BBRv3: Internet Draft Update

draft-cardwell-ccwg-bbr-00

Internet Draft Editors:

Neal Cardwell (Google), Ian Swett (Google), Joseph Beshay (Meta)

Speaker: Neal Cardwell

Outline

Overview

- Outline BBR Internet Draft updates

Goals for this talk:

- Provide a road map for...
  - Readers of the draft
  - Implementers of BBRv3 reading the draft
- Inviting the community to…
  - Read the draft and offer comments or edits (github pull requests)
- Gauging interest in making this a working group item

Overview of draft-cardwell-ccwg-bbr-00

- New BBR draft: draft-cardwell-ccwg-bbr
- Editors:
  - Neal Cardwell (Google)
  - Ian Swett (Google)
  - Joseph Beshay (Meta)
- High-level changes since draft-cardwell-iccrg-bbr-congestion-control-02:
  - Moved draft to github to ease IETF collaboration:
    - https://github.com/ietf-wg-ccwg/draft-cardwell-ccwg-bbr
  - Incorporated delivery rate draft (draft-cheng-iccrg-delivery-rate-estimation)
    - So all BBR material is in 1 draft, rather than split across 2 drafts [git commit]
  - Updated to reflect changes between BBRv2 and BBRv3
    - Most changes in BBRv3 were discussed at IETF 117 (SF) CCWG [slides]
    - BBRv3 changes are git commits with "BBRv3" in the commit title [git log]
  - Editorial wordsmithing for clarity and readability, updating references

BBRv3 changes in draft-cardwell-ccwg-bbr-00

- BBRv3 Changes already discussed at IETF 117 (SF) CCWG [slides]
  - Bug fix 1 [slide]: fix bandwidth convergence after loss
    - update BBR.full_bw logic to the more general BBRv3 logic
    - BBRv3 bug fix: avoid getting stuck with low throughput after loss/ECN
    - BBRv3 update: update prose for ProbeBW_UP logic to match BBRIsTimeToGoDown() pseudocode
  - Bug fix 2 [slide]: fix bandwidth convergence without loss
    - BBRv3 bug fix: update ProbeBW_UP cwnd_gain to 2.25
    - BBRv3 bug fix: use ProbeBW_DOWN pacing_gain of 0.9 instead of 0.75
  - Tuning [slide]: various parameter tunings to reduce loss without reducing throughput
    - BBRv3 tuning: Startup: cwnd_gain = 2.0, pacing_gain = 2.77
    - BBRv3 tuning: set BBR.inflight_hi upon exiting Startup due to loss
    - BBRv3 tuning: BBRStartupFullLossCnt=6

BBRv3 changes in draft-cardwell-ccwg-bbr-00

- Other misc minor changes; the primary changes to substance are:
  - BBRv3 change: remove BBR.packet_conservation, "packet conservation" behavior
  - BBRv3 change: make BBR.send_quantum floor 2 * SMSS at any bandwidth
  - BBRv3 change: in BBRUpdateACKAggregation, in Startup use filter of 1 round trip
  - BBRv3: clarify pacing_gain in Drain is still 0.35 and no longer 1/BBRStartupCwndGain
  - remove bw_hi
  - fix bug in UpdateRateSample() logic to decide if a packet is the newest
  - fix: add some missing pseudocode for BBR.loss_round_delivered

BBR deployment status at Google

- Google-internal traffic:
  - **BBRv3** is TCP congestion control for all internal **WAN traffic**
  - **BBR.Swift** is TCP congestion control used **within a datacenter**
- Google-external traffic:
  - **BBRv3** is TCP CC for all **Google.com** and **YouTube** public Internet traffic
  - A/B experiments: BBRv3 vs v1 for small % of users for:
    - QUIC for google.com and YouTube

Conclusion

- Inviting the community to read the draft and offer comments/edits, however you prefer:
  - CCWG mailing list emails
  - github issues [link]
  - github pull requests [link]
- Gauging interest in making this a working group item
- Thanks!