

A world map with a black background. The landmasses are outlined in white. Numerous small, multi-colored dots (red, orange, yellow, green, blue) are scattered across the map, representing data points. The dots are most densely clustered in North America, Europe, and East Asia.

# Preventing and Remediating Criminal Abuse of Online Infrastructure

Michel van Eeten



“Breaking into computers might be  
the bicycle theft of the future”

Netherlands Attorney General  
Gerrit van der Burg



# DDoS in Netherlands, 2015

> 30,000 attacks  
observed in  
honeypot data

vs.

86 reports filed  
with the police

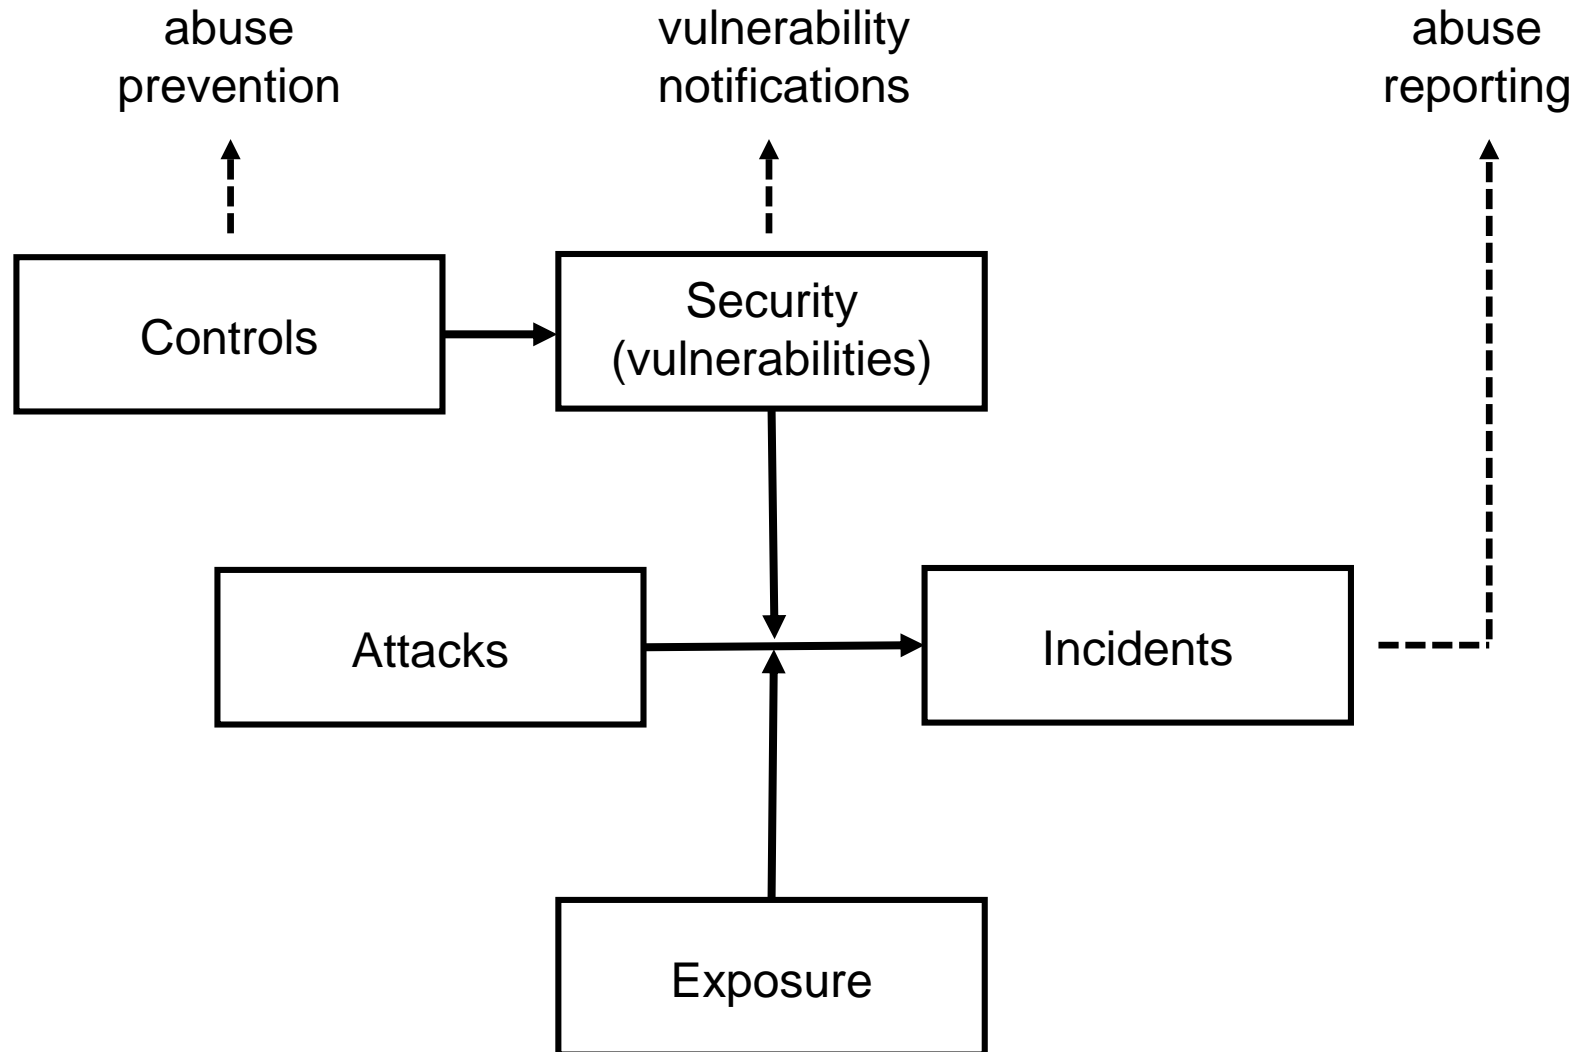
Source: Jan Koenders, The DDoS plague:  
Law enforcement view, 2016

