

# The Internet is Burning

Richard Bennett

<http://bennett.com>  
richard@bennett.com

ITIF Network Management Forum

Washington, DC  
March 12, 2008

# Who I Am

- Network architect and inventor for 25 years
  - Twisted-pair Ethernet
  - Wi-Fi
  - Wi-Media
- Political Activist in the 90s
  - Built one of the first activist blogs
  - Just as concerned about free speech as anyone
- Out to disentangle free speech concerns from network engineering necessities

# Outline

- Why We're Here
- Decennial Crisis Cycle
- Internet Technology
- Analyzing the Solutions
- Policy and Protocol
- Conclusion

# Why We're Here

- Understanding the crisis of “Video on the Net”
- Critiquing Proposed Solutions
- Developing Sound Policy Framework for Internet regulation

# The Internet's Decennial Cycle

- Every 10 years, the Internet has a traffic crisis
  - FTP caused “Congestion Collapse”
  - HTTP caused the “World Wide Wait”
  - P2P has caused an economic crisis
- Solutions have been similar:
  - Add more bandwidth
  - Re-engineer the offending application
  - Re-engineer the Internet Infrastructure
- This one is different
  - Interest groups taking solutions off the table
  - Faulty pictures of networking taking hold
  - Peer-to-peer very different from Web and FTP

# Cisco Traffic Predictions

- Business Internet traffic CAGR 23% to 2011
  - Consumer IP traffic CAGR 58% to 2011
  - 40% of consumer IP traffic will be Internet traffic, rest is intra-ISP video.
  - Internet traffic is dominated by peer-to-peer (P2P)
  - P2P traffic will more than quadruple
    - 649 petabytes per month in 2006
    - 2,836 petabytes per month in 2011.
  - P2P traffic share
    - 50% of consumer Internet traffic in 2006
    - 39% of consumer Internet traffic in 2011
- <http://www.computerworld.com.au/index.php/id;378881375>

# Japan's Experience

- 100 Mb/s to the home
- 2.5 X increase in traffic in 3 years
- Traffic mix dominated by P2P
- Top 10% of users account for 60-90% of traffic
- Top 10% of P2P users produce 60% of traffic
- Residential networks over 80% busy
  - Not healthy
- Policy debate about who pays

# Internet Technology


# Internet Bandwidth Management

- Back-pressure a fundamental notion
  - A network responsibility
  - Packet drop
  - ECN and Re-ECN
- Provisioning
  - Core links are Symmetrical
  - Residential links are Asymmetrical
- Pricing
  - Residential pricing is based on peak rate
  - Committed Information Rate much higher
  - Raises question of priorities and service levels

# Internet Transport Protocols

- User Datagram Protocol (UDP)
  - Unreliable but timely delivery
  - Lost packets not re-transmitted
  - Preferred by Real-time (VoIP, Videostreaming)
  - Does not participate in congestion management
- Transmission Control Protocol
  - Reliable but un-timely delivery
  - Lost packets re-transmitted
  - Responsible for congestion prevention
  - Multiplicative backoff, additive increase

# Delay vs. Offered Load



- Red = PSTN
- Blue = Shared
- Delay increases with load in shared-bandwidth networks
- Delay insensitive to load on the PSTN
- Prioritizing combines these two characteristics by limiting the effect of low-priority load on high-priority flows

# Router Behavior

- Discard packets when backlog builds up
- Don't discard short packets
- Don't discard TCP ACK packets
  - This would cause re-transmission
- Favor UDP over TCP
  - Discarding TCP has a bigger payoff for Congestion
- Create back-pressure where it's most effective

# Freedom, but not Fairness

- Original model provided equal bandwidth to each process running on a time-sharing system
- Current model only works if each user consumes the same bandwidth by choice
- **Persistent and Multiple streaming applications break this model: P2P**
- Assumes good behavior
  - No unresponsive Flows
  - No Defective TCP implementations
- Packet-drop highly inefficient
  - Also ambiguous in Wi-Fi settings

# Analyzing Solutions

# Can we build our way out?

- Japan's experience says fatter pipes are nice, but not the answer.
- Instead of **simply** adding bandwidth, we need to engineer our way out of the crisis:
  - More choice in customer agreements regarding volume, routing, and services
  - Assign costs appropriately
- This will require us to understand how things work

# Management Solutions

- User Quotas
  - Successful approach in business and campus networks
  - Some quota systems reach QoS levels
- Explicit Congestion Notification
  - Less re-transmission than packet drop
  - Held up by Microsoft
- Novel uses of TCP control bits such as Reset
- Prioritization

# Prioritizing Networks

- Superframe networks
  - DOCSIS
  - IEEE 802.15.3a UWB
  - WiMax
- Contention networks
  - VLAN Ethernet
  - WMM subset of 802.11e
  - WiMedia UWB

# Priorities Motivated by Modes of Communication

- Person-to-person:
  - Low delay + jitter, low bandwidth
- Machine-to-person:
  - Medium jitter, high bandwidth for short time
- Machine-to-machine:
  - Jitter doesn't matter, high bandwidth consumption
- Network maintenance:
  - Importance trumps all other applications because it keeps the network running.

