Defending Against Evolving DDoS Attacks: A Case Study Using Link Flooding Incidents

Min Suk Kang, Virgil D. Gligor(✉), and Vyas Sekar

Carnegie Mellon University, Pittsburgh, PA, USA
{minsukkang,gligor}@cmu.edu, vsekar@andrew.cmu.edu

Abstract. Distributed denial-of-service (DDoS) attacks are constantly evolving. Over the last few years, we have observed increasing evidence of attack evolution in multiple dimensions (e.g., attack goals, capabilities, and strategies) and wide-ranging timescales; e.g., from seconds to months. In this paper, we discuss the recent evolution of DDoS attacks and challenges of countering them. In particular, we focus on the evolution one of the most insidious DDoS attacks, namely link-flooding attacks, as a case study. To address the challenges posed by these attacks, we propose a two-tier defense that can be effectively implemented using emerging network technologies. The first tier is based on a deterrence mechanism whereas the second requires inter-ISP collaboration.

1 Introduction

Distributed denial-of-service (DDoS) has been and is a growing threat to critical services on the Internet. We have observed a dramatic escalation in the number and scale of DDoS attacks during the past few years. For instance, the maximum reported volume of a single attack has been doubled from 300 Gbps [32] to 600 Gbps [23] over the past couple of years. Aside from increasing attack volume, the evolution of DDoS attacks in other dimensions (e.g., number of hosts affected) have not been particularly noticeable during the past few years. In general, a DDoS attack targeted a single system resource (e.g., computation, memory, access bandwidth) for the duration of an attack, utilizing a static set of attack capabilities; e.g., traffic amplification capabilities. Although use of multiple capabilities can diversify an attack, their static use enables detection and blocking by current defense tools [3,5,31].

Recently, however, we have seen evidence of attack evolution in reported DDoS incidents. This ranges from changes of attack goals and capabilities to real-time flexible changes of attack strategy. For example, in 2013, an attack against Spamhaus [9] demonstrated that an adversary can adaptively change the attack targets from end-point servers to routers in Internet exchange points (IXPs) in response to the defense mechanism changes. In 2015, during an attack against ProtonMail [16], we noticed that the adversary also changed the attack strategies in *real-time* to react to defense strategy changes, which creates an interactive game between attackers and defenders [29].

J. Anderson et al. (Eds.): Security Protocols XXIV, LNCS 10368, pp. 47–57, 2017.
DOI: 10.1007/978-3-319-62033-6_7

The goal of this paper is to illustrate the evolution of DDoS attacks and discuss the challenges and opportunities in handling them. In particular, we observe the trend of attack evolution in three dimensions (i.e., goals, capabilities, and strategies) and on both coarse and fine timescales. For instance, attack capabilities evolve on a coarser timescale since they are typically a consequence of changes in technology and Internet economics; e.g., the widespread availability of inexpensive botnets enables the provision of DDoS capabilities as a service [6,22]. In contrast, on a more fine-grained timescale, we observe the evolution of the attack strategies employed using a given set of capabilities and goals; i.e., changing how the available capabilities are employed for a chosen attack goal [9,16,21,29].

In this paper, we focus on a particularly insidious type of DDoS attack, namely link-flooding attacks. Through this case study, we identify three major advantages of adversaries over defenders, which make the mitigation of link-flooding attacks especially challenging. However, we also also see opportunities to defend against these attacks and propose a two-tier defense approach using emerging network technologies.

In Sect. 2, we illustrate the evolution of DDoS attacks in multiple dimensions and timescales, using recent attack incidents. In Sect. 3, we present a new approach for countering these attacks. Section 4 concludes the paper.

2 Evolving DDoS Attacks

DDoS attacks have evoloved on multiple dimensions and timescales, as evidenced by recent incidents. If one defines an attack by the triple ⟨*goal, capabilities, strategies*⟩ [15], one can observe evolution on all three dimensions in the case of these attacks. Moreover, one can also observe evolution on a wide-ranging timescale: a coarse timescale of months or years, and another with a fine timescale of seconds, minutes, or hours. In the rest of this section, we review some of the patterns that we have observed in the case of DDoS attacks.

