

A Selective Survey of DDoS Related Research

Loránd Jakab
Universitat Politècnica de Catalunya
ljakab@ac.upc.edu

Jordi Domingo-Pascual
Universitat Politècnica de Catalunya
jordid@ac.upc.edu

Abstract

The open nature of the Internet makes distributed denial-of-service (DDoS) attacks relatively easy to mount and difficult to counter. In the ten years since the first attacks were observed several methods of defense were proposed, yet the problem still lacks a complete solution. We present a selective survey of proposed defense mechanisms and show possible research directions.

1 Introduction

As the Internet grew through the years to its current size and became a commodity, its user base changed from that of scientists, financial institutions and military to virtually anyone with a desire to be connected. The user base now resembles society in general, and social rules showed us that wherever there's a group of people there will be disagreement and conflict. On the Internet this manifested itself with flame wars on newsgroups and e-mail lists, web site defacement, viruses, worms, and so on. Some of the early attacks were only a means for hackers to get fame and attention, play pranks on each other, but it gradually reached the level of organized crime. Misbehavior by disrupting legitimate network activity can occur on different layers of the OSI reference model.

In the case of the *physical layer* the attacker needs physical access to the medium, or be able to remotely control links. These attacks range from cutting network cables to shutting down interfaces on routers to jamming radio frequencies in wireless networks.

Data-link (MAC) layer attacks are specific to each communication medium. They exploit weaknesses and shortcomings in the design of the MAC layer protocol in order to gain advantage over other users or to disrupt legitimate activity. Since the MAC protocol regulates access to a shared medium, mounting an attack at this layer usually requires physical access too.

The *network (IP) layer* makes it possible to mount attacks from any corner of the Internet (physical access to the

attacked resources not required). Another factor that makes misbehavior at the IP layer (and above) difficult to counter is the ability to spoof the source address of packets. The Internet architecture is based on the principle that *the core network should be kept as simple as possible, pushing complexity to the edge*. While this design principle allowed this packet switched inter-network to scale to its current size, it also introduces a number of problems, one of them being the above mentioned ability to spoof packets.

Transport/application layer misbehavior is specific to the application/protocol involved. However, the important applications and protocols that are at the base of the correct functioning of the Internet, like TCP or DNS are the most targeted by attackers, because they must be deployed on all connected hosts. The protection against these types of misbehavior is also application specific, with some of them easier to secure, other quite difficult.

One of the most common types of misbehavior is the *denial-of-service* attack.

1.1 Definition

In general terms *denial-of-service* occurs when an entity cannot perform an action, access a service that he is entitled to. In the networking world this usually means that a legitimate node on the network is unable to reach another node or their connectivity is severely degraded.

A *distributed denial-of-service* attack occurs when the attackers use several machines to launch the attack, making it more powerful. With a few thousands of DSL-connected home computers an attacker could saturate the well provisioned link of a major website.

In the following, we will focus on DDoS attacks at the IP layer.

1.2 Short History

The origins of DoS attacks are traced back to the early '90s, when it was a means for revenge or simply to show off to others in the online gaming and IRC (Internet Relay Chat) communities [10]. Typically it was a result of an

insult on the chat channels; the insulted person looked up the IP address of the other person, and used some of the available application layer DoS attacks or, if had a higher bandwidth connection, a simple flooding attacks. Some of the more “sophisticated” attackers had university accounts from which flooding a modem connection was a trivial task.

In 1996 an asymmetry in the TCP protocol handshake was found and exploited successfully for a long time. Known as TCP SYN-flood attacks, they were easy to perform, because using small packets one could deplete the TCP state buffer with half-open connections. And using spoofed packets (which was possible, because the TCP handshake didn’t have to complete) the source was very difficult to detect. Other protocol related vulnerabilities were also used to perpetrate attacks in this timeframe.

The next level for DoS was represented by reflection attacks. One attacker would send ICMP packets spoofed with the victims address to a broadcast address and several machines on the subnet would respond, thus amplifying many-fold the single attack packet (Smurf attack). Another type of amplification attack, the DNS reflector works by sending small recursive queries to DNS server spoofed with the victim’s address, which would get the (much larger) responses from the server.

When bandwidth between attackers and targets become more equal, using one machine for attacks wasn’t enough anymore, so attention was shifted to controlling several attack machines remotely. This gave birth the DDoS attacks. Initially it was restricted to very skilled attackers, and everything was done manually. Over the years attack tools become more and more automated and integrated, so less skilled attackers (know by the pejorative *skript kiddies*) would also be able to use them. These tools were traded amongst individuals and groups in the underground, first just for other tools, but more and more for money.