# And lest we forget...



[Search Torrents](#) | [Browse Torrents](#) | [Recent Torrents](#) | [TV shows](#) | [Music](#) | [Top 100](#)

[Preferences](#)  
[Languages](#)

All  Audio  Video  Applications  Games  Other  (search titles only)

## How do I download?

[Login](#) | [Register](#) | [Language / Select language](#) | [About](#) | [Legal threats](#) | [Blog](#)  
[Contact us](#) | [Usage policy](#) | [Search cloud](#) | [Doodles](#) | [Forum](#) | [Court Blog](#)  
[SlopsBox](#) | [Bayimg](#)

2.580.598 registered users. Last updated 04:50:04.  
5.580.993 peers (2.267.992 seeders + 3.313.001 leechers) in 728.636 torrents on tracker.

# Policy and Protocol

# Protocol Upgrades

- Adopt a multiple-service model
- Make better use of DSCP by standardizing a set of markings between home gateways and ISPs
- Enable ECN in Windows
- Pursue the **Re-ECN** experiment
- Allow ISPs to experiment with business models, pricing, and technical solutions
- Collect better data on residential networks

# One Standard Method: IEEE 802.11e Classifier Type 1

- *In a TCP or UDP header: Source Address, Destination Address, Source Port, Destination Port, and Version, plus:*
- *One of the following:*
  - *In an IPv4 header: Differentiated Services Code Point (DSCP) (IETF RFC 2472) and Protocol, or*
  - *In an IPv6 header: Flow Label. - p. 55*
- Uses both elements of Clark's dichotomy
- Discriminates based on protocol, among other things.
- By some definitions, this is “Deep Packet Inspection”

# Advanced Bandwidth Management

- Packet Scheduling
- Traffic Shaping & Policing
  - Organizing traffic by type
  - Implemented in end point, by network, or both
  - IEEE 802.11e example
- Admission Control
  - Matches resources with flows
  - Uses TSPEC
  - Comcast RST on excess streams
- Home routers do these things today, not just Great Firewalls in China.

# Disclosure

- Hard problem because terms are not well-understood
- Interest groups have deliberately created smoke
- ISPs don't want to disclose more than others do
- Academics want data
- The first step

# Conclusions

# Network Neutrality Harmful

- Tiered pricing is the norm and always has been
- The Internet is not purely an end-point network
  - Routing is wholly an internal function
  - AQM, RED, and Protocol knowledge essential
- Internet's fairness model has never been right
  - Every flow has an equal chance
  - Heavy users (lots of flows) impair light users
- Assumption of good behavior on the part of the end-user is naive

# FCC Guidelines

- 1 Does the practice support a rational goal, such as the fair distribution of bandwidth?
- 2 Is it applied, adapted, or modified by network conditions?
- 3 Does it conform to standard Internet practices, or to national or international standards, and if not, does it improve on them?
- 4 Has it been communicated to customers?
- 5 Has technical information that would allow for independent analysis been made available to the research community and the public at large?
- 6 Does the practice interfere with customer control of traffic priorities or parameters consistent with terms of service?
- 7 Is the practice efficient with respect to both the upstream and downstream data paths?
- 8 Does the practice accomplish its purpose with minimal disruption to the network experience of customers as a whole?

# Backup and Miscellaneous Slides

# RFC 3186, ECN for IP

*... there are serious concerns about the damage that can be done by non-compliant or unresponsive flows (that is, flows that do not respond to congestion control indications by reducing their arrival rate at the congested link). For example, an end-node could "turn off congestion control" by not reducing its congestion window in response to packet drops. This is a concern for the current Internet. It has been argued that routers will have to deploy mechanisms to detect and differentially treat packets from non-compliant flows [RFC2309,FF99]. It has also been suggested that techniques such as end-to-end per-flow scheduling and isolation of one flow from another, differentiated services, or end-to-end reservations could remove some of the more damaging effects of unresponsive flows.*

