TCB sharing: RFC 2140 vs. reality

Michael Welzl, **Safiqul Islam**, Joe Touch, Jianjie You

November 5, 2015

Why?

- Overview of the state of implementation
  - Reality vs Standard
- Update the standard
  - Coherent with the current implementations
  - For future developers?
  - Clarity?

RFC 2140: TCP Control Block Interdependence

- TCB sharing
  - Temporal vs. Ensemble Sharing
  - Only for initialization
- Suggested caching and sharing:
  Old_MSS, Old-RTT, Old-RTTvar, Old-snd_cwnd

Implementation - Reality vs Standard

L: Linux, F: FreeBSD

RFC 2140	**Description**	**Implementation**	**Status**
Old-MSS	Maximum Segment Size	F:rmx_mtu	Cached and shared
Old-RTT	Estimated Round-Trip Time	L:TCP_METRIC_RTT F:rmx_rtt	Cached (FreeBSD and Linux) Shared (FreeBSD)
Old-RTTvar	Estimated RTT Variance	L:TCP_METRIC_RTTVAR F:rmx_rttvar	Cached (FreeBSD and Linux) Shared (FreeBSD)
Old-snd_cwnd	Congestion Window	L:TCP_METRIC_CWND F:rmx_cwnd	Cached (Both) Not shared

Implementation - Reality vs Standard (Contd.)

L: Linux, F: FreeBSD

RFC 2140	Description	Implementation	Status
NA	Slow Start Thresold	L:TCP_METRIC_SSTHRESH F:rmx_ssthresh	Cached and shared
NA	Metric related to the extent of reordering	L:TCP_METRIC_REORDERING	Cached and shared (Linux)
NA	Estimated Bandwidth	F:rmx_bandwidth	Not in the TCP spec, and not set before cached
NA	Outbound Delay-Bandwidth Product	F:rmx_sendpipe	Not in the TCP spec, and not set before cached
NA	Inbound Delay-Bandwidth Product	F:rmx_recvpipe	Not in the TCP spec, and not set before cached

Q&A