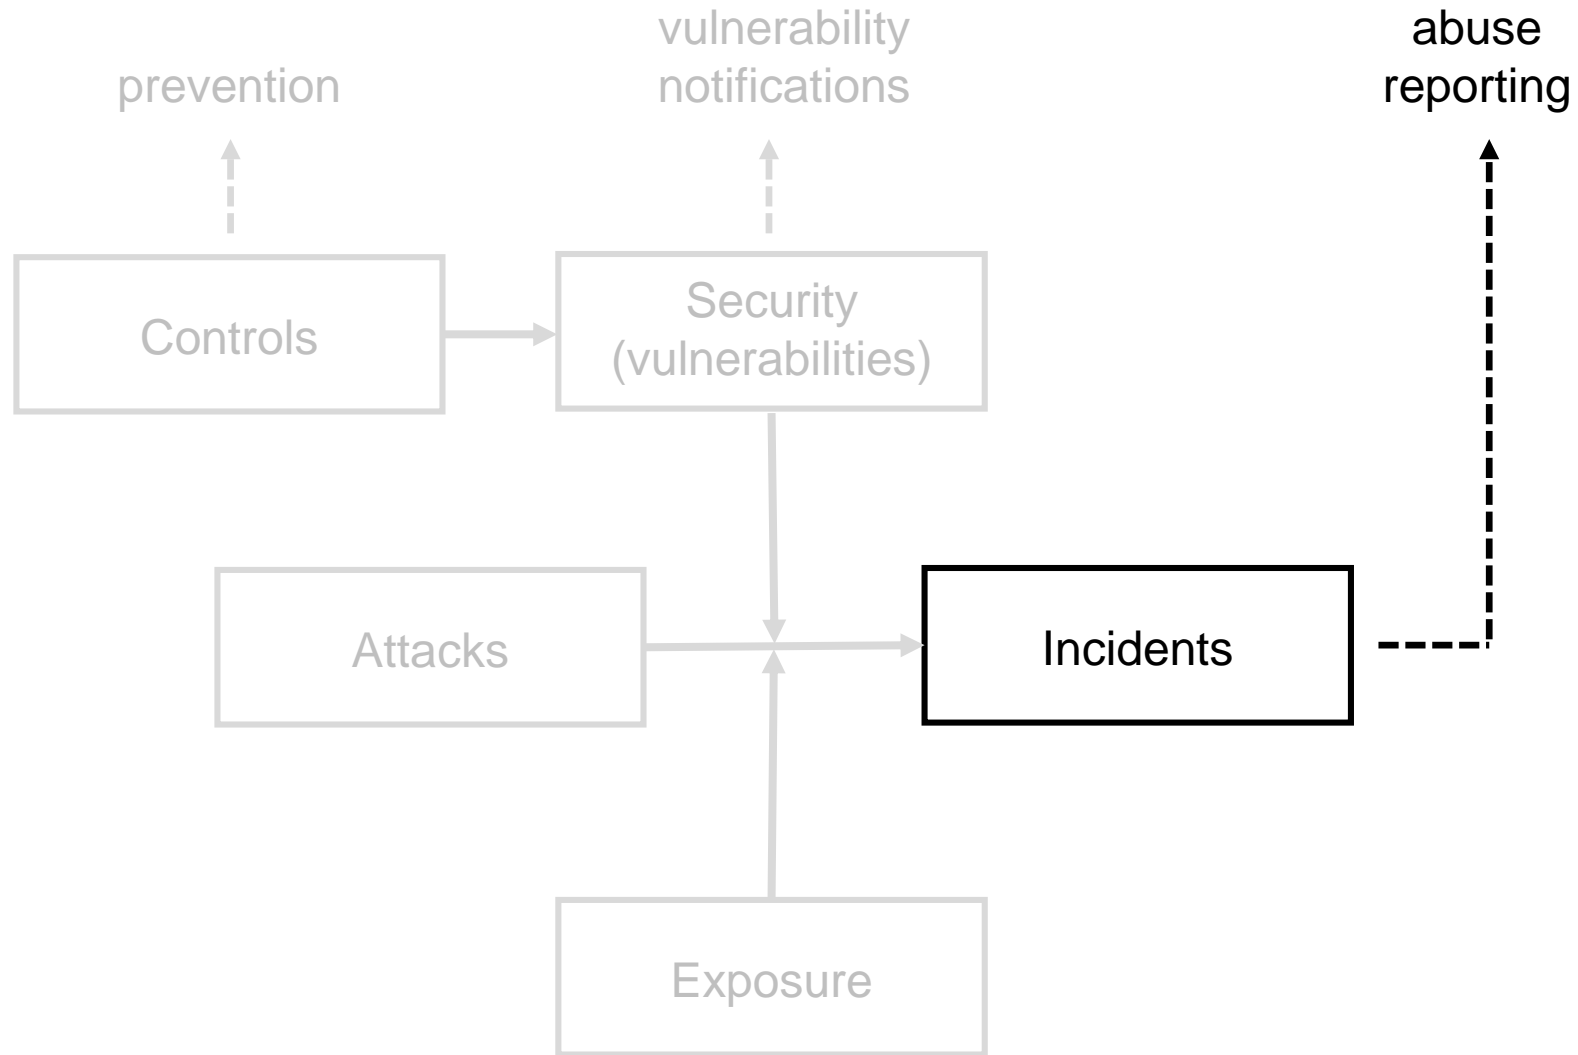




A lot of criminal  
abuse is handled by  
private actors on a  
voluntary basis

How well does  
this work?

