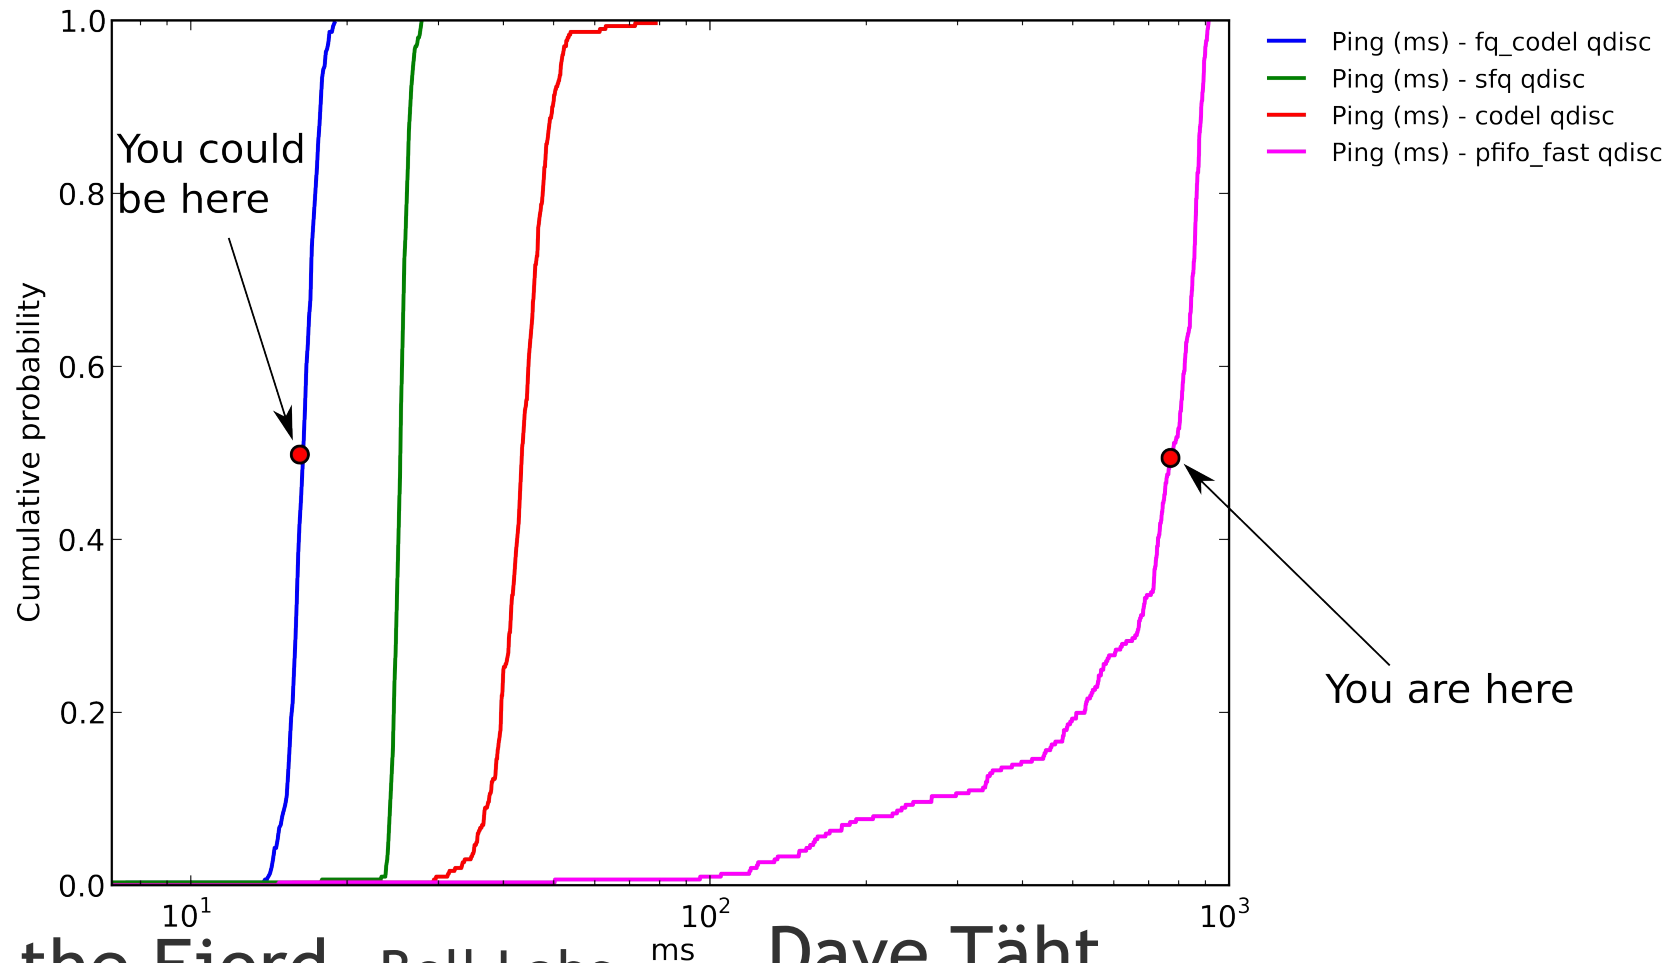


# CeroWrt against an Uncaring Universe

“Running Code” (and Rough Consensus)”



Jim Gettys  
Chairman of the Fjord, Bell Labs

[jg@freedesktop.org](mailto:jg@freedesktop.org)

Dave Täht  
Hindmost, Bufferbloat.net  
[dave.taht@gmail.com](mailto:dave.taht@gmail.com)

# March, 2011:

## Bufferbloat project starts CeroWrt

- We needed to attack the bufferbloat problem at the theory, OS, network stack and engineering level.
- In particular, we needed to **know** new algorithms were deployable on consumer hw and edge devices
- And we needed some ground truth about what was in the marketplace – in other words

Why did home gateways suck so badly?

# The sad Ground Truth on home CPE and home routers

- Vendors shipping 5+ year old kernels (2.6.32? WTF?) and often even older packages
- Operating on razor thin margins... with no perceived demand for features like IPv6
- Most coding outsourced... but usually dependent on ancient ports by the chip vendor
- Total disconnect between retail, academia, ISPs, IETF, and hw vendors
