Network Fingerprinting: Routers under Attack

Emeline Marechal
Montefiore Institute
Université de Liège
Liège, Belgium
emeline.marechal@uliege.be

Benoit Donnet
Montefiore Institute
Université de Liège
Liège, Belgium
benoit.donnet@uliege.be

***Abstract*—Nowadays, simple tools such as `traceroute` can be used by attackers to acquire topology knowledge remotely. Worse still, attackers can use a lightweight fingerprinting technique, based on `traceroute` and `ping`, to retrieve the routers brand, and use that knowledge to launch targeted attacks.**

In this paper, we show that the hardware ecosystem of network operators can greatly vary from one to another, with all potential security implications it brings. Indeed, depending on the autonomous system (AS), not all brands play the same role in terms of network connectivity. An attacker could find an interest in targeting a specific hardware vendor in a particular AS, if known defects are present in this hardware, and if the AS relies heavily on it for forwarding its traffic.

***Index Terms*—network fingerprinting, traceroute, ping, attack, connectivity**

1. Introduction

Fingerprinting [1, 2] is defined as the process of splitting network equipment into several disjoint classes. This is achieved by analyzing messages sent by equipment and their behavior, usually in response to some form of active probing. As such, fingerprinting is an expensive process as it could require many probes to be sent, and thus time consuming [1]. In addition, too many probes towards a network node or a subnet could easily appear as an attack and, consequently, be filtered by the target. Recently, Vanaubel et al. [3] have proposed a lightweight network fingerprinting technique that is based on inferring the initial TTL values in packets sent by routers. Vanaubel et al. have shown that it is enough to obtain the initial TTL of two ICMP messages (i.e., `time-exceeded` and `echo-reply`– the so-called *router signature*) to guess the router hardware vendor.

Providing such a fingerprinting is useful for several applications and studies. For instance, it has been used in alias resolution (i.e., the process of aggregating IP interfaces of a router into a single identifier) [4, 5]. It has also been used for revealing the content of MPLS tunnels hidden to `traceroute`, as some MPLS behaviors depend on the hardware vendor (mainly Cisco vs. Juniper) [6].

In this paper, we rely on routers signature to answer two research questions:

1) beyond the classic hardware vendor market share (i.e., proportion of Cisco vs. proportion of Juniper, etc.), we ask ourselves *what is the hardware ecosystem within Internet and operators?* In particular, we are interested in describing where are located the various hardware and the potential role they could play in the topology. Our findings in this paper show that, if Cisco largely dominates the overall market, this is not reflected when looking on a per autonomous system (AS) basis, where the distribution is more blurred.
2) if knowing the hardware ecosystem of an AS is straightforward and not that intrusive, we ask ourselves *what could happen if an attacker can easily identify router brands and target specific vendor with (known) security breaches?* This question is motivated by the recent five vulnerabilities found in various Cisco devices (four of them leading to remote code execution vulnerabilities, and one to a Denial of Service vulnerability) [7] and by the discovery of Bleichenbacher oracles [8] (i.e., an adaptative chosen ciphertext attack against some protocols based on RSA) in the IKEv1 implementations of four large network equipment manufacturers (Cisco, Huawei, Clavister, and ZyXEL) [9]. Those attacks are not limited to a few scattered devices, but could affect many different hardware models for each manufacturer, as the vulnerabilities are found in software common to many different products. If an attacker is able to easily identify unsecured equipment within an AS, they could easily target it and possibly disrupt the AS connectivity. Generally speaking, the attack could affect the AS connectivity or use the AS equipment for performing a larger-scale attack (e.g., DDoS). In this paper, we focus on connectivity loss, as result of the attack. In particular, we show that it is enough for an attacker to target an AS and a few devices (of a given brand) to affect its connectivity.

The remainder of this paper is organized as follows: Sec. 2 provides the required background for this paper; Sec. 3 discusses ASes hardware ecosystem (Research Question 1); Sec. 4 illustrates the impact of routers failures due to attacks (Research Question 2) and provides mitigation solutions to operators and hardware vendors; finally, Sec. 5 concludes this paper by summarizing its main achievements.