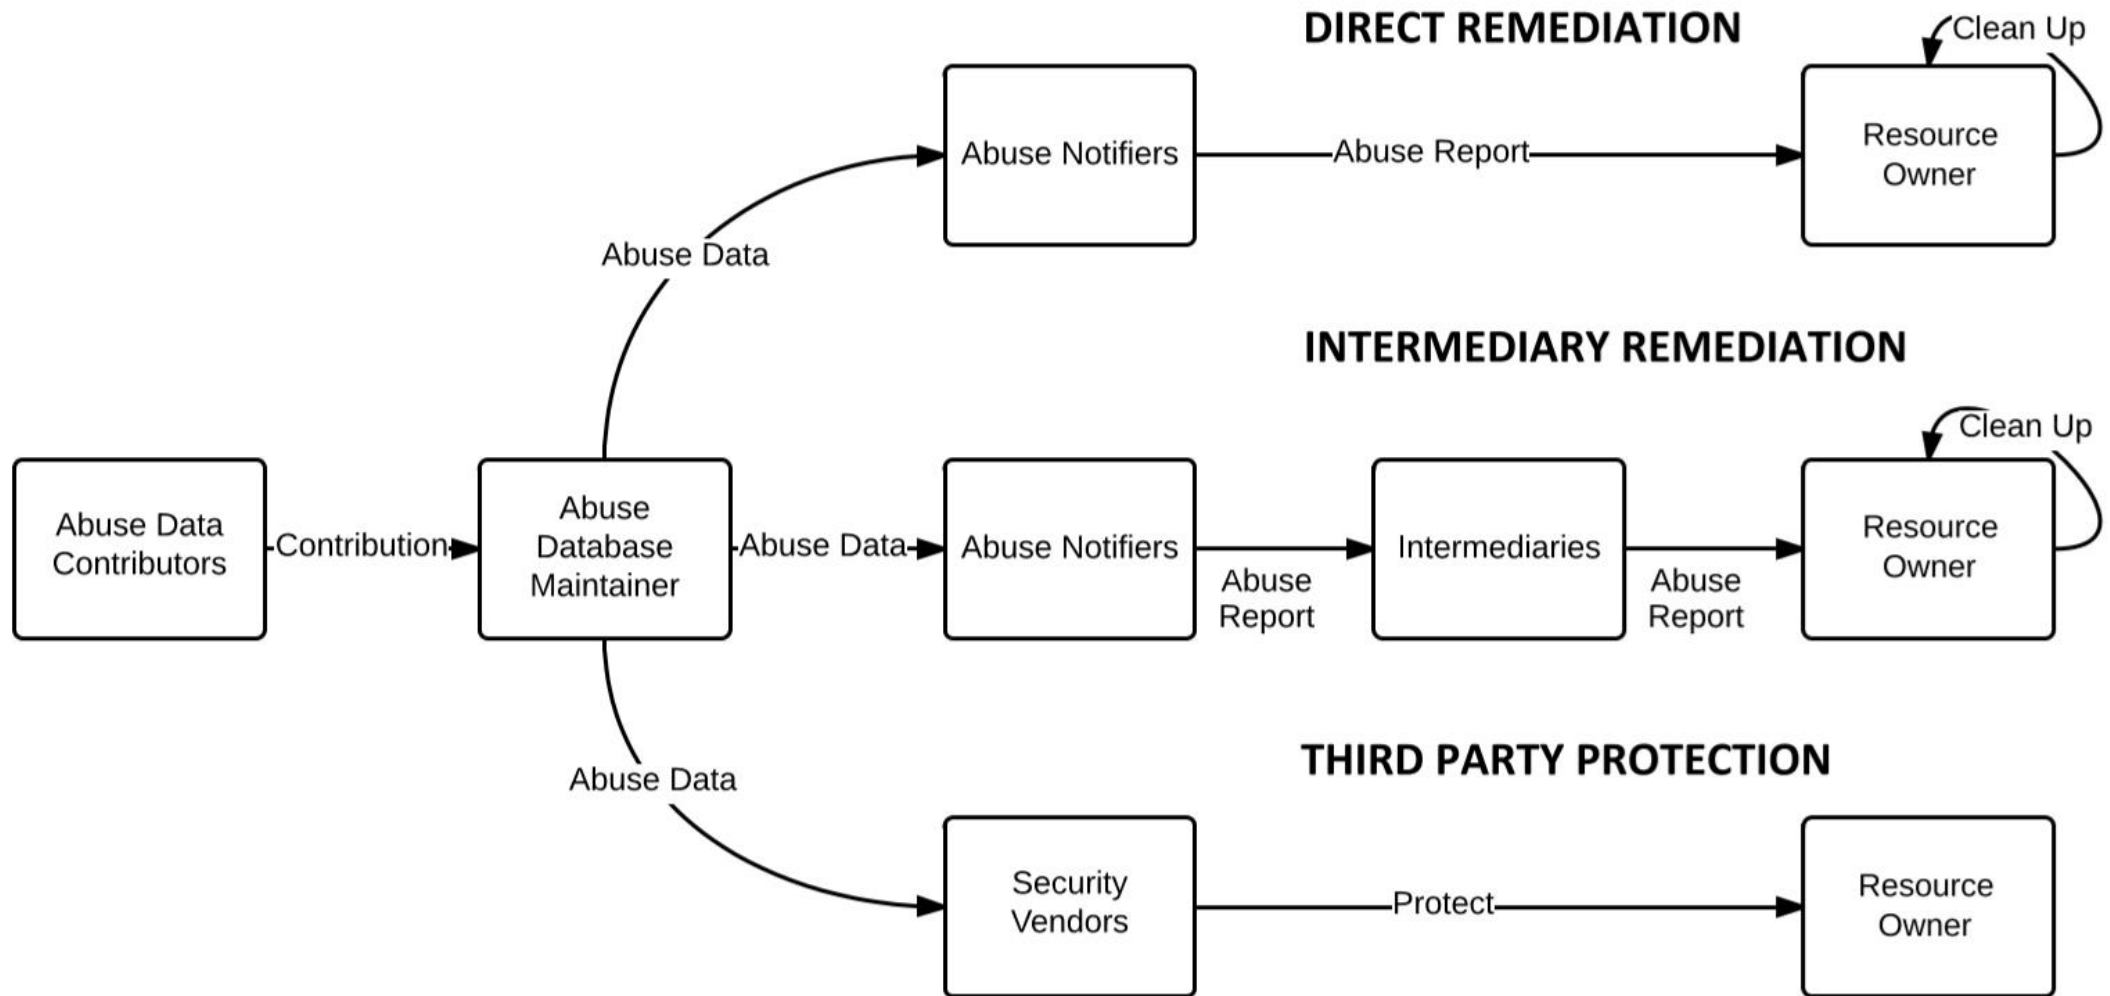






# I. Abuse Reporting







# Cleaning up compromised sites

- Most sites get cleaned by customer or hosting provider after receiving abuse report
- How to make abuse reporting more effective and reduce compromise levels?
- New experimental research (WEIS, USENIX, WWW...)

## Understanding the Role of Sender Reputation in Abuse Reporting and Cleanup

Orcun Cetin\*, Mohammad Hanif Jhaveri<sup>†</sup>, Carlos Gañán\*, Michel van Eeten\*, Tyler Moore<sup>‡</sup>

\*Delft University of Technology, Faculty of Technology, Policy and Management  
{f.o.cetin, c.h.g.hernandezganan, m.j.g.vaneeten}@tudelft.nl

<sup>†</sup>Southern Methodist University, Computer Science and Engineering Department  
{mjhaveri@alumni.smu.edu, tylerm@smu.edu}

**Abstract**—Participants on the front lines of abuse reporting have a variety of options to notify intermediaries and resource owners about abuse of their systems and services. These can include emails to personal messages to blacklists to machine-generated feeds. Recipients of these reports have to voluntarily act on this information. We know remarkably little about the factors that drive higher response rates to abuse reports. One such factor is the reputation of the sender. In this paper, we present the first randomized controlled experiment into sender reputation. We used a private dataset of Asprox-infected websites to issue notifications from three senders with different reputations: an individual, a university and an established anti-malware organization. We find that our detailed abuse reports significantly increase cleanup rates. Surprisingly, we find no evidence that sender reputation improves cleanup. We do see that the evasiveness of the attacker in hiding compromise can substantially hamper cleanup efforts. Furthermore, we find that the minority of hosting providers who viewed our cleanup advice webpage were much more likely to remediate infections than those who did not, but that website owners who viewed the advice fared no better.

