time=35. ms
64 bytes from www.altavista.com (204.152.190.16): icmp_seq=11. time=27. ms
64 bytes from www.altavista.com (204.152.190.16): icmp_seq=12. time=29. ms
64 bytes from www.altavista.com (204.152.190.16): icmp_seq=13. time=29. ms
64 bytes from www.altavista.com (204.152.190.16): icmp_seq=14. time=32. ms
64 bytes from www.altavista.com (204.152.190.16): icmp_seq=15. time=27. ms
64 bytes from www.altavista.com (204.152.190.16): icmp_seq=16. time=31. ms
64 bytes from www.altavista.com (204.152.190.16): icmp_seq=17. time=31. ms
^C
----altavista.com PING Statistics----
18 packets transmitted, 18 packets received, 0% packet loss
round-trip (ms)  min/avg/max = 27/30/35
```

Examples of this type of active monitoring program done on a systematic basis include:

— The NLANR AMP[30] project (round-trip measurement oriented)

— Advanced.Org’s Surveyor[31] project (one-way measurement oriented)

— Lawrence Berkeley Labs’ NIMI[32] (IP multicast) loss measurement project

— CAIDA’s Skitter[33] project (focusing on the path from a given site to many destinations spread across the Internet)

— IWR (Internet Weather Report) sites[34]

— Keynote Systems’ active applications-oriented monitoring program[35]

30. http://amp.nlanr.net/
31. http://www.advanced.org/csg-ippm/
32. http://www.ncne.nlanr.net/nimi/
33. http://www.caida.org/Tools/Skitter/
34. http://www.ad1440.net/~devnull/work/iwr/

The University of Oregon currently hosts an AMP box, a Surveyor box, and a NIMI box, and is due to receive a Skitter box shortly. UO also runs an Internet Weather Report-like monitoring box that watches ping times and packet loss to selected network peers. Among other OWEN/NERO partners, we know that Oregon State also participates in the NLANR AMP program.

Sample output from some of these measurement activities is included later in this report, however note that these active measurement activities have a number of limitations with respect to being applied to this study, including:

— many of the active measurement programs are limited to high performance "Internet2" partner sites, while our focus in this study is on commodity Internet transit bandwidth

— our interest is on inbound performance, which means that measurements need to be made from elsewhere to us, and public facilities for doing those sort of measurements are limited/non-existent as a general rule.[36]

What determines a choice from among those basic approaches?

In most cases, when it comes to categorizing network performance and understanding network bandwidth, a combination of network measurement approaches works best. For example, we monitor OWEN/NERO wide area bandwidth utilization on a macroscopic level using SNMP, we do flow studies when required for special reports such as this one, and we use active network monitoring tools to confirm that desired end-to-end performance is being obtained. In many cases, those approaches may need to be augmented with logs from individual systems (e.g., logs from web servers may be required to analyze what's being served from that system), and in some cases interviews with individual users will be the only way to ultimately determine what's actually going on. No single technical network monitoring approach can, nor should be expected to, tell the whole story.

The other point worth noting is that network measurement is still a very young discipline; for example, the first Passive and Active Measurement Workshop was only held in April 2000![37] Many techniques are still under active development and research, and many questions remain open at this time.

---

35. http://www.keynote.com/
36. The only common tool for doing this sort of things are traceroute gateways; see, for example: http://www.tracert.com/cgi-bin/trace.pl or http://www.boardwatch.com/traceroute.html
37. The First Passive and Active Measurement Workshop (PAM2000), April 3-4, 2000, Department of Computer Science, University of Waikato, Hamilton, NZ, ISBN 0-909007-20-9

Section 4. The OWEN/NERO Network

Section 4 Keypoints

- OWEN/NERO is Oregon's "network of networks."

- OWEN/NERO provides Internet connectivity for all public universities in the state, virtually all public elementary and secondary schools, and all state agencies, connecting in all over 620,000 Oregonians.

- OWEN/NERO has three hub sites, one in Portland, one in Eugene, and one in Corvallis.

- The three hub sites are connected via intrastate DS3 and OC3 point-to-point circuits, and that design protects against loss of connectivity due to failure of any single intrastate circuit or loss of any single hub site.

- OWEN/NERO's Internet transit is purchased from UUNet and Cable and Wireless, two major international network service providers; traffic to or from any of the OWEN/NERO partners may end up flowing to or from OWEN/NERO via either of those commodity Internet transit provider.

- OWEN/NERO also provides intra-consortia connectivity, which is particularly important for interconnecting the state's North and South LATAs, insuring non-Eugene OWEN/NERO partners access to the OIX, and protecting OWEN/NERO against catastrophic loss of transit connectivity.

What is OWEN/NERO? Who Are OWEN/NERO's Customers?

OWEN/NERO is Oregon's "network of networks," providing intrastate and Internet connectivity for all public (Oregon University System) 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 about 31,000-34,000 additional state agency employees — plus all members of the public who interact with those customers.

OWEN/NERO's Commodity Internet Transit Connectivity

OWEN/NERO's commodity Internet transit connectivity is purchased from UUNet[38] and from Cable and Wireless.[39]

The UUNet connectivity comes in via OWEN/NERO's Portland hub while the Cable and Wireless connectivity comes in via OWEN/NERO's Eugene hub.

It is important to note that traffic from anywhere in OWEN/NERO may enter or exit via EITHER transit connectivity provider — it would be incorrect to assume that North LATA/Portland area traffic exits or enters only via UUNet, or that South LATA/Eugene/Corvallis traffic exits or enters only via Cable and Wireless. Traffic to or from any of the OWEN/NERO partners may end up flowing to or from OWEN/NERO via either transit connectivity provider.

Multihoming OWEN/NERO this way protects the network from catastrophic loss of commodity Internet connectivity, as might occur if either of these Network Service Providers had a major failure.

OWEN/NERO's Intrastate Topology

OWEN/NERO has three hub sites, one in Portland, one in Eugene, and one in Corvallis. The Portland hub site is colocated at the ELI[40] colocate facility, the Eugene hub site is colocated at the University of Oregon, and the Corvallis hub site is colocated at Oregon State University. The Portland hub is in the state's North LATA ("local access and transport area," as defined by the Telecommunications Act of 1996[41]); the Eugene and Corvallis hubs are in the South ("Eugene") LATA.

38. http://www.uu.net/
39. http://www.cwix.net/
40. http://www.eli.net/
41. http://www.fcc.gov/telecom.html

The three hub sites are connected via three intrastate DS3 (45Mbps) leased line circuits (with the Portland-Eugene circuit moving to OC3 (155Mbps) soon).

Redundancy and load sharing is inherent in the way these circuits have been deployed (each hub site is connected to both of the other two hubs via separate leased circuits, thereby protecting OWEN/NERO against loss of connectivity associated with any single intrastate circuit failure).

Understanding the Value of OWEN/NERO's Intrastate Connectivity

In general, without OWEN/NERO's intrastate connectivity, inter-partner traffic, particularly traffic between sites in Portland and sites in Eugene and Corvallis, would all have to go via commercial Internet connections. Thus, instead of that traffic flowing over the OWEN/NERO backbone at no charge, an OWEN/NERO site originating a particular flow would have to buy sufficient additional Internet transit capacity to allow them to send that traffic over the public Internet, and the OWEN/NERO site receiving the flow would have to buy sufficient additional Internet transit capacity to allow them to receive that traffic over the public Internet.

In that type of scenario (if OWEN/NERO did not exist), if two current OWEN/NERO sites happened to buy their Internet connectivity from two different service providers, traffic between those two sites (which might be only miles apart in Oregon) might travel all the way to Washington State or all the way to California before being switched between the two providers, thereby adding unnecessary traffic delays and providing more opportunities for service interruptions.

Without OWEN/NERO's intrastate connectivity, at least some consortia-common bandwidth savings would also be lost. Why? Because currently some content is received over expensive commodity Internet transit links only once, and once it has been received it is shared internally with all interested members of the consortia. Examples include OWEN/NERO's Usenet News feeds and the consortia's web caching efforts. Without OWEN/NERO's intrastate connectivity, these services would be deployed on a redundant/ duplicative basis at each individual OWEN/NERO site or constituent network.

Another way of thinking about what OWEN/NERO provides via its intrastate links is to think about who would be impacted, or what resources would become less valuable, if we were to hypothetically delete OWEN/NERO's intrastate circuits:

— First, we should note that hypothetically deleting OWEN/NERO's intrastate circuits would differentially and dramatically affect some current OWEN/NERO partners (particularly OSU

and OPEN), far more than it would others. For example, both the University of Oregon and the Portland area universities will soon be able to fall back on Internet2 as a replacement inter-campus backbone if OWEN/NERO's intrastate connectivity were to go away, however Oregon State relies on OWEN/NERO's intrastate DS3's to tunnel its I2 connectivity to and from the Oregon Gigapop. Similarly, OPEN relies on OWEN/NERO's intrastate DS3's to carry cross-LATA traffic between OPEN North and OPEN South.

— Also, if OWEN/NERO's intrastate circuits were hypothetically to be deleted, Portland area and Corvallis OWEN/NERO partners (including DAS) would also lose access to peerage at the Oregon IX (UO and OPEN South would continue to have direct access). Portland may soon have peerage opportunities at the Pittock Block in Portland, but that peering point is still in the build-out phase, and there is no guarantee that any provider who may locate there will peer with any particular potential partner.

— Thirdly, we note that if OWEN/NERO's backbone circuits were torn down, all OWEN/NERO participants would also lose protection against commodity transit connectivity loss. Without OWEN/NERO's intrastate connectivity, traffic from the North LATA would only be able to exit from UUNet; traffic from the South (Eugene) LATA would only be able to exit via CWIX.

OWEN/NERO's intrastate circuits are one of those hidden resources, a real asset to the consortia that many OWEN/NERO participants don't understand and appreciate until you begin to think about what changes would take place if those links were to disappear.

Section 5. The Role of the Oregon IX

Section 5 Keypoints

- This section explains the role that the Oregon Internet Exchange plays for OWEN/NERO.

- In addition to purchasing commodity transit connectivity from UUNet and CWIX, OWEN/NERO also “peers” at the Oregon Internet Exchange (OIX) located at the U of O.

- Networks which “peer” agree to exchange customer traffic (and ONLY customer traffic) without paying each other any financial settlements.

- Peering can happen anywhere two networks agree to meet, but peering tends to occur at neutral exchange points where multiple networks are present.

- Peering serve to keep local traffic local, and also reduces the amount of expensive commodity transit bandwidth which must be purchased.

- Each DS3 class peering circuit brought into the OIX by a network represents a potential avoided cost of roughly $500,000/year.

- Major peers currently at the OIX include Verio (the world’s leading web hosting company, now part of NTT Japan) and a number of others, including Globix (with customers including Real Networks, Microsoft, the National Hockey League, Standard and Poors, and many others) and Akamai (a major new distributed web content delivery company) soon also to be live at the OIX.

Other OWEN/NERO Connectivity: Peerage at the Oregon Internet Exchange

In addition to commodity Internet transit connectivity, OWEN/NERO also peers at the Oregon Internet Exchange ("Oregon IX," "OIX"),[42] located at the University of Oregon in Eugene. The Oregon IX is one of a number of exchange points[43] in the country where network providers meet to exchange customer traffic (and ONLY customer traffic) without settlements, thereby keeping local traffic local, and reducing the amount of commodity transit connectivity each provider needs to provision.

In order to establish an exchange point, four things must exist:

1) An exchange point requires a suitable physical location, that is, a site with:

— Secure 24x7 access (typically via a cardkey system) with onsite security

— Incumbent and competitive local exchange carrier availability for provisioning local loops, plus fiber facilities for higher bandwidth needs

— Industry standard rackage for equipment, plus suitable cable raceways; many sites also offer locking cages for equipment deployment

— Conditioned power (e.g., large uninterruptable power supplies and/or backup generation capacity)

— Copious cooling capacity to prevent equipment from overheating

— A fire suppression system

2) An exchange point also requires networks who are interested in peering at that location, which is largely a function of questions such as:

— Where is the exchange point located? Are there any competing exchange points nearby which I might prefer?

— Will I be able to expand if I want/need to do so in the future?

---

42. http://www.antc.uoregon.edu/OREGON-EXCHANGE/
43. http://www.ep.net/

— Who else is at the exchange point? (e.g., what ASNs are represented?) How much of my traffic originates with or is destined for those ASNs?

— Do I have faith that those ASNs are competently engineered?

— Will the other ASNs who are there be willing to peer with me? (This is a function of the type of equipment and circuits that the other providers may have at the exchange point, their assessment of whether peering might foreclose an opportunity to sell commodity transit to that same party, relevant company policies, and a reciprocal assessment of engineering competence).

— What will it cost me in terms of circuits, fiber, equipment, engineering effort, travel, etc., to bring up a connection at this exchange point? Are there ongoing monthly fees? Is the business case for appearing at this exchange point sound?

3) An exchange point (unless it is to be based strictly on a mesh of peer-to-peer directly arranged private circuits) requires a central ethernet switch, router, or ATM switch into which peers can connect. In OIX's case, this is a negligible cost item, but at other sites it may represent an investment of hundreds of thousands of dollars.

4) Finally, an exchange point requires administration, including strategic planning, policy determination, network monitoring, marketing, coordination, etc.

For OWEN/NERO, one easily articulated advantage to peering at the Oregon IX is that OWEN/NERO has the ability to exchange customer traffic with other OIX peer network customers at no cost. If you assume that commodity transit costs (just for estimating purposes) $1000/Mbps/month, a peer who comes in with a DS3 (45Mbps circuit) represents an avoided cost (a value to OWEN/NERO) of over half a million dollars a year, assuming OWEN/NERO could fully utilize peerage circuits of that capacity. Connections between OWEN/NERO and peers at the Oregon IX are also free of local loop charges (because the Oregon IX is colocated with the OWEN/NERO Eugene hub site). Local loop charges can cost thousands of dollars per month or more, if local loop needs to be purchased from an ILEC or CLEC.

Finally, note that as providers come into the Oregon IX, a "critical mass" forms and the OIX becomes more attractive to additional providers. Peers currently at the Oregon IX currently include Verio[44] and other network service providers, and continues to increase. For example, Globix[45] is now coming to the OIX, and Akamai[46] is also in the process of colocating an Akamai distributed content server at the Oregon IX.

44. Verio is the world's #1 web hosting solutions provider, hosting more than 305,000 web sites for customers in 127 countries. Settlement free peerage with Verio means, for example, that all OWEN/NERO traffic to the Altavista search engine site (a Verio customer) and to the Excite search engine (another Verio customer) go via OIX Verio peerage at no charge.
45. Globix customers include Microsoft, Real Networks, Dow Jones, Standard and Poors, the NHL, and many others. See: http://www.globix.net/about_customers.html
46. http://www.akamai.com/
For a nice overview of Akamai, see: http://www.wired.com/wired/archive/7.08/akamai.html

Section 6. The Role of Internet2

Section 6 Keypoints

- This section discusses the role of Internet2 with respect to OWEN/NERO connectivity.

- Internet2 connectivity is a third type of connectivity available to eligible OWEN/NERO partners (in addition to commodity Internet transit and peerage at the Oregon Internet Exchange).

- Internet2 connectivity supplements (but does not eliminate) the need for commodity Internet transit connectivity since Internet2 connectivity can only be used to carry traffic between I2 sites, or between an I2 site and an approved I2 peer network.

- There are over 170 American Carnegie Research I and Research II universities which are I2 members at this time.

- I2 members can physically connect to Internet2 via either Abilene or the vBNS.

- Abilene is a high speed research and education network running on top of Qwest facilities; the vBNS is a high speed research and education network running on top of MCI Worldcom facilities.

- An Abilene OC3 connection cost $110,000/year (plus local loop charges). For comparison, a UUNet OC3 commodity transit connection costs $2,148,000/year. Thus I2 connectivity, while still expensive in absolute terms, is really quite a bargain compared to the cost of commodity Internet transit.

- Fortunately, most Internet2 sites have received federal, state, corporate and institutional financial support to help underwrite their Internet2 connectivity. In the case of Oregon, the NSF provided $436,320 to OSU, $350,000 to UO, and $542,979 to PSU/OHSU/OGI to help support establishing Internet2 connectivity.

- I2 member sites may elect to connect directly to Abilene or the vBNS, or multiple sites may connect to Abilene via a single shared connection managed by an entity called a “Gigapop.”

- At this time Oregon has one operational Gigapop, the Oregon Gigapop at the University of Oregon in Eugene. The Oregon Gigapop connects to I2 via two OC3 circuits. Those OC3 circuits are backhauled at no charge from the Oregon Gigapop to the Abilene Denver core node and to the Abilene Sacramento core node.

- PREN is a new Portland-area Internet2 Gigapop and metropolitan area network, and will have connectivity to Internet2 via the University of Washington’s Gigapop in Seattle. PREN traffic will get backhauled from Portland to Seattle at no charge via connectivity provided by WCI.

- OWEN/NERO customers who are Internet2 primary members at this time are UO, OSU, PSU, and OHSU; OGI is also a primary member of Internet2 (although OGI is not an OWEN/NERO customer).

- UO, OSU and PSU currently connect to Abilene via the Oregon Gigapop in Eugene.

- OHSU and OGI Internet2 connectivity is currently awaiting completion of PREN.

- OWEN/NERO has four Internet2 secondary participants, the first colleges granted this status in the country. Those secondary participants are EOU, OIT, SOU, WOU.

- EOU, OIT, SOU and WOU currently connect to Internet2 via the Oregon Gigapop in Eugene.

- In addition to interconnecting the 170 or so I2 member institutions, I2 also peers with a variety of federal mission networks, including DOE's ESNet, DOD's DREN, NASA's NREN/NISN, and over a dozen high speed foreign research networks.

Other OWEN/NERO Connectivity: I2 Connectivity

A third type of connectivity, in addition to commodity Internet transit and peerage at the Oregon IX, is Internet2[47] connectivity.

Internet2 is a sort of specialized high speed connectivity which supplements (but does not eliminate) the need for commodity Internet transit connectivity. Since Internet2 connectivity can only be used to carry traffic between Internet2 sites, or between an Internet2 site and an approved I2 peer network,[48] all Internet2 connected sites MUST maintain commodity Internet transit in addition to their Internet2 connectivity.

There are over 170[49] American Carnegie Research I[50] and Research II[51] universities which are members of Internet2 at this time. Is members can physically connect to either Abilene[52] or the vBNS.[53]

Abilene

Abilene is a high speed research and education network running on top of Qwest[54] facilities, with its network operations center at Indiana University.

Abilene OC3 connections cost $110,000/year, OC12 connections cost $320,000/year, and OC48 connections cost $495,000/year (in each case plus local loop charges). For comparison, note that a Qwest OC3 commodity transit connection costs $140,000/month,[55] or $1,680,000/year and a UUNet OC3 commodity transit connection costs $179,000/month,[56] or $2,148,000/year. Thus I2 connectivity, while still expensive in absolute terms, is really quite a bargain compared to the cost of commodity Internet transit.

In addition to connection and loop costs, other applicable fees include a $20,000/year Abilene participation fee, an annual UCAID membership fee of $25,000/year, and a Qwest Access Interconnect fee of $1,000 plus $1,000/month (OC3) or $2,000 plus $2,000/month (OC12).

---

47. http://www.internet2.edu/
48. http://www.ucaid.edu/abilene/html/cou.html
    http://www.vbns.net/vBNS+/vbns+faq.html
49. http://www.internet2.edu/html/universities.html
50. http://www.carnegiefoundation.org/OurWork/Classification/CIHE94/PartIfiles/ResearchI.htm
51. http://www.carnegiefoundation.org/OurWork/Classification/CIHE94/PartIfiles/ResearchII.htm
52. http://www.internet2.edu/abilene/
53. http://www.vbns.net/
54. http://www.qwest.net/
55. http://www.boardwatch.com/isp/summer99/bb/qwestpg5.html
56. http://www.boardwatch.com/isp/summer99/bb/uunetpg5.html

vBNS

The vBNS is a somewhat older high speed research and education network running on top of MCI Worldcom facilities.

vBNS (actually vBNS+, now) DS3's lists for $86,400/year, OC3's list for $259,200/year, and OC12's list for $1,036,800/year,[57] all plus local loop and other applicable fees. Because Abilene is significantly less expensive than the vBNS/vBNS+, we believe that most sites have (or will eventually) move their connections to Abilene unless they can negotiate a price for service that is significantly less than quoted list prices.

NSF Support

Fortunately given the magnitude of the costs mentioned, in addition to State, institutional, and corporate support for these connections, the National Science Foundation has also provided generous support for Internet2, including providing grants amounting to $436,320 to Oregon State University,[58] $350,000 to the University of Oregon,[59] and $542,979 to a Portland consortia comprised of Portland State University, Oregon Health Sciences University and Oregon Graduate Institute.[60]

Sites which accept NSF funding customarily must match that funding with funding from other sources, and must also commit to continuing connectivity at (or above) initial levels after the conclusion of the NSF support.

Gigapops

I2 sites can connect either directly to Abilene or the vBNS, or multiple sites may connect to Abilene via a single shared connection managed by a Gigapop.[61]

At this time, Oregon has one operational Gigapop, the Oregon Gigapop[62] at the University of Oregon in Eugene. The Oregon Gigapop connects to I2 via two OC3 circuits. Those OC3 circuits are backhauled at no charge from the Oregon Gigapop to the Abilene Denver core node and to the Abilene Sacramento core node.

---

57. http://www.vbns.net/main.html?q=5&t=69&i=170
58. NSF Award #9617043; see http://www.nsf.gov/cgi-bin/showaward?award=9617043
59. NSF Award #9729628; see http://www.nsf.gov/cgi-bin/showaward?award=9729628
60. NSF Award #9975992; see http://www.nsf.gov/cgi-bin/showaward?award=9975992
61. http://www.internet2.edu/html/gigapop_list.html
62. http://www.ogig.net/

PREN[63] is a new Portland-area Internet2 Gigapop and metropolitan area network, and will have connectivity to Internet2 via the University of Washington's Gigapop[64] in Seattle. PREN traffic will get backhauled from Portland to Seattle at no charge via connectivity provided by WCI Cable of Hillsboro, Oregon.[65] Connection speeds mentioned for that link to the University of Washington have ranged from fast ethernet (100 Mbps) all the way through OC12 (622Mbps).

OWEN/NERO Internet2 Memebers

OWEN/NERO customers who are Internet2 primary members at this time are UO, OSU, PSU, and OHSU; OGI is also a primary member of Internet2 (although OGI is not an OWEN/NERO customer). UO, OSU and PSU currently connect to Abilene via the Oregon Gigapop in Eugene. OHSU and OGI Internet2 connectivity is currently awaiting completion of PREN.

OWEN/NERO also has four Internet2 secondary participants,[66] the first colleges granted this status in I2. Those secondary participants are EOU, OIT, SOU, WOU, and they connect to Internet2 via the Oregon Gigapop in Eugene.

A map showing I2 site connectivity is available at http://www.abilene.iu.edu/images/logical.pdf

Internet2 Peer Networks: Federal Mission Networks and International Research and Education Networks

The value of Internet2 extends beyond just providing connectivity among its member institutions. Internet2 also has peering relationships with a number of government high performance networks, including:

— the Department of Energy's ESNet[67]

— the Department of Defense's DREN[68]

— NASA's NREN/NISN[69]

---

63. http://www.pren.net/
64. http://www.pnwgp.net/
65. http://www.wcicable.com/
66. http://www.ucaid.edu/abilene/html/secondary-application.html
67. http://www.es.net/
68. http://www.hpcmo.hpc.mil/Htdocs/DREN/
69. http://www.nren.nasa.gov/ and http://www.nisn.nasa.gov/

Foreign partner high performance research and education networks peering with Abilene include:[70]

— APAN (Japan)

— Canarie (Canada)

— CERN (Switzerland)

— Dante (Europe)

— DFN (Germany)

— IUCC (Israel)

— JANET (United Kingdom)

— Nordunet (servicing Norway, Sweden, Finland and Denmark)

— RedIRIS (Spain)

— Renater (France)

— Singaren (Singapore)

— Surfnet (the Netherlands)

— SWITCH (Switzerland)

— TAnet (Taiwan)

Traffic statistics for those peers are at: http://monon.uits.iupui.edu/abilene/peers.html

To see a map showing current general I2 traffic levels, see http://hydra.uits.iu.edu/~abilene/traffic/ (recall that the Oregon Gigapop connects directly to Denver and to Sacramento).

70. http://www.internet2.edu/international/

Section 7. Narrowing Our Focus: Understanding Why Inbound Commodity Transit Is Key

Section 7 Keypoints

- This section explains why this report is focussed on inbound commodity transit.

- Although all of OWEN/NERO's connectivity is valuable and important, in this report we've focused on OWEN/NERO's commodity transit connectivity, the largest cost associated with OWEN/NERO.

- When it comes to determining how much transit connectivity OWEN/NERO needs, you can't simply sum up the total capacity of all circuits that connect to OWEN/NERO -- that would result in purchase of too much capacity.

- If you just sum up the traffic transferred per day and try to just average that traffic over 24 hours, that would result in too little capacity (the load varies greatly during the course of the day, with significant peaks and deep troughs).

- Usage is greatest during the early to mid afternoon, and lightest in the early morning.

- Inbound usage dominates outbound usage, and inbound usage determines the amount of transit bandwidth required (inbound and outbound capacity cannot be independently provisioned at different levels).

- Usage isn't "flat topping" (e.g., demand isn't greatly exceeding the level of capacity which has been purchased), nor, conversely, does provisioned capacity greatly exceed peaking demand.

- By inference, outbound load isn't material at this time to bandwidth capacity requirement planning.

- Similarly, off peak load ins't material to bandwidth capacity requirement planning.

- **Sizing transit circuits to meet inbound peaking traffic loads is the key**.

- If you buy more capacity than you need, that's wasted capacity which you must pay for but can't use.

- If you buy too little capacity, the network will perform poorly and customers will be disatisfied.

- If OWEN/NERO customers are disatisfied with available bandwidth, they can, will, and previously have left the consortia to buy network access from another provider.

- OWEN/NERO currently protects itself against unbounded consumption by individual partners by capping inbound transit traffic at customer stipulated levels; major traffic in excess of that stipulated rate is dropped.

- OWEN/NERO partners pay $1000/megabit per second per month for their stipulated inbound commodity transit traffic level.

- Selected traffic which benefits all partners (e.g., a common news feed, web caching, IP multicast traffic) is excluded from that bandwidth cost structure.

- Partner traffic profiles are available at http://www.nero.net/cgi-bin/rrdcust.cgi/pritar=Traffic_Profile

So did you study all these different kinds of OWEN/NERO connectivity?

No. Although all the types of connectivity that OWEN/NERO provide are valuable and important, our focus for this bandwidth audit is on OWEN/NERO's commodity transit connectivity. It is the largest cost associated with OWEN/NERO, it is a directly variable cost, and we know that growth in commodity transit connectivity costs is the Legislature's greatest oversight concern.

Hence, we narrowed our focus solely to OWEN/NERO's commodity transit connectivity, the Internet connectivity OWEN/NERO buys from UUNet and Cable and Wireless.

So how did you gather data on the transit connectivity? Did you just collect an average week's worth of usage data?

No. To understand our bandwidth audit sampling plan, you need to understand OWEN/NERO's typical bandwidth profile. A typical MRTG graph for our UUNet transit connectivity for a month or so from earlier this year is shown below:

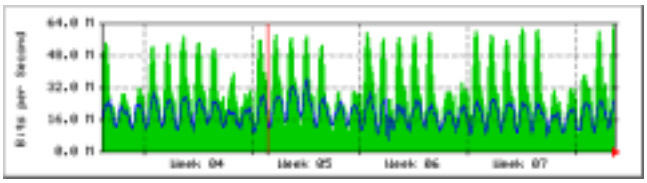

The tall peaks represent inbound (Internet to OWEN/NERO) circuit utilization, and the lower line represents outbound circuit utilization (OWEN/NERO to the Internet). The horizontal axis is time, with each green spike being roughly a day apart. The vertical axis is traffic measured in Mbps (megabits per second). The graph shown is a very typical graph for OWEN/NERO transit connectivity.

There are some important things we can glean from that graph:

— Observed transit bandwidth load was far below potential aggregate load — that is, even if there were tens of thousands of hosts connected at 10Mbps, hundreds connected at 100Mbps and scores connected at gigabit ethernet speeds, the aggregate inbound load over that circuit was only some 60Mbps at its peak. Obviously, OWEN/NERO couldn't size their Internet transit capacity by simply tallying up the capacity of all the hosts or all the circuits that they connect — they'd buy vastly too much capacity.

— If they were to have just taken the total traffic transferred per day and divide it by the time period involved, they'd have obtained an average amount of traffic transfered per hour — but that average would have been far below the capacity they would have required to meet observed peaking loads.

— Usage tended to peak dramatically during the day, with the greatest peaks occurring during early to mid afternoon; we know from other graphs that usage during the evenings for DAS and OPEN is nil (due to the fact that both of those OWEN/NERO partners have limited evening access via public labs or library computer pods, no residential networking or remote access (modem) capacity, limited international traffic and characteristic 8AM-5PM work patterns). Usage also exhibited periodicity associated with the day of the week, with weekend usage lower than weekday usage.

— For all intents and purposes, inbound usage dominated outbound usage, and inbound usage determined the amount of transit bandwidth required (inbound and outbound capacity cannot be independently provisioned at different levels)

— Usage wasn't "flattopping" — if they'd been grossly underprovisioned, usage would've exhibited a characteristic "flat top" or "plateau" area leading up to, through, and immediately after the period of peak demand, but with inbound usage peaking at roughly 60Mbps, OWEN/NERO should feel confident that they were actually obtaining the capacity they'd purchased from UUNet, while also being comfortable that they hadn't "overpurchased" capacity not required to meet peak demand.

Related to that, we can draw some further inferences:

— Outbound load (unless it flip-flops at some time, is uncontrollable, and begins to dominate inbound load), is NOT material

— Off-peak load in either direction isn't a material issue at this time (since OWEN/NERO needs to size to meet peaking loads, and off-peak loads just "rattle along" with plenty of headroom)

What does this all mean? Bottom line:

<u>Sizing transit circuits to meet inbound peaking traffic loads is clearly key</u>.

If OWEN/NERO buys more capacity than they need, they will spend more than they should, and for no benefit (clearly OWEN/NERO cannot “warehouse” excess network capacity or somehow carry that network capacity forward; if they cannot use provisioned capacity it is simply lost, a wasted expenditure).

If they buy less capacity than they need, particularly if they buy less capacity than they need to meet peaking loads, the network will perform poorly and their customers will become dissatisfied.

Why do you need to meet peaking load? Why not simply let things slow down during peak times of the day?

The answer is that OWEN/NERO is a consortia, operates like a utility, and needs to be responsive to its customer requirements. In this we are like any utility. Just as a utility can’t allow brownouts on hot summer days when they experience peak electrical loads, so, too, OWEN/NERO needs to have capacity available when its customers want to access the Internet.

OWEN/NERO customers DO differ from typical “captive” utility customers in one very important way. Unlike a typical utility’s customers, who are captive and only really have one potential provider for water or power, our consortia partners have the ability to choose who they rely on as an Internet service provider. If OWEN/NERO’s leadership and network engineers fail to operate OWEN/NERO in a professionally responsive way, OWEN/NERO’s customers have a clear option: they can leave the consortia for an alternative (more professionally responsive) network service provider.

In particular, OWEN/NERO’s OPEN’s K12 constituency is very adamant that network capacity must be increased to meet observed demand; if that isn’t done, ESDs (Education Service Districts) will “mutiny” and seek an alternative network service provider which can and will meet their capacity requirements.[71] OWEN/NERO really cannot compel any OWEN/NERO partner to accept slow service, nor would they want to.

At the same time, obviously they cannot allow unbounded consumption, as would be true if OWEN/NERO commodity transit capacity had no incremental cost. They are currently handling that problem by allowing individual participants to stipulate a level of inbound transit bandwidth they’d like to receive or which they believe they need to have, and charging them $1000/Mbps/month for their stipulated/desired level of capacity. In order to do that, it is necessary to differentially “color” or “tag” incoming commodity transit traffic (since it is directly related to incremental transit bandwidth costs) unlike peerage traffic and I2 traffic and intra-consortia traffic (which does not have a direct incremental cost). Traffic that facilitates consortia-wide services (e.g., centralized web caching support and inbound newsfeeds to the consortia’s news servers) are also excluded.

Currently traffic in excess of stipulated rates is dropped, subject only to instantaneous bursts briefly exceeding those levels.

You can see the traffic profile report for OWEN/NERO partners at:

`http://www.nero.net/cgi-bin/rrdcust.cgi/pritar=Traffic_Profile`

In looking at that traffic profile, compare the purple “transit usage” squiggly line against the flat red transit limit line.

---

71. For example, Northwest Regional Educational Service District decided that it wanted to independently procure its network services, and purchased service from Qwest rather than OPEN. This has a number of unfortunate consequences, including rather poor traffic routing between NWRESD and the rest of the OPEN community. For example, note how traffic to the NWRESD web site from Eugene currently goes to San Francisco, then back to Seattle before returning to northwestern Oregon!

    traceroute to www.nwresd.k12.or.us (198.236.4.100), 30 hops max, 40 byte packets
    1 cisco3-gw.uoregon.edu (128.223.142.1) 0.610 ms 0.572 ms 0.723 ms
    2 cisco7-gw.uoregon.edu (128.223.2.7) 0.465 ms 0.439 ms 0.464 ms
    3 eugene-hub.nero.net (207.98.66.11) 1.648 ms 1.643 ms 1.425 ms
    4 eugene-isp.nero.net (207.98.64.41) 71.921 ms 2.416 ms 227.423 ms
    5 xcore2-serial0-1-0.SanFrancisco.cw.net (204.70.32.5) 12.057 ms 11.477 ms 11.857 ms
    6 corerouter2.SanFrancisco.cw.net (204.70.9.132) 13.375 ms 17.075 ms 11.896 ms
    7 ngcore2.Seattle.cw.net (204.70.9.130) 27.446 ms 27.586 ms 28.998 ms
    8 core9.Seattle.cw.net (204.70.9.77) 28.612 ms 28.183 ms 27.709 ms
    9 sea-brdr-01.inet.qwest.net (205.171.4.77) 31.698 ms 30.843 ms 30.047 ms
    10 sea-core-01.inet.qwest.net (205.171.26.5) 31.113 ms 31.928 ms 30.630 ms
    11 sea-edge-03.inet.qwest.net (205.171.26.38) 32.395 ms 31.488 ms 30.115 ms
    12 205.171.45.166 (205.171.45.166) 34.554 ms 35.118 ms 36.787 ms
    13 205.171.45.166 (205.171.45.166) 35.631 ms 42.472 ms 37.367 ms
    14 ultra.nwresd.k12.or.us (198.236.4.100) 37.097 ms * 34.976 ms

    Contrast that with a traceroute to www.pps.k12.or.us (Portland Public Schools, an OPEN member), which never leaves the state, and is three times as fast...

    traceroute to qei.pps.k12.or.us (159.191.7.45), 30 hops max, 40 byte packets
    1 cisco3-gw.uoregon.edu (128.223.142.1) 0.560 ms 0.876 ms 0.437 ms
    2 cisco7-gw.uoregon.edu (128.223.2.7) 0.653 ms 5.786 ms 1.189 ms
    3 eugene-hub.nero.net (207.98.66.11) 1.751 ms 1.632 ms 1.242 ms
    4 eugene-isp.nero.net (207.98.64.41) 1.680 ms 1.583 ms 2.221 ms
    5 ptld-isp.nero.net (207.98.64.2) 5.817 ms 5.390 ms 8.894 ms
    6 ptld-hub.nero.net (207.98.64.177) 7.026 ms 5.410 ms 5.667 ms
    7 open-eli-gw.nero.net (207.98.68.6) 7.888 ms 8.893 ms 7.057 ms
    8 open-7507.k12.or.us (198.236.254.9) 6.747 ms 5.873 ms 7.005 ms
    9 pps-gw.k12.or.us (198.236.254.6) 9.081 ms 11.114 ms 10.175 ms

When traffic inbound to a OWEN/NERO partner exceeds the transit limit line, the partner then has three options:

— they can do nothing, in which case the excess traffic will automatically be dropped,

— they can adjust their transit limit upward by paying more (so that OWEN/NERO can in turn buy more transit bandwidth), or

— they may be able to internally manage their bandwidth so as to reduce their demand.

Currently, most partners are working below their contractual transit limits, with the possible exception of brief peak-usage periods during the middle of the day

Section 8. Bandwidth Standards

Section 8 Keypoints

✔ This section considers what might be done to establish a per user bandwidth standard, or a per user cost standard for bandwidth support.

✔ **With respect to establishing a bandwidth standard per serviced user, the current OWEN/NERO transit connectivity 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**.

✔ If we assumed that we wanted to guarantee at least 56,000 bits per second to all our users at the same time, we'd need to to buy 34,720 Mbps of transit, or 286 times OWEN/NERO's current bandwidth.

✔ Looking at it from a different perspective, if each OWEN/NERO customer paid a dollar a month for network bandwidth, OWEN/NERO would receive $14.8million/biennium (vs. its current budget of less than $2 million/biennium).

✔ Extreme usage by individuals at each partner site is theoretically limitable via purchase and installation of special hardware, however there are many practical reasons why that isn't recommended at this time.

What if you were to establish a bandwidth standard per user, and then provision bandwidth according to that standard?

This suggestion was heard from a committee member during the January JLCIMT hearing. Basically, the member wanted to know if we had established a standard for usage on a per-user basis, so that we could then provision our bandwidth accordingly, with the belief being that such a standard would tend to cap (lower) OWEN/NERO's bandwidth requirements.

We do not have such a standard right now, and after "running the numbers" it becomes clear why.

Let's do that now. Recall that OWEN/NERO services roughly 621,000-624,000 users (let's say 620,000 users just to use round numbers). OWEN/NERO has 121 Mbps worth of commodity transit bandwidth, 76 Mbps from UUNet and 45 Mbps from Cable and Wireless (CWIX). If you divide the 121,000,000 bits per second of bandwidth by our estimate of OWEN/NERO's 620,000 users, that works out to:

$$\frac{\text{121,000,000 bits per second}}{\text{620,000 users}} = \text{roughly 195 bits per second per user}$$

For context, note that dialup modem connects users at 56,000 bits per second ("56 Kbps"). Clearly 195 bits per second is a very modest level of commodity bandwidth by any standard, far below what would potentially be required if we were to guarantee even modem level throughput to all OWEN/NERO customers:

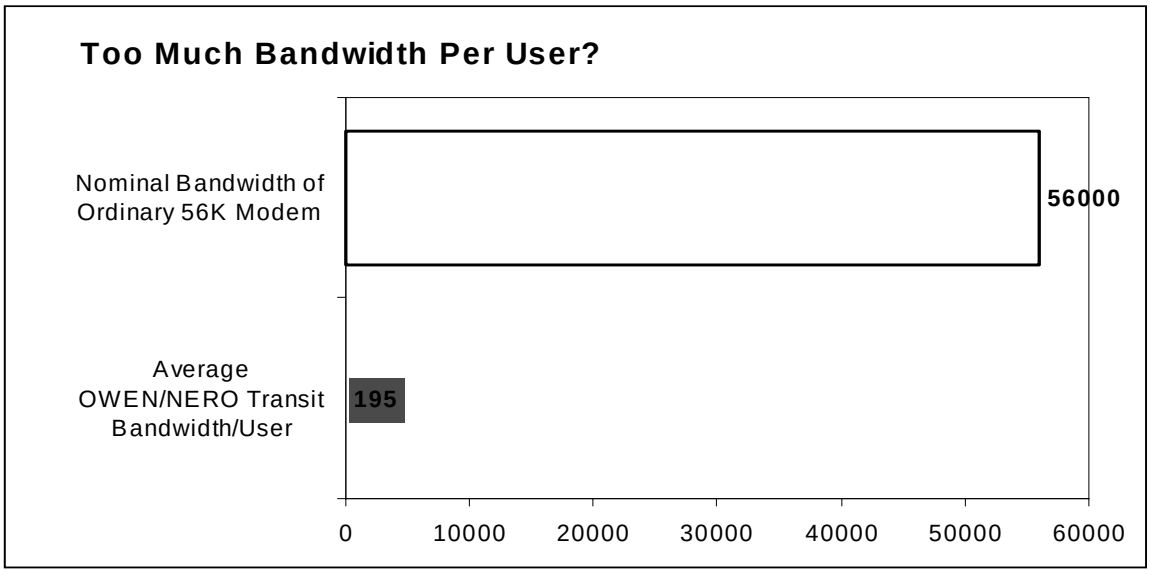

But what if we were to approach this problem in the reverse direction, computing the committed bandwidth we'd need to provision if we were to guarantee every user minimum simultaneous access at modem-like speeds of 56,000 bits per second?

Doing the math for that, we obtain:

(620,000 users) * (56,000 bps) = 34,720,000,000 bits per second or 34,720 Mbps

That's 286 times the current OWEN/NERO bandwidth. Even assuming that only 1% of our users are active at any time, that's still be nearly three times the currently deployed bandwidth.... Clearly OWEN/NERO is extremely closely provisioned and does unusually well when it comes to aggregating statewide demand in a cost-effective fashion.

Yet another way of thinking about this is to ask, "What would OWEN/NERO's budget be if every person serviced by OWEN/NERO paid a dollar a month for their Internet connectivity?"[72] The number is a rather staggering $14,880,000/biennium, nearly seven and a half times OWEN/NERO's total budget (including costs for personnel, equipment, intrastate circuits, transit bandwidth and other miscellaneous costs) of $1,955,000/biennium.

Similarly, since Internet service has become a "utility" like service, it may also make sense to compare it to the cost of providing electricity, natural gas, water, sewer, and related utilities. For the 1997-1998, we believe University of Oregon utility costs (NOT including any network costs) were $4,043,000,[73] and head count enrollment at UO in fall 1977 was 17,207.[74] Dividing that through, we see that utilities (NOT including networking) cost the University an average of $19.58 per student per month (coincidentally, $19.58 per student per month is very close to the per person per month price point of most network service providers). UO's cost per person per month for utilities is not going to hold for all OWEN/NERO partner institutions, obviously, not even for UO some three years later.... What that $19.58/student/month value does provide, however, is a benchmark order-of-magnitude check on whether or not positing a dollar/user/month for Internet bandwidth is sane. We believe that our $19.58/student value for other (non-network) utility costs shows that it is.

---

72. We would note that AOL and most other Internet service providers offer service at $15-$20/user/month, however they offer a variety of value-added services such as dialin access, host-based services (such as email and web page hosting), and end user support, over and above the network access that OWEN/NERO provides. We chose the dollar/user/month figure as a simple value representing only the network connectivity part of a typical user's monthly Internet costs, although obviously that number could be higher, or lower for a given user or a given provider.
73. Private communication, Director, UO Campus Operations, July 13th, 1998.
74. http://www-vms.uoregon.edu/~reoweb/facts/facts_f97.html

But is there really no way to cap extreme usage by individual users?

No, it is possible to cap extreme usage by individual users, and there are even some times when we'd definitely see some value in being able to automatically do that. For example, if a system were to be compromised by hacker/crackers, it would be useful to be able to automatically cap usage by that particular system.

The most common hardware product for implementing that sort of per user rate cap is Allot Communication's NetEnforcer boxes.[75] They have two models available which are relevant, the AC201 (optimized for network speeds up to 10Mbps, and with a list price of $7,500) and the more powerful AC301 (optimized for network speeds up to 100Mbps and with a list price of $13,000). Bandwidth management boxes covering speeds up to 45Mbps are also available from Packeteer[76] and others.

More than one bandwidth management box would likely be required per site. To understand why more than one bandwidth management box would be required/site, note:

— total drainage at some sites is in excess of 100Mbps, and the biggest bandwidth management box currently available is 100Mbps; thus, implicitly, handling more than 100Mbps worth of traffic will require more than one box, each such box installed at the edge of the network and covering only a single subnet running at 10Mbps or 100Mbps. At UO alone there are scores of subnets, although we suppose you could decide to traffic shape, but not others (athough obviously that sort of differential treatment would raise its own set of policy issues).

— since there are multiple exits at some sites, there is no single exit point at which policy enforcement could occur; if you were somehow going to multiple exits must be policed

— given that the bandwidth management box would be a mission-critical piece of hardware, it would need to be deployed in a redundant configuration to provide survivability in the event one NetEnforcer box were to fail.

— at least some boxes have maximum concurrent host/maximum concurrent flow limits

75. http://www.allot.com/products/ACfamily_DS.htm
76. See, for example://www.packeteer.com/products/packetshaper/index.cfm

In addition to the cost of the hardware, use of the NetEnforcer or PacketShaper technology would also probably require deployment of an LDAP directory at each institution, with a record for each user authorized to access the network, assuming the goal is to track usage per user (rather than per port), a daunting project given the size of the user base we're talking about.

The NetEnforcer or PacketShaper's activity would also materially increase the support load for user support staff as users with fast connections (10Mbps, 100Mbps, etc.) received only some fraction of that potential speed. It is very difficult to define and cap "bad" traffic (such as an attack launched from a compromised box) while not accidentally and undesirably curtailing "good" traffic from that system.

There's also the problem of coordinating and consolidating data received from multiple bandwidth management boxes if the goal is management of particular classes of traffic (e.g., game traffic) wherever it may happen to be originating to some specified maximum bandwidth limit.

For all these reasons, and because we believe that bandwidth management boxes of this sort actually mask problems (such as compromised systems) that we'd really rather be able to identify based on their traffic patterns. We do not anticipate recommending deployment of bandwidth management hardware for the forseeable future.

Section 9. OWEN/NERO Bandwidth Usage Compared to Other Consortia

Section 9 Keypoints

- It is uncommon for network consortia to make their bandwidth provisioning and utilization information public, however we were able to get data for nine consortia: CALREN-2 (California), the Great Plains Network (covering Arkansas, Kansas, Nebraska, North Dakota, Oklahoma and South Dakota), MichNet (Michigan), More.Net (Missouri), NCNE (serving five schools in Pennsylvania and West Virginia), NCREN (North Carolina), Net.Work.Virginia (Virginia), the Washington State K-20 Network, and WiscNet (Wisconsin).

- **Comparing OWEN/NERO bandwidth to bandwidth deployed by similar (or smaller) network consortia, OWEN/NERO's bandwidth is well under what comparable networks have installed**.

Is the consortia's transit bandwidth comparable to what other higher education network consortia have?

In our case, we purchase 76 Mbps worth of UUNet transit bandwidth, and 45Mbps worth of CWIX transit bandwidth. The CWIX bandwidth we obtained under a special promotional offer that allowed us to buy a full DS3 for the amount we would otherwise have had to pay for merely the next 3Mbps increment. Consequently, the CWIX circuit actually has more inbound capacity than we technically require right now, excess capacity which was literally free (obviously a great deal if we end up eventually needing it).

Thus, our total Internet transit bandwidth is 121 Mbps, although a more realistic value for comparison is 96 Mbps (76 Mbps inbound from UUNet plus the 20 Mbps inbound from CWIX that we actually use out of that 45 Mbps total MCI has provisioned under their special promotion).

Not all institutions or network consortia make commodity bandwidth utilization information publicly available, and obviously this sort of information is constantly subject to change. Some consortia that do have information available are:

— CALREN-2 and 4CNet:[77] These California networks don't have clean links to their transit provider bandwidth. However, you can see reported usage on a per-campus basis. Email from CALREN-2 in response to our query resulted in the information that CALREN-2 has two OC3s [e.g., 310 Mbps] worth of commodity transit capacity of its own, but that "this does not reflect all the transit bandwidth for all the CalREN-2 members. Many of the CalREN-2 members have their own transit providers and therefore don't use CalREN-2's. At present only the University of California campuses use this service, and those campuses also share a couple of DS-3's to Exodus."

— Great Plains Network:[78] GPN has 91 mbps worth of Internet connectivity[79] plus 1244mbps worth of Internet2 connectivity (2xOC12). GPN is the only consortia identified that runs with less commodity transit than OWEN/NERO does; interestingly, it has four times OGIG's I2 connectivity. It may also be worth noting that GPN hosts an Akamai box (as the OIX will be shortly), which has been delivering over 19Mbps worth of additional "capacity" for Great Plains at some peak times.[80]

77. http://www.calren2.net/router-stats/
78. http://nic-ks.greatplains.net/mrtg/index.html
79. http://nic-ks.greatplains.net/mrtg/index.html
80. http://nic-mn.greatplains.net/mrtg/MN/mn-1.r.greatplains.net.2.html

— MichNet:[81] Michnet has two CWIX DS3's in southeast Michigan and an additional 4.5Mbps to CWIX from the Upper Peninsula (at Houghton). They also have a Qwest DS3 and a WinStar Broadband DS3. Total commodity Internet drainage is thus 195Mbps. Michnet also has peerage at the Chicago NAP with Abovenet, AIS, Argonne, CAIS, Cetlink, Concentric, CRL, Digex, DRA, Exodus, FNSI, Globalcenter, IBM Global, IDT, Onvoy, and OARNet. These peering relationships materially reduce the amount of commodity transit Michnet needs to purchase. Michnet's 777 Mbps worth of Abilene connectivity (an OC12 plus an OC3) should also be factored in. Michnet also has T1 class peering arrangements at other locations around the state.

— MORENet:[82] Missouri's research and education network. Runs an OC3 statewide backbone with 225Mbps worth of Internet transit in February 2000, and "270Mbps this spring."[83] In 1999, had 96 FTE and a budget of $26,500,000.[84]

— NCNE[85] (servicing Penn State, CMU, Pitt and PSC and WVU): 45 Mbps to ATT Worldnet, 21 Mbps to ATT CERFnet, 22 Mbps to Sprintlink and 45Mbps to UUNet or 133 Mbps in total. While this is roughly equivalent to the current OWEN/NERO bandwidth, please notice that it services far fewer users.

— NCREN (North Carolina Research and Education Network[86]): full UUNet OC3 (155Mbps), plus a full Qwest OC3 (155Mbps), plus a full Sprint OC3 (another 155Mbps) — that's 465Mbps vs. OWEN/NERO's 121 Mbps

— Net.Work.Virginia: According to email from Net.Work.Virginia,[87] Net.Work.Virginia has an OC3 plus two DS3's from Sprint, for a total of 245Mbps worth of commodity internet transit.

— Washington K-20 Network: According to email from a Washington K20 network engineer, they currently have 4 DS3's, and will soon be adding a new OC3 while also upgrading one of their four existing DS3's to an OC3. (e.g., 3x45+2x155=445Mbps) Washington also has the first (and only, to date) Abilene OC48 (2.5Gbps).

---

81. http://www.merit.edu/michnet/maps/backbone.gif
82. http://www.more.net/
83. http://www.more.net/m3/
84. http://www.more.net/infoserv/tour_morenet/hisory.html
85. http://www.ncne.net/arch/genarch.html
86. http://mercury.ncren.net/
87. Private communication from 25 February 2000.

— WiscNet: "WiscNet is connected to Genuity by high speed (two OC-3) lines from both UW-Madison and UW-Milwaukee.[88]

We show this data graphically in the following horizontal bar chart (the vertical reference line is at 121Mbps, the current OWEN/NERO transit bandwidth):

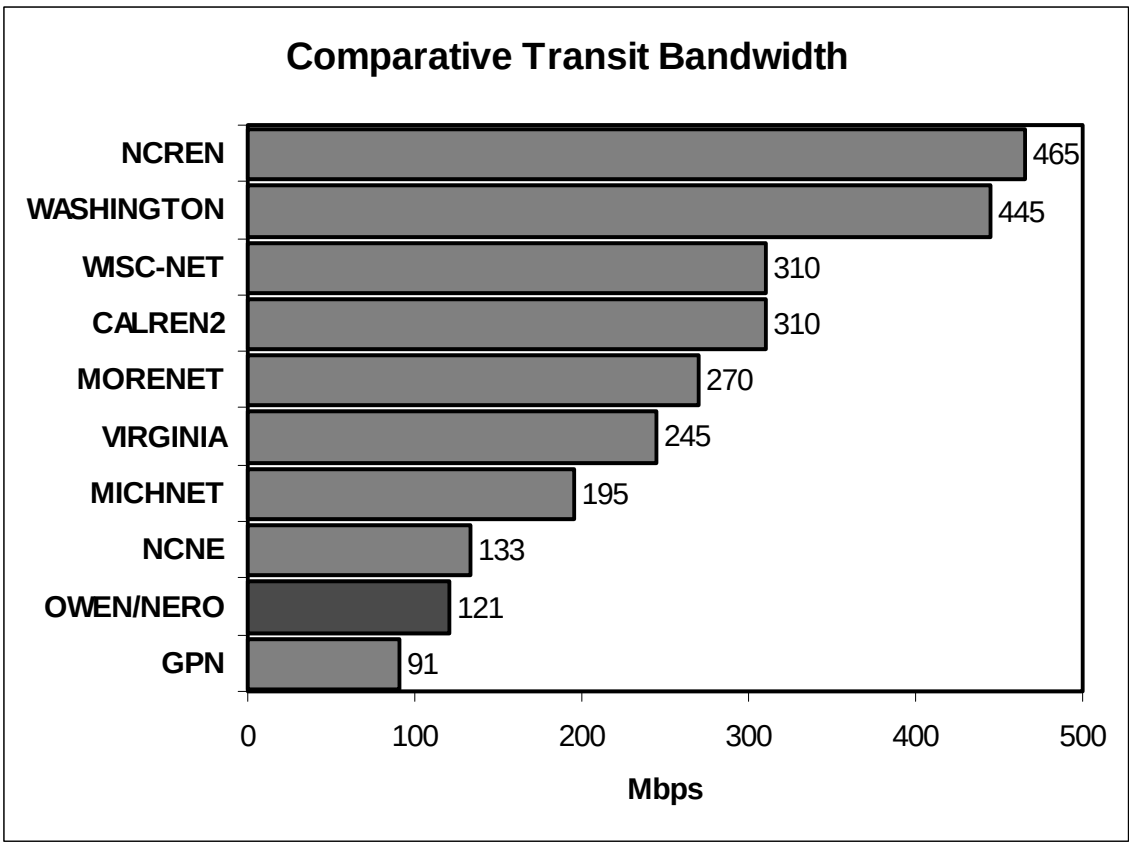

In summary: OWEN/NERO's commodity transit is generally MUCH less than what comparable networks typically have installed.

88. http://www.wiscnet.net/q&a.html

Section 10. OWEN/NERO Flow Data Sampling Plan and Basic Flow Descriptive Statistics

Section 10 Keypoints

- Given that inbound peaking loads drive the amount of Internet transit bandwidth required, we sampled inbound network traffic at peak demand time (2:00-2:15PM) on two successive days (April 6th and 7th).

- 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 being available for analysis**.

- **Descriptive statistics associated with the flow sample were consistent with those reported by Kevin Thompson for the MCI backbone as of 1997, and McCreary and Claffy's data for the Ames Internet Exchange based on a ten month period ending in March 2000.**

The Sampling Plan

So, given that we need to size our commodity connectivity to meet inbound peaking loads, and given that most OWEN/NERO members will be sizing their rate limits to accommodate observed inbound peaking flows, we decided to focus our sampling efforts on the time of day when those peaking flows occur.

We requested that samples be taken on two successive days during the period beginning at 2:00 PM and ending at 2:15 PM.

The requested samples were collected on April 6th and April 7th by a NERO/OWEN network engineer, drawn simultaneously from both the Eugene (CWIX) commodity transit pipe and from the Portland (UUNet) commodity transit pipe (see the cover of this report to see where the samples were drawn).

The April 6th and 7th dates represented dates from a normal school week, when both K12 and higher education were in session (obviously load patterns would be different on weekends, during inter-term breaks, or over the summer).

Samples were drawn without notice to reduce the possibility that users might intentionally or subconsciously alter their behaviors during the sampling period.

Details for the resulting data files look like:

```
Date, Location &
Data file size    Start Time  End Time    Elapsed     Flows    Flows/sec
-------------------------------------------------------------------------
04/06 EUG (CWIX)  955054800   955055597   797 ticks   237,659 ~292.2/sec
39,689,053 bytes  2:00:00 PM  2:13:17 PM

04/06 PDX (UUNet) 955054800   955055700   900 ticks   687,698 ~764.1/sec
114,845,566 bytes 2:00:00 PM  2:15:00 PM

04/07 EUG (CWIX)  955141200   955142100   900 ticks   246,705 ~274.1/sec
41,199,735 bytes  2:00:00 PM  2:15:00 PM

04/07 PDX (UUNet) 955141200   955142100   900 ticks   589,108 ~654.6/sec
98,381,036 bytes  2:00:00 PM  2:15:00 PM
                                                    ---------
                                                    1,761,170 flows
```

We noticed the 103 tick (103 second) discrepancy in the length of the 04/06 Eugene data file shown in the above summary table, and a query was made to the NERO engineer who collected the data. The engineer indicated that he believed that file was truncated because the buffers used to collect the final block of flow data had not completely filled in memory and had not been flushed to disk at the time the data collection was terminated, while it had done so in the other cases. We do not believe that this truncation will have a material impact on any aspect of our analysis.

What did the flow distribution look like?

The mean flow length in octets was 9,476.56 with a variance of 3.809×10^{10}. This average flow length is consistent with previously reported average flow lengths reported by Thompson.[89]

The median flow (50th percentile flow) was 554 bytes, and the modal (most common individual value) was 40 bytes. The smallest flow was only 20 bytes, and the largest flow was 62,920,569 bytes. Ninety-nine percent of all flows were 76,317 or less bytes in length.

Looking at flow durations in seconds, the mean flow time was just under 4 seconds (3.966794), with a variance of 422.343, a maximum duration of 306 and a minimum duration of less than a second (time values are recorded with a granularity of one second). Ninety-five percent of flows were 16 seconds long or less, with 99% of flows lasting for no more than 47 seconds.

Considering the number of packets per flow, the mean was 15.71456 (variance of 4,4062.3), with a minimum of 1 packet per flow and a maximum of 45,188 packets. (Thompson eported an average number of packets per flow of 16-20 for their TCP traffic.)

89. “Wide Area Traffic Patterns and Characteristics (Extended Version),” Kevin Thompson, http://www.vbns.net/presentations/papers/MCItraffic.pdf section 5.2 quoting 5-8KB and 5-9KB as typical average TCP flow sizes for various links as of 1997.

What protocols were seen?

In chapter two, we discussed the concept of TCP and UDP protocols and described some of the differences between the two. For reference, the distribution of protocols seen on OWEN/NERO during the sampling period (arranged in descending order by flow count) were as follows:

```
Protocols by Flows

PROTO   Frequency   Percent
-------------------------
TCP      1448700      82.3
UDP       275474      15.6
ICMP       36822       2.1
IPv6         137       0.0
GRE           28       0.0
IPIP           9       0.0
```

In addition to TCP and UDP, which we've already discussed in section 2 of this report, we saw a number of additional protocols show up in our samples. Those other uncommon protocols are:

— ICMP: Internet Control Message Protocol is defined in RFC 792.[90] ICMP is used by a variety of different applications/commands including the ping command and the traceroute command.

— IPv6: Internet Protocol Version Six is the next generation Internet Protocol with a vastly expanded address range, and is designed to help eliminate the IP address shortage that the Internet currently faces in IP version 4.[91]

— GRE: Generic Routing Encapsulation is defined in RFC 1701.[92] GRE permits any arbitrary protocol to be tunnelled (or "encapsulated") on top of IP.

— IPIP: IP in IP tunnelling is defined in RFC 1853.[93] It permits tunnelling with IP Security and other protocols.

Thompson's data had TCP at 75%-85% of the total number of flows (with TCP being a "higher percentage of overall traffic during business hours than in the evening or overnight") UDP at 20% of flows, and ICMP at 1.5% of flows. Our traffic is thus very consistent with Thompson's data.

---

90. ftp://ftp.isi.edu/in-notes/rfc792.txt
91. http://playground.sun.com/pub/ipng/html/ipng-main.html
92. ftp://ftp.isi.edu/in-notes/rfc1701.txt
93. ftp://ftp.isi.edu/in-notes/rfc1853.txt

Note that the above breakdown is by number of flows, not by number of octets. For information on the breakdown of total octets seen per protocol, please see the next table.

Protocols by Octets

PROTO	Frequency	Percent
TCP	1.556E10	93.2
UDP	1.1129E9	6.7
ICMP	18397369	0.1
IPv6	2595249	0.0
GRE	46891	0.0
IPIP	41340	0.0

[One note on the above table: “E notation” is used to represent very large numbers — interpret a number such as 1.556E10 as being 15,560,000,000; similarly, 1.1129E9 is the same as 1,112,900,000]

Again, our traffic pattern is very similar to that described in the literature; Thompson’s traffic is 95% TCP octets, roughly 5% UDP octets, and a half a percent ICMP octets.

A new paper[94] by McCreary and Claffy of CAIDA describing traffic seen over a period of ten months at the Ames Internet Exchange in Mountain View (based on a circuit with a median utilization of 85Mbps) provides an additional and more contemporaneous comparator (data collection for that project terminated in March of 2000). In their study, 91% of their traffic was TCP, 5.1% UDP, 2.7% GRE, and 0.7% ICMP (plus an assortment of other minor protocols, some of which we saw, and some of which we didn’t). Again, this data is reassuringly consistent. [Because of how McCreary and Claffy obtained their data sets, they do not have per flow statistics available]

94. “Trends in Wide Area IP Traffic Patterns: A View from Ames Internet Exchange,” Sean McCreary and kc claffy, http://www.caida.org/outreach/papers/AIX0005/

Section 11. Flow Application Breakdown

Section 11 Keypoints

- **On a per flow basis, nearly three quarters of all flows sampled were http (e.g., world wide web), which is very consistent with Thompson's reported value of 75%**.

- Roughly nine percent of all OWEN/NERO flows sampled were domain name system related; no other single application accounted for a significant number of flows (e.g., five percent or more 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)**.

- Free web email (Hotmail, etc.) continues to be quite popular for a variety of reasons, even when local web email is available as an option. Mail is an important application because virtually everyone uses email, email messages these days often include large attachments, and "free" web based email products are typically advertising driven. Addition of Akamai and Globix to the OIX will eliminate much of the bandwidth usage associated with most popular of these web based email products.

- A significant amount of inbound commodity transit may also be associated with hacker/cracker attacks; c.f., the EUNet Bulgaria DNS denial of service attack captured during the sampling interval.

What applications were seen?

Looking at a breakdown of application traffic by flows (and noting that categorizing traffic into application based on port/protocol data is really an inexact art at best), we saw:

Traffic Type	Flows	Percent	Thompson
http	1282974	72.8	75%
dns	154129	8.8	18%
half-life	43844	2.5	
ping	36822	2.1	<1.5%
starsiege tribes	27176	1.5	
https	21509	1.2	
AIM	18757	1.1	
smtp	18330	1.0	2%
napster (definite)	11766	0.7	
nntp	8828	0.5	<1%
ftp	7316	0.4	<1%
http proxy	6796	0.4	
identd	5212	0.3	
qt4/rtsp/realaudio	5072	0.3	
pop3	4697	0.3	
msn instant messngr	4379	0.2	
napster (likely)	4245	0.2	
realaudio	4063	0.2	negligible
netbios	3677	0.2	
NTP	3386	0.2	
shoutcast	3340	0.2	
telnet	3056	0.2	<1%
Sub 7 trojan	2533	0.1	
quake/quake2/quakewo	2371	0.1	
microsoft netshow	2314	0.1	
nfs	1864	0.1	
RC5 distributed.net	1729	0.1	
TheZone	1530	0.1	
ssh	1059	0.1	
misc 1024-1052	1039	0.1	
socks proxy	917	0.1	
bgp	882	0.1	
uncategorizable	31685	1.8	see footnote[95]
categorizable (but < 0.1% each)	62671	3.5	

95. Totaling up the reported percentage values for web, dns, nntp, ftp, telnet, etc accounts for virtually 100% of the traffic in Thompson's study — however, if you inspect figure 7f in his paper, it becomes quite clear that he has excluded "other" applications from what constitutes his base set of flows. Looking at 7f, it appears his "other" category actually runs somewhere in the neighborhood of 20% of all flows.

If we look at that same data, breaking it down by octets instead of by flows, we see:

Traffic Type	Octets	Percent	Thompson	McCreary&KC
http	1.056E10	63.3	75%	58.9%
nntp	1.6752E9	10.0	1-4%[96]	11.7%
realaudio	7.0359E8	4.2	0.5-2.5%	1.35%
ftp	4.9944E8	3.0	2-8%	4.3%
napster (definite)	3.3703E8	2.0	n/a[97]	3.0%
napster (likely)	2.0539E8	1.2	n/a	n/a
smtp	1.5633E8	0.9		3.4%
microsoft netshow	1.4489E8	0.9		n/a
Hotline	1.3648E8	0.8		0.43%
https	1.3613E8	0.8		0.99%
qt4/rtsp/realaudio	1.1091E8	0.7		0.33%
half-life	1.008E8	0.6		0.47%
netbios	90342295	0.5		0.26%
AIM	56746045	0.3		
starsiege tribes	48585007	0.3		
orbix	43141014	0.3		
http proxy	39271949	0.2		1.2%
dns	35834263	0.2	1-2%	1.1%
scour	29313719	0.2		
Gnutella	27511472	0.2		
citrix	20988849	0.1		
ping	18397369	0.1		
ICQ	15888021	0.1		
shoutcast	15341546	0.1		0.16%
misc 1024-1052	13559994	0.1		
quake/quake2/quakewo	12583117	0.1		0.12%
RC5 distributed.net	11517236	0.1		
telnet	10298055	0.1		
categorizable (but < 0.1% each)				
uncategorizable	1.388E9	8.4	see footnote[98]	

96. It is worth noting that we believe NNTP traffic is doubling at a rate that exceeds the rate of growth for the Internet as a whole, hence we are not surprised that NNTP traffic inbound has increased above levels from three years ago. In fact, with a doubling time of six months, we believe that current NNTP traffic volumes are, on average 2^6=64 times what they were in 1997; on the other hand not all sites do NNTP.
97. Napster didn't exist in 1997.
98. As in the Thompson flow data, the Thompson octet data appears to exclude miscellaneous/ uncategorized flows from the totals used as the base for percentages.

What are all those different applications?

We will attempt to summarize the major different applications, proceeding alphabetically. Note that in some cases we have used the description provided by the application developer's web site, which can be a bit florid at times.

AIM:

> AOL Instant Messenger.[99] Described by AOL as "a free software program that lets you: receive instant alerts; send instant messages; share photos, pictures and sounds; enjoy live conversations online - free; chat with friends and family or people with similar interests; stay on top of the news and stocks."
>
> You do not need to be a member of AOL to use AIM -- it comes integrated withNetscape,[100] for example, and interoperates (to some extent) with MSN Messenger[101] and Tribal Voice's PowWow[102] (the product that ATT chose to adopt for instant messenging purposes).
>
> Uses TCP ports 5190 and 5050.

bgp:

> Border Gateway Protocol.[103] BGP4 is used to control routing of network traffic between ASNs. Uses port 179.

categorizable (but < 0.1% each):

> This miscellaneous category includes all categories of traffic which we were successfully able to categorize by application, but which amounted to less than one tenth of one percent each. Distinguish this category from the uncategorizable category, which contains all traffic that is not for any known/identifiable application.

---

99. http://www.aol.com/aim/
100. http://home.netscape.com/communicator/aolinstant/v4.0/
101. http://messenger.msn.com/
102. http://www.tribal.com/
103. ftp://ftp.isi.edu/in-notes/rfc1771.txt Also see: http://www.cisco.com/warp/public/459/18.html

citrix:

Citrix Winframe.[104] "WinFrame (R) application server software provides access to virtually any Windows (R) application, across any type of network connection to any type of client. Based on our innovative ICA (R) and MultiWin technologies, [etc.]" Uses ports 2023 and 1494.

dns:

Domain Name System.[105] Resolves dotted quads to fully-qualified domain names and vice versa. BIND (Berkeley Internet Name Domain) is the default/dominant implementation of DNS services for most sites.[106] Uses port 53.

ftp:

File Transfer Protocol.[107] Used to move text and binary files from one system to another. Supports both authenticated transfers (login and password required) and anonymous ftp (login with userid=anonymous, password=<your email address> by convention). One of the basic protocols supported by Netscape and Internet Explorer (e.g., you can specify ftp:// as a URL and the browser will know how to handle that). Among the most popular dedicated ftp client programs are WS_FTP for Windows and Fetch for the Mac. ftp is now being replaced at many sites by scp (Secure Copy), which uses ssh to insure that plain text passwords aren't sent over the network. Uses ports 20 and 21.

gnutella:

Peer-to-peer file sharing program (like Napster), but without reliance on any single central server which might serve as a central point of failure, and without any limitation on the type of files which can be shared. Named by combining GNU (emblem of the Free Software Foundation) with Nutella (chocolate-hazelnut spread popular in parts of the world as a topping for bread). Uses port 6346[108] and others.

---

104. http://www.citrix.com/products/winframe/
105. http://www.dns.net/dnsrd/rfc/ has a nice set of pointers to the relevant DNS RFCs
See also: http://www.oreilly.com/catalog/dns3 ("DNS and BIND, 3rd Edition," by Paul Albitz and Cricket Liu, ISBN 1-56592-512-2, September 1998)
106. http://www.isc.org/products/BIND/
107. ftp://ftp.isi.edu/in-notes/rfc959.txt
108. Posting by Michael Pifer <pifer@GRINNELL.EDU> to RESNET-L@listserv.nd.edu, April 10, 2000

half-life:

Interactive online 3D "shoot-em-up" game[109] from Sierra On-Line, apparently named as "Game of the Year" by many game magazines.[110] Uses ports 27005 and 27015.[111]

hotline:

"Hotline enables private and public virtual community building and live interaction with real time chat, conferencing, messaging, data warehousing and file transfer and viewing — performed with minimal technical knowledge. More than two million people already use Hotline to exchange information and ideas."[112]

Note that Hotline is another "Napster-like" program in that many users use it to share MP3s or software on a peer-to-peer basis. Napster servers are listed in sites called "trackers," some public and listing thousands of servers, others private and listing only a handful of servers at most.[113] Uses ports 5500 and 5501.[114]

http/http proxy:

HyperText Transfer Protocol, also known as the World Wide Web. Defined in RFC 1945[115] (May 1996); see also RFC 2616[116] (June 1999) and RFC 2617[117] (June 1999); an excellent online resource is at http://www.w3.org/Protocols/ Uses port 80 (and other non-standard ports).

HTTP proxies[118] are servers that act as "intermediaries," allowing a given user to connect to a web site "through" them. HTTP proxies tend to use ports 8080, 81, and other ports (3128 and 3130 are reported separately as Squid; 1080 is reported separately as Socks). Proxies are also commonly installed on port 80; proxies installed on port 80 are difficult to separate from web servers running on port 80 from a traffic analysis viewpoint.

---

109. http://www.planethalflife.com/half-life/faq.shtm
110. http://www.sierrastudios.com/games/half-life/
111. http://half-life.pcgame.com/console/net.html
112. http://www.hotlinesw.com/
113. See: http://www.hotlinecentral.com/ or http://www.tracker-tracker.com/hotline/trackers.shtml
114. See, for example: http://www.sensei.com.au/macarc/apple-internet-providrrs
115. ftp://ftp.isi.edu/in-notes/rfc1945.txt
116. ftp://ftp.isi.edu/in-notes/rfc2616.txt
117. ftp://ftp.isi.edu/in-notes/rfc2617.txt
118. http://www.ijs.co.nz/proxies.htm

https:

Secure HTTP. Uses SSL encryption to prevent interception of web traffic between a web browser and web server (e.g., when transmitting credit card information to an online store). You know you're using https when the URL for a web document says https instead of http, and the little "lock" at the bottom of the browser window closes. An example of a secure web server is Apache-SSL.[119] Uses port 443.

ICQ:

"ICQ is a revolutionary, user-friendly Internet tool that informs you who's on-line at any time and enables you to contact them at will. No longer will you search in vain for friends or associates on the Net. ICQ does the searching for you, alerting you in real time when they log on. The need to conduct a directory search each time you want to communicate with a specific person is eliminated. With ICQ, you can chat, send messages, files and URL's, play games, or just hang out with your fellow 'Netters' while still surfing the Net."[120] Uses port 4000 UDP plus additional ports (which can make it very difficult to track).

identd:

identd is designed to allow a large system, like one of the University of Oregon's primary timesharing hosts, to respond to queries from a remote system requesting the identity of a person attempting to access that remote system.

To understand what this means, consider a hypothetical abuse incident which might be reported by a remote system administrator. Let us pretend that that system administrator has observed something inappropriate being done to his system from one of our systems, such as gladstone.uoregon.edu (a large Sun with over 15,500 users). Let's assume that at the time of the incident, "only" 10% of those users were connected. Without identd, we'd have no possible way of identifying which of those fifteen hundred logged in users was the hacker/cracker; with identd running, the remote system will be able to log not just the source of the attack, but may also potentially be able to identify the actual user, assuming his server is set up to query for identd information. Note that identd queries are routine, and are not a sign that anything has been hacked/cracked by OWEN/NERO users. See RFC 1413[121] for more information. Uses port 113.

---

119. http://www.apache-ssl.org/
120. http://www.icq.com/

IRC:

Internet Relay Chat, defined in RFC 1459.[122] “IRC provides a way of communicating in real time with people from all over the world. It consists of various separate networks (or "nets") of IRC servers, machines that allow users to connect to IRC. The largest nets are EFnet (the original IRC net, often having more than 32,000 people at once), Undernet, IRCnet, DALnet, and NewNet. * * * Once connected to an IRC server on an IRC network, you will usually join one or more "channels" and converse with others there. On EFnet, there often are more than 12,000 channels, each devoted to a different topic. Conversations may be public (where everyone in a channel can see what you type) or private (messages between only two people, who may or may not be on the same channel).”[123]

Note that IRC includes an ability to “DCC” files from user to user. Because of this, IRC is a popular means of sharing the same type of files as Napster.[124]

Different servers often run on different ports, or on ranges of ports, most notably on port 6667.[125]

microsoft netshow:

A streaming multimedia protocol incorporating RealAudio and RealVideo technology.[126] Netshow has now been supplanted by “Windows Media Technology.” Uses port 1755 plus others.[127]

misc 1024-1052:

Many non-privileged applications will automatically bind to transient ports immediately above the highest privileged port (normally 1023).[128] There is virtually no way to identify what application may be using these transient ports, and the same application may use multiple ports, or different ports, for successive runs.

---

121. ftp://ftp.isi.edu/in-notes/rfc1413.txt
122. ftp://ftp.isi.edu/in-notes/rfc1459.txt
123. http://www.irchelp.org/irchelp/new2irc.html#what
124. http://www.wired.com/news/culture/0,1284,35141,00.html
125. http://www.mirc.com/servers.html
126. http://serverwatch.internet.com/reviews/av-netshow.html
127. http://www.microsoft.com/ntserver/mediaserv/deployment/planning/firewall.asp
128. See, for example, Internetworking with TCP/IP, Volume III: Client-Server Programming and Applications, BSD Socket Version, Douglas E. Comer and David L. Stevens, Prentice Hall, Englewood Cliffs, 1993, ISBN 0-13-474222-2, pp. 65.

msn instant messenger:

> "With MSNIM Messenger Service, you can: See when your friends are online and send them instant messages! Have group conversations. Be notified when you receive new e-mail to your MSN Hotmail account. Add instant messaging capabilities to Outlook Express. Invite your friends to a Windows NetMeeting conference or to play a DirectPlay game. Control who can see when you are online and send you messages."[129] Uses port 1863[130] and others.

napster:

> "So, what the heck is Napster? Napster is a completely new way of thinking about music online. Imagine...an application that takes the hassle out of searching for MP3s. No more broken links, no more slow downloads, and no more busy, disorganized FTP sites. With Napster, you can locate and download your favorite music in MP3 format from one convenient, easy-to-use interface.
>
> "What else does it do? Quite a bit, actually. Some highlights include: CHAT - Allows users to chat with each other in forums based on music genre. AUDIO PLAYER - Plays MP3 files from right inside Napster, in case you don't have an external player or would prefer not to use one. HOTLIST - Lets you keep track of your favorite MP3 libraries for later browsing."[131]
>
> See also the OpenNap project.[132]
>
> Uses ports 6699, 8875, 3333, 4444, 5555, 6666, 7777[133]

netbios:

> Protocol used for sharing printers and files on PC networks; not routable over the wide area Internet in itself, but able to be routed when sent as netbios-over-TCP. (See RFC1001.[134]) Note that some music-sharing programs such as Scour have been known to probe for exposed netbios shares.[135] Uses ports 137, 138, 139 and sometimes others.

---

129. http://messenger.msn.com/
130. http://messenger.msn.com/support/firewall.asp
131. http://www.napster.com/
132. http://opennap.sourceforge.net/
133. Posting by Michael Pifer <pifer@GRINNELL.EDU> to RESNET-L@listserv.nd.edu, April 10, 2000
134. ftp://ftp.isi.edu/in-notes/rfc1001.txt
135. http://navasgrp.home.att.net/tech/netbios.htm#Scour

nfs:

Network Filesystem; permits a remote system to mount a filesystem from a local server. NFS Version 3 is defined in RFC 1813.[136] Uses ports 1110 and 2049.

nntp:

Network News Transfer Protocol, or "Usenet." Defined in RFC 977[137] and commonly implemented with a number of widely accepted extensions.[138] See also RFC 1036,[139] which defines the format of Usenet articles.

In a nutshell, NNTP, Usenet or "network news" is a global discussion network in which postings are stored on local shared servers and then exchanged between peers. Users read news articles using "news readers." A nice discussion of "What is Usenet?" can be seen at http://www.faqs.org/faqs/usenet/what-is/part1/ Uses port 119.

NTP:

Network Time Protocol. "The Network Time Protocol (NTP) is used to synchronize the time of a computer client or server to another server or reference time source, such as a radio or satellite receiver or modem. It provides client accuracy typically within a millisecond on LANS and up to a few tens of milliseconds on WANs relative to a primary server synchronized to Coordinated Universal Time (UTC) via a GPS receiver, for example."[140] Uses port 123 UDP.

orbix:

"... provides the robust, flexible and scalable middleware infrastructure needed to solve today's business integration problems. Orbix enables an organization's software development team to eliminate or reduce the time spent on solving integration issues that are created by the organization's variety of hardware platforms, network protocols, application tools, programming languages, operating system or compiler versions."[141] See also the Object FAQ.[142] Uses port 1571.

---

136. ftp://ftp.isi.edu/in-notes/rfc1813.txt
137. ftp://ftp.isi.edu/in-notes/rfc977.txt
138. See: http://www.tin.org/docs.html
139. ftp://ftp.isi.edu/in-notes/rfc1036.txt
140. http://www.eecis.udel.edu/~ntp/ntp_spool/html/index.htm
141. http://www.iona.com/products/orbix/orbixchoice.html

ping:

Network command used to see if a remote system is reachable, to see if there is packet loss en route to it, and how long it takes for packets to make the round trip. Named after the sound that a submarine sonar unit produces.A nice discussion of ping is available online in “Connected: An Internet Encyclopedia.”[143] Uses ICMP echo.

pop3:

Post Office Protocol Version 3. Protocol used by Eudora and other email clients to download a user’s email from a central server to the user’s desktop workstation. Defined in RFC 1725.[144] Uses port 110.

qt4/rtsp/realaudio:

QuickTime4/Real-Time Streaming Protocol, used to effect the delivery of audio and video over RTP (RFC 1889[145]). RTSP itself is defined in RFC 2326.[146] A nice discussion of RTSP can be found at the RTSP FAQ.[147] Apple’s QuickTime 4[148] uses the same port and cannot readily be distinguished from RTSP. Uses port 554.

quake/quake2/quakeworld/team fortress:

Interactive online 3D shoot-em-up games fromm Id Software.[149] Uses ports 27910, 27960, 26000, 27000, 27001, 27500, 26810, and 29000 among others.

---

142. http://www.cyberdyne-object-sys.com/oofaq2/
143. http://www.FreeSoft.org/CIE/Topics/53.htm
144. ftp://ftp.isi.edu/in-notes/rfc1725.txt
145. ftp://ftp.isi.edu/in-notes/rfc1889.txt
146. ftp://ftp.isi.edu/in-notes/rfc2326.txt
147. http://www.cs.columbia.edu/~hgs/rtsp/faq.html
148. http://developer.apple.com/quicktime/
149. http://www.idsoftware.com/quake/
http://www.idsoftware.com/quake2/
http://www.idsoftware.com/quakeworld/

RC5 distributed.net:

Distributed.Net[150] consists of people all around the Internet who work together to use spare processing power from idle (or underutilized) systems to attack a variety of computationally difficult problems. Uses port 2064.

real audio:

RealNetwork's streaming web based network audio and video product.[151] Uses UDP ports 6970-7170 plus TCP ports 7070 and 7071.

scournet:

"Scour Exchange (SX) is a revolutionary software program that will change the way you experience entertainment. With Scour Exchange you can share your favorite music, videos, and even your most embarrassing photos with users all around the wired world. Find other users who share your vibe, and add them to your hotlist for quicker access to their file collection. Share and share alike — we're all friends here."[152] See also the note under "netbios," above. Uses port 8311.[153]

shoutcast:

Shoutcast is a product from Nullsoft[154] that allows music to be streamed to users running Nullsoft's Winamp[155] music player, effectively delivering "Internet radio" to the desktop. Uses TCP ports 8000, 8001, 8600, 8700, 8800.

smtp:

Simple Mail Transfer Protocol[156] — the way email gets transferred from one server to another. See the ESMTP extensions in RFC 1869.[157] Uses port 25.

---

150. http://www.distributed.net/
151. http://www.real.com/
152. http://www.scour.com/Software/Scour_Exchange/
153. Posting by Zachary.J.Spalding@Marist.edu to RESNET-L@listserv.nd.edu, May 2, 2000.
154. http://www.shoutcast.com/
155. http://www.winamp.com/
156. ftp://ftp.isi.edu/in-notes/rfc821.txt
157. ftp://ftp.isi.edu/in-notes/rfc1869.txt

socks proxy:

> "What is a SOCKS Proxy Server? When an application client needs to connect to an application server, the client connects to a SOCKS proxy server. The proxy server connects to the application server on behalf of the client, and relays data between the client and the application server. For the application server, the proxy server is the client."[158] Defined in a number of RFCs.[159] Uses port 1080.

ssh:

> Secure Shell. "SSH Secure Shell is the standard for remote logins and file transfer over the Internet. It encrypts all traffic, and provides a high level of protection against hacker attacks. Main features of Secure Shell include secure remote logins, terminal emulation, fully integrated secure file transfers, secure tunneling of X11 traffic, and secure access to e-mail over the Internet."[160] Uses port 22.

starsiege tribes:

> Online 3D interactive game. "Starsiege TRIBES is a revolutionary first-person shooter set in the Starsiege Universe which pits different warring tribes against each other. This first-person 3D action shooter is designed from scratch to focus on cooperative multiplayer gaming. Players use single-player training missions to develop the skills required to become full-fledged warriors, but the real heat of TRIBES radiates from multiplayer combat between two to 32 players connected over the Internet or LAN."[161] Uses UDP ports 28000 through 28008.[162]

Sub 7 trojan:

> "SubSeven is a trojan for the windows platform. It comes at least in two parts a client and a server. The client is used by the hacker to connect to the victim' s machine. Once the server.exe is installed on the victim's machine the hacker has full access to the victim's machine."[163] Uses ports 1243 and 27374 among others.

---

158. http://www.socks.nec.com/aboutsocks.html
159. http://www.socks.nec.com/socksprot.html
160. http://www.ssh.org/
161. http://www.tribesplayers.com/tribesplayers/promo.html
162. http://www.tribesplayers.com/tribesplayers/faq.html
163. http://www.sans.org/y2k/subseven.htm

telnet:

Telnet allows you to login to a remote system over a local area network or the Internet. Being replaced by ssh at many sites. Defined in RFC854.[164] Uses port 23.

TheZone:

Microsoft Gaming Zone, http://www.zone.com. Online multiplayer Internet games. Uses port 28800 and others.

uncategorizable:

This residual category represents the applications that we've been unable to identify to date.

164. ftp://ftp.isi.edu/in-notes/rfc854.txt

Miscellaneous Application-Related Notes: 3rd Party Web-Based Email

We continue to see one particular category of traffic that we're a little surprised to still see as strong as it is: free web-based email from companies such as Hotmail.

Hotmail and kindred offerings appear to be durable even in the face of non-commercial web email available directly from OWEN/NERO partners! For example, just looking at some OWEN/NERO partners:

— UO offers its own widely publicized[165] secure (SSL'd) commercially-procured local web email interface (IHub's Webmail[166]).

— Oregon State[167] and Portland State[168] both use Imp[169] as a webmail interface to their Unix hosts. (Oregon State also offers a web front end interface to its Microsoft Exchange Server)

— The Eugene 4J School District[170] offers Mailspinner[171] for web-based email, yet they still found that they needed to clamp down on free web based email accounts.[172]

All of these webmail clients offer web based access to local email and yet users continue to use free web based email solutions instead (or in addition to) their local email account. Why? We believe that the continuing popularity of Hotmail and related free web email services can be attributed to a variety of factors, including:

— Many users may have more than one account[173] — they may have an OWEN/NERO partner account and one or more free web-based email accounts which they also use.

---

165. UO's home page offers a direct link to https://webmail.uoregon.edu/ for example.
166. http://www.ihub.com/
Other universities known to be running WebMail include Alabama A&M, Alaska Fairbanks, Auburn, Ball State, Boston College, Cal Poly, Catholic, Citadel, George Washington University, Harvard, Kansas State, McGill, Notre Dame, Rhode Island, Tennessee State, University of Tennessee, Utah State, Vanderbilt.
167. http://webmail.orst.edu/
168. http://webmail.pdx.edu/
169. http://www.horde.org/imp/
170. http://webmail.4j.lane.edu/
171. http://www.mailspinner.com/
172. "High Schools Discourage Use of Free E-Mail Sites," NY Times on the Web, http://partners.nytimes.com/library/tech/00/05/cyber/education/31education.html
173. None of these polls represent scientific representative samples, however see:
http://webreview.com/pub/1999/06/04/poll/results.html
http://www.survey.com/UCERI/junkemail.html
http://slashdot.org/pollBooth.pl?qid=emails&aid=-1

— Some OWEN/NERO partners may not offer a satisfactory local email solution for some/all of their constituencies; in other cases, a free web-based email account may offer greater functionality (or more storage) than a local partner's email offering.

— Most free web email services are virtually anonymous,[174] and enable users to engage in activity that is not allowed on institutional email accounts provided by some OWEN/NERO partner acceptable use policies (e.g., recreational use, commercial use, use for political advocacy, etc.); free web email messages are also generally believed to be less "accessible" to authorities.

— Some free web based email accounts offer cryptographic security end-to-end, e.g. Hushmail[175] and Lokmail[176]

— Free web email accounts can be kept in perpetuity; in the case of OWEN/NERO partner email accounts, access to those accounts typically ends when the individual is no longer employed by the OWEN/NERO partner (or is no longer attending school, etc.).

— Some web-based services offered by free web email providers are closely integrated with that provider's web-based email (e.g., you need to get an email account with that provider to be able to access their clubs).

Why do we focus on free email services here?

— Email is a basic service that virtually everyone uses, hence it is important and potentially a material part of Internet transit bandwidth load (in thinking about the traffic associated with email, do not think about short text messages, think about large MIME'd[177] attachments).

— In the case of web email, if the email is from a local user to another local user, in the ideal world the traffic should all be local, and have zero transit bandwidth cost. In the case of "free" web email offered by Hotmail or other providers however, both creating and receiving that email uses wide area bandwidth.[178]

---

174. Typically the originating dotted quad associated with a given message is included in the message header, but that's about it.
175. http://www.hushmail.com
176. http://lokmail.net
177. MIME is the "multipurpose internet mail extensions" and permits non-text content to be sent via email (for example, Excel spreadsheets). For more information about MIME, see: http://www.hunnysoft.com/mime/
178. For the message content and for advertising framing that message.

— Assessing the traffic associated with free web email products has become more difficult as web email sites branch out, offering a wide assortment of non-email “value-added services” (over and above just free email), and as sites not traditionally associated with free email services (such as Yahoo) choose to add free web-based email as an additional service.

— In some ways the issue is moot: it is clear that users will continue to embrace “free” web-based email regardless of what OWEN/NERO might try to do to lure them toward a less bandwidth intensive alternative. On the other hand, we believe that traffic associated with Hotmail may soon flow via new peers at the Oregon IX, in which case at least some email products will (truly) be costless for consortia partners.

Miscellaneous Applications Related Notes: Network Security

A significant amount of incoming commodity transit traffic may also be associated with hacker/cracker denial of service-related attack attempts.

For example, as we looked at the flow data, we were surprised to see EUNet Bulgaria show up ranked as high as it did (it was the 15th highest ranked ASN by number of flows during the sampling period).

Looking at the ports associated with those flows, 99.9% (19,720 of 19,745) of them were UDP dstport 53, (e.g., DNS), and all but three queries had an apparent srcaddr of sagittarius.viket.net (193.68.157.2).

Looking at the distribution of dstaddrs, we saw:

DSTADDR	Frequency	Percent
159.121.107.80 panther.state.or.us	892	4.5
128.193.4.20 fido.nws.orst.edu	891	4.5
207.98.65.2 ns1.nero.net	886	4.5
140.211.91.9 NS2.sou.edu	881	4.5
198.236.20.8 dns.clackesd.k12.or.us	878	4.5
159.121.107.82 lynx.state.or.us	874	4.4
140.211.99.5 ns.sou.edu	872	4.4
131.252.208.38 walt.ee.pdx.edu	869	4.4
140.211.141.11 alpha.OIT.OSSHE.EDU	867	4.4
207.98.103.10 did not resolve	867	4.4
207.98.102.10 ocatens.ocate.edu	866	4.4
128.193.64.33 dnsman.OCE.ORST.EDU	862	4.4
131.252.129.53 loki.cc.pdx.edu	856	4.3
158.165.1.26 edlane.lane.edu	849	4.3
128.193.128.3 rex.nws.orst.edu	842	4.3
140.211.60.10 ladd.eou.edu	841	4.3
140.211.10.20 OSSHE.EDU	840	4.3