2.1 New Capabilities

Recent changes in Internet technologies (e.g., the adoption of new protocols) and economics (e.g., pay-per-install botnet markets) have enabled new attack capabilities. In particular, we summarize three noticeable patterns of attack-capability evolution.

Low cost of botnets. Botnets have become an essential commodity of DDoS attacks and the maturity of botnet markets has led to their rapid cost reduction and availability. For example, renting 1,000 bots costs anywhere from a few U.S. dollars to a little more than 100 U.S. dollars [10]. Furthermore, as DDoS attacks begin to marshall emerging Internet-connected devices (e.g., sensors, refrigerators, and dryers) [17,38] and already inexpensive cloud resources, botnets are likely be even more affordable to a wide range of adversaries in the near future.

Attack-traffic amplification. The use of traffic amplifiers to dramatically increase traffic volume (e.g., amplification factors of tens to thousand times) has become popular during the past few years. The ability to amplify attacks is widely available due to the lack of security-aware management of public Internet services (i.e., DNS, NTP) and universally deployed countermeasures against IP spoofing (e.g., ingress filtering [13]); e.g., 20–30% of the Internet ASes cannot detect and block IP spoofing [8].

Flooding core network links. Another new attack capability is provided by the *routing bottlenecks*, which are links in the "middle" of the Internet (i.e., Tier-1/2 networks) that lie on a significant fraction of the traffic to targeted hosts. Flooding these links can severely disrupt connectivity of the host targets [19,21, 33]. Unlike the direct server-flooding attacks, these attacks are indirect as the flow of traffic may not even be destined for the server targets and thus can be stealthy [9,16,21].

2.2 New Strategies

Acquiring new attack capabilities typically occurs on a relatively long timescale. In contrast, changing attack strategies can occur on very short timescales. For example, the adversary observes the defensive posture that the victims adopt and responds by changing how the available attack capabilities are exploited to achieve the same goal; e.g., changing the locus of the attack while maintaining the same set of target hosts in a matter of minutes.

There is already plenty of reported evidence of attack-strategy evolution within real attack campaigns during the past couple of years. In 2013, we witnessed the first large-scale Internet attack where an adversary changed the locus of attack and adapted on a short timescale. That is, after the Spamhaus service was moved to the cloud service (i.e., CloudFlare) in response to a massive DDoS attack, end-host flooding became infeasible. In response, the adversary changed his attack strategy by flooding a few links of four major Internet exchange points (IXPs) in Europe and Asia to degrade the connectivity of the cloud service and implicitly of Spamhaus [9].

An adversary's rapid strategy change was recently observed in a large-scale DDoS attack against ProtonMail, an email provider in Switzerland. Here is a quote from an Internet Service Provider who helped mitigate the attack [29]:

> *"First we moved the BGP IP prefix," said Gargula as he detailed the attack, "I tried to isolate legit human traffic from bot traffic and not to mix it up. We sacrificed one of their three BGP uplink layers as a 'canary' to test the sophistication of the attack. Then we changed the configuration for the IP uplink." The new attackers were incredibly advanced, Gargula explained, and became more sophisticated through the week. "Every time we made a change in tactics, they responded with a change," he said. "It was like Chess: you move a piece, they move a piece. At this point, it became clear that we had a very serious situation on our hands."*

2.3 New Goals

The evolution of DDoS adversary goals is particularly visible in the selection of attack targets. Three new types of critical-infrastructure targets appear to be particularly vulnerable; i.e., emergency, cellular, and power-grid services.

Emergency Services. Increasingly, we see DDoS attacks targeting emergency services [28] by automatically initiating bogus calls. In particular, we see a trend where the emergency networks are utilizing the public Internet since the standardization of the Next Generation 911 (NG911) by National Emergency Number Association (NENA) in the US in 2011 [27]. Unfortunately, by embracing IP technologies and the public Internet, the emergency networks in the US inherit new vulnerabilities. Specifically, the gateways that interconnect the public Internet and the traditional 911 emergency networks can be targets of traditional flooding attacks. This kind of threat is real and significant; e.g., a recent 911 outage in April 2014 showed that even a single device failure (due to a software bug in this case) can cause a 911-service outage affecting about 7 million people in seven US states for six hours [12]. Considering that a single device failure caused a severe outage in emergency services, a well-crafted targeted attack can possibly endanger the public safety of an entire country.

