

Multipath TCP for FreeBSD: WG Update

Nigel Williams, Lawrence Stewart, Grenville Armitage
Swinburne University of Technology

Background

- Started implementation around May/June 2012
 - Based on FreeBSD-HEAD
- Project goals
 - Create flexible platform for research
 - Use for congestion control experiments
- Developers: Nigel Williams and Lawrence Stewart
- Project page: <http://caia.swin.edu.au/newtcp/mptcp/>

Current Status

- Version 0.1 released in March this year
- Currently at Version 0.3
 - Completed most of the big architectural changes
 - <http://caia.swin.edu.au/newtcp/mptcp/tools.html>
 - Technical Report: High-level design overview
 - <http://caia.swin.edu.au/reports/130424A/CAIA-TR-130424A.pdf>
- Version 0.4 to be released later this week.
 - Greatly improved stability (though still 'alpha')
 - Can now focus more on becoming feature complete

Working

- Compatible with standard TCP implementations
- FreeBSD <-> FreeBSD multi-subflow MPTCP connections
- MP_CAPABLE, ADD_ADDR, MP_JOIN, DSS options
- Socket isolated from individual subflows
 - Basic scheduler; Send buffer to subflow mapping
 - Multi-packet, per-packet DS-Maps

Working

- Modified reassembly logic and data structures
 - Rx queues merged with receive buffer
- Fast Recovery, SACK, RTOs, window probes (with multiple subflows), timeouts.
- Static path manager (via sysctl)
 - Specify 'slave' subflows
- Clean connection teardown (sans DFIN)

Upcoming

- Sender side
 - Hooks for congestion control, module to use them
 - Improved algorithm for send socket-to-subflow mapping
- Receiver side
 - Tweaks to receive window advertisements
- DFIN, Data-level RTO
- Coupled CC
- Profiling/performance testing (stack changes have been extensive...)
- ISO image

Known Issues

- IPv4 only
- Basic packet scheduler
 - Just tries to distribute bytes between subflows
 - Subflows can starve if application does not fill send buffer
- No fallback
- Still Alpha, not performance optimised
- Need testers, feedback

Projects

- MPTCP for persistent Vehicle to Infrastructure (V2I) connections
 - Delay-tolerant 'Infotainment' applications
 - Embedded On Board Unit uses a combination of 3/4G and 802.11p
 - MAC layer feedback to aid scheduling decisions