

Processing of RST segments by Multipath TCP

Olivier Bonaventure

Christoph Paasch

Gregory Detal

RST with regular TCP

- Transmission of a RST segment
 - Sending host terminates abruptly the connection and **does not have state anymore**
- Reception of a RST segment
 - Remote host **does not have state** for this connection anymore and it must be terminated

RST with Multipath TCP

- Transmission of a RST segment
 - The subflow over which the RST is sent is terminated
 - No segment can be sent/received over **this** subflow
 - The sender of the RST segment **still maintains state** for the Multipath TCP connection that remains active
 - How should the remote host react upon reception of this RST ?

RST with Multipath TCP

- Reception of a RST segment
 - The subflow over which the RST is received is terminated
 - No segment can be sent/received over **this** subflow
 - The receiver of the RST segment **still maintains state** for the Multipath TCP connection that remains active
 - How should the receiver react beyond terminating the subflow ?
 - Restart a subflow over the same interface ?
 - Terminate other subflows ?

Proposed RST option

- Objective
 - Provide information on the reason for the RST
- Reasons
 - Lack of resources
 - Sending host is running out of memory
 - Administratively prohibited
 - Response to SYN received with invalid port or on invalid interface
 - Too many already acknowledged data
 - See discussion with lossy interfaces in experience draft

Proposed RST option (2)

- Reasons
 - Unacceptable performance
 - Subflow is too slow or has too many retransmissions
 - Lifetime expired
 - Subflow has been active for too long
 - Removed address
 - Address used by subflow has been removed
 - Middlebox interference
 - Multipath TCP specific error, Unspecified TCP error
 - Fast Close

Middleboxes

- RST generated by a middlebox
 - Without a RST option
 - would lead to termination of subflow
 - With a proposed RST option
 - could lead the host to react in a specific way
 - But on-path middleboxes can already change window and other fields of the TCP header

Are there alternatives to a new option ?

- RFC793 allows the transmission of data in the RST to indicate a reason
 - Opaque and unstructured ASCII data
 - Seems to have been used by some middleboxes
- Do nothing
 - Will need to define complex heuristics in MPTCP implementations to determine why a subflow has been terminated
 - Server needs to be able to terminate a subflow and “expect” that the client will recreate it

Conclusion

- RST option can aid Multipath TCP implementations by providing info on reason for subflow termination
 - Very useful when server needs to terminate a subflow and only clients create subflows
- Small extension to the protocol that would be very useful for subflow path managers