

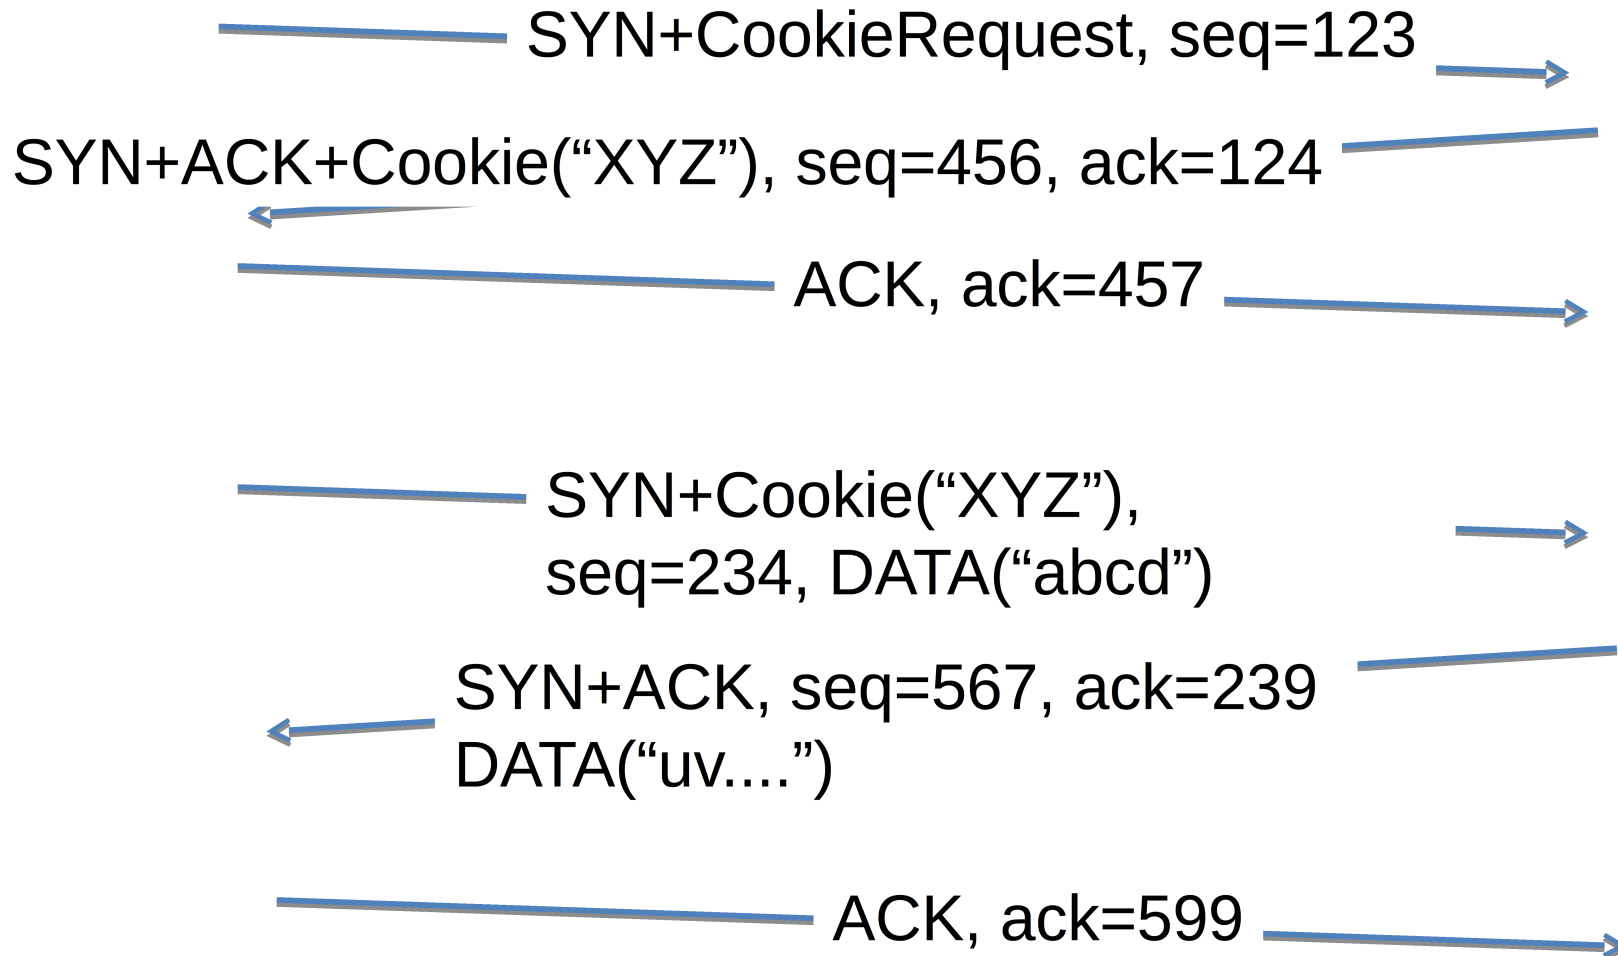
TFO support for Multipath TCP

Sébastien Barré

Grégory Detal

Olivier Bonaventure

TFO in a nutshell



Issue #1

Space in the SYN segments

- Initial handshake : SYN
 - Cookie request consumes 2 bytes
 - MP_CAPABLE consumes 12 bytes
- SYN+ACK
 - MP_CAPABLE consumes 12 bytes
 - Cookie consumes 6-18 bytes
 - Server needs a short cookie to support RFC1323s and other options
- Subsequent handshake
 - SYN has same length as SYN+ACK above
 - SYN+ACK uses fewer options

Issue #2

Data sequence mapping

- When does the DSS mapping starts
 - Does IDSN starts with SYN ?

