

The Case for Traffic Shaping At Internet2 Schools

Internet2 Member Meeting

Arlington, Virginia

May 6-8, 2002

Joe St Sauver, Ph.D.

(joe@oregon.uoregon.edu)

Computing Center

University of Oregon

Background/Context

- We've always been "QoS skeptical" (at least when it comes to traditional wide area premium QoS service), but when Ben Teitelbaum asked, "Does anyone think there'd be interest in I2 for a BOF about traffic shaping boxes?" I made the mistake of replying, "Yes, you bet!"
- As punishment for my sins, Ben talked me into doing this "pro-shaper-deployment" part of today's talk.


There Are Probably 3 Groups of Folks in the Audience...

- For many, the issue of traffic shaping is old news, and something you may already be doing (or not doing) . That's great, either way -- you've at least thought about the issue.
- I also recognize that some of you may consider traffic shaping inherently evil, and the underlying cause of all that's wrong in the world. Absolutely, sure, yep, I agree. :-)
- I'm really here to talk to the remaining group, folks who really haven't made up their mind.

Why Do Traffic Shaping?

Commodity Transit \$\$\$

Those are \$\$\$, folks!



You Can't Go Fast Only On I2

- While the focus of I2 is properly on Abilene and related high performance networks initiatives, what one does to enable high network throughput for I2 purposes has an impact on commodity Internet transit traffic levels, too -- you can't go fast only on Internet2.
- I made this point back during my talk at Joint Techs in Minneapolis in May 2000, and it continues to be true today.

If You Give A Kid a Hatchet, He's Gonna Go Chop Something

- So having ignored that fact, where are we today? Pretty much everybody has engineered and deployed fast campus networks that enable high throughput to Internet2 -- and to the commodity Internet. Those networks usually encompass not only faculty offices and classrooms, but also student residence halls (aka “the dorms”). And students are using what’s been built...

For Example...

- Many sites may have noticed surprising traffic levels associated with dorm users of peer to peer applications such as Kazaa/Morpheus/FastTrack, (running on 1214/TCP) although Gnutella clients (such as BearShare/LimeWire) are also often seen.
- Typically, when such applications begin to be used in an unshaped residence hall environment, they routinely use all or most of the available commodity bandwidth.

Can You Write the Big Check?

- Few sites can afford to offer an ongoing AYCE (“all you can eat”) commodity transit usage policy -- it costs too much to provision that commodity transit bandwidth.
- By implication, if a site cannot continue to grow its commodity transit bandwidth (and also fails to explicitly control its bandwidth usage), wide area congestion will arise... and that congestion may have diverse real and perceived effects (general slowness, domain name service may get erratic, etc.)

Highest and Best Use of Funds

- Even if you could find the funds to continually buy more bandwidth, do you think you could possibly also find other/better uses for that money if somehow you could avoid having to continually feed the insatiable appetite of the “bandwidth monster?” Maybe some new equipment? More support staff? New services? Longer hours?

“What About Just Doing Some Sort of Charge Back?”


- -- Users are not willing to pay (if they were willing to pay, they'd just buy their own T1 or fractional DS3 to do their file sharing)
 - Housing folks usually have no budget to buy bandwidth for their dorm users to burn
 - Billing is a huge pain to administer (think back to those bad-old mainframe charge-back days...)
 - How do you handle “unintentional” high bandwidth usage by naïve/careless users?

Bottom Line, It's the Bottom Line


- Clearly, the desire to spend limited commodity transit dollars more efficiently is one reason why traffic shaping makes sense for most Internet2 schools -- it provides a viable alternative to simply continuing to throw money at boundless demand for commodity transit bandwidth.
- But controlling commodity transit costs is not the only reason why traffic shaping makes sense...

Copyrighted Traffic...

**Lots of traffic
both in AND out**



**In and out both
traffic shaped...**



**Outbound traffic
under control,
but inbound still
running hot...**

File Sharing Apps and Copyrighted Materials

- While file sharing applications can be used in non-infringing ways, simple observation tells us that for many P2P users a work's copyright status is simply moot.
- Out of several hundred Kazaa/Morpheus/FastTrack users we personally checked during Fall 2001, only a couple were offering solely non-infringing materials -- the remainder were distributing copyrighted music, movies, software, etc.

