



Internet Congestion Management

Presented at the 2009 Phoenix Center Annual U.S. Telecoms Symposium
December 3, 2009

Jason Livingood

Comcast
National Engineering & Technical Operations
Internet Systems Engineering

© 2009 Comcast



Overview of Remarks – An Operator's View



- ISP motivations for congestion management
- Comcast's congestion management system
- What the Internet technical community is doing

ISP Motivations for Congestion Management

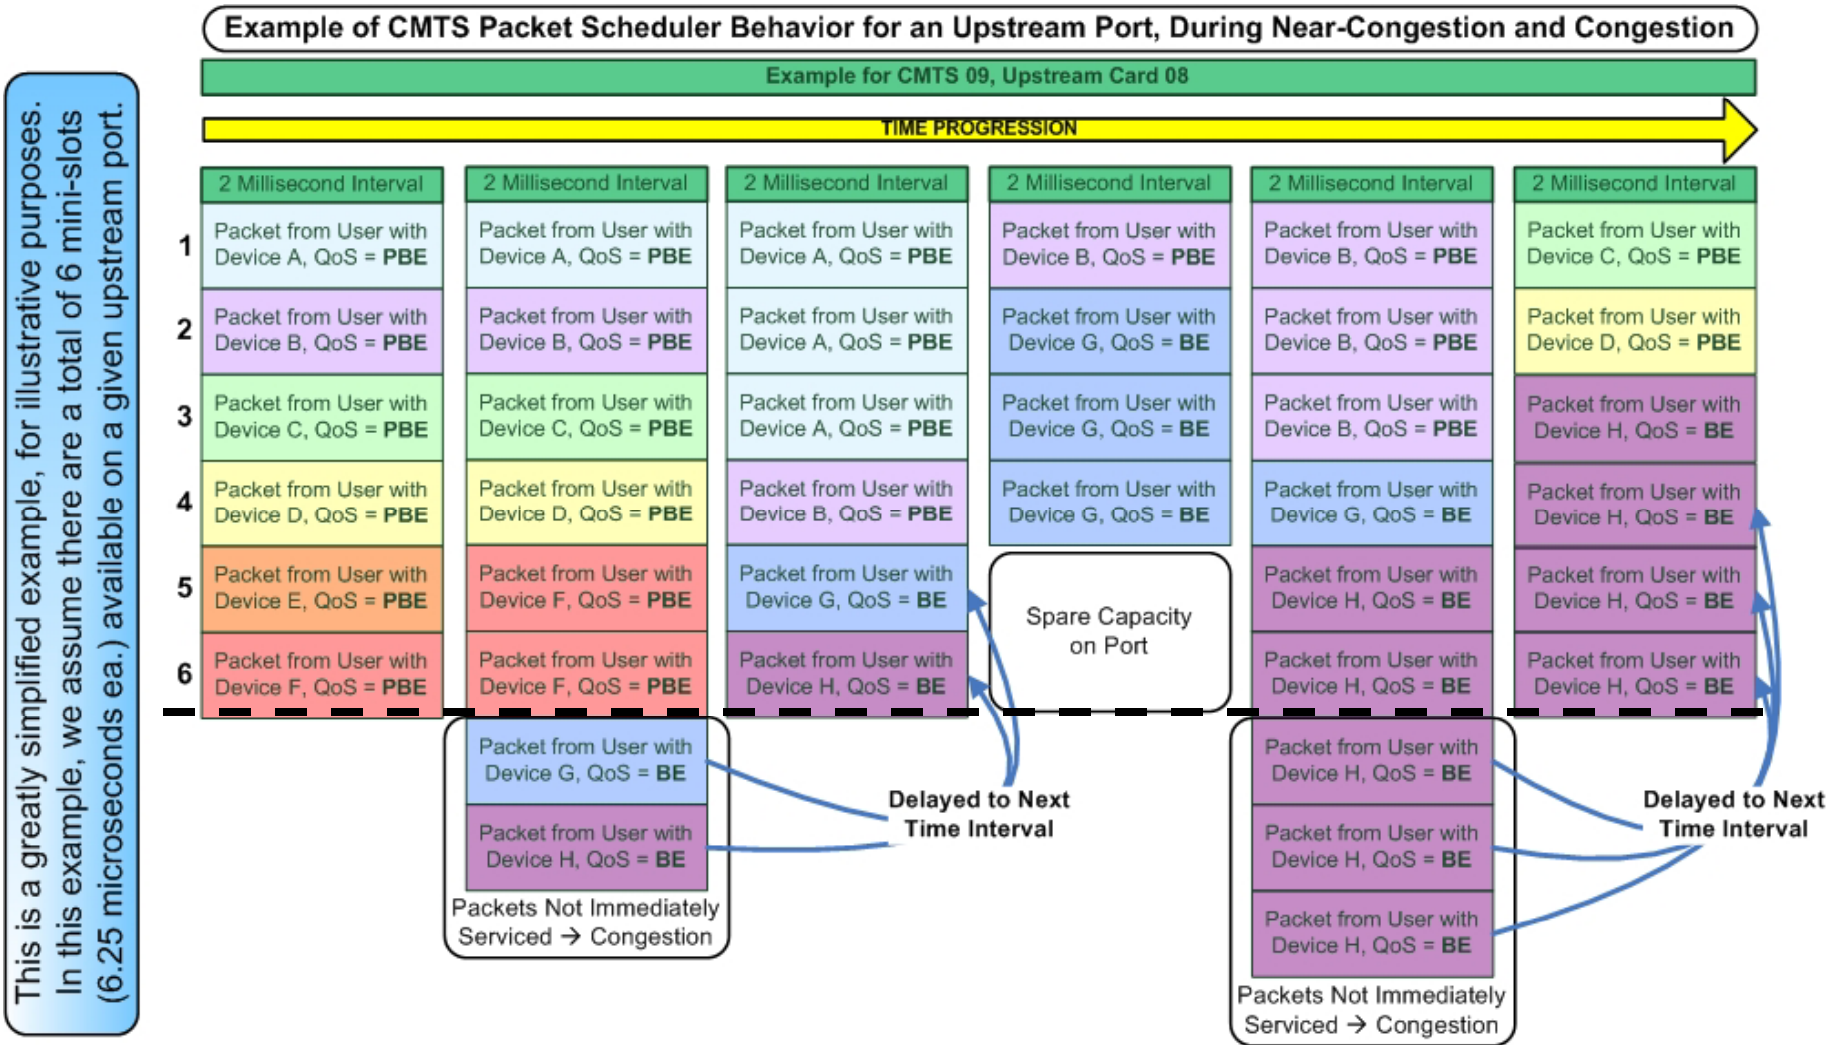
- ISP must be responsive to dissimilar customer application demands
 - Customer care call volume is an obvious indicator of customer dissatisfaction, as well as its own support cost
 - Mix of popular customer applications tends to vary according to demographics, e.g., higher P2P usage in college environments
 - Interactive applications (VoIP, web, streaming video, online gaming) tend to have much stronger diurnal consumption patterns than bulk file distribution (P2P)
- ISP must balance multiple external concerns
 - Internet community, government regulators, different traffic sources & sinks, sustainable business models, etc.
- Network capacity increases are not instantaneous
 - DOCSIS bandwidth augmentation usually requires fiber node splits and CMTS port allocations; it sometimes requires new fiber runs, additional CMTS blades and chassis, and occasionally the allocation of additional RF spectrum
 - Additional access network capacity can be consumed quickly

