A Taxonomy of Criteria for Evaluating Defence Mechanisms against Flooding DoS Attacks

Jarmo V. E. Mölsä

Communications Laboratory, Helsinki University of Technology,
P.O. Box 3000, FI-02015 TKK, Finland
jarmo.molsa@tkk.fi

Abstract. This paper describes a set of criteria for evaluating defence mechanisms against flooding denial of service (DoS) attacks. Effectiveness and usefulness of a defence mechanism in mitigating a DoS attack depends on many issues which are presented here in the form of a taxonomy. The primary goal of this taxonomy is to help in getting a comprehensive view on both the strengths and weaknesses of a specific defence mechanism. A good defence mechanism should not disturb legitimate traffic when there is no attack, should mitigate the amount of attack traffic well enough, should increase the quality of service (QoS) available to legitimate traffic during an attack, and should use as little resources in this task as possible. In addition, any defence mechanism should be robust against changes in attack characteristics and intentional misuse.

1 Introduction

Flooding denial of service (DoS) attacks [1] cannot be mitigated in any straightforward way. These attacks use brute force to exploit normal behaviour of protocols and services in an excessive manner. As attack traffic looks similar to legitimate traffic it is not possible to simply block all attack traffic, because some legitimate flows will be misclassified as attack flows (false positives), and some attack flows will not be detected at all (false negatives) [2–4]. Another difficulty in mitigating flooding DoS attacks is the collateral damage from using a defence mechanism, such as performance degradation in routers. A defence mechanism may cause collateral damage even if an attack is not present. Defences also increase the complexity of a system, which results in more vulnerabilities available for an attacker to exploit. All these issues should be considered when evaluating effectiveness and usefulness of defence mechanisms against flooding DoS attacks. Distributed DoS (DDoS) attacks using multiple sources at the same time are considered as subtypes of DoS attacks in this paper.

Evaluation of defence mechanisms is typically based on very limited criteria. Sometimes evaluation is carried out under ideal conditions, where there are no false positives and no collateral damage, such as when studying rate-limiting in [5]. Three existing preventive defence mechanisms were compared in [6], and it was found out that they were originally evaluated according to restricted criteria, ignoring the risk of attacks that they cannot solve. In [7] rate-limiting is carried out against any excessive traffic aggregate, regardless of being malicious or benign, but it emphasizes the difficulty in differentiating between innocent and attack traffic. In [8] it is argued that a relatively complex overlay structure to mitigate DoS attacks does not introduce any new vulnerabilities, which is not a reasonable claim. [9] describes a defence against TCP SYN attacks and considers shortly the possibility of false positives and false negatives, but still the final evaluation is based on the easiest test cases. Some commercial products promise to "prevent DoS attacks", which implies that difficult issues in mitigating flooding DoS attacks are forgotten. All these examples show that it is difficult to make an extensive evaluation of a defence mechanism. It should be emphasized here that all well-specified defences have definitely their application areas regardless of a restricted evaluation. The point here is that an extensive evaluation will help to identify the limitations and the best application areas, and to get the best benefit out of a defence mechanism.

An organized approach for evaluating defence mechanisms is missing. As expressed in [6], very little has been done to compare, contrast, and categorize the different ideas related to DoS attacks and defences. A user of a defence mechanism should be able to consider both advantages and disadvantages in an objective way. A taxonomy of DDoS attack and DDoS defence mechanisms is presented in [10]. This taxonomy, however, lacks all evaluation issues.

The main contribution of this paper is to present a taxonomy of criteria for evaluating defence mechanisms against flooding DoS attacks. This taxonomy will point out important issues in evaluating a defence mechanism, which will make it easier to carry out a more realistic evaluation. The taxonomy can be considered as an extension to the taxonomy defined in [10]. Together these two taxonomies will help to understand how attacks and defences operate, and what issues have an effect on the effectiveness and usefulness of defence mechanisms.

2 Taxonomy of Evaluation Criteria

This section presents the taxonomy of criteria for evaluating defence mechanisms against flooding DoS attacks. It contains important issues in estimating the effectiveness and usefulness of a defence. The taxonomy is shown in the Fig. 1.

