Peer(ing) Pressure:
Achieving Social Action at Scale in the Internet Infrastructure

Ben Collier, University of Edinburgh and Richard Clayton, University of Cambridge

Abstract

We evaluate a rare successful intervention in the management of Internet infrastructure – an anti-spoofing campaign which has achieved genuine traction against an issue that has dogged the network engineering community for more than thirty years. While much scholarship in the security literature has sought to establish the perverse commercial incentives frustrating action against cybercrime and identify possible ways to alter these, in this case we observe a community acting to short-circuit them entirely. We develop the concept of *infrastructural capital* to explain how key actors were able to relocate the issue of spoofing away from the commercial incentive structures of a decentralised community of competing providers with little motivation to solve the issue and into the incentive structures of a far more densely networked and centralised professional community of network engineers. This extends previous work applying theory from infrastructure studies to cybercrime economies, developing a new account of how power can be asserted within infrastructure to achieve change, apparently against the grain of other long-standing incentives.

1 Introduction

The field of security economics is populated by a wealth of examples that show things going badly for the health and security of the global Internet. Users sacrificing privacy and security for minor convenience; market actors failing to coordinate vital mitigations; regulation and insurance setting up perverse incentives; legislation perpetually out of date; and protracted inaction on problems and vulnerabilities are all widely bemoaned within industry, policy, and academic study.

In this context, we identify a rare 'win'; a recent and transformative case in which a major security fix to core Internet infrastructure has been achieved against the flow of individual incentives, market structures, and economic imperatives. Despite these barriers, the decentralised global network of autonomous systems has managed to undertake concerted global action against spoofing – the mechanism underpinning a large part of the market for Denial of Service attacks – resulting in the biggest step-change in effective disruptive action in the cybercrime-as-a-service ecosystem in the past two decades. Exploring this case in detail, we try to determine how it was possible to 'move the needle' – achieving wide-scale coordinated change in the decentralised private markets which dominate the security and management of the core Internet infrastructure.

We combine our own experience of events with auto-ethnographic research and interviews with three key players in the intervention to produce an account and analysis of the main dynamics at play. Interviews were conducted in line with standard principles for qualitative fieldwork and all ethnographic contributions conformed with best ethical practices as required by our institutions. Particular care was taken with commercial sensitivity for participants, and quotes are presented without explicitly identifying the contributor.

Workshop on the Economics of Information Security (WEIS), 2024.

2 A short history of spoofing

There are many different types of Denial of Service (DoS) attack which can impact the availability of an Internet resource. This paper is concerned with DoS attacks which involve spoofing the source address of Internet Protocol (IP) packets.

Absent any controls it is possible to send an IP packet with a source IP address which is not your own, 'spoofing' the actual sender. The recipient of the packet will consider the sending address to be legitimate and, having swapped source and destination IP addresses, will send a reply to the spoofed address rather than to the true source of the traffic.

A distributed DoS attack (DDoS) is a DoS attack where there are many senders of traffic to the target. A reflected amplified DDoS attack is one where spoofed packets are sent to servers whose response is larger than the request – hence a relatively modest amount of spoofed traffic can be amplified into a very substantial amount of traffic at the target. For example, a Network Time Protocol (NTP) MONLIST request can be just 76 bytes long, but can elicit a response of 46800 bytes. Thus a server generating a 1 Gbit/s data stream of spoofed packets, reflecting these off sufficient NTP servers, can in principle create a 500 Gbit/s DDoS attack – something that only a handful of specialist services can mitigate. It is also possible to produce (rather smaller) amplified attacks using TCP since a server will typically send three ACK responses to a SYN packet.

Spoofing was considered by Morris in 1985 (who was concerned with attacks on TCP handshakes) [59] and Bellovin in 1989 [13] (who was mainly concerned with the ability to spoof TCP connections with non-random Initial Sequence Numbers). In January 1995 CERT issued an advisory on "IP Spoofing Attacks and Hijacked Terminal Connections" which was again concerned with issues that might arise with fully open TCP connections [22]. In January 1996 they warned of a "UDP Port Denial-of-Service Attack" involving CHARGEN and ECHO services [23].

By September 1996 the TCP advisory had been updated to address the issue of "half-open" connections (referencing proof of concept code in "two underground magazines") where TCP SYN packets are sent with spoofed source addresses in a so-called SYN flood. Servers reply to a SYN with a SYN ACK and consume server resources in anticipation of an ACK packet arriving to complete the TCP "three way handshake" to open a connection [24]. When the source address on the SYN is spoofed there will never be a handshake – at best the spoofed source replies with an RST packet ("I wasn't expecting you to talk to me") and the resources are freed, at worst the server keeps 'state' for a long period until the (presumed) connection attempt times out. This attack can be mitigated with 'SYN cookies' [14] or a 'SYN cache' [52]; Korn surveys the residual issues with these approaches and others [48].

In 1998 Ferguson and Senie wrote RFC2267 on "Network Ingress Filtering" which says that spoofing should not be permitted – instructing ISPs and hosting companies to configure the devices at the edge of the network to only be able to use the source IP address(es) they had been assigned [39]. The document mainly addresses the issue of SYN floods, but also mentions the UDP issues with CHARGEN and ECHO. RFC2267 was superseded in May 2000 by the almost identical RFC2827 which is better known by its identity in the IETF "Best Common Practice" series: BCP38 [72].

RFC2267 led directly to efforts to optimise and streamline the ability to deploy and maintain what became known as Source Address Validation (SAV). Access Control List (ACLs) were hard to deploy on thousands of ISP customers and so Chandra created Unicast Reverse Path Forwarding (uRPF) to leverage the forwarding tables within routers as a way to check if the source IP is valid [9]. The original RFC3704 was updated in February 2020 by Sriram and Montgomery's RFC8704 [74].

There was considerable activity by ISPs and some hosting companies from 1998 onwards to use the newly available SAV features in their network devices to prevent spoofing. However, a 2005 paper by Beverly and Bauer describing their 'Spoofer Project' estimated that around a quarter of all networks still allowed spoofing and that many networks disallowed only some, but not all, spoofing [15]. The Spoofer Project relied on volunteers, often the network engineers trying to deploy SAV, to download the test software and run it on their networks. It was operated by Beverly until 2015 when it was taken over by CAIDA.[1] The Spoofer Project drove a further increase in deployment of SAV, with Greene claiming in 2012 that rollout was now 80% complete (with a much more difficult 20% remaining) [42].

However, in 2013 a longitudinal analysis of the Spoofer reports up to that time concluded that there was no discernible trend in the amount of spoofing, either up or down [16], but the way in which the project collected data meant that reports tended to come in bursts when publicity about the project caused more volunteers to come forward. In 2018 an initiative by Lone et al. increased Spoofer coverage 15% by recruiting over 1500 volunteers in six weeks... the resultant paper discusses the lessons learned in recruiting and remunerating these workers [55].

In 2020 Lone et al. combined Spoofer results with other scan results and then examined whether there were any common characteristics to the sites that permitted spoofing. They conclude that what portion of an ISPs address space permits spoofing is influenced by network complexity, security efforts, ISP characteristics and the institutional environment [54].

The most recent results in the academic literature are found in Dai and Shulman's 2021 paper. They scanned the whole Internet using some new tricks to determine if various servers could be used to show that spoofing was allowed. They found 63522 ASs[2] where spoofing was possible – substantially more than any previous survey – and typically more than three quarters of all the ASs in a particular country [37]. However, unlike the Lone paper, they do not discuss the extent to which spoofing is possible within each AS. This means their numbers should be carefully interpreted: it would be wrong to assume that any end-user on a consumer network can spoof just because an ISP has not locked down their DNS server.

A completely orthogonal approach to preventing spoofing has been attempts to reduce the number of devices that will reflect (and amplify) traffic. The rise in importance of reflected amplified UDP for DDoS attacks from around 2013 led to efforts to reduce the number of amplifiers that were available to be used in these attacks. Various efforts were started around that time to perform scans of the Internet and then report the existence of systems that could be used as amplifiers to the networks that hosted them. Best known of these are the scans performed by the Shadowserver Foundation[3] – which makes the results available not only to the hosting network but also to the relevant national CERT. The German CERT has published a report showing how efforts by German ISPs and hosting companies reduced the number of reflectors by 75% or more for a number of relevant protocols over a six year period [20].

3 The DDoS traceback initiative

The size of DDoS attacks grew over the years and by 2018 only specialist anti-DDoS providers could be certain of being able to cope with them. In February 2019 Barry Greene, who had many years involvement in initiatives to improve Internet security, organised an invite-only full day workshop on Denial of Service topics. He was taking advantage of both M3AAWG (Messaging Malware Mobile Anti-Abuse Working Group – an industry group tacking abuse of the Internet) and NANOG (North American Network Operators Group – a industry group for networking professionals) holding meetings in San Francisco in the same week. This meeting of many of the most active people in anti-DDoS work sowed the seeds of significant later action.

In June 2019 Damian Menscher, who had been at the February event, gave a talk at the next NANOG meeting in which he set out a new approach to dealing with reflected amplified DDoS attacks [57]. He observed that for many types of attack there were so many amplifiers that even a totally implausible

[1] https://www.caida.org/projects/spoofer/

[2] An AS is an Autonomous System which can be thought of as an ISP, hosting company or network infrastructure provider. With notable exceptions (in both directions) there is one-to-one relationship between an AS and these entities.

[3] https://www.shadowserver.org/

99.9% success rate in fixing them would not mitigate the problem. He also pointed out the incentive issues with BCP38, in that preventing spoofing helps others rather than yourself. He then suggested that "problem" sources of spoofed traffic could be identified by examining Netflow/Sflow data (sampled information about traffic flows across networks) or even just packet counters on routers. What would be looked for would be the flows of spoofed traffic on the way to the reflectors. These flows could be, in principle, be traced back across networks to their source and mitigation then applied.

Menscher's proposal did not gain immediate traction but in the summer of 2020 there were a series of high profile DDoS extortion attacks on banks and stock exchanges across the world. Although many targets were not affected there were short-lived banking outages in several countries and the New Zealand stock exchange had to suspend trading on several days [77] – and in due course, their regulator issued a damning report on their lack of preparation for such attacks [60].