Router Signature	Router Brand and OS
$< 255, 255 >$	Cisco (IOS, IOS XR)
$< 255, 64 >$	Juniper (Junos)
$< 128, 128 >$	Juniper (JunosE)
$< 64, 64 >$	Brocade, Alcatel, Linux (BAL)

TABLE 1. SUMMARY OF MAIN ROUTER SIGNATURE, THE FIRST INITIAL TTL OF THE PAIR CORRESPONDS TO ICMP `TIME-EXCEEDED`, WHILE THE SECOND IS FOR ICMP `ECHO-REPLY`.

	IP	Router
coherent	84,7%	93,8%
weakly incoherent	15,1%	5,5%
incoherent	0.2%	0,7%

TABLE 2. SIGNATURE COHERENCE FOR BOTH IP ADDRESSES AND ROUTERS (AFTER ALIAS RESOLUTION).

2. Background

The IP packet header contains a *time-to-live* (TTL) field used to avoid packets looping forever when a routing loop occurs. This 8-bit field is set by the originating host/router to an *initial value* (iTTL) that is usually and nearly always a power of 2 in the list 32 (or 30), 64, 128, and 255. RFC1700 [10] recommends to use 64 as iTTL value but in practice, this is not followed by most router manufacturers, each one having its own policy that may also depend on the protocol used [3].

Based on that, Vanaubel et al. [3] have proposed a *router signature* made of a n-tuple of n iTTLs, those iTTLs being retrieved from different ICMP messages received from routers.[1]

Vanaubel et al. have demonstrated that it is sufficient to consider the iTTL of two different messages (i.e., $n = 2$) to discriminate hardware vendors basic pair-signature: a `time-exceeded` message elicited by a `traceroute` probe, and an `echo-reply` message obtained from an `echo-request` probe. The advantage, here, is that router signatures can be easily retrieved with basic `traceroute` and `ping` exploration.

Table 1 summarizes the main router signatures, with associated router brands and router OSes.

3. Hardware Ecosystem

In this section, we discuss our first research question: *what is the hardware ecosystem within Internet and operators?* We first present how data has been collected and pre-processed (Sec. 3.1). Next, we check the signature coherence for both IP interfaces and routers (Sec. 3.2). We finally provide some insight into hardware distribution (Sec. 3.3) and discuss the limits of our approach (Sec. 3.4).

3.1. Data Collection

We collected data using `TNT` [11, 12], a Paris-traceroute [13] extension that is able to reveal the content of MPLS tunnels hidden to `traceroute` exploration [6], revealing so more links and IP interfaces than standard `traceroute` exploration. `TNT` comes with the advantage that it also automatically collects the signature for each collected IP interface.

We deployed `TNT` on the Archipelago infrastructure [14] between November 1st and 13th, 2019 over 28 vantage points, scattered all around the world: Europe (9), North America (11), South America (1), Asia (4), and Australia (3). The overall set of destinations, over 10 million IP addresses, is inherited from the Archipelago dataset and spread over the 28 vantage points to speed up the probing process. `TNT` data is freely available on CAIDA's website.

A total of 1,280,291 distinct unique IP addresses (excluding `traceroute` targets) have been collected, with 45,763 being non-publicly routable addresses (and thus excluded from our dataset).

We then used MIDAR [15], a tool based on similarities in the IP-ID field, to perform alias resolution on our set of addresses. Alias resolution [16] is the process of identifying IP addresses that belong to the same router, and can thus yield a router-level topology from the address-level topology that `traceroute` gives. This more concrete topology can then be used, among other purposes, to study more precisely the physical infrastructure of routers, their diversity, and the resiliency of the infrastructure. Out of the 1.2 million addresses discovered by `TNT`, MIDAR found 65,778 routers involving 221,464 addresses. Note that, despite considerable progress in this domain, alias resolution remains an imperfect process. We discuss in Sec. 3.4 the limits of this approach, and why it is relevant to employ it despite its shortcomings.