```
DSTADDR            Frequency   Percent
-------------------------------------
158.165.1.20            840       4.3
  rogue.lane.edu
140.211.117.7           836       4.2
  gilligan.wou.edu
140.211.91.8            835       4.2
  did not resolve
128.223.32.35           829       4.2
  phloem.uoregon.edu
140.211.135.12          827       4.2
  internet.OIT.EDU
128.223.21.15           818       4.1
  ruminant.uoregon.edu
128.193.232.138           1       0.0
  schumann.RCN.ORST.EDU
206.99.0.184              1       0.0
  did not resolve
```

All of those hosts are OWEN/NERO systems which do DNS service.

We are quite confident that this was a denial of service attack aimed at sagittarius.viket.net[179]

For comparison, and as an indication that the results above really are quite anomalous, the sites with the next highest number of DNS-related queries during the study period were:

```
2,716 from *.root-servers.net (normal)

2,430 from *.akamaitechnologies.com (normal)

2,368 from 216.228.201.67 (a system at Eastern Oregon Net, a level
       of traffic which strikes us as somewhat anomalous)

1,723 from *.adknowledge.com (normal)

1,713 from *.aol.com  (normal)
```

For background purposes: you should know that denial of service attacks try to do one of two things:

— render a targeted resource unusable directly by “overloading it” with spurious queries, e.g., in this case, to potentially overload our DNS server with bogus DNS requests, thereby rendering it less able to answer real DNS queries (however, this does not appear to be the case here since the queries moved from one host to another on a periodic basis)

179. http://www.ciac.org/ciac/bulletins/j-063.shtml

> — attempt to use the targeted host as an instrumentality for attacking some third party host, typically by forging or “spoofing” that third party host’s address into bogus requests that the attacker generates... We believe that’s what was going on here.

Thus, while this attack may appear to be coming “from” sagittarius.viket.net, we believe that sagittarius.viket.net is probably the *target* of this denial of service attack, not the origin of the attack. This analysis is consistent with the fact that the attack shifted from DNS server to DNS server during the course of the attack, a tactic that is designed to prevent the targeted host from blocking the attack on a systematic basis at a firewall or packet filter.

One example of a simple DNS-oriented denial of service attack is DOOMDN;.[180] more sophisticated denial of service attacks may employ an “army” of compromised “zombie” systems to launch a coordinated denial of service attack from many different locations using many different hosts and many different networks at once.[181]

Why did the attacker use OWEN/NERO resources to wage this attack? The answer is simple: OWEN/NERO has fast, well connected systems; it was just one of many networks used for this attack (and hence we shouldn’t feel too “special”); and we didn’t/couldn’t prevent this particular attack.

As this example illustrates, security related issues can be an integral part of any bandwidth audit, and security incidents can skew the applications which appear to be “widely used.”

---

180. http://www.securityfocus.com/templates/archive.pike?list=1&date=1997-07-29&msg=199907310000.AA154206596@sail.it
181. Examples include Trinoo, TFN, Stacheldraht, TFN2K and Shaft.

Section 12. Flow ASN Breakdown

Section 12 Keypoints

- This section looks at flows by automous system.

- Looking at the source ASNs of the flows we saw, 13.3% 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 a per octet 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, and AOL had 4.3%)

- Commodity transit traffic ASN distributions are expected to change in the weeks and months ahead for a variety of reasons; what's reported here should be taken as a snapshot at this point in time, only.

- A number of Oregon ASNs are showing up as connecting with OWEN/NERO via commodity transit; those Oregon providers should be strongly encouraged to consider peering at the OIX.

- Sources of traffic have changed over time; in the past, it was "well known" that most traffic would come from large network service providers; now major sources of traffic also include web hosting companies running large colo farms, broadband access companies (xDSL and cable modem companies), mega ISPs (e.g., AOL, Microsoft, etc.), and web advertising companies (IMGIS, DoubleClick, Web Side Story, etc.)

- We were also surprised at the number of flows associated with N2H2 (makers of the Bess Filtering System); that product appears to be popular with some OPEN customers.

What ASNs had traffic for us?

We will begin by showing the autonomous systems with the most inbound flows over our commodity transit links. Following that, we will show a second table listing the ASNs with the most inbound octets over our commodity transit links.

Remember that because we are only looking at our two commodity transit links, these tables will NOT show traffic going to ASNs that we access via the Oregon Internet Exchange, nor will they show traffic from OWEN/NERO I2 members to I2-connected ASNs, nor will they show intra-OWEN traffic.

We provide descriptions in footnotes for entities with at least 1% of flows (or octets); we provide citations to web sites for entities with less than 1% (but with at least 0.1%).

Source Autonomous System	Flows	Percent
[19 ASNs each had at least 1% of all flows]		
Exodus[182]	234042	13.3
AOL[183]	122639	7.0
GlobalCenter/Primenet[184]	112963	6.4
Abovenet[185]	99390	5.6
Hotmail[186]	61858	3.5
DoubleClick[187]	44170	2.5
Microsoft[188]	42687	2.4
UUNet[189]	38514	2.2

182. International network backbone and web hosting company. http://www.exodus.net/faq.html states that Exodus customers include “Yahoo!, DoubleClick, Excite, CBS MarketWatch.Com, CBS SportsLine, USA Today.com, Storage Networks, and American Greetings.com.”

183. http://www.corp.aol.com/whowhere.html states: “Founded in 1985, America Online, Inc., based in Dulles, Virgina, is the world’s leader in interactive services, Web brands, Internet technologies, and e-commerce services. America Online, Inc. operates: two worldwide Internet services, America Online, with more than 22 million members, and CompuServe, with more than 2.5 million members; several leading Internet brands, including ICQ, AOL Instant Messenger and Digital City, Inc.; the Netscape Netcenter and AOL.COM portals; and the Netscape Navigator and Communicator browsers; and AOL MovieFone, the nation’s largest movie listing guide and ticketing service. [continues]”

184. http://www.boardwatch.com/isp/summer99/bb/frontier.html states: “Frontier GlobalCenter emphasizes digital distribution for Web customers, with mirrored Web sites and software and hardware to determine download paths. Frontier GlobalCenter says it currently handles over 7.5 billion page views every month, 1.8 billion hits per day and has scaled to a sustained peak of 180 Mbps in file downloads for a single customer. Through Frontier GlobalCenter’s eight domestic and two international Media Distribution Centers and national combined optical IP, Packet Over SONET [POS] and ATM Backbone, the company provides direct Internet connectivity, web hosting and collocation services. Frontier GlobalCenter’s customers include such web sites as, Yahoo!, Washington Post, Newsweek Interactive, Motley Fool and eToys. The company claims more than 8,000 business Internet customers.”

Source Autonomous System	Flows	Percent
IMGIS[190]	36363	2.1
Internap.com[191]	26217	1.5
Excite[192]	24620	1.4
Netscape[193]	22869	1.3
Home.Net[194]	21187	1.2
BBN[195]	19756	1.1
EUNet Bulgaria[196]	19745	1.1
Genuity[197]	18347	1.0
Level3[198]	18273	1.0
softaware.com[199]	17151	1.0
Worldcom Advanced Networks[200]	17059	1.0

185. http://www.abovenet.net/company/overview.html states: “AboveNet Communications, a wholly owned subsidiary of Metromedia Fiber Network, is the architect of a global one-hop net work. Through its extensive peering relationships, it has built a network with the largest aggregated bandwidth in the world. In centralized, co-location facilities in San Jose, Vienna, VA, and New York City AboveNet brings together ISPs and Content Providers for maximum Internet performance.”
186. Web-based free email company owned by Microsoft.
187. Web-based advertising and market research company which also does customer profiling. See: http://www.doubleclick.net:8080/company_info/
188. Operating systems, application software, a national network, major portals, etc., etc., etc. See: http://www.microsoft.com/presspass/corpprof.htm
189. “UUNET, an MCI WorldCom company, is a global leader in Internet communications solutions offering a comprehensive range of Internet services to business customers worldwide. Providing Internet access, web hosting, remote access and other value-added services, UUNET offers service in over 100 countries, to more than 70,000 businesses, and owns and operates a global network in thousands of cities throughout North America, Europe and Asia Pacific. * * * Products & Services. UUNET's product portfolio contains cost-effective IP-based services including: Internet access: Dial-up and dedicated access from 56 Kbps to OC-12 speeds, and wholesale Internet access provisioned for Internet and online service providers 'UUcast' multicast services.” (see: http://www.uu.net/about/)
190. Web advertising and customer profiling company. (see: http://www.imgis.com/)
191. “InterNAP Network Services Corporation is a leading provider of fast, reliable and centrally managed Internet connectivity services targeted at businesses seeking to maximize the performance of mission-critical Internet-based applications. Customers connected to one of the Company's Private-Network Access Points ("P-NAPs") have their data optimally routed to and from destinations on the Internet in a manner that minimizes the use of congested public network access points and private peering points. This optimal routing of data traffic over the multiplicity of networks that comprise the Internet enables higher transmission speeds, lower instances of packet loss and greater quality of service. * * * Major e-commerce companies and networks served by InterNAP include Amazon.com, Datek Online, Go2Net, ITXC, MindSpring, The NASDAQ, TheStreet.com, WebTV and many others - including many local and regional ISPs.” (see http://www.internap.com/who.htm)
192. Traditionally, a major search engine, now part of Excite@Home (the cable modem services company). See http://www.home.net/about/facts.html
193. The browser and web portal company. For a list of their current product lines, see: http://home.netscape.com/products/index.html
194. See the footnote above for Excite@Home.

Source Autonomous System	Flows	Percent
[12 additional ASNs each have at least 0.5% of all flows, but < 1%]		
Sprintlink[201]	15041	0.9
Turner Broadcasting[202]	14248	0.8
ATT Worldnet[203]	13122	0.7

195. Boardwatch (http://www.boardwatch.com/isp/summer99/bb/gte.html)describes the company as: “In 1998 Bell Atlantic and GTE Corp announced a multi-billion dollar merger. Combined, the two companies say they would have a total of 63 million access lines providing last mile links in 38 states as well as an international presence in more than 30 countries. * * * GTE Internetworking, a unit of GTE Corp., includes the former BBN Corp. which 28 years ago developed the ARPANET, the forerunner to today’s Internet. * * * in December 1997 GTE Internetworking acquired Genuity Inc. increasing its hosting infrastructure to 12 distributed data centers * * * The former BBN Corporation had been involved with TCP/IP networking from its conception. The design firm, then Bolt Beranek and Newman, essentially built the Advanced Research Projects Administration (ARPANET) network, predecessor to the Internet. The firm invented the first IP packet router and devised the @ symbol convention for e-mail addressing.” A list of customers can be found at http://www.genuity.com/about/more_info/customers/list.htm
196. http://www.digsys.bg/company/ states, “Digital Systems, known also as EUnet Bulgaria is the first and leading Internet provider in Bulgaria. The company offers a wide range of network-based services to many companies and individual users. Digital Systems' infrastructure is covering the whole country. The company is managed from the headquarters in Varna, from where the communication links, traffic flows and network planning are controlled.”
197. See footnote above for BBN.
198. Boardwatch (http://www.boardwatch.com/isp/summer99/bb/level_3.html) describes the company as: “Level 3 Communications, Inc. is a communications and information services company that was originally founded in 1985 as Kiewit Diversified Group Inc. (KDG). KDG is the wholly-owned subsidiary of Peter Kiewit Sons, Inc. (PKS), a 114-year-old construction, mining, information services and communications company headquartered in Omaha, Nebraska. * * * Level 3 is building an international network designed for Internet Protocol technology. The Level 3 network will combine local and long distance networks, connecting customers through Gateways across the U.S. and in Europe and Asia. The company expects to complete the U.S. intercity portion of the network during the first quarter of 2001. In the interim, Level 3 has leased a national network over which it offers services. Level 3 will provide a full range of communications services, including local, long distance,international and Internet services.”
199. Softaware Networks customers include Fox Kids, Virgin Records, USA Today, Paramount, Universal Studios, and many others. See: http://www.softaware.com/customers.html
200. Boardwatch (http://www.boardwatch.com/isp/summer99/bb/mci.html) describes the company as: “MCI WorldCom Advanced Networks, one of two data and Internet units of MCI WorldCom Communications, is a provider of networking and hosting services in 114 countries. In January 1998, WorldCom acquired the CompuServe Network Services division from H&R Block in a stock-for-stock transaction valued at about $1.2 billion. Additionally, WorldCom agreed to acquire ANS Communications from America Online (AOL). WorldCom Advanced Networks was formed in May 1998, from the combined forces of CompuServe Network Services, ANS Communications, GridNet International and the Web services business unit of UUNET Technologies. The division was renamed MCI WorldCom Advanced Networks in September 1998, following WorldComís acquisition of MCI Communications. MCI WorldCom Advanced Networks serves more than 3,300 global companies through fully integrated, supported and managed Internet, intranet and extranet services.”
201. http://www.sprintlink.net/
202. http://www.turner.com/
203. http://www.att.net/

Source Autonomous System	Flows	Percent
Qwest[204]	12346	0.7
N2H2[205]	12074	0.7
DIGEX[206]	12047	0.7
MCI[207]	11477	0.7
US West[208]	11075	0.6
PSI[209]	10324	0.6
Amazon[210]	9765	0.6
Flycast[211]	9614	0.5
Flying Crocodile[212]	8713	0.5
[... and an additional 164 ASNs with 0.4%-0.1% of all flows]		
IBM[213]	7755	0.4
Bell Advanced Communications[214]	7708	0.4
Advance Publications[215]	6936	0.4
Ebay[216]	6725	0.4
Simple Network Communications[217]	6523	0.4
Progressive Networks[218]	6150	0.3
Pacific Bell Internet[219]	6099	0.3
digitalNation[220]	5437	0.3
Web Side Story[221]	5406	0.3
CERFnet[222]	4674	0.3
InterNAP[223]	4640	0.3
Concentric[224]	4626	0.3
Interland Inc[225]	4509	0.3

204. http://www.qwest.net/
205. http://www.n2h2.com/ (the Bess Filtering System Folks)
206. http://www.digex.net/
207. http://www.wcom.com/ (MCI merged with Worldcom)
208. http://www.uswest.net/
209. http://www.psi.net/index1.html
210. http://www.amazon.com/
211. http://www.flycast.com/ (web advertising and online profiling company)
212. http://www.flyingcroc.net/ (web hosting company at the Westin Bldg in Seattle)
213. http://www.ibm.com/
214. http://www.bell.ca/
215. http://www.advance.net/
216. http://www.ebay.com/
217. http://www.simplenet.com/ (web hosting company)
218. http://www.real.com/ (e.g., the RealAudio folks)
219. http://www.pacbell.net/
220. http://www.dn.net/ (web hosting company; a Verio company)
221. http://www.websidestory.com/ (web advertising and online profiling company)
222. http://www.cerf.net/
223. http://www.internap.com/
224. http://www.concentric.net/
225. http://interland.net/ (“ranked #1 web hosting provider for small- to medium-sized businesses”)

Source Autonomous System	Flows	Percent
NYSERNet[226]	4169	0.3
Mindspring[227]	4146	0.2
Electric Lightwave[228]	3931	0.2
USA.Net[229]	3729	0.2
Digital Island[230]	3688	0.2
Cable and Wireless[231]	3686	0.2
IACNet.com[232]	3588	0.2
Interliant[233]	3587	0.2
Infoseek[234]	3405	0.2
Arc Four[235]	3345	0.2
Road Runner[236]	3303	0.2
HP Britain[237]	3274	0.2
Eastern Oregon Net Inc[238]	3187	0.2
Arizona Tri Univ Network[239]	3100	0.2
Internic[240]	3044	0.2
Continental Cablevision[241]	2935	0.2
Globecomm[242]	2846	0.2
CRL[243]	2845	0.2
Cal State Univ[244]	2823	0.2
pair Networks[245]	2822	0.2
Time Warner[246]	2810	0.2
PFM Communications[247]	2789	0.2

226. http://nysernet.org/
227. http://www.mindspring.com/
228. http://www.eli.net/ (OIX participant)
229. http://www.usa.net
230. http://www.digitalisland.com/
231. http://www.cwix.net/
232. ASN 1830. http://www.iacnet.com/ (note that page may have been hacked/cracked; at the time it was checked, it consisted solely of the text, “This is a web page!?!...”)
233. ASN 5697. http://www.interliant.com/ (formerly clever.net and sagenetworks)
234. http://infoseek.go.com/
235. ASN 7964. http://www.arcfour.com/ (minimal page; note that “arcfour” is a cryptographic-related term)
236. http://www.rr.com/
237. http://www.hp.com/
238. http://www.eoni.com/
239. ASN 2900. No web site per se; see: http://www.asu.edu/
240. http://www.internic.net/
241. http://www.mediaone.com/ (Continental Cablevision was acquired by Mediaone)
242. http://www.iaf.iname.com/info/company/aboutus.html (GlobeComm, Inc. divisions include iName, BestDomains, and GlobalDomains)
243. http://www.crl.com (now acquired by AppliedTheory Corporation)
244. http://www.calstate.edu/
245. http://www.pair.com/ (Pittsburgh PA web hosting company hosting “over 91,000 sites”
246. http://www.timewarner.com/

Source Autonomous System	Flows	Percent
Navisite Internet Services[248]	2766	0.2
NASA Internet[249]	2729	0.2
NERO[250]	2723	0.2
Shore.Net[251]	2675	0.2
Japan NIC[252]	2607	0.1
UUNet Canada[253]	2575	0.1
Network Solutions[254]	2530	0.1
Online Computer Library Center[255]	2527	0.1
Maxim Computer Systems[256]	2501	0.1
Data Research Group[257]	2471	0.1
CheckOut.Com[258]	2460	0.1
Yahoo Broadcast Services[259]	2454	0.1
Hurricane Electric Internet[260]	2361	0.1
Earthlink[261]	2332	0.1
NBC Internet[262]	2323	0.1
OLM LLC[263]	2308	0.1
UUNet Customer[264]	2265	0.1
Easystreet[265]	2189	0.1
BTnet UK Regional Network[266]	2187	0.1

247. http://www.pfmc.net/ (now part of Globix Corporation)
248. http://www.navisite.com/
249. ASN 297. http://www.nisn.nasa.gov/ (NASA Science Internet --> NASA Internet --> NISN)
250. ASN 3701. http://www.nero.net/ -- our own OWEN/NERO. Should not show up on the “outside” of our transit links coming inward.
251. http://www.shore.net/
252. http://www.nic.ad.jp
253. http://www.uunet.ca/
254. http://www.networksolutions.com/
255. http://www.oclc.org/
256. http://www.maxim.net (web hosting company; nice map showing their Bay Area connectivity at http://www.maxim.net/map/index.html)
257. http://www.willamette.net/
258. http://www.checkout.com/ (“The Entertainment Network” -- includes CheckOutMusic.com, CheckOutMovies.com, CheckOutGames.com, The Lounge and Warehouse Online Stores.)
259. http://www.broadcast.com/
260. http://www.he.net/ (web hosting company in Fremont California)
261. http://www.earthlink.net/ (following the combination of Mindspring and Earthlink, they are now the “second largest Internet service provider in the United States.”)
262. http://www.nbci.com/ (“NBC Internet (NBCi) is the publicly traded company that would combine Xoom.com, Snap.com, three Internet businesses contributed by NBC -- NBC.com, NBC’s Interactive Neighborhood and Videoseeker.com -- and a 10 percent ownership stake in the new CNBC.com” (quoting http://www.xoom.com/about/nbci/))
263. http://www.olm.net/ (web hosting company in Lisle, Illinois, hosting over 50,000 domains)
264. Customer of UUNet (http://www.uu.net/)
265. http://www.easystreet.com/ (“Portland’s largest Independent Internet service provider.”)
266. http://www.bt.net/ (British Telecom)

Source Autonomous System	Flows	Percent
Alchemy Communications[267]	2185	0.1
pogo.com[268]	2168	0.1
NETCOM[269]	2146	0.1
Apple Computer[270]	2123	0.1
Savvis[271]	2109	0.1
Semaphore Corporation[272]	2036	0.1
LightRealm Communication[273]	2017	0.1
Splitrock[274]	2000	0.1
Cymitar Network Systems[275]	1958	0.1
CommuniTech.Net Inc[276]	1915	0.1
Bellsouth[277]	1892	0.1
Innovative Access Inc[278]	1876	0.1
MichNet[279]	1850	0.1
Deutsche Telekom[280]	1836	0.1
MIXNet[281]	1789	0.1
Transport Logic[282]	1779	0.1
ANS[283]	1716	0.1

267. http://www.alchemyfx.com/ (web hosting company in Los Angeles, claims to be “building a network with the largest aggregated bandwidth in the world”)
268. http://www.pogo.com/ (formerly the “Total Entertainment Network”; operates the “leading online games service targeting the rapidly expanding ‘family games’ market—games that appeal to everyone.” They partner with @Home, Altavista, CNET, Excite, GO Network, MediaOne, Netscape, Road Runner, Snap.com, WebCrawler, Xoom, etc.)
269. http://www.netcom.com/ (now part of Mindspring/Earthlink)
270. http://www.apple.com/
271. http://www.savvis.net/
272. http://www.semaphore.com/ (web hosting company at the Westin Building in Seattle)
273. http://www.lightrealm.com/ (owned by Micron Electronics, Inc.; offers “ultra-secure data center in former military command center” located in Moses Lake, Washington, plus additional data centers in LA, Seattle and Boise; Micron (dba as Lightrealm, HostPro and Micron Internet) hosts over 70,000 sites, making it “the fourth largest hosting company in the country.”)
274. http://www.splitrock.net/ (“Splitrock currently provides nationwide Internet dial access and related services to Prodigy, our primary customer and one of the largest Internet service providers in the United States.” Splitrock also sells dialup capacity, and is purchasing fiber capacity nationally in a network covering 15,000 route miles across the United States)
275. http://www.rackspace.com/ (web hosting company located in San Antonio Texas)
276. http://www.communitech.net (“rated as one of the top 25 webhosting companies in the world by C|Net and HostIndex”)
277. http://www.bellsouth.net/
278. http://www.inaxx.net/ (web hosting company located in Atlanta Georgia)
279. http://www.merit.edu/michnet/
280. http://www.dtag.de/english/
281. Minnesota Internet Exchange, formerly part of MR.Net (which became Onvoy). See: http://www.onvoy.com/
282. http://www.transportlogic.com/
283. http://www.ans.net/ (showing that ANS has become part of UUNet, an MCI Worldcom company)

Source Autonomous System	Flows	Percent
BCTEL Advanced Communications[284]	1712	0.1
Whole Earth Networks[285]	1689	0.1
Pointcast[286]	1661	0.1
Data Return[287]	1647	0.1
Colorado Internet Cooperative[288]	1643	0.1
Tera-byte Online Services[289]	1635	0.1
UMI[290]	1626	0.1
SCN Research Inc[291]	1625	0.1
Pipex[292]	1617	0.1
Web Professionals[293]	1567	0.1
Time Inc[294]	1548	0.1
Washington State K-20[295]	1534	0.1
Telus Advanced Communications[296]	1520	0.1
Management Analysis Inc[297]	1512	0.1
Online Career Center[298]	1504	0.1
InfoStructure[299]	1492	0.1
Sprint IP Dial[300]	1486	0.1
aracnet.com[301]	1484	0.1
State of WA Info Services[302]	1478	0.1
Lexis-Nexis[303]	1475	0.1

284. http://www.bctel.net/ (Telus, Canada's second largest telecommunications company)
285. http://www.gstwenet.net/, part of GST Telecommunications (http://www.gstcorp.com/)
286. http://www.pointcast.com/ (becoming Entrypoint http://www.entrypoint.com/)
287. http://www.datareturn.com/ (web hosting company located in Irving Texas)
288. http://www.coop.net/ (Member owned and operated Colorado Internet cooperative; amazing rates for DS3 connectivity -- full T3 ("43Mbps") is $16,500/month plus $13,000 (one time costs) plus loop, and T1 (1.5Mbps) is $550/month plus $3,100 (one time costs) plus loop.
289. http://www.tera-byte.com/ (web hosting company in Edmonton, Alberta; also offers the Spaceports free web hosting service with over 100,000 users; "Canada's largest web host")
290. http://www.umi.com/ (Bell and Howell Information and Learning, formerly University Microfilm International)
291. http://www.scn.rain.com/
292. http://www.pipex.net/ ("the UK's first commercial Internet Service Provider..." "still the preferred choice amongst business executives and professionals, with no fewer than 83 of The Times Top 100 Companies and literally tens of thousands of small to medium-sized businesses..."
293. http://www.professionals.com/ (web hosting company located at the Palo Alto IX)
294. http://www.time.com/
295. http://www.wa-k20.net/
296. http://www.telus.com/ (aka BCTel, see above)
297. http://www.mainet.com/
298. http://www.occ.com/ (aka http://www.monster.com/)
299. http://www.mind.net/
300. http://www.sprintbiz.com/ (now Earthlink)
301. http://www.aracnet.com/
302. http://www.wa.gov/dis/

Source Autonomous System	Flows	Percent
Northwest Link[304]	1469	0.1
Epoch[305]	1467	0.1
Servint.com[306]	1464	0.1
NCREN[307]	1439	0.1
Digiweb[308]	1416	0.1
Internet Broadcasting System[309]	1397	0.1
ACSI[310]	1386	0.1
DACOM Korea[311]	1386	0.1
University of Dortmund[312]	1378	0.1
9 Net Avenue[313]	1377	0.1
OARNet[314]	1377	0.1
RealSelect Inc[315]	1346	0.1
InfoSpace.Com Inc[316]	1334	0.1
WebGenesis Inc[317]	1330	0.1
Sabre Group[318]	1322	0.1
Smartnap.com[319]	1308	0.1
Empire Net[320]	1307	0.1
National Supervisory Network[321]	1307	0.1
UUNet Netherlands[322]	1293	0.1
Unknown (ASN=0)	1290	0.1
Erols[323]	1284	0.1

303. http://www.lexis-nexis.com/lncc/
304. http://www.nwlink.com/
305. http://www.epoch.net/ (“the nation’s largest privately held ISP”)
306. http://www.servint.net/home.html (“the only privately-held and internally-funded backbone” based in McLean Virginia and run by 24 year old Reed Caldwell)
307. http://www.ncren.net/
308. http://www.digiweb.com/ (web hosting company in College Park, Maryland with over 15,000 customers)
309. http://www.ibsys.com/
310. http://www.acsint.net/ (a subsidiary of e.spire Communications, Inc, http://www.espire.net/)
311. http://www.dacom.co.kr/english/home_e.html
312. http://www.uni-dortmund.de/UniDo/Index_en.html
313. http://www.9netave.net/ (a web hosting company located in Secaucus, NJ, which Entrepreneur Magazine listed as the “fastest growing company in the industry” in their June ‘99 issue)
314. http://www.oar.net/
315. http://www.realtor.com/
316. http://www.infospace.com/ (home of “ActiveShopper” web price comparison agent)
317. http://www.webgenesis.com/
318. http://www.sabre.com/
319. http://www.smartnap.com/
320. http://www.empnet.com/
321. http://nsn.net/ (“NSN is a wholly owned subsidiary of Clear Channel Communications...”) Satellite capacity reseller.
322. http://www.nl.uu.net/
323. http://www.erols.com/

Source Autonomous System	Flows	Percent
CAIS[324]	1282	0.1
Telianet Sweden[325]	1265	0.1
Sonera Finland[326]	1260	0.1
Teleglobe[327]	1259	0.1
Winstar[328]	1236	0.1
Conacyt[329]	1223	0.1
CDS Internet[330]	1221	0.1
Intel[331]	1218	0.1
UGO Networks Inc[332]	1216	0.1
Micron Internet Services[333]	1192	0.1
Cybercon Inc[334]	1191	0.1
Southwestern Bell[335]	1184	0.1
Southern Online Systems[336]	1180	0.1
Nabisco Foods[337]	1164	0.1
McAfee Associates[338]	1163	0.1
Cascade Communication[339]	1159	0.1
NWNexus[340]	1147	0.1
Demon[341]	1128	0.1
Virtualis Systems[342]	1112	0.1
American Digital Network[343]	1109	0.1
Vector Internet Services[344]	1102	0.1
Software Partners[345]	1098	0.1

324. http://www.cais.com/
325. http://www.telia.net/
326. http://www.sonera.fi/
327. http://www.teleglobe.com/ (“...with the most extensive global Internet network, Teleglobe provides service to ISPs worldwide and connects over 60,000 businesses...” “Teleglobe was the first service provider to offer high-speed bandwidth services up to 10 Gbit/s on an intercontinental and intracontinental basis.”)
328. http://www.winstar.net/
329. http://www.conacyt.mx/
330. http://www.cdsnet.net/
331. http://www.intel.com/
332. http://www.ugo.com/ (“UnderGroundOnline: Your source for Games, TV, Film, ...”)
333. http://www.micron.net/
334. http://www.cybercon.com/ (web hosting company located in St Louis)
335. http://www.swbell.com/
336. http://www.socomm.net/ (web hosting company located in Memphis)
337. http://www.nabisco.com/
338. http://www.nai.com/
339. http://www.pond.net/
340. http://www.nwnexus.net/
341. http://www.demon.net/
342. http://www.virtualis.com/ (web hosting company located in Studio City California)
343. http://www.adnc.com/index.html
344. http://www.visi.com/

Source Autonomous System	Flows	Percent
Shaw Fiberlink[346]	1092	0.1
Sprint Canada[347]	1090	0.1
Videotron Telecom Ltee[348]	1083	0.1
CompleteWeb.Net[349]	1076	0.1
C I Host[350]	1054	0.1
CDnow Inc[351]	1052	0.1
Cove Software[352]	1048	0.1
Rogers Network Services[353]	1044	0.1
CMG Direct Interactive[354]	1034	0.1
EDS[355]	1019	0.1
Internet Connect Inc[356]	1009	0.1
Berkeley[357]	995	0.1
Netlimited LLC[358]	993	0.1
Aces Research Inc[359]	985	0.1
You Tools Corp/Fast.Net[360]	983	0.1
Digital Telemedia[361]	981	0.1
Critical Path[362]	971	0.1
Delta.net[363]	969	0.1
Bell Atlantic[364]	964	0.1
Dell Computer[365]	962	0.1
US Dept of Agriculture[366]	959	0.1
Myriad Corp[367]	932	0.1

345. ASN 7750 (http://www.onsale.com/, http://www.egghead.com/)
346. http://www.fiberlink.net/
347. http://www.sprint.ca/
348. http://telecom.videotron.com/en/
349. http://www.completeweb.net/ (web hosting company in Columbus Ohio)
350. http://www.cihost.com/ (web hosting company in Bedford Texas)
351. http://www.cdnow.com/
352. http://www.covesoft.com/ (web hosting company in Annapolis Maryland)
353. http://www.rogers.com/
354. http://www.cmgi.com/
355. http://www.eds.com/
356. http://www.inc.net/
357. http://www.berkeley.edu/
358. http://www.netlimited.net/ (aka http://www.hostpro.net/ a web hosting company)
359. http://www.aces.com/
360. http://www.fast.net/main.html
361. http://www.dti.net/
362. http://www.cp.net/
363. http://www.delta.net/