- Binary blobs everywhere resulting in field non-updateable firmware.
- Result: GRIDLOCK.

# How to get obsolete software fixed on CPE?

- **Open Resolver**

- 350051 'skbroadband'
- 236311 **'dnsmasq-2.40'** Released 10 YEARS AGO
- 212406 '.'
- 187012 'Nominum Vantio 5.3.0.0'
- 169557 'TelecomItaliaDNS'
- 168460 'PowerDNS Recursor 3.3 \$Id:
- 159537 'Go away!'
- 158940 'Cyber World Leader Kornet!'
- 151151 'DNS server'
- 136084 'N/A'
- 125645 'Nominum Vantio 5.3.2.2'
- 123798 '9.7.3'
- 119653 'cdns5-asd2-1-100823-1441'
- 104603 'Dns Server'
- 95402 'unknown'
- 87458 **'dnsmasq-2.52'**
- 76437 'x.x.x'
- 74282 'None of your business'
- 65724 'yamutech-bind'
- 60922 **'dnsmasq-2.55'** Released 2 years ago
- 60240 'Microsoft DNS 6.1.7601 (1DB14556)'
- 57484 'dnsmasq-2.45' Released 4 years ago
- 55315 'CTL'
- 54800 '9.8.1-P1'
- 54201 **'dnsmasq-2.48'** – Released 5 years ago!

- Note: DNSMasq runs on about a billion boxes ! now – home and commercial routers, a ton of CPE, used heavily in Virtual Machines, it's the default in ubuntu, and, um, android uses it.
- 
- And it's not broken here, just misconfigured by the vendor...

