DNS Rate Limiting

W. Matthijs Mekking

matthijs@NLnetLabs.nl

One slide DNS

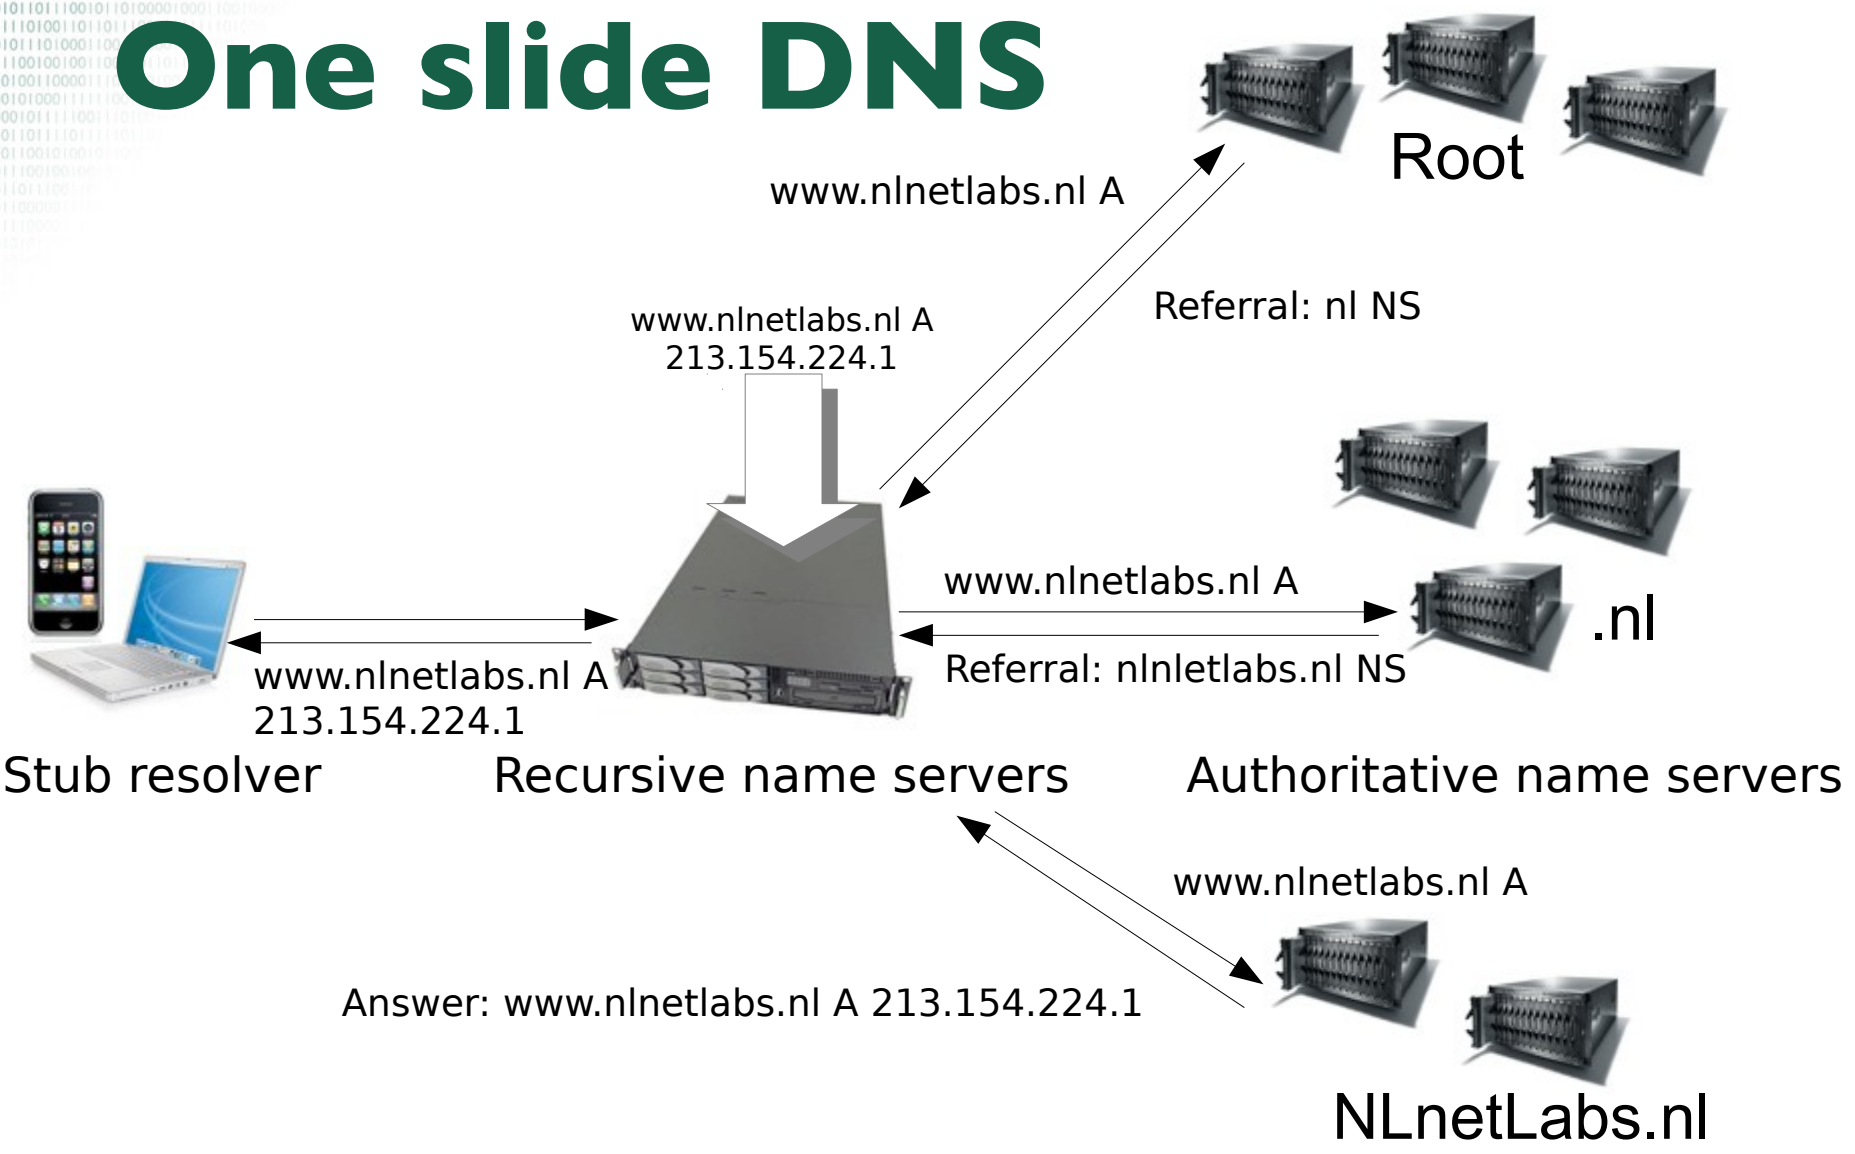

Authoritative name servers

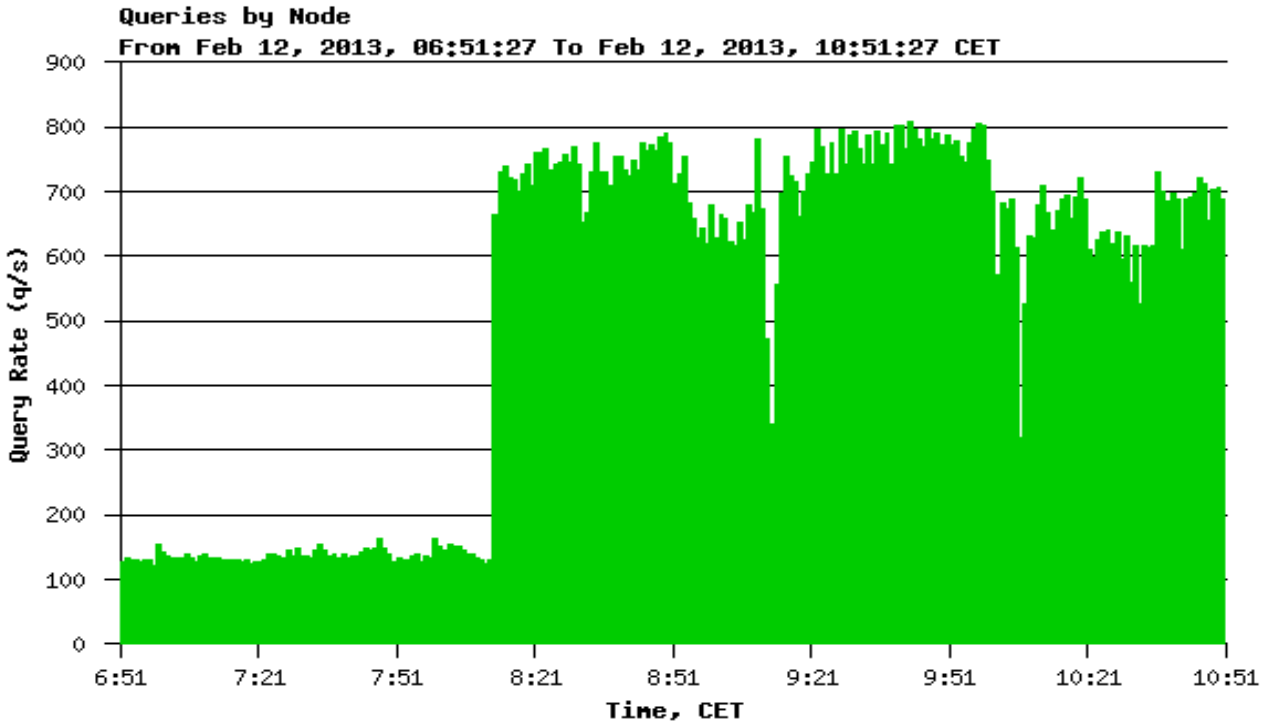

Inhaltsübersicht

- DNS DDoS attacks
- Defense mechanisms
- Response Rate Limiting
- Future developments

Background

- I am. Matthijs Mekking
- I work at. NLnet Labs
- I maintain. NSD

- And. OpenDNSSEC, Unbound, ldns, shim6, ...

DNS DDoS attacks

Attack properties

- "Reflection"
  - IP address spoofing

- Solution: BCP 38 (Ingress filtering)
  - No deployment

Attack properties

- "Amplification"
  - Large responses
  - UDP

"These two characteristics make DNS a hot target."

DNS amplification

- ANY
  - Apex, NS RRset
  - ≈ 1:80 (with DNSSEC)
- NXDOMAIN + DNSSEC
  - NSEC(3) + RRSIG records
  - ≈ 1:18 (NXDOMAIN, NSEC)
  - ≈ 1:25 (NXDOMAIN, NSEC3)