Finally, from the router dataset obtained with MIDAR, we applied `bdrmapIT` [17], a tool for annotating routers with AS ownership. The objective here is to delimit as accurately as possible ASes in order to better study their hardware infrastructure. Studying the Internet at the scale of ASes, rather than at a global scale, is more meaningful because each AS is an independent network, operated by different people with different policies and technologies. This scale provides thus a more refined vision of the hardware distribution, with all potential implications for network resiliency and security. Moreover, focusing on individual ASes when discussing routers under attack is more realistic than envisioning a world-wide attack.

3.2. Address and Router Signature Coherence

In order to evaluate the coherence of the collected fingerprints, we first check if the same data collected multiple times by `TNT` always corresponds to the same signature. Following Vanaubel et al. [3] vocabulary, IP addresses can either be (i) *coherent* (i.e., the same signature is always observed for that address), (ii) *weakly incoherent* (i.e., two signatures are observed for a given IP address, one being an incomplete version of the other – e.g., $< x, y >$ and $< x, * >$), and, (iii) *incoherent* (i.e., several different signatures are observed for a given address). We extend this vocabulary to routers, where the coherence of a router is determined by the signature of each of its addresses. Table 2 reports the results for both IP addresses and routers.

We observe that the majority of IP addresses signatures are coherent, which is the perfect case for us, as there is no

1. To estimate the iTTL forged by the router, it is enough to find the smallest number in $32, 64, 128, 255$ that is larger than the received value in the TTL field of the IP packet encompassing the ICMP message.

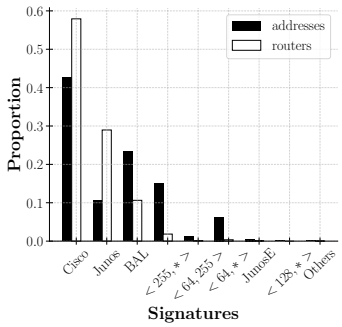

Figure 1. Big picture of hardware distribution (complete dataset).

ambiguity in the fingerprint. Weakly incoherent signatures are slightly significant, probably due to overloading, rate limiting, or filtering on routers, which prevent the device to answer to one of the two probes. As such, we can simply consider the signature of these addresses to be the complete one of the two. Finally, incoherent IP signatures are very infrequent. Only 0.2% of the addresses cannot be classified into a router brand, and have therefore been removed from our data.

With respect to routers, an even greater majority of routers are coherent, meaning that all interfaces of the router show the same signature. A small portion of routers are weakly incoherent, and, as for addresses, we consider the signature of a router to be the complete one among its interfaces. Finally, we cannot conclude anything regarding the brand of a router for only 0.7% of the routers. These fingerprinting results are consistent with the alias resolution process. Indeed, it is expected, by definition, that addresses showing different signatures cannot share a router, which is confirmed in our results.

3.3. Hardware Distribution

The big picture. Leveraging the fingerprinting method, we can have a look at the hardware distribution in the network. Fig. 1 illustrates the global signature distribution for both addresses and routers in the Internet.

Regarding addresses, we notice that Cisco signatures are largely dominant, with more than 40% of addresses in that class. The second most important class is the *BAL* class (i.e., Brocade, Alcatel, and Linux machines) with around 25% of addresses. After that, with a share of 15% of the addresses, comes the signature $< 255, * >$. This signature is an incomplete one, meaning that the device answered to the `traceroute` probe, but not to the `ping` one. This class is probably made of addresses that actually belong either to $< 255, 255 >$ (Cisco) or $< 255, 64 >$ (Junos) classes. Unfortunately, we are unable to discern between them, because those routers did not answer to the second probe. Finally, the fourth most important class is the Junos one, with approximately 11% of the market share. The remaining signatures (including JunosE) are quite rare.

Looking at routers now, we can observe that the signature distribution is extremely different. Cisco routers are still largely dominating the other brands, with almost 60% of the devices. Next comes the Junos class, followed by BAL, followed by the $< 255, * >$ incomplete signature, which has become almost non-existent. At first sight, we may think that the router topology, derived from the address topology with alias resolution, provides the actual hardware distribution in the Internet and corrects previous believes about the share of the market, when looking only at addresses. However, we found that the alias resolution introduces a bias in our data, as discussed in Sec. 3.4, and that these results must be considered with hindsight.