Greene had been developing ideas for some new community initiatives to tackle DDoS, looking to use the Anti-DDoS group within M3AAWG as a focal point. In August 2020 Greene talked with Richard Clayton, a Cambridge University academic with a longstanding interest in DDoS measurement, who had also been at the February 2019 meeting. Clayton suggested a scheme for organising traceback along the lines that Menscher had suggested. Each week he would use data from his DDoS monitoring network to provide a short list of the biggest reflected amplified DDoS attacks from the previous week, and network operators would be invited to inspect their Netflow (etc.) to trace back the source of the spoofed traffic. The key idea was that instead of trying to deal with every source of spoofed traffic at the same time, attention would be concentrated on the most important sources of abuse and a 'rhythm of action' would develop where small amounts of effort each week from the participants in the trace back would lead to change.

Tracebacks started in earnest in April 2021 using a private chat channel on Slack. The initiative was initially kept on a need-to-know basis because a traceback had been performed by the FBI during the previous summer's extortion attacks and it was felt unwise to inform the criminals, who perceived themselves as impossible to find, how easily their infrastructure could be located. Over time, as it became clear that this was not a tool that Law Enforcement would be regularly using, the initiative was allowed to become more widely known.

After some initial variation, the procedure became that shortly after midnight UTC on a Monday morning Clayton would post details of attacks using the UDP protocols CHARGEN (still an issue despite the CERT advisories from 25 years earlier), DNS, NTP, LDAP, PORTMAP and SSDP. The attacks chosen were the two largest attacks for each of the protocols (by packet count hitting his sensors) and fastest attacks (by rate of packet arrival over the attack period). The attacks considered were for the previous four days, so that a one week log retention period (thought likely to be implemented by some networks) would not prove an impediment. The use of Clayton's data in an objective manner was intended to avoid any impression that any particular DDoS targets were getting favourable treatment by having their attackers traced – and the use of 'big' and 'fast' was intended not only to make people feel that important attacks were being investigated, but also to increase the chance that data would be visible in Netflow logs.

4 An expectation of limited success

Achieving communal action in the commercialised and (nominally) decentralised infrastructures of the Internet has long been recognised as a major problem [43, 70]. The competitive and marketised nature of the provision of Internet services like hosting and network connectivity have generally been characterised as frustrating communal action through well-established incentive structures. Where issues affect the whole Internet, this competitive design can often create a hypothecated zero-sum game, in which companies are incentivised to pursue their own benefits rather than global goods [3]. Where abuse is directed outside a service, rather than internally, there is often little incentive to tackle it directly. The security economic and social scientific literature replicates this account, suggesting that in many cases, action against cybercrime is frustrated by incentive structures – whether in insurance [76, 81], security behaviours [58], patching [7, 69], passwords [47, 66], privacy behaviours [2], investment in security [44] or abuse moderation [79]. However, as we argue throughout this paper, incentives do not only operate at the level of commercial relationships, individual users, or material infrastructure. Underneath this, much of the Internet backbone relies on rather more human relationships within professional communities, which, though not free from their own incentives, can be shaped by rather different dynamics.

Concern about incentives meant that when the traceback project started there was limited expectation that it would have very much success. A hosting company that permitted spoofing, perhaps inadvertently, would attract miscreants that would send traffic that damaged other parts of the Internet, but the impact on the hosting company would be negligible. There would be rather more outgoing traffic than normal, but that might well have to be paid for. Similarly, a transit network that provided connectivity to the hosting company would be able to charge for the traffic – and the victims would generally be elsewhere, so again there was no incentive to look for evidence of spoofing, let alone to filter the traffic (assuming that the relevant network gear was capable of this) or threaten to terminate customer relationships if no action was taken to curb it. As one of our interviewees colourfully put it:

> So basically people wrote this RFC saying that spoofing should not be allowed, and they expected everybody to fix this. And basically, they didn't, because of economics, as usual. Which is, why would you fix it? Because it doesn't damage you, people are using your machines, people are coming onto your network, buying services off you, giving you money, and then attacking somebody else. So why should you care? Now those hosting providers fall into two categories. Well, they fall into several categories, but a simple categorisation is, they don't know what they're doing, and therefore people come along and the hosting provider doesn't notice that they're spoofing. And then you have the people who know damn well you're spoofing, and they're basically making that a feature and possibly charging a premium for it.

Peering – the provision of network connectivity – is a complex ecosystem, as described by Hall et al. [45]. Hosting companies and ISPs can purchase connectivity from 'transit providers' or they can 'peer' directly with other networks, usually at Internet Exchange Points (IXPs), and exchange traffic destined for each other's networks for free. Peering has evolved dynamically over time, with significant differences in approaches and policies over the years. In general, the trend noted by the literature has been that Internet interconnection evolved from an 1980s hierarchical model based initially around academic and military providers passing traffic to and from the NSF backbone, to a 1990s model centralised around dominant major telcos, then in the 2000s to a far more decentralised, flatter, and more 'rhizomatic' model of densely interconnected thickets of peering relationships [49]. Peering is generally a win-win for the participants because any traffic going over the peering link does not have to go over expensive transit connections. Once a network has set up a link to an IXP they will generally peer with as many networks as possible,[4] to save as much money as possible, and peering will generally occur provided one of the companies doesn't think it more appropriate for the other to become a customer and pay.

Significant attention has been given to the business relationships and economic incentives which drive these partnerships [61, 62], but far less has focused on the social dynamics of those who actually agree these relationships. Peering arrangements are seldom set up with formal contracts, but ad hoc by in person handshakes at IXP meetings or even just by swapping emails. This creates some incentives to be seen as 'competent' and 'one of the good guys' when dealing with network abuse – such as permitting spoofing. At one level higher, the transit networks themselves peer with one another in order to establish access for their clients to the global Internet. Although the commercial relationships involved in inter-Tier 1 peering are not generally public knowledge, they too involve a significant social component.

Hence there was some expectation that the peer pressure (with apologies for the pun) generated by the publication of the weekly list, concerted action by the members of the spoofing group, and leveraging the informal networks that underpin peering work would mean that there would be some limited incentives for

[4] For example, as reported by Norton in conversation with Brokaw Price, in the early 2000s, Yahoo's peering policy was a single word – "Yes!" [63]

network engineers to want to do something about spoofing when told about it by someone they knew. It was also expected, however, that this would sit in tension with the sales and marketing people (representing the interests of the provider itself) who would expect some customers to leave if they could no longer spoof. In general, actions relying on appealing to the professional standards of the community (such as setting out RFCs to establish best practice) had not previously succeeded in generating sustained progress against spoofing – the skills and resources involved in detecting spoofing are not widely available to smaller providers, tackling spoofing had not traditionally been a core responsibility of the network engineer role (and there was a significant time cost to individuals in making it one), and there was little downside to not taking action, especially for these providers who were not part of the core social community.

5 What happened in practice

In the event, the traceback initiative has made a significant difference to the prevalence of spoofed DDoS attacks. There is no quantitative data available as to which networks have made changes, and commercial confidentiality (and the ephemeral nature of the communications involved in getting action to be taken) means that it is unlikely that this will ever be known. However, it is possible to measure a change in the type of DDoS attacks that occur. NETSCOUT, a provider of DDoS mitigation services, with data from over 400 of their customer's networks, produces a regular report on DDoS attacks.[5] Figure 1 shows a graph they have shared with us showing that reflected amplified attacks used to be about half of all DDoS attacks, the other half being 'direct path' attacks where traffic is sent directly to the target. That has changed, it is believed because of the traceback work, so that spoofing has dropped to only about a third of all attacks. Although the continued rise in direct path attacks paints an ambiguous picture of the effects of this intervention on DDoS attack levels more generally, this does not detract from the importance of this case – which has shown real success in both achieving a form of communal action long thought near-impossible, and in apparently significantly changing the practices and behaviours of attackers. Now that a strategic 'playbook', a sense of collective efficacy, and a set of collaborative structures have been developed, and success has been demonstrated, the landscape for future interventions looks brighter.

We interviewed three engineers working on the traceback initiative, conducting open-ended sociological interviews discussing the intervention, the motivations of different players, and the experiences and perspectives of key actors. One of the authors of this paper was also involved directly in the intervention as a key coordinating actor, and has contributed significant auto-ethnographic reflections on their experiences to this piece. The interviews were semi-structured and analysed through a qualitative coding process incorporating both inductive (allowing conceptual findings to emerge through the coding process) and deductive (approaching the interviews with a clear analytical framework in mind) aspects. In this case, the inductive process was sharpened by reflections from the auto-ethnographic component, which indicated that Bourdieusian concepts of *social capital* could be potentially relevant. As a result, in addition to more open-ended coding, particular care was taken to identify and assess the role, if any, that capital (in this case, social capital, i.e. the networks of links and relationships between players), habitus (professional practices, knowledge, and values) and the broader field of network engineering might have played in this intervention. The commercial sensitivities involved mean that significant review has been required by some of the corporate actors involved, and the key participants and interviewees gave substantial review comments on this piece before submission.

In the subsequent section, we discuss the reasons why we believe that this intervention has worked, apparently acting against the incentive structures built into the distributed design of the global Internet as well as of the companies who deliver and maintain this infrastructure. We argue that important design features of this intervention have served to shift the arena of action away from the interests of individual companies, instead reorienting it towards the complex networks of social capital – i.e. the human peering networks – on which the infrastructure relies. In these networks, which are much more centralised around