Nowadays not only is it possible to buy these toolkits, it is also possible to buy or rent large *botnets* [8]. And as several cases demonstrate, the motivations have also evolved from pranks and “supremacy wars” between underground hackers to extortion and political pressure. One high-profile case is that of anti-spam firm Blue Security, who actually stopped its anti-spam activities due to a large and very long DDoS attack.

2 Previous Work

An good starting point in this research field is *The DDoS* book [10] by Mirkovic *et al.* To our knowledge this is the only book that exclusively deals with this topic. It is primarily aimed at computer system and network administrators and describes the origins, motivations, methodology and effects of these types of attacks, information necessary to defend against them. Several practical defense approaches are

then presented, both proactive and reactive.

Proactive approaches focus on securing end hosts on the network, designing good firewall rules and well balanced network design, which includes redundancy for management tasks. With good proactive measures one will not become the source of attacks and can better bear the effects of an attack.

The reactive approaches consider actions one can do after detecting an attack. These may include filtering attacks packets, if they can be identified, asking help from upstream ISPs, make use of backup equipment and so on.

While not the main intended audience, researchers will find a detailed chapter on the solutions proposed by the academic community, which were either theoretical or in a prototype phase at the time of writing. The authors point out that the seriousness of these attacks will only increase as they are now mostly economically motivated.

Two of the authors published a taxonomy of DDoS attacks and defense mechanism in [11], a short version of which is included in the book. It offers a good overall picture of this research field and facilitates a deeper understanding of the phenomenon. See below a short overview of the defense mechanisms’ classification, from three points of view:

- **Activity level**

- *Preventive*. This concerns mostly system and network administrators as it’s related to optimizing network configuration and securing end systems.
- *Reactive*. This is of most interest to the academic research community as it relates to pattern and anomaly detection followed by an adequate response.

- **Cooperation degree**. In case of using several different defense mechanisms (which is recommended by the authors) these could be autonomous, cooperative or interdependent.

- **Deployment location**. Defense mechanisms can be deployed at the victim network, in the intermediate network, or at the source network. Each of these locations has it’s advantages and disadvantages and researchers debate strongly which one is the best. In [10] the authors suggest that a larger cooperative defense system could be the best solution, which would combine components in all three locations. A project in this direction was proposed in [12].

Defenses at the **victim network** seem to be the most popular. This is understandable, since it doesn’t require the cooperation of an external entity, being the most feasible and cost-effective solution. It is at this location that it’s easiest

to know when an attack is ongoing as the farther the defense location, the more incomplete its view about what is going on related to that site will be. Keeping complete control of the defense system is also a definite plus.

These systems however may themselves become targets of DoS attacks, but most importantly they can just be overwhelmed by the sheer amount of traffic coming their way. Imagine a client who's link to his ISP is saturated by an attack; his defense system will not be able to help him.

Other defense systems focus on the **source network**. These mechanisms should be deployed at all locations where possible attacks could originate from, practically in all networks connected to the Internet. At this place attack flows are not so aggregated yet, so it would put less burden on the defense systems to analyze them. And since they would be cut off at the source, it would save transit networks from transporting malicious traffic.

This approach however requires a very large scale deployment in order to be effective. And since attack streams in the source network usually are small in volume, they may be more difficult to detect. Deployment motivation is low since by deploying these systems the owners of the network actually help others to protect their networks, and need everyone else to also deploy it to protect their own.

Intermediate network deployment usually means ISP and/or the core of the Internet. As most Internet traffic has to go through a relatively small number of core Internet routers, it would seem the ideal place for a DDoS defense system. The authors of [10] suggest that *if core defenses were effective, accurate, cheap, and easy to deploy, they could thus completely solve the problem of DDoS attacks.*

The core routers are able to do what they do (route a huge amount of traffic) because of the above mentioned "complexity at the edges" paradigm. There is a general reluctance to increase their complexity and the benefits have to be very convincing. Distinguishing attack traffic inside the huge volume of traffic they have to handle is not a trivial task, and the risk of causing collateral damage, *i.e.*, filtering legitimate packets in the core is not acceptable. Designing an DDoS defense system for the core is very challenging task.