AS	# IP	# Router	# Traces
$\mathcal{A}$	2,322	629	901,795
$\mathcal{B}$	447	104	339,184
$\mathcal{C}$	1,425	363	417,391
$\mathcal{D}$	1,049	236	2,177,628

TABLE 3. BASIC DATA ABOUT ASES OF INTEREST. AS NUMBERS HAVE BEEN ANONYMIZED.

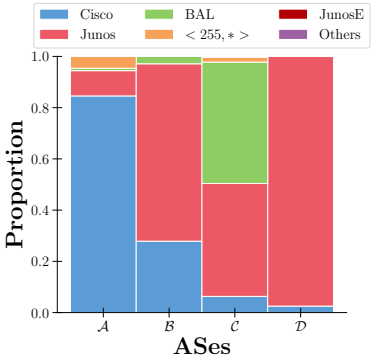

Figure 2. Hardware distribution, per AS.

Large ASes. For the remainder of this paper, we restrict ourselves to four ASes, the largest in our dataset, in order to study the hardware distribution at a finer granularity. To do so, as mentioned in Sec. 3.1, we used `bdrmapIT` for annotating routers with their AS number. For security reasons (see Sec. 4), we have anonymized AS numbers. Table 3 provides high level statistics on those ASes: the number of routers, as well as the number of IP addresses involved in those routers. The column labeled "# Traces" provides the number of `TNT` traces crossing each AS.

If the global hardware distribution (see Fig. 1) stated that Cisco largely dominates the market, the situation differs within our four ASes of interest, as illustrated by Fig. 2. Indeed, only AS $\mathcal{A}$ is largely dominated by Cisco, while AS $\mathcal{D}$ is relying nearly only on Junos. AS $\mathcal{C}$ uses a nearly equal mix of BAL and Junos. Finally, AS $\mathcal{B}$ deploys essentially Junos with Cisco. It is now obvious that different operators can have very different hardware infrastructure, with all potential implications for network resiliency and security.

Fig. 3 provides a deeper view of the hardware distribution in our ASes of interest by splitting the hardware according to its role (core vs. edge) in the AS topology. We can distinguish between core and edge devices thanks to `bdrmapIT` annotations that also provide the location of a device in the AS. Regarding AS $\mathcal{A}$ and AS $\mathcal{B}$, both have very similar distributions in the edge and in the core. The operators do not seem to privilege a brand or another for a particular usage in their network. AS $\mathcal{C}$ presents a slight dissimilarity in its distributions: core devices are composed at 60% of BAL routers and 30% of Junos routers, while edge devices present an equal mix of BAL

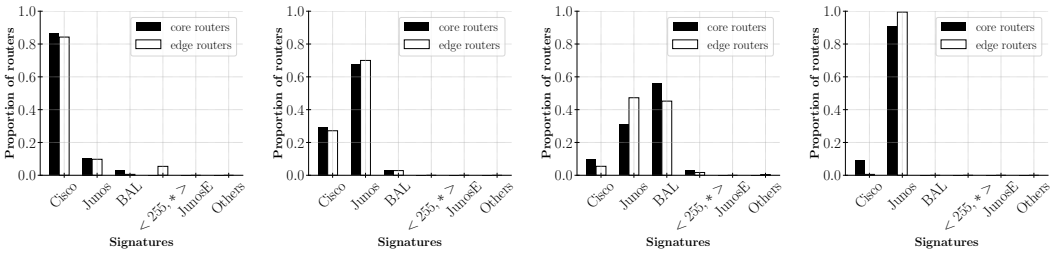

(a) AS $\mathcal{A}$ (b) AS $\mathcal{B}$ (c) AS $\mathcal{C}$ (d) AS $\mathcal{D}$

Figure 3. Hardware distribution, with respect to core and edge routers.

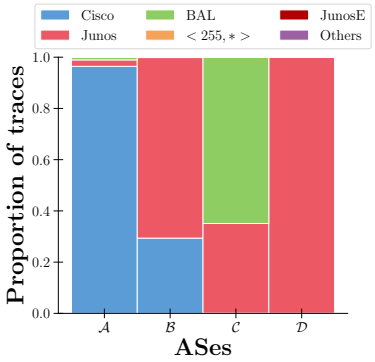

Figure 4. Hardware *popularity*, per AS.