[5] https://www.netscout.com/threatreport

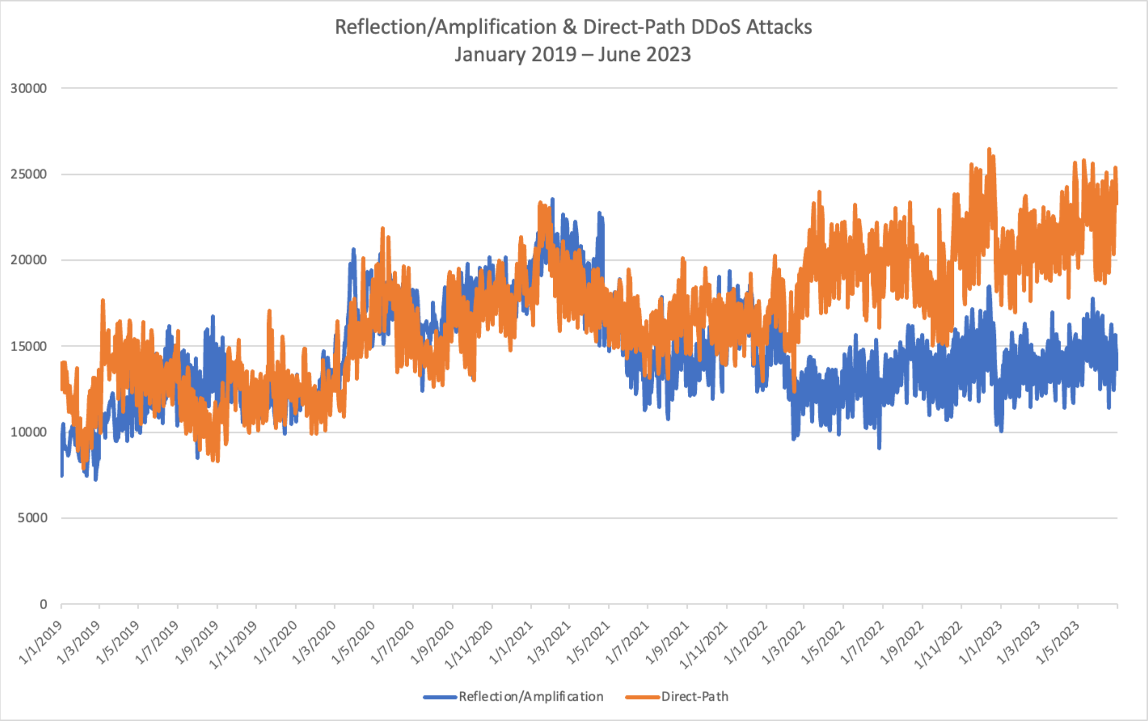

Figure 1: Number of DDoS attacks per day. Blue are reflected/amplified attacks and Red are direct path attacks. Graph courtesy of NETSCOUT

key players with significant interests in common, the incentive structures are substantially different, and permit far more effective, coordinated social action (though not without substantial enabling work by a number of key actors).

5.1 Social capital in the field of network engineering

While the incentives pertaining to hosting companies and service providers are important in shaping the dynamics of the global Internet, the human social dynamics of the people who administer these networks – the network engineers – are rather different. These complex networks of social capital and professional practice create their own topology of 'human peering', and these interpersonal relationships exert significant power in permitting or frustrating action.

Although network engineering has its roots in a long history of work dating back to the earliest foundations of the Internet [11, 50, 51], this array of practices developed into a coherent global professional community across the era of Internet commercialisation following 1993. In the mid and late 1990s, network engineers were the frontline pushing forward the digital frontier, connecting businesses, consumers, and government services up to the global Internet. As they did so, they negotiated peering relationships, established friendships, and developed a core set of principles and practices which still define the network engineering profession.

This shared work of spreading the Internet backbone across the world in the late 1990s contributed to the development of network engineering *qua* professional identity – one with (at the time) significant social capital and the power to shape the wider field of Internet engineering, given the centrality of this work to the expansion of Internet connectivity which was driving much of the associated economic growth. They shared a loyalty to this growing professional community – forged in conferences, shared IRC chats, and less formal settings. Some of this reflected a shared working life and set of challenges; as they developed this new profession,its skills and practices, they were encountering shared problems with hardware, frustrations with their employers, and developing a shared set of core values.

This core set of 'network engineer' values and practices has been far less studied than the related professional networks of system administrators – not least because they are at least an order of magnitude fewer in number – but shares similar commitments to facilitating the free flow of information, promoting security, customer privacy, and the availability of services, and a wider connection to techno-libertarian political ideals. A key distinction, however, relates to the practices of peering. Incentivised as they are to establish and maintain these peering relationships (as they save the company considerable money), and by the need to promote connectivity not only to their own customers but to the entire rest of the world, they are inherently more externally-facing than sysadmins or cybersecurity professionals. The cultivation of these external relationships with their peers is therefore a crucial component of their work. As an interviewee explained:

> Within a large organization of a service provider, there are so many roles that do not require external collaboration with entities on the Internet. But in the networking domain, it's basically required that you have to engage with external organizations. On top of that, there are various escalations that drive having to collaborate with external parties and various internal teams will look to you as the one with the rolodex to facilitate that communication. There may not be clear rules on how you handle that communication, but you use high judgement to determine what information you can share to help diagnose a particular problem. You refrain from sharing specific customer details or configuration, and just focus on the problem at hand.

Remarkably, the core of this community is composed of many of the same people who developed the profession in the late 1990s and early 2000s, who still play crucial roles in the peering ecosystem now. Due to consolidation around a smaller number of big players, the slowing of Internet expansion, technological changes which centralised this work around a number of key firms, and the increased prominence of software engineering (and later data science) as high-status technical professions that were more attractive to young computer scientists, the community has largely aged with the infrastructure. They have generally failed to recruit newer members, remaining central and fairly high-status players within the corporate providers of Internet infrastructure. Outside this dense network of the 'usual faces', there are a panoply of medium- or small-level providers, who tend to be less professionalised, and cultivate peering relationships one-to-one or buy access into larger networks directly. However, they still generally have – by the necessity of interconnection – at least some direct connection into this central network.

Thus, the networks of social capital built across the period in which the Internet backbone was growing across the world still largely persist to this day – and still form a functional part of the contemporary practices of network engineering.

> A lot of those relationships that formed over that era still pay off today. You'll find that you'll need to engage with a particular company and can call upon folks you've dealt with in the past to help with something. Even if it may not be directly networking related, it could be some common customer issue or challenge where you require the help of someone that you worked at before to help navigate an organization to get to the right people. For example, finding a team that handles DNS, abuse or some other particular issue. The result is that you leverage those relationships to help solve an issue and get to the right resources.

Within this profession, key social events such as the NANOG and M3AAWG conferences enabled the developing social links and self-image of the network engineer to proliferate. These settings facilitate what Coleman describes as the "celebration of a lifeworld" [32] – and the strengthening of these crucial bonds of trust and reputation. This allowed these spaces to also become an informal social infrastructure for facilitating the collaborative labour which made this peering work in practice. They also, along with workshops like APRICOT, provide an important space for bringing new generations of network engineers into the profession, linking them up to these networks of social capital, skills, and core values.

> So ok, you're coming in from KDDI, this is your first time at a NANOG. Here, here's Vince Fuller, Vince Fuller is one of the old timers of the Internet and you're going to go around with Vince. I convinced Vince, and so you're going to go around and… hang out with Vince

> and let Vince know if you need to meet people to peer with. Right? And so that is now the newcomers, and so you know, IETF and RIPE and NANOG has the newcomers. And as part of the newcomer sign-up, I'm going to go to a newcomers meeting, so you have to like, be paired up. So that's all institutionalised, to do that sort of stuff. Because we know this works.

Crucially, these spaces permitted actions nominally going against the incentives or policies of their employers, but which worked for the benefit of the network engineering profession, and hence (by extension) the Internet as a whole.

> So at forums like NANOG, you'd have a lot of informal collaboration with your peers within the industry. Some folks would refer to it as the hallway track or beer track, where you could discuss topics and be a bit less guarded. For example, folks would talk about common challenges with specific types of network hardware and OS versions and learning what things were or weren't working well. Other folks would talk about specific networks that they were having challenges with from an operational and support perspective. Examples of that may include which networks were running congested in certain locations or to specific networks. You also got to learn about what people or networks mis-represented themselves about their infrastructure as well and the techniques and methodology on how to audit how a network operates and what their relationship was with their peers. It was also helpful to understand the truth behind outages and impairments with the understanding that you wouldn't share that broadly, which was valuable because you could learn from those experiences to help make yourself operate better.

This practice – of going against the immediate motivations of one's employer in service of a wider community, longer-term economic incentives around accessibility, and a wider set of global public goods around the free flow of information – has long been a feature of the core community of network engineers. This partly reflects the nature of the work – the availability of assets across the entire Internet benefits each individual provider, and hence the incentive structures are generally more communitarian than might be expected from the competitive and marketised design of Internet services.

Thus, the closeness of the social networks which characterise this ecosystem suggested that communal action against spoofing might indeed be possible. However, two issues remained which may have accounted for the initial lack of movement. Firstly, there was a need for a concrete intervention to make this an issue for the profession as a whole, not just an avid community of anti-spoofers with an RFC document. Secondly, there was the issue of spreading this action beyond the dense network that accounted for the major providers, out to the loose, disparate, periphery of the network engineering community; of operators with little connection to the profession. For the intervention to work, action would have to spread well into the 'long tail' of providers to achieve near total coverage – otherwise DDoS would simply cluster among these smaller and less socially-networked actors, who would often have direct incentives to cater to this new customer base, and little loyalty to the wider networks.

5.2 Re-framing the problem

Although the network engineering community was indeed characterised by a dense and strong social network, DoS had never historically been seen as a core part of their professional responsibilities. The damage from DoS attacks was largely not being felt by providers – rather, major Internet services such as Amazon and Google were far bigger targets. Historically, DDoS had been considered to be an issue of 'abuse', crime or security, and hence the terrain of other professions (sysadmins, law enforcement, or cybersecurity staff) whose core practices were either outside the control of the organisation or focused largely on internal priorities and motivations (rather than a common 'greater good'). Thus, the intervention relied on significant work being done to define the problem of DDoS as not only falling within the domain of 'network engineering' but being an issue of professional concern for the community and its networks of collaborative work.

> For us specifically, we never really engaged with external networks from a DDoS perspective. Certainly for networking related topics we had, and we were big on pushing for BGP RPKI but DDoS was an area we just didn't get all that involved other than from a defense perspective. So we started to channel some of that same energy that we had done in other areas into the DDoS space and acknowledged that we were new to doing this and wanted to do what was within our power to try raise the bar here.

Accordingly, the novel approach was developed initially not within 'network engineering' but within the security community, and not at a transit provider, but at a major web service – namely, in the espousing of novel traceback methods by Damian Menscher at Google. Denial of Service had been a long-standing issue for major Internet services going back decades, however attempts to coordinate solutions (such as enforcing anti-spoofing practices within transit providers and peers) had been unsuccessful. Direct action from Google – mobilising its significant heft as a major Internet player through submitting abuse reports to offending transit providers – had historically proven ineffective. There were significant reputational risks at play in 'throwing one's weight around' too openly as a major global private sector service provider, and Google had significant commercial incentives to avoid being seen as bullying other providers into action through threats of de-peering. Further, the skills-base and resources even in the larger transit providers was not always there to tackle a problem which had generally been seen as a security (and hence victim-focused), rather than network (and hence intermediary-focused) issue.