2.1 Effectiveness

A defence mechanism should be effective in mitigating a flooding DoS attack. This attack time effectiveness, however, is not enough because a defence mechanism should also be effective when there is no attack, i.e. during normal time a defence should disturb legitimate traffic as little as possible.

Normal Time. Some defence mechanisms are active continuously both during normal and attack time. In the taxonomy of DDoS defence mechanisms specified in [10] these are called preventive mechanisms. Normal time effectiveness of a defence mechanism is very important because attacks against a specific target are rather rare after all.

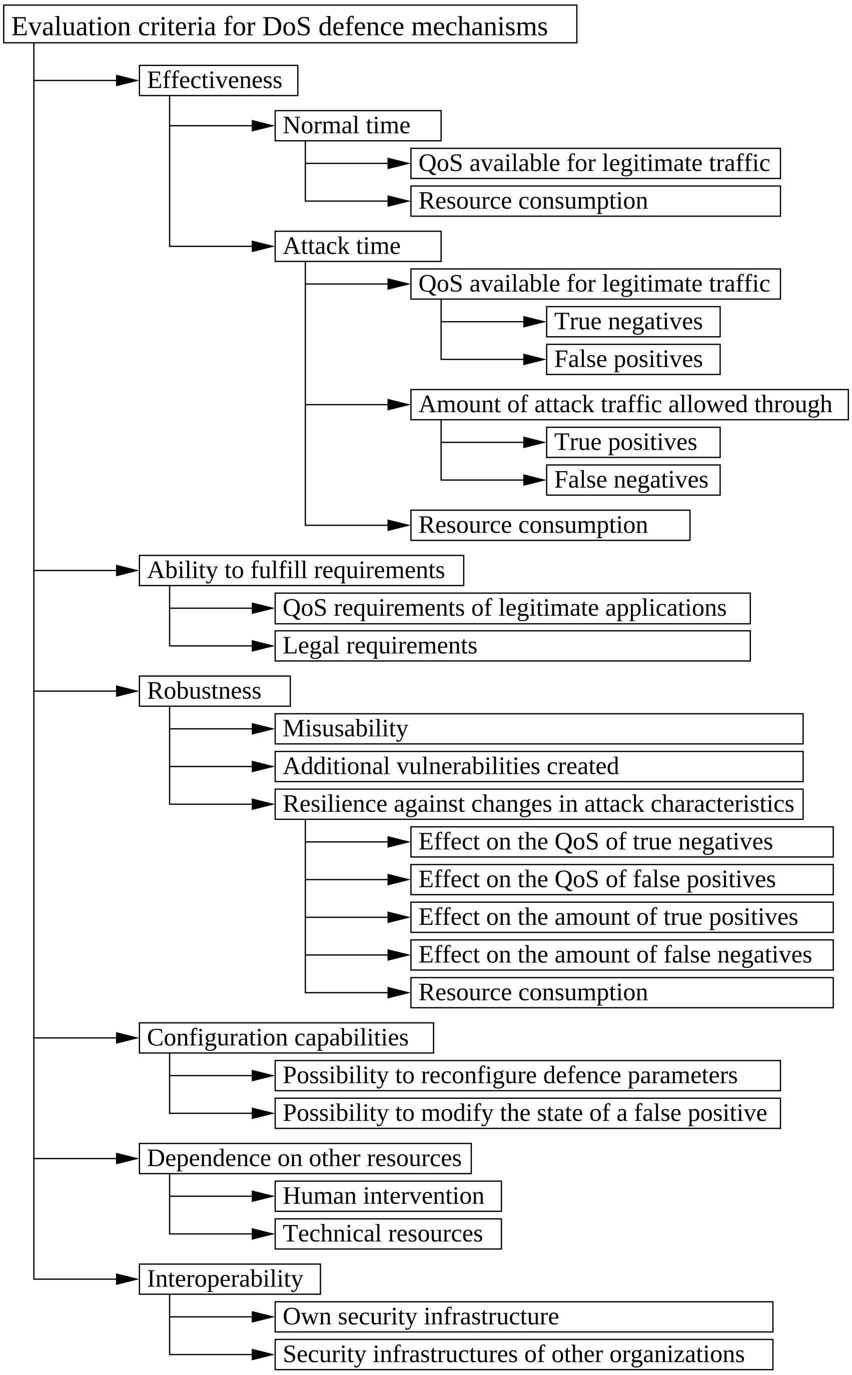
Evaluation criteria for DoS defence mechanisms
Effectiveness
Normal time
QoS available for legitimate traffic
Resource consumption
Attack time
QoS available for legitimate traffic
True negatives
False positives
Amount of attack traffic allowed through
True positives
False negatives
Resource consumption
Ability to fulfill requirements
QoS requirements of legitimate applications
Legal requirements
Robustness
Misusability
Additional vulnerabilities created
Resilience against changes in attack characteristics
Effect on the QoS of true negatives
Effect on the QoS of false positives
Effect on the amount of true positives
Effect on the amount of false negatives
Resource consumption
Configuration capabilities
Possibility to reconfigure defence parameters
Possibility to modify the state of a false positive
Dependence on other resources
Human intervention
Technical resources
Interoperability
Own security infrastructure
Security infrastructures of other organizations

