Fingerprinting of Host Based Firewalls

Andrej Šimko

andrej.simko@mail.muni.cz

Faculty of Informatics
Masaryk University
Brno, Czech Republic

Abstract

This work describes new procedures for fingerprinting widely used Windows personal firewalls that are integrated into endpoint security system solutions. In order to represent majority of the market, we chose the most widespread solutions for testing. Free open source tool Nmap was used for this research. All tests were designed in order to be easily repeated. We describe two approaches to personal firewall fingerprinting – port state based and time based. Both approaches exploit various inconsistencies between IPv4 and IPv6 ports of certain firewalls. The port state based method observes in which states are top 1000 used ports reported by Nmap. It observes various differences between "open", "open|filtered", "unfiltered", "filtered" and "closed" ports detected by network probes. Special attention is given to TCP/0 port states. Scanning the TCP/0 port can separate firewalls into 4 groups. The time based method exploits differences on how long it takes for different Nmap scanning techniques to finish. More "exotic" Nmap techniques, for example SCTP cookie echo, TCP Maimon, or IP protocol scans are used with both approaches. We will show that it is rather easy to fingerprint all 18 personal firewalls selected in our testbed.

Keywords: firewall, fingerprinting, port scanning, port TCP/0.

1 Introduction

There are plenty of products which are advertising they are "the best" protection for the end user workstations. In the recent years most of them, if not all, have advanced into complex all-in-one packages. These packages contain antivirus, anti-malware, anti-spam, HIPS (Host Intrusion Prevention System), firewall and many other modules. This evolution into a single product is easy and comfortable approach for average computer users. They have many protections integrated into one system which is regularly updated with new signatures, and may even be a part of a cloud-based detection network. On the other hand, endpoint protection systems themselves are just another applications which are vulnerable to zero-day exploits themselves. See *Figure 1: Number of vulnerabilities reported in the National Vulnerability Database (NVD) for ten antivirus vendors between 2005 and November 2007*. Provided that the attacker could, with some probability, discover which endpoint security protection system is installed on the host, he could easily misuse it. Either by using exploits to that particular software, or by using such attacks/viruses/malware, which would not be detected in the first place.

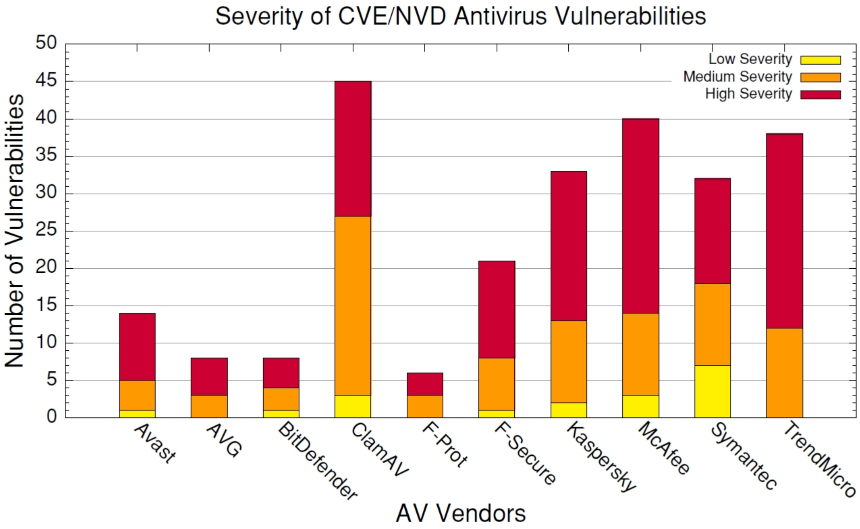

Figure 1: Number of vulnerabilities reported in the National Vulnerability Database (NVD) for ten antivirus vendors between 2005 and November 2007 [8].

This paper presents personal firewall fingerprinting with using widely available tool Nmap 6.47 [2] by performing different port scans on the victim and observing responses to scan probes. Although Nmap can detect “open”, “open|filtered”, “closed”, “filtered”, or “unfiltered” ports on the victim, used operating system, and sometimes fingerprint particular services which are running on selected ports, it lacks the capability to fingerprint different firewalls as such. It integrates various firewall/IDS (Intrusion Detection System) evasion techniques, for example fragmentation or address spoofing.

Protection against the port scanning is still a valid security concern. Even though there are plenty of countermeasures and security devices/features that should render port scanning useless (firewalls, IPS, NAT, proxy servers…) it is nevertheless widely used. When new vulnerabilities are discovered, there is a spike in worldwide port scanning of IPv4 address space [9] which happened for Heartbleed, NTP DDoS and many others. Another case is the HACIENDA program by NSA/GCHQ [10], which describes the common ground of port scanning of entire countries. Since there is obviously a risk connected to this attack, endpoint security solutions should be able to counteract it and protect its users.

This research is based on a master thesis [5], where more details, tests, tables and all Nmap log files and results can be found online.