Initial attempts by Menscher to submit abuse reports led to some well-meaning engineers at other networks suggesting that they cut off traffic to Google services to staunch the flow of DoS traffic (which would in reality simply cause a major outage of Google's services for their customers). Finally, given the incentives to allow spoofing as a method of profit generation, the easiest action for transit providers and peers was simply to directly refuse (leveraging their own central position in the infrastructure) or to disallow spoofing to Google's networks, but permit it to all others, thus preventing Google from gathering intelligence on DoS, while doing little to solve the wider issue. Thus, several effectively commercial incentives meant that Google was unable to simply mobilise its central position and power in the Internet infrastructure; a wider coalition was necessary.

The impetus for a broader push came in 2017, with a series of state-sponsored attacks which demonstrated to Menscher and others the need for wider action:

> And so I largely gave up on the idea of independently approaching [the transit providers], and trying to get [our DDoS problem] solved. Then around 2017, we received an attack from the Chinese government. Fairly confidently from the Chinese government. It was two and a half Tb of spoofed, amplification attack traffic. This was before memcached, so this is like, you know, DNS amplification. The Chinese government was spoofing at least like a hundred gig outbound, in order to generate that amplification. Uh, and this was coming from China Telecom, China Unicom, China Mobile. It came from everything associated with China. And nowhere else in the world. And it was associated with other attacks that we'd been seeing at the time from the government. And so that sort of got me to revisit, like, OK, we need to get people to act in order to cut this off, because this is a matter of national security for every nation other than China.

As described above, a subsequent series of major attacks in 2020 then led to the coordination of Greene and Clayton, who set up the Slack channel and began the development of the Weekly List and the community traceback initiative with Menscher. This attempted to move the action beyond this small group to enrol the core network of key players at the transit providers and major ISPs. The Slack channel and the List were designed to make DDoS as an issue a concern of the network engineering community – initially, inviting a small number of long-standing figures at core transit providers and ISPs, and then extending this to a much wider network of providers. Together, this initial network wielded significant authority within the network engineering community.

A number of these key players had a history as coordinating figures within the anti-abuse ecosystem, and had developed over time a library of distinctive strategies, experiences, and relationships since the early 1990s. These networks extended beyond the network engineering community, to wider partnerships with law enforcement and private sector actors. In addition to more general strategies for team-working and leadership practices, they together had significant experience of consciously designing social relationships and professional networks in the various communities that maintain Internet infrastructure; working around the bottlenecks, jurisdictional issues, legal hurdles and information gaps which typified the work of managing abuse.

> So that story [of our attempts to develop strategies in the mid-1990s] went into Microsoft who co-founded with FBI, the NCTFA. So NCTFA is the National Cyber-Forensics Task Force. That's the police side. Since [NCTFA] been in operations for so long, they've got a huge network, because what you do is you go, any law enforcement agency around the world can apply, and they can send an officer that they fund, and do a six month or a one year stint with NCTFA in Pittsburgh. With lots of other law enforcement people, and an alumni of lots of law enforcement people. And they walk away, and they know each other, and they see operations, they see how to do things, they've learned all sorts of techie stuff, right? And how to do different trackdowns of different sort of technical investigations. And they go back to the home country and they have this Rolodex of power, right? And this is one of the checks and balances that we set up in the system. And that worked, right?

Thus, the traceback action began not only with a technical innovation, but also with a set of strategies to mobilise the social networks of network engineers and a small network of key players from different professional communities; security, academia, and network engineering. As we discuss in later sections of this paper, the fairly dense social and professional networks which exerted informal governance over the peering ecosystem meant that enrolling these core actors – with legitimacy and authority in this community – was able to generate far wider action.

5.3 Direct effects, intervention design, and incentives

The initiative itself – focused on reporting the biggest and fastest attacks identified by Clayton's sensors every week and then engaging in community attempts to traceback the source of the traffic and 'plug the hole' – had a number of direct effects. As initially designed, a key aspect of this intervention was intended to involve social shaming — the idea that, within this community, an established (or even more junior) engineer wouldn't want to see their own network featured on the list. This was intended to play on the wider social dynamics of the network engineering community – namely that people wouldn't want to be out in the cold when the time next came for a post-conference beer. This, coupled with the general camaraderie associated with a group effort spearheaded by respected community leaders, was itself intended to shift the incentives out of the domain of commercial economics and into the domain of professional community. In addition to a purported 'shaming' effect, the List also acted to reward participants, who could visually observe effects, generating a clear reward, a feeling of progress and group membership.

This social cost and benefit was coupled with a broader coordination effect generated by the List, which incorporated a number of elements designed tactically to reduce the risks, reframe the issue as a 'neutral' technical matter for the community, and make the action more manageable, moving from a large number of attacks and possible spoofing sources to concentrating action on a small number of key targets every week. One of the main blockers to moving the action into the professional domain were commercial issues – namely, the sharing of sensitive information (which we discuss in the subsequent section) and perceptions of commerical bias. The participants were keenly aware that it couldn't appear as through the major corporate players were throwing their weight around – this would move things back into the domain of commercial power, in which the terrain of incentives were well-established as dangerous. The enrolment of Clayton as a neutral academic party moved the intervention away from the field of inter-organisational competition, reframing it as a matter of professional practice. The weekly 'cadence' of action was similarly important – keeping it a live issue and confining participation to a small window of time every week.

> The lag is, can we actually force the movement. Not make it go away, force the movement. Alright? That means the lead measure needs to be, what we decided on is, every Monday, we just used one of the HoneyNet systems, we've got a lot out there, so we used one of the HoneyNet systems to say, here's a list of, this is what Richard put forward, to say here's a list, every week, of DoS attacks. And then we started backtracing there. And we just need thirty minutes from an operator, every single week, you go out there, let's, OK, hit this AS, this AS, this AS. And you trace it back, and you go out there and plug the hole. Right?

Beyond the immediate effects of the list in countering the targets named every week, it appears that the intervention only required the partial enrolment of the wider social networks of the engineers in order to work. In the event, while many of the core major network providers were convinced of the benefits of the intervention and were happy to help, they lacked the engagement and ability to devote serious resources to hunting down spoofing sources. Further, as the Slack opened up, large numbers of engineers joined who had little tangible to contribute other than social solidarity – they were keen to help, but had only a restricted local view of the global Internet and so could not trace traffic to its source, and often lacked core skills and resources. A single central engineer – Tom Scholl – went far further, taking up the cause more widely, committing serious resources to research and intelligence coordination, and used his own influence and relationships to cascade this central effort out to the wider network.

> Richard works on his weekly report, which I used to perform traceback exercises. But I didn't want to just rely upon that alone, so I started building my own internal reports to find as many networks as I could that were spoofing into us and started to engage them directly. I put together some standard email scripts and reporting formats and would identify external networks and specific abuse and NOC aliases to engage. I'd pull in some of our folks who handle peering negotiations/interacts to help find the right individuals at specific companies to raise awareness and escalate with those external networks to action these reports. I just couldn't stand seeing report after report of the same networks showing up, so I kept at chasing down each network as it showed up.

Thus, rather than solely a direct shaming effect (though this undoubtedly did play a part in compelling immediate compliance), equally important was the role which this played in activating wider networks of social capital and professional influence. Given the topology of the ASs which make up the Internet, this influence propagates through direct peer-to-peer relationships, i.e. even though you might not personally be a direct peer with a network, they will be a peer with someone who is one of your peers, so if you can enrol all your peers then they can propagate the signal more widely.

> There's certainly a set of networks that somewhat ignore or don't prioritize some of the complaints. For those that we directly peer with, its generally hard to avoid since we have a relationship there. I've never threatened to de-peer anyone either. There are certainly upstream providers that we pay that haven't always handled these well which has required in escalation to our account teams to drive attention on our complaints. With settlement-free peering though, they don't necessarily owe you anything, so you just have to be persistent and draw more attention on their side. You really have to amplify the message that they've got a customer downstream doing this behavior and that you want them to do their part to mitigate this behavior. It really comes down to being persistent and being able to influence without having authority and trying your best to be persuasive.

This necessitated significant wider resources coordinated from the central group – walking people through the problem, upskilling them, showing them what to do, developing proficiency (much as with the kinds of peer learning, education, and problem solving which characterised the growth of the profession in its early days), getting them used to thinking about and engaging on this issue, and generally making it a part of the rhythms of their job.

> With some of the larger networks we've helped build that muscle with them to be able to be proficient at tracing back spoofing activity. They know how to look it up, they know what

> actions they need to take and how to deal with a downstream customer who isn't responding. That's a real success story where you can convey the specific IPs, ports and timestamps and they can run with it to completion. Some other networks, you have to help explain to them what spoofing is use pictures to explain exactly what is going on and provide additional logs and help them. This is where it can be a challenge because you may be dealing with staff not familiar with how to use their netflow related tooling and have to help them depending on what product they're using to navigate it and build the right queries. Sometimes they'll struggle with separating out peer vs. customer traffic flows or mis-identify what is the pre-amplification vs. attack traffic. In some cases, networks struggle with even the concept of spoofing and will look at their routing table and claim they aren't announcing the prefixes which requires to explain what spoofing is. There is a bingo card that has been shared around that outlines some of the responses one can observe when engaging with networks that explains the range of answers when engaging with networks fairly well. There's also one network that certainly is a standout that will self-identify these customers spoofing on their own and are on top of their game. Those are the ones where you reach out to them and they're already aware and are working with their customers already. They'll even go above and beyond to help work with their customers and do interactive screen sharing sessions to assist them with using their own netflow tools to identify the traffic.

This led further into not only compelling action but slowly incorporating these practices into the core *habitus* and skills of the network engineer profession. While most of the providers may not have had the resources to look pro-actively, and may not have had the interest to become a part of the pro-active 'core' of anti-spoofers, this wider alignment of the profession as a whole enabled a smaller number of highly-engaged engineers at large providers to take on this role for the wider network and facilitate their action. Their trust and authority, built up over years, was crucial to this role.

> The challenges with delegating this out to others within the organization has to deal with the level of trust involved when engaging with external networks. For example, external networks may not be comfortable sharing details of specific customers involved in spoofing with someone they don't know for fear that that information may be mishandled. There is also a level of familiarity with what specific patterns you learn when you see specific networks spoofing in a particular geographical area to give you hints of which specific network is behind the spoofing as well that isn't trivial to figure out unless you've dealt with it before. It really requires someone who wants to dive into the whole problem and obsess around tracking specific networks, educating others and being very persistent.

