

This Green Revolution

Improving the yield of your network investment

Internet Reloaded
March 2009

Eric Klinker
eklinker@bittorrent.com



 **BitTorrent** =



What We Do

“we make the Internet more efficient”

For Consumers:
Efficient access
to content

For Publishers:
Lower delivery
costs, less
infrastructure

For Operators:
More efficient
use of the
network

*Core Technology to
organize User-Contributed Bandwidth
into a Viable Resource*

BitTorrent Consumer Clients

Consumer BitTorrent freeware:

- The world's 2 most popular BitTorrent clients:

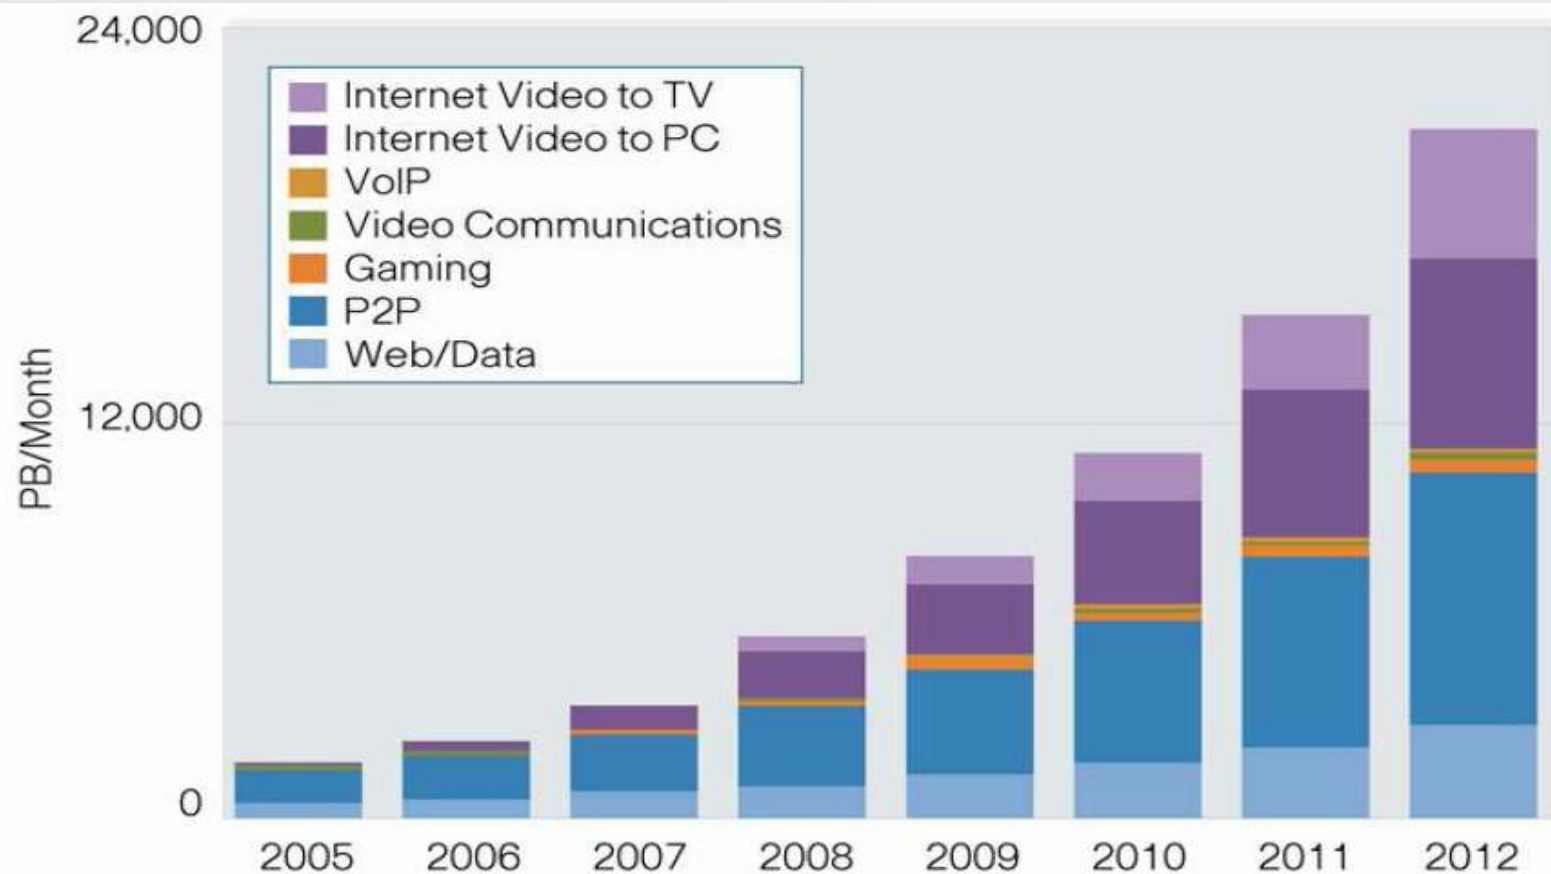


- >200m downloads to date
- >40m active clients monthly
- High performance
- Minimal resource utilization



Internet Traffic Growth

Publishers will soon be pushing traffic that rivals P2P usage today



So what is an Exabyte?


- Roughly = 1 thousand PetaBytes = 1 million TeraBytes
= 1 billion GigaBytes = 1 trillion MegaBytes

3,507,000	months of TV (stream encoded at 1 Mbps)	(292,000 years)
64,944	months of Blu-Ray DVD (max std 54Mbps)	(5,412 years)