Comcast Congestion Management

- Concepts behind Comcast congestion management
 - Goal: consistent performance of Internet applications even with heavy traffic, e.g., from P2P file sharing
 - Both “protocol agnostic” and “application agnostic”
 - Acts on *current* network conditions and *recent* user traffic
 - Compatible with Internet standards and supportive of Internet innovation
- Congestion management details
 - Use two different Quality of Service (QoS) levels for best effort traffic over DOCSIS network
 - Priority Best Effort (PBE), which is the default QoS
 - Best Effort (BE)
 - Customer traffic in congestion-managed state is forwarded with a lower priority (BE)
 - Upstream and downstream DOCSIS networks are managed separately
 - Only impact on the traffic of users marked with BE QoS is when congestion actually occurs
- Very few customers (< 1%) impacted by congestion management

Traffic Impact Using Congestion Management

- The CMTS scheduler treats PBE and BE traffic according to priority during congestion. BE traffic marking is determined by recent user traffic volumes.



Summary of ISP Motivations & Comcast System

- The Comcast congestion management solution is effective, is application-neutral, and supports future Internet innovation.
- Comcast is collaborating with the IETF on new protocols as part of future solutions for end-to-end congestion management and related Internet capabilities.
- Which leads me to what the technical community is doing regarding congestion...

Growing Community Consensus on Congestion

- Prior to 2008: ISPs trying to solve congestion on their own
- May 2008: IETF Workshop on Peer to Peer Infrastructure (P2Pi) challenges, held at MIT, with wide Internet community participation
- 2008 – 2009: Relevant Community Activity
 - Numerous IETF BoFs / Working Groups, with significant participation:
 - Low Extra Delay Background Transport (LEDBAT), re. latency-friendly file transfer
 - Application-Layer Traffic Optimization (ALTO), re. P2P localization
 - Congestion Exposure (CONEX), re. E2E congestion signaling
 - De-Coupled Application Data Enroute (DECADE), re. P2P caching
 - Multi-Path TCP (MPTCP), re. using multiple paths to avoid mid-net congestion
 - Broadband Home Gateways (HOMEGATE), re. congestion control extending to home LAN, among other things
 - IETF 75: Technical Plenary, “Network Neutrality; Placing the Debate in IETF Context”
 - IETF 76: Internet Society (ISOC) discussion, “Internet Bandwidth Growth: Dealing with Reality”
 - NANOG 47: “2009 Internet Observatory Report” (data from 110 ISPs)
 - MITAS: MIT Internet Traffic Analysis Study (data from Comcast, TWC, Telus, Rogers, Liberty Global, Clearwire)

Growing Community Consensus on Congestion

- Current congestion control mechanisms in TCP designed in late 1980s when the backbone speeds were roughly what many home users can now buy from their ISP.
- Continuous and persistent congestion, on-going for a long time, can be addressed by adding capacity. Capacity addition is time-consuming.
- Short-lived, momentary congestion requires other remedies:
 - There's a need for the community to focus on this problem space (reference the number of IETF efforts in this area).
 - Best managed via protocol-agnostic methods, as compared to application-targeted methods.
 - You can't "build your way out" of it, and many in the community believe that this is how TCP is designed to behave.
 - Quoting a CONEX BoF presentation: "If some place along the path isn't congested then the transport protocol is doing something wrong."

Thank You



I look forward to Q&A

Feel free to follow-up with me later:
jason_livingood@cable.comcast.com

Jason Livingood
Internet Systems Engineering
National Engineering & Technical Operations