

NAME

tcp_bbr - TCP Bottleneck Bandwidth and Round-Trip Time Algorithm

SYNOPSIS

To load the driver as a module at boot time, place the following line in loader.conf(5):

```
tcp_bbr_load="YES"
```

To enable the TCP stack you must place the following line in the sysctl.conf(5):

```
net.inet.tcp.functions_default=bbr
```

DESCRIPTION

Bottleneck bandwidth and round-trip time (BBR) is a congestion control algorithm which seeks high throughput with a small queue by probing BW and RTT. It is a round-up redesign of congestion control, which is not loss-based, delay-based, ECN-based or AIMD-based.

The core design of BBR is about creating a model graph of the network path by estimating the maximum BW and minimum RTT on each ACK.

MIB Variables

The algorithm exposes the following scopes in the *net.inet.tcp.bbr* branch of the sysctl(3) MIB:

<i>cwnd</i>	Cwnd controls, for example "target cwnd rtt measurement" and "BBR initial window".
<i>measure</i>	Measurement controls.
<i>pacing</i>	Connection pacing controls.
<i>policer</i>	Policer controls, for example "false detection threshold" and "loss threshold".
<i>probertt</i>	Probe RTT controls.
<i>startup</i>	Startup controls.
<i>states</i>	State controls.
<i>timeout</i>	Time out controls.

Besides the variables within the above scopes the following variables are also exposed in the *net.inet.tcp.bbr* branch:

<i>clrlost</i>	Clear lost counters.
<i>software_pacing</i>	Total number of software paced flows.
<i>hdwr_pacing</i>	Total number of hardware paced flows.
<i>enob_no_hdwr_pacing</i>	Total number of enobufs for non-hardware paced flows.
<i>enob_hdwr_pacing</i>	Total number of enobufs for hardware paced flows.
<i>rtt_tlp_thresh</i>	What divisor for TLP rtt/retran will be added (1=rtt, 2=1/2 rtt etc).
<i>reorder_fade</i>	Does reorder detection fade, if so how many ms (0 means never).
<i>reorder_thresh</i>	What factor for rack will be added when seeing reordering (shift right).
<i>bb_verbose</i>	Should BBR black box logging be verbose.
<i>sblklimit</i>	When do we start ignoring small sack blocks.
<i>resend_use_tso</i>	Can resends use TSO?
<i>data_after_close</i>	Do we hold off sending a RST until all pending data is ack'd.
<i>kill_paceout</i>	When we hit this many errors in a row, kill the session?
<i>error_paceout</i>	When we hit an error what is the min to pace out in usec's?
<i>cheat_rxt</i>	Do we burst 1ms between sends on retransmissions (like rack)?
<i>minrto</i>	Minimum RTO in ms.

SEE ALSO

cc_chd(4), cc_cubic(4), cc_hd(4), cc_htcp(4), cc_newreno(4), cc_vegas(4), h_ertt(4), mod_cc(4), tcp(4), tcp_rack(4), mod_cc(9)

Neal Cardwell, Yuchung Cheng, Stephen Gunn, Soheil Hassas Yeganeh, and Van Jacobson, "BBR: Congestion-Based Congestion Control", *ACM Queue*, Vol. 14, September / October 2016.

Dominik Scholz, Benedikt Jaeger, Lukas Schwaighofer, Daniel Raumer, Fabien Geyer, and Georg Carle, "Towards a Deeper Understanding of TCP BBR Congestion Control", *IFIP Networking 2018*, <http://www.net.in.tum.de/fileadmin/bibtex/publications/papers/IFIP-Networking-2018-TCP-BBR.pdf>, May 2018.

HISTORY

The **tcp_bbr** congestion control module first appeared in FreeBSD 13.0.

AUTHORS

The **tcp_bbr** congestion control module was written by Randall Stewart <rrs@FreeBSD.org> and sponsored by Netflix, Inc. This manual page was written by Gordon Bergling <gbe@FreeBSD.org>.