Good News/Bad News...

- Those few rare non-infringing users are a sort of “good news/bad news” phenomena:
 - they weren't infringing (and that's good, we're glad to see them obey the law)
 - BUT, because those legitimate non-infringing users do exist, it would not be appropriate to block P2P file sharing applications outright (and that's bad because it makes things more complicated, but c'est la vie).

An Aside: Why Is So Much P2P Content Infringing?

- I believe its because copyrighted content tends to be what's out there and what users know and like.
- Non-copyrighted content tends to be less common (by sheer volume), less known to the public, often more experimental in nature, and sometimes produced without the benefit of professional quality equipment.

When It's Raining, It's Hard to Avoid Getting Wet...

- Copyrighted content, on the other hand, tends to:
 - be the default norm (99% of all content?)
 - receive extensive marketing attention/airplay (so users become familiar with it),
 - be “well dialed in” to popular tastes, and
 - be technically well produced (big budgets obviously help artists get access to first class gear and talented technical production staff)

Users are Only Humans, and The Temptations Are Great...

- Why dwell on the copyrighted vs. non-copyrighted issue? Simple: it may not be realistic to expect users of P2P applications to have the self-discipline to restrict themselves to just non-copyrighted materials. [It would be great if they would, but that really isn't very realistic -- an easy test of that assertion is the drop in Napster's popularity once it was forced to filter copyrighted materials...]

Here's Your Ticket; Y'all Have a Nice Day Now, You Hear?

- When faced with cases involving documented copyright violations, institutions can (and do) deal with those violations via established student conduct or faculty/staff managerial channels. However, identifying, investigating, and shepherding each such case through “the system” can be quite time consuming, can create substantial ill will with users, and the sheer volume of cases involved can put real strains on staff.

Inbound vs. Outbound

- There is also the pragmatic reality that while it is easy to document copyrighted files that a user may be serving “outbound” (e.g., Kazaa routinely provides a listing of files being distributed from a given server), it is an order of magnitude more difficult to document the copyright status of files being downloaded “inbound.”
- Thus, at many sites, copyright management of inbound P2P traffic is non-existent.

Minimizing the Magnitude of the Problem One (May) Have

- Traffic shaping provides a mechanism whereby an institution can balance the need to (a) permit legitimate non-infringing uses while simultaneously (b) at least minimizing the magnitude of potentially infringing traffic, both outbound and inbound.

Bad is Better Than Worse (Even Though It Is Still Bad)

- Does this mean we're saying that copyright infringement is "okay" as long as it "only" happens a "little bit?" No. Any infringement is wrong.
- What we are saying is that less copyright infringement is preferable to more copyright infringement, given that you may be unable to stop 100% of all infringement from occurring (unless you overreach and completely block even non-infringing uses).

Digital Millennium Copyright Act

- About this point people always seem to ask, “So what about the DMCA?”

I am not a lawyer, this is not a talk for lawyers, nor is this legal advice.

However, if you want to read the DMCA, I’ve included a brief and positively poetic excerpt from it on the following slide.

17 U.S.C. 512(c)

c) Information Residing on Systems or Networks At Direction of Users. -

(1) In general. -

A service provider shall not be liable for monetary relief, or, except as provided in subsection (j), for injunctive or other equitable relief, for infringement of copyright by reason of the storage at the direction of a user of material that resides on a system or network controlled or operated by or for the service provider, if the service provider -

- (A) (i) does not have actual knowledge that the material or an activity using the material on the system or network is infringing;
 - (ii) in the absence of such actual knowledge, is not aware of facts or circumstances from which infringing activity is apparent; or
 - (iii) upon obtaining such knowledge or awareness, acts expeditiously to remove, or disable access to, the material;
- (B) does not receive a financial benefit directly attributable to the infringing activity, in a case in which the service provider has the right and ability to control such activity; and
- (C) upon notification of claimed infringement as described in paragraph (3), responds expeditiously to remove or disable access to, the material that is claimed to be infringing or to be the subject of infringing activity.

