

BBR Updates:

Internal Deployment, Code, Draft Plans

TCP BBR: Neal Cardwell, Yuchung Cheng, Kevin Yang

Soheil Hassas Yeganeh, Priyaranjan Jha, Yousuk Seung, Luke Hsiao, Matt Mathis

Van Jacobson

QUIC BBR: Ian Swett, Bin Wu, Victor Vasiliev

<https://groups.google.com/d/forum/bbr-dev>

Outline

- BBRv2 deployment status at Google: nearing completion for internal TCP traffic
- BBRv2 alpha open source release updates
- BBRv2 Internet Draft plans
- BBR.Swift status

Target for this talk:

- Sharing our experience with experiments
- Inviting the community share feedback, test results, issues, patches, or ideas

Update on BBR v2 deployment status at Google

- Google-internal traffic:
 - BBRv2 being deployed as default TCP congestion control for internal Google traffic
 - Used as the congestion control for **>98%** of internal TCP traffic at Google
 - Currently using bandwidth * min_rtt, ECN, loss as signals
 - Seeing reductions in tail latency for RPC traffic
- Google-external traffic: YouTube, google.com: deployed for a small percentage of users
 - Reduced queuing delays: RTTs lower than BBR v1 and CUBIC
 - Reduced packet loss: loss rates closer to CUBIC than BBR v1
- Continuing to iterate using production experiments and lab tests

Status of BBR v2 algorithm and code

- TCP BBRv2 "alpha/preview" release:
 - Linux TCP (dual GPLv2/BSD): github.com/google/bbr/blob/v2alpha/README.md
 - Recent updates:
 - Rebased to Linux v5.10
 - Various minor bug fixes
- QUIC BBR v2 "alpha/preview" release:
 - Chromium QUIC (BSD): on chromium.org in bbr2_sender. { [cc](#), [h](#) }
- BBR v2 alpha release is ready for research experiments; we invite:
 - Ideas for test cases and metrics to evaluate
 - Test results and traces
 - Algorithm/code ideas
 - Patches
- BBR v2 algorithm: IETF 104 [[slides](#) | [video](#)]; alpha code: IETF 105 [[slides](#) | [video](#)]

BBR v2: Internet Draft plans

- Current BBR Internet Drafts cover BBRv1:
 - Delivery rate estimation: [draft-cheng-iccr-g-delivery-rate-estimation](#)
 - BBRv1 congestion control: [draft-cardwell-iccr-g-bbr-congestion-control](#)
- Planning on updating drafts to reflect BBRv2 for July IETF

BBR.Swift: Status and Plans

- BBR.Swift
 - Leverages approaches from Swift [[SIGCOMM-2020](#)]
 - Uses Network_RTT as primary congestion signal
 - Motivation: Provides richer information about the current degree of queuing
 - Allows faster reaction to long queues
 - Avoids overreaction and underutilization with short queues
- Status: Preparing for production testing of BBR.Swift for Google internal traffic
- Plan: deploy, release BBR.Swift code as open source, document the algorithm in detail
 - Including implementation of [draft-yang-tcpm-ets-00](#)
- Goal: we want transports to be able to use BBR.Swift as their CC
 - On connections where...
 - Target Network_RTT is known
 - It is known that other traffic sharing bottlenecks is using Swift or BBR.Swift
 - On physical machines or virtual machines

Conclusion

- Actively working on BBR v2, BBR.Swift at Google
 - Finishing roll-out for internal TCP traffic
 - Tuning performance to enable roll-out for external Google traffic
 - Improving the algorithm to scale to larger numbers of flows
 - Planning on BBRv2 Internet Drafts for July IETF
 - We invite the community share test results, issues, patches, or ideas

<https://groups.google.com/d/forum/bbr-dev>

Internet Drafts, paper, code, mailing list, talks, etc.

Special thanks to Eric Dumazet, Nandita Dukkipati, C. Stephen Gunn, Jana Iyengar, Pawel Jurczyk, Biren Roy, David Wetherall, Amin Vahdat, Leonidas Kontothanassis, and {YouTube, google.com, SRE, BWE} teams.