

Characterizing the IPv6 Security Landscape by Large-Scale Measurements

Luuk Hendriks^(✉), Anna Sperotto, and Aiko Pras

Design and Analysis of Communication Systems (DACS),
University of Twente, Enschede, The Netherlands
{luuk.hendriks,a.sperotto,a.pras}@utwente.nl

Abstract. Networks are transitioning from IP version 4 to the new version 6. Fundamental differences in the protocols introduce new security challenges with varying levels of evidence. As enabling IPv6 in an existing network is often already challenging on the functional level, security aspects are overlooked, even those that are emphasized in literature. Reusing existing security solutions for IPv4 might seem easy and cost-effective, but is based on the unproven assumption that IPv6 attack traffic features the same characteristics. By performing network measurements and analyzing IPv6 attacks on the network-level, we determine the current state of security in the IPv6 domain. With the inevitable switch to the new protocol version, assessing the applicability of existing security approaches and determining the requirements for new solutions becomes a necessity.

1 Introduction

In this paper we aim to describe our plans on researching the status of IPv6 security by performing measurements. After introducing the subject and the motivation for the work, the goal and research questions are stated and explained in Section 2. The approaches planned in order to answer the questions are explained in Section 3, followed by brief final considerations in Section 4.

The new version 6 of the Internet Protocol (IPv6) is gradually being adopted by the Internet. The successor of the currently most-used IP version 4 (IPv4) is often seen as an expansion in terms of address space, but that is just a one of many changes. Once designed with security in mind, after 20 years of developments and implementations, some question whether IPv6 is indeed more secure than its predecessor. As IPv6 is gaining traction, the amount of malicious traffic transferred over it increases [12]. Besides attacks that are based on aspects of the new protocol, we can expect traditional attacks from IPv4 occurring over IPv6 as well, as many types of these threats abuse features or phenomena on higher layers (e.g., the transport or application layer). In addition to that, the technologies designed to aid in the transitional phase from IPv4 to IPv6 (e.g., tunneling techniques *6to4* and *TEREDO*, among others) come with their own issues [3,6]. Circumvention of firewalls [5] in certain scenarios is a severe example of this. Lastly, another likely source of security issues are end-users not aware

of IPv6 connectivity on their systems and network [10], making them prone to attacks via unexpected ways.

The new version of the IP protocol itself, as well as its supporting protocols (e.g., the Neighbour Discovery Protocol), have been subjected to research and insecurities have been pointed out in the literature [1,2,6]. These works have described issues and weaknesses, but did not focus on actual occurrences of attempts to exploit them in the Internet. Some research has been done in that area, mainly by measuring so-called *darknets*. A darknet is address space that contains no actual services, but is advertised and routed. Any traffic arriving in such a space can be considered malicious [4]. Although a darknet simplifies the classification of traffic as being either benign or malicious, the fact it contains no actual services is likely to lower the interest of those with bad intents. These studies [4,9,12] do however show an increase in the amount of observed traffic throughout the years.

With this work we intend to research the actual state of security in the IPv6 Internet, and how it can be improved. Besides the shown increase of malicious traffic, our study is also motivated by the imbalance between availability of tools (e.g., the THC hacker toolkit [8] for IPv6, first released in 2005) and countermeasures for weaknesses these tools exploit: the first RFC describing RA-Guard [11] on how to mitigate Router Advertisement-based attacks is dated 2011, and the first version of a BCP concerning rogue DHCPv6 servers [7] was presented in 2012. A comprehensive study of Ullrich et al. [14] divides the known IPv6 problems in 36 security vulnerabilities and 14 privacy issues. With that, 44 possible countermeasures are listed. The large number of countermeasures shows that most issues are technically surmountable, but the possibility of configuration errors or omissions is significant. Furthermore they point out that several concepts within the IPv6 domain are being deprecated now or in the foreseeable future, while many implementations are already running in production. Updates to those implementations are not a given, so even with deprecation of security-impairing aspects in mind, there is no guarantee on how fast the improvements will actually be functional. This further emphasizes the need for real-world measurement-based studies to complement theoretical conclusions.

2 Goal and Research Questions

The goal of the research is *to characterize the IPv6 landscape from a security perspective*, as motivated in the previous section.

The following research questions will be answered:

1. *What types of attacks can be observed over IPv6 on the network-level, and how do they relate to attacks over IPv4?*
2. *Which fraction of the Internet is susceptibility to IPv6-based attacks?*
3. *How can detection of IPv6-based threats be performed?*

The answer to the first research question will give an overview of attacks that are occurring in the Internet, and the attacks that are possible based on attack

tools openly available. The gained knowledge is used in the approach to answer question 2. By answering question 2, the level of security in the IPv6 Internet is determined, and the likelihood and severeness of the attacks found in question 1 are assessed. Lastly, by answering question 3, ways of detecting threats in the IPv6 domain are researched, in order to improve the overall level of security of the IPv6 Internet.

3 Approach

Our approach is mainly based on performing measurements. Answering the first question will comprise of two separate, simultaneous forms of measurements. By doing large-scale passive measurements, attacks over IPv6 are observed. The vast address space reduces the probability of malicious traffic entering the monitored networks. By focussing on address space that has been allocated for several years (e.g., the SURFnet network) the chances of obtaining useful data increase. To further increase possibilities of acquiring data, a reactive system might be used, i.e., a system acting on incoming IPv6 traffic regardless of the destination. A similar approach in [12], based on dynamic instantiation in honeynets, showed promising results. Traffic will be collected in forms of both packet captures and flow records, increasing the flexibility in the analysis. As this process involves a certain waiting time, and the possibility of not resulting in sufficient data to analyze, lab experiments will be conducted parallel to the passive measurements. By collecting and analyzing attack tools openly available, signatures can be created to aid in both the analysis of the data obtained via the passive measurements, as well as answering research questions 2 and 3.