(2) Designated agent. -

The limitations on liability established in this subsection apply to a service provider only if the service provider has designated an agent to receive notifications of claimed infringement described in paragraph (3), by making available through its service, including on its website in a location accessible to the public, and by providing to the Copyright Office, substantially the following information:

- (A) the name, address, phone number, and electronic mail address of the agent.
- (B) other contact information which the Register of Copyrights may deem appropriate.

The Register of Copyrights shall maintain a current directory of agents available to the public for inspection, including through the Internet, in both electronic and hard copy formats, and may require payment of a fee by service providers to cover the costs of maintaining the directory.

(3) Elements of notification. -

- (A) To be effective under this subsection, a notification of claimed infringement must be a written communication provided to the designated agent of a service provider that includes substantially the following:
 - (i) A physical or electronic signature of a person authorized to act on behalf of the owner of an exclusive right that is allegedly infringed.
 - (ii) Identification of the copyrighted work claimed to have been infringed, or, if multiple copyrighted works at a single online site are covered by a single notification, a representative list of such works at that site.
 - (iii) Identification of the material that is claimed to be infringing or to be the subject of infringing activity and that is to be removed or access to which is to be disabled, and information reasonably sufficient to permit the service provider to locate the material.
 - (iv) Information reasonably sufficient to permit the service provider to contact the complaining party, such as an address, telephone number, and, if available, an electronic mail address at which the complaining party may be contacted.
 - (v) A statement that the complaining party has a good faith belief that use of the material in the manner complained of is not authorized by the copyright owner, its agent, or the law.
 - (vi) A statement that the information in the notification is accurate, and under penalty of perjury, that the complaining party is authorized to act on behalf of the owner of an exclusive right that is allegedly infringed.
- (B) (i) Subject to clause (ii), a notification from a copyright owner or from a person authorized to act on behalf of the copyright owner that fails to comply substantially with the provisions of subparagraph (A) shall not be considered under paragraph (1)(A) in determining whether a service provider has actual knowledge or is aware of facts or circumstances from which infringing activity is apparent.
 - (ii) In a case in which the notification that is provided to the service provider's designated agent fails to comply substantially with all the provisions of subparagraph (A) but substantially complies with clauses (ii), (iii), and (iv) of subparagraph (A), clause (i) of this subparagraph applies only if the service provider promptly attempts to contact the person making the notification or takes other reasonable steps to assist in the receipt of notification that substantially complies with all the provisions of subparagraph (A).

[etc., etc., etc.]

Not Liable, But Still Guilty... And Maybe Embarrassed Too

- Schools which meet the requirements of that law may be able to minimize their liability as a service network provider.
- However, DMCA doesn't eliminate a University's moral/ethical responsibilities, nor does it eliminate their potential "PR" exposure if a University's users have lots of infringing content online and a journalist or legislator elects to make a crusade out of it.

Copyright and Your AUP

- The other negative side effect of relying on the DMCA is that the DMCA is probably far more rigorous (at least with respect to complaint format) than your AUP may be.

So should you, will you, ignore AUP copyright violations if the reports aren't made in a strictly DMCA compliant format? And what sort of message does it send to your users if you do so?

Sledgehammer or Scalpel?



Force or Finesse?

- Another advantage of using a traffic shaper is that it provides you with the ability to deal with issues “surgically” rather than with sheer brute force. For example, rather than choking all traffic from a given subnet equally (whether for recreational or for academic purposes; whether to on-campus or to off-campus destinations), using a shaper lets you control only selected types of traffic. Shapers provide you finesse.

Sliders Rather Than Toggles

- Without shaping technology, you know your basic options for dealing with problematic usage you discover:
 - you completely block a particular class of traffic, or you don't
 - you turn off a user's port, or you don't.With a shaper, you have a continuously adjustable rheostat rather than a switch, and you can dial in whatever level of usage you can live with. Shapers give you flexibility.