SYN+Cookie("XYZ")+MP_CAPABLE
seq=234, DATA("abcd") →

← SYN+ACK+MP_CAPABLE
seq=567, ack=239 DATA("uv....")

ACK MP_CAPABLE, ack=599 →

seq=240, DSS[5->240,len=4] →
DATA("efgh")

Issue #2

Data sequence mapping

- When does the DSS mapping starts
 - Does IDSN starts with SYN ?

SYN+Cookie("XYZ")+MP_CAPABLE
seq=234, DATA("abcd") →

← SYN+ACK+MP_CAPABLE
seq=567, ack=239 DATA("abcd")

What if middlebox
changes/removes data
inside payload ?

ACK MP_CAPABLE, ack=599 →

seq=240, DSS[5->240,len=4]
DATA("efgh") →

Issue #2

Data sequence mapping

- When does the DSS mapping starts ?
 - IDSN must start with first byte in data segment

SYN+Cookie("XYZ")+MP_CAPABLE
seq=234, DATA("abcd") →

← SYN+ACK+MP_CAPABLE
seq=567, ack=239 DATA("uv....")

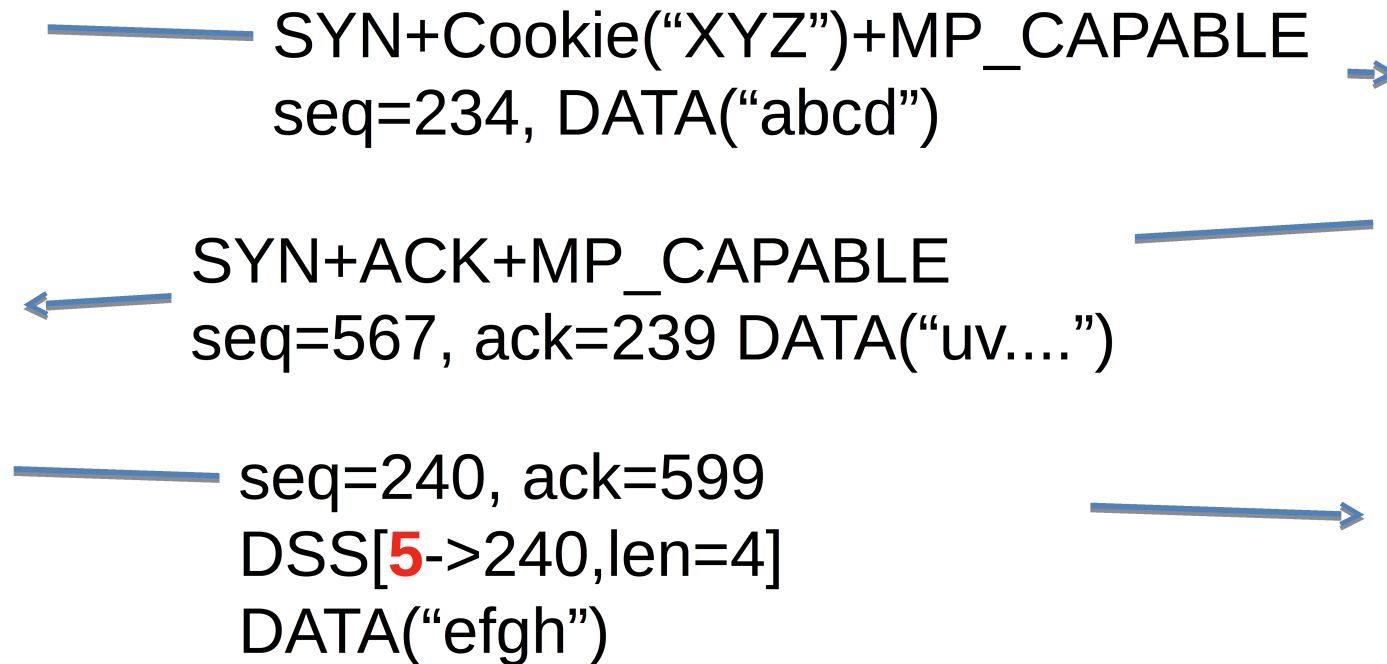
ACK MP_CAPABLE, ack=599 →

seq=240, DSS[0->240,len=4]
DATA("efgh") →

Issue #3

Stateless server

- With TFO, server has created state for connection upon reception of SYN
 - No need to replicate keys in third ack



Issue #4

- Should we use TFO for the additional subflows ?
 - Does not seem to be useful

RFC6824 : It is not permitted to send data while in the PRE_ESTABLISHED state”

Conclusion

- Multipath TCP and TFO can be combined
 - Small improvements are possible, the two are clearly not incompatible
 - Prototype implementation in Linux kernel exists and will be released soon