

# Is it still possible to extend TCP?

Authors: Michio Honda, Keio University, Fujisawa, Japan  
Yoshifumi Nishida, Keio University, Fujisawa, Japan  
Costin Raicu, Universitatea Politehnica Bucuresti, Romania  
Adam Greenhgh, University College London, United Kingdom  
Mark Handley, University College London, United Kingdom  
Hideyuki Tokuda, Keio University, Fujisawa, Japan

Presenter: Alexandros Tsalidis  
Technische Universität München  
Munich, 1. June 2017

# Content

1. Introduction
2. Methodology and Datasets
3. Tests and Results
4. Impact of findings on TCP extensions
5. Conclusion

# Introduction

- Internet was designed to be extensible but things changed
- TCP options to extend functionality
- Invasion of middleboxes
  - Interaction with TCP options?

## Contribution

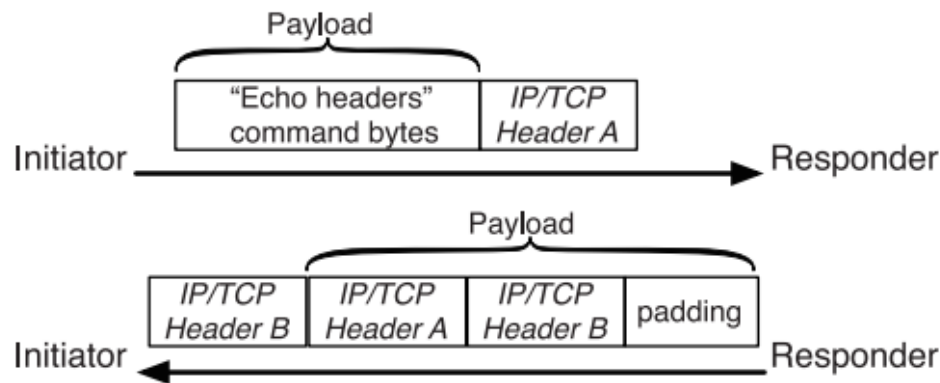
- ✓ Effects of middleboxes on TCP options
- ✓ What middleboxes are doing to traffic
- ✓ Implications for protocol designers

# Methodology and Datasets

Testing tool: TCPEXposure

## ➤ Client and Server tool

Parameter	Initiator	Responder
Initial Sequence Num (ISN)	252001	11259375
Window Size	8064	32768
MSS	512	512
Window Scale	-	6
SACKOK	-	1
Timestamp (TS_val)	-	12345678



# Methodology and Datasets

Port 80 (http)

Port 443 (https)

Port 34343 (unassigned)

Measurements target access networks, where ISPs deploy middleboxes to improve the experience of the customers

Six types of access networks:

- Home
- Hotspot
- Cellular
- University
- Enterprise
- Hosting

# TCP Options test

The following tests were performed:

1. **Unknown option in SYN.** The SYN and SYN/ACK segments include an unregistered option.
2. **Unknown option in Data segment.** The test includes unknown options in data segments sent by client and server.
3. **Known option in Data segment.** The test includes a well-known option in data segments sent by client and server.

We tested for:

- SYN is passed unmodified.
- SYN containing the option is dropped.
- SYN is received, but option was removed.
- Connection is reset by the middlebox.

**Table 3: Unknown Option in Syn**

Observed Behavior	TCP Port		
	34343	80	443
<i>Passed</i>	129 (96%)	122 (86%)	133 (94%)
<i>Removed</i>	6 (4%)	20 (14%)	9 (6%)
<i>Changed</i>	0 (0%)	0 (0%)	0 (0%)
<i>Error</i>	0 (0%)	0 (0%)	0 (0%)
<b>Total</b>	<b>135 (100%)</b>	<b>142 (100%)</b>	<b>142 (100%)</b>

**Table 4: Known Option in Data**

Observed Behavior	TCP Port		
	34343	80	443
<i>Passed</i>	129 (96%)	122 (86%)	133 (94%)
<i>Removed</i>	6 (4%)	9 (6%)	6 (4%)
<i>Changed</i>	0 (0%)	4 (3%)	3 (2%)
<i>Error</i>	0 (0%)	7 (5%)	0 (0%)
<b>Total</b>	<b>135 (100%)</b>	<b>142 (100%)</b>	<b>142 (100%)</b>

**Table 5: Unknown Option in Data**

Observed Behavior	TCP Port		
	34343	80	443
<i>Passed</i>	129 (96%)	122 (86%)	133 (94%)
<i>Removed</i>	6 (4%)	13 (9%)	9 (6%)
<i>Changed</i>	0 (0%)	0 (0%)	0 (0%)
<i>Error</i>	0 (0%)	7 (5%)	0 (0%)
<b>Total</b>	<b>135 (100%)</b>	<b>142 (100%)</b>	<b>142 (100%)</b>

# Sequence number modification

We examine whether middleboxes modify the sequence numbers sent by the end systems

Behavior	TCP Port		
	34343	80	443
<i>Unchanged</i>	126 (93%)	116 (82%)	128 (90%)
<i>Mod. outbound</i>	5 (4%)	5 (4%)	6 (4%)
<i>Mod. inbound</i>	0 (0%)	1 (1%)	1 (1%)
<i>Mod. both</i>	4 (3%)	13 (9%)	7 (5%)
<i>Proxy (probably mod. both)</i>	0 (0%)	7 (5%)	0 (0%)
Total	135 (100%)	142 (100%)	142 (100%)

Finally it is unsafe to embed sequence numbers in TCP options!