Remember the Innocents

- Some may ask, “What does it matter if we just choke all traffic from a particular subnet? It ‘fixes’ the problem, doesn’t it?” I would urge you to “remember the innocents.” You do have students who are trying hard to get an education, who are “playing by the rules,” and who need to be able to use the network. If you let a few people hog most of the network resources -- or if you choke back everyone too crudely -- it’s the innocents who will suffer most.

“Apply Only to Afflicted Area”

- This is also an appropriate point to mention that if you aren't careful, out of a misplaced sense of “fairness” you may be tempted to apply bandwidth controls “everywhere,” rather than only where a problem exists. Shape traffic only where problems exist, and then as close to the problem as possible. For example, if your dorms are generating too much Kazaa traffic, shape just those subnets, not all campus users.


The Risks of Over Control

- Why would you want to avoid “over controlling” your bandwidth by shaping areas where you don’t have a problem?
 - (1) Your hardware costs increase, as you install shaping where you don’t need it
 - (2) The number of people whose network experience you affect (for no good reason) goes up, making it more likely that you’ll end up antagonizing innocent parties and
 - (3) You don’t want to unnecessarily add YADITP (yet another device in the path)

An Off-Topic Thought About Copyright-Related Issues

- Regardless of what else you may do or may not do, you really should give strong consideration to a copyright education program. Yes, you shouldn't need to tell people about copyright issues, but the need is probably there. There are a lot of myths out there such as “it's okay as long as no money's changing hands,” or “it's sharing, not stealing,” or “its okay to download stuff as long as you're not serving stuff,” etc.

Fairness



Skewed vs. Flat Resource Distribution Curves

- In some parts of the developing world, it is common for there to be a few who are very wealthy and many who are very poor.
- In America, while there are still income inequities (very wealthy and very poor people), most people think of themselves as “middle class.” Most aren’t homeless and unemployed, nor are they so well off that they no longer need to work whatsoever.

Welcome to the Third World (Online Edition)

- As it is in life, so it is online. In an unshaped online environment, there tend to be a few people who take lots of bandwidth, and many people who may be forced to make do with what dregs of bandwidth may be left... An unshaped online environment is thus very analogous to the third world in its fundamentally inequitable distribution of a scarce resource (e.g., bandwidth).

No New Porsches; No Junkers; Lots of Used Year-Old Camrys

- In the alternative online model, bandwidth shaper are used to define and deliver a uniform “middle class” level of access to bandwidth, e.g., let’s hypothetically say a an adequate if not particularly exciting 256Kbps DSL-like service model. There will be no speedsters doing a sustained six or seven Mbps, but equally there will be no jalopies crawling along at only a couple of Kbps per second.

We Need To Learn from ISPs

- When I buy service from a commercial ISP (such as 256Kbps DSL), that service is carefully spelled out. The ISP will (hopefully) give me exactly what I ordered and paid for, no more and no less. Expectations and reality are congruent. Because the ISP has told me what I'm supposed to be getting, and because I am limited to that usage, there are no surprises, no hard feelings. But when have YOU ever told YOUR users what to expect?

If I Never Promise Anything, I Can Never Fail to Deliver...

- The typical higher education model, which typically consists of offering a “best efforts” undefined service, is nice in some ways for those of us who are providers: if I never promised you anything in particular, I can never fail to live up to those (non)promises.
- On the other hand, if I never do define what I’m supposed to be providing to you, does that mean you’re entitled to simply take whatever you can?

Time To Get Businesslike

- One of the side effects of defining a fair network usage policy is that at the same time you define what's appropriate, you end up defining what users should expect. This clarification of expectations and the articulation of a campus' bandwidth delivery policy is one of the unexpected benefits of deploying shaper technology. Defining a shaper policy forces you to quit screwing around and pick a number.

Enforcing “The Number” You Pick

- Among the responsibilities most of us have collectively shirked to date is responsibility for keeping our users from using more bandwidth per user than they should.
- Shaper boxes provide a very convenient way of doing that per-user bandwidth usage limitation, and of delivering fair access to what we’ve said we’re going to provide and no more.