Fig. 1. A taxonomy of criteria for evaluating defence mechanisms against flooding DoS attacks.

Quality of service (QoS) Available for Legitimate Traffic. A preventive defence mechanism can have an effect on the QoS available for legitimate traffic even if there is no attack due to the overhead from running the defence. The term QoS is understood here in a technical way (intrinsic QoS in [11]). QoS experienced by legitimate flows can be deteriorated if a mechanism introduces additional network security devices increasing the transmission latency, new security procedures requiring additional steps for accessing services, new security information increasing the length of legitimate packets etc. An opposite example providing better QoS during normal time is a preventive defence based on resource multiplication [10] which will enhance the QoS experienced by legitimate flows.

QoS degradation during normal time has been analyzed, for example, in [12] which describes a proactive secure overlay service structure (WebSOS) to prevent DDoS attacks against web servers. On average the basic WebSOS increases the end-to-end communication latency between a browser and a web server by a factor of 7 when compared to normal routing used in the current Internet. In the worst-case this latency is increased by a factor of 11 when compared to normal routing. In addition to this permanent increase in latency there is an additional security procedure in WebSOS when initiating an access to a web server. The goal of this initial security procedure is to verify that a human is trying to use the web server. This Completely Automated Public Turing test to tell Computers and Humans Apart (CAPTCHA) increases the time before a user is granted access to a web server.

Resource Consumption. Preventive defence mechanisms will consume resources, such as processing power, memory, and transmission capacity of a network.

For example, if IP packets are encrypted/authenticated, each packet will consume more processing power and memory in nodes initiating or terminating a secured path. Also, the increased packet size will consume more network bandwidth.

Attack Time. The primary goal for any flooding DoS defence mechanism is to limit the volume of malicious traffic during an attack. In addition to preventive mechanisms, also reactive defence mechanisms can be used, but they require a separate detection mechanism. Reactive mechanisms do not generally degrade QoS of legitimate flows during normal time, i.e. when there is no attack. According to the taxonomy of DDoS defences defined in [10] there are four different basic types of reactive defence mechanisms: agent identification (source traceback), rate-limiting, filtering (blocking), and reconfiguration.

QoS Available for Legitimate Traffic. Legitimate traffic is divided in two groups. True negatives are legitimate flows that are classified as legitimate. False positives are those legitimate flows that are classified accidentally as attack flows. Depending on the requirements of the applications used in a network, these two different traffic groups may have similar or different requirements for the QoS. In any case, these both types of legitimate traffic must be considered when evaluating the effectiveness of a defence mechanism.

The QoS available for true negatives is enhanced by preventing part or all of the detected attack traffic from entering a victim network or host. This results in more resources (e.g., network bandwidth or processing power at a server) available for true negatives.

In case of reactive defences, true negatives will experience a short period of time of low QoS at the beginning of an attack. There is an inherent delay between the time when an attack is detected and the corresponding reactive defence begins to mitigate the attack. For example, if the core nodes of the WebSOS overlay network [12] are attacked, the system will heal itself within 10 seconds. This reaction delay should be included in the evaluation of the available QoS during an attack.

