

TCB Control Block Sharing: 2140bis

draft-ietf-tcpm-2140bis-00

IETF 105 - Montreal



← Joe Touch, consultant
Michael Welzl, U. Oslo →
← Safiqul Islam, U. Oslo



Changes from draft-touch-tcpm-2140bis-06

- Re-issued as draft-ietf-tcpm-2140bis due to WG adoption
- Cleaned orphan references to T/TCP, removed incomplete refs
- Moved references to informative section and updated Sec 2
- Updated to clarify “no impact to interoperability”
- Updated appendix B to avoid 2119 language

Next steps

Currently, document discusses:

1. Ensemble sharing

- Sharing (initializing from cache) while connections are ongoing

2. Temporal sharing

- Making cached TCB information available to later connections

- ...but what is "later"? Not much discussion of timescales...
- ...and what defines success? Currently, TCP convergence resolves it

→ We plan to add a short section on these issues.

Timescale considerations

- How long is TCB information valid? Depends...
- E.g.:
 - **RTTVar** reflects current traffic behavior
 - **TFO state** reflects state from the peer + path state that may last longer
 - Subject to routing changes and equipment configuration / upgrades
- Very long lasting state may benefit from special consideration
 - E.g., some constants that never work may need to be adjusted

Auto-adjusting over long timescales (think, e.g.: months, years)

- **Consider IW**: we now have Exp IW 10 [RFC 6928]
 - Some hosts use much larger values
 - Some use smaller ones
- Could we be a little bit more dynamic?
 - If a large IW to one destination never worked, it's pointless to keep?
 - If `snd_cwnd` always became much larger to a destination, why limit to 10 ?
- This would automatically do an 'experiment', and auto-adjust to environment
 - I.e., as initially suggested in draft-touch-tcpm-automatic-iw
 - A node on a slow connection or with a small device (IoT) might never increase its IW
 - My office machine (network always upgraded with the rest of the backbone) might

Types of feedback

- Current:
 - Convergence-based
 - Connection end values affect future predictions by how statistics aggregate
- Additional opportunities (also as noted in draft-touch-tcpm-automatic-iw)
 - More direct management of those statistics
 - E.g., discounting values that didn't work out, amplifying those that do
- Two kinds of additional feedback :
 - Implicit
 - Tracking whether initial conditions persist over a connection
 - Explicit
 - looking for packet exchanges that more directly indicate success