Cellular Services. DDoS attacks against the cellular datacenters can impair cellular service over a large area. According to a recent study on national cellular infrastructure, major cellular carriers in the US in general have only a handful of datacenters throughout the nation [34]. A successful attack against one cellular datacenter (e.g., covering the east coast of the US) could disable the majority of cellular connectivity (both voice and data) of tens of millions of people. Similarly, an adversary could launch large-scale connection degradation (e.g., link-flooding) attacks against cellular-network gateways to the Internet and VoIP servers causing major communication disruption.

Power Grids. Increasing deployment of Internet-of-Things devices, where traditional embedded systems can be programmed and controlled via Internet connections (e.g., Google NEST or Samsung's Smart Home Automation System), can pose significant DDoS challenges for the power grid [25]. Such embedded systems are based on powerful compute platforms with non-trivial processing and network capabilities. Their increased sophistication and features also cause greater threat of compromise. For example, the operation of power grids could be broken by simple on/off cycling HVAC systems over the Internet [11]. On a smaller scale, power surges triggered by attack-induced, server-rack power demand could trip circuit breakers disrupting data center operations [35]. In the simplest case, the attacker overloads the grid by increasing energy consumption. More subtle attacks can lead to cascaded failures or induce persistent load oscillations. In an even more insidious attack, the adversary could use a combination of the grid-overloading attack together with one where access to the

pricing server is denied to magnify the impact by preventing legitimate users from being able to scale back their consumption.

3 Evolving Defenses: A Case Study of Handling Link-Flooding Attacks

Evolving DDoS attacks create significant *advantage* for DDoS adversaries over defenders, which makes it challenging to handle the attacks. To understand the challenges using concrete examples and to illustrate specific opportunities for emerging network technologies as evolving defenses, we focus on a case study of handling link-flooding attacks, one of the most powerful DDoS attacks. First, we identify the three basic challenges of handling link-flooding attacks in Sect. 3.2. Then we present a particular defense strategy based on emerging network technologies that can effectively handle link-flooding attacks in Sect. 3.3.

3.1 Link-Flooding Attacks

A new class of *link-flooding attacks* appeared recently, which have several characteristics that make them hard to handle: (1) they make attack flows *indistinguishable* from legitimate flows[1], and hence become undetectable; (2) they can attack targets *indirectly* in the sense that the locus of attack is different from the actual targets, and hence they cannot be easily detected by the actual targets since they do not receive attack flows; and (3) they are *adaptive* in the sense that they adopt evolving and changing attack postures while achieving the same DDoS goals with the same capabilities; e.g., botnet hosts attack targets as soon as targets recover or deploy a specific mitigation mechanism.

To launch these attacks, an adversary carefully maps the network connectivity infrastructure of the target(s). Having constructed this network map, the adversary identifies routing bottleneck links in the Internet core (e.g., Tier-1/2 networks) that lie on a significant fraction of the traffic to targeted hosts [19,21].

The two real-world attacks mentioned in the previous sections (e.g., attacks against Spamhaus and ProtonMail) utilized different types of link-flooding attacks in different degrees. Moreover, several academic studies have shown the feasibility of link-flooding attacks [21,33].

3.2 Basic Challenges

An effective defense against the link-flooding attacks must address three fundamental challenges, namely *inability to distinguish* attack flows from legitimate ones, the *adversary's asymmetric advantage* over the defender in the Internet, and the *defender's dilemma*.

[1] A flow is defined by 5-tuple, which is a stream of packets having the same source and destination IP addresses, same source and destination port numbers, and same protocol number.