3 Commercial Solutions

As the goal of researchers is to find solutions that can be deployed in practice, it is interesting to see which ideas have been implemented in commercial solutions so far. While the inner workings of their devices are not released publicly, we can usually know which type of defense they employ. As you will see, these are only partial solutions, that improve the resistance of the protected systems or networks to the attack, and they have to be tuned for a particular customer. None of these devices offers a complete solution.

Arbor Networks' [1] appliance uses signature-based detection, where attack packets must be precisely characterized in order to detect an ongoing attack. It is fairly easy for attackers to modify the attack traffic packets, so keeping a signature up to date is essential. It shares the disadvantages of the software virus scanners, not being able to detect new attacks and targeted attacks.

Most commercial solutions are based on statistical filtering [2, 5, 3], creating a normal traffic model of the site where they are installed and generating alerts when observed traffic is diverging from this baseline model. Due to the probability of false positives they don't usually automatically take filtering actions, instead network engineers respond to the generated alerts.

Another technology adopted by industry [4, 6] is packet scoring and selective discarding, based on some dynamic score thresholds.

4 Recent Work

Before presenting some of the recent research papers, we would like to point out the concern of the IETF on the matter: RFC4732 [9] was published November 2006, in order to encourage protocol designers and network engineers towards more robust designs. They reveal some partial solutions to reduce the effectiveness of attacks, and highlight how some of these partial attacks can be used by attackers to perpetrate alternative attacks. In light of all this, the editors conclude that proper *architectural* solutions are lacking, and encourage research into architectural solutions that might be feasible and cost-effective for deployment. To us, this means that solutions that patch one problem and possibly introduce another is not desirable, however radical solutions are also to be avoided, as they will not be feasible and cost-effective to deploy.

The document contains suggestions for mitigation strategies that are well documented in the literature:

- **Protocol design.** Avoiding protocol asymmetry, making difficult to simulate a legitimate user by introducing proof of work tests, graceful routing degradation on overload, and careful design of of autoconfiguration and authentication protocols.
- **Network design and configuration.** Network administrators should consider redundancy and distributed service in order to protect their valuable resources.
- **Router implementation issues.** The most important issues are to correctly determine who sent any given message, if the content conforms to protocol formats and correctly handle resource exhaustion.
- **End-to-end implementation issues.** Minimize state lookup complexity, and have a good monitoring frame-

work as to observe network activity especially anomalies.

In the following paragraphs we will present some very recent research work in this field.

A recent proposal by Walfish *et al.* has an interesting and maybe radical idea at its base. DDoS defense by offense tries to mitigate the effect of an ongoing attack by asking clients to send higher volumes of traffic. The authors start from the presumption that attackers already use all their available bandwidth in order to perform the attack, while the traffic of good clients, due to congestion avoidance rules decreases or even connections get dropped. By sending higher traffic volume, good clients can claim a much larger share of the server's resources.

Yaar *et al.* propose a packet marking scheme called *StackPi* in [15], improving on a previous proposal (called simply *Pi*) in [14]. This packet marking scheme, and the related filtering mechanisms intend to solve the problem introduced by IP spoofing, identifying packets that come from the same source, regardless of the source address field of the IP header. In order to do so, they use the 16 bits of the Identification field of the IP header on packets to mark the path the packet has travelled. Each packet following the same path would receive the same marking. However, collisions are possible and this may mean collateral damage, if filtering is applied.

In this approach the detection is left to a different mechanism. The *StackPi* filtering system would be notified by an independent detection engine that a given packet is part of an attack. This would result in applying a filter to packets bearing the same marking as the attack packet.

In order to work, the system needs wide deployment. With analytical models the authors show that a minimum of 20% of all routers on the Internet should deploy the packet marking scheme in order to have some degree of DDoS protection.

Tan *et al.* propose a statistical filtering framework tailored to the special needs of MANETs [13]. They first simulate DDoS attacks on the ad-hoc network infrastructure, originating from within the network (a case they don't consider is when this network is connected to the wired Internet and the attacks originate there), and conclude that packet delivery ratio and average end-to-end delay decrease significantly, especially with high mobility. The proposed filtering framework is not evaluated.

In interesting architectural work [7] by Allman *et al.* proposes a high-level misbehavior-detection framework. It is based on entities called *detectives* and *witnesses*, distributed in the Internet, and an entity that gathers and disseminates information. Witnesses are virtually everywhere and the system doesn't necessarily trust them. Instead, it relies on the detectives, which are few in number, all known, and trustworthy to "judge" the information from witnesses, who

only "log the facts". Since detectives cannot assume that witnesses are trustworthy, just like in the real world they may query several witnesses in the same area to decide what "actually happened" and if that event is important or not.

