

Bufferbloat

Premature optimization is the root of all evil

- Donald Knuth


What is Bufferbloat?

- An optimization the went wrong...
- <https://www.youtube.com/watch?v=NuHYOu4aAqg>.

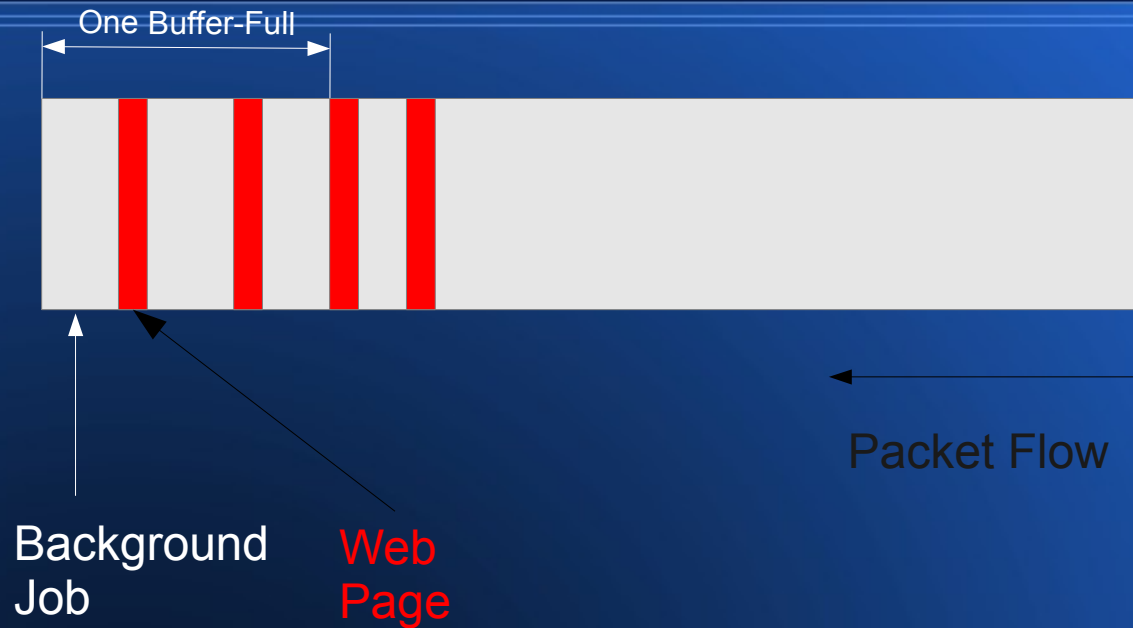
What We Just Saw

- Both routers' networks had a client using up all their bandwidth (http download, ftp, bittorrent, etc)
- Both routers tried to share bandwidth equally among all clients on their net
- The bad router delayed all other traffic by the time it takes to flush an entire buffer, defeating the sharing
- The good one let them share properly

Bad Router



Good Router



How it Happens

- How buffers help
 - Take bursts of traffic
 - Hold data if the link slows
- How they hinder
 - If they're always full, you have to wait from them to empty before putting something in
- Congestion control
 - The sender should send just fast enough to keep busy

TCP Congestion Control

- Sender ramps up until link is congested
 - Constantly increases “send window”
 - Congestion detected when a packet is dropped by IP
- Sender doesn't get the ack, cuts rate
 - Ramps up until another drop happens
- Constantly tries to speed up, but backs off when it can't

How We Fix It

- Make the buffers smaller?
 - Actually no! They are still needed for bursts
- Use Trickle?
 - Semi-manually limit the amount of data sent
 - Yes, but that's an application-level hack
- Drop packets
 - Sounds bad, but it's really good

That Given, What Do We Do?

- For the heaviest flow(s), drop packets until the buffer stays empty
 - This triggers IP congestion control
 - Once that is done, normal fairness mechanisms work.
- Dave Taht, Jim Gettys, Van Jacobsen and friends have a fix
 - fq_codel in CeroWRT
 - Now in the default Linux kernel
 - Migrating to all the other WRTs

What Should We Do Today

- Put fq_codel in every network-capable device
 - A little hard
- Put in limits, so stupid devices can't bloat
 - Wherever you connect to a slower network
- Put fq_codel in front of the stupid device
- Keep putting it everywhere

<http://www.netvibes.com/bufferbloat>

What Comes Next?

- Cell Phones
 - Pessimial case of the previous generation of sw
 - Research problem for fq_codel folks
- Problem Space
 - Multiple towers, constantly changing distances
 - Interference from tall buildings, thunderstorms
 - Constantly changing protocols / backward compatibility
 - Telcos (!)