This meant that, where the core large-scale providers were generally simpler to align, they could then push this action through to mid-size and smaller networks. The substantial additional work required to trace some of the more persistent offenders – including purchasing illicit datasets to track down providers with abusive users – could then be handled by this relatively small central community and then the intelligence radiated out through these networks of contacts.

> When engaging networks, it takes a fair amount of patience to convince them and work through some of the challenges with the effort. In some cases, you have to give hints of which specific customers you want a network to look at. Sometimes it's a case where you have other telemetry pointing to the specific customer involved, sometimes it's a case where a network was multi-homed and you worked with a separate peer to identify that specific company. In some situations, we've gotten VMs at specific hosting providers to verify we could actually spoof from a specific hosting shop as well. One of the interesting things is if you don't announce all your prefixes consistently across the globe, you can build a "fingerprint" of a specific network and how it ingresses into your network infrastructure via particular peers and transit. That gives you more options to be able to identify a source via multiple peers and increase the likelihood that one of those peers will provide additional information into the network spoofing. From that, you can use that information to point to a specific peer who is struggling and

> telling them what specific customer they should dive into further. I've found that works really well when a network is struggling to use netflow and identify a specific source.

Thus, we observe a combination of (1) a process of direct professionalisation (i.e. upskilling key members of staff at ISPs), (2) actively working to develop productive social capital and links within the global community, bringing these outlying parts of the network into the fold, and linking them up more strongly to the professional community and core networks of capital and (3) relying on the economics of peering to facilitate action.

5.4 Supportive social infrastructure

In addition to social networks and peer relationships, key 'social infrastructures' of the network engineer 'scene' played an important part in facilitating this intervention. The Slack channel itself was a key factor in allowing coordination among the central network and providing a centralised site for information diffusion, deploying Clayton's List where it would shape the behaviour and perceptions of key actors. As we identify above, both online and offline community sites had long been a core aspect of the network engineering profession.

> Since very early 2000s the Internet network operator scene had its own methods to facilitate communication between operators. For example, there is one IRC network and its servers that has been around which has been useful over the years for being able to communicate with other operators. You'd have folks from the various network operator communities (NANOG, RIPE, APRICOT, GPF, etc.) there. These forums have been useful for networks to communicate with each other when dealing with large scale Internet outages or needing to engage with specific networks to directly engage in troubleshooting. Other topics as well such as general network design, or airing of issues with specific network hardware or software have been commonplace as well. Again, the benefit is that you have people of various networks being able to directly engage with one or multiple parties rapidly and cut through red tape, or NOC email aliases or heirarchy to get ahold of a resource when you need it.

This was supported by its own infrastructure – of conferences, IRC (and later Slack) channels, and forums. The incorporation of these sites into the heart of the profession has long allowed for release valves – permitting flexibility where information needs to flow outside of corporate official channels. The historic and present use of these settings were crucial to the success of the intervention; combining both informal settings – pub chats, IRC servers, and Slack channels – where operators could talk 'out of school' with underwriting by formal legal protections from professional bodies to protect engineers from being compelled to disclose the contents of these private chats to their employers.

> So I call a meeting with M3AAWG, we anchored it in M3AAWG so it had... the anchor in M3AAWG is important because if you don't anchor in some sort of institution that has good legal bylaws of confidentiality that you can hide behind, what happens is that if you're in a T[raffic] L[ight] P[rotocol]-realm, TLP-red or TLP-amber, then you can have the lawyers inside your company say, no, you have to tell us everything. You break trust. But um, that's one of the reasons of finding an anchor organisation. So we got M3AAWG as the cover, so if [anybody's company] says, you have to tell us what's going on, [you say] like, this is a M3AAWG project. These are the M3AAWG bylaws. This is how it works. Talk to the M3AAWG lawyers. Right? The TLP-red part, the reason why we anchored the law enforcement operations for other industry with the NCTFA is they have the same thing. Which is more confidential because it's part of the investigation, because there they're trying to arrest people on the TLP-red side. So you've got coverage. So that's one of the techniques around us, is making sure you're anchored to an association, to give you, to protect the TLP trust. Right?

Thus, through a coordinated set of carefully designed strategic interventions, a small number of players were able to successfully 'short-circuit' the landscape of incentives preventing action in the networks of commercial power, aligning the power of the Internet infrastructure and the professional community who maintain it to achieve a genuine social good against apparently intractable issues.

5.5 Attacker-side economics and incentives

Finally, a series of incentives relating to the attacker ecosystem contributed to the success of the intervention. As has been previously noted in the research literature, the same incentives which cause centralisation in wider Internet services also apply to illicit ones [30]. The availability and cost of hosting, particular provider norms and rules around user behaviour, and the dynamics of information and social learning, in which customers tend to share information about markets in peer networks, all mean that illicit users tend to concentrate in a small set of particularly propitious hosting providers. There is an element of path dependency here – illicit user markets tend to cluster around the providers which are identified early by the community rather than spreading out to explore the full market. This itself can lead to lock-in at providers, for whom these users suddenly become a major new customer base to which they need to cater.

Much as the centralised social infrastructure of the Slack channel permitted social action within the anti-spoofing community, the booters and DDoS-as-a-service providers are heavily reliant on their own rather centralised social infrastructures for information diffusion. New providers are found through sharing paid spooflists of providers that allow spoofing, through peer recommendations within booter communities, and through discussions on a relatively small number of Telegram channels. These themselves, as discussed above, also provide a centralised source of intelligence for the spoofing initiative to find and shut down potential sites for illicit hosting. Over time, the incentive has been for the booter market to progressively centralise due to the limited appeal and profit of running a small-scale service, increasing risk to smaller providers, the proliferation of scams and the difficulty of establishing a reputation and customer base (the 'cold start' problem), the preferable economics for mid-level providers to simply resell the attack capacity of the major services, and classic economies of scale for the larger providers.

All this means that where these pockets of concentration are removed, this can achieve significant disruption across a large swathe of the market. One of our interviewees noted this directly, where following successful enrolment of a major Russian provider, large numbers of booters lost their main hosting provider.

> In February 2022, there were a number of booters that they lost some of their spoofing power in a single shot and mentioned that loss of power publicly. That was the result of directly engaging with an upstream provider in Russia who finally closed the door on that particular hosting provider that a number of booters were using. There's been a few other cases where we'd observe multiple booters generating spoofing from the same general vicinity and would chase down the spoofing source particularly when we were seeing a significant amount that was basically like a hot spot. We'd engage each of their transit providers to get anti-spoofing controls applied on them and the booters overall power would slowly diminish until it was completely shut down. It was clear that the booters would share information about which particular hosting provider to gravitate to once they found a place that was working well. Then you'd get the spoofing shutdown and they'd scatter and have to find new locations to move to.

This then has wider effects on the (far looser) networks of peer trust, capital and expertise in the booter community – degrading the trust in recommendations as people percieve that they have been scammed. The same cadence of action that facilitates ongoing action on the hosting side then produces a further disruptive effect as network engineers observe and intervene in the attempts of the booter community to find and exploit new spoofing sites.

6 Applicable theory

In this paper we have outlined a successful intervention in which collective action was achieved against apparent organisational incentives. We now discuss some theoretical resources which we argue can provide some explanatory power in this case. These can further contribute to expanded frameworks for making sense of these higher-order dynamics of incentives and action in global security markets.

Security economics tends to view incentives through the lens of rational action, in which actors make individual decisions based on their perceptions of personal risk, benefit, and the actions of other players [5, 6]. These often take the organisation as the unit of analysis in assessing the likelihood of action. That is, focusing on the costs and benefits to the company as a whole of patching vulnerabilities, in paying ransoms, or in the incorporation of particular features. However, security decisions within these environments are made by individuals whose motivations and incentives may be more complex, and not necessarily entirely aligned with those of their organisation. In some cases, as we identify here, it is possible, using well-tailored interventions, to shift the frame to allow a different set of incentives to predominate. The cultures and kinds of work of the actors involved – the structure of their human networks, technological enablers, and well-designed strategic interventions can serve to shift these calculations to permit sustained coordinated action apparently against immediate individual incentives.

This wider argument – studying the effects of culture and infrastructure in shaping economic arenas of action – has been developed across a series of recent papers in WEIS, including evaluations of maintenance work in cybercrime ecosystems [34], the entrepreneurial dynamics of cybercrime groups [4] and technical sophistication in cybercrime practices [33]. These argue for a wider industrial ecosystem perspective, integrating culture, practices, and infrastructural factors alongside and interacting with economics for understanding the dynamics of offending in cybercrime markets. In this paper, we extend this approach to include the wider factors shaping action within security ecosystems from the defender side.

6.1 Informal governance in Internet infrastructure

Scholarship on Internet governance historically understood Internet systems as shaped by predominately market forces following the commercialisation of the Internet, in tandem with regulation by formal processes housed in key international institutions like ICANN and the IETF which forged the influence of practitioners, academics, and policymakers [26, 28]. This literature focused more generally on competitive business dynamics in stimulating or preventing change, and on the importance of standard-setting as a key regulatory tool for shaping markets. The wider political dynamics of action within the commercialised Internet infrastructure were indeed initially shaped in this era of globalisation, with geopolitical factors mostly (though not entirely) conforming to the global free market politics which themselves underpinned bodies like ICANN and the ideology of the early Internet [25, 10]. However, in recent years, this terrain has somewhat shifted. The central forces and dynamics shaping the Internet are decidedly multi-polar, reflecting a resurgent technological nationalism (or at least, national industrial strategy), online abuse and crime as a key informal dimension of interstate conflict, and an international standards landscape in which US interests are more open to challenge [8, 67].