Abusing resolvers

Service only to the intended clients

- RFC 5358:
  Preventing Use of Recursive Nameservers in Reflector Attacks

Abusing auth. servers

Service to many, many recursive name servers

- Wide audience, RFC 5358 not applicable

Abusing auth. servers

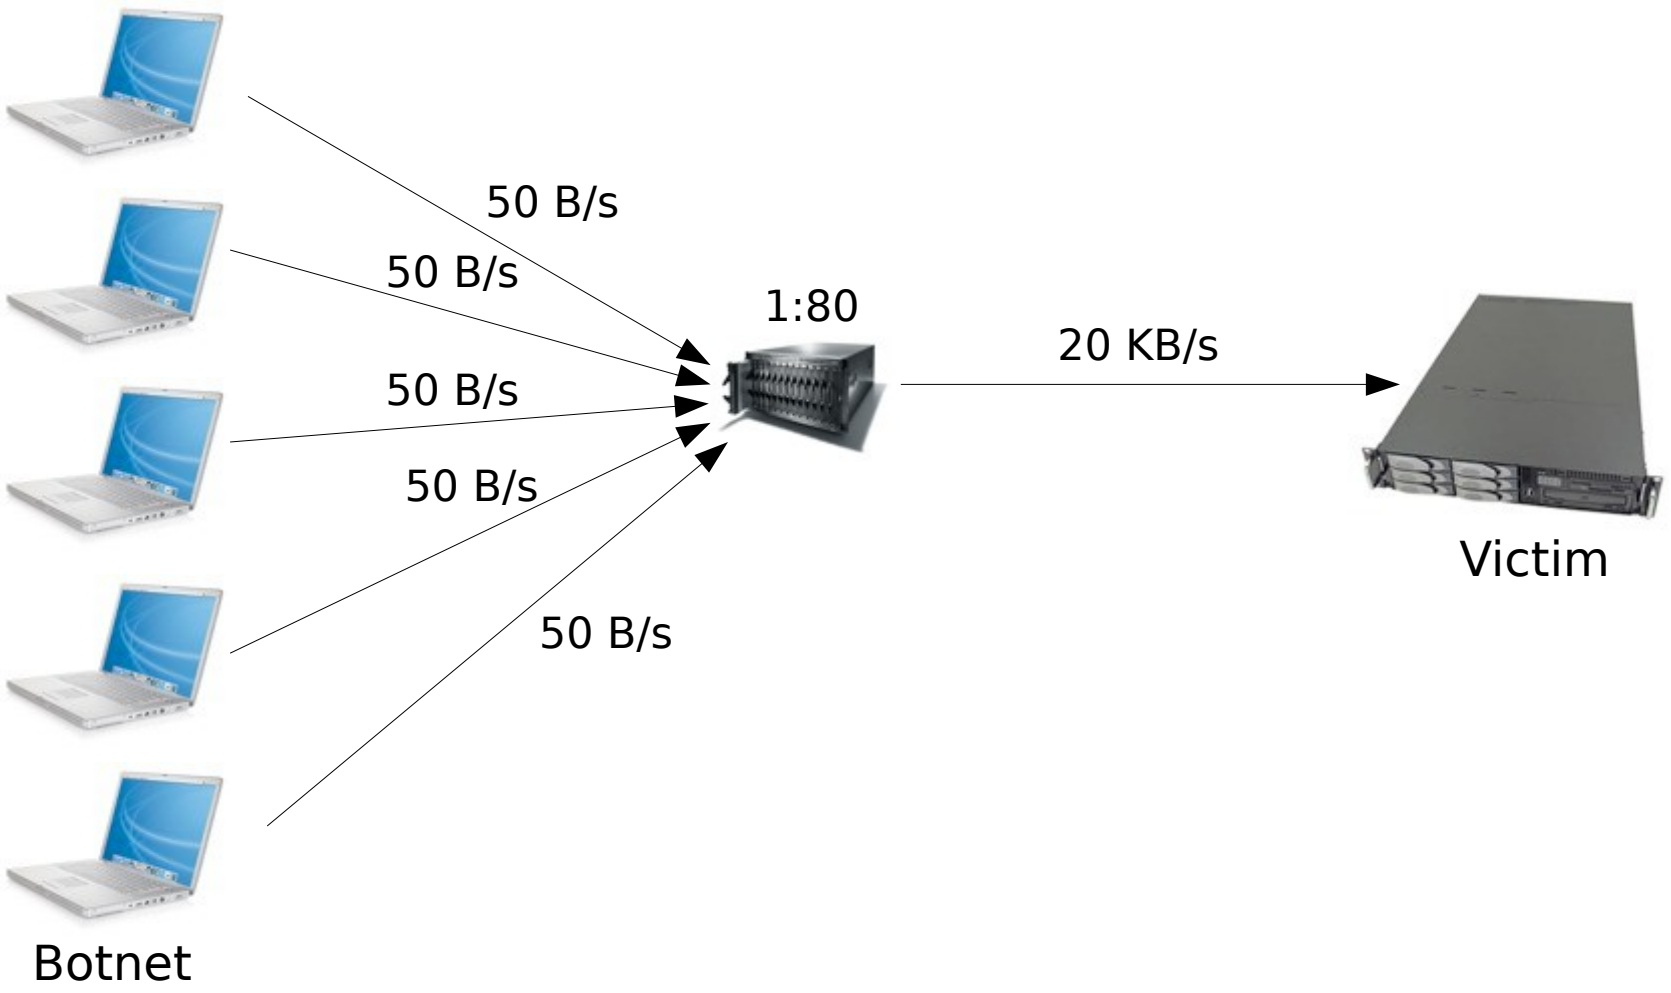

Impact of the attack

Defense mechanisms

Defense mechanisms

- Adding capacity
  - But more capacity means more abuse
- Prevent reflection
  - BCP 38
- Mitigate amplification
  - Limit bandwidth
  - DNS Rate Limiting

DNS Rate Limiting

- Blocking UDP ANY queries
- DNS Firewall
- DNS Dampening
- Response Rate Limiting

DNS Rate Limiting

- Blocking UDP ANY queries
  - ANY queries are legitimate requests
    - DNS "Health checks"
    - Mail forwarders (collecting addresses, data)
  - Quick fix for one attack scenario

DNS Rate Limiting

- DNS Firewall
  - Can be configured to block specific packets or address range
  - Drawbacks:
    - (Manual) reactive approach
    - No flexibility

DNS Rate Limiting

- DNS Firewall (cont.)

```
iptables -A INPUT -p udp --dport 53 -m hashlimit \
 --hashlimit-name DNS --hashlimit-above 20/second \  --hashlimit-
mode srcip --hashlimit-burst 100 \
 --hashlimit-srcmask 28 -j DROP
```

http://www.bortzmeyer.org/files/generate-netfilter-u32-dns-rule.py

DNS Rate Limiting

- DNS Dampening
  - Proposal by Lutz Donnerhacke
  - Based on BGP Route Dampening
  - Self learning service
    - Collect penalty points per address range (network)
    - Hysteresis: Start at high level, stop at a much lower value
    - During dampening no processing occurs

DNS Rate Limiting

- DNS Dampening (cont.)
  - Unofficial patch for BIND9
  - Effective, but aggressive:
    - False positives during attack
    - Configuration defaults

http://lutz.donnerhacke.de/eng/Blog/DNS-Dampening

DNS Rate Limiting

- Response Rate Limiting (RRL)
  - Proposal by Paul Vixie and Vernon Schryver

http://www.redbarn.org/dns/ratelimits

RRL

RRL - Overview

- Limit responses instead of queries
  - Resolvers have caches
  - Drop answers that exceed rate limit
- False positive mitigation
  - TCP fallback
- Performs well in existing cases
- Implemented in BIND9 (official patch) and NSD 3.2.15
  - … This just in: Also in Knot 1.2.0

RRL – How it works

- State blob:
  - Address buckets:
    - IPv4 /24, IPv6 /56
  - Response name:
    - QNAME or Wildcard or Apex (NXDOMAIN)
  - Error status (RCODE)
    - NOERROR or NXDOMAIN or ERROR
  - <bucket, name, status>

RRL – How it works

- When generating response
  - Look up state blob:
    - <mask(ip), impute(name), status(rcode))
  - Increment state blob:
    - <bucket, name, status> ++
  - If state blob threshold is reached, drop
  - Or slip...

RRL – How it works

- SLIP
  - False positive mitigation
  - Allow DDoS victim to contact DNS servers over TCP
  - Approximately 1/SLIP responses with TC (instead of dropped)

RRL - Configuration

- RESPONSES-PER-SECOND
- NXDOMAINS-PER-SECOND
- ERRORS-PER-SECOND
- WINDOW
- IPv{4,6}-PREFIX-LENGTH
- SLIP
- MAX-TABLE-SIZE
- MIN-TABLE-SIZE

RRL – Measurements

"Defending against DNS reflection amplification attacks."

Thijs Rozekrans, Javy de Koning
Universiteit van Amsterdam
February 2013

http://www.nlnetlabs.nl/downloads/publications/report-rrl-dekoning-rozekrans.pdf

RRL – Measurements

- Use cases:
  - ANY
    - Targeted at zone apex
  - NXDOMAIN
    - Trigger NSEC(3)s in response
  - Varying queries
    - 25%, 50%, 75%, 100% positive responses
    - Guessing, indexing, zone walking, ...