The QoS available for false positives, on the other hand, is generally not very high, because these flows cannot be differentiated from attack flows. In other words, both false positives and true positives (attack traffic classified as malicious) are associated with the same QoS level.

If all legitimate traffic, including false positives, require a reasonable QoS even during a flooding DoS attack, both preventive defences (such as resource multiplication) and some of the reactive defences (such as rate-limiting [13] and reconfiguration) can be useful. For example, flooding DoS attacks have mostly failed against the root servers of the Domain Name System (DNS) due to required overprovisioning [14, 15]. Reactive defence mechanisms based on filtering are not suitable, if the QoS of false positives is important.

Amount of Attack Traffic Allowed Through. Attack traffic is divided in two groups. True positives are those attack flows classified as malicious. False negatives are those attack flows classified accidentally as legitimate flows. At least all false negatives are allowed to pass through, and possibly part of the true positives (such as in the case of rate-limiting).

The probabilities for a true positive and a false positive are inter-related [16]. If true positives must be detected with a high probability, then also the probability for a false positive grows. The same holds also between false negatives and true negatives, i.e. if the probability for a false negative must be lowered, then also the probability for a true negative will get lower. In practice this means that the less we allow undetected attack traffic to get in (false negatives), the less we allow legitimate traffic to get in (true negatives). It is not possible to adjust the attack detection so that exactly only legitimate traffic would get in. This fact has been recognized in existing DDoS toolkits which generate attack traffic looking very similar to legitimate traffic [5].

Due to these realities, both true positives and false negatives must be considered in the evaluation of a defence mechanism against flooding DoS attacks. Especially reactive defences do not mitigate the volume of false negatives.

Resource Consumption. Mitigating a detected flooding DoS attack is resource consuming. Limiting the volume of attack traffic may require, for example, several filters for classifying attack flows at routers. There are practical limits on the amount of different filters in routers. If attack traffic is highly variable, it may not be possible to install filters for all detected types of attack traffic due to resource limitations.

At least enough processing power and memory should be available for mitigating an attack at routers and other network security devices.

2.2 Ability to Fulfill Requirements

Different applications have different requirements for QoS. Legislation, standards, specifications, recommendations about best current practices, and other documents may also dictate requirements for the QoS of important applications. All these requirements should be fulfilled as well as possible, even during attack.

QoS Requirements of Legitimate Applications. Using a defence mechanism against a flooding DoS attack increases the level of QoS during attack. To see whether the available QoS is enough, one must compare this available QoS with the requirements of the most important application (or applications) used in a network. There should be a reasonable match between the available and required level of QoS.

Legal Requirements. Legislation or other official rules may require organizations to provide a reasonable resistance to known security vulnerabilities. For example, [17] describes recommended security services and procedures for Internet service providers. DoS attacks found frequently in real-life may have to considered as known security vulnerabilities. Real-life DoS attacks have been investigated, for example, in [1] and [18].

2.3 Robustness

Any defence mechanism should be robust for not opening any new possibilities for carrying out DoS attacks.

Misusability. It is possible to use some defence mechanisms as the ultimate tool for the DoS attack. Such a defence mechanism results in more damage than the original attack itself. Using intelligence in selecting the contents of DoS attack traffic, it may be possible to force a defence mechanism to fail in its most important task.

Additional Vulnerabilities Created. Defence mechanisms increase the complexity of a system and thus result in new vulnerabilities to be exploited by attackers. As attacks can exploit any weaknesses in protocols and services, any additional security protocol, security service, or network security device may provide a possible avenue for carrying out an attack.

Resilience against Changes in Attack Characteristics. Many attacks have varying attack characteristics. Source address validity (how source address spoofing is used), attack rate, possibility of characterization (how easy it is to identify attack packets), persistence of attack sources, and victim type (application, host, resource, network) can be modified on the fly [10].

A defence mechanism should be able to adapt to changes in attack properties. One should also evaluate how varying attack properties affect the QoS of legitimate traffic and the amount of attack traffic allowed to reach a victim.