364. http://www.bellatlantic.com/
365. http://www.dell.com/
366. http://www.usda.gov/
367. http://www.cox-internet.com/

Source Autonomous System	Flows	Percent
BBC Internet NY[368]	926	0.1
IDT[369]	925	0.1
EBSCO Publishing[370]	924	0.1
Sandbox Entertainment[371]	924	0.1
Alabanza Inc[372]	921	0.1
Clippernet[373]	897	0.1
Cisco Systems[374]	891	0.1
Infonautics Corp[375]	885	0.1

368. http://www.bbc.co.uk/
369. http://www.idt.com/
370. http://www.epnet.com/
371. http://www.sandbox.net/
372. http://www.alabanza.com/ (a web hosting company servicing over 65,000 domains)
373. http://www.clipper.net/
374. http://www.cisco.com/
375. http://www.infonautics.com/

That was transit traffic by ASN ranked by number of flows — how about traffic per ASN on a per-octet basis?

The following table shows inbound transit usage per ASN on a per-octet basis. [We have not refootnoted ASNs which were footnoted in the previous section]

Source Autonomous System	Octets	Percent
[18 ASN's, each accounting for at least 1% of all octets received]		
Exodus	2.0811E9	12.5
GlobalCenter/Primenet	1.2141E9	7.3
Abovenet	8.1985E8	4.9
AOL	7.1614E8	4.3
Hotmail	4.3146E8	2.6
Microsoft	3.9253E8	2.4
Home.Net	3.2092E8	1.9
Level3	2.943E8	1.8
UUNet	2.8368E8	1.7
Internap.com	2.7734E8	1.7
MIXNet	2.3562E8	1.4
Yahoo Broadcast Services	2.202E8	1.3
Stupi[376]	2.1533E8	1.3
Genuity	2.1409E8	1.3
US West	1.8017E8	1.1
McAfee Associates	1.7913E8	1.1
CERFNet	1.7133E8	1.1
MCI	1.7083E8	1.0
[19 additional ASN's have < 1% but >= 0.5% of all octets]		
BBN	1.4473E8	0.9
CRL	1.4429E8	0.9
Netscape	1.3962E8	0.8
Qwest	1.2337E8	0.7
Worldcom Advanced Netwk	1.1909E8	0.7
Excite	1.1268E8	0.7
Interland Inc	1.0856E8	0.7
Maxim Computer Systems	1.0638E8	0.6
Turner Broadcasting	1.0485E8	0.6
ConXion Corp[377]	1.0484E8	0.6
Ebay	1.0006E8	0.6
Videotron Telecom Ltee	87675858	0.5
DIGEX	86224881	0.5
Cal State Univ	85279856	0.5
ATT Worldnet	80653437	0.5
Sprintlink	77907753	0.5

376. http://www.stupi.se/ (This is the site of the senior network engineer for Sprint in Europe)
377. http://www.conxion.net/ (web hosting company with egress capacity in excess of 17Gbps and network presence in a variety of sites including Portland, Seattle, Sacramento and SFO. Very interesting company.)

Source Autonomous System	Octets	Percent
PCWorld Online[378]	76353791	0.5
Concentric	75996485	0.5
[...and an additional 181 ASNs each account for 0.4%-0.1%]		
Time Warner	72990278	0.4
softaware.com	72885139	0.4
Wirehub[379]	70195546	0.4
ValuServe.Com[380]	67717864	0.4
Cablevision Systems[381]	66022974	0.4
Sun Microsystems[382]	63702499	0.4
DoubleClick	61081192	0.4
Empire Net	60762678	0.4
Continental Cablevision	60269828	0.4
PSI	59821442	0.4
Simple Network Comm	59421213	0.4
IBM	57227736	0.3
NYSERNet	57018023	0.3
HE.Net	56069014	0.3
Progressive Networks	54212961	0.3
Indiana University[383]	52464567	0.3
Cable and Wireless	52368726	0.3
Amazon	50854315	0.3
Fibrcom/Fibrnet[384]	50232661	0.3
Advance Publications	50196986	0.3
Bell Advanced Comm	50134622	0.3
Pacific Bell Internet	47198724	0.3
Williams Comm[385]	45555438	0.3
digitalNation	44879688	0.3
Iowa Network Services[386]	38462679	0.2
Infoseek	37647789	0.2
Apple	37039862	0.2
PFM Communications	36979942	0.2
Good.net	35889744	0.2
CANet[387]	35846718	0.2
University of Tulsa[388]	35772243	0.2
DACOM Korea	35306310	0.2

378. http://www.pcworld.com/ (computer magazine site)
379. http://www.wirehub.net/
380. http://www.valuserve.com/ (however note that we were unable to reach that site when testing)
381. http://www.cablevision.com/
382. http://www.sun.com/
383. http://www.indiana.edu/
384. http://www.fibr.net/ (however note that Fibrcom is now part of TimeWarner Telecom, see: http://www.twtelecom.com/)
385. http://www.wcg.net/
386. http://www.isintouch.com/
387. http://www.canet3.net/
388. http://www.utulsa.edu/

Source Autonomous System	Octets	Percent
Digital River[389]	35086206	0.2
Delta.net	34473662	0.2
Interactive Telecom[390]	32649179	0.2
Flying Crocodile	32493121	0.2
ViaNet[391]	32109667	0.2
US Air Force[392]	31848124	0.2
Alchemy Communications	31841588	0.2
WebGenesis Inc	31312310	0.2
Flycast	30318002	0.2
pair Networks	30132201	0.2
Cymitar Network Systems	30067166	0.2
Nabisco Foods	29451266	0.2
Proxad ISP[393]	29023941	0.2
Mindspring	28615300	0.2
SkyCache Inc	28378583	0.2
OLM LLC	28343463	0.2
IDT	28129573	0.2
CommuniTech.Net Inc	27811077	0.2
IACNet.com	27360251	0.2
InterNAP	27303243	0.2
Fiber Network Solutions[394]	27289218	0.2
American Digital Network	27214275	0.2
Servint.com	27115619	0.2
Internet Connect Inc	26604526	0.2
Earthlink	25904241	0.2
SiteStream Inc[395]	25740008	0.2
University of Arizona[396]	25395366	0.2
Electric Lightwave	25352313	0.2
Telepac Portugal[397]	24822599	0.1
Road Runner	24783138	0.1
Pipex	24726698	0.1
National Supervisory Network	24291679	0.1
Cogeco Cable[398]	23229352	0.1
SGI[399]	23076526	0.1
Data Research Group	22924987	0.1
OARNet	22828361	0.1
World Online France[400]	22721794	0.1

389. http://www.digitalriver.com/
390. http://www.itninc.net/
391. http://www.via.net/
392. http://www.af.mil/
393. http://www.proxad.com/
394. http://www.fnsi.net/
395. http://www.sitestream.net/
396. http://www.arizona.edu/
397. http://www.telepac.pt/
398. http://www.cgocable.net/
399. http://www.sgi.com/

Source Autonomous System	Octets	Percent
Savvis	22515979	0.1
U Minnesota[401]	22470632	0.1
Berkeley	22213036	0.1
New Mexico Technet[402]	22181365	0.1
Excalibur Group[403]	21598559	0.1
LightRealm Communication	21226232	0.1
HP Britain	21172854	0.1
9 Net Avenue	20830712	0.1
BCTEL Advanced Communications	20805824	0.1
SingNet Singapore[404]	20417205	0.1
Shore.Net	19756285	0.1
Lincoln Telecommunication[405]	19638199	0.1
UUNet Customer	19627454	0.1
Japan NIC	19204180	0.1
CheckOut.Com	19139787	0.1
Shaw Fiberlink	19101432	0.1
Texas Networking Inc[406]	19016049	0.1
Adelphia Corp[407]	18780309	0.1
Semaphore Corporation	18465339	0.1
Netlimited LLC	18079306	0.1
CDnow Inc	18031321	0.1
CM2.Com[408]	17910892	0.1
UUNet Netherlands	17722314	0.1
Ebone[409]	17529676	0.1
Teleglobe	17402503	0.1
Erols	17308222	0.1
Data Return	17281803	0.1
Dell Computer	17102404	0.1
Whole Earth Networks	17090007	0.1
Novell[410]	16936181	0.1
N2H2	16866827	0.1
Online Computer Library Center	16445256	0.1
Colorado Internet Cooperative	16387964	0.1

400. http://www.worldonline.fr/
401. http://www.umn.edu/
402. http://www.technet.nm.org/
403. ASN 10311, record last updated June 23, 1997. Since that it time, it looks as if the Excalibur Group has become part of MediaOne (according to a December 1997 article available at http://www.zdnet.com/zdnn/content/inwk/0444/264100.html)
404. http://www.singnet.com.sg/
405. ASN 3850, record last updated September 26, 1994. Unable to find a relevant web site for this autonomous system number (I believe it should be/used to be http://www.letc.net, but that site is down/unreachable)
406. http://www.texas.net/
407. http://www.adelphia.net/
408. http://www.cm2.com/
409. http://www.ebone.net/
410. http://www.novell.com/

Source Autonomous System	Octets	Percent
Digital Island	16053525	0.1
NBC Internet	16049784	0.1
Digiweb	15850868	0.1
Myriad Corp	15674369	0.1
US Dept of Agriculture[411]	15563635	0.1
Arc Four	15489887	0.1
IMGIS	15438740	0.1
SOVAM Teleport Moscow[412]	15241882	0.1
UUNet Canada	14947135	0.1
Rogers Network Services	14714604	0.1
Globecomm	14561022	0.1
Five Colleges Network Mass[413]	14538534	0.1
ArosNet Inc[414]	14171124	0.1
NWNexus	13968921	0.1
Estpak Data Ltd Estonia[415]	13822166	0.1
RealSelect Inc	13709189	0.1
Execpc.com[416]	13546495	0.1
SWIPnet[417]	13541135	0.1
Diamond Multimedia[418]	13491139	0.1
UT Austin[419]	13384604	0.1
MichNet	13353923	0.1
Omnilink Germany[420]	13209594	0.1
Innovative Access Inc[421]	13199452	0.1
Utah Education Network[422]	13137290	0.1
Management Analysis Inc	13052482	0.1
HANARO Telecom Korea[423]	12937231	0.1
Chello Broadband Europe[424]	12932646	0.1
Sandbox Entertainment	12881051	0.1
Bellsouth	12794673	0.1
Hitter Communications[425]	12761835	0.1
Sprint Canada	12596665	0.1
CMG Direct Interactive	12466524	0.1

411. http://www.usda.gov/
412. http://www.goldentelecom.ru/eng/
413. ASN 1249. Associated with the University of Massachusetts (http://www.umass.edu/)
414. http://www.aros.net/
415. http://www.estpak.ee/
416. http://www.execpc.com/ (now becoming voyager.net)
417. http://www.swipnet.se/
418. http://www.diamondmm.com/
419. http://www.utexas.edu/
420. http://www.omnilink.net/
421. http://www.inaxx.net/
422. http://www.uen.org/
423. http://www.hanaro.com/english/main.html
424. http://www.chello.nl/
425. http://www.hitter.net/

Source Autonomous System	Octets	Percent
aracnet.com	12371790	0.1
Southwestern Bell	12272753	0.1
You Tools Corp/Fast.Net	12105693	0.1
I3S, Inc[426]	12091106	0.1
NASA Internet	11999562	0.1
AINet[427]	11956810	0.1
CONNECTnet Internet[428]	11945945	0.1
Geological Survey[429]	11396310	0.1
UMI	11155143	0.1
Internet Direct Canada[430]	11070240	0.1
NCREN	11039137	0.1
Online Career Center	11037573	0.1
Advanced Internet Tech[431]	11025224	0.1
Lexis-Nexis	10989963	0.1
CDS Internet	10977127	0.1
Sabre Group	10875431	0.1
InfoSpace.Com Inc	10858634	0.1
University of Idaho[432]	10818355	0.1
USA.Net	10482467	0.1
HINet[433]	10479949	0.1
IHETS Indiana[434]	10415218	0.1
Arachnitec, Inc[435]	10173719	0.1
Epoch	10163427	0.1
Interactive Classified Ntwk[436]	10153382	0.1
Cisco Systems	10099078	0.1
CAIS	9885011	0.1
UGO Networks Inc	9824255	0.1
Megsinet[437]	9806684	0.1
Internet Direct[438]	9692445	0.1
Software Partners	9635855	0.1
C I Host	9598819	0.1

426. http://www.i3s.com/ (which redirects to http://www.bbnow.com/ announcing their name change to BroadbandNOW)
427. http://www.ai.net/
428. http://www.connectnet.com/
429. http://www.usgs.gov/
430. ASN 7271. http://www.idirect.com/ (which redirects eventually to http://www.looktown.com)
431. http://www.aitcom.net/
432. http://www.uidaho.edu
433. http://www.hinet.net/
434. http://www.ihets.org/
435. ASN 6921. Looks like it should be arachnitec.net, but that domain doesn’t seem to currently have a web site available.
436. AS 10404. Several search engines could not find “Interactive Classified Network” but it is apparently affiliated with http://www.dataway.com/
437. http://www.corecomm.net/
438. ASN 3812. http://www.direct.ca/ (which redirects eventually to http://www.looktown.com)

Source Autonomous System	Octets	Percent
Nacamar Data Comm[439]	9591116	0.1
pogo.com	9580311	0.1
Navisite Internet Services	9419763	0.1
Easystreet[440]	9310148	0.1
Brooks Fiber[441]	9076415	0.1
Bell Atlantic	9070039	0.1
Cove Software	9021807	0.1
Smartnap.com	8848011	0.1
Web Side Story	8837649	0.1
ONet Ontario[442]	8712287	0.1
Pointcast	8702084	0.1
Micron Internet Services	8567504	0.1
The Grid[443]	8547105	0.1
Xmission LLC[444]	8532118	0.1
NTIS[445]	8528813	0.1
Seoul National Univ[446]	8490774	0.1

439. http://www.nacamar.net/
440. http://www.easystreet.com/
441. http://www.brooks.net/
442. http://www.onet.on.ca/
443. http://www.thegrid.net/ (which redirects to http://hometown.onemain.com/weblogic/TownSquare.jsp)
444. http://www.xmission.com/
445. http://www.ntis.gov/
446. http://www.snu.ac.kr/

Discussion of ASN Data

In reviewing the ASN data just presented, there are some points that you should note.

1) Commodity transit traffic will changes in the weeks and months ahead:

— Peering is currently in a state of flux, and addition of new peers at the OIX will dramatically change what ASNs we see via our commodity transit links.

— Internet2 participation is currently in a state of flux, and when that is completed it too will have an affect on what ASNs we see via our commodity transit links.

We should also note that a variety of I2-connected destinations are showing up via commodity transit connectivity, including NYSERNet, Arizona Tri University Networks, Berkeley and others; we believe those I2 connected sites are showing up because one (or more) of the following holds:

— The ASN advertises only some, but not all of their address space via I2 (for example, SUNY Buffalo, part of NYSERNet, doesn't advertise its entire network block to I2). The non-HPC-advertised portions of those sites will obviously be reached, then, via the commodity Internet (even though all traffic from an I2 site to an I2 site meets the I2 acceptable use policy).

— A non-I2 OWEN/NERO partner is accessing the I2 connected ASN, which means by definition their traffic cannot flow via I2.

— Some I2-eligible addresses may be getting spoofed as part of attacks

— Internet2 routing may be broken at selected locations. See, for example, Hank Nussbacher's paper, "The Asymmetry of Internet-2"[447] reporting results obtained by IUCC in conjunction with the University of Oregon.

— Some Oregon domains are showing up as coming in via commodity transit rather than via peerage at a local exchange point.

This includes, for example, US West, Eastern Oregon Net, Data Research Group, Easystreet,

447. http://www.internet-2.org.il/i2-asymmetry/index.htm

Transport Logic, Clippernet, and others. These Oregon companies should be strongly encouraged to peer with OWEN/NERO rather than exchanging traffic over commodity transit links.

— A number of networks are actively engaged in merger and acquisitions, consolidations, or other business reorganization activities. As a result of that activity, at least some ASNs should probably be combined for the purposes of recognizing aggregation that is occuring. Other ASNs have names which no longer correctly describe the organization the ASN represents.

2) Major sources of inbound traffic today aren't what they used to be.

That is, in the past, major sources of inbound commodity transit traffic were typically major networks aggregating traffic for lots of smaller customers, and there is still some of that happening today.

Now, however, we believe that major sources of traffic also include:

— Web hosting companies which run colo farms (Exodus, Abovenet, etc.). This is not unexpected since companies which are deliverying large amounts of traffic often favor colo farms offering discounted bandwidth costs as well as professionally managed server space.

— Broadband access companies (xDSL, cable modem companies, etc.). Again, this is not surprising since we know that many xDSL and cable modem users are taking advantage of their connectivity to run servers (whether or not those servers are compatible with the broadband access company's AUP.

— The new mega ISPs (e.g., AOL, Microsoft, etc.). Again, this is not unexpected since these operations are simply huge, and even a tiny amount of traffic from a huge number of customers eventually results in material aggregated traffic.

— We are also seeing a large number of flows from advertising and web profiling-related companies such as IMGIS, Flycast, DoubleClick, Web Side Story, etc.

Users who wish to reduce the quantity of banner ads and cookies they see may want to consider using a product such as the Internet JunkBuster.[448]

448. http://www.internet.junkbuster.com/

— We were also surprised at the number of flows which were associated with N2H2 (the makers of N2H2 Network filter and the Bess filtering service, and the operator of "Bess, The Internet Retriever" search engine). Literally 99.5% of the N2H2 flows were associated with OPEN's destination ASs, 92.2% of flows we from srcprt 80 (e.g., web server traffic),[449] and 79.7 % were associated with a single IP address, 216.32.10.110. We do not believe that this traffic is driven by "Bess, The Internet Retriever" (which runs at 206.129.0.160); although we are unable to resolve a host name or connect to that IP address, we suspect that this traffic is Internet filtering related. If so, this is interesting because Bess has been the subject of sharp criticism[450] yet may perhaps be the most widely used Internet filtering technology.[451]

449. Other ports seen were 8080 (2.2%), 53 (1.7%) and 9018 (1.3%). All other ports were < 1%.
450. http://www.peacefire.org/censorware/BESS/
451. http://www.n2h2.com/solutions/school_products.html

Section 13. Analysis of Web Flows

Section 13 Keypoints

✔ If we narrow our focus to just web sites, a variety of national rankings are available for the "N most popular" web sites, however there is only rough agreement between those rankings for a variety of methodological reasons.

✔ If we want to compare those national rankings to the web sites that are most popular among OWEN/NERO users, we first must pre-process the raw OWEN/NERO flow data according to a multi-step procedure (described in detail in the body of the report).

✔ Related to that processing, there are some important notes, including the fact that it is not always possible to map a dotted quad source address to a web site.

✔ Similarly, we adopted a one tenth of one percent cutoff for individual addresses -- individual dotted quad addresses that didn't account for at least a tenth of a percent were excluded (but obviously, that means that sites with a bunch of related addresses, all individually less than 0.1%, would get overlooked).

✔ We also freely acknowledge that our content categories are arbitrary, and in many cases a particular site might arguably be categorizable in more than one way.

✔ For a variety of reasons relating to a) web caching, b) content providers contracting with Akamai to deliver all or part of their pages, c) exclusion of traffic flowing over the OIX or Internet2, and d) behavioral factors related to the time the sample was drawn, we should also caution that there may be material sites or categories of content not reflected in our OWEN/NERO web traffic categorization. For example, note that Verio, an OIX partner, is the world's largest web hosting company, but their traffic doesn't go over OWEN/NERO's commodity transit links, and hence will not be reflected in our analysis.

✔ Notwithstanding those caveats, **on a per flow basis, the 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 second largest category on a per flow basis, at nearly 9% of all web flows, consisted of distributed content delivery, virtually entirely associated with Akamai. Akamai traffic will soon flow via the OIX, and will no longer require use of OWEN/NERO transit bandwidth.

- **The only other categories accounting for at least 5% of web flows or more were the "mega Internet Service Providers" or default 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 had less than five percent of web flows.**

- On a per flow basis, web flows associated with identifiable adult sites amounted to less than 1%.

- Web flows associated with identifiable hacker/cracker web sites amounted to only a tenth of one percent.

- Of all sampled web flows, 58.4% of them were classifiable and are reflected in the categories reported in our report; the remainder consists largely of dotted quad addresses that individually accounted for less than 0.1% of all flows.

- **On a per octet basis, the largest category of web traffic was distributed content delivery (e.g., Akamai) at just over 7%. (Recall that Akamai traffic will soon be eliminated from OWEN/NERO transit usage).**

- **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%.

- On a per octet basis, web traffic associated with identifiable hacker/cracker sites amounted to only a tenth of a percent.

- Of all sampled web traffic, 59% was classifiable on a per-octet basis; the remainder consists largely of dotted quad addresses that individually account for less than 0.1% of all octets.

Let's come back to something a little more basic: what web sites are people going to?

Given that web traffic accounts for nearly 73% of all flows, it is natural to want to know what's being accessed on the world wide web.

On a global basis, you may be familiar with various web pages that offer rankings of the "n Most Popular Web Sites," such as the Alexa 1000[452] which attempt to list the 1000 most popular web sites. For example, for March 2000, Alexa's top ten sites are:

1. msn.com
2. yahoo.com
3. ebay.com
4. aol.com
5. excite.com
6. microsoft.com
7. altavista.com
8. go.com
9. geocities.com
10. yahoo.co.jp

The inclusion of yahoo.co.jp in that list illustrates an important issue — yahoo.co.jp is the Japanese version of Yahoo, and presumably seldom used by most OWEN/NERO customers (since most of us do not read Japanese, nor even have computers equipped to display Japanese fonts). The inclusion of such globally important (but locally irrelevant) sites points out that global rankings really need to be viewed cautiously when it comes to using them as a guide to local web usage trends. But what about other lists?

Another available web ranking comes from 100Hot.[453] Their top 10 sites are:

1. Yahoo and Four11
2. Microsoft Corp., MSN.com and LinkExchange
3. AOL.com and Netscape
4. Lycos Search Engine, Point and WhoWhere
5. Excite, Magellan, City.Net and WebCrawler
6. Altavista Search Engine, Compaq, and Tandem
7. Go.com World Network
8. Quote.com
9. Xoom
10. Amazon

---

452. http://www.alexa.com/
453. http://www.100hot.com/directory/100hot/ (excludes Go2Net, Metacrawler, Dogpile, 100hot from ranking consideration; also uses a non-proportional sampling methodology and excludes gifs, jpegs and frames at this time)

In this case, note that each ranking is actually a collection of related sites rather than a single domain name. While that correctly recognizes the intertwined nature of many web sites, it does make it harder to pigeonhole the type of content that's being accessed.

MediaMetrix[454] quotes the top ten sites (based on unique visitors for March 2000, and ignoring repeat visits from the same user):

1. AOL Network (Proprietary & WWW)
2. Yahoo Sites
3. Microsoft Sites
4. Lycos
5. Excite@Home
6. Go Network
7. NBC Internet
8. Amazon
9. Time Warner Online
10. Real.com Network

Nielsen/NetRatings[455] ranks the Top 10 Web Properties for March 2000 on a per household basis as:

1. AOL Websites
2. Yahoo!
3. MSN
4. Lycos Network
5. Excite@Home
6. GO Network
7. Microsoft
8. NBC Internet
9. Time Warner
10. AltaVista

Internetweek magazine[456] ranks the top ten web sites for 3/2000 (based on 120,000 home Internet users) as:

1. Yahoo
2. AOL
3. Geocities
4. MSN
5. Go.com
6. Lycos
7. Passport (e.g., Microsoft's Hotmail)
8. Angelfire
9. Microsoft
10. Netscape

454. http://www.mediametrix.com/TopRankings/TopRankings.html
455. http://209.249.142.16/nnpm/owa/NRpublicreports.toppropertiesmonthly
456. Internetweek, May 1, 2000, pp. 74.

PC Dataonline Reports[457] ranks the top ten web sites for April 2000 as:

1. yahoo.com
2. aol.com
3. msn.com
4. geocities.com
5. microsoft.com
6. AOL Proprietary.aol
7. passport.com
8. lycos.com
9. angelfire.com
10. netscape.com

There are some common themes there, but clearly no complete agreement.

What web sites (or categories of web sites) are most popular for our users?

457. http://www.pcdataonline.com/reports/topmonthlyfree.asp

Compiling Our Web Statistics

The summaries presented in this section were compiled by doing the following:

1) First, we selected only flows that had been tagged as http. Our http-related flows include both flows from a remote web server and flows to a local web server (and note that since we focus on the remote end of the connection, this section will not provide any indication of the most popular OWEN/NERO web servers).

2) Next, looking only at the flows from step 1, we selected only the individual dotted quads that accounted for at least a tenth of a percent of our traffic.

3) We then attempted to do dotted-quad-to-FQDN domain name reverse lookups on those addresses.

4) Next, we reduced each resolved FQDN down to just two levels — that is, if a site was abc.xyz.aol.com, we rewrote it so it became just *.aol.com

5) We then summed up all entries for each of those two-level domains.

6) We then assigned each summarized two-level domain name to a content category.

7) We then tackled the dotted quads that didn't resolve. Because these sites should be[458] web-related, in many cases it is possible to identify what a site is about simply by putting the dotted quad into a web browser and looking to see what popped up. In some cases, though, that strategy doesn't work, in which case we used ipw to attempt to determine ownership of the network address. When we could identify a dotted quad this way, those data were then slotted in with the summarized two-level domain names.

8) Finally, we totalled up the traffic associated with each category, and arranged the categories in descending order by flow count or octets, respectively.

---

458. To understand why some remote sites (srcaddrs) won't be associated with web servers, remember that we are looking at flows from remote web servers, and flows to local web servers — flows to local web servers are almost certainly FROM non-web servers (e.g., individual users, provider cache boxes, etc.)

Interpreting the Web Statistics

There are a number of important interpretive notes that you should review before looking at our web statistics summaries:

— Content categories

We freely admit that our content categories are purely arbitrary, and yes, some web sites could easily be put into more than one category. You should feel free to re-categorize the sites into whatever categories make sense for you.

An excellent example of a tough categorization decision is the non-resolvable dotted quad 199.172.146.52 — if you actually go to that address in your web browser, you find that it is the "My Excite Start Page." Looking at that page, several possible categorization options for that page come to mind, including:

— Search engine (that's what excite.com is best known for, and if the dotted quad had successfully resolved, that's where we'd have automatically put it; search functionality is also the centerpiece of the default page at that site).

— Portal (clearly that's what they intend the My Exite Start Page to be)

— Free email (a prominently featured and important bit of functionality)

— News (another prominent portal constituent component)

— Financial services (site includes stock tracking and financial news)

— Sports (another area that is included)

— Online Consumer E Commerce (you can click on links to buy stuff)

and there are many more constituent components which, if they were the only thing present, would result in "this" particular flow being associated with "that" particular category — and we have no way of knowing from outside their relative importance.

In our case, however, for sites of this sort, we endeavored to treat the sites as we would have had the dotted quad correctly resolved to the site name. Thus, in this example, the My Excite Start Page traffic is slotted under “Search engines” rather than under “Portals” or one of the other possible categories.

It is crucially important to recognize that these categorizations are arbitrary, imprecise at best, and are just meant to help render the mass of data these sites represent somewhat more comprehensible.

— The “sweeping up crumbs” effect

Another important thing to note: had we been able to resolve all the dotted quads associated with the over 1.7 million flows that were in this sample, it is quite probable that a number of individually-less-than-0.1% dotted quads would, when consolidated into two level domains, accumulate sufficient traffic to meet our 0.1% cutoff. Unfortunately, doing even tens of thousands of domain name lookups takes a phenomenal amount of time;[459] and doing hundreds of thousands or millions of domain name lookups simply isn’t feasible.

There is also the pragmatic issue that SAS, the statistical package used to process this data, has a hard limit on the number of dotted quad to FQDN mappings which can be accommodated, thereby precluding us from mapping an abritrarily huge number of dotted quads.

— Unique page impressions vs. multiple page elements/flows per page

We should also stress the fact that many web pages are comprised of numerous small graphical elements, each of which may show up as a separate flow. Because we count flows, not unique page views, sites that use lots of small graphical elements may rank disproportionately high relative to a (comparably graphical) site constructed using only a few large graphics.[460]

459. To understand why, note that while most DNS lookups are very fast, DNS lookups for sites that do NOT correctly resolve can take many seconds, if not minutes, to time out.

460. An easy way to get a count of the number of elements in a given web page is to use Bobby (http://www.cast.org/bobby/) -- besides checking for accessibility issues, Bobby will also give you a summary of page elements and the size of each of those elements.

— Web cache-related effects

We do not — we technically cannot — consider the true level of interest associated with pages retrieved via a local web cache. A single hit on a popular web page, via a web cache, might satisfy ten or a hundred users who'd otherwise go out and retrieve that page directly.

— Akamaization effects

Another factor to keep in mind is that many of the most popular web sites are now delivering much or some of their web content via Akamai's distributed content delivery service — those pages "count" as Akamai pages, not as Yahoo pages or Hotmail pages, for example, for the purposes of our analysis.

— Commodity-transit-only effects

Similarly, because we are only looking at web traffic that took place over our commodity Internet transit links, web traffic that flowed via Internet2 or web traffic that went via the Oregon Internet Exchange, or web traffic that was internal to OWEN/NERO will not appear. Because at least one OIX peer has targeted web hosting as a core business (Verio is the #1 web hosting company in the world), this will also have a material impact.

— Peak usage time effects

We also are looking at content retrieved during "peak usage" time, not content retrieved in the evenings, or on weekends, or during vacation periods — that traffic would no doubt look quite different.

Similarly, we believe that content popularity will vary with a host of factors ranging from the occurrence of holidays (near Mother's Day, one might expect to see an upsurge in traffic to flower stores on the web, for example), to press coverage of breaking news and events, to factors as mundane as weather (when the weather is good, we expect there to be less interest in checking out weather web sites, and more interest in actually getting outside to enjoy a nice day).

All of these factors should make you interpret the following results carefully.