20	months of 2007 YouTube traffic	(1.7 years)
1	months of BitTorrent traffic	



How to Deliver an ExaByte

Past



Content Delivery – Past

- Content delivered directly from company server
- Delivery time slow, often interrupted and expensive

Present



8% of the capacity of a large modern nuclear power plant (e.g. Japan's 1100 MW Higashidōri



Roughly equal to all revenues of CDN market leader Akamai

CDNs – Present


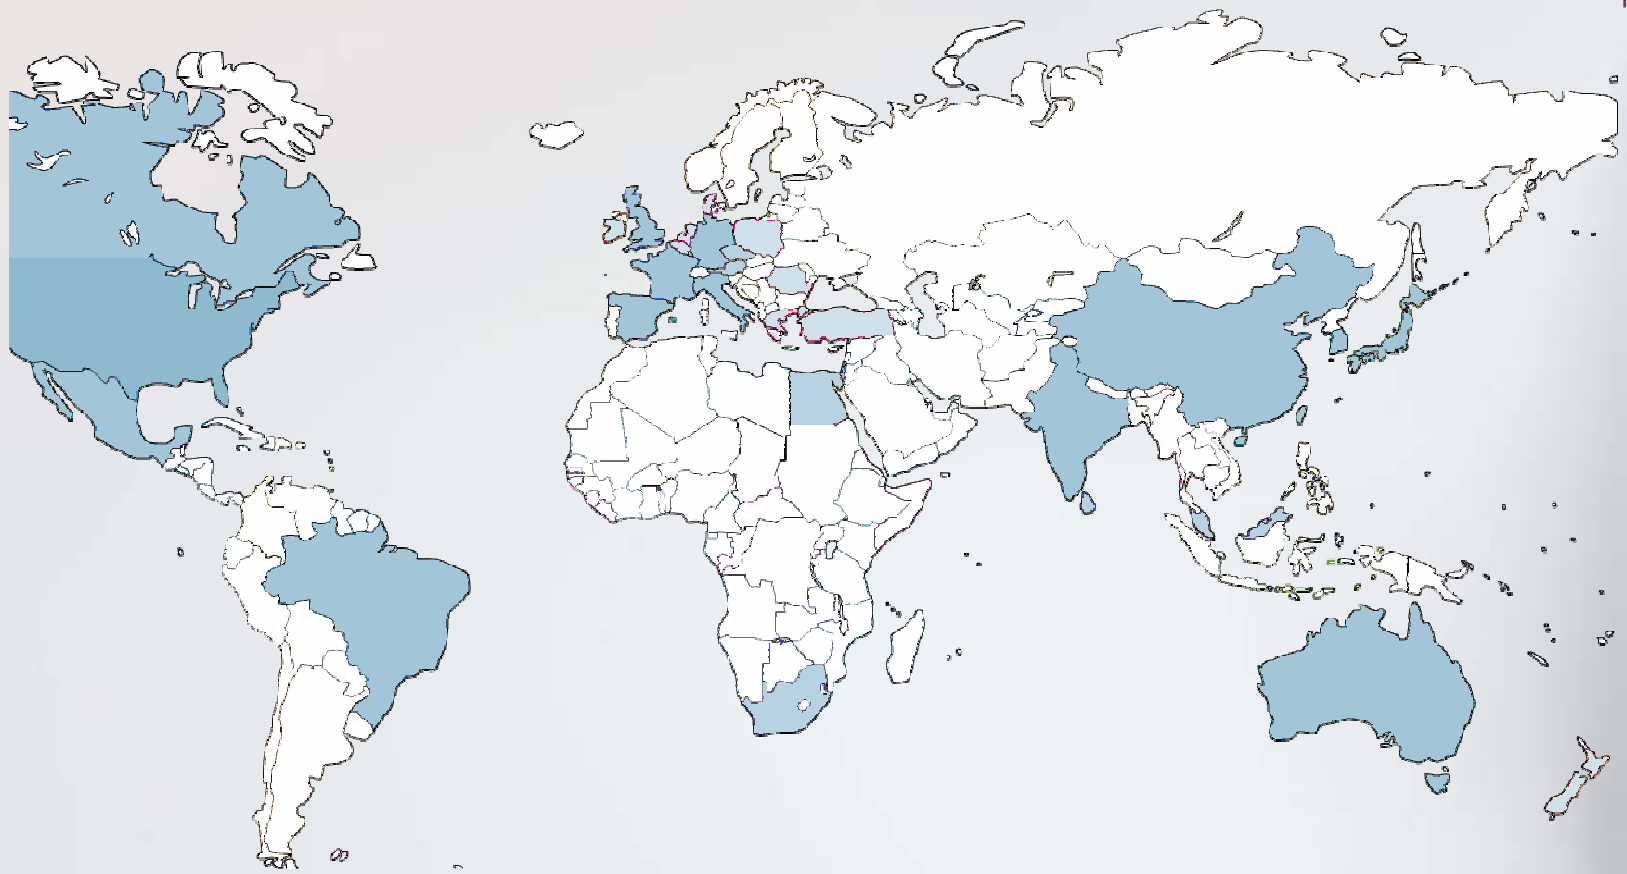
1	Exabyte/month
~70,000	Servers
~1.2	KWatt/ server (power + HVAC)
84	MWatts