Despite these factors, the Internet has nonetheless continued largely to work, thanks in no small part to the efforts of sets of global communities of maintainers and administrators working in their own professional domains, of which peering and network transit is only one. Recent scholarship has foregrounded the importance of *informal* governance processes in promoting stability and compelling action in the Internet infrastructure, with some going as far as to describe a *techno-scientific elite* of key players with an outsize role in coordinating Internet infrastructure [25]. In this scholarship, the concept of 'normfare' [68] has recently come to the fore – outlining processes of what Becker [12] and Cohen [31] separately describe in the classic sociological literature as *moral entrepreneurship*, in which individuals and small groups seek to establish a change in norms in a wider community through strategic action across multiple different spheres. Within the wide range of different levels in which normative entrepreneurship operates in the Internet infrastructure (including political negotiation, formal institutional standard-setting, academic research and commercial practices), there is an important place for informal governance processes [29, 40] – including the role of social capital and informal agreement within these expert communities in establishing a terrain of action somewhat separate from the corporate, formal institutional, and political forces long assumed to characterise action in Internet infrastructure [1, 80]. There is a particular importance here given to working groups and initiatives of the kind discussed in this paper, in which these projects of 'normative entrepreneurship' are mobilised to make changes to the wider practices of the network [46, 53, 71, 73]. Some of the descriptions in this literature, for example, of key actors 'internalising the externalities' of these issues [78], are particularly salient in our case.

This literature, though valuable, still largely focuses on the policy-setting and international relations perspectives which characterised the previous literature, though extended to less formal settings – for example, showing how professional communities exert power outside established political and institutional frames and through their own personal networks. These are focused more on larger-scale standardisation decisions (i.e. how standards are negotiated, drafted, and accepted) than the lower-level, more practice-focused efforts to extend compliance with these standards out into the wider professional community that we observe here (with some exceptions, such as around patching or botnet mitigation) [21, 64]. However, some important examples of this scholarship describe exactly these kinds of interpersonal networks of action [28]. In these areas, these studies bear a debt to Bourdieu, concerned as they are with the dense networks of social capital which define professions, how these link to *habitus* (the core shared professional practices and deep conscious and unconscious dispositions which define a community), and the wider field – the domain of action in which they try to assert themselves. Bourdieu's theoretical work is effectively concerned with *power* – how people work strategically and deploy capital of different kinds within professional communities to change the 'rules of the game' [18]. In addition to the traditional forms of material economic capital, this extends to a wider set of resources that can be 'spent' to achieve social action – including social capital (the dense networks of links, alliances, and connections that people develop across their careers and can draw on to achieve action), cultural capital (one's alignment with and knowledge of key cultural ideas and values which signal belonging to a community) and symbolic capital (associated with status and authority) [17, 19].

This allows us to theorise a distinctive network of people supporting the infrastructure from below, arguing, as we observe in this paper, that the nominally distributed network of the Internet backbone in fact relies on highly coordinated peer relationships within these professional communities. Tracing the networks of *social* capital (rather than the network of infrastructural peering connections discussed in much of the literature on peering), we see a distinctive structure, in which a core of highly networked engineers maintains significant authority to compel action through persuasion, shame, sanction, and reward within a wider network. At the more diffuse outer reaches of this network remain a cluster of poorly-enrolled engineers who are marginal to the network, and whose action is hard to compel in this terrain.

6.2 Infrastructural capital

The idea of social and symbolic capital as central to processes of informal governance are clearly relevant to our case – we are particularly interested in how capital of different kinds can be deployed by key actors to set off cascades of action within these networks of people, companies, and technologies. However, our research suggests that the technical centrality of particular players in the Internet infrastructure – as De Nardis describes, their situation within the Internet's material and technical networks of 'control points' [38] – is a crucial aspect that cannot be ignored in favour of a solely humanistic account. Additionally, we note that there are aspects of the 'control points' view which do not account for some key dynamics at play in our case – namely, many of our actors leverage their position in the Internet infrastructure not to dominate or control action, but instead, for forms of *visibility* which permit crucial aspects of maintenance. Star's account of infrastructure puts a substantial emphasis on the importance of practices as forms of work which connect people up to infrastructure and bind them together in shared social worlds [75]. Extending this, Bourdieusian conceptualisations of practice show these as themselves bound up in professions and institutions, and defined by their own network topology of social capital. For our infrastructure providers, the topology of these networks of human capital is strongly linked (though, crucially, does not directly conform to) the physical topology of the Internet infrastructure. The incentives, economics, and forces of concentration work as much in these networks as they do in networks of companies or cables; these domains are densely interconnected. The professional field of the Internet infrastructure – and the mobilisations of actions we observe – is strongly shaped by the macro-scale technical structures and networks of the Internet itself. These infrastructural networks are crucial in setting the relationships between players and the incentives of different actors, and they confer substantial explicit and implicit power and authority.

The exercise of this power is not straightforward, and requires significant nuance. We observe action taking place in *three main domains*, each defined by their own rather different economics, network structures, relationships, and incentives. These can be thought of as (1) *the infrastructural domain*, relating to positioning within the material dynamics, networks, and topology of the Internet itself. Centrality and reach of influence within the infrastructural networks confers not only significant power in the form of *infrastructural capital* but also a distinctive view of the Internet that can itself form a key enabler of action. (2) *the commercial domain* in which business interests and power predominate and (3) *the professional domain* defined by personal relationships, networks of social capital, skills, clout, and symbolic power.

Important players in the peering ecosystem have power in all three domains. The core individuals we study here are able to draw on infrastructural, commerical, and social capital – especially when allied together in a small team – to align and coordinate these forms of power. They move more or less symmetrically between the three types of network and are able to mobilise action in all three of these domains, and tactically switch between them when needed. When action in one domain (here, the commercial) is blocked, they work to replace it with action in another layer.

We propose the concept of *infrastructural capital* as a form of capital accrued and spent by actors in these Internet networks (and in other forms of infrastructure), and which allows them to leverage their own position in these infrastructural networks in a variety of ways. In particular, we observe actors using their position in the technical networks, or a set of interventions which rely in other ways on this infrastructural domain, in tandem with action in the networks of social capital in the network engineering profession. A small number of key engineers and other players were able to use their view of the Internet and centrality within the relevant social and technical networks to achieve action. Infrastructural capital can be seen here playing much the same role as other forms of capital (such as organisational capital) in Bourdieusian accounts – a resource to be cashed in, traded, and mobilised with other forms of capital in tactical and strategic games to achieve action. Infrastructural capital acts in both directions – allowing actors to draw on the power of their links to the infrastructure to push social action, and to draw on their social power within the professional communities of infrastructural actors to compel material changes to the infrastructure. Thus we arrive at a depiction of the network engineering field as structured by networks of both infrastructural and social capital, which produce the topology along which action can flow.

6.3 Mobilising change in Internet infrastructure

Taking these theoretical resources together, we note the following features which proved to be crucial aspects of this intervention. Firstly, it is clear – as Chenou notes [27] – that key individuals and their networks play a crucial role in achieving transformations in practice and in the wider governance of Internet infrastructure. Beneath the level of governments, policy bodies, companies, institutions, and technical committees, these individuals wield significant power and agency.

This action took the form of *persuasion and alignment*; brute force at the level of infrastructure or business was rarely effective. Attempts to assert power directly in any of these domains had to be carefully negotiated, and generally needed to be consensual – if the actors pushed any one of these domains too far then they would lose the subtle mechanisms of influence and provoke resistance. Even powerful companies like Google with a huge purchase on the global Internet weren't able to just directly bully others, instead needing to do significant work to mobilise action and make this an issue for the profession.

Equally, the social domain didn't act on its own either – although the list did provoke some 'shaming' effects in combination with the wider coordination of action which it provided, in fact attempts to solely and directly mobilise symbolic capital among the network engineers *on its own* were rarely successful. The initial RFC led to many major players implementing the fix on their own networks, however outside this core, little activity was generated; in these cases, the commercial incentives or power structures proved able to resist, and the aligned players were initially reluctant to attempt to compel action outside their own networks. Instead, careful and tactical alignment of action between the infrastructural and professional networks by a few key players was key to success. A great deal of the material dynamics of the network infrastructure needed to align to make the social action possible, including: the core actors' central view of the the Internet, resources like Clayton's sensors and access to spooflists, their power and position, the material necessities and power dynamics of peering, and the wider tactical design of the intervention. However, these in turn relied on the dense social networks of the engineers to achieve wider action; these obviated the need to get the hundreds of actors in the peering community 'on board' – if the network is strong then just a few key people can then activate this wider network into action.

In this case, the major blockers and perverse incentives were, and still are, in the commercial domain. For action to work, there was a need to de-link the incentives and power structures of the commercial networks from the arena of action. Here, the dense social networks of the network engineers and a series of established institutional and informal spaces allowed the professional and infrastructural networks of action to short-circuit these commercial ones. This was coupled with careful design of the intervention to maintain this separation – the objective and academic nature of the List, the work done by core actors in conducting traceback and finding sources, and the legal protections and support conferred by organisations like M3AAWG. Although the commercial muscle of some of these key players doubtless contributed, they more often found themselves asserting infrastructural or professional power.

Finally, it is important to note that the dense professional networks of the network engineer profession do not emerge from nowhere, neither are they immune from entropy. Considerable work has been done over decades to grow and sustain this social network, from the more formal work of organising conferences, meetings, and training, to the more informal maintenance of the social life of the profession. This is strategically vital to the continuation of initiatives like this one, and needs continual maintenance as younger engineers drift in, and older engineers are promoted.

7 Related community initiatives

In November 2008 malware, dubbed Conficker A, started to infect machines which had not installed a recent patch for their Windows software. It rapidly became clear that the malware author was in a position to control a multi-million machine botnet – a significant threat to Internet connected systems. At that time botnet command and control systems (C&C) generally used a handful of domain names (infected machines would resolve these and then report in to the C&C) so that mitigation could be achieved by seizing these domains. The original version of Conficker generated 250 new, random, domain names every day spread across eight top-level domains – all would need to be seized to prevent the malware author from controlling the botnet. The response to this threat is documented in a white paper [36]:

> In an unprecedented act of coordination and collaboration, the cybersecurity community, including Microsoft, ICANN, domain registry operators, anti-virus vendors, and academic researchers organized to block the infected computers from reaching the domains – an informal group that was eventually dubbed the Conficker Working Group (CWG).

The CWG effort was a success in that the Conficker botnet never did significant damage.

The white paper describes the evolution and operational aspects of the CWG but also reports on a number of 'debrief' interviews (much as this paper does for the traceback initiative). Of particular interest is the list of reasons provided for people to become involved with the group.

- General altruism for stopping a threat
- Cooperation rather than competition on a threat
- Coordination of research and sinkholing of data
- Public relations benefits
- Networking with other professionals
- Monetization of research and data
- Damage control
- Concerns of being on the outside

Many of these reasons resonate with our interviews, although a very significant difference is that the CWG was based around organisations joining the effort (hence the PR topic) rather than individuals stepping up because of the data that they had to hand as part of an organisation.