- **Source:** <http://openresolverproject.org/version.bind.report.txt>

# Why CeroWrt? - Do research on a *real* home router

- CeroWrt is our OpenWrt based test platform for bufferbloat, IPv6 and DNS, and Routing work (and whatever else comes up)
  - For fixes to go upstream easily, experience says to stay within 2 releases of kernel.org: five year old vendor code need not apply!
  - Even supporting **one** hardware platform instance is lots of work, since drivers and the OS itself have problems
- Needed an attractive nuisance
  - something that people like you would like to run, not just a one off test test platform for a single student/research group using one (stale) Linux release; we knew fixing everything we wanted to fix will take several years.
  - Something to point to when a vendor says nonsensical things
  - Realization: that OpenWRT could be such a platform, and that *we must also address IPv6, routing, security and naming* as first class issues...

# Problems in Home Routers and CPE

- Business model: 5 year old code is as good as it gets today
  - Go from WG to running, tested, deployed code takes even longer
- Bufferbloat: AQM experiments vital
  - ECN, good classification & diffserv also desirable
- Bridging ethernet and wireless networks is a very bad idea
- Wanted “Grandma proof” configuration: need to be able to plug routers in arbitrarily and have them “just work”
- IPv6: why the heck was it taking so long to deploy?
- Multi homing networks
- Name services: IPv6 brings naming to a critical impasse
- Autonomous operation when your Internet link is down
- Security: secure, ongoing update needed – and DNSsec!

# CeroWrt's Plan


- What if bufferbloat.net focused on a single platform, with entirely open source software, tackled bufferbloat, dns, and ipv6 head-on... convinced a lot of top people to help, and sat down to fix every problem that came up?
- And spun new releases *really* fast, prove that new algorithms would work on common CPE and home gateways?
- And then pushed the results upstream to Openwrt, the mainline linux and the package maintainers? Heck, what would happen if we convinced those guys that the bloat, DNS, and IPv6 problems were cool, interesting and urgent?
- And if we succeeded – if we then made the home router and CPE feature set so compelling that we could get bufferbloat fixes out there AND slip in a major IPv6 upgrade along with it?

*We thought that would be Cool!*

# And, now, with fq\_codel...

## Bufferbloat is looking rather fixed!

Realtime Response Under Load  
Download, upload, ping (scaled versions)