RRL – Measurements

- Setup:

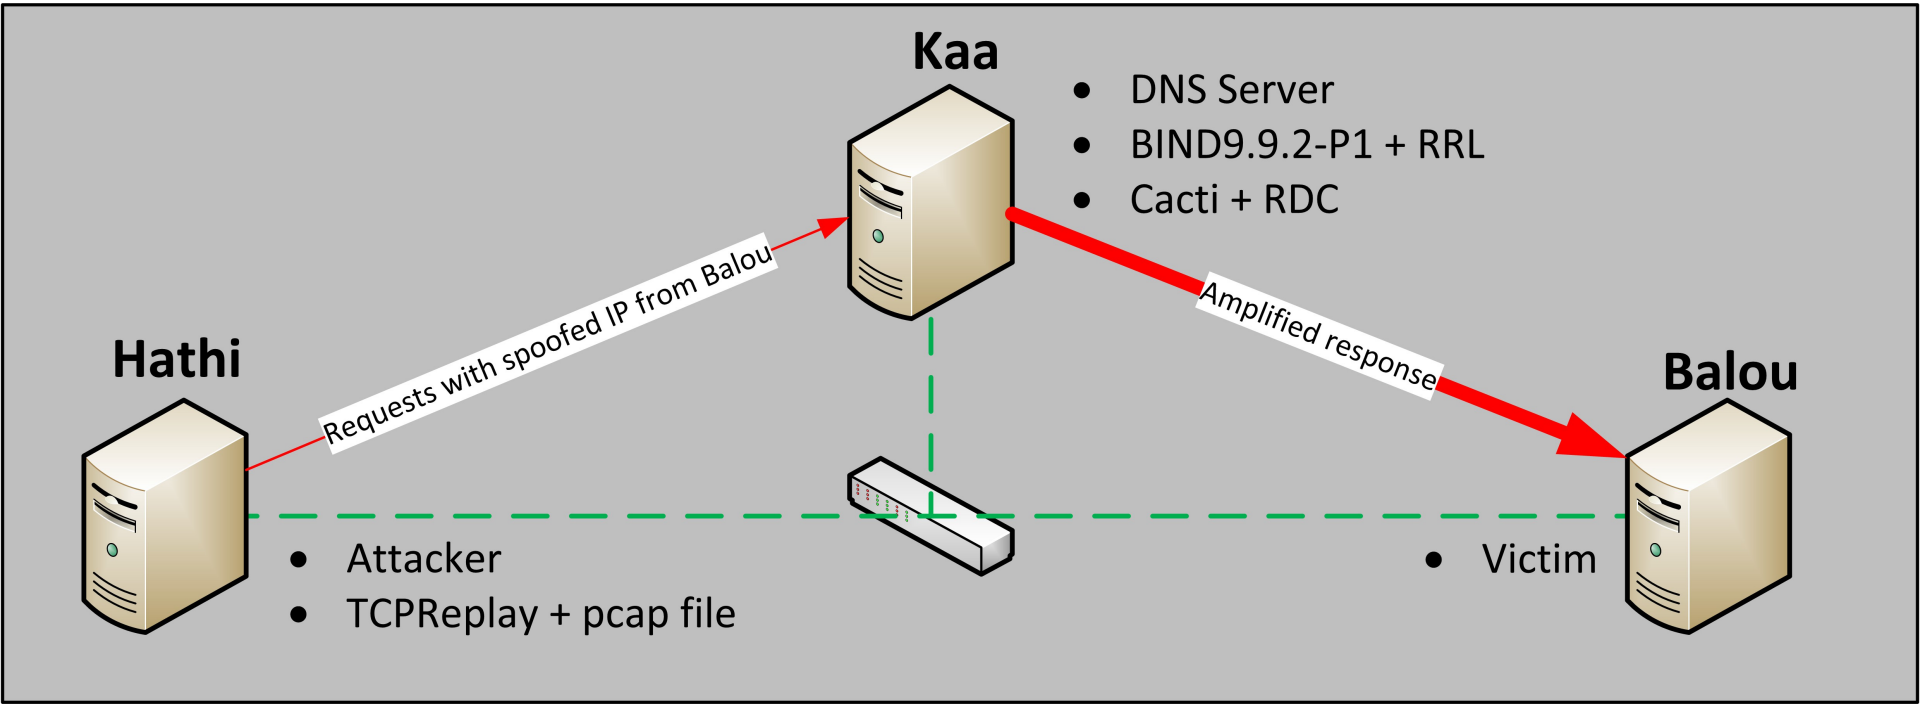

RRL – Measurements

- ANY attack at 11:20

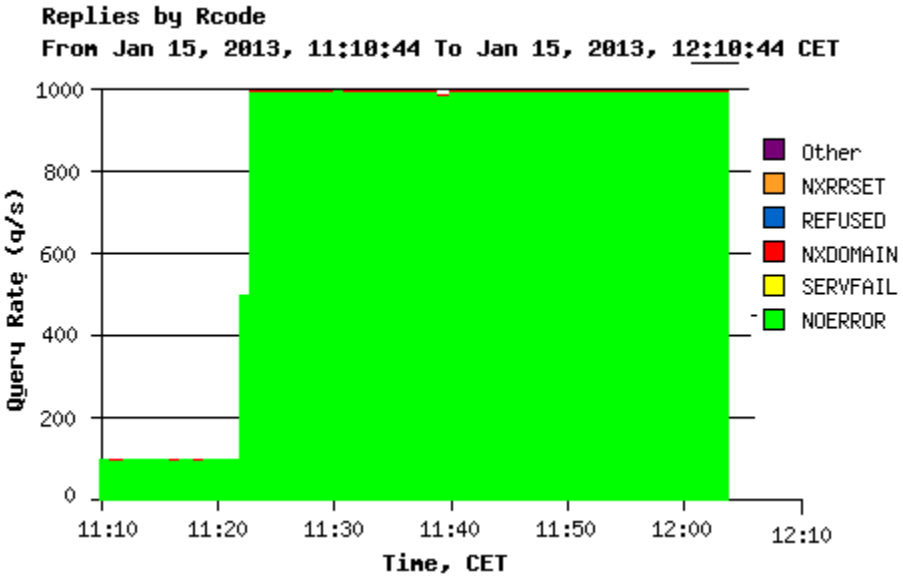

RRL – Measurements

- ANY attack at 11:20
  - RRL enabled at 14:15

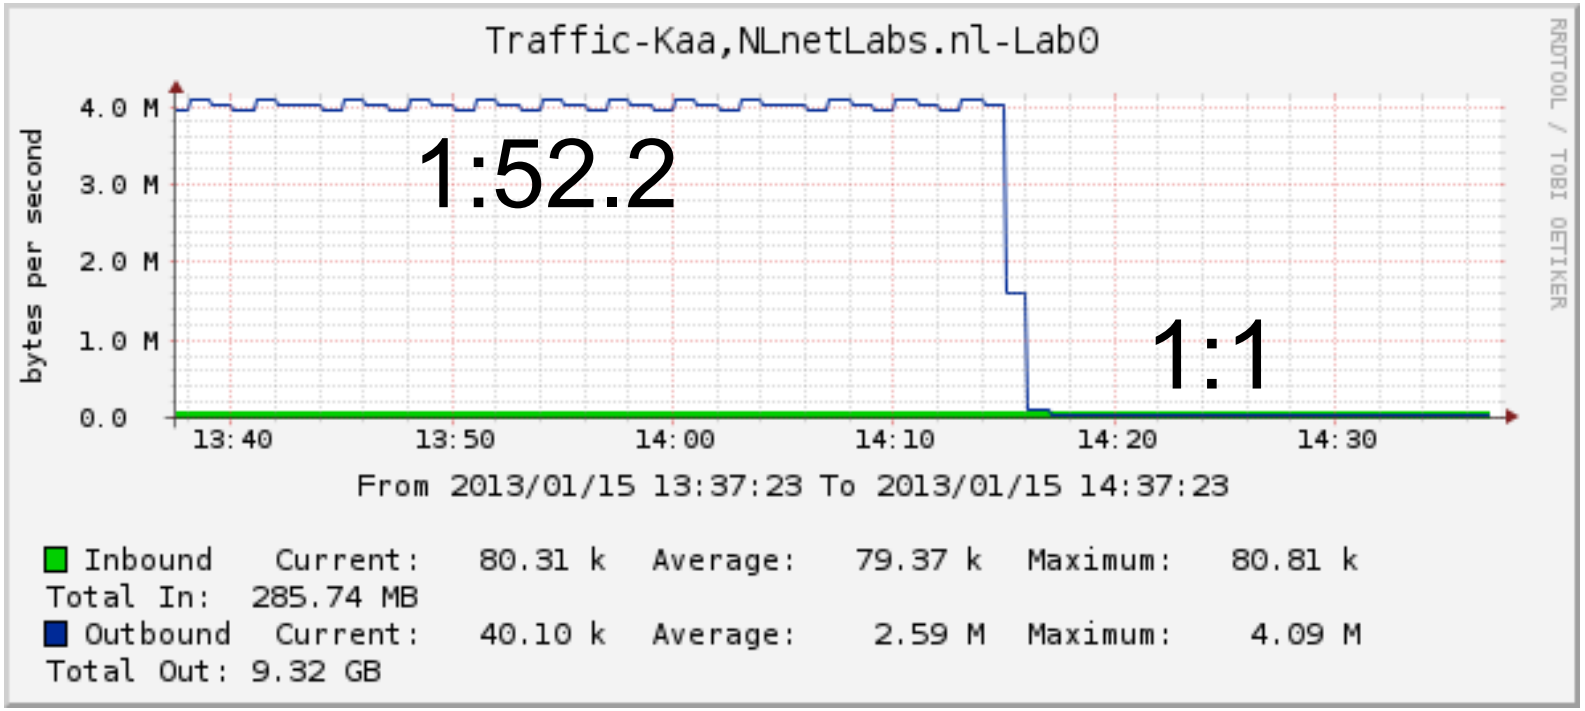

RRL – Measurements

- ANY attack at 11:20
  - RRL enabled at 14:15

SLIP	In	Out	Amp. Ratio	False positives*	TCP responses
1	80 KB/s	81 KB/s	≈1:1	0%	100 %
2	79 KB/s	39 KB/s	≈1:0.5	50%	87.5 %
3	79 KB/s	26 KB/s	≈1:0.3	66.6%	66 %
5	80 KB/s	16 KB/s	≈1:0.2	80%	49 %
10	80 KB/s	8 KB/s	≈1:0.1	90%	27 %

* Possible fps, assuming 3 tries

RRL – Measurements

- NXDOMAIN attack
  - 0% positive answers, 100% nxdomains
  - Similar results as ANY attack