### 1. INTRODUCTION

Advances in detecting and predicting malicious activity on the Internet, impressive as they are, tend to obscure

and recipient. This voluntary action is an under-appreciated component of the fight against cybercrime.


Remarkably little research has been undertaken into what factors drive the chances of a recipient acting upon an abuse report (notable exceptions are [1]–[4]). One factor, the reputation of the sender, clearly plays an important role in practice. Not all reports are treated equal, as can be seen from the fact that some recipients assign a trusted status to some senders ('trusted complainer'), sometimes tied to a specific API for receiving the report and even semi-automatically acting upon it.

The underlying issue is a signaling problem, and therefore, an economic one. There is no central authority that clears which notifications are valid and merit the attention of the intermediary or resource owner. This problem is exacerbated by the fact that many intermediaries receive thousands of reports each day. One way to triage this influx of requests for action is to judge the reputation of the sender.

We present the first randomized controlled experiment to measure the effect of sender reputation on cleanup rates and

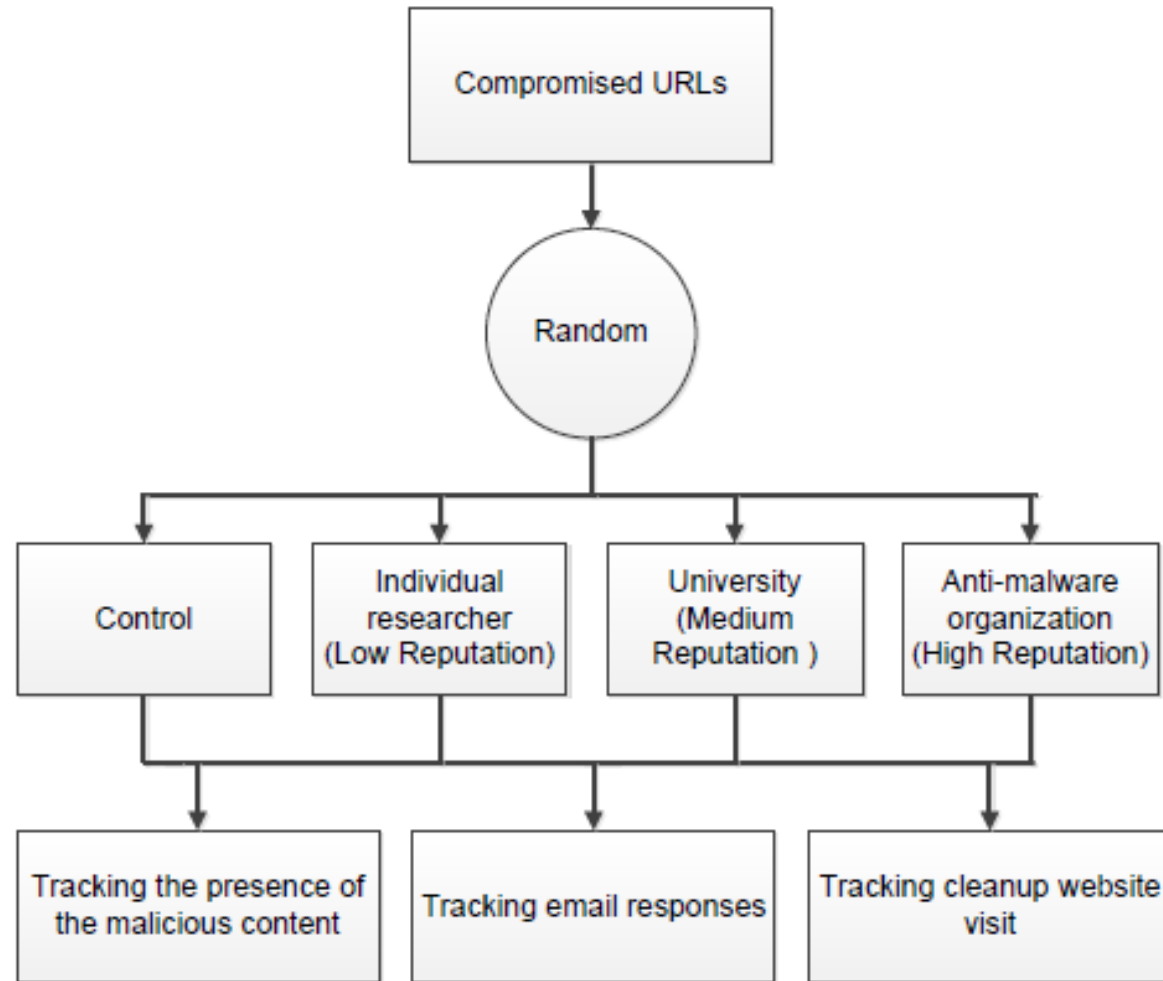
# Asprox compromised servers

- Active since 2007
- Uses thousands of compromised websites for spreading malware and redirects to phishing websites
- Deploys countermeasures to tracking and takedown
  - Centralized IP based blacklisting
  - Only serves malware to certain User-Agents
  - Fake error messages to suggest malicious URL is removed