1	Exabyte/month
~7,000	Gbps capacity
~\$10	Price Per Mbps
\$70M	Monthly

pure software solution

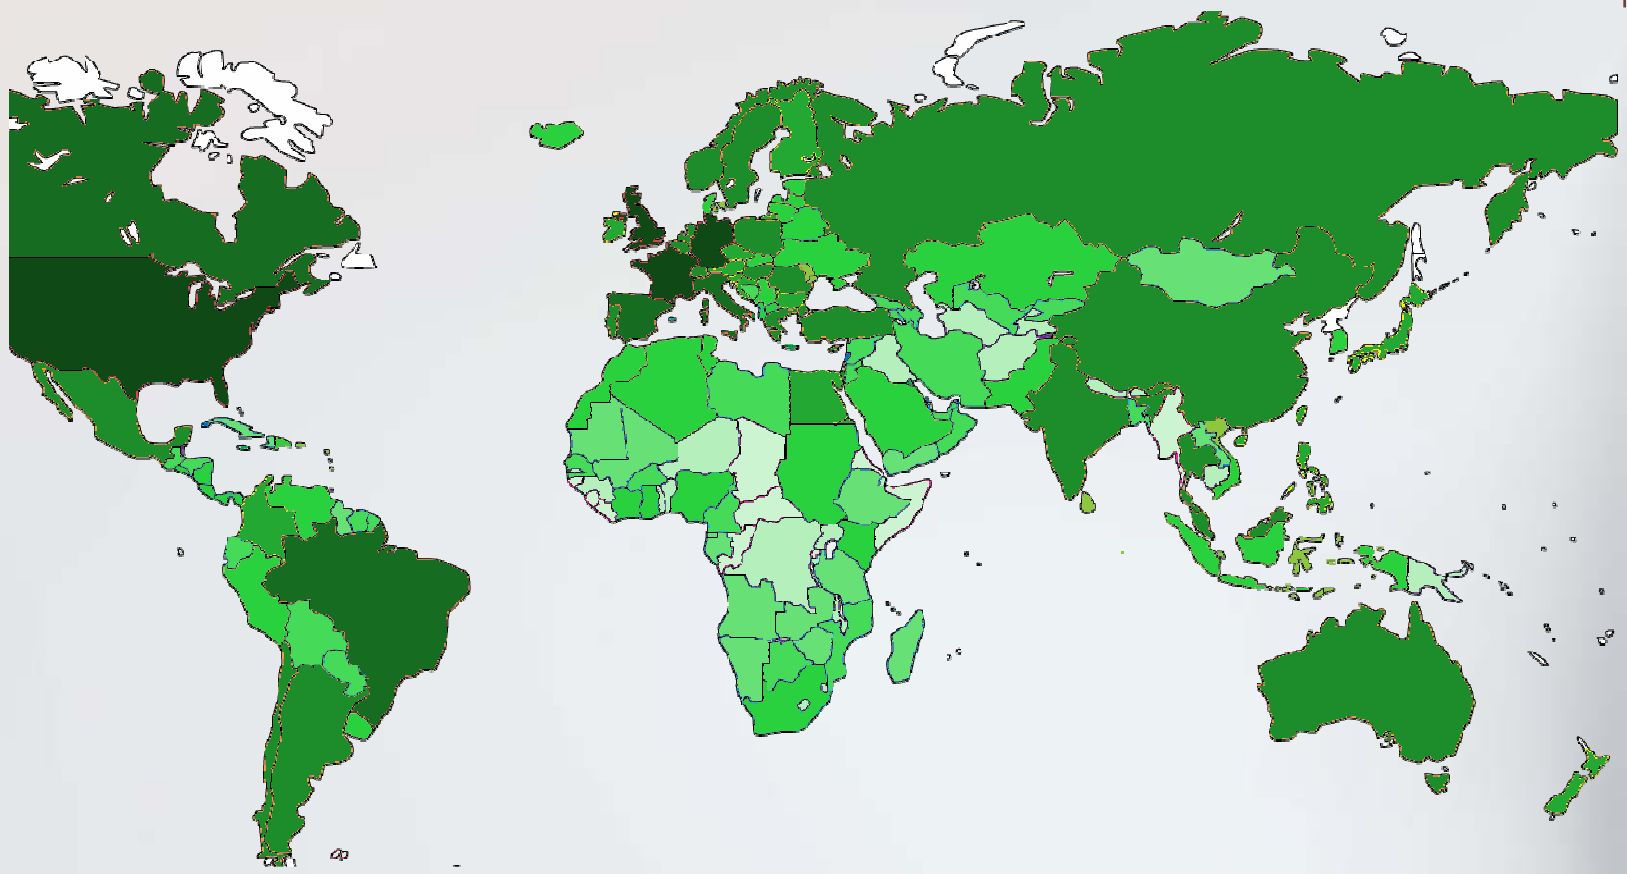
Market Coverage – Big CDN

Publishers




Market Coverage – BitTorrent

Publishers



Market Saturation



Improving Efficiency for Operators

- Policy and Routing
- Congestion Control



Policy

- Applications make ultimate peer selection based on peer performance
 - This does not consider potential efficiencies of
 - Capacity
 - Cost
 - Traffic Engineering
- Policy gives operators a say in these decisions
 - More effective traffic engineering
 - Help guide applications in the overall use of the network
- Can you approach Multicast efficiencies in network utilization?



