

Non Queue Building (NQB) Per Hop Behavior

[draft-ietf-tsvwg-nqb-09](#)


Greg White, CableLabs

Thomas Fossati, ARM

TSVWG Interim 1 - 2022

Feb. 16, 2022

Status

- Draft-05 published March 8, 2021
 - Draft-06 published July 12
 - Draft-07 published July 28
 - Draft-08 published Oct 25
 - Draft-09 published Feb 11, 2022
 - Milestone: Submit as Proposed Standard RFC by April 2022
- 
- September

Changes in draft-09 based on list discussion

1. Updated recommendation and informational text in sender requirements section
2. Recommendation NOT to use 45 across interconnects
3. New section 5.3 on Guidance for Very Low Rate Links
4. Updated text on EDCA configuration for Wi-Fi

§4.1 Sender Requirements

Guidance on Sender Rates

- Previous recommendation:

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.

- Previous list discussion suggested changing to, e.g.

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.

- List comments pointed to issues with using global average access link capacity.

- Revised recommendation in draft-09:

If the application's traffic exceeds more than a few packets per RTT, or exceeds **the expected path capacity** on an instantaneous (e.g. inter-packet or a suitably short time interval) basis, the application SHOULD NOT mark its traffic with the NQB DSCP.

Continued on next slide....

§4.1 Sender Requirements

Guidance on Sender Rates

- New informational text

At the time of writing, it is believed that 1 Mbps is a reasonable expectation of path capacity, but this value is of course subject to the context in which the application is expected to be deployed.

An application that marks its traffic as NQB but happens to exceed the available path capacity (even on an instantaneous basis) runs the risk of being subjected to a Traffic Protection algorithm (see [Section 5.2](#)), which could result in the excess traffic being discarded or queued separately as default traffic (and thus potentially delivered out of order). As a result, applications that aren't clearly beneath the threshold described above would need to weigh the risk of additional loss or out-of-order delivery against the expected latency benefits of NQB treatment in determining whether or not to mark their packets as NQB.

Moved mention of 1 Mbps to informational text

Added explanation of the risks that an app developer should consider if their app is borderline.

Recommendation NOT to use 45 across interconnects

- Previous text:

To facilitate the default treatment of NQB traffic in backbones and core networks discussed in the previous section (where IP Precedence may be deployed), networks that support NQB SHOULD remap NQB traffic (DSCP 45) to DSCP 5 prior to interconnection, unless agreed otherwise between the interconnecting partners.

- New text:

To facilitate the default treatment of NQB traffic in backbones and core networks discussed in the previous section (where IP Precedence may be deployed), networks that support NQB SHOULD NOT use the value 45 for NQB at network interconnects unless that usage is explicitly documented in the TCA (Traffic Conditioning Agreement, see [[RFC2475](#)]) for that interconnection. Rather, networks SHOULD remap NQB traffic to DSCP 5 prior to interconnection, unless agreed otherwise between the interconnecting partners.

...also swapped the order of the next 2 sentences in that paragraph

New §5.3: Guidance for Very Low Rate Links

The NQB sender requirements in [Section 4.1](#) place responsibility in the hands of the application developer to determine the likelihood that the application's sending behavior could result in a queue forming along the path. These requirements rely on application developers having a reasonable sense for the network context in which their application is to be deployed. Even so, there will undoubtedly be networks that contain links having a data rate that is below what is considered "typical", and even potentially below the instantaneous sending rate of certain NQB-marked applications.

To limit the consequences of this scenario, operators of such networks SHOULD utilize a traffic protection function that is more tolerant of burstiness (i.e. a temporary queue). Alternatively, operators of such networks MAY choose to disable NQB support on these low speed links. In particular, for links that are far below "typical" path rates, it is RECOMMENDED that NQB support be disabled.

Updated §8.3.1 on Interoperability with existing WiFi networks

- **Previous text:**

In order to preserve the incentives principle for NQB, WiFi systems SHOULD **configure** the EDCA parameters for the Video Access Category to match those of the Best Effort Access Category.

- **New text:**

In order to preserve the incentives principle for NQB, WiFi systems SHOULD **be configured such that** the EDCA parameters for the Video Access Category ~~to~~ match those of the Best Effort Access Category. **This recommendation is presumably most applicable to WiFi systems deployed in a managed environment or those deployed by an ISP. In other situations (e.g. consumer-grade WiFi gear deployed by an ISP's customer) this configuration may not be possible, and the requirements and recommendations in Section 4.3.1 would apply.**

Next Steps

- Remaining list suggestions
 - In sender reqs: Add example of when the RTT limit is a tighter bound
 - In Traffic Protection, add guidance:
 - ... should be based on **actual** queuing, not just arrival rate
- Other comments prior to WGLC?