**You have exceeded the maximum number of downloads allowed for your IP. Please try again later.**

# Experimental design



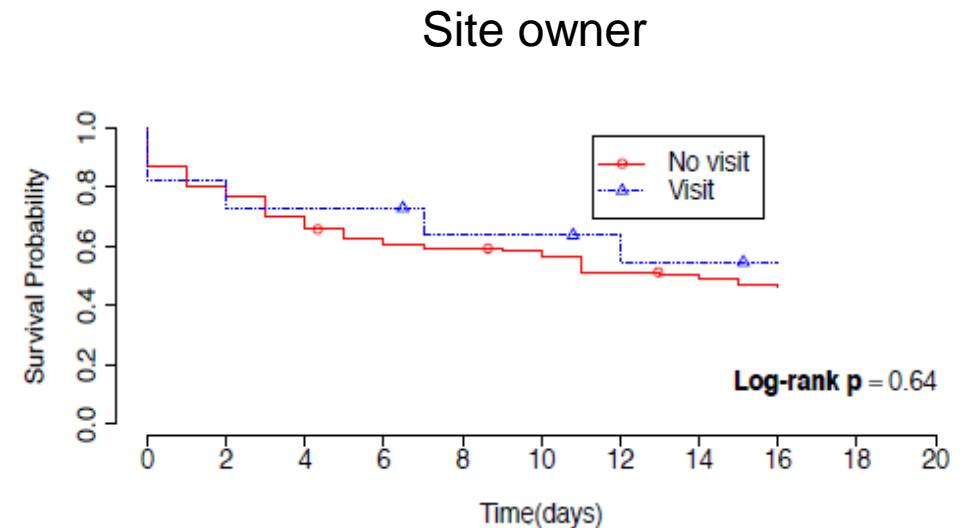
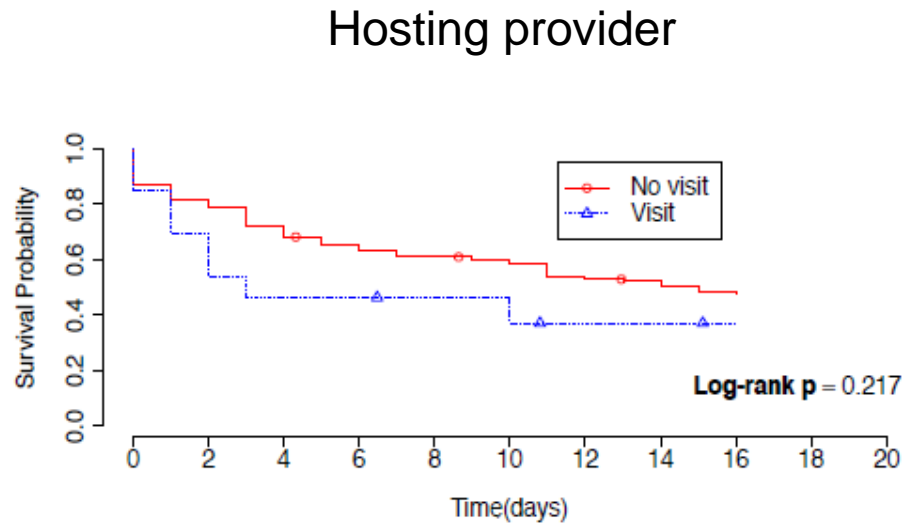


# Does sender reputation matter?




- Treatment groups have similar remediation rates (44%-49%)
- Reputation of the sender did not significantly affect cleanup

# Does cleanup advice help?



- Only 9% of the hosting providers and 7% of the site owners visited our cleanup advice website
- Unlike site owners, hosting providers that visited the site achieved higher cleanup rates

# Do hosting providers make a difference?



- Some providers do substantially better than others, from barely any cleanup to total removal
- Suggests discretion: provider policies make a difference

