Avoiding Interactions of Quick-Start TCP and Flow Control

Michael Scharf <michael.scharf@ikr.uni-stuttgart.de>
Sally Floyd <floyd@icir.org>
Pasi Sarolahti <pasi.sarolahti@iki.fi>

IETF 68 - TSVWG

Experimental Quick-Start TCP Extension

One Usage Scenario of RFC 4782

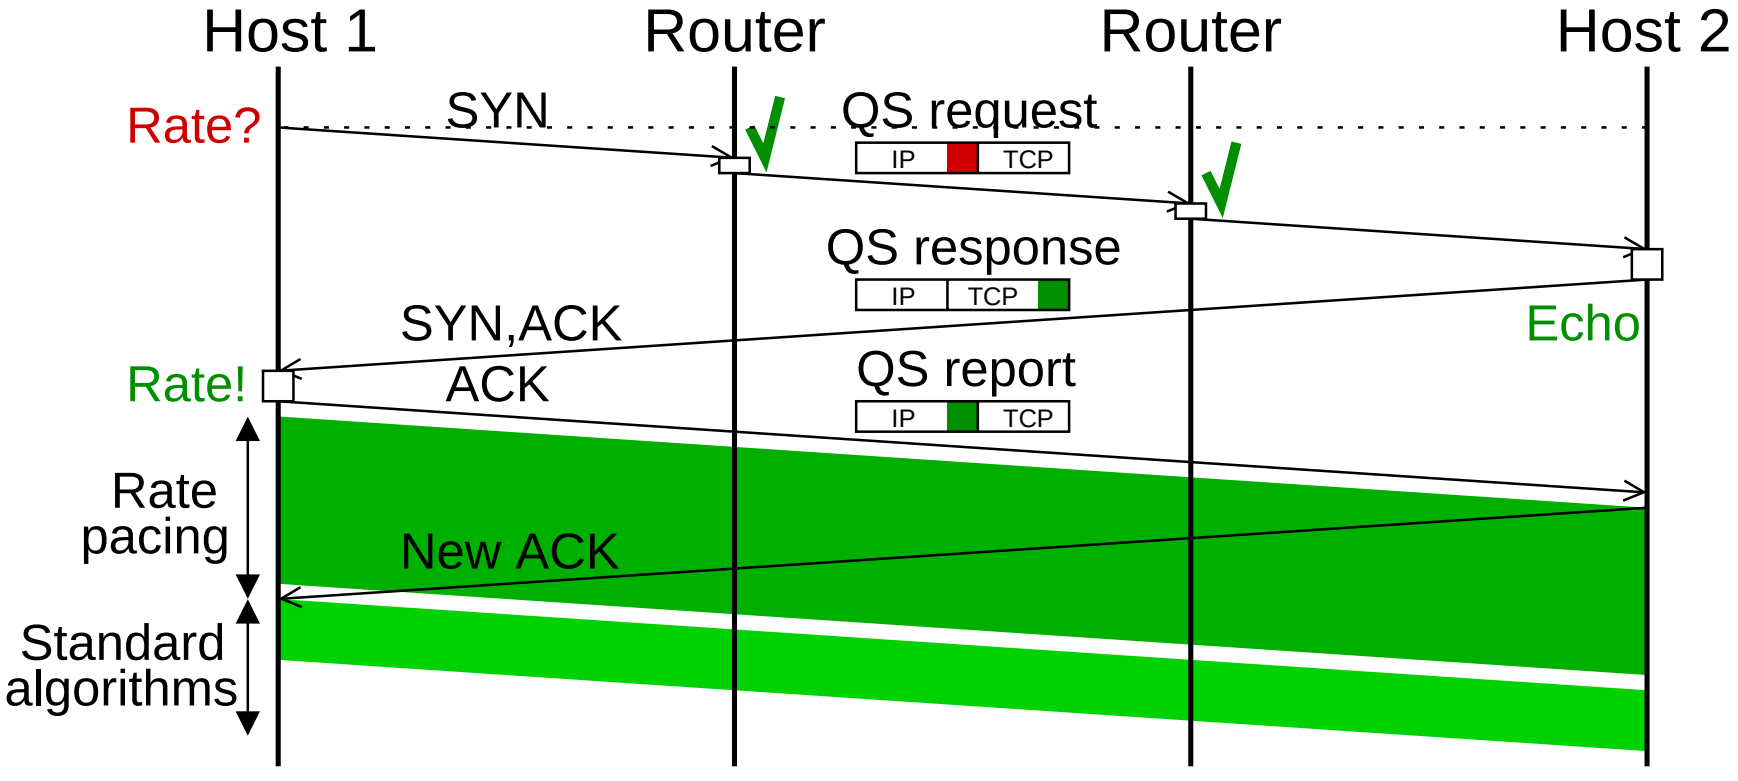

Flow Control Issue #1: Buffer Allocation

Buffer Space Autotuning

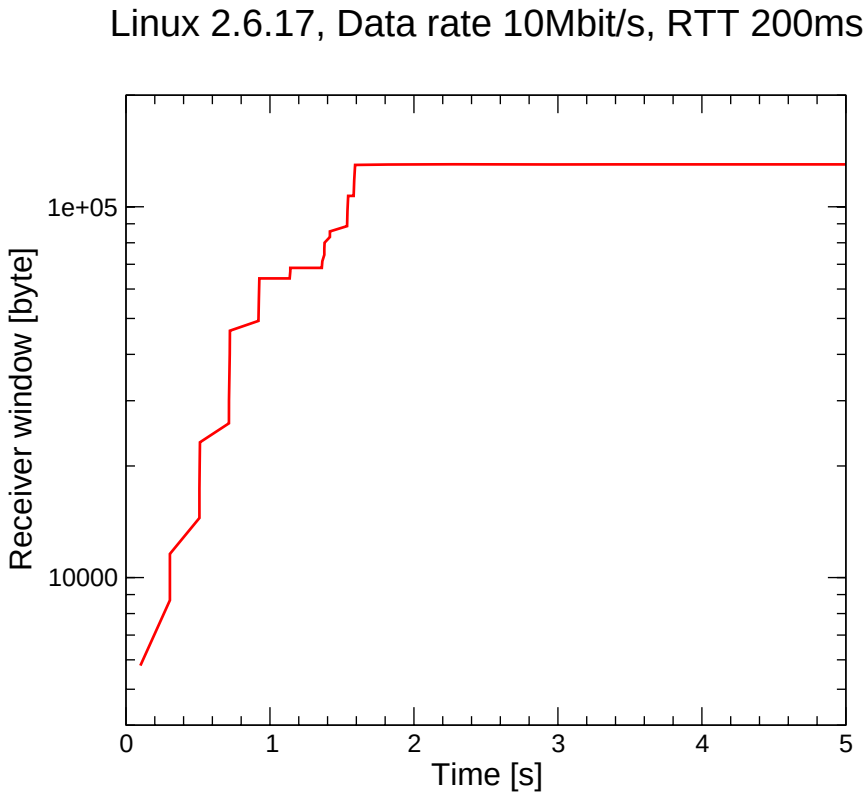

- **Quick-Start only effective for immediate large receiver buffer**
- **Reasonable buffer size can be determined from approved rate**

➥ **Recommendation for modified buffer allocation**

Flow Control Issue #2: RWND Scaling

Problem: RFC 1323 Window Scaling

- **Window scaling required when receive window is larger than 65kB**
- **BUT: "Window field in a SYN (i.e., a <SYN> or <SYN,ACK>) segment itself is never scaled."**

➥ **Maximum receive window of 65kB in <SYN> and <SYN,ACK>**

Consequences for Quick-Start

- **If Quick-Start is included in <SYN> segment**
  - At most 65kB sent during rate-pacing phase
  - Maximum congestion window after Quick-Start phase is 65kB
- **No problem otherwise**

Possible Solutions

1. **Scale RWND with Quick-Start options, thus violate RFC 1323**
2. **Signal true value of RWND, if required**