Frequent changes in attack characteristics may cause excessive resource consumption, for example, by overloading a router when it receives frequent descriptions of attack traffic.

2.4 Configuration Capabilities

Any defence mechanism should incorporate reasonable reconfiguration capabilities. This is required when attack characteristics change, or when a critical false positive is identified.

2.5 Dependence on Other Resources

To operate effectively, a defence mechanism may require extensive human intervention and several other security devices, such as intrusion detection systems. All these dependencies affect the cost-effectiveness of attack detection and response [19].

Human Intervention. If a defence mechanism is dependent on human interaction, this will increase the delays in operating a mechanism. Autonomous defences should be preferred when primary applications cannot tolerate breaks in the availability of services.

Technical Resources. A defence mechanism is more prone to malfunction if it is dependent on availability of other security devices. For example, a reactive defence mechanism is useless without a correctly operating detection system.

A requirement for a large number of other security devices or for a widespread deployment of a distributed defence mechanism has implications on implementation issues. Incremental deployment is needed for these more complex defence systems.

2.6 Interoperability

Defence mechanisms are never separate entities in organizations. A prerequisite for any defence mechanism is that it must fit with an existing security infrastructure of an organization and be able to co-operate with other existing defence mechanisms, such as intrusion detection systems and security management tools. When global or distributed defence infrastructures are used, even higher demands for interoperability exist.

3 Related Work

Only few organized approaches for analyzing denial of service attacks and defence mechanisms have been published in addition to the taxonomy presented in [10].

A framework of criteria for evaluating proactive DoS defence mechanisms is presented in [6] which compares three existing preventive solutions according to the following requirements: incremental deployment, resistance to traffic analysis, resistance to compromised infrastructure routers, and resistance to DoS attack on the infrastructure. In that paper it is expected that attacks are distributed only in a limited fashion, i.e. attack traffic can be mostly distinguished from legitimate traffic. Reactive defences were not considered at all in the paper, and the included set of evaluation criteria omitted many issues included in the taxonomy presented in this paper.

A cost-based framework for analyzing the resistance of cryptographic protocols against DoS attacks is presented in [20]. It provides a formalized mechanism for comparing the costs of a DoS attack for both an attacker and a victim. The goal is to make the cost of carrying out a DoS attack as expensive as possible when compared to the costs of the victim who is required to process attack packets.

The number of attacking source hosts (DDoS agents) is important for initiating source traceback for locating the real source of attack traffic. If there is only one host transmitting attack traffic, source traceback can be a useful reactive defence mechanism against this kind of an attack. The more there are hosts sending attack traffic against a single victim, the more difficult it is to mitigate an attack by using source traceback. A framework for classifying DoS attacks as either single- or multi-source is presented in [18], and it is based on analyzing packet headers, the ramp-up behaviour of attack traffic intensity, and the spectral content analysis of the inter-message time.

4 Conclusion

Evaluation of defence mechanisms against flooding DoS attacks has often concentrated on easy or simple test scenarios, where the possibility to circumvent or defeat a defence mechanism has been either underestimated or completely forgotten. There are no existing papers providing any framework or taxonomy for this important subject.

This paper presented a taxonomy which classifies evaluation criteria for defence mechanisms against flooding DoS attacks. This taxonomy can be treated as an extension to the taxonomy of DDoS attack and DDoS defence mechanisms presented in [10].

The presented taxonomy emphasizes effectiveness when there is no attack, effectiveness during an attack, ability to fulfill requirements on application QoS, robustness against misuse, resilience against changes in attack characteristics, possibility for dynamic configuration especially for removing critical false positives, dependence on technical resources and human interaction, and interoperability with existing security infrastructures.

As evaluation of DoS defence mechanisms is fairly complicated, the presented taxonomy will provide a structured list of things to be considered during an evaluation process.

Acknowledgement

This work was partially funded by the Finnish Defence Forces.