2 Testing environment

A total of 18 different firewalls (see *Table 1: List of tested firewalls*) were tested, which authors believe to be among the most widely used Windows based endpoint protection systems. All products were downloaded as trial versions, activated if necessary, and updated to the latest version. The testing was done on virtual machines with a fully updated Windows 8.1 64-bit operating system to reflect the most up-to-date system a normal user could have installed to protect himself. Each particular endpoint security system was installed on Windows with clean installation and no additional software nor any services. This includes the absence of SCTP protocol, although we used both Nmap SCTP scanning techniques. Static IPv4 and IPv6 addresses were configured both for the attacker and the victim. If the firewall asked about configuration settings, it was set to "work" profile (out of "public", "work" and "private" profiles of network) and "automatic" (out of "automatic", "interactive" and "learning" modes).

It is important to note that no other options were changed on firewalls. As this research focused on quantitative rather than qualitative analysis, and most endpoint protection users are not making any changes to default settings, no product-specific options were pursued.

To observe the difference across all firewalls, following 12 Nmap scanning techniques were used: TCP SYN (*-sS*), TCP Connect (*-sT*), TCP ACK (*-sA*), TCP Window (*-sW*), TCP Maimon (*-sM*), UDP (*-sU*), TCP Null (*-sN*), TCP FIN (*-sF*), TCP Xmas (*-sX*), SCTP Init (*-sY*), SCTP cookie echo (*-sZ*), and IP protocol scan (*-sO*). Default Nmap scan of 1000 most widely used ports was applied on every firewall. Note that both SCTP scans have only 52 ports in total which are scanned by the Nmap. On the IP protocol scan, all 256 possibilities in the 8-bit IP protocol field in IP header were scanned. Every port technique was used both on both IPv4 and IPv6 protocols.

Company	Product	Tested version
Agnitum	Outpost Pro Security Suite	9.1
Avast!	Internet Security	2014.9.0.2021
AVG	Internet Security 2015	2015.0.5315
Avira	Antivirus Pro	14.0.7.306
Bitdefender	Internet Security 2015	18.17.0.1227
COMODO	Internet Security Premium	7.0.317799.414
Emsisoft	Internet Security	9.0.0.4570
ESET	Smart Security	8.0.301.0
F-Secure	SAFE Internet Security	14.115 build
Gdata	Internet Security	24.4727
Kaspersky	Internet Security 2015	15.0.0.463 (a)
McAfee	Total Protection	12.8.988
Microsoft	Windows 8 Firewall	-
Norton	Security	22.0.1.14
Panda	AntiVirus Pro 2015	15.0.4
Quick Heal	AntiVirus Pro	15.00 (8.0.8.0)
TrustPort	Internet Security	14.0.5.5273
ZoneAlarm	Free Antivirus + Firewall 2015	13.3.209.000

Table 1: List of tested firewalls.

3 Related work

To the best of our knowledge, there is no published research on extensive host-based firewall fingerprinting prior to this work. There have been only few papers on firewall fingerprinting in general, usually focusing on enterprise firewalls. In [1], the method of observing firewall decisions based on TCP packets with unusual flags along with the machine learning techniques were used. However, because of the budget issues, the testbed consisted only of 3 hardware firewalls whose brands were not disclosed due to privacy reasons.

On the other hand, fingerprinting of operating systems, along with particular services running on the victim is fairly common. Tools like Nmap [2], Xprobe2++ [3], or p0f [4] can be used for these purposes. However, none of them is capable of firewall fingerprinting. Their only capability is to distinguish that a firewall could be installed on the victim's workstation.

4 Port state based fingerprinting

When Nmap sends its probes to the victim, the victim's firewall can respond with various kinds of responses. It can choose not to respond at all, send TCP packets containing particular flags (e.g. SYN+ACK to TCP SYN scan, RST to TCP ACK scan, …), or generate an ICMP message. Based on the firewall's response, Nmap can differentiate between various port *states*. All responses, along with different ICMP messages have different interpretation. For example on UDP scan, the ICMP port unreachable error (type 3, code 3) shows that the port is "closed", while ICMP unreachable errors (type 3, codes 1, 2, 9, 10 or 13) designate ports as being "filtered". Many firewalls differ on how they respond to particular probes and scanning techniques, which makes port state based fingerprinting possible.

Nmap assigned state to every scanned port. Since every technique has different possibilities of port states, they were taken into an account. For example, on TCP SYN scan, the Nmap differentiates between the "open", "closed", and "filtered"; while for the TCP Maimon scan it can differentiate only between "open|filtered" and "closed" states. The summation numbers of all ports in particular states were observed. For example: 7 open ports, 986 closed and 7 filtered ports (1000 in total) were seen on TCP SYN scan on IPv4 on Kaspersky. To put down all results, we created a table with 70 columns. Each of the 12 scanning techniques with all possible port states on both IPv4 and IPv6 was written down for further analysis. Unfortunately, such a table can't be shown here in full scope. See *Table 2: Port states on TCP ACK scan on IPv6* for an example of abbreviated version.

