




Lessons Learnt from Middlebox Measurement

Michio Honda

Advanced Technology Group,
NetApp

michio@netapp.com



With acknowledge to Yoshifumi Nishida,
Costin Raiciu and Mark Handley

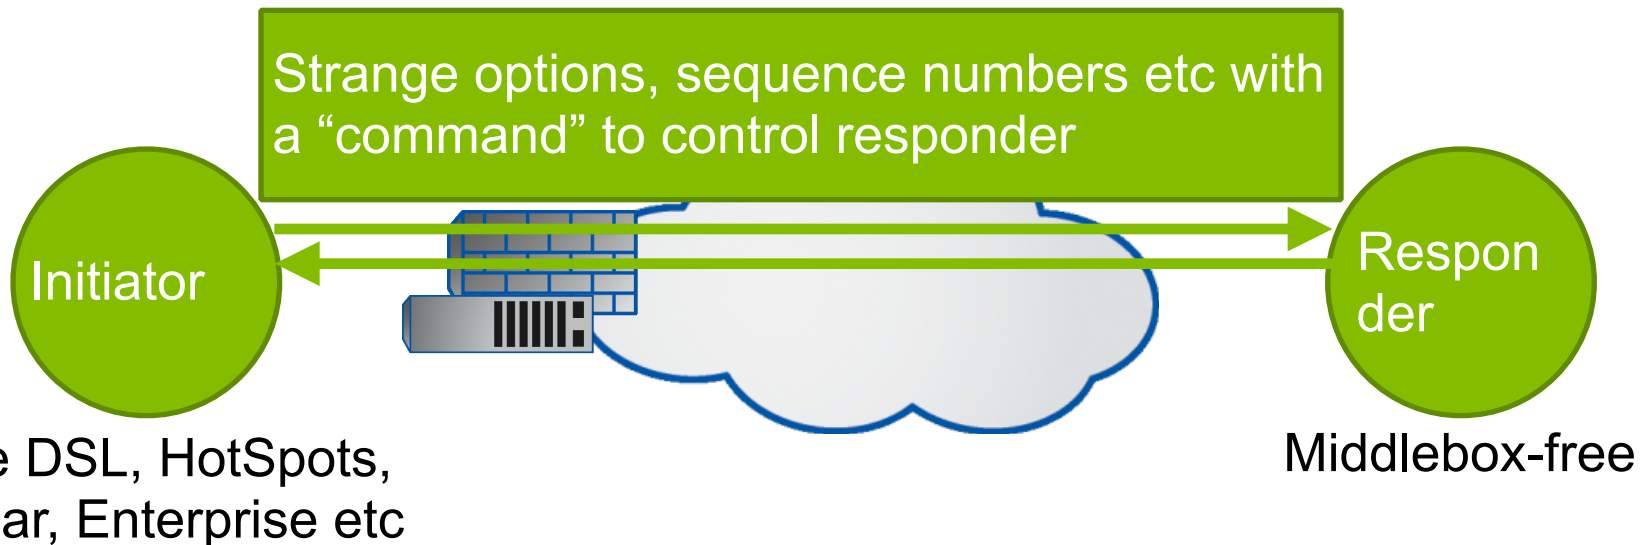


Summary of the Last Measurement

- “Is it still possible to extend TCP?” (IMC’11)
 - Identified whether potential TCP extension design works or not
 - Identified deployability of potential TCP extensions

Summary of the Last Measurement

- “Is it still possible to extend TCP?” (IMC’11)
 - Identified whether potential TCP extension design works or not
 - Identified deployability of potential TCP extensions





Rerunning experiments is very hard



Before Measurement

- Carefully define **objectives**
 - Experiments and methodology depend on objectives
 - Our case: extension design (not deployment status)
- Carefully design **experiments**
 - Identify what middlebox behavior is really useful to examine
 - Our case: mimicking possible TCP extensions (not testing whether a middlebox conforms standards)



Designing Tools

- Carefully design and implement tools
 - Generating raw packets needs root privilege
 - Cannot be just a smartphone app
 - People might not want to “install” the tool
 - Should work just with downloading and executing it
- We should have supported Windows to get more contributors...



Designing Tools (2)

- Responder
 - Stateless server helps a lot but has some limitations
 - Inability to see if option is removed from SYN-ACK



Gathering results

- Important to reward contributors
 - We should have done better...



Other know-hows

- App-level (e.g., HTTP) proxy is difficult to pass
 - Test traffic really need to pretend real traffic between a real HTTP client and server



Comparison with Tracebox

	<i>Measurement</i>	<i>Install</i>	<i>limitations</i>
<i>TCPEXposure</i>	<i>Both ends are controlled</i>	<i>No</i>	<i>measuring server-side middlebox behavior is limited</i>
<i>Tracebox</i>	<i>Only initiator is controlled (rely on ICMP response)</i>	<i>Needed</i>	<i>limited ability for qualitative measurement</i>



Comparison with Tracebox

	<i>Measurement</i>	<i>Install</i>	<i>limitations</i>
<i>TCPEXposure</i>	<i>Both ends are controlled</i>	<i>No</i>	<i>measuring server-side middlebox behavior is limited</i>
<i>Tracebox</i>	<i>Only initiator is controlled (rely on ICMP response)</i>	<i>Needed</i>	<i>limited ability for qualitative measurement</i>

We should combine both methodologies for RG-wide activity



Summary

- **Rerunning experiments is really hard**
 - Carefully define objectives
 - Carefully identify what is really useful
 - Carefully design experiments