Analyzing the CoAP DDoS Amplification Attack Ecosystem: A Honeypot Study

Anton Kalpakchiev
Supervisor: Dr. Harm Griffioen
Delft University of Technology

TU Delft

1) Motivation

- The CoAP protocol is used in DDoS amplification attacks, reaching **320 Gbps of bandwidth**
- **Unknown attacker behavior** - CoAP differs from the studied protocols that threat models are based on
- **No guidance on defense thresholds** - aggressive thresholds can hurt performance for real CoAP users

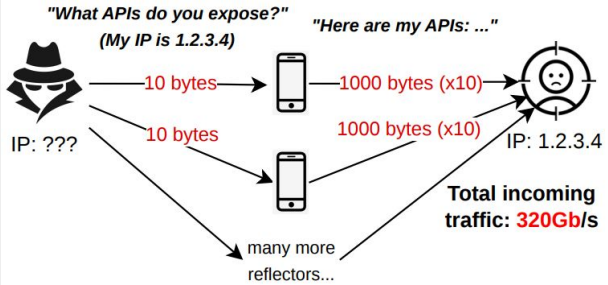

Figure 1: A DDoS amplification attack with CoAP

Questions:

How does CoAP compare to studied algorithms in 1) scanning behavior and 2) actor sophistication?

How effective are defenses (blockwise transfer and ratelimitting) at different thresholds?

2) Methodology

- We deploy a **honeypot** - a VM that pretends to be an amplifier and records all traffic for analysis:

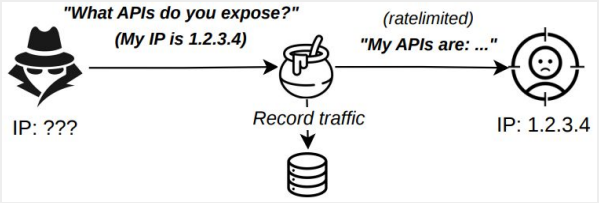

Figure 2: A honeypot captures actor traffic

- We deploy 8 honeypots with varying thresholds for 1) ratelimitting and 2) response block sizes
- We later look for differences in actor behavior

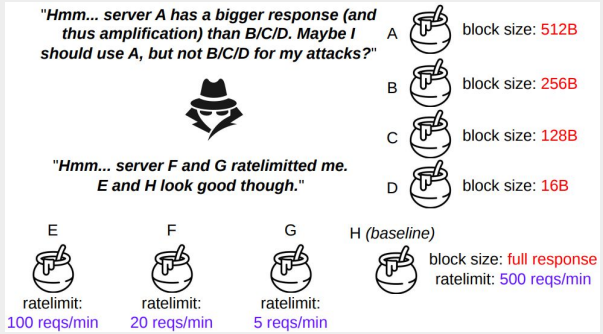

Figure 3: The 8 honeypots differing across two defense dimensions (block size & ratelimit)

3) Results & Conclusions

Actor sophistication is higher than in previously studied protocols

- **Device fingerprinting** - actors look for a specific CoAP implementation (QLC Chain)

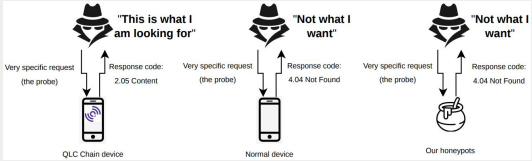

Figure 4: Actors probe for a QLC Chain device

- **Defense detection** (not seen before) - actors probe for blockwise transfer support

No attacks occurred and actor behavior was uniform across honeypots

- Honeypots filtered, maybe due to fingerprint
- Answering Q2 is therefore left as future work

Scanner Composition

- CoAP has more research scanners - 70% vs 36% for other protocols

Table 1: Distribution of actors

Category	Scanners	%
Research	26	70%
Abuse-flagged	1	3%
Unknown	10	27%
Total	**37**	

a.d.kalpakchiev@student.tudelft.nl
github.com/Anton-Kalpakchiev
23nd June 2026