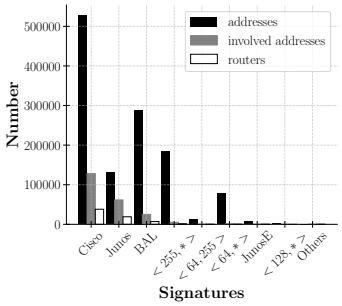

Figure 5. MIDAR performance depending on the hardware.

and Junos. Finally, AS $\mathcal{D}$ uses only Junos routers for their edge devices. The only Cisco routers in AS $\mathcal{D}$ are seen in the core of the network.

Finally, anticipating on Sec. 4, we present for each AS the hardware *popularity* in Fig. 4. Similarly to Sanchez et al. [18], we measure the hardware *popularity* as the proportion of TNT traces crossing each hardware brand. While the hardware distribution is already a first indicator of the topological importance of a brand in terms of connectivity, the *popularity* of a brand reflects more accurately the notion of topological importance than the distribution does. Indeed, and this is particularly true for Internet networks, not all nodes play the same role in terms of connectivity [19, 20], and some are more important than others. As such, it could turn out that, although a brand is largely represented in the network, the role it plays in terms of connectivity and in terms of traffic volume it carries is not as important.

However, we find that the hardware *popularity* is actually quite close to the hardware distribution. For each AS, the dominating brand in terms of hardware share is also the brand that plays a major role in terms of connectivity.

3.4. Limits

Despite considerable progress in this domain, alias resolution remains an imperfect process. All techniques present the risk of false positive, and all of them also have significant incompleteness. The inability to draw an accurate router-level map of the Internet limits what we can study, leading in this section on a lower bound of statistics discussed.

In our case, out of the 1.2 million addresses discovered by TNT, MIDAR found 65,778 routers involving 221,464 addresses. This represents only 18% of the initial dataset. The majority of aliases missed by MIDAR are due to routers not responding to probes, because of, e.g., overloading, rate-limiting, or filtering. The completeness of MIDAR is also limited by routers that do not use monotonic counters, or do not share a counter across interfaces, making them, by definition, undetectable by any IP-ID based technique [15].

Furthermore, we found that the completeness of the alias resolution process is not the same depending on the brand of routers. Indeed, as can be seen in Fig. 5, the proportion of addresses involved in routers varies drastically whether we are facing Cisco, Junos, or BAL routers. We see that Junos routers are more likely to be found by MIDAR, as the proportion of involved addresses reaches almost 50% (resp. 20% for Cisco and 10% for BAL). For the signature $< 255, * >$, the proportion of involved addresses is extremely small, which can also explain our router hardware distribution and why the $< 255, * >$ signature, previously quite important for address distribution, almost disappeared in the router distribution.

This observation corroborates a finding by Grailet and Donnet [4], in which they discovered that an IP interface with a healthy counter (i.e., that is susceptible to be found with an IP-ID resolution technique) is very likely to use the initial TTL value 64 (or, more rarely, 128) for the `echo-reply`. The value 64 for the `echo-reply` corresponds precisely to Junos routers, explaining that its proportion of involved addresses is the best among all brands.

This inequality of the alias resolution process regarding the brands introduces a bias in the router hardware distribution. However, despite its shortcomings, we believe the router level better captures hardware brand distribution than the IP level does. In fact, not using alias resolution and performing the analysis directly on IP addresses would lead to irrelevant results, and would also introduce a bias: based only on the number of interfaces we see, a