Solutions to the Policy Problem

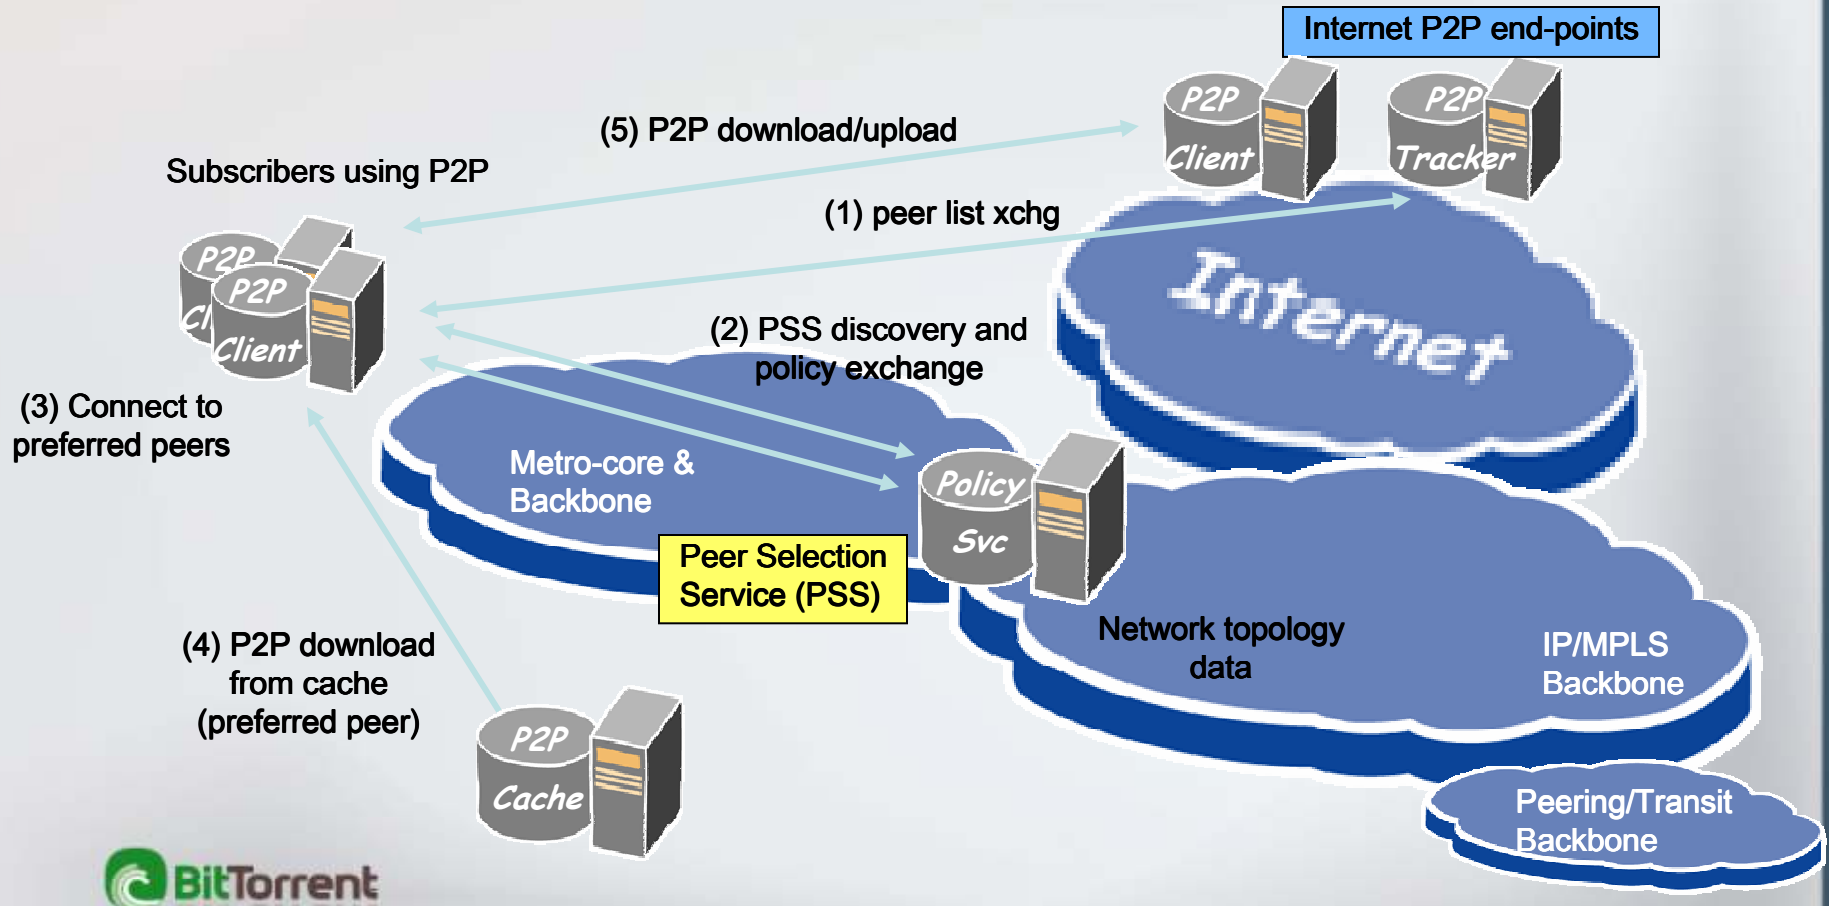
- “P4P” working group
 - Research sponsored by Yale University
 - Promising results from limited trials
 - However, the architecture is not suited for consumer BitTorrent
 - Trackers in this ecosystem are generally uncooperative
 - Clients learn about peers using decentralized means
 - DHT
 - PEX