RRL – Measurements

- Varying query attack
  - 25% positive answers, 75% nxdomains

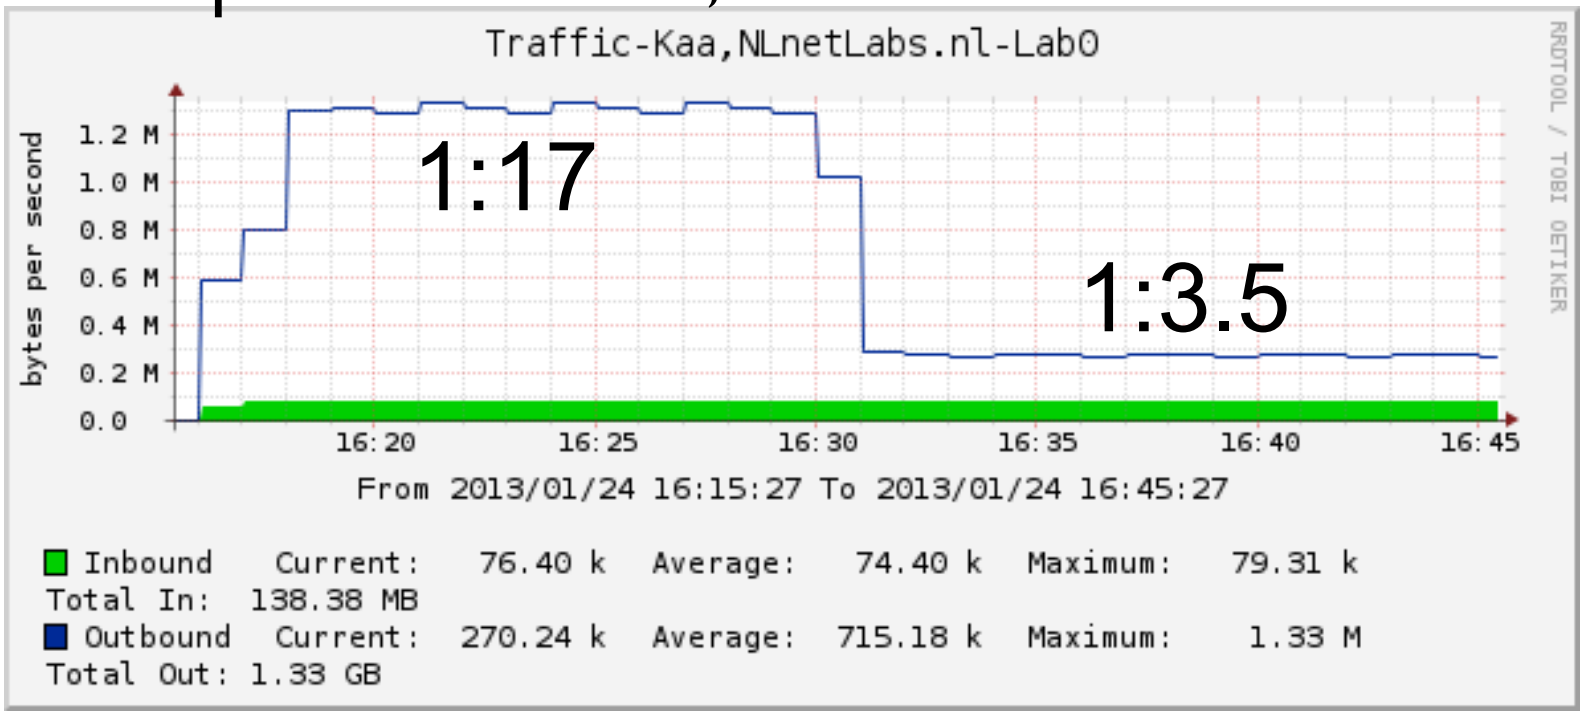

RRL – Measurements

- Varying query attack
  - 50% positive answers, 50% nxdomains

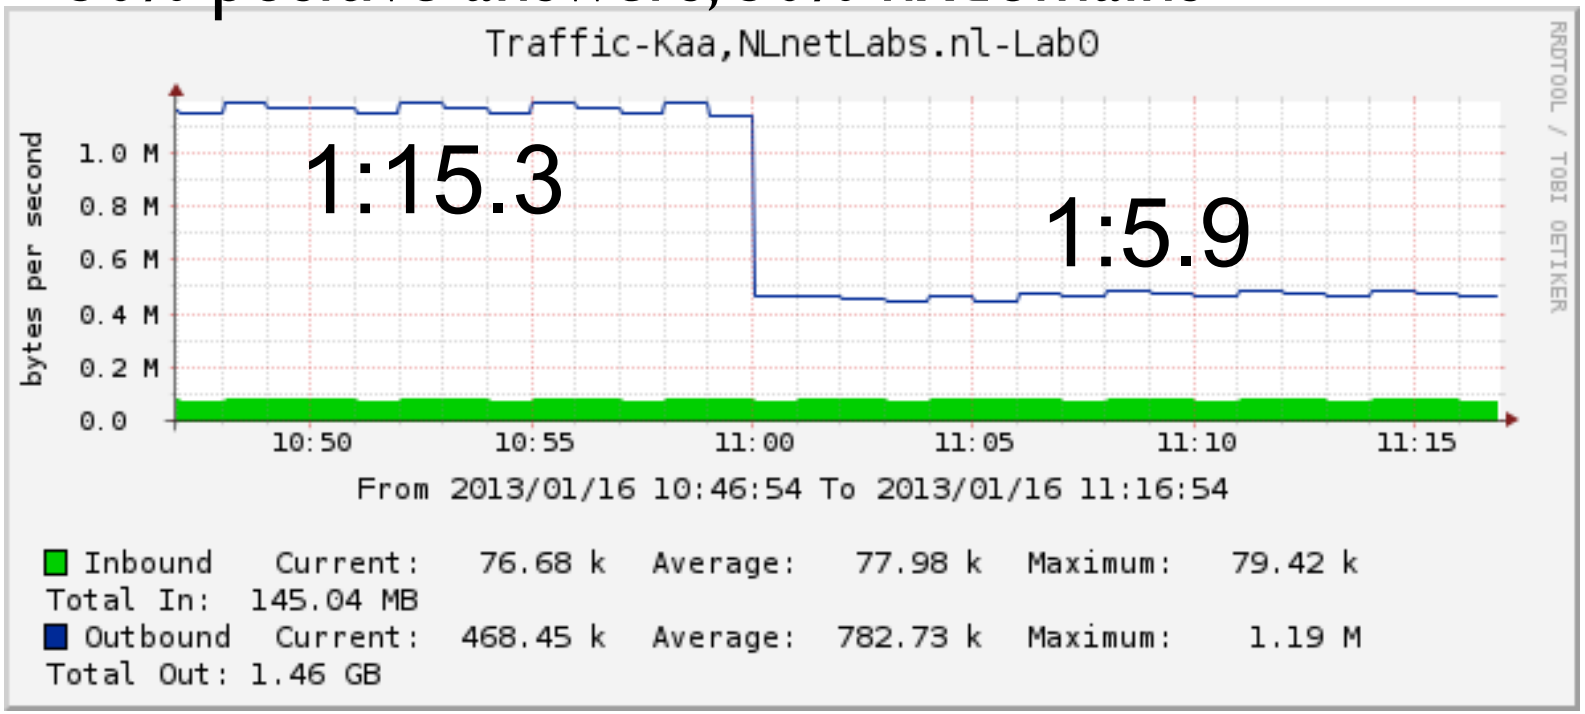

RRL – Measurements

- Varying query attack
  - 75% positive answers, 25% nxdomains

