Non Queue Building (NQB) Per Hop Behavior

draft-ietf-tsvwg-nqb-08

Greg White, CableLabs

Thomas Fossati, ARM

TSVWG @ IETF112

Nov. 12, 2021

Status

- Draft-05 published March 8, 2021
- Draft-06 published July 12
- Draft-07 published July 28
- Draft-08 published Oct 25
- Milestone: Submit as Proposed Standard RFC by April 2022

Changes in draft-08

1. Indicates that it "Updates: rfc8325"
2. Includes a clarification that one of the SHOULD statements applies to RFC8325 equipment
3. Moves the text on safeguards for legacy WiFi up into a new subsection that is generalized to also cover unmanaged networks that happen to implement IP Precedence.
4. Minor wordsmithing

List discussion – guidance on sender rates

Current Text

§4.1 Non-Queue-Building Sender Requirements

Non-queue-building (NQB) flows are typically UDP flows that don't seek the maximum capacity of the link (examples: online games, voice chat, DNS lookups, real-time IoT analytics data). Here the data rate is limited by the application itself rather than by network capacity - these applications send, at most, the equivalent of a few well-spaced packets per RTT, even if the packets are not actually RTT-clocked. In today's network this corresponds to an instantaneous data rate (packet size divided by packet inter-arrival time) of no more than about 1 Mbps (e.g. no more than one 1250 B packet every 10 ms), but there is no precise bound since it depends on the conditions in which the application is operating.

Note that, while such flows ordinarily don't implement a traditional congestion control mechanism, they nonetheless are expected to comply with existing guidance for safe deployment on the Internet, for example the requirements in [RFC8085] and Section 2 of [RFC3551] (also see the circuit breaker limits in Section 4.3 of [RFC8083] and the description of inelastic pseudowires in Section 4 of [RFC7893]). To be clear, the description of NQB flows in this document should not be interpreted as suggesting that such flows are in any way exempt from this responsibility.

Applications that align with the description of NQB behavior in the preceding paragraphs SHOULD identify themselves to the network using a Diffserv Code Point (DSCP) of 45 (decimal) so that their packets can be queued separately from QB flows. The choice of the value 45 is motivated in part by the desire to achieve separate queuing in existing WiFi networks (see Section 8.3). In networks where another (e.g. a local-use) codepoint is designated for NQB traffic, or where specialized PHBs are available that can meet specific application requirements (e.g. a guaranteed-latency path for voice traffic), it may be preferred to use another DSCP.

If the application's traffic exceeds more than a few packets per RTT, or exceeds approximately 1 Mbps on an instantaneous (inter-packet) basis, the application SHOULD NOT mark its traffic with the NQB DSCP. In such a case, the application has to instead implement a relevant congestion control mechanism, for example as described in Section 3.1 of [RFC8085] or [I-D.ietf-tsvwg-ecn-l4s-id].

Describes NQB flows as sending at most a few well-spaced packets per RTT, e.g. no more than one 1250B packet every 10ms (1 Mbps)

NQB flows still need to be “safe” per existing RFC/BCP guidance

Recommendation to NOT use NQB if instantaneous rate exceeds a limit: ~1Mbps

List discussion (slide 2 of 3)

- Sender Requirements
  - Current text:

"If the application's traffic exceeds more than a few packets per RTT, or exceeds approximately 1 Mbps on an instantaneous (inter-packet) basis, the application SHOULD NOT mark its traffic with the NQB DSCP."

  - Suggested replacement:

"If the application's traffic exceeds more than a few packets per RTT, or on an instantaneous (inter-packet) basis exceeds 10% of the global average access link capacity at the time, the application SHOULD NOT mark its traffic with the NQB DSCP. At the time of writing this document, the global average access link capacity is 63 Mbps down and 13 Mbps up for mobile networks, 113 Mbps down and 62 Mbps up for fixed networks [https://www.speedtest.net/global-index]. For a typical server application, this implies a 6.3 Mbps maximum instantaneous rate. For a typical client application, this implies a 1.3 Mbps maximum instantaneous rate."

Questions for WG:

1) Is this formulation better?
2) If so, what percentage?

List discussion (slide 3 of 3)

- Additional items if text on previous slide is agreed:
  - Add example of when the RTT limit is a tighter bound.
    - I'll probably use "a few" = 5.
  - Add a stronger statement on the importance of Traffic Protection on low rate (less than 10% of global average) links that support NQB, with and without L4S.
- Strengthen language around Traffic Protection algorithms
  - … should operate on a per-flow basis
  - … should be based on *actual* queuing, not just arrival rate

Question from IANA

- Should both NQB DSCPs be named “Non-Queue-Building”?
- Maybe:
  - 5: “Non-Queue-Building (interconnection)”
  - 45: “Non-Queue-Building”

Next Steps

- Finalize changes arising from list discussion
- WGLC?