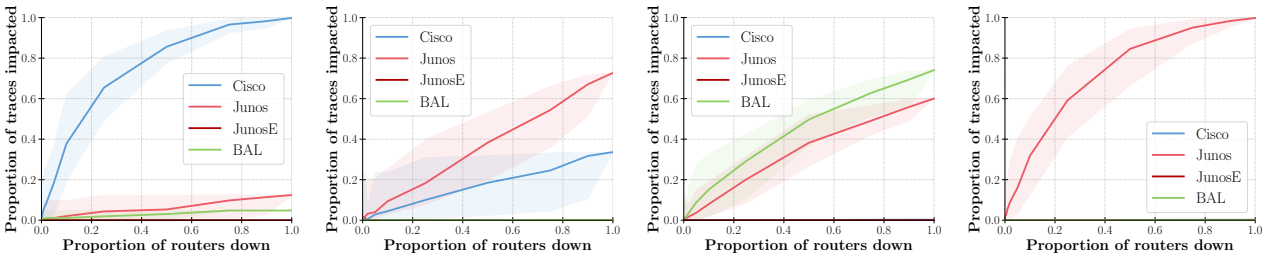

(a) AS $\mathcal{A}$ (b) AS $\mathcal{B}$ (c) AS $\mathcal{C}$ (d) AS $\mathcal{D}$

Figure 6. Impact of the attack on *traces* crossing each AS.

brand could be over-represented. Alias resolution allows to avoid this pitfall and aggregates multiple interfaces into fewer routers.

4. Routers Under Attack

Internet networks are known to belong to a certain class of networks called scale-free networks [19]. Those networks have the property of being extremely robust to random node removals, but also to be particularly vulnerable to targeted attacks on a few (highly connected) nodes playing a vital role in maintaining the network's connectivity [21]. This vulnerability is intrinsic to scale-free networks, in which the global connectivity of the whole network is ensured by a few nodes.

Therefore, if a certain router brand contributes heavily to network connectivity in a particular AS, and if known defects are present in this hardware, an attacker could cause great damage with little effort by targeting this router brand in the AS. Knowing that an attacker can fairly easily fingerprint a device (with only two simple probes), it becomes even easier for them.

Knowing this, and leveraging our new understanding of the ASes hardware infrastructure, we now look back at our second research question: *what could happen if an attacker can easily identify router brands and target specific vendor with (known) security breaches?* Indeed, we already saw in Sec. 3.3 that different operators have different hardware infrastructure, and thus most likely have different levels of vulnerability to a brand-targeted attack as well.

The hardware distribution is a first coarse indicator of the sensibility to brand-targeted attacks. However, it does not a priori reflect the topological importance of a particular brand in terms of connectivity. Therefore, to assess the vulnerability of an AS, we will review our hardware *popularity* approach, and study the number of *traces* (`traceroute` paths) impacted when a fraction of the nodes are removed from the network. We consider this metric to reflect the topological importance[2] and to estimate which brands of routers carry traffic to a significant fraction of the Internet. Indeed, Sanchez et al. [18] showed that the *popularity* of a link or a router is strongly related to the amount of traffic being carried.

2. A node that is topologically important will be sampled redundantly (multiple times) by `traceroute`. This is due to the `traceroute` exploration process that statistically focuses on topologically important nodes and links [20]. Therefore, the topological importance is reflected in the number of *traces* crossing a node.

To simulate an attack, we consider each AS separately and remove different proportions (0.001, 0.005, 0.01, 0.02, 0.05, 0.1, 0.25, 0.5, 0.75, 0.9, 1.0) of routers for four brands: Cisco, Junos, JunosE, and BAL. Each time, we count the total number of traces that go through those downed routers. For each percentage, we performed the simulation 30 times, averaged the results, and built confidence zone around the mean.

4.1. Results

Fig. 6 presents, for each AS, the number of *traces* impacted given a percentage of removed routers, for each brand. The number of *traces* impacted has been normalized by the total number of *traces* for that AS, in order to compare the four ASes together.

As expected given their hardware distribution and *popularity*, different ASes are sensitive to the removal of routers from different brands. Regarding AS $\mathcal{A}$, we see that it is enough to remove 20% of the Cisco nodes to impact nearly 60% of the *traces*, while other brands do not seem to have much of an impact. In this case, Cisco plays a major role in network connectivity. For AS $\mathcal{B}$ and AS $\mathcal{D}$, the removal of Junos is more harmful to network connectivity than for the other brands. AS $\mathcal{D}$ is also more vulnerable: for 20% of removed Junos nodes, more than 40% of *traces* are impacted, while this number barely reaches 15% for AS $\mathcal{B}$. Finally, for AS $\mathcal{C}$, we see that both attacks on BAL or on Junos routers have a significant impact on connectivity.