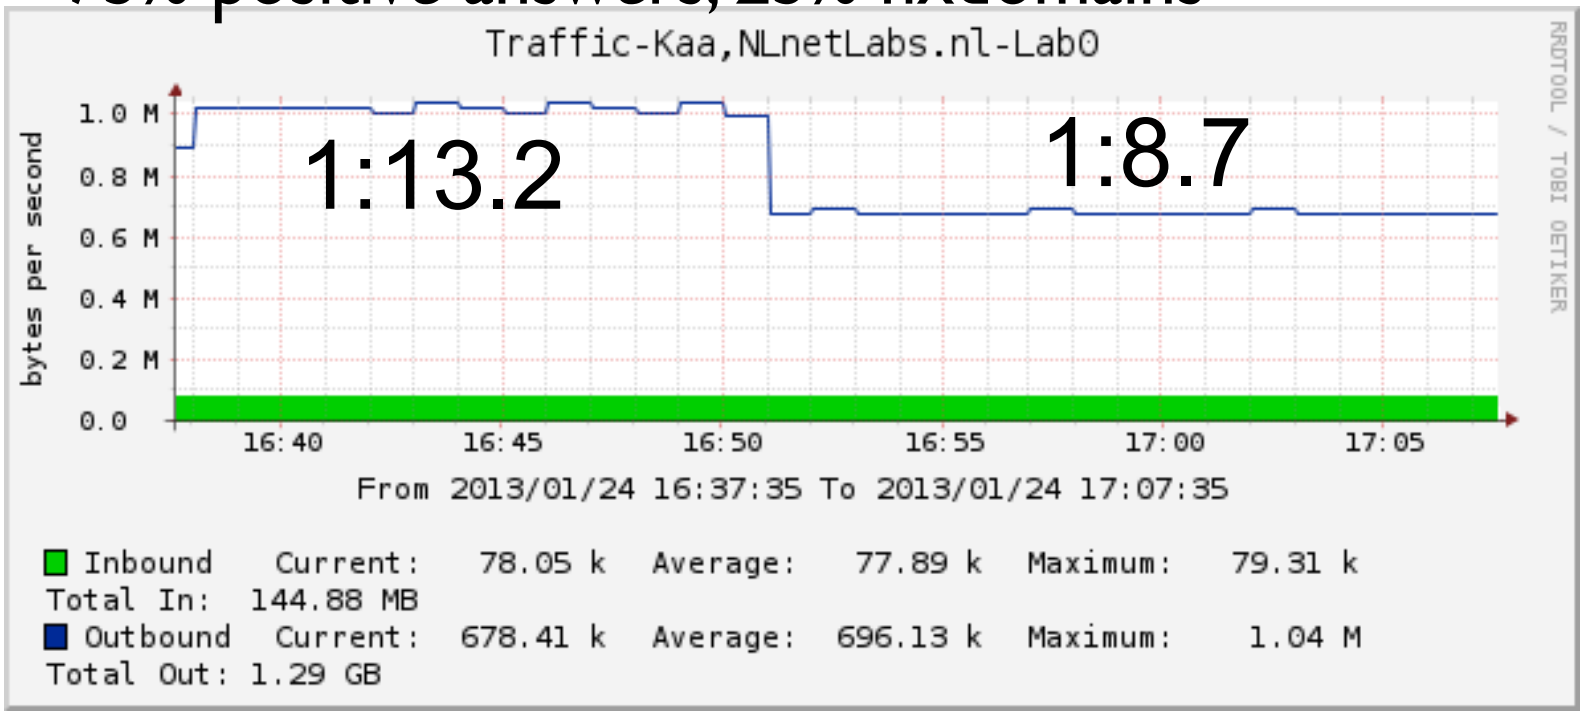

RRL – Measurements

- Varying query attack
  - 75% positive answers, 25% nxdomains

SLIP	In	Out	Amp. Ratio	False positives*	TCP responses
1	79 KB/s	689 KB/s	≈1:8.72	0%	100 %
2	78 KB/s	680 KB/s	≈1:8.72	50%	87.5 %
3	79 KB/s	677 KB/s	≈1:8.57	66.6%	66 %
5	79 KB/s	673 KB/s	≈1:8.52	80%	49 %
10	79 KB/s	665 KB/s	≈1:8.42	90%	27 %