- ALTO working group in the IETF
 - Recently chartered, first meeting in Minneapolis Dec '08
 - Envisions something closer to a centralized policy service



Architecture

- Goals:**
- Reduce Operator expense
 - Improve the manageability of the network



An Example Policy Format

cidr:10/8:10

asn:0:5

cidr:10.1/16:20

cidr:10.2/16:-10

cidr:[de:ad:be:ef:fe:ed]/48:20

- Positive weights are “Preferred”
- Negative weights are “To be avoided”
- Zero weight (or IPs not covered) are the default set
- The greater the weight, the greater the preference



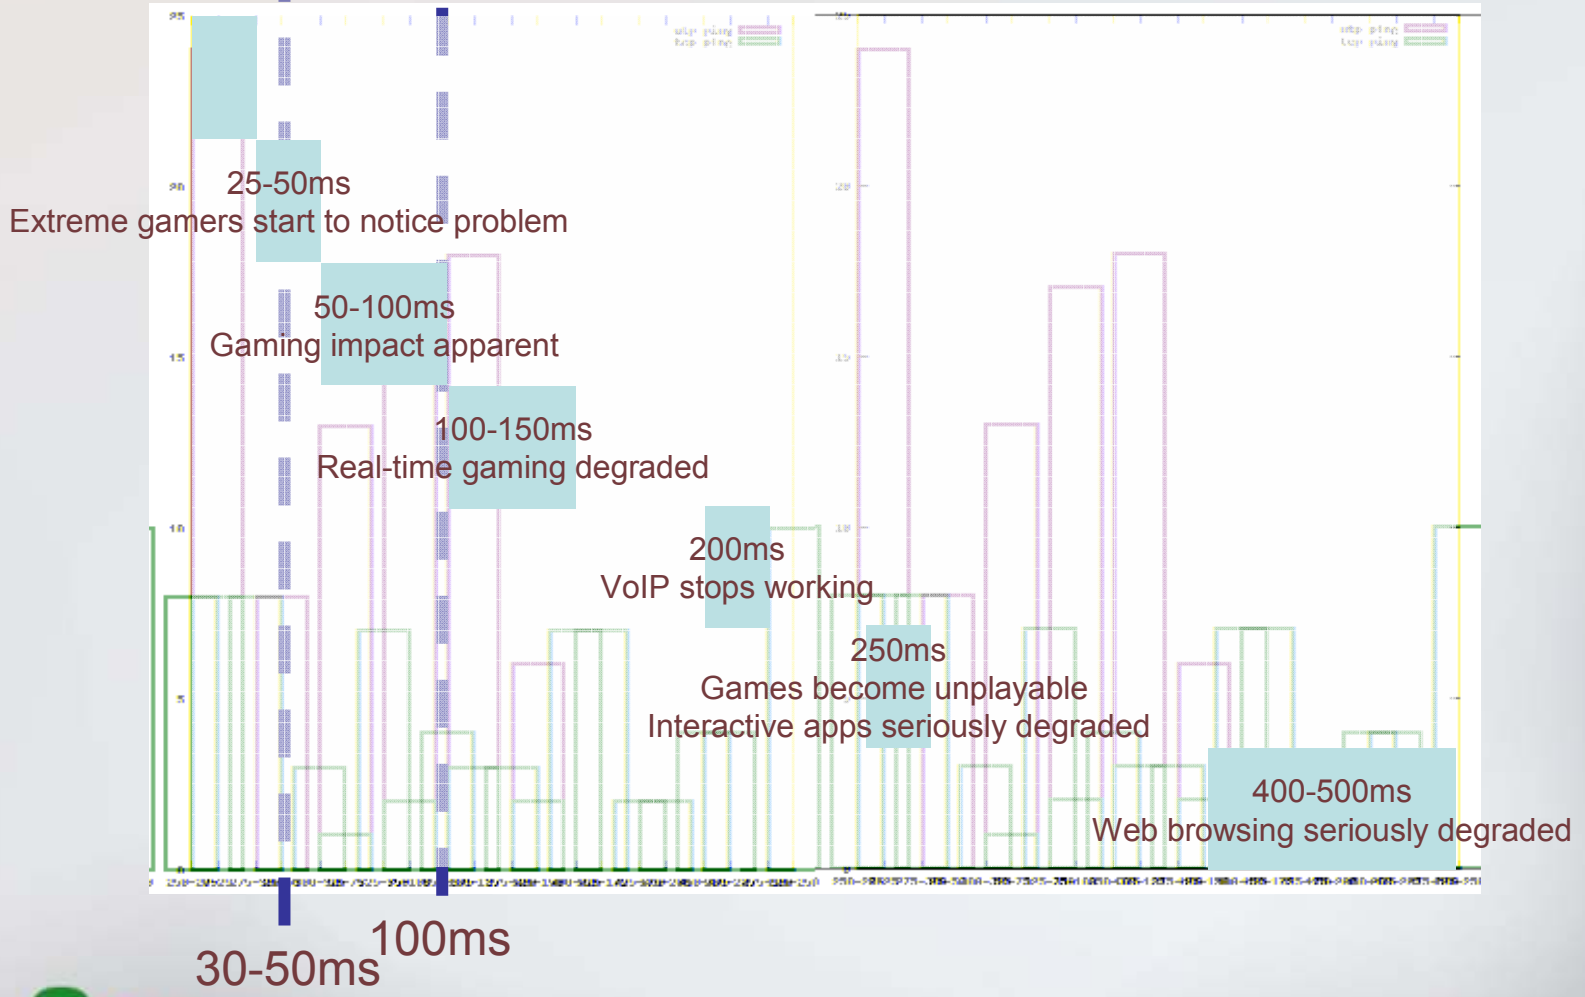
We Need a New Congestion Control Protocol