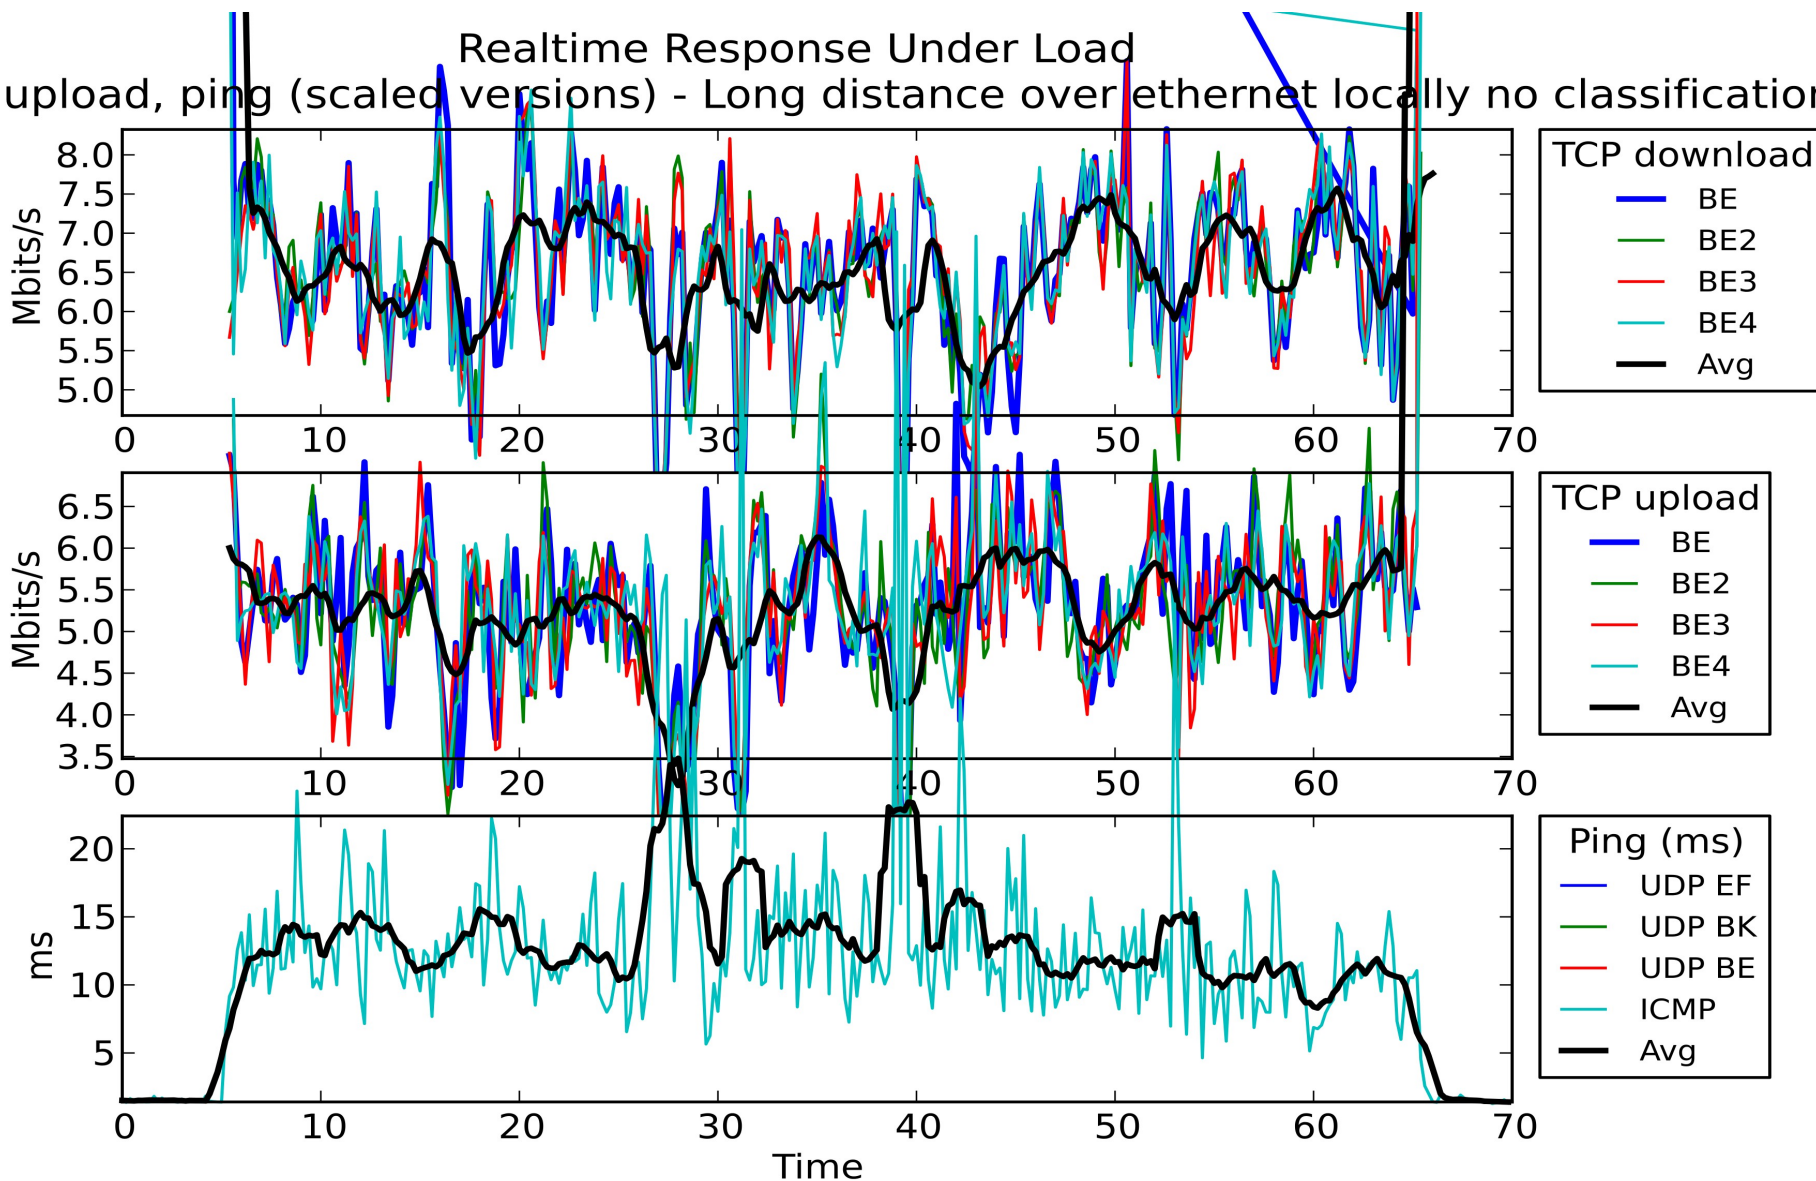




# Here's a 4 hop mesh on wifi

Realtime Response Under Load

Download, upload, ping (scaled versions) - Long distance over ethernet locally no classification



# All fast->slow transitions on the Net benefit from (fq\_)code actually

- Rate limited user access and provider interconnects
- Load Balancers
- Servers in general
- Servers under Virtual Machines
- Desktops, handhelds and wireless devices
- And yea, home routers and CPE
- And dslams, CMTSes, wireless, LTE...
- And it was really nice to test them on a single well understood platform...

# Bufferbloat.net Resources

Bufferbloat.net: <http://bufferbloat.net>

Email Lists: <http://lists.bufferbloat.net>

IRC Channel: #bufferbloat on chat.freenode.net

CeroWrt: <http://www.bufferbloat.net/projects/cerowrt>

Other talks: <http://mirrors.bufferbloat.net/Talks>

Jim Gettys Blog – <http://gettys.wordpress.com>

*Various videos of the fq\_codel world tour*

***<http://www.bufferbloat.net/projects/cerowrt/wiki/Bloat-videos>***

# What is CeroWrt? - Hardware

Modern, well supported home router with fully open drivers for all hardware and interest from the developers

- Atheros wireless driver is the most mature: 802.11n aggregation makes wireless drivers really, really, **really** “interesting”