RRL – Measurements

- Varying query attack
  - 100% positive answers, 0% nxdomains
  - Amplification ≈1:11.14
  - RRL has no effect
  - Attack only works for large zones (TLD)

RRL - Measurements

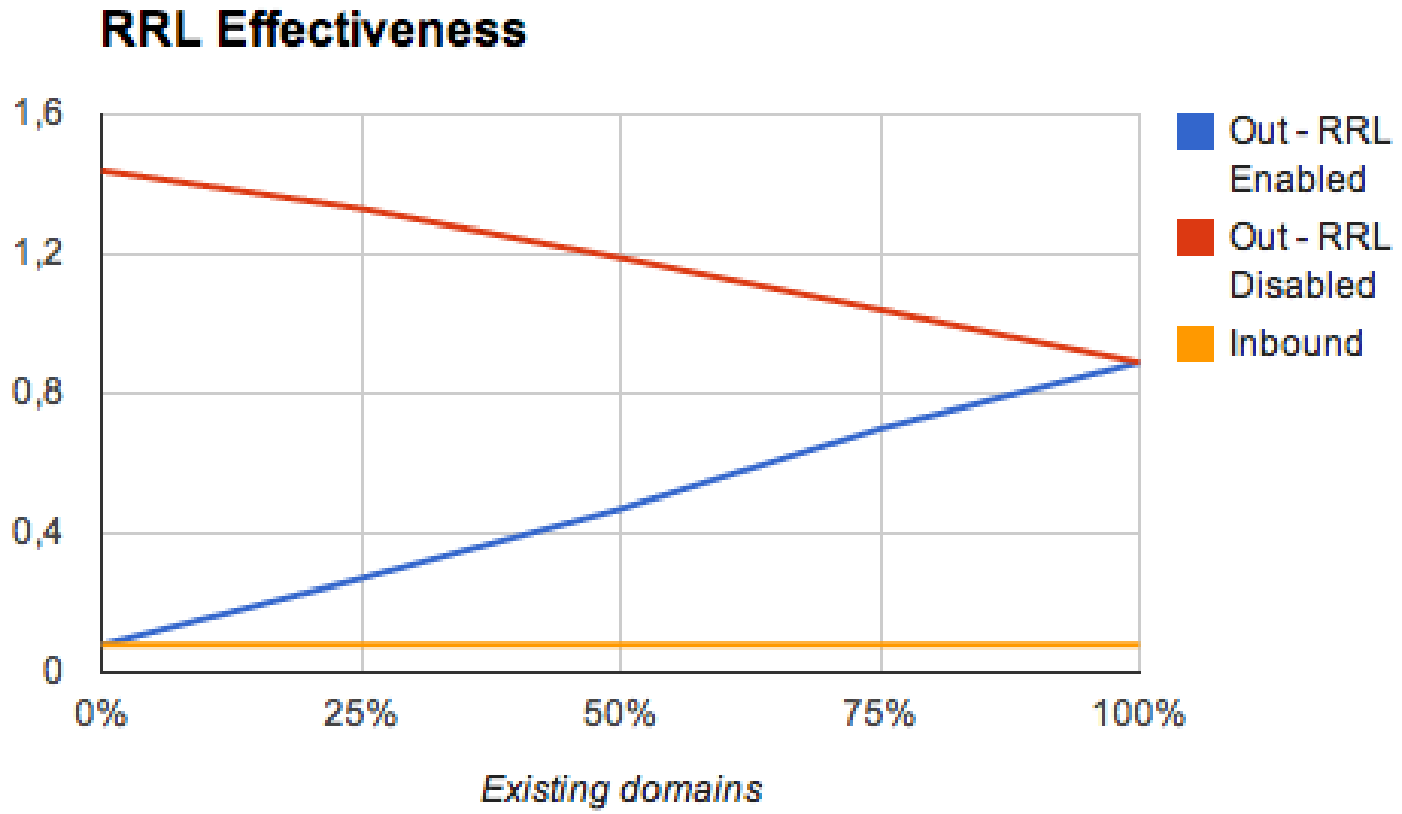

RRL – Settings

- But for the short term, a good solution

RRL – Settings

- SLIP
  - Trade off between #false positives and #TCP sessions
  - Default 2

RRL – Settings

- Rate limit
  - Default: 5
  - M: Maximum acceptable outgoing traffic per subnet
  - L: Largest response found
  - Configure threshold to be M/L
  - Example:
    [30KB/s / 4KB/s] = 8

RRL – Settings

- Window
  - Keep small: back to normal operation asap
  - But not too small: flapping
  - Default 15s

RRL – Settings

- IPv4 prefix length
  - Default /24
- IPv6 prefix length
  - Default /56

RRL – Settings

- Max table size
  - max(qps) * window
  - Default 10000 (≈1MB)
- Min table size
  - Default 1000
  - Advised to set higher

Future developments

Cat-and-mouse game

- DNS Rate Limiting is a cat-and-mouse game.
- RRL is currently an effective defense mechanism...
- … but attacks are getting more sophisticated.
  - Varying query attack

Cat-and-mouse game

- Dampening does actually do well with the varying query attack:

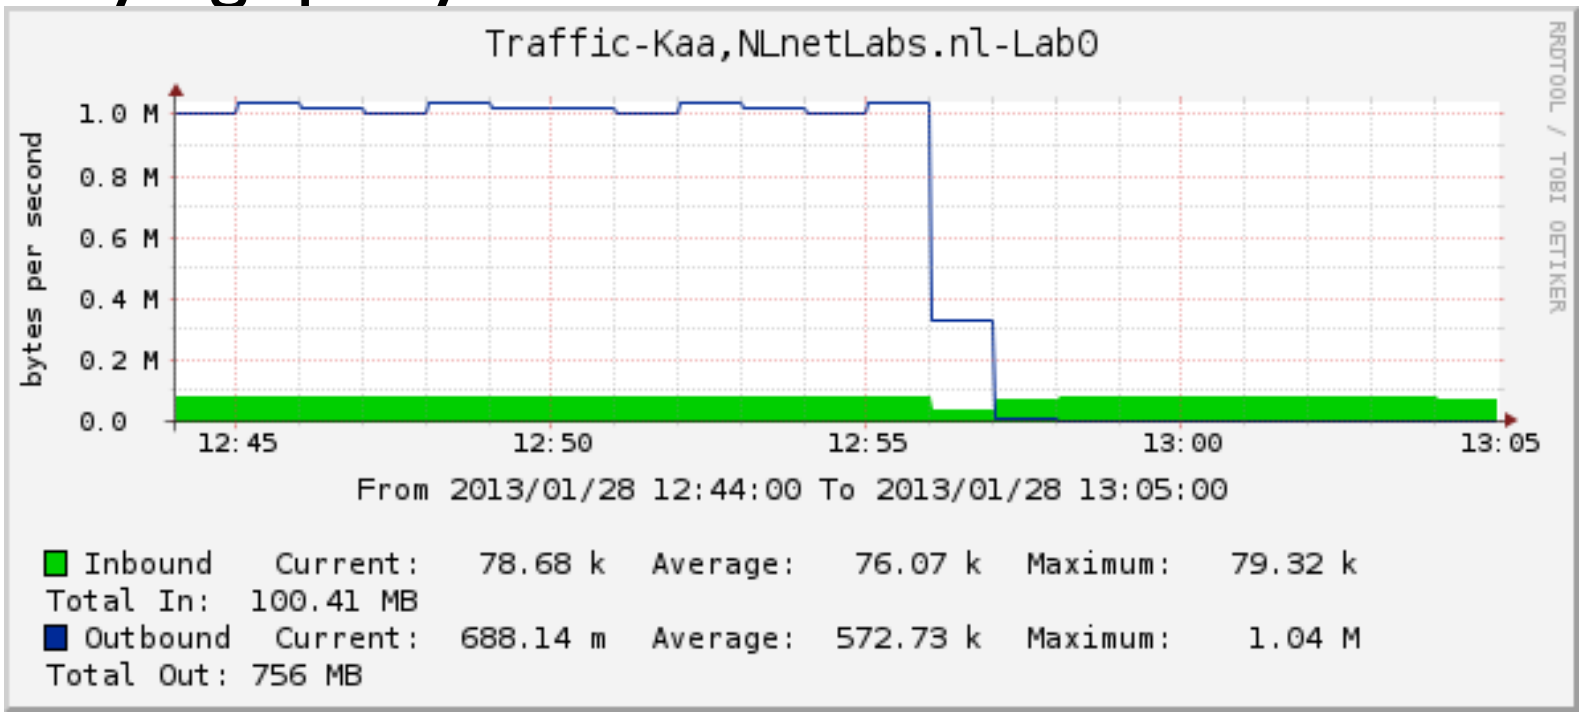

Cat-and-mouse game

- Dampening does actually do well with the varying query attack.
- It needs SLIP and parameter tweaking...

Cat-and-mouse game

- … but not anymore if distributed attacks will happen

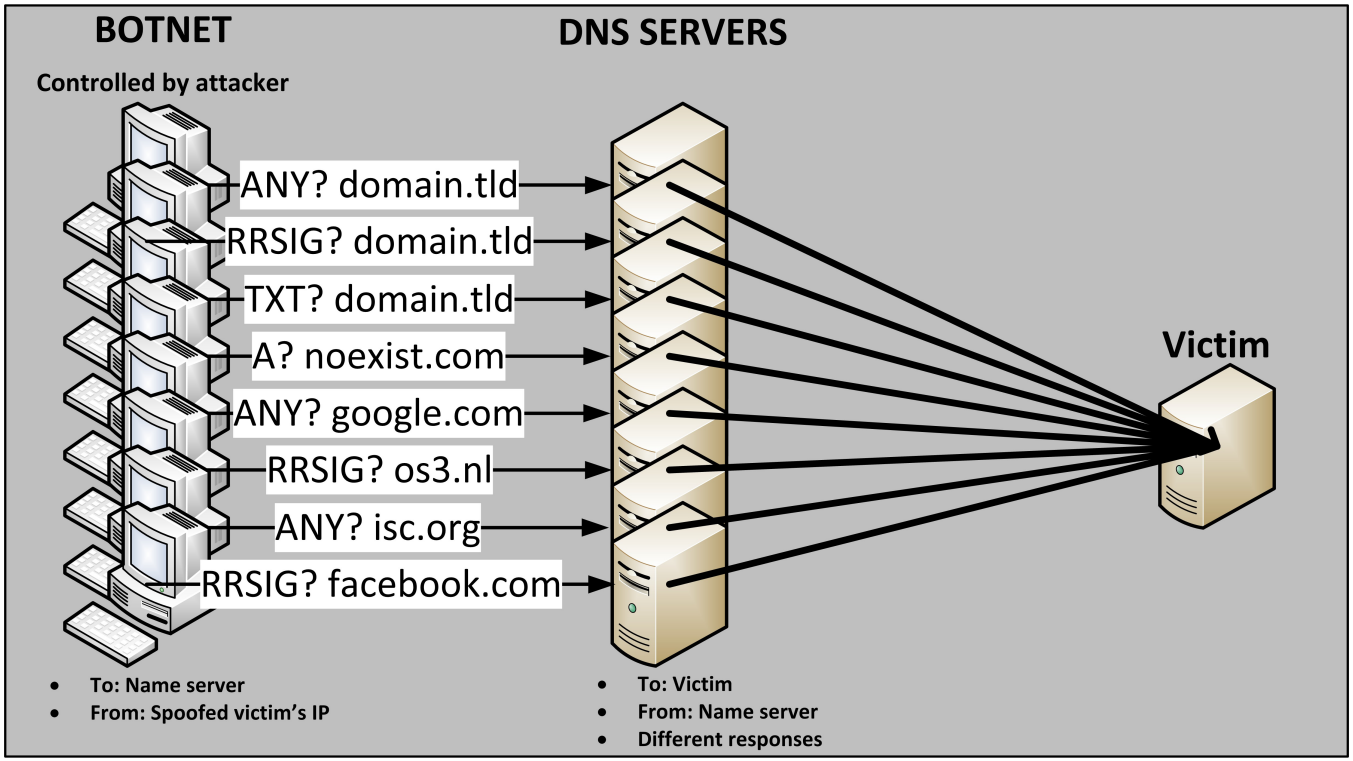

BCP 38

- Amplification is not the biggest issue, reflection is.
- Network vendors need to implement source validation (BCP 38, Ingress filtering)
  - For the good of the Internet
  - But no economic or legal incentive
  - No direct benefit, operational costs, risks

Conclusions

- DNS DDoS attacks are real practical problems
  - Reflection and amplification
- Amplification mitigation
  - Resolvers: RFC 5358
  - Authoritative nameservers: Rate Limiting
    - Cat-and-mouse game
- Reflection prevention
  - BCP 38 is the real solution
    - How to stimulate adoption?

Questions?

W. Matthijs Mekking
matthijs@nlnetlabs.nl

If you like our work, please consider sponsoring us