- Replace TCP for P2P transfers
- Design Goals for this new protocol:
 - Keep the bottleneck full (maximize throughput)
 - Keep the delay lower than the unloaded delay + ϵ
 - Yield to TCP on the forward path
 - Separate reverse path congestion
 - React in 1 RTT
- Congestion control approach:
 - Continuously estimate one-way delay
 - Separate queuing delay from propagation delay
 - Target a small value for queuing delay

We call this new transport “uTP”

Impact of uTP with target of 100ms

Legend:
[Red Box] UTP
[Green Box] TCP



Implications for Network Operators

- uTP creates a “Scavenger Service”
 - Only idle capacity is utilized
 - Other applications have priority
- Networks can run “hotter” (with revised cap planning)
- Networks can now use capacity previously reserved
 - VPNs and MPLS
 - Operators enjoy improved profitability or reduced capex
- User complaints from P2P greatly diminish
- No ability for P2P from one user to harm another user

Potential for greater and more efficient use of the network

Status of uTP

- Implemented inside the 'DNA' commercial client
- Instrumented to minimize queuing delay
- Extensively tested in the lab
- Tested on the Internet in over 14M DNA clients worldwide
- Works as designed
- LEDBAT working group formed in the IETF
- Currently in testing with 400k beta consumer clients
- Planned implementation in over 40M active consumer clients.

What about Users??

- Users will have to “opt-in” to policy usage
 - ISPs can’t implement policies that hurt users
 - ISPs should provide an incentive to users
 - E.g. local, “on-net” bytes are *free* or don’t count against a cap
 - “Performance benefits” alone are likely insufficient
- Better congestion control improves the user experience
 - Eliminates common local congestion effect on the home network
 - Equivalent to automatic bandwidth management settings
 - ISP offered incentives for users work here as well



A Fork in the Road Ahead



Scavenger transport is a failure

Voracious clients emerge that break from the friendly posture

A return to the zero sum game of measure and countermeasure

The tragedy of the commons is confirmed

Scavenger transport is successful

Shaping/Blocking becomes an anachronism

Networks enjoy a period of profitable growth

The myth of the tragedy of the commons is confirmed

Thank you