In the light of those results, we can definitely conclude that not all router brands contribute equally to network connectivity, and that some of them are topologically more important, depending on the AS. Newly discovered vulnerabilities could be exploited by those seeking to damage networks. Even though an attack targeting a defect on a particular brand could be launched blindly, without fingerprinting devices beforehand, if the attack requires heavy resources from the attacker, they can find a benefit in fingerprinting to focus the scope of their attack on carefully selected nodes.

4.2. Risk Mitigation

Network operators can protect themselves from fingerprinting, and all its potential security implications, in different manners. The first one is extremely simple: use a standard initial TTL (*iTTL*) in ICMP packets, as recommended by RFC1700 [10]. As the iTTL is not

configurable by operators, this suggestion is addressed to hardware vendors who must anonymize their routers by ensuring that each packet is forged with the same iTTL.

A more sophisticated technique is to obfuscate the topology to prevent attackers from discovering potential targets. An example would be to adapt, e.g., NetHide [22] for obfuscating the links but also anonymizing the routers.

Finally, as a last-resort solution, an operator could decide to completely turn off ICMP packets (or at least filter them at the edge, as done with IGMP [23]), effectively hiding the topology and the hardware infrastructure. However, this solution brings more drawbacks than advantages, and is most often seen as completely impracticable, as `traceroute` and `ping` are essential network debugging tools (error messages, connectivity checking, PMTU discovery, ...). Furthermore, the case is even worst for IPv6, where ICMPv6 cannot be treated as an auxiliary function, like its IPv4 counterpart, with packets that can be dropped in most cases without damaging the functionality of the network [24].

5. Conclusion

In this paper, we investigated two research questions: (i) what is the hardware ecosystem in the Internet and operators and (ii) what could happen if an attacker can easily identify router brands and target specific vendor with (known) security breaches?

For the first question, we showed that, if Cisco largely dominates the overall market, the hardware distribution appears more colorful when looking on a per AS basis. Different ASes can have very different hardware ecosystems, with all potential implications for network resiliency and security.

With respect to the second question, we demonstrated that not all brands contribute equally to network connectivity, depending on the AS. An attacker seeking to cause a lot of damage, with the least amount of effort, could target a specific brand that plays a vital role in network connectivity, and do so very easily given the simplicity of our fingerprinting technique.

Ethical Considerations

We are aware that someone with bad intentions (hacker, unscrupulous competitor, ...) could easily replicate what has been described here. To avoid this situation, this paper also includes simple schemes that could be applied by hardware vendors to "anonymize" their hardware, while still allowing `traceroute` and `ping`, that are valuable monitoring tools.

In addition, for security reasons, we have anonymized the four ASes of interest described in this paper.

References

[1] G. F. Lyon, *Nmap Network Scanning: The Official Nmap Project Guide to Network Discovery and Security Scanning*. Nmap Project, 2009, see http://nmap.org/book/toc.html.

[2] T. Kohno, A. Broido, and k. claffy, "Remote physical device fingerprinting," *IEEE Transactions on Dependable and Secure Computing*, vol. 2, no. 2, pp. 93–108, May 2005.

[3] Y. Vanaubel, J.-J. Pansiot, P. Mérindol, and B. Donnet, "Network fingerprinting: TTL-based router signature," in *Proc. ACM Internet Measurement Conference (IMC)*, October 2013.

[4] J.-F. Grailet and B. Donnet, "Towards a renewed alias resolution with space search reduction and IP fingerprinting," in *Proc. IFIP Network Traffic Measurementand Analysis Conference (TMA)*, June 2017.

[5] K. Vermeulen, S. Strowes, O. Fourmaux, and T. Friedman, "Multilvel MDA-lite paris traceroute," in *Proc. ACM Internet Measurement Conference (IMC)*, October 2018.

[6] Y. Vanaubel, P. Mérindol, J.-J. Pansiot, and B. Donnet, "Through the wormhole: Tracking invisible MPLS tunnels," in *Proc. ACM Internet Measurement Conference (IMC)*, November 2017.