More directly relevant to DDoS mitigation is the so-called 'Big Pipes' initiative where industry, law enforcement and academics have been cooperating to tackle denial of service problems since 2014. The traceback initiative spun out of this group – whose discussions and membership are not public. There is a overview of their activity written by a journalist in 2023 [41]. The impact of takedowns of DDoS-for-hire systems (booters), many of which were initiatives of this group, were studied by Collier et al. [35, 65].

8 The way ahead

Since the traceback initiative is still going on, still evolving and still has much to do, it seems more appropriate to conclude this paper by looking forward rather than backward. The NETSCOUT graph shows that though spoofing is much reduced there is still some way to go.

As we have described, spoofing was first documented as a problem almost forty years ago. The appearance of BCP38 twenty-five years ago undoubtedly made for improvements, but there has been little change in the past twenty years. Activity to reduce the number of reflective amplifiers has been going on for fifteen years, but since damaging attacks may use less than a hundred amplifiers the reduction must be almost complete before any benefit is seen.

Menscher's 2019 proposal to exploit the widespread availability of traffic flow logs to trace back to the source of spoofed traffic was eventually taken up in 2021. The initiative's approach of trying to find only a handful of sources at a time, but with a regular 'rhythm of action' to keep prodding people to act has driven the project forward and this has in turn encouraged various people to be proactive and identify and mitigate sources that never appeared on the weekly list.

To sum up what we believe is going on: the role of hosting companies in spoofing and the wider booter ecosystem, as with many areas of cybersecurity, is characterised by commercial incentives which pose significant hurdles to action for the common good. One approach to tackling this might be in changing these incentives – increasing the risk and costs of inaction through regulation or enforcement, or producing tangible commercial benefits for action.

However, here, we have observed a different strategy, in which key actors have been able to shift the action into other domains – domains which are determined by the dynamics and incentives of professional communities and backed up by the strategic power of the infrastructure itself. Cutting the commercial incentives out of the loop takes substantial effort and depends on the dynamics of the people themselves – the networks of social capital which define the network engineering profession present a rather different topology of power and incentives to those of the networks of companies and providers.

M3AAWG recently awarded its prestigious 'J D Falk' award to Tom Scholl of Amazon for his contributions to the project [56]. One of the characteristics of J D Falk's contributions to email standards and the wider security community, before his untimely death in 2011, was the way in which he brought humour into his work. There is humour in this space as well – a widely circulated 'bingo card' (see Figure 2) provides a set of the various technically inept responses that are regularly seen when hosting companies and networks are first approached to get them to resolve a spoofing incident.

Perhaps then, the most interesting observation we can make is that, besides being amusing – and resonating with the experiences of many of the network engineers who have part of the traceback project – the bingo card is now being used as a guide to developing a curriculum and creating training material so that the current generation of network engineers can pass on their knowledge, experience and ethos on to the next generation.

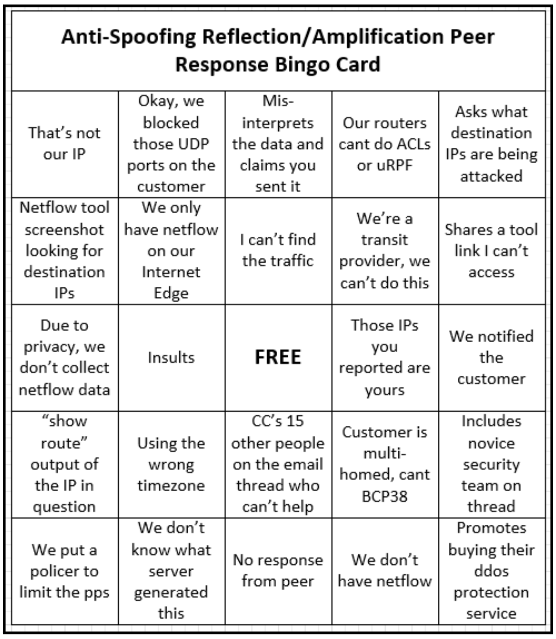

Anti-Spoofing Reflection/Amplification Peer Response Bingo Card

That's not our IP	Okay, we blocked those UDP ports on the customer	Mis-interprets the data and claims you sent it	Our routers cant do ACLs or uRPF	Asks what destination IPs are being attacked
Netflow tool screenshot looking for destination IPs	We only have netflow on our Internet Edge	I can't find the traffic	We're a transit provider, we can't do this	Shares a tool link I can't access
Due to privacy, we don't collect netflow data	Insults	**FREE**	Those IPs you reported are yours	We notified the customer
"show route" output of the IP in question	Using the wrong timezone	CC's 15 other people on the email thread who can't help	Customer is multi-homed, cant BCP38	Includes novice security team on thread
We put a policer to limit the pps	We don't know what server generated this	No response from peer	We don't have netflow	Promotes buying their ddos protection service

Figure 2: A 'bingo card' of the technically inept responses received when trying to track down and mitigate sources of spoofed traffic used in attacks.

9 Acknowledgements

We are enormously grateful to Barry Greene (senki.org), Damian Menscher (Google) and Tom Scholl (Amazon) for agreeing to be interviewed at length; and for all their other helpful comments on drafts of this paper. Richard Clayton is supported by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 949127).

References

[1] Kenneth W Abbott and Benjamin Faude. Hybrid institutional complexes in global governance. *The Review of International Organizations*, 17(2):263–291, 2022.

[2] Alessandro Acquisti. Privacy and security of personal information: Economic incentives and technological solutions. In *Economics of Information Security*, pages 179–186. Springer, 2004.

[3] Ross Anderson. Why information security is hard – an economic perspective. In *Seventeenth Annual Computer Security Applications Conference*, pages 358–365. IEEE, 2001.

[4] Ross Anderson, Richard Clayton, Rainer Böhme, and Ben Collier. Silicon den: Cybercrime is entrepreneurship. In *Workshop on the Economics of Information Security (WEIS)*, 2021.

[5] Ross Anderson and Tyler Moore. The economics of information security. *Science*, 314(5799):610–613, 2006.

[6] Ross Anderson and Tyler Moore. Information security economics – and beyond. In *Annual International Cryptology Conference*, pages 68–91. Springer, 2007.

[7] Terrence August and Tunay I Tunca. Network software security and user incentives. *Management Science*, 52(11):1703–1720, 2006.

[8] Julia Bader. To sign or not to sign. Hegemony, global Internet governance, and the International Telecommunication Regulations. *Foreign Policy Analysis*, 15(2):244–262, 2019.

[9] Fred Baker and Pekka Savola. Ingress filtering for multihomed networks. RFC 3704, 2004.

[10] Gabriele Balbi and Andreas Fickers. *History of the International Telecommunication Union (ITU): Transnational techno-diplomacy from the telegraph to the Internet*, volume 1. Walter de Gruyter GmbH & Co KG, 2020.

[11] S Bartholomew. The art of peering. *BT Technology Journal*, 18(3):33–39, 2000.

[12] Howard S Becker. *Outsiders*, volume 1973. Free Press New York, 1963.

[13] S. M. Bellovin. Security problems in the TCP/IP protocol suite. *SIGCOMM Comput. Commun. Rev.*, 19(2):32–48, Apr 1989.

[14] D. J. Bernstein. Welcome to syncookies, 1996. URL: https://cr.yp.to/syncookies/archive.

[15] Robert Beverly and Steven Bauer. The Spoofer Project: Inferring the extent of source address filtering on the Internet. In *Steps to Reducing Unwanted Traffic on the Internet (SRUTI) Workshop*, pages 53–59. USENIX, 2005.

[16] Robert Beverly, Ryan Koga, and kc claffy. Initial longitudinal analysis of IP source spoofing capability on the Internet. Briefing paper, 2013. URL: https://www.internetsociety.org/resources/doc/2013/initial-longitudinal-analysis-of-ip-source-spoofing-capability-on-the-internet/.

[17] Pierre Bourdieu. What makes a social class? on the theoretical and practical existence of groups. *Berkeley Journal of Sociology*, 32:1–17, 1987.

[18] Pierre Bourdieu. *The logic of practice*. Stanford University Press, 1990.

[19] Pierre Bourdieu. Distinction & the aristocracy of culture. *Cultural theory and popular culture: A reader*, pages 498–507, 2009.

[20] Bundesamt für Sicherheit in der Informationstechnik (BSI). Reports on openly accessible server services. CERT Bund Report, 2019. URL: https://www.bsi.bund.de/EN/Themen/Unternehmen-und-Organisationen/Cyber-Sicherheitslage/Reaktion/CERT-Bund/CERT-Bund-Reports/Offene-Server-Dienste/offene-server-dienste_node.html.

[21] Corinne Cath. The technology we choose to create: Human rights advocacy in the Internet Engineering Task Force. *Telecommunications Policy*, 45(6):102144, 2021.

[22] CERT. CERT* Advisory CA-95.01: Topic: IP spoofing attacks and hijacked terminal connections, 1995. URL: https://web.archive.org/web/20000815082526/http://www.cert.org/advisories/CA-95.01.IP.spoofing.attacks.and.hijacked.terminal.connections.html.

[23] CERT. CERT* Advisory CA-96.01: Topic: UDP port Denial-of-Service attack, 1996. URL: https://web.archive.org/web/20000815082414/http://www.cert.org/advisories/CA-96.01.UDP_service_denial.html.

[24] CERT. CERT* Advisory CA-96.21: Topic: TCP SYN flooding and IP spoofing attacks, 1996. URL: https://web.archive.org/web/20000815082252/http://www.cert.org/advisories/CA-96.21.tcp_syn_flooding.html.

[25] Jean-Marie Chenou. Multistakeholderism or elitism? the creation of a transnational field of Internet governance. In *GigaNet: Global Internet Governance Academic Network, Annual Symposium*, 2010.

[26] Jean-Marie Chenou. From cyber-libertarianism to neoliberalism: Internet exceptionalism, multi-stakeholderism, and the institutionalisation of Internet governance in the 1990s. *Globalizations*, 11(2):205–223, 2014.

[27] Jean-Marie Chenou. *The role of transnational elites in shaping the evolving field of Internet governance.* PhD thesis, Université de Lausanne, Faculté des Sciences Sociales et Politiques, 2014.