Flow Control Issue #2: RWND Scaling

Proposed Solution: Additional Acknowledgement

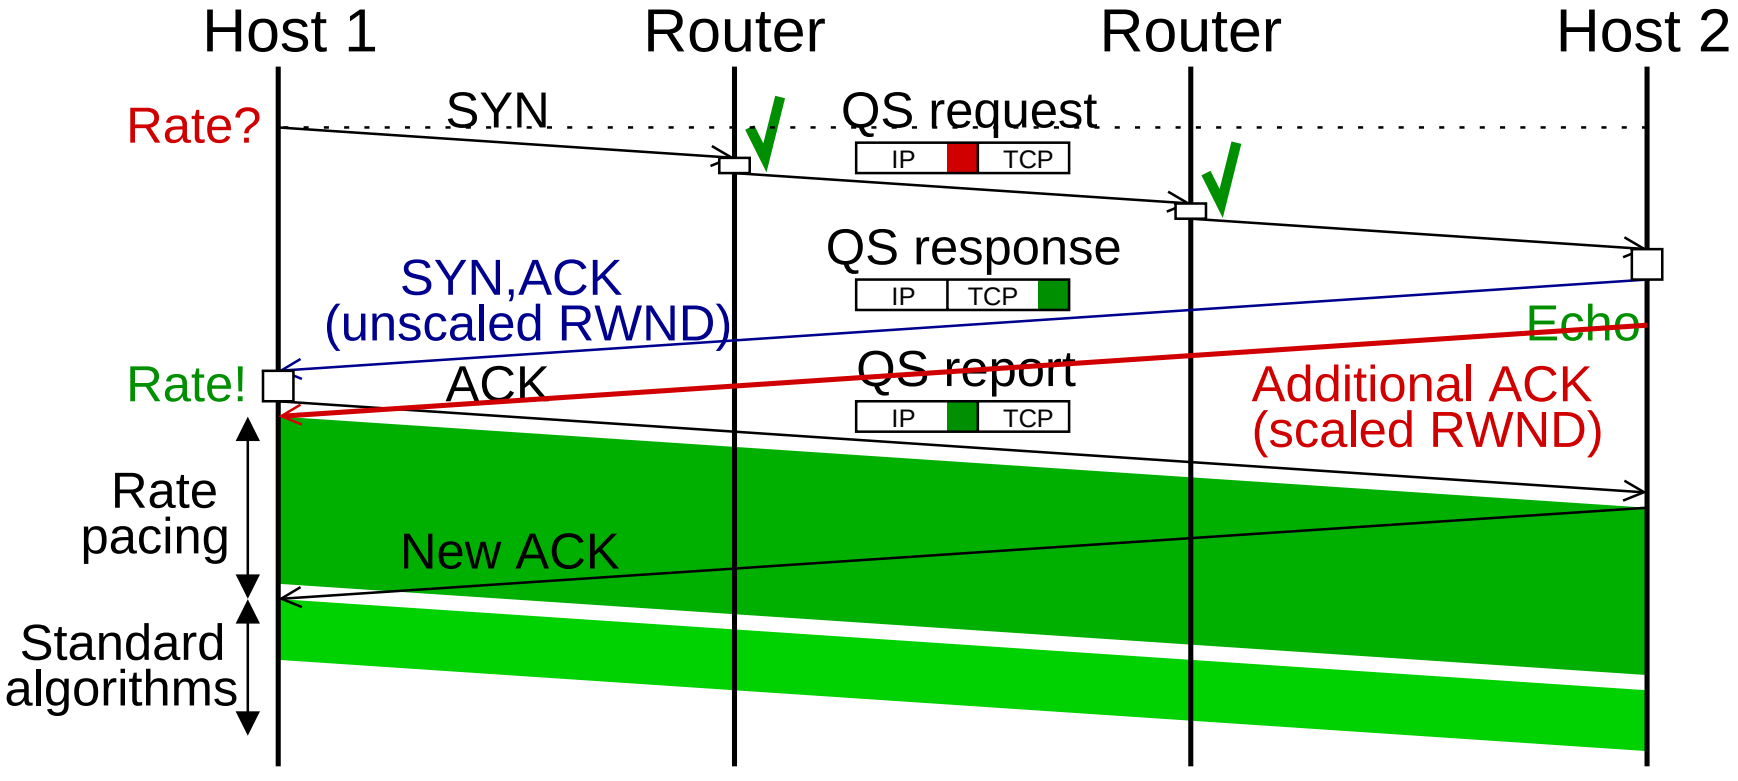

- **Send additional acknowledgement if RWND is larger than 65kB**
- **Not necessarily back-to-back with <SYN,ACK>**

Conclusion and Next Steps

Conclusions

- **TCP flow control should be optimized when using Quick-Start**
- **Possible interactions**
  1. Currently deployed buffer size auto-tuning mechanisms
  2. RFC 1323 window scaling signaling

➥ **Not discussed in RFC 4782**

Next Steps

- **Further elaborate on buffer dimensioning for Quick-Start**
- **Experimental RFC?**
- **Errata to RFC 4782?**
- **Discuss effects of adjusting RFC 1323 <SYN,ACK> behavior in future**