1. **Inability to Distinguish Flows.** The first challenge a defender faces is that an adversary can craft link-flooding flows that are *indistinguishable* from legitimate flows in the targeted routers. The main reason for this is that the immediate targets of the attacks are the router, *not* the end-point hosts. That is, routers are supposed to forward *all* Internet traffic while end-point servers usually are intended to receive *only* certain types of traffic; e.g., web servers expect to see mostly web traffic. Therefore, it is much harder for routers to define and filter out protocol *non-conforming* traffic than end-point servers. For example, the Crossfire attack generates low-rate, protocol-conforming flows and their flows can remain indistinguishable from legitimate traffic in routers [21].
   When attacks use flows that are indistinguishable from legitimate ones, handling link flooding at a target router reduces to a *resource-sharing problem*, where multiple indistinguishable resource requesters (i.e., both legitimate and malicious) contend for the same resource; i.e., the network link bandwidth. In this case, the operation of any requester becomes dependent upon the operation of other—often malicious and unknown—requesters of that resource. The existence of this type of *undesirable dependency* among requesters is the necessary condition for all denial of service in resource-sharing problems [14], and can be countered only by enforcing *agreements* among requesters (i.e., constraints placed on requester behavior) *outside* the shared-resource service [36].
2. **Adversary's Cost Advantage.** Whenever a countermeasure to link flooding reduces to finding a solution to a resource-sharing problem, the cost of the resource for both adversaries and defenders (i.e., the cost of targeted routers bandwidth) becomes a key factor in determining the effectiveness of both attacks and defenses. For example, if the cost of generating attack traffic (i.e., the cost of the shared resource requests) is extremely high, or if the cost of available bandwidth (i.e., resource provisioning) at the target network (i.e., resource manager) is negligible, attacks would become very unattractive.
   Unfortunately, in the current Internet, the opposite cost relation prevails, which makes link-flooding attacks very attractive. That is, the cost of bandwidth for generating attack traffic is *orders of magnitude* lower than that of provisioning backbone-link bandwidth.[2] In other words, a severe *cost asymmetry* exists that favors the adversary over the defender. Furthermore, removing this cost asymmetry is not only a matter of changing the Internet design. Instead, whether the asymmetry can be removed depends on two independent markets: the botnet markets and backbone bandwidth markets. The former is an underground online e-commerce market [10], where bot buyers can demand and sellers supply attack bots, whereas the latter is a legitimate

[2] For example, the adversary's cost of flooding a 10 Gbps with bots whose uplink bandwidth is only 1 Mbps averaged about $920 with a minimum of about $80 in the US in 2011 [10]. In contrast, the cost of 10 Gbps bandwidth in Internet transit was about $6,300 in 2015 [1]. This represents a cost advantage of 7–80 times of the adversary over the defender.

network-infrastructure market, where many entities compete by well-established rules that determine the market price of backbone bandwidth.[3] We note that removing the cost asymmetry, and even reversing it to favor the defender, does *not* completely deter link-flooding attacks; e.g., cost-insensitive adversaries, such as those sponsored by a state, could still launch link-flooding attacks. However, it would change today's severely imbalanced cost structure and would certainly deter cost-sensitive (e.g., rational) adversaries. Hence, it would yield an effective first line of defense as argued in the later this section.

3. **Defender's Dilemma.** Many link-flooding attacks rely on the existence of *a few* link targets whose congestion would disrupt the majority of routes that pass the Internet core from a set of sources to a set of destination hosts. These links are called the *routing bottleneck* of a set of sources and destinations and shown that its existence is an *undesirable artifact* of Internet design [19]. Although in the attack-free mode of operation these bottlenecks are not an operational hazard, we seek to remove them since they can constitute an Internet vulnerability in the presence of a link-flooding adversary.
   However, as shown in a recent measurement study [19], removing routing bottlenecks to prevent link flooding is impractical in the current Internet because they are the result of employing a *cost-minimizing* (or revenue-maximizing) policy of the Internet routing and topology designs. In other words, the source of many link-flooding vulnerabilities is, in fact, a very *desirable* feature of the Internet business model. Hence, a defender faces the following dilemma: *how can one remove a vulnerability of a system when it is caused by a very desirable feature of the system's design and operation?*
   As long as the causality between a route-cost minimization policy and the existence of flooding targets holds, any attempt to remove the latter would necessarily affect the former. However, in the highly competitive Internet transit markets ISPs would naturally be very reluctant to adopt any countermeasure that would increase the operating cost.