# ECN and Deterrents to Evil Applications

*I don't think that there are **\*\*any\*\*** deterrents in the current Internet to keep connections from ignoring congestion control (other than the fact that they will know that they are being socially irresponsible). Somehow, it is still the case that roughly 95% of the traffic seems to be using some form of vaguely-conformant TCP, the steady-state packet drop rate is not spiralling up to 10 or 20 percent, and congestion collapse has not yet hit. I think that we shouldn't press our luck, and I think that deterrents are indeed necessary, but I don't think that the addition of ECN makes the job of providing such deterrents any more difficult, or any more pressing. I think that the need for such deterrents is already pretty pressing.... - Sally Floyd, <http://www.icir.org/floyd/email/sf.97dec01.txt>*

# Flow Rate Fairness: Dismantling a Religion

- *Resource allocation and accountability keep reappearing on every list of requirements for the Internet architecture. The reason we never resolve these issues is a broken idea of what the problem is. The applied research and standards communities are using completely unrealistic and impractical fairness criteria. The resulting mechanisms don't even allocate the right thing and they don't allocate it between the right entities. - Bob Briscoe, UCL and BT Research*

[http://www.theregister.co.uk/2007/12/13/bennett\\_eff\\_neutrality\\_analysis/](http://www.theregister.co.uk/2007/12/13/bennett_eff_neutrality_analysis/)

# The Trouble with BitTorrent

- An economic innovation, not a technical one
- No relationship between traffic and cost
  - Loads most expensive links
  - No economic limit on bandwidth consumption
- Hides itself from network management
  - Header obfuscation
  - Tracker obfuscation
- Long-running flows with multiple connections
  - Less responsive to back-pressure than HTTP
- Makes traditional means of management ineffective

# Beyond End-to-End

*What is needed is a set of principles that interoperate with each other—some build on the end to end model, and some on a new model of network-centered function...from the beginning, the end to end arguments revolved around requirements that could be implemented correctly at the end-points; if implementation inside the network is the only way to accomplish the requirement, then an end to end argument isn't appropriate in the first place. -*

“Rethinking the design of the Internet: The end to end arguments vs. the brave new world,” Clark and Blumenthal, 2000

# FCC Harvard Hearing

- Free Press/Public Knowledge “Deadwood System” of traffic management not accepted
  - Instead of “degrade all streams equally” there was a recognition that priority organization makes sense
- Debate about methods of classification
  - ISP vs. User dichotomy
  - In reality, it must be a co-operative effort
- No “bright line rule” emerged
  - Multi-part test is needed
- Data collection and disclosure needed

# Good vs. Bad Traffic Management

- An empirical question, not a philosophical one
  - What immediate purpose does it serve?
  - What are the network conditions?
  - Who benefits from it?
- Digital packet networks are unique
  - Providing appropriate service to applications is critical
  - Assigning bandwidth from a flexible pool is fundamental
  - Not just a series of fatter tubes

# Dangers of Traffic Management

- Degradation of a competitive service
  - Can be active or passive
  - Bandwidth demands of some applications have this effect.
  - Failure to manage can be discriminatory in effect.
- Over-management may reduce system throughput
  - Not a live problem any more due to advances in hardware
- Hard to explain to the user

# Dangers of Regulation

- Some proposed regulations make it impossible to have the enhancements we need:
  - Per-user quotas to manage bandwidth hogs
  - Truly competitive VoIP and Teleconferencing services enabled by Bandwidth-on-demand
- Proposed regulations break new ground
  - Internet **switching** has always been free of regulation
  - The Internet is more than telco lines
- There is no such thing as a “light touch” on a hair trigger

# The Virtual Network

- The Internet is less a network than an abstraction of the data links that are the real networks.
- As the data links change, so must the abstraction
- The dumb Ethernet of the 1970s has long been obsolete
- New data links support QoS and active management
- The Virtual Network needs to support its constituent parts.

# Origins of the Debate

- Blog debate on symmetrical pipes, network capacity, and ISP management started in 2003.
  - Lessig
  - Gillmor
  - Searls
- Occasioned by media ownership rule changes
- Opponents of the new ownership rules argued that first-mile links could never support “New Media.”
- Does free speech require application neutrality in the first mile?

# Prophecies of Doom

- *Make no mistake: The freewheeling Internet as we know it could very well become history. - Save the Internet*
- *If an ISP does not like Christian Coalition of America or the ACLU's emails, they can and will block our emails from going out over their server. - Christian Coalition*
- *Craigslist has been blocked for three months from Cox customers because of security software malfunctions. - Matt Stoller*
- *We're looking at the beginnings of a private police state - Tim Wu*

# A Mis-framed Debate

*Network neutrality is a completely mis-framed debate. It assumes that the user has access to a single telco product: pre-paid (by the user) ISP access. The real market will be vastly more complex, with users having access to many “virtual” networks — some overlaid on the Internet, some private. All the bogeymen making noise about blocking and throttling are just the shadows of welcome improvements in the wholesale markets. - Martin Geddes, Jan 15, 2008*