The goal of question 2 is to determine to what extent the observed attacks can be effective in reality, or, how well protected the Internet is from them. The question is answered by performing active measurements, where the exact form of measurements is determined by the outcome of the first research question. If a large amount of attacks is observed via the passive measurements, these attacks will be characterized. If on the other hand the results from the passive measurements are not substantial, characteristics from the collected and analyzed attack tools are used. By smartly [14] scanning parts of the IPv6 address space, systems connected via IPv6 are found. The share of vulnerable hosts is determined based on the attacks' characteristics. For example, we might find attacks based on abusing services on the application layer. Operators might use specific addressing schemes within their network, where part of the address represents the service running on that system [13]. Examples are DNS services being deployed on 2001:db8::100:53, or HTTP on 2001:db8::100:80, thus the last field representing the transport layer port. Performing smart scans possibly provides insights in this scenario. In the case that attacks are tailored towards exploiting vulnerabilities in implementations of IPv6 network stacks, vulnerability is related to (the version of) the operation system on possible target hosts. Fingerprinting of operating systems will result in more useful information to determine the share of hosts possibly subjected to this kind of attacks.

The third question finally, assesses the possibilities of detecting the attacks. In this study, the applicability of approaches and solutions from the IPv4 domain is researched. The goal of this question is to research whether new detection technologies are required to ensure security within the IPv6-adopting Internet. For attacks newly introduced with IPv6, we intend to design adequate detection algorithms. The form of input for these algorithms depends on the form of the attacks. If flow-level data is insufficient to perform accurate detection, e.g. because specific headers or payload are key in detecting the attack, packet-level input will be used. The algorithms will be tested by means of implementing a prototype, to be validated with a ground truth. The ground truth again depends on the form of the attack: if attacks have an analogue in the IPv4 domain, existing detection algorithms and tools can provide the ground truth. Otherwise, analysis of logfiles on attacked or compromised end-hosts will provide insights to validate the results of detection on the network-level. If validity is proven, the algorithms might be tested for scalability and performance in larger set-ups, e.g., by deployment on National Research and Educational Network (NREN) links, or at Computer Security Incident Response Teams (CSIRTs).

4 Final Considerations

Within our research, measurements will be conducted in various forms. Performing active measurements will likely raise ethical or perhaps legal questions. Extra consideration or adjustment of plans might be needed if these questions create legitimate limitations. The main research goal as described is to be achieved within the duration of four years, as parts of Ph.D. research. The research is partly funded by the European FLAMINGO¹ project (ICT-FP7 318488) and SURFnet².

References

1. Beck, F., Cholez, T., Festor, O., Chrisment, I.: Monitoring the neighbor discovery protocol. In: The Second International Workshop on IPv6 Today-Technology and Deployment-IPv6TD 2007 (2007)
2. Caicedo, C.E., Joshi, J.B., Tuladhar, S.R.: IPv6 security challenges. *Computer* (2), 36–42 (2009)
3. Elich, M., Velan, P., Jirsik, T., Celeda, P.: An investigation into teredo and 6to4 transition mechanisms: Traffic analysis. In: 2013 IEEE 38th Conference on Local Computer Networks Workshops (LCN Workshops), pp. 1018–1024 (October 2013)
4. Ford, M., Stevens, J., Ronan, J.: Initial Results from an IPv6 Darknet. In: International Conference on Internet Surveillance and Protection, ICISP 2006, pp. 13–17 (2006)
5. Giobbi, R.: Bypassing Firewalls with IPv6 Tunnels (2009), <http://www.cert.org/blogs/certcc/post.cfm?EntryID=37> (accessed March 2015)

¹ <http://www.fp7-flamingo.eu/>

² <http://surfnet.nl>

6. Gont, F.: Security implications of IPv6 on IPv4 networks, RFC 7123, Internet Engineering Task Force (2014)
7. Gont, F., et al.: DHCPv6-Shield: Protecting Against Rogue DHCPv6 Servers (2012)
8. Heuse, M.: THC IPv6 attack tool kit, <https://www.thc.org/thc-ipv6/> (accessed March 2015)
9. Huston, G.: Background Radiation in IPv6. The ISP Column, APNIC (2010)
10. Krishnan, S., Hoagland, J., Thaler, D.: Security Concerns with IP Tunneling, RFC 6169, Internet Engineering Task Force (2011)
11. Levy-Abegnoli, E., Van de Velde, G., Popoviciu, C., Mohacsi, J.: IPv6 Router Advertisement Guard. Tech. rep., RFC 6105, Internet Engineering Task Force (2011)
12. Schindler, S., Schnor, B., Kiertscher, S., Scheffler, T., Zack, E.: IPv6 network attack detection with HoneydV6. In: Obaidat, M.S., Filipe, J. (eds.) ICETE 2013. CCIS, vol. 456, pp. 252–269. Springer, Heidelberg (2014)
13. SURFnet: Preparing an IPv6 addressing plan (2013), <https://www.surf.nl/en/knowledge-and-innovation/knowledge-base/2013/white-paper-preparing-an-ipv6-address-plan.html> (accessed March 2015)
14. Ullrich, J., Krombholz, K., Hobel, H., Dabrowski, A., Weippl, E.: IPv6 security: attacks and countermeasures in a nutshell. In: Proceedings of the 8th USENIX Conference on Offensive Technologies, pp. 5–16. USENIX Association (2014)