Priorities

- Another way that shapers help us be fair is by allowing us to institutionalize priorities. As universities, it is completely appropriate for us to put teaching and research uses of the network ahead of recreational/personal uses. For example, it would not be fair for someone homework or research to suffer in order to accommodate online gamers. Shapers allow us to establish traffic priorities that mirror institutional priorities.

Mehr Licht!

- At a minimum, deploying shapers helps achieve fairness by bringing a little light to the issue of what's going on on the network. Even if you do nothing else with them, most shapers do a fine job of passively analyzing and reporting on the traffic they see.

Once you know what's going on, you at least have a fighting chance of managing that load equitably.

FAQ's /Excuses :-)



* image courtesy <http://www.cartooncottage.com/images/>

1. “But they’ll just port hop!”

- **Q:** “I’ve considered deploying a traffic shaping device but I’m afraid my users will just run the same old bad applications on random hard-to-identify TCP ports!”

A: Most current generation shapers look at the content of traffic streams (rather than just port numbers), and can readily identify most common P2P protocols regardless of what TCP port a user tries to run them on. [Besides, our experience is that most users *won’t* run P2P apps on non-standard ports.]

2. “Shapers won’t be able to keep up with our traffic!”

- **Q:** “We’d probably install a traffic shaper, except they wouldn’t keep up with our OC12/gigabit/OC48/10gig transit links!”
A: Don’t deploy traffic shapers on your I2 or commodity transit links, install them where you are having bandwidth problems, e.g., just upstream of your residence halls, for example. You CAN easily buy multiple DS3/fast ethernet/OC3-class traffic shapers on the commercial market today.

3. “Traffic shapers will ruin end-to-end performance!”

- **Q.** “I’d install traffic shapers, but I’m afraid they’d ruin end-to-end performance for my faculty and staff users!”

A. So don’t install traffic shapers in front of all users, just in front of incorrigible/chronic abusers. [Moreover, based on the I2 weekly netflow reports (see: netflow.internet2.edu) median I2 TCP throughput is only roughly 1.8 Mbps for single stream bulk transfers -- that’s not really very fast as it is...]

4. “I’ll just charge users for their network usage...”

- **Q:** “After thinking about it, I’m just going to charge users for the network bandwidth they’re using.”

A: Okay, try it and report back on how that experiment goes! [or let’s all see how Time Warner does with this approach for its cable modem bandwidth hogs (e.g., see the article <http://www.nwfusion.com/net.worker/news/2002/0408networker.html> and the discussion that that article spawned...)]

5. “I’ll just shape the traffic on the routers...”

- **Q:** “I really don’t see the need for a separate traffic shaping appliance -- I’ll just shape traffic on our routers.”

A: Bad idea. Let routers concentrate on forwarding packets as fast as they can. Recognize that routers don’t know how to look at the content of flows to identify underlying traffic types. Keep router geeks out of the ugly business of bandwidth policy administration and enforcement.

6. “Shaping traffic would be an invasion of user privacy...”

- **Q:** “Wouldn’t shaping traffic somehow be an invasion of user privacy?”
A: First, few schools explicitly guarantee users any formal level of privacy. Second, shaping traffic is a non-intrusive intervention relative to many other options, such as turning copyright infringers over to the authorities. Third, traffic shaping can be done via technical means and on an aggregated/anonymous basis, if privacy is an issue.

7. “I’m buying a new OC12, I’ll have plenty of headroom.”

- Q: “If I had traffic issues, I’d worry about traffic shaping, but I’ve got plenty of headroom -- they’ll never fill an OC12.”

A. Consider the case of TUDelft, in the NL. Their traffic has gone from roughly 50Mbps in August 2000 to over 450Mbps in November 2001 (http://www.ncne.nlanr.net/training/techs/2002/0127/presentations/200201-bos1_files/v3_document.htm)... are your users that different from theirs?

Conclusion

Traffic shaping isn't perfect, but...

- In conclusion, hopefully one of the forgoing arguments may make you think further about deploying traffic shaping technology. Traffic shaping technology may not be perfect, but it can offer some fundamental advantages you really shouldn't forgo.
- Questions?