[7] B. Hadad, B. Seri, and Y. Sarel, "CDPwn: Breaking the discovery protocols of the entreprise of things," Armis, Inc., Technical White Paper 20200205-1, February 2020, see https://www.armis.com/cdpwn/ for additional details.

[8] D. Bleichebacher, "Chose ciphertext attacks against protocols based on the RSA encryption standard PKCS#1," in *Proc. International Cryptology Conference on Advances in Cryptology (CRYPTO)*, August 1998.

[9] D. Felsch, M. Grothe, and J. Schwenk, "The dangers of key reuse: Practical attacks on IPsec IKE," in *Proc. USENIX Security Symposium*, August 2018.

[10] J. Postel, "Assigned numbers," Internet Engineering Task Force, RFC 1700, October 1994.

[11] Y. Vanaubel, J.-R. Luttringer, P. Mérindol, J.-J. Pansiot, and B. Donnet, "TNT, watch me explode: A light in the dark for revealing MPLS tunnels," in *Proc. IFIP Network Traffic Measurement and Analysis Conference (TMA)*, June 2019.

[12] J.-R. Luttringer, Y. Vanaubel, P. Mérindol, J.-J. Pansiot, and B. Donnet, "Let there be light: Revealing hidden MPLS tunnels with TNT," *IEEE Transactions on Network and Service Management (TNSM)*, vol. 17, no. 2, pp. 1239–1253, June 2020.

[13] B. Augustin, X. Cuvellier, B. Orgogozo, F. Viger, T. Friedman, M. Latapy, C. Magnien, and R. Teixeira, "Avoiding traceroute anomalies with Paris traceroute," in *Proc. ACM Internet Measurement Conference (IMC)*, October 2006.

[14] k. claffy, Y. Hyun, K. Keys, M. Fomenkov, and D. Krioukov, "Internet mapping: from art to science," in *Proc. IEEE Cybersecurity Application and Technologies Conference for Homeland Security (CATCH)*, March 2009.

[15] K. Keys, Y. Hyun, M. Luckie, and k. claffy, "Internet-scale IPv4 alias resolution with MIDAR," *IEEE/ACM Transactions on Networking*, vol. 21, no. 2, pp. 383–399, April 2013.

[16] K. Keys, "Internet-scale IP alias resolution techniques," *ACM SIGCOMM Computer Communication Review*, vol. 40, no. 1, pp. 50–55, January 2010.

[17] A. Marder, M. Luckie, A. Dhamdhere, B. Huffaker, J. Smith, and k. claffy, "Pushing the boundaries with bdrmapIT: Mapping router ownership at internet scale," in *Proc. ACM Internet Measurement Conference (IMC)*, November 2018.

[18] M. Sanchez, F. Bustamante, B. Krishnamurthy, W. Willinger, G. Smaragdakis, and J. Erman, "Inter-domain traffic estimation for the outsider," in *Proc. ACM Internet Measurement Conference (IMC)*, November 2014.

[19] M. Faloutsos, P. Faloutsos, and C. Faloutsos, "On power-law relationships of the internet topology," in *Proc. ACM SIGCOMM*, September 1999.

[20] L. Dall'Asta, I. Alvarez-Hamelin, A. Barrat, A. Vásquez, and A. Vespignani, "A statistical approach to the traceroute-like exploration of networks: Theory and simulations," in *Proc. Combinatorial and Algorithmic Aspects of Networking (CAAN) Workshop*, August 2004.

[21] R. Albert, H. Jeong, and A.-L. Barábsi, "Error and attack tolerance of complex networks," *Nature*, vol. 406, pp. 378–382, July 2000.

[22] R. Meier, P. Tsankov, V. Lenders, L. Vanbever, and M. Vechev, "NetHide: Secure and practical network topology obfuscation," in *Proc. USENIX Security Symposium*, August 2018.

[23] P. Marchetta, P. Mérindol, B. Donnet, A. Pescapé, and J.-J. Pansiot, "Quantifying and mitigating IGMP filtering in topology discovery," in *Proc. IEEE Global Communications Conference (GLOBECOM)*, December 2012.

[24] E. Davies and J. Mohacsi, "Recommendations for filtering ICMPv6 messages in firewalls," Internet Engineering Task Force, RFC 4890, May 2007.