

Provisioning for Peer-to-Peer Applications and Their Impact on Our Campuses

Comments on the topics submitted for the panelists for the Internet2 Spring 2001 Member Meeting, Washington DC

Joe St Sauver, Ph.D. (joe@oregon.uoregon.edu)
Assistant Director, Academic User Services
University of Oregon Computing Center

(1) What techniques/tools have been easily deployed and are used for traffic shaping?

A popular approach is to use CAR. For background info, see:

- http://www.cisco.com/warp/public/cc/pd/iosw/tech/carat_wp.htm
- <http://www.cisco.com/univercd/cc/td/doc/product/software/ios111/cc111/car.htm>
- <http://www.nanog.org/mtg-9811/ppt/witt/sld001.htm>

See also Charley Kline's excellent talk on "Usage-Pattern-Adaptive Rate Limiting" from the Hawaii NLANR/Internet2 Joint Techs Conference:

- http://www.ncne.nlanr.net/training/techs/2001/0128/presentations/200101-kline1_files/v3_document.htm

Or one could use a bandwidth management appliance:

- http://www.allot.com/html/products_netenforcer.shtm
- http://www.intel.com/network/idc/products/bandwidth_management.htm
- <http://www.packeteer.com/products/packetshaper/>

(2) What (or should) cost models be established for charging for bandwidth? what should these be based on?

- Most university-connected users do not have pre-defined bandwidth guarantees... nor predefined bandwidth caps. Why? Casual use by university community members has not (generally) resulted in large volumes of traffic, hence there has been little to motivate hard usage caps. And because most university networks are grossly underprovisioned, bandwidth guarantees are risky endeavors to make (particularly when users traditionally aren't paying anything directly for connectivity).
- To understand what I mean by this, do the math. Assume, for example, that you have 3,000 ethernet connected users and you wanted to guarantee them all 256Kbps worth of bandwidth (e.g., DSL-class capacity) round the clock. If we were to assume that they would ALL use full capacity at the same time, in order to deliver that capacity you'd have to provision $3000 * 0.256\text{Mbps} = 768\text{Mbps}$. Even guaranteeing them all 56Kbps/user (e.g., modem-class capacity) implies buying 168Mbps worth of capacity for those 3,000 users... few (if any) of us are willing to do that. We all rely on the fact that our statistically multiplexed "real life" network load won't require that sort of 1-to-1 buildout.

"Well, what if we just charge them for what they actually *use*?"

- Doing accounting on a per octet or per flow or per user isn't much fun. (The connection-oriented phone company is good at tracking and charging for few penny/minute flows; the packet oriented network world is not.)
- Rebuttable assumption #1: It costs more to track the usage of average users than it would be worth. ("Why don't we track how much water the average person uses on campus?")
- Rebuttable assumption #2: High bandwidth users don't have the inclination or the wherewithal to pay the cost of the bandwidth they're actually using (let's call "high bandwidth" users those users who are averaging at least 1.5Mbps round the clock) -- if those users were willing to pay "full freight," they'd buy their own connection and enjoy full discretion in the use of that bandwidth.

Complicating factors if you try to build a cost model:

- The bandwidth picture for university users differs from the bandwidth picture for traditional commercial customers because universities may have a mixture of expensive connections (commodity transit), comparatively inexpensive connections (HPC connectivity), and "costless" connections (local traffic plus exchange point peering). In the ideal world, usage of inexpensive connections would be lightly throttled (if at all), and usage of costless connections would be unchoked. Unfortunately, it can be difficult to differentiate the different types of traffic in real world environments, and if a site goes fast to one type of pipe, it usually goes fast to all types of connections. See:

<http://darkwing.uoregon.edu/~joe/how-to-go-fast.ppt>

- It is also important to recognize that most universities have asymmetric traffic patterns, with inbound traffic dominating outbound traffic. Universities typically provision capacity to match inbound peaking loads (e.g., inbound peaking loads are the controlling factor determining transit requirements and costs). See, e.g.: <https://web-vms.uoregon.edu/~joe/bw2/owen/index.html>
- Traffic tends to be time phased, with peaks and troughs — will you charge the same amount at all times?
- Allocation of outbound capacity is a management question. What does the institution WANT to do with that capacity? Try to sell it? Use it altruistically to run an anonymous ftp archive? Ignore it? (Beware large and powerful unused network resources -- I guarantee they will end up being misused/abused.)
- What about "accidental" use (or network mischief)? If a user doesn't know that running his streaming multimedia application will result in a big bill, and does so naively, will you "comp" him his error? What if he's the target of a denial of service attack? (Those who remember charged use of time sharing systems can probably remember run away numerical models bankrupting a student's or professor's "funny money" account for the term.)
- What if at least some users (typically students) are already paying a computer usage fee? Should we (DARE we?) "double charge" them by charging them a usage fee in addition to a flat fee? Do we charge ALL users (including faculty and staff and administrators) a comparable fee? (ugh, *not* a popular move...)

