

Middlebox Investigation in The Internet for Designing TCP Extensions

Michio Honda, Keio University / UCL

Yoshifumi Nishida, WIDE Project

Costin Raiciu, UCL

Mark Handley, UCL

27.July, 2010

78th IETF @ Maastricht, Netherlands



Keio University
1858
CALAMVS
GLADIO
FORTIOR

Introduction

- Middleboxes interfere TCP connections/segments
 - This prevents TCP extensions getting deployed
 - We need to know how middleboxes affect TCP
- This study is very helpful to design TCP extensions

Middlebox Behavior We Concern


- New/unknown TCP options
 - How middleboxes affect unknown TCP options in SYN and data segments
- Segment splitting
 - How middleboxes affect full-sized segments (including how options are treated)
- Segment coalescing
 - How middlebox affect small (less than MSS) segments (including how options are treated)
- Transparent proxies
 - How transparent proxies exist in the Internet, and how they affect TCP

Middlebox Behavior We Concern

- Initial Sequence Numbers
 - How frequently middleboxes rewrite sequence numbers
- Retransmission
 - How middleboxes behave against retransmitted segments
- Sequence holes
 - How middleboxes behave against sequences including holes


Experimental Setup

- We initiate/transmit TCP traffics to 476 PlanetLab nodes
 - 476/991 nodes responded regular TCP traffic
- Source node:
 - sfc.wide.ad.jp (Fujisawa, Japan)
- Server-side behavior
 - Send back Ack for every segment
 - Put DATA_ACK into every Ack for segments including MP_DATA



SYN Option Experiment


- Transmit a SYN segment with Multipath CAPable option



- 479/487 paths (~98.36%) passed MPCAP
- 0/487 paths dropped packet
- 8/487 paths (~1.64 %) removed MPCAP
- We got same results when we place options with padding each of them

Data Option Experiment

- Transmit a full-sized segment with MultiPath DATA Sequence Number option



- 441/487 paths (~90.55 %) passed MPDATA
- 0/487 paths dropped packet
- 46/487 paths (~9.45 %) removed MPDATA
 - 50 % of paths removing MPCAP also removed MPDATA
 - Another 50 % of paths removing MPCAP didn't remove MPDATA


Retransmission Experiment

- Retransmit sequences with different payload


Seg. 1 (seq. 0-511, 512 Bytes) Seg. 2 (seq. 512-1010, 499 Bytes)



Server always sends back
dup Ack for 499 Bytes data




Seg. 2 (seq. 512-1023, 512 Bytes)



- 437/438 paths (~99.77 %) passed retransmissions with different payload
- 1/438 paths (~0.23 %) of paths rejected retransmissions with different payload

Sequence Hole Experiment

- Transmit 2 non-consecutive segments



- 465/467 paths (~99.57 %) allowed sequence hole
- 2/467 paths (~0.43 %) of paths rejected sequence hole

The Other Experiment


- Initial Sequence Number
 - 477/479 paths (~99.58 %) didn't rewrote ISN
 - 2/479 paths (~0.42 %) rewrote ISN
- Transparent proxy
 - 479/483 paths (~99.17 %) couldn't have transparent proxy
 - 4/483 paths (~0.83 %) could have transparent proxy
- Segment Coalescing
 - 0/467 paths coalesced small segments
 - 0/438 paths coalesced small segments including unknown option

Conclusion


- Explore middlebox existence in the Internet by using PlanetLab
- Middleboxes in front of PlanetLab nodes likely don't affect TCP so much
 - Many of them are located in academic network
- But we figured out some of paths affect TCP behavior
- Ongoing work:
 - Opposite direction test
 - Well-known port test
 - Big-option test (e.g., more than 40 Bytes)\
 - TCP flags treatment

Segment Coalescing Experiment

- Transmit small segments




- 0/467 paths (0%) coalesced segments
 - In all trials, we got 3 acks for each segment




- 0/438 of paths (0%) coalesced segments

Segment Coalescing Experiment (2)

- Transmit small out-of-order segments to make middleboxes queue them



- 0/467 paths (0%) coalesced segments



0/438 (0%) paths coalesced segments