- Enough flash so we don't spend time engineering around its lack
- Initial target: Netgear WNDR3800;
  - A modern abgn dual radio gigabit router,
  - Fully open source including the 802.11 driver
  - 680Mhz MIPS cpu, 4 port gigE switch, gigE WAN, 64M RAM
  - 16M flash to speed development
  - Costs \$119 (obsolete now, getting hard to find on used market)

# What is CeroWrt? - Software

- A development build of OpenWrt with unusual features
  - 3.8.13 Linux kernel + westwood+, reno, veno, cubic, bic; ECN enabled  
ISC BIND 9.9.X + DNSsec available optionally, ECN enabled
  - Bufferbloat fixes throughout
  - AQM's – SFB, DRR, Codel, FQ\_Codel, PIE
  - Extensive network test tools to ease research and development
  - IPv6 (native, 6in4, 6to4, 6rd, ULA border, dhcp-pd, etc)
  - Lighttpd for web, polipo web proxy (ipv6 to ipv4 translation)
  - Driver fixes for ath9k buffer management, ECN & other fixes if not yet upstream in OpenWrt and/or kernel.org
  - Routing, not bridging, between radios and ethernet
- Plenty of space left in flash: squashfs build is 8.9MB.
- It's a RESEARCH project – results go upstream. Get your code from there unless you want to be WAY beyond the bleeding edge....