(3) What policies are reasonable and how do you manage them?

Without commenting on the reasonableness of any of the following policies, I'll just mention a few policies that some schools are known to have tried (and some of the problems with them):

- do nothing (everything slows down, or you go broke buying commodity transit)
- blocking (or rate limiting) specific ports at the border router (but users will bring the same service up on a different port)
- rate limiting total traffic per user (sort of a shame to do this, however, particularly if you set a relatively low limit such as 256Kbps but would like to encourage MPEG1 IP multicast at 1.5Mbps or H.323 video at 384Kbps, or if you really have bandwidth constraints only during part of the day)
- establishing a “no server” policy (and possibly enforcing it with a network address translation box, periodic NMAP scans, etc.)
- tinkering with DNS to blackhole particular DNS requests (but users can run their own BIND's on Linux boxes, or work by raw IP address in many cases)
- watching traffic levels, and investigating consistently hot users, ports, etc. (sort of a post hoc approach; potentially manpower intensive; not much fun to run around policing bandwidth abusers)
- banning particular applications by name (but users just shift to a new product with comparable functionality that hasn't been banned)
- allowing natural bottlenecks to occur (e.g., running with 10Mbps hubs rather than 100Mbps switches or gigabit switches)
- deploying content delivery network devices (e.g., Akamai, iBeam, etc.) to accelerate at least some content
- capping general Internet traffic but exempt (for example) traffic sent via an institutional web cache
- using passive caching to force all web access through a web cache (but note that if a web cache halves your traffic, that still only buys you maybe six months worth of breathing room)
- substituting cheapo transit providers for first tier transit providers (quality may (or may not) suffer)
- trimming the tail of the distribution (e.g., kicking out the few percent of all users who consume the vast majority of the resources)
- partitioning your bandwidth so that residential networking runs on a separate pipe from the rest of campus (but then you lose the ability to share that resource during off peak times, and routing gets trickier, and costs increase)
- work hard at building out a presence at exchange points so that you can substitute free peering for expensive transit bandwidth (this is the **ONLY** policy that scales in the long run)

(4) What are the students doing and the impact of what we are doing as administrators/etc.?

Some of the applications which we've seen or which have popped up include Napster, Gnutella, Freenet, Scour Exchange, CuteMX, iMesh, HotLine, IRC/DCC, Etree — you name it.

Of these, I believe the biggest emerging potential bandwidth issues may be associated with etree.org, and other private ftp-based file sharing circles.

(5) Have educational campaigns made a difference?

We have noticed a reduction in traffic levels following enforcement activity (e.g., law enforcement seizing equipment and/or felony arrests and convictions), but general requests for moderation (even with justification/explanation) seem to receive less attention and have little practical effect.

(6) Should we provision our campus networks with p2p in mind? Or are we too late for that? Are we still on time?

Absolutely.

Once you grow beyond roughly the OC3 level of commodity transit, you *really* need to be planning how you will execute a strategy that will let you peer with your major traffic sources and sinks. (Unless you can get insanely cheap commodity transit, ala Cogent Communications, Yipes, etc.)

If you fail to execute a strategy that will let you peer, you will either face congestion or unsupportably high ongoing commodity transit costs.

(7) Beyond Napster's recent ruling, how will universities deal with other p2p applications?

See (3) above.

(8) What role can Gigapops and service providers play when addressing these issues?

Internet2 needs to move beyond just a focus on high performance networking via Abilene and vBNS+, and recognize that the greatest challenge higher education faces is in the area of provisioning run of the mill commodity Internet transit.

Internet2 needs to proactively help higher education move toward a peering based model, with a presence at all the major exchange points/NAPs, and fiber (or WDM lambda) connectivity between those facilities.