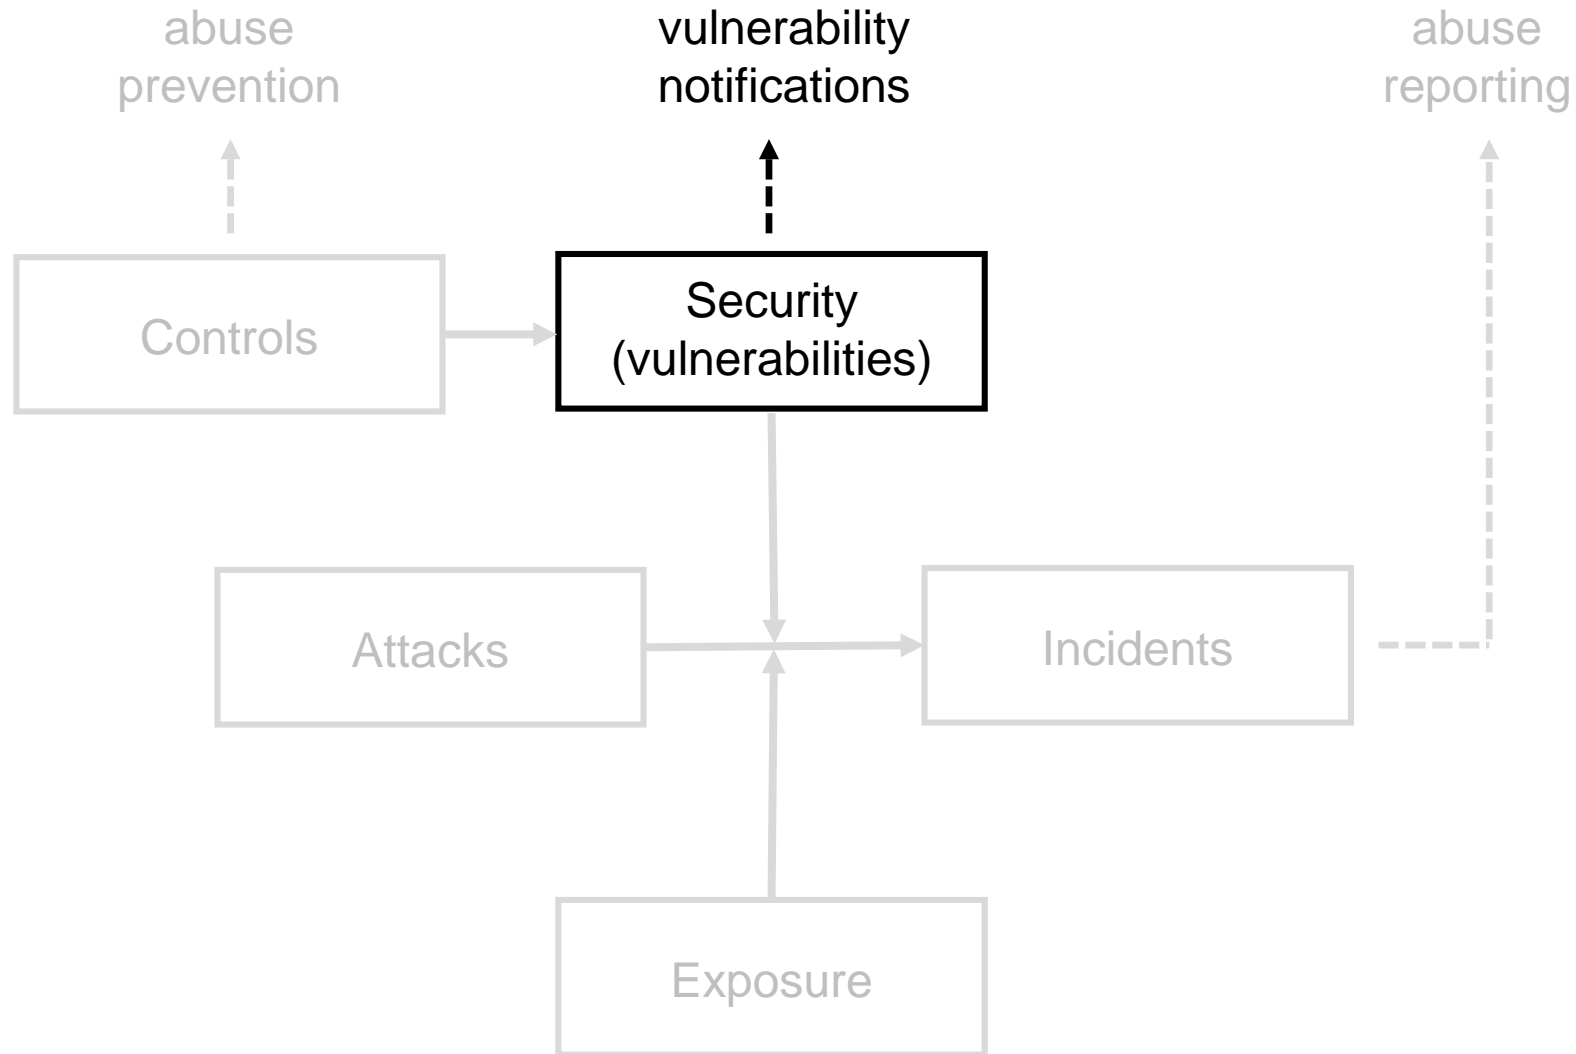


# Some lessons from related work

- ~30-60% hacked sites cleaned up in two weeks after notification
- Open channel to resource owner (e.g., Google console) is most effective (Li et al 2016)
- Full technical report works better than short report with key info (Vasek and Moore 2012)
- Getting ISPs to clean up infected customers shows high variance, orders of magnitude difference in infection rates
- Effective incentives: soft regulatory pressure, benchmarking, reduced cost (e.g., centralized clearinghouse, automatic quarantine)

## II. Vulnerability Notifications







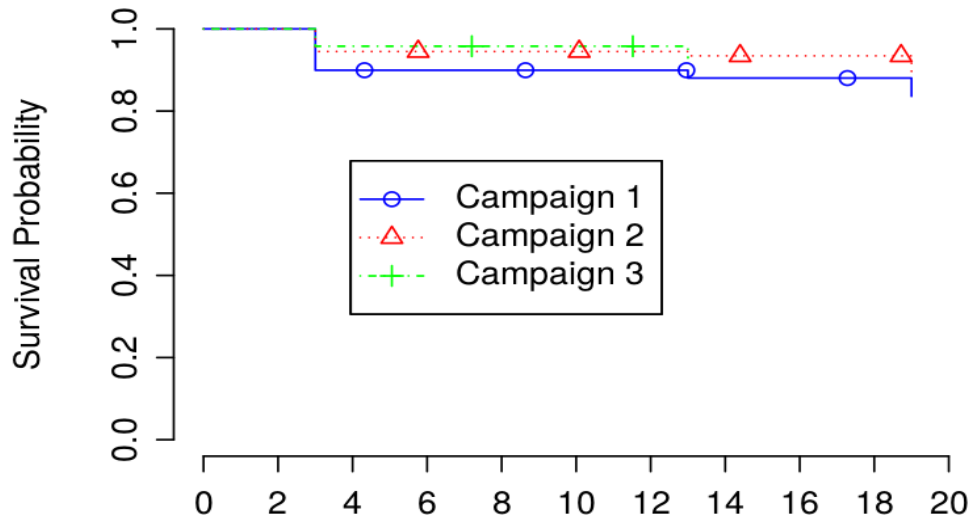

# Age of ZMap and Shodan

- Finding vulnerable devices/systems at scale has become cheap
- How can you reach resource owners at scale?
- Which channel contains the strongest incentive for remediation?
- What factors make notifications more effective?

# How to reach relevant actor at scale?

- Follow standards (RFC 2142, IP WHOIS abuse mailbox, domain WHOIS registrant email)
- Different degrees of failure for different mechanisms
- Network operators are the most reachable, but are further removed from the resource

Campaign	Treatment type	Number of emails sent	Rate of undelivered emails
1	Demonstration	669	70.40%
	Conventional	657	67.73%
2	Demonstration	940	44.68%
	Conventional	1111	35.64%
3	Demonstration	208	12.01%
	Conventional	209	5.2%

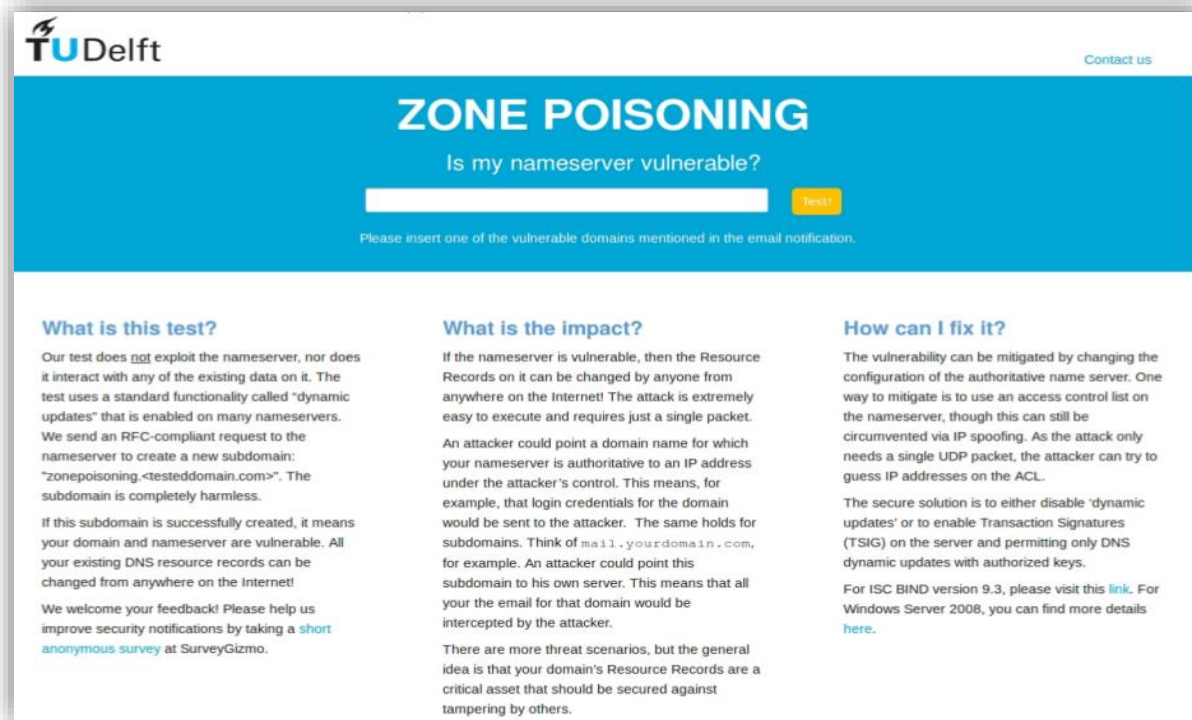


## Which channel mobilizes the strongest incentive for remediation?

- All notified groups did better than the control group
- Still, overall remediation rates were low
- No clear difference between the channels