3.3 Evolving Defense: Two-Tier Approach

We argue that new defense mechanisms become necessary to counter the basic challenges of handling link-flooding attacks. To that end, we present a *two-tier* defense approach. In the two-tier approach, a *first-line* defense uses low-cost, light-weight, and readily-deployable mechanisms to handle frequently-used attacks while a *second-line* defense is invoked to perform high-cost defense mechanisms only for infrequent attacks that have not been handled by the first-line defense. Considering the defender's dilemma, the use of first-line defenses is very desirable *even if they do not counter all possible attacks* because they render the use of complex high-cost mechanisms for handling uncommon adversaries necessary only infrequently.

[3] The market involves many layers of businesses, including equipment companies, optical cable companies, undersea cable companies, Internet exchange points (IXPs), etc.

First-Line Defense: Deterrence. In our first-line defense, we focus on the *attack deterrence*, particularly targeting *rational* adversaries *only*; i.e., *cost-sensitive* adversaries who wish to remain *undetected*. All other (e.g., irrational, cost-unbounded) adversaries would not be deterred by the first-line defense and need to be countered by a second line of defense.

We believe that the majority of link-flooding attackers are rational in the current DDoS attack landscape. According to a behavioral economics study [30], there is strong evidence that cyber criminals are economically motivated. Also, rational adversary behaviors in DDoS attacks are observed in a recent study that analyzed real DDoS attack incidents in 240 countries over 5 years [18]. Note that if the cost of bots drastically decreases to become almost negligible in the future, most adversaries could become cost-insensitive, making the invocation of a second-line of defense mechanism necessary. However, even when fewer cost-sensitive adversaries are deterred, a first-line defense mechanism would be useful since it has low-cost deployment and operation cost.

Cost-detectability tradeoff. To deter economically motivated (or cost-sensitive) adversaries, we focus on the aforementioned adversary's cost-asymmetry advantage with respect to defenders. Our approach is to reduce (or even reverse) the cost asymmetry to deter cost-sensitive adversaries. We create an *untenable tradeoff* between the cost and detectability. By definition, any countermeasure that can either substantially increase the attack cost relative to the defense cost or induce detectability will deter attacks by rational adversaries.

In recent work we showed that it is possible to force a link-flooding adversary into an untenable tradeoff by using only intra-domain network operations. The proposed system, which we called SPIFFY [20], implements a mechanism to *logically* increase the bandwidth of a targeted link by a large factor (e.g., 10 times) temporarily utilizing flexible intra-domain route control capability implemented by software-defined networking capabilities; e.g., OpenFlow [2], OpenSketch [37]. After the increase, SPIFFY attempts to distinguish attack traffic sources from legitimate ones by observing their response to the temporary bandwidth increase: legitimate sources running TCP-like flows will naturally see a corresponding increase in their throughputs as the bandwidth of their bottleneck link (i.e., the targeted link) has increased; however, attack sources will not observe this increase as a rational cost-sensitive attacker would have chosen to *fully utilize* the available bandwidth of the upstream links of the attack sources in the first stage. Alternatively, to avoid detection, the attacker could choose to keep each bot's attack traffic rate much lower than the available bandwidth of its upstream link. However, that this will increase the number of required bots and thus *increase attack cost* proportionally. In essence, adversaries are forced to either allow their attack sources to be detected or accept an increase in attack cost.

Note that adversaries *cannot* predict *when* a temporary bandwidth expansion will be executed since its operation is determined by defenders; e.g., network operators of link targets. The unpredictability of bandwidth expansion makes it difficult for the adversaries to temporarily purchase additional bots at a low cost to avoid the SPIFFY's cost-detectability tradeoff. This is because the botnet

markets *cannot* provide low-cost temporary bot purchase since the markets are required to reserve large numbers of bots that are *always ready* for the temporary bot demands.

