CableLabs

The NQB PHB & DSCP

Greg White | Distinguished Technologist
g.white@cablelabs.com

What is NQB? History/genesis

CableLabs®

- Can we allow low-data-rate, latency-sensitive applications to avoid the latency/jitter/loss caused by bulk traffic?

- Previous attempts at this:
  - Prioritize “Latency-Sensitive” traffic
    - All sorts of policy and system complexity problems[1]
  - Heavy-Hitter Filter[2,3] / Dynamic Packet Prioritization[4]
    - “elephant” vs “mice” flows <- wrong metric
  - “RD”[5] / “Loss-Latency Tradeoff (LoLa)”[6]
    - Secondary shallow queue (low priority and/or high packet loss)

- Inspired by FQ-CoDel’s[7] prioritization of “sparse” flows
  - Low-rate, non-bursty flows get high priority
  - Can we achieve similar results without FQ?
    - i.e. optimized for high-performance/low-cost hardware

1. Claffy, KC. and D. Clark, "Adding Enhanced Services to the Internet: Lessons from History", TPRC , 2015.
2. Estan, C. and G. Varghese, "New directions in traffic measurement and accounting: Focusing on the elephants, ignoring the mice", ACM Transactions on Computer Systems Vol.23, Iss.3, August 2003.
3. Lam, T., "net-qdisc-hhf: Heavy-Hitter Filter (HHF) qdisc", December 2013.
4. Cisco, "Intelligent Buffer Management on Cisco Nexus 9000 Series Switches White Paper", June 2017.
5. Podlesny, M. and S. Gorinsky, "Rd Network Services: Differentiation Through Performance Incentives", SIGCOMM , 2008.
6. Fossati, T., Fairhurst, G., Gutierrez, P., and M. Kuehlewind, "A Loss-Latency Trade-off Signal for the Mobile Network", draft-fossati-tsvwg-lola-00, December 2018.
7. Hoeiland-Joergensen, T., McKenney, P., Taht, D., Gettys, J., and E. Dumazet, "The Flow Queue CoDel Packet Scheduler and Active Queue Management Algorithm", RFC 8290, January 2018.

NQB network behavior

- At network bottlenecks, provide two, equal-priority, “best-effort” queues:
  - One with a deep buffer (Default)
  - One with a shallow buffer (NQB)
- Allow each application to choose* which queue works best for it using a DSCP

Additional requirements/recommendations:

- Strongly recommend “Traffic Protection”
  - Automatic means to monitor for misuse and mitigate it
- In multi-user network gear (e.g. CMTS, BNG), keep users separate

NQB Application Requirements

- Per-microflow requirements:
  - Low-rate:
    - Rate is less than about 1% of “typical” network path capacity
    - e.g. < 500 kbps today
  - Non-bursty:
    - bytes sent (in any time interval T) ≤ R * T + 1 MTU

- If the above are true, microflow can be marked NQB
  - DSCP 45 has been assigned by IANA

Relationship to Diffserv

- Diffserv has traditionally NOT been used on the “public” Internet
  - DSCP is frequently bleached (i.e. cleared to 0) on ingress (~50% of ASes)

- Historically, standardized (IETF) definitions revolved around QoS wants & needs
  - An application uses DSCP to request a certain treatment by routers:
    - Expedited Forwarding, Assured Forwarding, Best Effort, Scavenger
    - Commonly equates to priority
  - Infeasible to support end-to-end
    - Non-verifiable
    - Zero-sum game incentivizes lying, requires control/policing/enforcement/mgmt.

NQB DSCP

- DSCP to express “Non-Queue-Building” application behavior
  - Verifiable behavior
  - No incentive to lie
    - Non-compliant traffic is likely to get worse performance if it mis-marks its traffic as NQB.
  - Traffic protection as a backstop (and to enforce the above)
- Metro/core networks treat NQB traffic no differently from BE (default) user traffic
  - NQB is only differentiated from QB on dual-queue bottleneck links

Relationship to L4S

NQB

NQB is intended only for low-data-rate applications

L4S is intended for high-data-rate applications

- A Dual-Queue implementation could implement support for both NQB and L4S, with both types sharing the “low-latency” queue.
- An application could mark its packets both NQB and L4S(ECT1), if it complies with both sets of prerequisite requirements

Why not use ECT1 instead of a DSCP?

- Standards/Architecture reason:
  - The “Prague L4S requirements” aren’t written for very low-rate flows
    - i.e. “MUST provide CE feedback”, ”MUST be replaceable with Classic CC”, “MUST detect non-L4S ECN”, “MUST converge to RTT independent rate”, etc.
- Pragmatic reasons:
  - Wi-Fi, implementation & adoption complexity

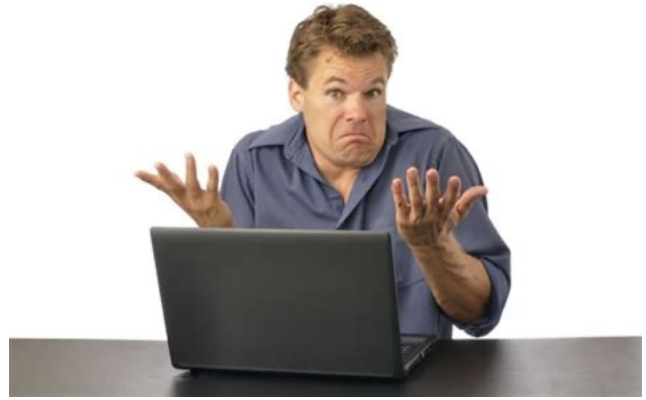

IETF process / current status

- Lengthy review process in the IETF Transport Working Group
  - Many opinions expressed
  - Many revisions to the draft
- Finished Working Group stage, on to final IESG/IETF approvals
  - https://www.ietf.org/archive/id/draft-ietf-tsvwg-nqb-27.html
- DSCP 45 is now listed by IANA as being assigned to NQB.
- Comcast announced DSCP-45 peering in February 2024

Q&A