

RDDP: TCP Framing

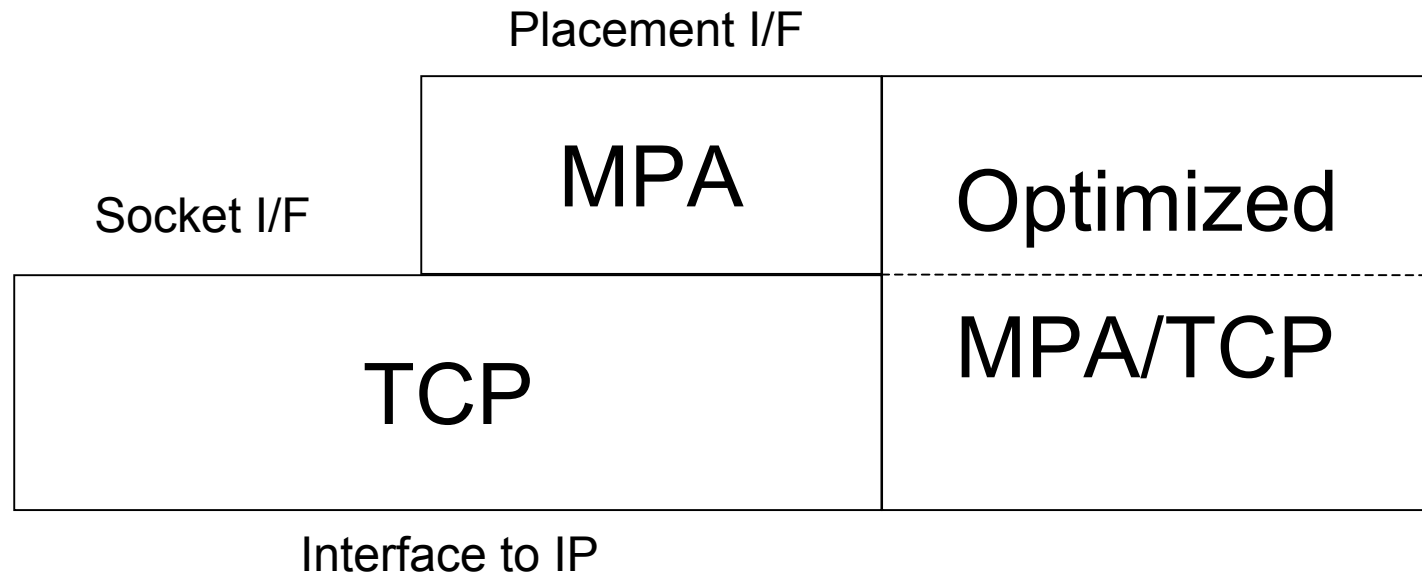
draft-ietf-rddp-mpa-05.txt

David L. Black
(rddp WG chair)

RDDP (Partial) Overview

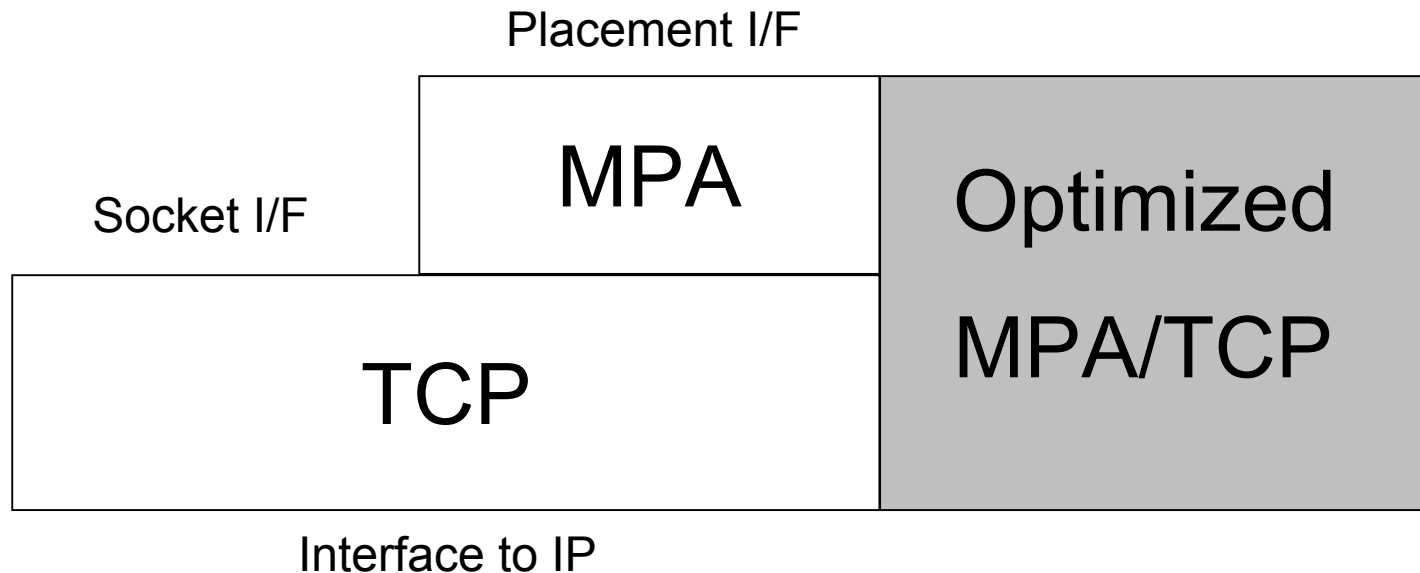
- Function: Remote Direct Data Placement (RDDP)
 - E.g., 4k filesystem block needs to land in 4k aligned pre-allocated buffer, even if it shows up as 6 packets.
- Goal: Aligned zero-copy placement at receiver
- Method: Framed payloads
 - Optimized receiver implementation looks at frame header to select receive buffer
- MPA draft defines framing for TCP
 - TCPM review concern: RDDP is not supposed to change the TCP protocol (RDDP WG agrees)
- All other concerns believed to be under control
 - E.g., allow overwrite of partially received data

MPA and TCP: The Big Picture



- Zero copy requires optimizing MPA with TCP
 - Socket interface is optimized away
 - Socket interface remains available to other applications
 - No change to other interfaces
- MPA draft specifies “optimized” socket i/f (oops)

MPA and TCP: Proposal



- Specify MPA as strictly layered on TCP
 - Straightforward implementation copies in receive path
- Explain zero-copy optimizations in Appendix
 - No normative requirements at MPA/TCP interface
 - Optimized MPA/TCP interface is private to MPA

Questions?