References

1. Moore, D., Voelker, G.M., Savage, S.: Inferring Internet denial-of-service activity. In: Proceedings of the 10th USENIX Security Symposium, Washington, D.C. (2001)
2. Ptacek, T.H., Newsham, T.N.: Insertion, Evasion, and Denial of Service: Eluding Network Intrusion Detection. Secure Networks, Inc. (1998)
3. Durst, R., Champion, T., Witten, B., Miller, E., Spagnuolo, L.: Testing and evaluating computer intrusion detection systems. Communications of the ACM **42** (1999) 53–61
4. Lippmann, R.P., Fried, D.J., Graf, I., Haines, J.W., Kendall, K.R., McClung, D., Weber, D., Webster, S.E., Wyschogrod, D., Cunningham, R.K., Zissman, M.A.: Evaluating intrusion detection systems: The 1998 DARPA off-line intrusion detection evaluation. In: Proceedings of the DARPA Information Survivability Conference and Exposition. (2000)
5. Sterne, D., Djahandari, K., Wilson, B., Babson, B., Schnackenberg, D., Holliday, H., Reid, T.: Autonomic response to distributed Denial of Service attacks. In: Proceedings of Recent Advances in Intrusion Detection, 4th International Symposium, Davis, California (2001) 134–149
6. Mulligan, J.A.: A comparison framework for proactive Internet denial of service solutions. Technical report, Massachusetts Institute of Technology (2005)
7. Mahajan, R., Bellovin, S.M., Floyd, S., Ioannidis, J., Paxson, V., Shenker, S.: Controlling high bandwidth aggregates in the network. ACM SIGCOMM Computer Communication Review **32** (2002) 62–73
8. Adkins, D., Lakshminarayanan, K., Perrig, A., Stoica, I.: Towards a more functional and secure network infrastructure. Technical Report UCB/CSD-03-1242, University of California, Berkeley (2003)
9. Schuba, C.L., Krsul, I.V., Kuhn, M.G., Spafford, E.H., Sundaram, A., Zamboni, D.: Analysis of a Denial of Service attack on TCP. In: Proceedings of the IEEE Symposium on Security and Privacy, Oakland, California (1997) 208–223
10. Mirkovic, J., Reiher, P.: A taxonomy of DDoS attack and DDoS defense mechanisms. ACM SIGCOMM Computer Communication Review **34** (2004) 39–53
11. Gozdecki, J., Jajszczyk, A., Stankiewicz, R.: Quality of service terminology in IP networks. IEEE Commun. Mag. **41** (2003) 153–159
12. Morein, W.G., Stavrou, A., Cook, D.L., Keromytis, A.D., Misra, V., Rubenstein, D.: Using graphic Turing tests to counter automated DDoS attacks against web servers. In: Proceedings of the ACM conference on computer and communications security, Washington, DC, USA (2003) 8–19

13. Mölsä, J.V.E.: Effectiveness of rate-limiting in mitigating flooding DoS attacks. In Hamza, M.H., ed.: Proceedings of the Third IASTED International Conference on Communications, Internet, and Information Technology at St. Thomas, US Virgin Islands, Anaheim, California, USA, ACTA Press (2004) 155–160
14. Vixie, P., Sneeringer, G., Schleifer, M.: Events of 21-Oct-2002. Technical report, ISC/UMD/Cogent (2002)
15. Bush, R., Karrenberg, D., Kosters, M., Plzak, R.: Root name server operational requirements. Request for Comments RFC 2870, Internet Engineering Task Force (2000)
16. Wickens, C.D., Hollands, J.G.: Engineering Psychology and Human Performance. 3 edn. Prentice Hall, Upper Saddle River, New Jersey, USA (2000)
17. Killalea, T.: Recommended Internet Service Provider Security Services and Procedures. RFC 3013 (2000)
18. Hussain, A., Heidemann, J., Papadopoulos, C.: A framework for classifying denial of service attacks. In: Proceedings of ACM SIGCOMM, Karlsruhe, Germany (2003)
19. Lee, W., Fan, W., Miller, M., Stolfo, S.J., Zadok, E.: Toward cost-sensitive modeling for intrusion detection and response. Journal of Computer Security **10** (2002)
20. Meadows, C.: A cost-based framework for analysis of denial of service in networks. Journal of Computer Security **9** (2001) 143–164