# Cerowrt's IPv6 test sites

- ISC.org's Cerowrt LAB: IPv6 native
- Yurtlab: Various tunnels and native backbone
- Comcast: DHCP-PD fully tested
- Jim Gettys: 6to4
- Multiple users: 6in4
- Users worldwide: ipv6 all formats
- Netperf based Bandwidth and analysis sites setup worldwide for RRUL testing...

# IPv6 related fixes (all upstream now)

- We treated IPv6 as a first class object throughout all development and testing
- Many encapsulation bugs fixed
- ECN bugs fixed
- 2X speedup on IPv6 forwarding over manufacturer's hardware
- Whole bunch of unaligned instruction traps fixed
- IPv6 routing bugs fixed
- Ipv6 native, 6rd, 6to4, 6in4, and dslite landed last week
- Tiny DHCP-PD client and server, AHCP, RA support via dnsmasq
- Various RFCs far more fully supported
- I've kind of lost track, but by Linux 3.6 I was pretty happy with IPv6 on CeroWrt... and as of last month, everything major in Cerowrt landed in OpenWrt
- OpenWrt has enabled IPv6 *by default in the pending Barrier Breaker release*. On every (100s) hw platform they support.

# Results of spinning the code fast

- Bufferbloat is Beaten
  - Went from ns2 code to cerowrt to linux mainline kernel in 1 week flat last june...
  - FQ\_Codel landed a week later.
  - We've been analyzing it ever since...
- IPv6 “just works” to a very large extent
- Bind9 DNSSEC proven to work on a home router, DNSSEC on dnsmasq work commencing..
- And... Incredible stability throughout the continuous integration and development process...



# One of my favorite bugreports

*“Uptime as of this morning was 259 days. Sometime during the night, all the computers connected by wifi fell off the net.*

*I once worked on an OS that had a really weird bug that caused its network stack to crash after exactly 259 days uptime; it turned out to be a 32-bit counter that went negative when that much time had passed. The fact that this glitch also happened at 259 days uptime could easily be a coincidence. But maybe it isn't. If it happens again on October 6 I'll definitely let you know.”*

# IPv6 Integration improvements in openwrt barrier breaker”

- TWO tiny dhcpv6 implementations now exist
  - DNSmasq gained tight naming support for SLAAC and DHCPv6 devices... in 294k (including full primary dns resolver support, secondary zone transfer and full ipv6 support)
  - Odhcpv6 supports (in 30K!) PD client and Server mode...
- Much improved firewalling and RFC compliance
- 3.8.X Linux kernels has vastly better ipv6 support
- Still missing a working IPv6 enabled shaper (cero has several)
- Way more stuff than I can remember...

# CeroWrt today


- Pretty darn stable “Modena” release:  
<http://www.bufferbloat.net/projects/cerowrt/wiki>
- All major patches pushed into OpenWrt, the linux kernel, various other distros and the core package maintainers.
- Got Cisco's PIE friday, had it in a test release 12 hours later, next release is gated on pie
- Improvements to ipv4 and IPv6 traffic shaping

# Ongoing Work

- Simon Kelly is adding DNSSEC to dnsmasq
- Architectural rework of Linux WiFi nearly complete, work commencing this summer
- Cisco's PIE AQM
- Research into torrent, vpn, and openflow behavior
- RRUL test enhancements
- Hardening further against attacks
- SRC/DST routing
- Following IETF Homenet working group's efforts
- MDNSext?
- What else is needed?

# Go forth and try this stuff out!

CDF of Web Page Load Time under Tested Conditions



## Cablelabs Sim Study

# Acknowledgements

*A big thanks to the bloat mailing list, Jim, Kathie, Van, and Eric, the ISC open home gateway project, the CeroWrt contributors, OpenWrt, Google, the Internet Civil Engineering Institute, and the Comcast Technology Research and Development Fund for their interest and support in the work!*

*And special thanks to:*

**Tom Sawyer!**

*(It's going to be a great fence when it's done...)*



We take paypal:

<http://www.teklibre.com/cerowrt/subscribe.html>