# Does it help to demonstrate the vulnerability?



The screenshot shows a web page from TU Delft titled "ZONE POISONING". It features a blue header with the TU Delft logo and a "Contact us" link. Below the header, the title "ZONE POISONING" is displayed in large white letters, followed by the question "Is my nameserver vulnerable?". A white input field is provided for the user to enter a domain, with a yellow "Test!" button to its right. Below the input field, a small note says "Please insert one of the vulnerable domains mentioned in the email notification." The main content area is divided into three columns: "What is this test?", "What is the impact?", and "How can I fix it?". Each column contains detailed text explaining the test, the potential impact of a successful attack, and mitigation strategies.

**What is this test?**

Our test does not exploit the nameserver, nor does it interact with any of the existing data on it. The test uses a standard functionality called "dynamic updates" that is enabled on many nameservers. We send an RFC-compliant request to the nameserver to create a new subdomain: "zonepoisoning.<testdomain.com>". The subdomain is completely harmless.

If this subdomain is successfully created, it means your domain and nameserver are vulnerable. All your existing DNS resource records can be changed from anywhere on the Internet!

We welcome your feedback! Please help us improve security notifications by taking a [short anonymous survey](#) at SurveyGizmo.

**What is the impact?**

If the nameserver is vulnerable, then the Resource Records on it can be changed by anyone from anywhere on the Internet! The attack is extremely easy to execute and requires just a single packet.

An attacker could point a domain name for which your nameserver is authoritative to an IP address under the attacker's control. This means, for example, that login credentials for the domain would be sent to the attacker. The same holds for subdomains. Think of `mail.yourdomain.com`, for example. An attacker could point this subdomain to his own server. This means that all your email for that domain would be intercepted by the attacker.

There are more threat scenarios, but the general idea is that your domain's Resource Records are a critical asset that should be secured against tampering by others.

**How can I fix it?**

The vulnerability can be mitigated by changing the configuration of the authoritative name server. One way to mitigate is to use an access control list on the nameserver, though this can still be circumvented via IP spoofing. As the attack only needs a single UDP packet, the attacker can try to guess IP addresses on the ACL.

The secure solution is to either disable 'dynamic updates' or to enable Transaction Signatures (TSIG) on the server and permitting only DNS dynamic updates with authorized keys.

For ISC BIND version 9.3, please visit [this link](#). For Windows Server 2008, you can find more details [here](#).

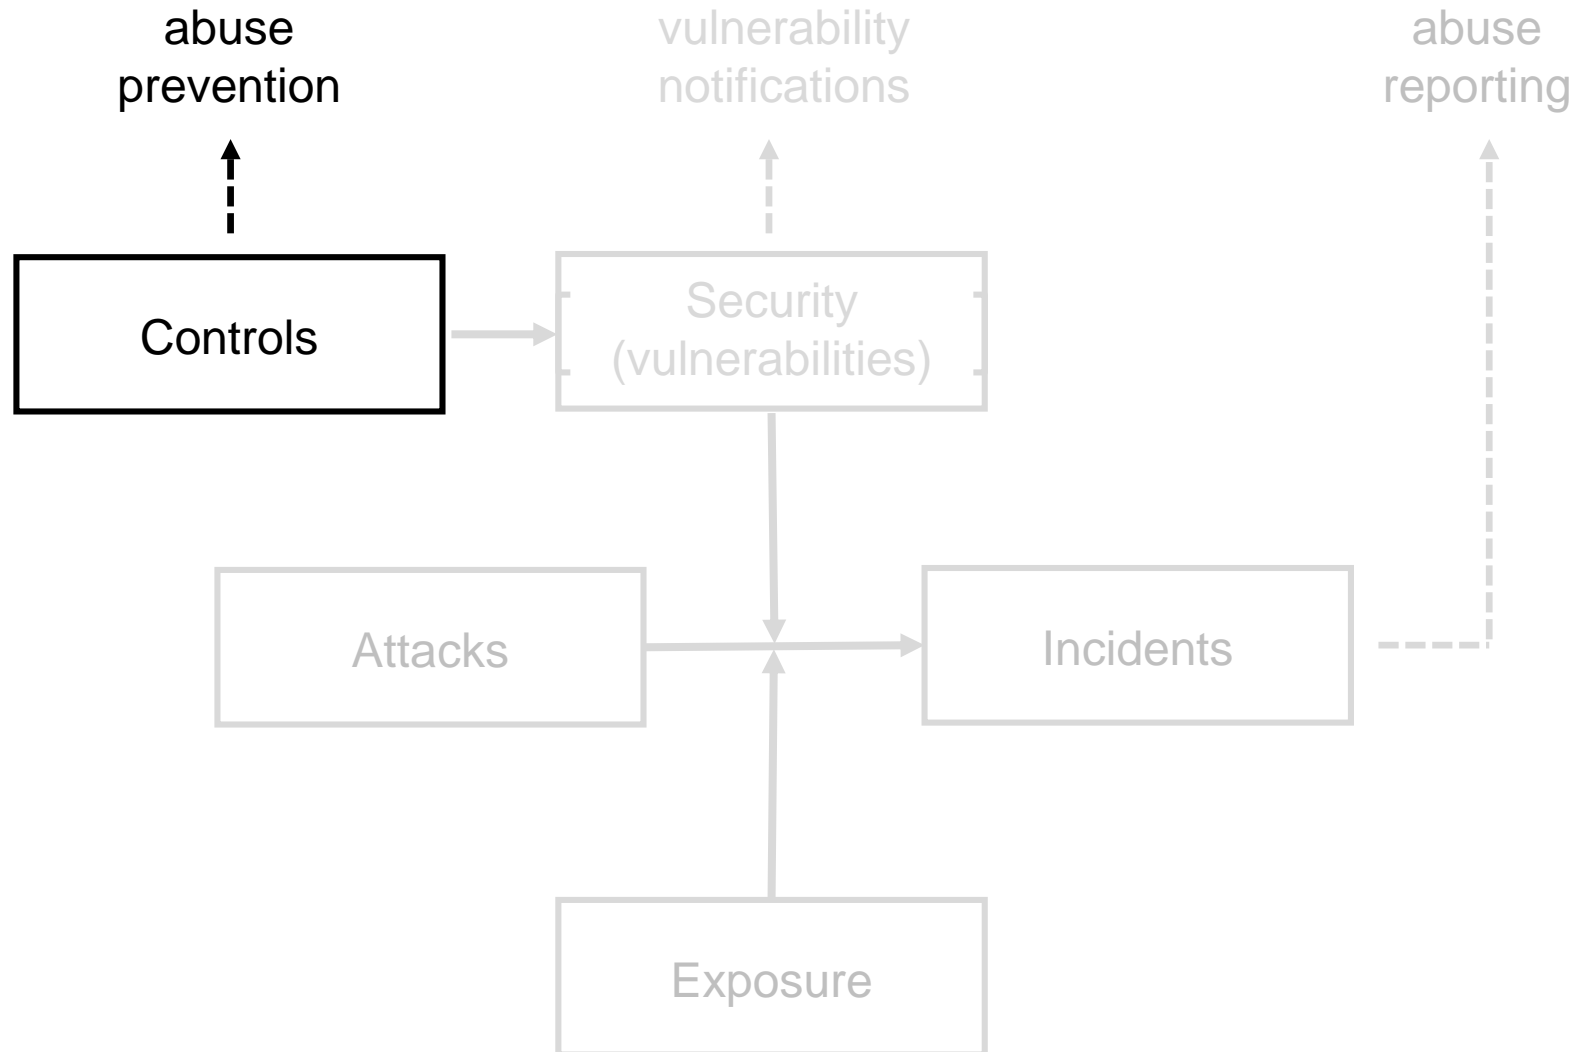
- Short answer: no.