# Sequence space holes

TCP is a reliable protocol. Its cumulative ACK does not move forward unless all segments are received

Two ways for a middlebox to observe if that does not happen:

- Data-First
- Ack-First

**Table 10: Data-First Sequence Hole Test**

Behavior	TCP Port		
	34343	80	443
<i>Passed</i>	131 (97%)	120 (85%)	135 (95%)
<i>No response</i>	2 (1%)	6 (4%)	2 (1%)
<i>Duplicate Ack</i>	1 (1%)	9 (6%)	5 (4%)
<i>Test Error</i>	1 (1%)	7 (5%)	0 (0%)
<b>Total</b>	<b>135 (100%)</b>	<b>142 (100%)</b>	<b>142 (100%)</b>

**Table 11: Ack-first Sequence Hole Test**

Behavior	TCP Port		
	34343	80	443
<i>Passed</i>	102 (76%)	95 (67%)	105 (74%)
<i>No response</i>	28 (21%)	28 (20%)	29 (20%)
<i>Ack fixed</i>	4 (3%)	5 (4%)	3 (2%)
<i>Retransmitted</i>	1 (1%)	7 (5%)	5 (4%)
<i>Test Error</i>	0 (0%)	7 (5%)	0 (0%)
<b>Total</b>	<b>135 (100%)</b>	<b>142 (100%)</b>	<b>142 (100%)</b>

The conclusion of this test was that TCP extensions relying on sequence number holes are unsafe!



# Inconsistent Retransmission

What happens if a TCP sender retransmits a packet, but includes different data than the original in the retransmission?

**Table 12: Results of Retransmission Test**

Observed Behavior	TCP Port / Retransmitting size								
	34343			80			443		
	same	smaller	larger	same	smaller	larger	same	smaller	larger
<i>Passed</i>	134 (99%)	134 (99%)	132 (98%)	124 (87%)	124 (87%)	123 (87%)	138 (97%)	138 (97%)	136 (96%)
<i>No response</i>	0 (0%)	0 (0%)	1 (1%)	0 (0%)	0 (0%)	1 (1%)	0 (0%)	0 (0%)	1 (1%)
<i>Ack adv'ced</i>	1 (1%)	1 (1%)	1 (1%)	10 (7%)	10 (7%)	10 (7%)	4 (3%)	4 (3%)	4 (3%)
<i>Reset conn</i>	0 (0%)	0 (0%)	0 (0%)	1 (1%)	1 (1%)	1 (1%)	0 (0%)	0 (0%)	0 (0%)
<i>Error</i>	0 (0%)	0 (0%)	1 (1%)	7 (5%)	7 (5%)	7 (5%)	0 (0%)	0 (0%)	1 (1%)
<b>Total</b>	135 (100%)			142 (100%)			142 (100%)		

Any extension that wished to use inconsistent retransmissions would encounter few problems, so long as it did not matter greatly whether the original or the retransmission actually arrives.

# Re-segmentation

- To test segment splitting, we simply send a full-sized segment
- The opposite of segment splitting is segment coalescing, where a middlebox combines two or more segments into a larger segment

**Table 13: Results of Segment Coalescing Test**

Observed Behavior	34343	TCP Port 80	443
<i>Passed</i>	132 (98%)	123 (87%)	138 (97%)
<i>Coal. ordered</i>	1 (1%)	1 (1%)	0 (0%)
<i>Coal. queued</i>	1 (1%)	3 (2%)	1 (1%)
<i>Coal. both</i>	0 (0%)	6 (4%)	3 (2%)
<i>Error</i>	1 (0%)	9 (6%)	0 (0%)
Total	135 (100%)	142 (100%)	142 (100%)

Although middleboxes do split and coalesce segments, none did so while passing unknown options. So it seems relatively safe to assume that if an option is passed, it arrives with the segment on which it was sent

# Intelligent NICs

- Even Network Interface Cards can have embedded TCP knowledge
- We are concerned in particular with TCP Segmentation Offload (TSO), where the host OS sends large segments and relies on the NIC to resegment to match the MTU or the receiver's MSS
- 12 NICs were tested and they all correctly copied the options to all the split segments
- The implication is that TCP options must be designed so that when they are duplicated on consecutive segments, this does not adversely affect correctness or performance

# Implications for protocol designers

- Multipath TCP
  - Decisions were dictated by middleboxes
- TcpCrypt
  - Tightly constrained by middleboxes for its solutions
- Extending TCP option space
  - MPTCP and TcpCrypt use relatively large option space and leave very little free. So Long Option can be deployed with some concerns

# Conclusion

- It is still possible to extend TCP using the intended mechanism – TCP options although there are some guidelines and assumptions we should not take for granted
- Based on that information the three TCP extensions we checked had made sensible choices and are tightly constrained by middlebox behaviors to their chosen solutions
- We urge middlebox designers to consider explicitly whether they want to allow new TCP extensions when implementing certain functionality

Thank you for your attention!