Second-Line Defense: Collaborative Defenses. Countering link-flooding attacks by cost-insensitive, irrational adversaries requires collaboration among multiple ISPs. For example, CoDef [24] requires coordination between the ISPs hosting the attack sources and targets to defend against link-flooding attacks. SENSS [4] requires coordination between ISPs hosting the attack target and the intermediate ISPs that control the incoming attack traffic. SIBRA [7] utilizes global coordination among ISPs to create entire end-to-end Internet paths to protect a user's traffic. Although ISP collaboration-based defenses are generally harder to orchestrate in a climate of competitive relationships between ISPs in the current Internet [26], when they are used as a second-tier defense they can be effective for the less frequent cases where the adversary is cost-insensitive or irrational.

4 Conclusions

Evolving DDoS attacks that target the critical-infrastructure services (e.g., emergency, power grids, and cellular communications services) require new countermeasures that are currently unavailable on the Internet. As a case study, we investigate link-flooding attacks and discuss the particular challenges and opportunities in handling them. We argue that defenses against link-flooding attacks should be multi-tiered. We provide an example of a two-tier defense scheme where the first line of defense deters cost-sensitive rational adversaries, who appear to be responsible for the vast majority of DDoS attacks. The second tier is a collaborative defense intended to counter attacks by cost-insensitive or irrational adversaries, which can be more costly since it is infrequent in practice.

References

1. Internet transit pricing: historical and projected. http://drpeering.net/white-papers/Internet-Transit-Pricing-Historical-And-Projected.php
2. Open flow. https://www.opennetworking.org
3. Akamai: The state of the internet 2nd quarter. Report (2012)
4. Alwabel, A., Yu, M., Zhang, Y., Mirkovic, J.: SENSS: observe and control your own traffic in the Internet. In: Proceeding of ACM SIGCOMM (2014)
5. Arbor Networks: Worldwide infrastructure security report, volume IX. Arbor Special Report (2014)
6. Barker, I.: 2016 will see the rise of DDoS-as-a-service. In: BetaNews (Dec 28 2015). http://betanews.com/2015/12/28/2016-will-see-the-rise-of-ddos-as-a-service/
7. Basescu, C., Reischuk, R.M., Szalachowski, P., Perrig, A., Zhang, Y., Hsiao, H.C., Kubota, A., Urakawa, J.: SIBRA: Scalable internet bandwidth reservation architecture. In: Proceeding of NDSS (2016)

8. Beverly, R., Koga, R., Claffy, K.: Initial longitudinal analysis of IP source spoofing capability on the Internet (2013)
9. Bright, P.: Can a DDoS break the Internet? Sure.. just not all of it. In: Ars Technica (2 April 2013). http://arstechnica.com/security/2013/04/can-a-ddos-break-the-internet-sure-just-not-all-of-it/
10. Caballero, J., Grier, C., Kreibich, C., Paxson, V.: Measuring pay-per-install: The commoditization of malware distribution. In: Proceeding of USENIX Security (2011)
11. Cerf, V.: The freedom to be who you want to be: strong authentication and pseudonymity on the internet. In: RSA Conference (2013)
12. FCC: April 2014 Multistate 911 Outage: Cause and Impact. Public Safety Docket No. 14–72, PSHSB Case File Nos. 14-CCR-0001-0007 (2014)
13. Ferguson, P.: Network ingress filtering: Defeating denial of service attacks which employ IP source address spoofing. IETF RFC2827 (2000)
14. Gligor, V.D.: A note on the denial-of-service problem. In: Proceeding of IEEE Security and Privacy (1983)
15. Gligor, V.: Dancing with the adversary: a tale of wimps and giants. In: Christianson, B., Malcolm, J., Matyáš, V., Švenda, P., Stajano, F., Anderson, J. (eds.) Security Protocols 2014. LNCS, vol. 8809, pp. 100–115. Springer, Cham (2014). doi:10.1007/978-3-319-12400-1_11
16. Goodin, D.: How extorted e-mail provider got back online after crippling DDoS attack. In: Ars Technica, (10 November 2015). http://arstechnica.com/security/2015/11/how-extorted-e-mail-provider-got-back-online-after-crippling-ddos-attack/
17. Greene, T.: Bot-herders can launch DDoS attacks from dryers, refrigerators, other Internet of things devices. In: NetworkWorld (24 September 2014)
18. Hui, K.-L., Kim, S.-H., Wang, Q.-H.: Marginal deterrence in the enforcement of law: evidence from distributed denial of service attack. In: Workshop on Analytics for Business, Consumer and Social Insights (BCSI). Singapore, August 2013
19. Kang, M.S., Gligor, V.D.: Routing bottlenecks in the internet: causes, exploits, and countermeasures. In: Proceeding of ACM CCS (2014)
20. Kang, M.S., Gligor, V.D., Sekar, V.: SPIFFY: Inducing Cost-Detectability Tradeoffs for Persistent Link-Flooding Attacks. In: Proceedings of NDSS (2016)
21. Kang, M.S., Lee, S.B., Gligor, V.D.: The Crossfire Attack. In: Proceeding of IEEE S and P (2013)
22. Karami, M., McCoy, D.: Understanding the emerging threat of DDoS-as-a-service. In: Proceeding of USENIX Workshop on Large-Scale Exploits and Emergent Threats (LEET) (2013)
23. Khandelwal, S.: 602 Gbps! This may have been the largest DDoS attack in history. In: NetworkWorld (8 January 2016)
24. Lee, S.B., Kang, M.S., Gligor, V.D.: CoDef: collaborative defense against large-scale link-flooding attacks. In: Proceedinf of ACM CoNEXT (2013)
25. Mo, Y., Kim, T.H.J., Brancik, K., Dickinson, D., Lee, H., Perrig, A., Sinopoli, B.: Cyber-physical security of a smart grid infrastructure. Proc. IEEE **100**(1), 195–209 (2012)
26. Mortensen, A.: DDoS Open Threat Signaling Requirements. IETF draft-mortensen-threat-signaling-requirements-00 (2015)
27. NENA: NENA i3 Technical Requirements Document. NENA VoIP/Packet Technical Committee Long Term Definition Working Group (2006)