# Some lessons from related work

- No good mechanism to distribute wealth of vulnerability data
- Or to incentivize remediation
- Similar problems with poor reachability and low remediation rates reported by Li et al. (2016) and Stock et al. (2016)
- CERTs don't help
- ...


# III. Abuse Prevention





# Providers adopting best practices

- BCP38 (anti-spoofing) is a cost to the provider, while all benefits go to the rest of the Internet
- The question is not Why aren't some providers adopting BCP38, but Why would anyone adopt it at all?
- Remarkably, lot of providers are compliant. Why? Social norms within provider community (M3AAWG, NANOG, etc)



Source:  
<https://www.caida.org/projects/spoofer/>





# Voluntary action against cybercrime

- ▶ Glass half full...  
Many thousands of compromised machines are cleaned every day
- ▶ Reputation effects help  
Less naming & shaming than benchmarking, a.k.a. correcting self image
- ▶ So do social norms  
Many providers do adopt good practices
- ▶ Better mechanisms  
Reduce friction, solve reachability, clearinghouses and exchanges
- ▶ Role for governments?  
Pressure concentration points, soft regulation, duty to care, liability
- ▶ Externalities from the long tail  
Lack of incentives, lack of accountability, out of reach



Thank you!

More info:  
[m.j.g.vaneeten@tudelft.nl](mailto:m.j.g.vaneeten@tudelft.nl)

# More info on underlying studies

- M. Korczynski, S. Tajalizadehkhoob, A. Noroozian, M. Wullink, C. Hesselman, and M. van Eeten, "[Reputation Metrics Design to Improve Intermediary Incentives for Security of TLDs](#)", *IEEE European Symposium on Security and Privacy* (Euro S&P 2017), April 2017
- Tajalizadehkhoob, S., Böhme, R., Gañán, C., Korczyński, M., & Van Eeten, M. (2017). [Rotten Apples or Bad Harvest? What We Are Measuring When We Are Measuring Abuse](#). *ACM TOIT*
- Tajalizadehkhoob, S., Gañán, C., Noroozian, A., & Van Eeten, M. (2017). [The Role of Hosting Providers in Fighting Command and Control Infrastructure of Financial Malware](#). In *12th ACM Asia Symposium on Computer and Communications Security (AsiaCCS 2017)*, Abu Dhabi, April 3-8, 2017.
- Jhaveri, M. H., Cetin, O., Gañán, C., Moore, T., & Eeten, M. V. (2017). [Abuse Reporting and the Fight Against Cybercrime](#). *ACM Computing Surveys (CSUR)*, 49(4), 68.
- Lone, Q., Luckie, M., Korczyński, M., & van Eeten, M. (2017). [Using Loops Observed in Traceroute to Infer the Ability to Spoof](#). In *International Conference on Passive and Active Network Measurement* (pp. 229-241). Springer.
- van Eeten, M., Lone, Q., Moura, G., Asghari, H., & Korczyński, M. (2016). [Evaluating the Impact of AbuseHUB on Botnet Mitigation](#). *arXiv preprint arXiv:1612.03101*.
- Asghari, H. Cybersecurity via Intermediaries: [Analyzing Security Measurements to Understand Intermediary Incentives and Inform Public Policy](#). *Diss. TU Delft*, Delft University of Technology, 2016
- Tajalizadehkhoob, Samaneh, Maciej Korczynski, Arman Noroozian, Carlos Gañán, and Michel van Eeten. "[Apples, Oranges and Hosting Providers: Heterogeneity and Security in the Hosting Market](#)." In *IEEE Network Operations and Management Symposium (IEEE-NOMS 2016)*, Istanbul, 25-29 April 2016
- Asghari, Hadi, Michel JG van Eeten, and Johannes M. Bauer. "[Economics of Fighting Botnets: Lessons from a Decade of Mitigation](#)." In *IEEE Security & Privacy* 5, 16-23, 2015.
- Noroozian, Arman, Maciej Korczynski, Samaneh TajalizadehKhoob, and Michel van Eeten. "[Developing security reputation metrics for hosting providers](#)." In *Proceedings of the 8th USENIX Conference on Cyber Security Experimentation and Test*, pp. 5-5. USENIX Association, 2015.
- Asghari, Hadi, Michael Ciere, and Michel JG Van Eeten. "[Post-mortem of a zombie: conficker cleanup after six years](#)." In *24th USENIX Security Symposium (USENIX Security 15)*, Washington DC. 2015.