In the networked world, these entities would translate into sophisticated intrusion detection systems (detectives) and traffic capture/monitoring points (witnesses). It would be interesting to build a prototype and test such a system in an IP network testbed.

Considering that the IETF suggests architectural solutions [9] and that we believe in a distributed solution for the DDoS problem, we are interested in finding a proper implementation for the above mentioned proposal. In fact, this architecture need not be directly implemented for the whole Internet: e.g., a community network may set up its own internal system for protection limited to internal threats. We are most interested in this scenario.

4.1 Other considerations

A lot of research has gone into defending against packet spoofing, considered by many the most important reason why DDoS attacks are so difficult to counter. Attackers send spoofed packets to make setting up filtering rules more difficult. While solving the spoofing problem would certainly make DDoS attacks more difficult to perform, it would only raise the bar: building large enough botnets is just a question of strong enough motivation and with such botnets no spoofing is required to mount a successful attack.

5 CONTENT

Within CONTENT activities our work is related to misbehavior detection in community networks (at the network layer). Misbehavior can mean a lot of things. In our view, network activity which has an adverse effect on community resources is undesirable and is considered misbehavior. The goal of the project is to advance research in the field of "easy-to-install and easy-to-use AV services in and between homes." As pointed out before, wherever there is a group of people, there will be disagreements, and several times personal interest is above group interest. An easy-to-use service must offload the user from security concerns, thus the design must include security from the start. While protocol design is of utmost importance, in a networked environment monitoring key points of the network is a must, for detecting anomalous activity and responding to it.

Intrusion Prevention/Detection Systems serve this purpose, and we are interested in adapting experience in this field to the world of community networks.

6 Conclusions

An important point to make (and the existence of RFC4732 [9] underlines this) is that protocols of emerging network technologies should be designed with security in mind, not by adding a security layer later. This especially holds true for susceptibility to DoS attacks, which can only be diminished by a carefully designed architecture.

References

- [1] Arbor Networks. <http://www.arbornetworks.com>.
- [2] Cisco Guard, Cisco Systems. <http://www.cisco.com/en/US/products/ps5888/index.html>.
- [3] Cyber Operations. <http://www.cyberoperations.com>.
- [4] Lancope. <http://www.lancope.com>.
- [5] Mazu Networks. <http://www.mazunetworks.com>.
- [6] Webscreen Technology. <http://www.webscreen-technology.com>.
- [7] M. Allman, E. Blanton, V. Paxson, and S. Shenker. Fighting coordinated attackers with cross-organizational information sharing. In *Proc. ACM HotNets-V*, Nov. 2006.
- [8] P. Barford and V. Yegneswaran. *An Inside Look at Botnets*. Advances in Information Security. Springer, 2006.
- [9] M. J. Handley and E. Rescorla. Internet denial-of-service considerations. RFC 4732 (Informational), Nov. 2006.
- [10] J. Mirkovic, S. Dietrich, D. Dittrich, and P. Reiher. *Internet Denial of Service: Attack and Defense Mechanisms*. Prentice Hall PTR, Dec. 2004.
- [11] J. Mirkovic and P. Reiher. A taxonomy of DDoS attack and DDoS defense mechanisms. *ACM SIGCOMM Computer Communication Review*, 34(2):39–53, Apr. 2004.
- [12] J. Mirkovic, M. Robinson, and P. Reiher. Alliance formation for DDoS defense. In *NSPW '03: Proceedings of the 2003 workshop on New security paradigms*, pages 11–18, Aug. 2003.
- [13] H.-X. Tan and W. K. G. Seah. Framework for statistical filtering against DDoS attacks in MANETs. In *ICISS '05*, Dec. 2005.
- [14] A. Yaar, A. Perrig, and D. Song. Pi: A path identification mechanism to defend against DDoS attacks. In *Proc. IEEE Symposium on Security and Privacy*, pages 93–107, May 2003.
- [15] A. Yaar, A. Perrig, and D. Song. StackPi: New packet marking and filtering mechanisms for DDoS and IP spoofing defense. *IEEE Journal on Selected Areas in Communications*, 24(10):1853–1863, Oct. 2006.