28. Nussman, C.: DHS Bulletin on Telephony Denial of Service (TDOS) attacks on PSAPs. In: National Emergency Number Association (NENA), (17 March 2013). https://www.nena.org/news/119592/DHS-Bulletin-on-Denial-of-Service-TDoS-Attacks-on-PSAPs.htm
29. Patterson, D.: Exclusive: inside the ProtonMail siege: how two small companies fought off one of Europe's largest DDoS attacks. In: TechRepublic, (13 November 2015). http://www.techrepublic.com/article/exclusive-inside-the-protonmail-siege-how-two-small-companies-fought-off-one-of-europes-largest-ddos/
30. Png, I.P., Wang, C.Y., Wang, Q.H.: The deterrent and displacement effects of information security enforcement: International evidence. J. Manag. Inf. Syst. **25**, 125–144 (2008)
31. Rossow, C.: Amplification hell: revisiting network protocols for DDoS abuse. In: Proceeding of NDSS (2014)
32. Storm, D.: Biggest DDoS attack in history slows Internet, breaks record at 300 Gbps. In: ComputerWorld (27 March 2013)
33. Studer, A., Perrig, A.: The coremelt attack. In: Backes, M., Ning, P. (eds.) ESORICS 2009. LNCS, vol. 5789, pp. 37–52. Springer, Heidelberg (2009). doi:10.1007/978-3-642-04444-1_3
34. Xu, Q., Huang, J., Wang, Z., Qian, F., Gerber, A., Mao, Z.M.: Cellular data network infrastructure characterization and implication on mobile content placement. In: Proceeding of ACM SIGMETRICS (2011)
35. Xu, Z., Wang, H., Xu, Z., Wang, X.: Power attack: An increasing threat to data centers. In: Proceeding of NDSS (2014)
36. Yu, C.F., Gligor, V.D.: A formal specification and verification method for the prevention of denial of service. In: Proceeding of IEEE Security and Privacy (1988)
37. Yu, M., Jose, L., Miao, R.: Software defined traffic measurement with opensketch. In: Proceeding of USENIX NSDI (2013)
38. Yu, T., Sekar, V., Seshan, S., Agarwal, Y., Xu, C.: Handling a trillion (unfixable) flaws on a billion devices: Rethinking network security for the Internet-of-Things. In: Proceeding of HotNets (2015)