Keeping in mind those disclaimers, let's now look at web flow counts.

Web Traffic Flow Counts By Category

1. Web Advertising/Online Profiling (13.3%)

SRCADDR	Frequency	Percent
*.doubleclick.net	54347	4.2
*.adforce.com	16459	1.3
*.flycast.com	8792	0.7
*.hitbox.com	4398	0.3
*.valueclick.com	4005	0.3
*.adknowledge.com	3058	0.2
*.mediaplex.com	2305	0.2
*.burstnet.com	1500	0.1
*.extreme-dm.com	1233	0.1
*.admaximize.com	757	0.1
*.datais.com	753	0.1
216.111.248.10	9175	0.7
[the above dotted quad is] adforce.com		
216.34.88.200	8139	0.6
216.34.88.240	713	0.1
[the above dotted quads are] avenue a		
204.71.191.220	3898	0.3
204.71.191.251	2856	0.2
206.132.79.68	839	0.1
206.132.79.21	779	0.1
linkexchange		
199.172.144.24	3178	0.2
199.172.144.25	2431	0.2
206.41.20.6	744	0.1
208.178.169.6	696	0.1
mathlogic		
209.67.38.105	2416	0.2
205.138.3.202	2277	0.2
205.138.3.162	1810	0.1
205.138.3.102	1700	0.1
205.138.3.182	1224	0.1
205.138.3.82	1136	0.1
205.138.3.142	1022	0.1
205.138.3.42	916	0.1
205.138.3.62	743	0.1
208.32.211.215	2082	0.2
208.32.211.230	2067	0.2
208.32.211.200	1596	0.1
209.67.38.103	1956	0.2
209.67.38.104	1592	0.1
209.67.38.106	1541	0.1
209.67.38.101	1484	0.1
209.67.38.102	1446	0.1
204.178.112.180	1367	0.1
204.178.112.170	1334	0.1
doubleclick		

216.35.185.140	1188	0.1
209.1.218.220 admonitor.net	1035	0.1
216.34.56.10 comscore	708	0.1
216.35.211.245 Teknosurf	666	0.1
204.176.36.72 spinbox.net	658	0.1
216.35.210.89 realmedia ad network	656	0.1

2. Distributed Content Delivery (8.6%)

SRCADDR	Frequency	Percent
*.akamaitechnologies.com	108860	8.5
*.digisle.net	1580	0.1

3. Mega ISP/Default Portal (7.8%)

SRCADDR	Frequency	Percent
*.aol.com	46103	3.6
*.netscape.com	26579	2.1
*.msn.com	13905	1.1
207.46.208.196	2810	0.2
207.46.208.197	2597	0.2
207.46.208.198	2403	0.2
207.46.133.14 msft.net	1310	0.1
*.msimg.com	2832	0.2
*.passportimages.com [microsoft]	855	0.1

4. Search/Web Directory (7.3%)

SRCADDR	Frequency	Percent
*.yahoo.com	46129	3.6
*.excite.com	12608	1.0
*.go2net.com	4740	0.4
*.infoseek.com	3129	0.2
*.askjeeves.com	3064	0.2
*.snap.com	2111	0.2
*.googlebot.com	1821	0.1
*.lycos.com	1760	0.1
*.google.com	1738	0.1
*.looksmart.com	1713	0.1
*.hotbot.com	1650	0.1
*.about.com	1563	0.1
*.infospace.com	1302	0.1
*.alexa.com	1217	0.1
*.citysearch.com	737	0.1
*.tmcs.net	721	0.1
209.185.108.203	2941	0.2
Google		
199.172.146.52	1520	0.1
199.172.146.210	764	0.1
excite.com		
206.132.152.250	1251	0.1
goto.com		
216.32.10.26	738	0.1
searchopolis.com		
216.200.22.192	731	0.1
yahoo.com		

5. Free Email (4.0%)

SRCADDR	Frequency	Percent
*.hotmail.com	46984	3.7
*.usa.net	1572	0.1
*.mail.com	1391	0.1
204.68.24.107	953	0.1
netaddress.com (usa.net email)		

6. News/Publishing Companies Online (2.9%)

SRCADDR	Frequency	Percent
*.cnn.com	8373	0.7
*.advance.net	6568	0.5
*.uproar.com	3665	0.3
*.zdnet.com	2623	0.2
*.cnet.com	2295	0.2
*.msnbc.com	1699	0.1
*.newsdigital.net	1314	0.1
*.entrypoint.com	1105	0.1
*.pathfinder.com	884	0.1
*.pbs.org	874	0.1
*.registerguard.com	796	0.1
*.discovery.com	760	0.1
*.kgw.com	685	0.1
216.33.87.17 usatoday.com	935	0.1
209.185.191.239 offspring magazine	899	0.1

7. Free Web Pages (2.1%)

SRCADDR	Frequency	Percent
*.geocities.com	14546	1.1
*.angelfire.com	5211	0.4
*.tripod.com	3277	0.3
*.xoom.com	1775	0.1
*.homestead.com	1191	0.1
*.theglobe.com	771	0.1

8. Streaming Media/Online Music/Television/Video/Games/Entertainment (1.9%)

SRCADDR	Frequency	Percent
*.real.com	4161	0.3
*.mtv.com	2591	0.2
*.spinner.com	2090	0.2
*.scour.net	1975	0.2
*.sony.com	1517	0.1
*.ign.com	1123	0.1
*.sandbox.com	845	0.1
*.mp3.com	797	0.1
208.218.3.6 disney.com	1103	0.1
208.45.172.106 koin.com (tv)	955	0.1
207.16.139.50 abc.com	900	0.1
216.68.76.131 z100portland.com (radio)	741	0.1
205.229.74.190 ugo.com	738	0.1
216.246.6.3 foxworld	659	0.1

9. Education/Reference/Governmental (1.2%)

[NOTE: Many higher education/governmental flows will occur via Internet2 and hence will not be reflected in this commodity-transit-only analysis]

SRCADDR	Frequency	Percent
*.thinkquest.org	3413	0.3
*.oclc.org	1986	0.2
*.worldbookonline.com	1370	0.1
*.umi.com	1279	0.1
*.lexis-nexis.com	911	0.1
*.ovid.com	791	0.1
*.berkeley.edu	748	0.1
148.208.100.38 SEIT (Mexico)	1125	0.1
216.35.120.28 africam virtual game reserve	705	0.1

10. Software (1.1%)

SRCADDR	Frequency	Percent
*.eudora.com	7169	0.6
*.macromedia.com	1315	0.1
*.microsoft.com	898	0.1
*.adobe.com	819	0.1
*.bonzi.com	804	0.1
216.35.148.103 macromedia	1217	0.1

11. NSP/ISP/Web Hosting/Broadband (1.1%)

[Note: may reflect flows from remote users to a local web server]

SRCADDR	Frequency	Percent
*.above.net	2240	0.2
*.bbnplanet.com	1643	0.1
*.flyingcroc.net	1352	0.1
*.ihost.com	1051	0.1
*.home.com	899	0.1
*.crl.com	835	0.1
*.uswest.com	786	0.1
216.33.46.173 Reach Communications (logical.net)	1748	0.1
207.138.178.52 globalcrossing	654	0.1
209.58.150.61 shore.net	783	0.1

12. Online Consumer E-Commerce (0.9%)

SRCADDR	Frequency	Percent
*.amazon.com	5393	0.4
*.cdnow.com	960	0.1
*.travelocity.com	862	0.1
*.columbiahouse.com	723	0.1
208.33.218.15 amazon.com	1501	0.1
208.216.181.15 amazon.com	1158	0.1

13. Adult Content (0.8%)

SRCADDR	Frequency	Percent
*.maxhardcore.com	8179	0.6
*.sextracker.com	3156	0.2

14. Portal (0.7%)

SRCADDR	Frequency	Percent
*.collegeclub.com	2128	0.2
*.thirdage.com	871	0.1
*.iwon.com	657	0.1
209.67.39.223 www.bolt.com (portal)	652	0.1
216.35.123.108 snowball (portal)	1651	0.1
209.132.14.123 collegeclub.com	812	0.1

15. Content Filtering (0.7%)

SRCADDR	Frequency	Percent
216.32.10.110 n2h2.com	9622	0.7

16. Computer/Network/Electronics Companies (0.6%)

SRCADDR	Frequency	Percent
*.hp.com	2223	0.2
*.apple.com	1290	0.1
*.cisco.com	762	0.1
*.intel.com	733	0.1
*.x10.com	675	0.1

17. Online Auctions (0.6%)

SRCADDR	Frequency	Percent
*.ebay.com	6600	0.5
*.auctionwatch.com	769	0.1

18. Photo/Graphics/ClipArt/Fonts/Web Stuff (0.5%)

SRCADDR	Frequency	Percent
*.webshots.com	1299	0.1
*.free-graphics.com	1197	0.1
*.mediabuilder.com	826	0.1
207.138.36.163 webshots.com	1106	0.1
204.71.191.241 bcentral.com (microsoft counter site)	861	0.1

19. Financial (0.4%)

SRCADDR	Frequency	Percent
*.cnnfn.com	2317	0.2
*.quicken.com	845	0.1
216.34.178.251 fastweb.com (scholarships, etc.)	737	0.1

20. Unable to categorize (0.3%)

SRCADDR	Frequency	Percent
*.iacnet.com	3152	0.2
*.transoftcorp.com	801	0.1

21. Weather (0.2%)

SRCADDR	Frequency	Percent
*.weather.com	720	0.1
216.34.4.77 weather underground	702	0.1

22. Miscellaneous Corporations (0.2%)

SRCADDR	Frequency	Percent
*.chrysler.com	799	0.1
*.cnf.com	694	0.1

23. Sports (0.2%)

SRCADDR	Frequency	Percent
*.cnnsi.com	2401	0.2

24. Electronic Greeting Cards (0.2%)

SRCADDR	Frequency	Percent
*.bmarts.com	2960	0.2

25. Instant Messaging (0.2%)

SRCADDR	Frequency	Percent
*.icq.com	2628	0.2

26. Finding People (0.1%)

SRCADDR	Frequency	Percent
*.whowhere.com	1494	0.1

27. Job Search (0.1%)

SRCADDR	Frequency	Percent
*.monster.com	1253	0.1

28. Food and Drink (0.1%)

SRCADDR	Frequency	Percent
*.pyramidbrew.com	1010	0.1

29. Internet Phone (0.1%)

SRCADDR	Frequency	Percent
64.14.212.30 dialpad.com (free internet phone)	865	0.1

30. Hacker/Cracker Sites (0.1%)

SRCADDR	Frequency	Percent
*.wwwhack.com	804	0.1

31. File Sharing (0.1%)

SRCADDR	Frequency	Percent
*.bigredh.com (hotline)	773	0.1

Percent of all web related flows allocated via above categories: 58.4%

Web Traffic Octets by Category

1. Distributed Content Delivery (7.1%)

SRCADDR	Frequency	Percent
*.akamaitechnologies.com	7.3471E8	7.0
*.digisle.net	7869937	0.1

2. File Sharing (5.4%)

SRCADDR	Frequency	Percent
*.juston.com	3.2543E8	3.1
*.i-drive.com	1.4986E8	1.4
*.xdrive.com	35860641	0.3
*.freedrive.com	34770312	0.3
*.scour.net	17325773	0.2
*.myplay.com	10911503	0.1

3. Mega ISP/Default Portal (5.2%)

SRCADDR	Frequency	Percent
*.aol.com	1.748E8	1.7
*.netscape.com	1.4845E8	1.4
*.msn.com	1.3024E8	1.2
*.msft.net	72286726	0.7
207.46.133.14 (microsoft)	5761353	0.1
*.msimg.com (microsoft)	5639701	0.1

4. Search/Web Directory (5.0%)

SRCADDR	Frequency	Percent
*.yahoo.com	2.3342E8	2.2
*.excite.com	65469407	0.6
*.infoseek.com	35978549	0.3
*.go2net.com	34186123	0.3
*.askjeeves.com	27110137	0.3
*.about.com	21149094	0.2
*.lycos.com	20610345	0.2
*.snap.com	15773921	0.1
*.hotbot.com	10879300	0.1
206.132.152.250 goto.com	10859571	0.1
*.infospace.com	10773771	0.1
*.looksmart.com	9653579	0.1
*.google.com	8217396	0.1
209.185.108.203 google.com	7903408	0.1
*.citysearch.com	7346264	0.1
199.172.146.52 excite.com	7272490	0.1

5. Streaming Media/Online Music/Television/Video/Games/Entertainment (4.4%)

SRCADDR	Frequency	Percent
*.mp3.com	79239636	0.7
*.spinner.com	59666649	0.6
*.blizzard.com	22634592	0.2
*.real.com	18597604	0.2
208.218.3.6 Disney	17290703	0.2
*.mtv.com	16775679	0.2
*.ibeam.com	15428967	0.1
*.sony.com	13817219	0.1
*.sandbox.com	12590115	0.1
*.mtvn.com	12468707	0.1
*.www.com	11230718	0.1
*.nausicaa.net (anime)	10804800	0.1
206.190.53.195 Yahoo Broadcast Services	10333780	0.1
*.peeps.com (music)	10133026	0.1
*.startrek.com	9823928	0.1
*.warnerbros.com	9165055	0.1
*.arzach.com (comic book)	8284359	0.1
*.kittykatstew.com (band)	5343458	0.1
*.musicmatch.com	7701829	0.1
208.178.163.59 (napster)	7571572	0.1
*.ign.com	7491212	0.1
208.247.156.172 (Six Flags amusement parks)	6499638	0.1
*.entrypoint.com (‘push’ content delivery)	6171143	0.1
*.broadcast.com	5957881	0.1
*.foxkids.com	5890018	0.1
*.seganet.com	5756365	0.1
208.49.53.50 everstream.com	5521150	0.1
*.bigfoot4x4.com (monster truck site)	5406257	0.1
205.188.246.24 spinner.com	5285717	0.1

6. Free Web Pages (4.1%)

SRCADDR	Frequency	Percent
*.homestead.com	1.6736E8	1.6
*.geocities.com	1.6126E8	1.5
*.tripod.com	32829724	0.3
*.angelfire.com	31589941	0.3
*.free.fr	26415766	0.2
*.xoom.com	23405760	0.2
*.fortunecity.com	8752711	0.1
205.134.183.177 justfree.com	6883034	0.1
*.hypermart.net	6687568	0.1
209.185.176.10 tripod	5378773	0.1

7. Free Email (4.1%)

SRCADDR	Frequency	Percent
*.hotmail.com	4.0902E8	3.9
*.mail.com	7325579	0.1
*.onelist.com (mailing lists)	6184861	0.1

8. News/Publishing Companies Online (3.9%)

SRCADDR	Frequency	Percent
*.pcworld.com	75181821	0.7
*.cnn.com	59107000	0.6
*.advance.net	48829218	0.5
*.msnbc.com	18816335	0.2
*.cnet.com	18239717	0.2
*.zdnet.com	11543968	0.1
*.uproar.com	10968374	0.1
*.newsdigital.net	10114945	0.1
216.33.87.17 USA Today	9498282	0.1
*.discovery.com	9358038	0.1
*.usatoday.com	7955745	0.1
*.startribune.com	7207397	0.1
*.pbs.org	6376012	0.1
*.registerguard.com	5509249	0.1
*.kgw.com	7387499	0.1
216.68.77.23 jacor.com (art bell, etc.)	6976134	0.1
216.71.18.229 (Petite Magazine (fashions))	6302002	0.1
209.221.152.202 kgw.com	6243256	0.1
*.lee.net (publisher/media company)	6009397	0.1
208.45.172.106 (channel 6000/koin.com)	5750638	0.1
207.16.139.50 abc.com	5516995	0.1
208.48.26.200 ny times	5443654	0.1

9. NSP/ISP/Web Hosting/Broadband (3.3%)

[Note: may reflect flows from remote users to a local web server]

SRCADDR	Frequency	Percent
*.conxion.com	90625007	0.9
*.concentric.net	49713168	0.5
*.telepac.pt	24673789	0.2
*.level3.net	21264760	0.2
*.above.net	20547663	0.2
*.exodus.net	17500336	0.2
*.ihost.com	13912652	0.1
*.cw.net	10418289	0.1
*.creative-webs.com	9787065	0.1
*.bbnplanet.com	6472356	0.1
*.iuinc.com (hostme.com)	7205850	0.1
*.unet.com.mk ('the first internet provider in Macedonia')	5994684	0.1
205.231.82.39 uunet	5884512	0.1
*.bk.ru	5782805	0.1
*.hopefx.com	5596651	0.1
206.151.164.3 cw.net	5594191	0.1
*.skyweb.net	5499478	0.1

10. Web Advertising/Online Profiling (2.7%)

SRCADDR	Frequency	Percent
*.doubleclick.net	1.3485E8	1.3
*.flycast.com	29136583	0.3
*.adknowledge.com	21151828	0.2
*.valueclick.com	19324279	0.2
199.172.144.25 MathLogic	12513873	0.1
*.hitbox.com	6054134	0.1
*.adforce.com	8557070	0.1
*.mediaplex.com	8059847	0.1
204.178.112.170	6886835	0.1
204.178.112.180 doubleclick.net	6689978	0.1
216.34.88.200 Avenue A	5630459	0.1

11. Software (1.9%)

SRCADDR	Frequency	Percent
*.macromedia.com	59657139	0.6
*.microsoft.com	28581627	0.3
*.allaire.com	21732020	0.2
*.symantec.com	14812145	0.1
*.eudora.com	11075878	0.1
*.novell.com	9406414	0.1
*.neoland.com (shareware)	8748382	0.1
*.adobe.com	5793016	0.1
*.shannontech.com (productivity software for the Mac)	8371576	0.1
*.winsite.com (shareware)	8016071	0.1
*.sierra.com	7398034	0.1

12. Portal (1.8%)

SRCADDR	Frequency	Percent
*.nettaxi.com	1.2796E8	1.2
*.snowball.com	25922460	0.2
*.collegeclub.com	15096154	0.1
*.iwon.com	13226242	0.1
*.bla-bla.com	8980662	0.1
209.132.14.123 collegeclub	5840664	0.1

13. Education/Reference/Governmental (1.7%)

[NOTE: Many higher education/governmental flows will occur via Internet2 and hence will not be reflected in this commodity-transit-only analysis]

SRCADDR	Frequency	Percent
*.berkeley.edu	20172421	0.2
*.oclc.org	15222161	0.1
208.248.180.232 Persian Teaching System	12433665	0.1
*.noruae.net	11901891	0.1
*.umi.com	11012437	0.1
*.worldbookonline.com	9998654	0.1
*.mapquest.com	9142216	0.1
*.nasa.gov	7662136	0.1
*.thinkquest.org	7379493	0.1
*.utexas.edu	6478459	0.1
*.epa.gov	7100804	0.1
*.westlaw.com	7046867	0.1
*.lexis-nexis.com	6507654	0.1
*.odci.gov	5919977	0.1
*.ovid.com	5683055	0.1
*.uibk.ac.at	5439812	0.1

14. Unable to categorize (1.5%)

SRCADDR	Frequency	Percent
*.sbusiness.com	66684732	0.6
*.consumptionjunction.com	36656827	0.3
*.iacnet.com	24819477	0.2
*.in-addr.arpa (didn’t resolve to a symbolic name)	13530468	0.1
*.onsale.com (redirects to egghead.com, and is password protected)	8284694	0.1
*.transpect.net (‘this site has been temporarily taken off-line’)	6807949	0.1
*.transoftcorp.com	6312039	0.1

15. Computer/Network/Electronics Companies (1.1%)

SRCADDR	Frequency	Percent
*.apple.com	29334413	0.3
*.creaf.com	19483071	0.2
*.sun.com	15248023	0.1
*.dell.com	9748582	0.1
*.panasonic.com	6980261	0.1
*.hp.com	5824954	0.1
*.amd.com	7150199	0.1
198.133.17.62 ibm.com	5742409	0.1

16. Online Auctions (1.0%)

SRCADDR	Frequency	Percent
*.ebay.com	99716094	0.9
*.auctionwatch.com	13312508	0.1

17. Online Consumer E-Commerce (0.8%)

SRCADDR	Frequency	Percent
*.amazon.com	18646182	0.2
*.cdnow.com	15458398	0.1
208.202.218.15	10091182	0.1
208.33.218.15	8045625	0.1
208.216.181.15 amazon.com	6504062	0.1
*.travelocity.com	8081312	0.1
*.columbiahouse.com	7342704	0.1

18. Adult Content (0.7%)

SRCADDR	Frequency	Percent
*.sextracker.com	8591551	0.1
*.erosvillage.com	5996062	0.1
*.playboy.com	5730138	0.1
*.maxhardcore.com	5529645	0.1
208.48.35.249 troma.com	8317765	0.1
216.218.222.247 bunkasha.com	6519109	0.1
204.178.96.75 porncity.net	5352342	0.1

19. Photo/Graphics/ClipArt/Fonts/Web Stuff (0.5%)

SRCADDR	Frequency	Percent
*.webshots.com	30463168	0.3
207.138.36.163 webshots.com	11644735	0.1
*.mediabuilder.com	8884134	0.1

20. Sports (0.4%)

SRCADDR	Frequency	Percent
*.cnnsi.com	19645827	0.2
206.79.229.15 PGA Tour (golf)	12214473	0.1
209.67.111.78 rivals.com	9077608	0.1

21. Job Search (0.4%)

SRCADDR	Frequency	Percent
*.itnnet.com	27382761	0.3
*.monster.com	9552510	0.1

22. Financial Services (0.4%)

SRCADDR	Frequency	Percent
*.cnnfn.com	13077399	0.1
*.realtor.com	10378387	0.1
*.quicken.com	6467153	0.1
216.46.236.77 freescholarships.com	6189532	0.1

23. Food and Drink (0.3%)

SRCADDR	Frequency	Percent
162.117.132.154	11338762	0.1
159.164.183.154 candystand.com (nabisco's lifesavers)	8816710	0.1
64.30.22.120 dietsite.com	6780253	0.1

24. Electronic Greeting Cards (0.3%)

SRCADDR	Frequency	Percent
*.bmarts.com	15285042	0.1
*.americangreetings.com	5653959	0.1
*.egreetings.com	6159366	0.1

25. Weather (0.3%)

SRCADDR	Frequency	Percent
*.weather.com	22248520	0.2
206.79.180.105 weather.com	6795796	0.1

26. Wireless Services (0.2%)

SRCADDR	Frequency	Percent
209.67.75.202 proxinet	21047840	0.2

27. Content Filtering (0.1%)

SRCADDR	Frequency	Percent
216.32.10.110 n2h2	6580088	0.1

28. Instant Messaging (0.1%)

SRCADDR	Frequency	Percent
*.icq.com	10507945	0.1

29. Miscellaneous Corporations (0.1%)

SRCADDR	Frequency	Percent
*.chrysler.com	6010781	0.1

30. Online Classified Ads (0.1%)

SRCADDR	Frequency	Percent
*.classifieds2000.com	8514237	0.1

31. Hacker/Cracker Sites (0.1%)

SRCADDR	Frequency	Percent
*.wwwhack.com	8112583	0.1

Percent of all web related octets allocated by above: 59.0%

Section 14. Other Potentially Significant Traffic Categories

Section 14 Keypoints

✔ This section endeavors to describe other applications that may constitute a material amount of traffic, other than the World Wide Web (which was addressed in the preceding section).

✔ Other than web, the other individual application that superficially appears to use a lot of banwidth is Usenet News, however in comparison to expected volumes, we believe that OWEN/NERO's total Usenet News traffic is only some 69% of expected daily volume (the difference believed to be due primarily to two factors: feeds received via Internet2 and via the OIX, and the fact that OWEN/NERO does NOT take a "full" feed (e.g., OWEN/NERO doesn't carry pirated software newsgroups, pirated music, or other newsgroups prima facie violative of federal or state law.

✔ We could also artificially construct a third "composite" application category of online multimedia/file sharing encompassing RealAudio, Napster, and a variety of other applications; that composite category collectively would encompass just over 10% of all octets. Again, while over 10% of all octets may sound like a substantial volume of traffic, this is actually far below what many sites have observed for Napster alone (reports range twenty through over sixty percent of all traffic at some sites).

✔ OWEN/NERO bandwidth usage is lower in the online multimedia/file sharing category than at many sites for a variety of reasons including the fact that Oregon State and some other OWEN/NERO partners have banned Napster outright, and the fact that University of Oregon has a reputation for aggressive prosecution of copyright infringers, having cooperated in the first federal felony criminal prosecution under the No Electronic Theft (NET) Act, as well as for other reasons such as educational campaigns/publicity asking users to restrain themselves.

✔ If we artificially construct a category for online games and chat, that category has only 1.4% of all traffic measured in octets, a generally trivial amount of traffic, we believe.

✔ Despite all our best efforts, a residual category of some 8.4% (by octets) remained uncategorizable by application, a very low value, we believe, as these things go.

The other application, besides the web, that seems to use quite a bit of bandwidth is NNTP — is OWEN/NERO's traffic for that application normal?

We know from other sites that inbound NNTP traffic is currently running about 116 GB/day for a full feed.[461] That daily volume translates to 928 gigabits, or a little under 10.7 Mbps worth of traffic around the clock, assuming a full feed, non-bursty flows, no articles sent more than once, all articles coming in via commodity transit connectivity, etc.[462]

The consortia's measured commodity Internet transit NNTP traffic was 1.6752×10^9 octets for the thirty minute study period. There are 1,440 minutes in a day, which implies that the reported traffic represented 30/1440 or 1/48th of a day. Multiplying the observed 1.6752×10^9 octets * 48 = 8.04096×10^{10} octets or only ~69% of the theoretical volume of 116 GB/day — OWEN/NERO is actually doing LESS inbound commodity transit NNTP-related traffic than most sites at this point in time. How can this be, you ask? A couple of the more important factors include:

— Non-Commodity-Transit News Sources

Some NNTP traffic comes in via Internet2 connectivity, and still other NNTP traffic is coming in via Oregon IX connectivity. To the extent that OWEN/NERO can build a well connected NNTP server mesh in I2 and among our Oregon IX peers, OWEN/NERO can avoid taking NNTP traffic inbound over commodity Internet connectivity.

— The Feed OWEN/NERO Takes Is Actually Not a Full Feed

OWEN/NERO doen't carry (nor feed) all groups. For example, like many academic sites, the consortia does not carry nor feed "warez" groups (pirated software), mp3 groups (replete with pirated music), child pornography (defined to include "teen" binary groups), nor other groups which are *prima facie* violative of federal, state or local law.[463]

461. For examples of total daily volume statistics, see:
http://www.bc.net/news/stats/statsin.20000424.html (116,863,495,339 octets)
http://www.bc.net/news/stats/statsin.20000423.html (116,027,400,905 octets)
http://www.bc.net/news/stats/statsin.20000422.html (115,991,158,048 octets)
Those values are consistent with what is being reported at:
http://dca1-hub1.news.digex.net/stats/
http://feed1.news.rcn.net/local/stats/ci.04-24-2000.html

462. As a check on this, note that Cidera, a satellite-based news provider, states "A full Usenet newsfeed consumes more than 12 megabits a second on average — and that rate increases daily." (http://www.cidera.com/services/usenet_news/index.shtml)

463. For a discussion of the Usenet News Service guidelines applicable to UO reader boxes, see http://cc.uoregon.edu/docs/news_policy.html

Online multimedia/file sharing

In addition to the world wide web and news, we believe that there is one other "composite category" of applications that also merit attention: online multimedia/file sharing.

Online multimedia/file sharing applications include realaudio (4.2%), napster (3.2%), microsoft netshow (0.9%), hotline (0.8%), qt4/rtsp (0.7%), scour (0.2%), and gnutella (0.2%), for a total of 10.2% of all total octets.

Given that some institutions have reported that Napster alone has consumed "20% or more of their bandwidth,"[464] we believe that a value of 3.2% for Napster and 10.2% for all online multimedia and file sharing applications is actually quite modest with respect to the observed popularity of this composite class of applications.

Why is the OWEN/NERO consortia's values for this class of applications lower than might otherwise be expected? We believe there are a number of reasons for the low observed values in the OWEN/NERO data, including:

— Not all schools have seen equal network traffic associated with Napster and related applications; for example, Harvard has stated that it has not seen Napster-related problems.[465]

— Oregon State University (and, we believe, a number of other OWEN/NERO partners) have made the decision to try employing technical means to attempt to ban Napster outright (although we believe that Napster users may have simply segued to other non-blocked file sharing programs)

— The University of Oregon, while it has not banned Napster, has demonstrably low tolerance for copyright infringement, having had the first felony conviction of a student charged with illegally sharing copywritten music under the federal No Electronic Theft (NET) Act.[466]

---

464. See http://news.cnet.com/news/0-1005-200-1527930.html ("[Northwestern University] estimates that Napster traffic was sucking up more than 20 percent of its bandwidth...") or http://www.wired.com/news/culture/0,1284,34382,00.html ("Bucknell University, for example, says Napster is responsible for approximately 40 percent of its network's overall traffic. Indiana University cites even higher figures. 'At one point, 61 percent of our Internet connection was being consumed by the use of Napster [...]'" or http://www.studentadvantage.com/article_story/1,1075,c1-i42-t197-a22762,00.html quoting Yale as having seen Napster traffic ranging from "five to 33 percent of network traffic, depending on time of day."
465. http://www.studentadvantage.lycos.com/lycos/article/0,1534,c1-i42-t197-a22780,00.html
466. http://www.usdoj.gov/criminal/cybercrime/levy2rls.htm

— At some OWEN/NERO partner sites, residence hall networking support staff will counsel residents whose network ports seem to be generating unusually high amounts of traffic, urging them to reduce their usage — whatever it might be — before that traffic becomes a problem.

— Some traffic in this category no doubt flows via other exits, including Internet2.[467] Interestingly enough, Napster is one of the first "connection aware" applications, e.g., one of the first Internet protocols sophisticated enough to prefer high performance connectivity over more expensive and less capable commodity Internet connectivity.

— There may be less traffic of this type during peak weekday periods than during evenings, weekends and other times when leisure activities are more common; we sampled during the middle of the afternoon on a weekday.

— Some sites may fail to distinguish inbound and outbound traffic, and may be reporting total in- and out-bound Napster-related traffic.

— Many OWEN/NERO partner sites have endeavored to educate their users about issues related to online multimedia and file sharing, including issues related to copyright and issues related to network load.[468]

All of these reasons, we believe, help explain why OWEN/NERO's traffic in this composite category may be lower than the traffic that other sites are seeing from Napster alone.

We should also mention that McCreary and Claffy's recent study[469] (based on data from the Ames Internet Exchange) put Napster traffic at ~2.25 — ~4% over the course of their ten month study; that would be entirely consistent with our observed value.

---

467. http://www.time.com/time/digital/daily/0,2822,41844,00.html
468. See, for example: http://cc.uoregon.edu/cnews/spring2000/napster.html
http://www.dailyemerald.com/texis/scripts/vnews/newspaper/+/ART/2000/05/26/392ea6cf7?inarc=1
http://osu.orst.edu/dept/Barometer/2000/winter2000/week5/mon/gettingoutthetruth.html
469. http://www.caida.org/outreach/papers/AIX0005/

What about online games/chat?

We can also create an online games/chat composite category from our data by looking at half-life (0.6%), AOL Instant Messenger (AIM) (0.3%), starsiege tribes (0.3%), ICQ (0.1%), quake/quake2/quakeworld (0.1%) plus we know that there should be some additional traffic associated with IRC (Internet Relay Chat), however it is difficult to tease out of network flow data for a variety of reasons.

If we ignore IRC-related traffic, that leaves us with 1.4% worth of total inbound octets associated with one or another online games or chat. This is consistent with McCreary and Claffy's[470] value (their game related traffic (not including any chat traffic) varied between less than a percent and just under two percent). We believe this level of activity is not high enough to merit concern.

Uncategorizable traffic

Finally, we freely concede that we haven't been able to allocate all the traffic that we saw in the sample we drew -- in our case, 8.4% of all octets were unclassifiable. For comparison, the last publicly available vBNS monthly report had 24% of all octets flowing over the vBNS in an "other" category, and McCreary and Claffy's recent study[471] had 10.9% worth of traffic unallocated.

470. http://www.caida.org/outreach/papers/AIX0005/
471. http://www.caida.org/outreach/papers/AIX0005/

Section 15. Active Measurements

Section 15 Keypoints

- This section describes active network measurement techniques.

- Earlier sections were all based on passively collected data. An alternative approach to examining network performance is to perform active measurements yourself, to conduct network "experiments" or network "tests" and then use those measurements as the basis for your analysis, rather than doing an observational *post hoc* analysis of traffic that happened to be naturally present at a particular time.

- Active measurements normally focus on three characteristics of the network: delay (and the variation therein), packet loss, and delivered bandwidth.

- One active measurement program is NLANR's AMP project, focussed on round trip delay. UO and OSU both participate in (e.g., host measurement points for) the AMP project.

- Another active measurement program is Advanced.Org's Surveyor project, focussed on one way delay measurements and packet loss measurements. UO participates in (hosts a measurement point for) the Surveyor project.

- An example of how to use Surveyor to diagnose network problems between Oregon and New Brunswick is shown to illustrate the utility of this sort of active measurement program.

- Another active measurement program is NIMI, which focusses on IP multicast traffic propagation issues.

Active Measurement Programs: An Overview

The preceding analyses were all based on flow data — observations of traffic between particular sources and particular destinations that happened to be active during the sampling period.

An alternative approach to examining network performance is to perform active measurements yourself, to conduct network "experiments" or network "tests" and then use those measurements as the basis for your analysis, rather than doing an observational *post hoc* analysis of traffic that happened to be flowing at a particular time.

Active measurements normally focus on three characteristics of the network: delay (and the variation therein), packet loss, and delivered bandwidth.

Delay

Some network delay is unavoidable: transmission of packets from one location to another occurs very fast, but is not completely instantaneous. For example, it takes about 72 ms to go from UO to NYU in Manhattan and back; it takes about 600 ms (via satellite) to get to the Federated States of Micronesia in the South Pacific (e.g., traceroute to www.fm).

Network delay is important for several reasons.

Delay needs to be tightly controlled if want to do voice telephony over IP. If you fail to stay within your delay budget, communication quality may become unacceptable.[472]

For another thing, it is very hard to transmit data fast using TCP if you have large network delays due to something known as the "bandwidth delay product." An excellent resource discussing the bandwidth delay product problem can be found at the Pittsburgh Supercomputer Center web site.[473] A number of companies are actively marketing boxes intended to solve this problem; see, for example, Mentat Performance Networking SkyX Gateway[474] or Flash Networks SatBooster.[475]

---

472. http://cc.uoregon.edu/cnews/summer1999/ip_phone.html
473. http://www.psc.edu/networking/perf_tune.html
474. http://www.mentat.com/
475. http://www.flash-networks.com/html/f-satellite.htm

Variation in Delay (Jitter)

Related to packet delay is variation in packet delivery delay, or “jitter.” To understand the concept of jitter, think about a stream of packets being sent in sequence, packet after packet, launched onto the network with metronome-like regularity:

```
packet ... packet ... packet ... packet ... packet ... packet ...
```

If the packets (which had been launched at a constant rate) arrive at their destination still equally spaced, they have no jitter.

If the packets arrive with varying spacing between packets, that is, with some packets “clumpped together” we say they have “jitter:”[476]

```
packet ... packet-packet-packet ........ packet-packet ..........
```

Packet jitter is quite undesirable, and when jitter gets sufficiently large, it can cause degraded audio and video playback.

One solution to the problem of jitter is to use a buffer to smooth out variations in packet delivery rate (for example, that is how Cisco’s IP/TV product[477] insures that it can overcome jitter-related artifacts). Unfortunately, use of buffering is problematic in interactive applications (e.g., video conferencing), because it takes a period to “load” the buffer with content before anything gets transmitted, thereby causing parties to a conversation to “step on” each other if they don’t maintain “CB style” conversational discipline and pace their conversations by saying, “Over...” when they are done talking.

Packet Loss

Packet loss occurs when packets that have been released onto the network don’t get delivered to their destination, due to congestion, network errors, as a normal part of the functioning of certain network protocols (such as RED[478]) or for other reasons.

---

476. http://ns.uoregon.edu/~ursula/results/measuring.html or,
for a thesis on the topic which was done by an employee of the UO Computing Center, see:
http://network-services.uoregon.edu/~ursula/diplom_lt.ps (PostScript format)
477. http://www.cisco.com/iptv/
478. http://www.aciri.org/floyd/red.html

Unfortunately, as is true in the case of large network delays, it is also very hard to have a network that goes fast in the face of significant packet loss. An excellent discussion of this point can be found in Curtis Villamizer's NANOG[479] presentation "TCP Response Under Loss Conditions."[480]

Delivered Bandwidth

A third characteristic that often is measured as part of an active mesurements program is delivered network bandwidth. That is, how much traffic can we cram down the network we've built?

The normal way to test this sort of question is by deploying a pair of traffic generator boxes, using them to generate simulated traffic with known characteristics to load the network. The market standard for this sort of thing is probably the Netcom Systems SmartBits boxes.[481] NERO/OWEN does not currently have any boxes of this sort deployed around the network, thus we cannot report on delivered bandwidth measurements, although we are generally quite confident that we are obtaining the throughput we've architected.[482]

Let us now consider some of the active measurement projects OWEN/NERO participants are involved with.

Please note that these active measurement programs are not (at least at this time) deployed in a way which will let us use them to make useful observations about the state of OWEN/NERO's inbound commodity transit bandwidth, although they do illustrate the sort of program that could be deployed at sites around the commodity internet if sufficient interest and funding were available.

We mention them here to illustrate the sort of things that could be done in the way of an active mesaurement program on the commodity Internet site of our connectivity, and to help illustrate the general quality of the connectivity OWEN/NERO users are getting via Internet2.

---

479. http://www.nanog.org/
480. http://www.academ.com/nanog/feb1997/tcp-loss/index.html
481. http://www.netcomsystems.com/solutions/products/products.html
482. Note that because traffic generators typically are able to saturate a given circuit, you really shouldn't do testing over production networks anyway — this is one case where it is trivially easy to confirm that "measuring something changes it." (In this case, measuring network capacity with a traffic generator eliminates its ability to carry any production traffic!)

NLANR/AMP

The NLANR AMP (Active Measurement Program, see http://amp.nlanr.net/) does round trip time measurements between a set of roughly 100 monitoring boxes, including AMP boxes located at the University of Oregon and Oregon State.

AMP measurement data includes minimum, mean, maximum, standard deviation and packet lossage in percent between each AMP site and each other AMP site. For example, from here at the University of Oregon to all other sites reachable by high performance connectivity,[483] round trip times and packet lossage statistics look like:

UO to ...		Min (ms)	Mean (ms)	Max (ms)	SD (ms)	Loss (%)
1	Alabama	67.00	68.63	83.99	2.18	0.14
2	Alabama Birmingham	65.00	67.23	90.00	2.28	0.07
3	Alabama Huntsville	68.00	70.18	93.00	2.09	0.07
4	Alaska	62.00	63.82	76.00	1.81	0.00
5	Arizona	50.00	51.24	75.00	1.97	0.00
6	Arizona State	28.00	29.57	48.0	2.01	0.00
7	Boston U	76.0	78.38	243.00	4.98	0.14
8	California Berkeley	13.00	15.46	82.00	2.82	0.07
9	California Irvine	19.00	22.18	330.00	11.11	0.00
10	California LA	18.00	20.35	54.00	2.17	0.00
11	California San Diego	21.00	29.90	4197.00	119.26	0.00
12	California Santa Cruz	16.00	18.64	47.00	2.43	0.00
13	CalTech	22.00	23.52	89.00	2.98	0.00
14	Case Western	69.00	75.72	117.00	4.74	0.07
15	Central Florida	77.00	79.01	103.00	2.39	0.07
16	Cincinnati	68.00	76.65	352.00	12.36	0.21
17	Clemson	77.00	79.62	246.00	7.26	0.56
18	Colorado	33.00	35.63	588.00	22.68	0.00
19	Colorado State	34.00	36.85	599.00	20.95	0.00
20	Columbia	70.00	72.46	105.00	2.44	0.14
21	Connecticut	74.00	75.92	122.00	3.08	0.07
22	Cornell	76.00	78.30	113.00	2.26	0.07
23	Dartmouth	83.00	85.35	520.00	12.23	0.14
24	Delaware	73.00	74.29	167.00	3.44	0.14
25	Duke	72.00	73.58	178.00	3.71	0.14
26	Emory	63.00	64.76	119.00	2.69	0.14
27	FNAL	59.00	68.56	133.00	18.25	0.14
28	Florida	73.00	75.06	102.00	2.17	0.14
29	FIU	81.00	82.55	107.00	2.15	0.07
30	Florida State	79.00	81.33	301.00	8.42	0.14
31	George Mason	79.00	90.37	207.00	12.26	0.14
32	Georgetown	76.00	79.82	338.00	9.92	0.14
33	Georgia	66.00	73.89	466.00	19.68	0.07
34	Georgia Tech	62.00	63.45	85.00	1.95	0.14
35	Harvard	80.00	87.09	172.00	7.29	4.10

483. The AMP box is actually located on a subnet of the Oregon Gigapop, and that subnet is only advertised via high performance connectivity routes.

UO to...	Min (ms)	Mean (ms)	Max (ms)	SD (ms)	Loss (%)
36 Hawaii	88.00	90.27	119.00	3.46	0.00
37 Illinois Chicago	55.00	57.25	127.00	2.92	0.14
38 Illinois Urbana	59.00	61.13	103.00	2.55	0.07
39 Indiana	53.00	57.53	99.00	7.13	0.14
40 Iowa	63.00	68.94	119.00	13.13	0.42
41 Iowa State	65.00	67.49	107.00	2.29	10.07
42 Kansas	44.00	47.26	143.00	4.86	0.14
43 Kansas State	45.00	47.26	365.00	9.46	0.14
44 Maryland	76.00	78.65	117.00	4.16	0.14
45 Maryland Balt. County	77.00	79.57	115.00	4.04	0.14
46 Massachusetts	78.00	83.48	166.00	8.74	0.62
47 MIT	80.00	84.62	147.00	6.24	8.33
48 Miami	82.00	84.03	216.00	4.02	0.14
49 Michigan	72.00	73.95	257.00	6.39	0.14
50 Michigan State	74.00	76.31	119.00	2.80	0.14
51 Michigan Tech	67.00	95.01	5405.00	249.79	0.14
52 Mississippi State	73.00	77.10	637.00	18.85	0.14
53 Missouri	69.00	74.43	125.00	12.92	2.92
54 Montana State	37.00	38.93	52.00	1.78	0.07
55 NCAR	33.00	35.18	591.00	20.79	0.00
56 NCSA	60.00	66.13	114.00	12.11	0.28
57 NCSA DC Access	74.00	77.83	126.00	4.42	0.07
58 New Mexico State	102.00	107.94	173.00	11.84	0.97
59 North Carolina	71.00	75.22	103.00	5.03	0.28
60 North Carolina State	70.00	71.24	92.00	1.96	0.14
61 North Dakota State	88.00	90.34	131.00	2.44	0.14
62 Northwestern	56.00	57.37	98.00	2.65	0.14
63 Norwegian Univ S&T	182.00	197.62	284.00	19.78	0.35
64 Oklahoma	49.00	50.81	67.00	1.95	0.28
65 Oklahoma State	51.00	53.12	83.00	2.10	0.35
66 Old Dominion	79.00	81.66	116.00	4.02	0.14
67 Oregon State	2.0	4.36	17.00	1.92	0.00
68 Pennsylvania	71.00	73.69	108.00	2.43	0.14
69 Penn State	65.00	67.92	114.00	2.73	0.14
70 PSC	62.00	64.67	106.00	3.19	0.14
71 Princeton	78.00	80.20	123.00	2.71	0.21
72 Rice	60.00	75.38	382.00	30.13	0.28
73 Rochester	77.00	79.05	157.00	3.06	0.14
74 SDSC Ramona (HWB home)	74.00	94.97	611.00	44.97	0.28
75 SDSC	21.00	23.74	111.00	4.90	0.00
76 SD School of Mines	78.00	80.00	134.00	2.92	0.14
77 South Florida	79.00	80.66	133.00	3.71	0.14
78 Southern Methodist	67.00	70.17	275.00	10.65	27.99
79 SLAC	15.00	17.77	195.00	5.79	0.42
80 Stanford	13.00	15.14	26.00	1.60	47.29
81 STARTap	56.00	57.25	101.00	2.58	0.14
82 SUNY Buffalo	79.00	81.23	115.00	2.37	0.07
83 Tennessee	94.00	96.78	263.00	9.05	0.07
84 Utah	42.00	43.37	57.00	1.76	0.00
85 Vanderbilt	67.00	70.81	147.00	4.34	0.07
86 Virginia	77.00	80.75	131.00	4.41	0.14
87 Virgina Polytechnic	80.00	85.76	139.00	6.76	0.07
88 Washington	24.00	25.59	77.00	2.18	0.00
89 Washington State	40.00	41.50	76.00	2.19	0.07
90 Washington U St Louis	68.00	74.04	1062.00	44.85	10.07
91 Wayne State	73.00	74.43	123.00	2.49	0.07

UO to...		Min (ms)	Mean (ms)	Max (ms)	SD (ms)	Loss (%)
92	West Virginia	71.00	75.33	115.00	7.33	0.21
93	Wisconsin	60.00	61.66	139.00	3.68	0.14
94	Wisconsin Milwaukee	48.00	59.65	198.00	2.72	0.14
95	Wyoming	37.00	39.97	1007.00	36.01	0.00
96	Yale	80.00	83.17	125.00	3.39	0.56

Looking at that tabulated data, it is easy to tell that there aren't any material bandwidth related problems between UO and those sites — everything appears to be connected via fast circuits which have nil packet loss and virtually no jitter. Not surprisingly, all of those destinations are destinations to which we connect via Internet2.

Regretably, a comparable monitoring program isn't yet available on the commodity Internet transit side of things, although obviously there is no reason why such a program couldn't be deployed between state government networks, for example, or between the leading hundred or two hundred K12 school districts in the country, or between city government networks. [This is an obvious measurement project opportunity.]

In addition to making tabular data of that sort available, AMP also provides graphics to help with visualization of the network performance that you're seeing. For example, the UO to Berkeley round trip times for a recent day looked like:

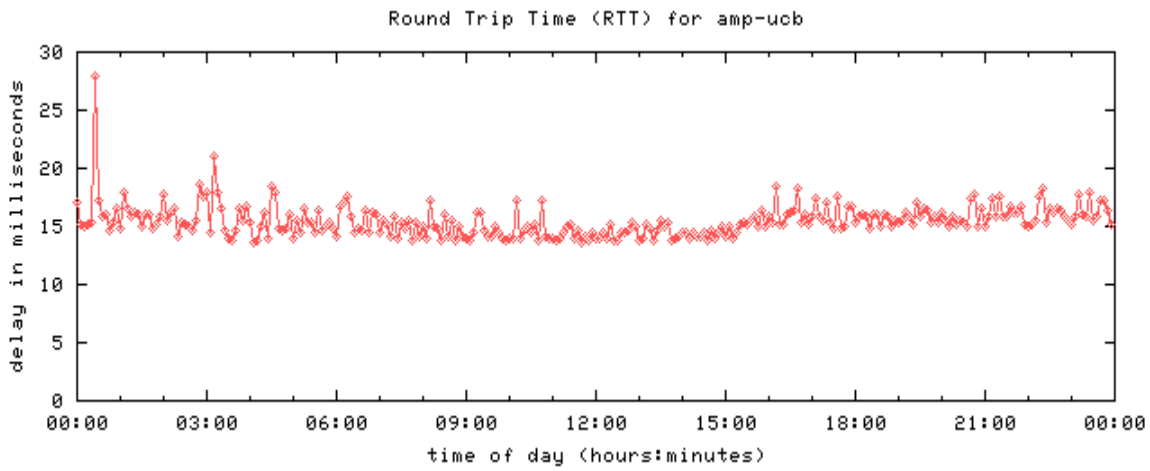

This is excellent performance, very fast, and with very little jitter (except for the one transient spike at the far left edge of the graph, most of the round trip times were within a millisecond or two of each other).

Some might say, “Well, that’s to Berkeley, and your Internet2 connectivity goes to Sacramento and to Denver... one would expect sites close to those connections to have good performance.” What’s nice is that Abilene in fact offers comparably good performance even going to the other end of the country. For example, consider the following graph between Oregon and Florida International University (located in Miami):

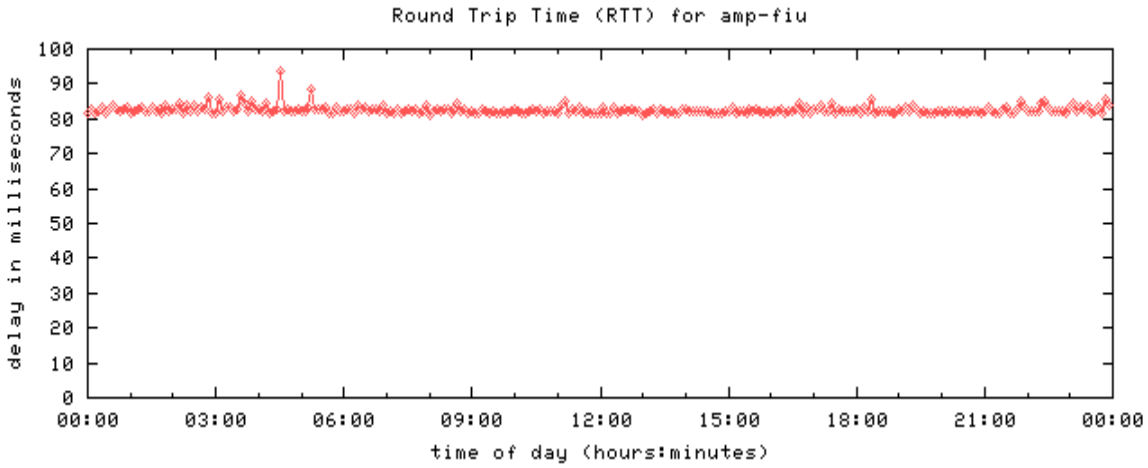

Yes, it does take a little longer to get all the way down to southeast Florida (rather than just down to Sacramento), but 80 msec round trip times are excellent, and there’s very little variation in this graph, just as there was very little variation in the Berkeley graph.

Surveyor

Although NLANR’s AMP is a well established measurement activity, Advanced.Org’s Surveyor[484] project can in some ways result in more directly useful diagnostic data (although the interface is arguably more opaque).[485]

Unlike NLANR’s AMP, which tracks round trip times, Surveyor looks at one way measurement times, relying on GPS time measurements for inter-machine time synchronization.

Having one-way plots can sometimes be very helpful when it comes to isolating the source of a problem. For example, let’s consider a set of eight Surveyor plots:

484. http://www.advanced.org/surveyor/

Plot Number:	From:	To:	Measuring...
1	Oregon	Nova Scotia	Delay
2	Nova Scotia	Oregon	Delay
3	Oregon	New Brunswick	Delay
4	New Brunswick	Oregon	Delay
5	Oregon	Nova Scotia	Packet Loss
6	Nova Scotia	Oregon	Packet Loss
7	Oregon	New Brunswick	Packet Loss
8	New Brunswick	Oregon	Packet Loss

We believe that the Nova Scotia plots show excellent connectivity (little jitter and little packet loss), while the New Brunswick plots show a problem in at least one direction.

Let’s go ahead and look at the plots. (As you look at the plots, you can ignore any gaps you see — those are just times when measurements weren’t taken.)

---

485. To see relevant graphs, go to http://www.advanced.org/surveyor and move down the left hand frame to “Daily Performance Reports.” Click on it.

Select “Calendar” from the View drop down menu in the top frame.

Set either “Source” or “Destination” to be “Oregon Gigapop,” pick a different site for the other end, and then click the “Show Calendar” button.

When the calendar is displayed, pick one of the underlined days (those are the days for which data are available).

When the Daily Reports Display comes up, you’ll see one way packet transmission time between the two sites by default. If you want to select something different, like packet loss, you can chose that in the “Plot Types” framelet in the upper right hand corner.

If you like, you can also look at recorded traceroutes between the two boxes by clicking on one of the “numbered routes” down in the lower right hand framelet.

Plots 1 and 2: Oregon to Nova Scotia and Nova Scotia to Oregon, Delay. (Normal)

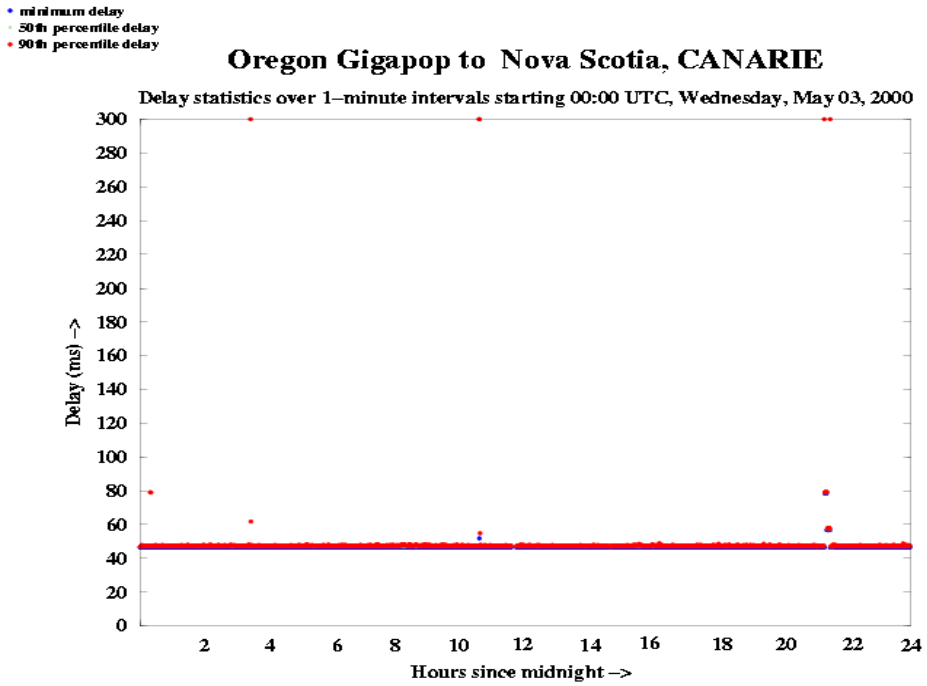

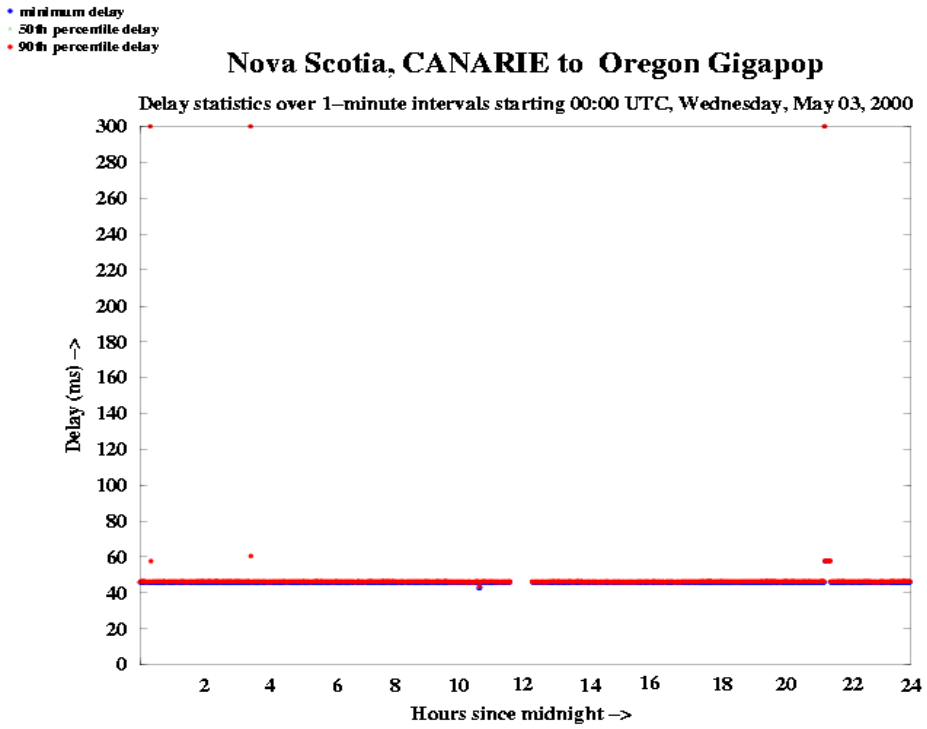

Notice that the above plots are linear... observed delay is low, and quite constant.

Plots 3 and 4: Oregon to New Brunswick and New Brunswick to Oregon, Delay. (Ugh!)

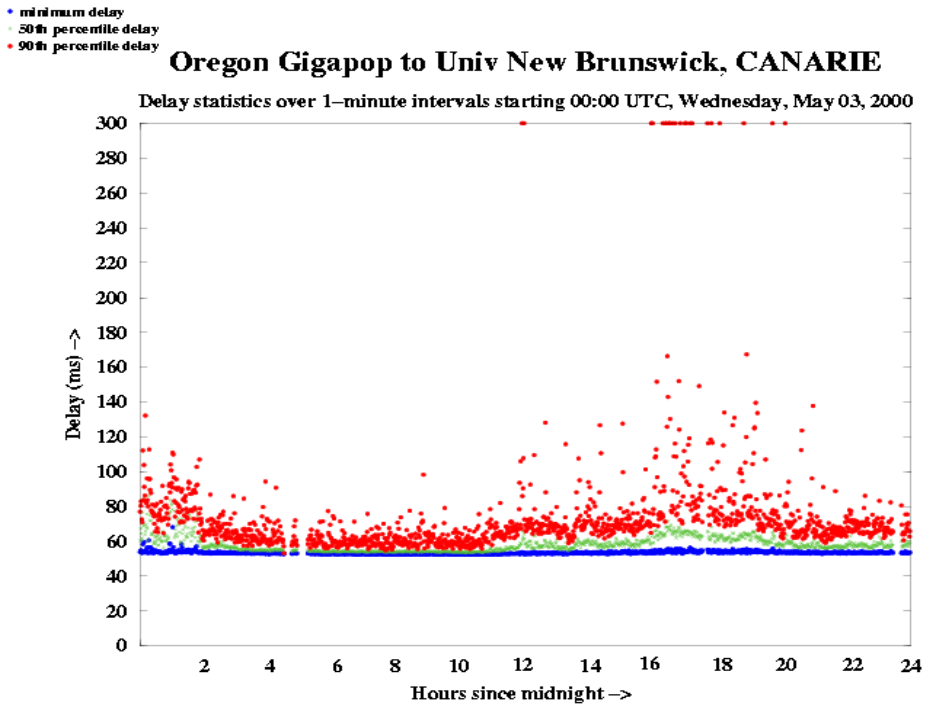

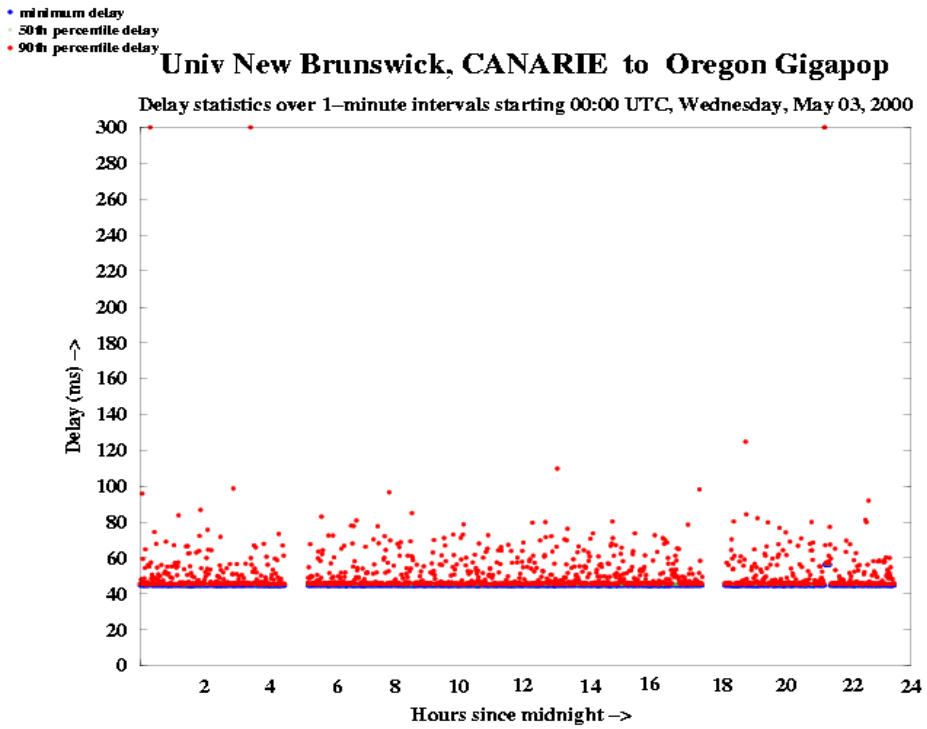

But now look at these plots! The cloud of measurements rising up from the baseline (like a cloud of mosquitos rising up off a swamp), represents jitter and lost/delayed packets....

Plots 5 and 6: Oregon to Nova Scotia and Nova Scotia to Oregon, Loss. (Normal)

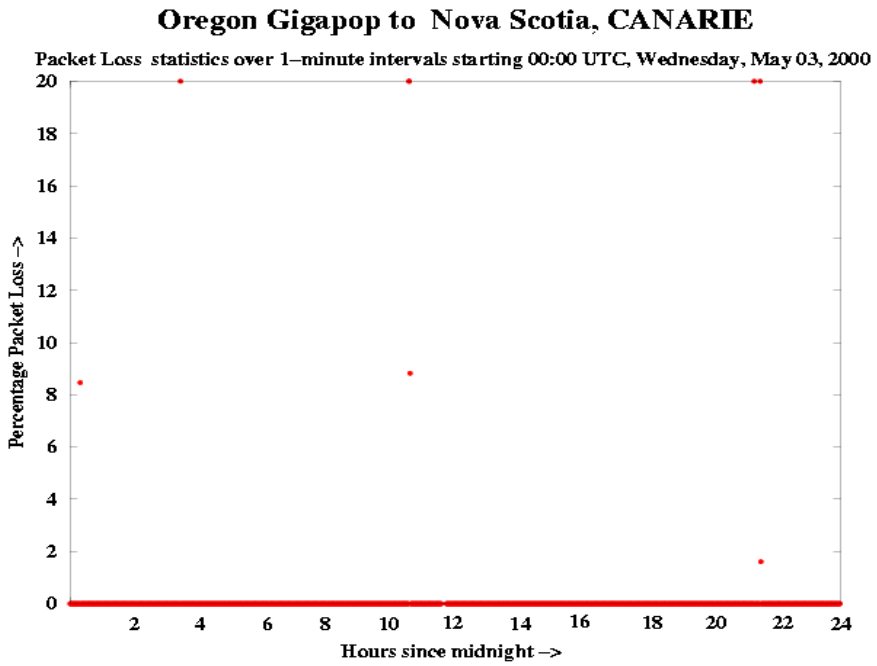

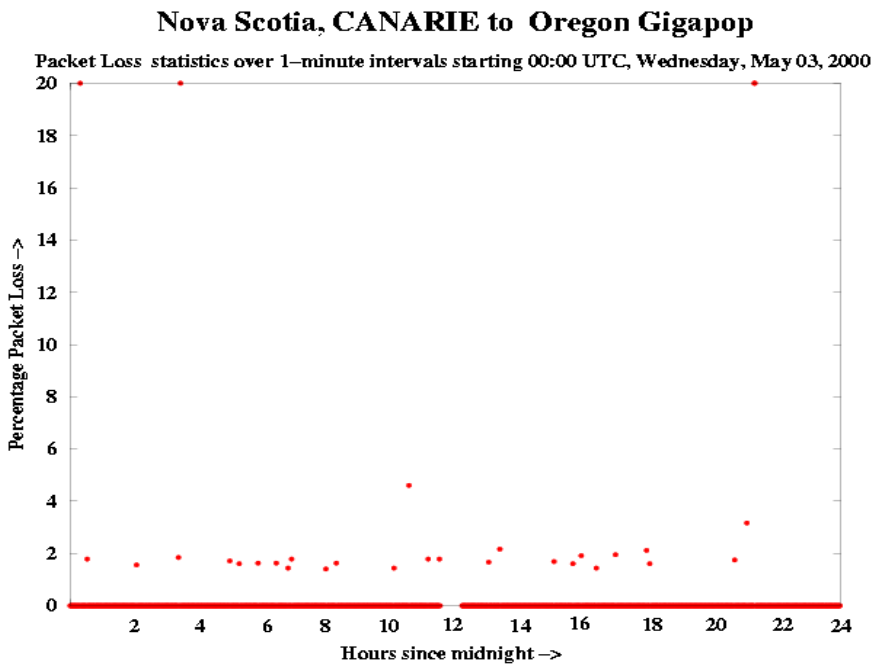

Now we're looking at packet loss. Packet loss is rare between UO and Nova Scotia.

Plots 7 and 8: Oregon to New Brunswick, New Brunswick to Oregon, Loss. (Ugh!)

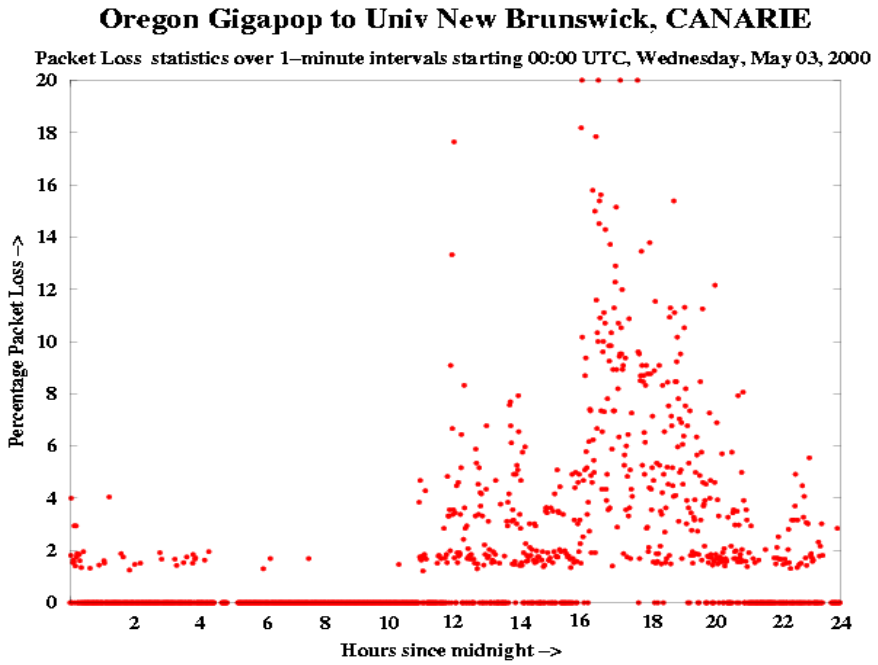

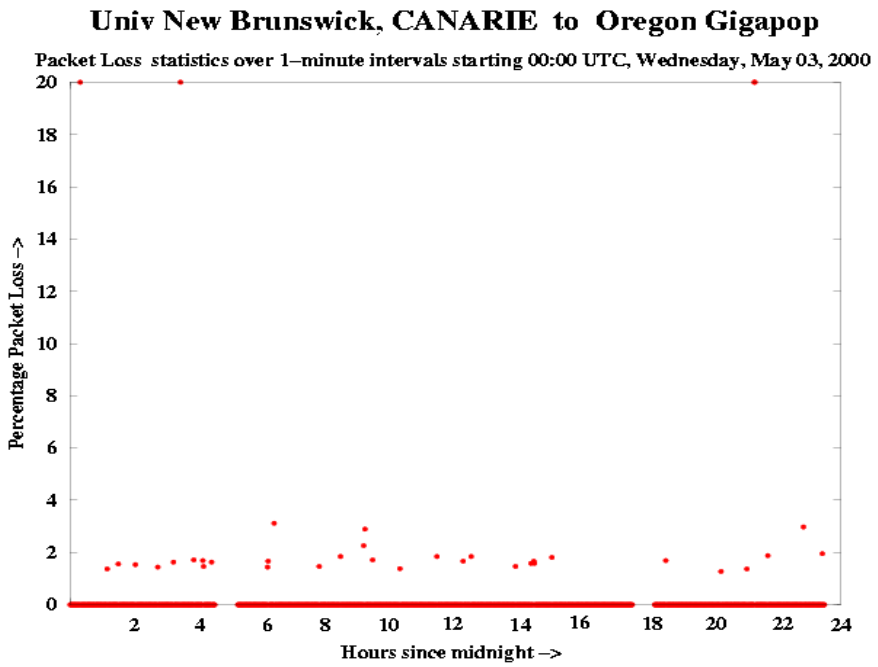

But now look at the loss from UO to New Bruswick! Wow! (The other way looks okay)

We're starting to narrow in on the problem. What does the packet loss from other sites into New Brunswick look like? Let's check Indiana...

Ahah! Indiana is seeing the same sort of packet loss...

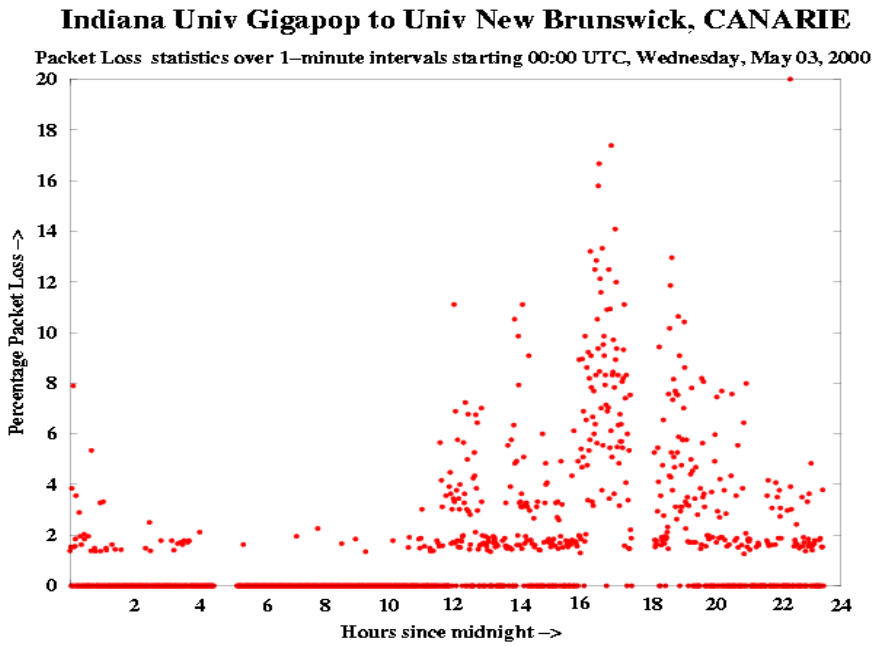

The availability of deployed Surveyor boxes allows us to study this sort of problem from multiple locations, and to isolate the direction in which the problem exists (which can be very useful given that routes out to a site do not necessarily have to match routes back from a site). Let's see if our current problem here is due to some sort of asymmetric routing.

How do packets get from Oregon to New Brunswick? If we check Surveyor's traceroutes for the path, we can see that traffic from Oregon to New Brunswick goes via the commodity Internet! New Brunswick is not advertising (or is not correctly pref'ing) its high performance connectivity route announcements!

Even though New Brunswick has a great Internet2 connection via CANet II, they are telling the world, "Hey, send traffic you've got for me via my [congested, slow] commodity Internet connection!"[486]

486. For information about New Brunswick's Gigapop, see: http://www.ncne.nlanr.net/news/workshop/99607/Talks/kaye/index.htm

Looking at the traceroutes for the other direction, traffic flowing from New Bruswick to Oregon, we can see that Oregon's routes are advertised correctly and flows are going over Internet2 for that part of the trip.

The lesson? Routes may indeed be asymetric, and splitting up the flows can really help isolate and identify this sort of problem.

Oh yes, one additional substantive comment.

One might think that the packet loss shown on the New Brunswick graphs, a few percent ranging up to ten or even twenty percent loss, is "no big deal" — after all the majority of packets are making it through, right? The ones that don't make it through just get resent anyhow, so what's the big deal? Well, it turns out that packet lossage of even a few percent makes it virtually impossible to go very fast. Loss of even a few percent of packets is, or should be, a cause for great concern.

New Brunswick has since fixed this problem (they'd had a router fail on their high performance circuit).

NIMI

We should also mention that unicast active measurement projects do NOT look at IP multicast[487] network performance.

Unlike unicast IP, where packets flow from a single source to a single destination, IP multicast allows multiple receivers to share a single source of packets. For instance, consider Internet video. In the (unicast) streaming "video on demand" model (VOD), ten individual viewers sitting in the same computer lab will each get their own individual video stream when they click on a web video program, even if they are all watching the same video clip.

If they were using IP multicast, as Cisco's IP/TV does, however, a single IP multicast video stream would service all of those users, whether there is one of them, or ten of them, or ten thousand of them. Because of this property, we say that IP multicast "scales to Internet size audiences." For example, UO worked with Cisco to multicast the UN's Netaid Benefit Concert[488] throughout Internet2 in MPEG1 format, and the load was negligible.

So why does everyone know about streaming video on demand, while IP multicast is still a rarity? There are a number of reasons, including:

— In order for you to be able to get IP multicast traffic, your network needs to be "multicast enabled," and so do all the networks that exist between your network and the content that you're interested in viewing. Fortunately, a growing number of networks are now becoming IP multicast enabled by default.

— Multicast has traditionally been a "Unix thing" using programs such as vic, vat and sdr which weren't available for PCs running Windows or for Macs. Again, we are making progress: Cisco's IP/TV is available ONLY for PCs running windows, and the University of Oregon has released a new free tool[489] that is compatible with the higher quality video formats that IP/TV is famous for so that both PC users and Unix users can watch the same content now.

---

487. For a nice tutorial introduction to IP multicast, see: ftp://ftpeng.cisco.com/ipmulticast/multicast_training.html
488. http://www.cisco.com/netaid/
489. mim is available from http://darkwing.uoregon.edu/~htran/projects/mim/

But even as we make steps in the right direction, we know that IP multicast reachability is still spotty at best. Even though UO IP multicast content is routinely at or near the top of the UCSB SDR Global Session Monitoring Effort Report,[490] we know that there are still many locations where IP multicast traffic is not getting through. To help identify IP multicast problem areas, we participate in the NASA JPL NIMI project.[491]

Several times a day, our NIMI box tries to offer IP multicast traffic to all the other NIMI boxes, and all the other NIMI boxes try to offer IP multicast traffic to us.

Sometimes, between some sets of partners, the traffic gets through and sometimes some of it is lost. Sample out from a typical NIMI report is shown below:

```
Date: Wed, 03 May 2000 15:25:10 -0700 (PDT)
From: [nimi multicast reporting daemon]
Subject: 10-min NIMI multicast connectivity test

[snip]

nimi.uoregon.edu (49 pkts):
        changes: +umass +isi-e
        isi-e: delay = 52 sec, loss = 0%, clk = 0.2 secs
        kaist: delay = 15 sec, loss = 0%, clk = -185.9 secs
        lbl: delay = 0 sec, loss = 0%, clk = 0.1 secs
        ucb: delay = 0 sec, loss = 0%, clk = 0.0 secs
        unipi: delay = 0 sec, loss = 0%, clk = 39.0 secs
        sics: delay = 0 sec, loss = 2%, clk = 0.1 secs
        umass: delay = 0 sec, loss = 0%, clk = 106.9 secs
        gatech: delay = 52 sec, loss = 2%, clk = -38.1 secs
        fnal: delay = 0 sec, loss = 0%, clk = 0.0 secs
        luth: delay = 0 sec, loss = 2%, clk = 187.8 secs
        att: delay = 250 sec, loss = 0%, clk = -402.4 secs
        uoregon: delay = 0 sec, loss = 0%, clk = 0.0 secs
        psc1: delay = 0 sec, loss = 0%, clk = 0.0 secs
        nether: delay = 0 sec, loss = 4%, clk = 206.0 secs
        ucsb: delay = 0 sec, loss = 0%, clk = 83.8 secs
        sony: delay = 0 sec, loss = 0%, clk = 83.5 secs
```

490. http://imj.ucsb.edu/sdr-monitor/
491. http://www.ncne.nlanr.net/nimi/

[As part of that message, each of the other peers also report on the IP multicast reachability of each of their NIMI peers, including UO, so we get to see “both directions”]

```
club.bmrc.berkeley.edu (51 pkts):
        changes: +umass +isi-e
        isi-e: delay = 71 sec, loss = 8%, clk = 0.0 secs
        kaist: delay = 71 sec, loss = 8%, clk = -185.9 secs
        lbl: delay = 0 sec, loss = 0%, clk = 0.0 secs
        ucb: delay = 0 sec, loss = 0%, clk = 0.0 secs
        unipi: delay = 0 sec, loss = 4%, clk = 39.1 secs
        sics: delay = 99 sec, loss = 22%, clk = 0.1 secs
        umass: delay = 0 sec, loss = 10%, clk = 106.9 secs
        gatech: delay = 134 sec, loss = 17%, clk = -38.1 secs
        fnal: delay = 0 sec, loss = 0%, clk = 0.0 secs
        luth: delay = 0 sec, loss = 20%, clk = 188.0 secs
        att: delay = 269 sec, loss = 0%, clk = -402.4 secs
        uoregon: delay = 0 sec, loss = 0%, clk = 0.0 secs
        psc1: delay = 0 sec, loss = 10%, clk = 0.0 secs
        nether: delay = 0 sec, loss = 12%, clk = 205.9 secs
        ucsb: delay = 0 sec, loss = 0%, clk = 83.8 secs
        sony: delay = 0 sec, loss = 10%, clk = 84.2 secs
```

If you inspect the first section of the NIMI report shown, you will notice that IP multicast connectivity between UO and other sites is generally excellent, with nil loss.

“Delay” (as used in this report) represents the time which elapsed between creation of the test sessions, and the time they were available at the remote site; some delay is fairly routine in this context.

“Clk” represents the difference in time between the clock on our NIMI box (which is sync’d to a highly accurate GPS-anchored time source) and the clock on the other boxes — obviously some of them are fast or slow.

If you compare our IP multicast connectivity to that of Berkeley’s, you’ll see that our IP multicast connectivity generally has less loss than their’s does, particularly for traffic going to SICS (2% loss vs. 22% loss), UMass (no loss vs. 10% loss), and selected other partners.

Obviously it would be very helpful if more sites would consider running NIMI IP multicast performance monitors.

Section 16. Miscellaneous Issues

Section 16 Keypoints

- This section handles three miscellaneous issues not dealt with elsewhere in the report:

- What (if anything) can be done to use OWEN/NERO's transit bandwidth more efficiently? We believe that peak shaving (shifting demand from peak times to off peak times), promoting voluntary use of web caching by partner sites, and use of satellite based web cache preloading and Usenet news distribution all have some potential for increasing OWEN/NERO's efficiency.

- Is OWEN/NERO's bandwidth demand going to continue to grow without limit? If so, how does OWEN/NERO propose to meet that demand, given limited funding? The answer to this question is that yes, we do believe that OWEN/NERO will continue to see an increase in demand for bandwidth, although it is exceedingly difficult to accurately forecast that demand given discontinuities caused by emerging new network applications and large structural changes such as the deployment of SOEN statewide. We do believe, however, that OWEN/NERO's charge of $1000/Mbps/month will largely underwrite the demand that proves to be required.

- What about OWEN/NERO and its relationship to the new State of Oregon Enterprise Network? OWEN/NERO has been identified by the State as the SOEN network service provider. In turn, OWEN/NERO anticipates buying level 2 circuits and other network services from the SOEN-developed statewide contract schedules.

This section

In this section we handle remaining “open issues” not addressed elsewhere in the report.

Can OWEN/NERO use its inbound transit bandwidth more efficiently?

OWEN/NERO already does a good job of matching its network capacity to its contractually established partner traffic requirements. However, can OWEN/NERO or its partners use its inbound transit bandwidth more efficiently?

We believe that yes, there are some opportunities for improvement in this area. The three strategies that we believe particularly merit consideration are:

— **Promoting off peak usage (“peak shaving”):** recall that OWEN/NERO needs to size its bandwidth to meet observed peak incoming loads, and that peak loads tend to occur during the middle of the afternoon. If some part of that peak load can be shifted to a lower usage period (such as early evening hours), transit bandwidth requirements can be reduced. Two of many ways to promote off peak usage is to offer network dialin access (so people can work from home), or to implement a flex time program (so that employees can start earlier than normal or work later than normal).

— **Encouraging voluntary use of web caching to eliminate redundant retrieval of popular pages over commodity transit links:** given that web traffic dominates OWEN/NERO’s transit bandwidth load, any measures which can be taken to modulate that traffic obviously have a high potential payoff. Encouraging voluntary use of web caching by OWEN/NERO partners is an obvious example. A number of OWEN/NERO partners already offer web caching services for their users, but web caching should be offered by all OWEN/NERO partners. For more information about web caching, see the IRCache web site,[492] or visit the Cache Now! project web site.[493]

— **Purchasing satellite-based web cache preloading and Usenet News distribution services:** a number of satellite-based service providers offer a web cache preloading service and Usenet news distribution service via satellite, including Cidera[494] (formerly SkyCache). Subscription to this sort of service might improve cache hit rates and also offload some inbound bandwidth.

492. http://www.ircache.net/
493. http://www.vancouver-webpages.com/CacheNow/
494. http://www.skycache.com/

Is OWEN/NERO's bandwidth demand going to continue to grow? If so, how does OWEN/NERO propose to meet that demand, given limited funding?

We do indeed believe that bandwidth demand will continue to grow, as it clearly has in the past. It is difficult to accurate forecast future demand since bandwidth demand growth tends to occur in discontinuous steps as new applications are introduced or downstream bottlenecks are eliminated, however we do know that growth is inevitable, and at a rate of speed that's likely to be exponential rather than incremental.

With OWEN/NERO's $1000/Mbps/month charging model, we believe that OWEN/NERO partners now have incentives to intelligently manage their use of bandwidth, and that funding helps OWEN/NERO to pay for the incremental capacity that is needed.

We would also note that while demand for bandwidth is increasing, we believe bandwidth costs have dropped, and will continue to do so, as fiber continues to get deployed throughout the country and around the world, and improvements in long haul fiber optics (particularly in the area of WDM) reach the marketplace.

Perversely, too, the larger OWEN/NERO becomes, the greater the incentive for network service providers to peer with us, thus increasing the consortia's connectivity while not necessarily increasing our costs.

At this point, it is simply impossible to project bandwidth demand or bandwidth pricing with any confidence over a multiyear time period. One source of uncertainty is how deployment of the State of Oregon Enterprise Network (SOEN) will affect demand for bandwidth across the state.

Speaking of SOEN, what is the relationship between it and OWEN/NERO?

OWEN/NERO anticipates that it will be working closely with SOEN in the future.

For example, as of May 11th, 2000, OWEN/NERO has been designated by the State of Oregon Department of Administrative Services as SOEN's network service provider.[495] In turn, OWEN/NERO has indicated its intent to purchase scheduled layer two circuits and various other services from SOEN where those services become available.

---

495. SOEN Transport and Value-Added Baseline Requirements (Phase II) - v20.doc, prepared by Network Evolutions, Inc., dated May 11, 2000, page 30.

Section 17. Findings

Findings

Finding 1: We considered OWEN/NERO's activities relative to its mission. We found that OWEN/NERO is doing what it has been charged with doing, that is, it provides high quality and cost effective aggregated Internet connectivity for Oregon's public universities, Oregon's public K12 schools, and Oregon's state agencies.

Finding 2: OWEN/NERO's connectivity to the Internet occurs via a variety of different paths, the most important of which (and the most expensive of which) is called "commodity Internet transit."

Finding 3: OWEN/NERO's level of commodity Internet transit connectivity is quite low in comparison to most state network consortia. That is, OWEN/NERO is "thinly provisioned" or "tightly engineered," with Internet transit bandwidth capacity maintained at the minimum levels required to meet OWEN/NERO partner contractual demand.

Finding 4: Bandwidth usage by partner is limited by OWEN/NERO (via technical means) to the capacity that each partner has purchased. Ongoing monitoring of those limits indicates that those controls are functioning as intended.

Finding 5: Looking at average OWEN/NERO bandwidth on a per user basis, OWEN/NERO has an average of 195 bits per second per user worth of Internet transit bandwidth. Compared to nominal dialin modem speeds of 56000 bits per second, an average of 195 bits per second per user is obviously a very modest level of transit bandwidth.

Finding 6: OWEN/NERO's commodity Internet transit bandwidth usage (in terms of the network protocols observed, applications used, web sites visited and all other observable macroscopic characteristics), is consistently in line with values previously reported in the network measurement literature.

Finding 7: For OWEN/NERO (as for the Internet as a whole), the World Wide Web continues to be the single most popular network application, both on a per flow (count) and on a per octet (volume) basis.

Finding 8: Excluding advertising-related flows and other incidental/infrastructural categories of web traffic, the most popular web sites for OWEN/NERO customers are generally the same ones that are known to be popular Internet wide (e.g., megaportal sites such as aol.com/netscape.com and msn.com; Internet directory and search sites such as Yahoo, Excite, Infoseek, etc.; free email sites; free web page sites; file sharing sites; news and publishing sites; streaming media sites, etc.).

Finding 9: Free commercial web-based email services (such as Microsoft's Hotmail) continue to be popular with OWEN/NERO users, despite the fact that many sites now offer their own web-accessible email interfaces, in part because of the great flexibility and virtual anonymity such free email accounts provide.

Finding 10: Identifiable adult web content and identifiable hacker/cracker web content was negligible (totalling together less than one percent of all OWEN/NERO web traffic).

Finding 11: Other than the world wide web, the only other individual application that accounts for a material amount of incoming bandwidth is Usenet News, comprising 10% (by volume) of incoming traffic. Comparing that value to known/expected Usenet News traffic volumes, observed OWEN/NERO Usenet News traffic received via commodity transit connectivity was actually only 69% of what would be expected. This is believed to be due to receipt of some Usenet News traffic via Internet2 and other non-commodity Internet connectivity, and due to OWEN/NERO's policy of not carrying newsgroups whose content is prima facie violative of federal law (such as the warez groups which are available at some sites for the purpose of facilitating trafficking in pirated commercial software).

Finding 12: Online MP3-based music sharing applications (such as Napster) have received much press lately (for example, Napster was featured on the cover of the June 5th Newsweek magazine), with some sites making annecdotal reports that Napster and related applications were accounting for phenomenally large fractions of their total Internet bandwidth (some sites have claimed that Napster was accounting for over 60% of all their Internet traffic). For comparison, based on our OWEN/NERO flow data, 3.2% worth of OWEN/NERO's inbound commodity transit traffic was Napster related, a value consistent with Napster traffic levels measured by McCreary and Claffy of CAIDA (the NSF's Cooperative Association for Internet Data Analaysis) at the Ames Internet Exchange (they saw 2.25% to 4% over a ten month period).

Finding 13: We also looked at OWEN/NERO bandwidth usage associated with online games, finding that 1.4% worth of total inbound octets were associated with games and instant messaging applications. This traffic level was consistent with the CAIDA study, which measured game related traffic at levels ranging from less than a percent to just under two percent.

Finding 14: OWEN/NERO obtains significant technical and financial benefits from exchanging customer traffic at no-charge with a variety of network service providers who are at the Oregon Internet Exchange ("OIX"), which is run by the University of Oregon and located in Eugene. The number of networks participating in the OIX continues to grow, thereby increasing its value to OWEN/NERO partners.

Finding 15: OWEN/NERO is making effective use of Internet2 ("I2") as a cost effective and high performance alternative to the commodity Internet to the full extent currently permitted by I2 policies. UO, OSU, PSU, EOU, OIT, SOU, and WOU are all connected and passing traffic on Internet2 at this time. Discussions are taking place within Internet2 at the I2 board level with respect to modifying I2 policies so as to allow connecting the remaining OWEN/NERO partner communities (e.g., OPEN and DAS) to Internet2.

Finding 16: OWEN/NERO partners participate in a wide variety of active network measurement programs designed to quantify network performance, however typically the scope of those active measurement programs is limited to research networks such as Internet2. Examples of active measurement programs include NLANR's AMP, Advanced.Org's Surveyor and NIMI.

Finding 17: OWEN/NERO and its partners may (or may not) be able to use their existing inbound transit capacity more efficiently if they can: (a) shift network traffic from peak periods to off peak periods, for example by offering remote access via modem or by implementing flex time programs for employees; (b) encourage voluntary use of web caching technology to eliminate redundant retrieval of popular web pages over transit links; and/or (c) purchase satellite-based web cache preloading and Usenet distribution services.

Finding 18: OWEN/NERO will virtually certainly be facing a growth in demand for commodity Internet transit bandwidth, although it is not currently possible to forecast the exact demand which will be seen, or the rate at which it will grow. Forecasting the likely increase in bandwidth demand is impossible to do with any accuracy because of unanticipable new network applications, and because of discontinuities associated with major events such as the roll-out of SOEN, the State of Oregon Enterprise Network.

Finding 19: At the same time, however, we do believe that while demand for bandwidth will continue to increase, we believe that the cost of providing that bandwidth will continue to drop due to a variety of national and international fiber projects coming to fruition, and due to widespread deployment of dense wave division multiplexing, a technology which effectively multiplies how much traffic a given piece of fiber optic cable can carry.

Finding 20: The relationship between OWEN/NERO and the State of Oregon Enterprise Network (SOEN) is becoming clear as time goes by. The most recent SOEN Transport and Value-Added Baseline Requirements (Phase II) document dated May 11, 2000, now identifies OWEN/NERO as the Internet service provider for the SOEN project. OWEN/NERO, in turn, anticipates that it will buy various circuits and network services from the SOEN contract.

Finding 21: Specific recommendations relating to campus/WAN networking and future OWEN/NERO operations (e.g., recommendations outside the scope of this report) will be forwarded to the Vice Chancellor for Administration for Oregon University System for his consideration and disposition.

Index