[28] Jean-Marie Chenou and Roxana Radu. Global Internet policy: A fifteen-year long debate. *The Evolution of Global Internet Governance: Principles and Policies in the Making. Heidelberg, New York, London, Springer*, pages 3–22, 2014.

[29] Thomas Christiansen and Christine Neuhold. *International handbook on informal governance*. Edward Elgar Publishing, 2012.

[30] Richard Clayton, Tyler Moore, and Nicolas Christin. Concentrating correctly on cybercrime concentration. In *Workshop on the Economics of Information Security (WEIS)*, 2015.

[31] Stanley Cohen. *Folk devils and moral panics*. Routledge, 2011.

[32] Gabriella Coleman. The hacker conference: A ritual condensation and celebration of a lifeworld. *Anthropological Quarterly*, pages 47–72, 2010.

[33] Ben Collier and Richard Clayton. A "sophisticated attack"? innovation technical sophistication and creativity in the cybercrime ecosystem. In *Workshop on the Economics of Information Security (WEIS)*, 2022.

[34] Ben Collier, Richard Clayton, Alice Hutchings, and Daniel R Thomas. Cybercrime is (often) boring: maintaining the infrastructure of cybercrime economies. In *Workshop on the Economics of Information Security (WEIS)*, 2020.

[35] Ben Collier, Daniel R Thomas, Richard Clayton, and Alice Hutchings. Booting the booters: Evaluating the effects of police interventions in the market for denial-of-service attacks. In *Proceedings of the Internet Measurement Conference (IMC)*, pages 50–64. ACM, 2019.

[36] Conficker Working Group. Lessons learned. White Paper, 2011. URL: https://www.senki.org/wp-content/uploads/2020/11/Conficker-Working-Group-Lessons-Learned-June-2010-Published-January-2011-whitepaper_76813745321.pdf.

[37] Tianxiang Dai and Haya Shulman. SMap: Internet-wide scanning for spoofing. In *Proceedings of the 37th Annual Computer Security Applications Conference*, ACSAC '21, pages 1039–1050. ACM, 2021.

[38] Laura DeNardis. Hidden levers of Internet control: An infrastructure-based theory of Internet governance. *Information, Communication & Society*, 15(5):720–738, 2012.

[39] Paul Ferguson and Daniel Senie. Network ingress filtering: Defeating denial of service attacks which employ IP source address spoofing. RFC 2267, 1998.

[40] Robert Gorwa. The platform governance triangle: Conceptualising the informal regulation of online content. *Internet Policy Review*, 8(2):1–22, 2019.

[41] Andy Greenberg. The team of sleuths quietly hunting cyberattack-for-hire services. https://www.wired.com/story/big-pipes-ddos-for-hire-fbi, 2023.

[42] Barry Raveendran Greene. Everyone should be deploying BCP 38! wait, they are. https://www.senki.org/everyone-should-be-deploying-bcp-38-wait-they-are, 2012.

[43] Alok Gupta, Dale O Stahl, and Andrew B Whinston. The Internet: A future tragedy of the commons? In *Computational approaches to economic problems*, pages 347–361. Springer, 1997.

[44] J Alex Halderman. To strengthen security, change developers' incentives. *IEEE Security & Privacy*, 8(2):79–82, 2010.

[45] Chris Hall, Ross J Anderson, Richard Clayton, Evangelos Ouzounis, and Panagiotis Trimintzios. Resilience of the Internet interconnection ecosystem. In *Workshop on the Economics of Information Security (WEIS)*, 2011.

[46] Mikko Hypponen. The Conficker mystery. BlackHat Talk, 2009. URL: https://www.blackhat.com/presentations/bh-usa-09/HYPPONEN/BHUSA09-Hypponen-ConfickerMystery-PAPER.pdf.

[47] Shipi Kankane, Carlina DiRusso, and Christen Buckley. Can we nudge users toward better password management? an initial study. In *Extended Abstracts of the 2018 CHI Conference on Human Factors in Computing Systems*, pages 1–6, 2018.

[48] András Korn. *Defense mechanisms against network attacks and worms.* PhD thesis, Budapest University of Technology and Economics, 2011.

[49] Bill Krogfoss, Marcus Weldon, and Lev Sofman. Internet architecture evolution and the complex economies of content peering. *Bell Labs Technical Journal*, 17(1):163–184, 2012.

[50] Barry M Leiner, Vinton G Cerf, David D Clark, Robert E Kahn, Leonard Kleinrock, Daniel C Lynch, Jon Postel, Larry G Roberts, and Stephen Wolff. A brief history of the Internet. *ACM SIGCOMM computer communication review*, 39(5):22–31, 2009.

[51] Barry M Leiner, Vinton G Cerf, David D Clark, Robert E Kahn, Leonard Kleinrock, Daniel C Lynch, Jon Postel, Lawrence G Roberts, and Stephen S Wolff. The past and future history of the Internet. *Communications of the ACM*, 40(2):102–108, 1997.

[52] Jonathan Lemon. Resisting {SYN} flood {DoS} attacks with a {SYN} cache. In *BSDCon 2002*, 2002.

[53] Nanette S Levinson. Idea entrepreneurs: The United Nations open-ended working group & cybersecurity. *Telecommunications Policy*, 45(6):102142, 2021.

[54] Qasim Lone, Maciej Korczyínski, Carlos H Gañán, and Michel van Eeten. SAVing the Internet: Explaining the adoption of source address validation by Internet Service Providers. In *Workshop on the Economics of Information Security WEIS*, 2020.

[55] Qasim Lone, Matthew Luckie, Maciej Korczyński, Hadi Asghari, Mobin Javed, and Michel van Eeten. Using crowdsourcing marketplaces for network measurements: The case of Spoofer. In *2018 Network Traffic Measurement and Analysis Conference (TMA)*, pages 1–8, 2018.

[56] M3AAWG. Scholl receives 2023 M3AAWG J.D. Falk award for IP spoofing mitigation. https://www.m3aawg.org/blog/2023JDFalkAward-TomScholl, 2023.

[57] Damian Menscher. Practical solutions for amplification attacks. NANOG talk, 2019. URL: https://pc.nanog.org/static/published/meetings/NANOG76/1976/20190610_Menscher_Practical_Solutions_For_v1.pdf.

[58] Alain Mermoud, Marcus Keupp, Kévin Huguenin, Maximilian Palmié, and Dimitri Percia David. Incentives for human agents to share security information: a model and an empirical test. In *Workshop on the Economics of Information Security (WEIS)*, 2018.

[59] Robert T Morris. A weakness in the 4.2BSD Unix TCP/IP software. Computer Science Technical Report 117, Bell Labs, 1985.

[60] New Zealand Financial Markets Authority. Market operator obligations targeted review – NZX. https://www.bloomberg.com/news/articles/2021-02-04/how-a-dated-cyber-attack-brought-a-stock-exchange-to-its-knees, 2021.

[61] William B Norton. Internet service providers and peering. In *Proceedings of NANOG*, volume 19, pages 1–17, 2001.

[62] William B Norton. A business case for ISP peering. *White Paper (v1. 3), February*, 2002.

[63] William B Norton. The evolution of the US Internet peering ecosystem. *Equinix white papers*, 2004.

[64] Nicola Palladino. The role of epistemic communities in the "constitutionalization" of Internet governance: The example of the European Commission High-Level Expert Group on Artificial Intelligence. *Telecommunications Policy*, 45(6):102149, 2021.

[65] Elliott Peterson and Cameron Schroeder. Dismantling DDoS: Lessons in scaling. BlackHat Talk, 2023. URL: https://www.youtube.com/watch?v=9YK7Ugx1MOs.

[66] Sören Preibusch and Joseph Bonneau. The password game: negative externalities from weak password practices. In *International Conference on Decision and Game Theory for Security*, pages 192–207. Springer, 2010.

[67] Roxana Radu. *Negotiating Internet governance.* Oxford University Press, 2019.

[68] Roxana Radu, Matthias C Kettemann, Trisha Meyer, and Jamal Shahin. Normfare: Norm entrepreneurship in Internet governance. *Telecommunications Policy*, 45(6):102148, 2021.

[69] Mohammad S Rahman, Karthik N Kannan, and Mohit Tawarmalani. The countervailing incentive of restricted patch distribution: Economic and policy implications. In *Workshop on the Economics of Information Security (WEIS)*, 2007.

[70] Chris Rose, Jean Gordon, et al. Internet security and the tragedy of the commons. *Journal of Business & Economics Research (JBER)*, 1(11), 2003.

[71] Andreas Schmidt. At the boundaries of peer production: The organization of Internet security production in the cases of Estonia 2007 and Conficker. *Telecommunications Policy*, 36(6):451–461, 2012.

[72] Daniel Senie and Paul Ferguson. Network ingress filtering: Defeating denial of service attacks which employ IP source address spoofing. RFC 2827, 2000.

[73] Gary Shiffman and Ravi Gupta. Crowdsourcing cyber security: a property rights view of exclusion and theft on the information commons. *International Journal of the Commons*, 7(1), 2013.

[74] Kotikalapudi Sriram, Doug Montgomery, and Jeffrey Haas. Enhanced Feasible-Path Unicast Reverse Path Forwarding. RFC 8704, 2020.

[75] Susan Leigh Star. The ethnography of infrastructure. *American Behavioral Scientist*, 43(3):377–391, 1999.

[76] James Sullivan and Jason RC Nurse. Cyber security incentives and the role of cyber insurance. *RUSI Emerging Insights Paper*, 2021.

[77] Jamie Tarabay. How a dated cyber-attack brought a stock exchange to its knees. https://www.bloomberg.com/news/articles/2021-02-04/how-a-dated-cyber-attack-brought-a-stock-exchange-to-its-knees, 2021.

[78] Michel van Eeten. Patching security governance: an empirical view of emergent governance mechanisms for cybersecurity. *Digital Policy, Regulation and Governance*, 19(6):429–448, 2017.

[79] Michel van Eeten and Johannes M Bauer. Emerging threats to Internet security: Incentives, externalities and policy implications. *Journal of Contingencies and Crisis Management*, 17(4):221–232, 2009.

[80] Carla L Wilkin, John Campbell, and Stephen Moore. Creating value through governing IT deployment in a public/private-sector inter-organisational context: A human agency perspective. *European Journal of Information Systems*, 22(5):498–511, 2013.

[81] Daniel W Woods and Rainer Böhme. How cyber insurance shapes incident response: A mixed methods study. In *Workshop on the Economics of Information Security (WEIS)*, 2021.