Company	unfiltered	filtered
Agnitum	11	989
Avast!, AVG, Avira, COMODO, ESET, F-Secure, Gdata, McAfee, Microsoft, Norton, Quick Heal	0	1000
Bitdefender	-	-
Emsisoft	977	23
Kaspersky	993	7
Panda, TrustPort, ZoneAlarm	1000	0

Table 2: Port states on TCP ACK scan on IPv6.

Very interesting result was achieved by observing port states of port TCP/0, which is reserved by IANA [6] and should not be used in any applications. The scan of this port was not detected by any firewall, which makes it very stealthy, fast, and thanks to unique results between firewalls, it can be easily used for the base of fingerprinting. See *Table 3: TCP/0 port states across firewalls* for more details.

Company	IPv4	IPv6
Agnitum, AVG, Avira, COMODO, ESET, F-Secure, Gdata, McAfee, Microsoft, Norton, Quick Heal	filtered	filtered
Avast!, Emsisoft, Kaspersky, Panda	closed	closed
Bitdefender, ZoneAlarm	filtered	-
TrustPort	filtered	closed

Table 3: TCP/0 port states across firewalls.

All firewalls have a default behavior how to respond to certain packets. Their response can be different – for example with a packet with RST flag, generating ICMP error, or by sending SYN+ACK packet. Nmap will then categorize responses from firewalls into the port states “open”, “open|filtered”, “closed”, “filtered”, or “unfiltered”. The most used behavior observed across our testbed with TCP SYN scan is “filtered”. The only exceptions to this rule on IPv4 were Emsisoft and Kaspersky, which had most of the ports designated as “closed”. TrustPort changed the default behavior – on IPv4 it was “filtered”, while on IPv6 it was “closed”. Nmap on their webpages states that Microsoft Windows doesn’t fully follow RFC 793 [2] (in chapter: Port Scanning Techniques) and sends RST packet to TCP FIN, TCP Null and TCP Xmas probes which are therefore designated as “closed”. However the observed behavior was different across firewalls all using Windows 8.1. For the IPv4, only Emsisoft and Kaspersky designated most of the ports as “closed”, whereas all other firewalls made them “open|filtered”. For the IPv6 it was even more different – Emsisoft, Kaspersky, Panda, TrustPort and ZoneAlarm designated them as “closed”, while the scan was entirely unsuccessful on Bitdefender. Not used SCTP protocol had default behavior on McAfee, TrustPort and ZoneAlarm as “filtered”, whether on all other firewalls it was “open|filtered” on IPv6. IP protocol scan on IPv6 marked AVG, Emsisoft, Kaspersky, Panda, TrustPort and ZoneAlarm to have default port states “closed”, whether all other (apart from BitDefender) were stated as “open|filtered”.

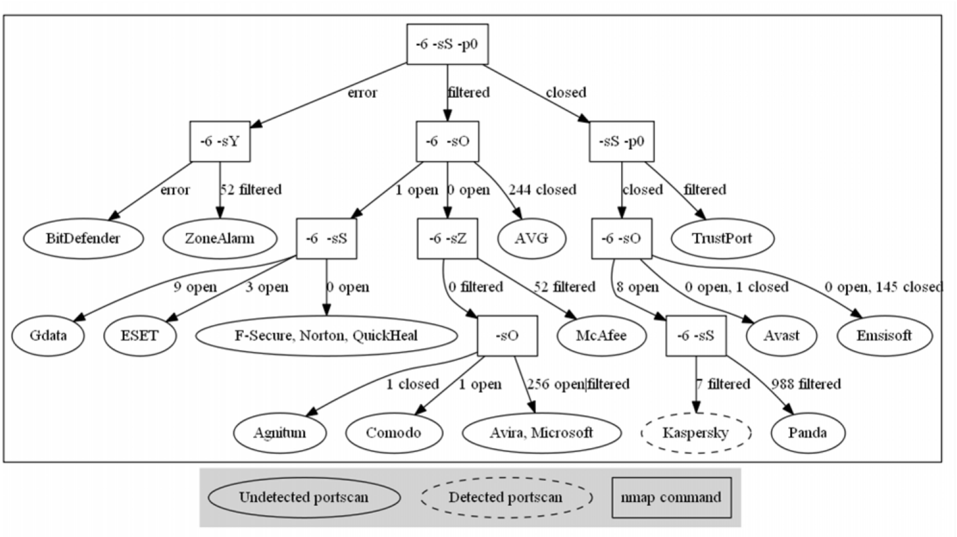

Figure 2: Port state based fingerprinting decision tree.

The observed default behavior varies across firewalls and IP protocol version. By performing different Nmap port scanning attacks and observing the number of ports in certain states, we can effectively fingerprint different firewalls. We've created a decision tree – a graph of Nmap scans, which can be followed to reach the decision which firewall from our testbed is installed on the victim's computer. It tries to use undetectable port scanning techniques first and using this approach, only Kaspersky had the ability to detect this fingerprinting scans. See *Figure 2: Port state based fingerprinting decision* tree. As you can see, it is not possible to distinguish between (F-Secure, Norton, QuickHeal) and (Avira, Microsoft) firewalls with this simple method.

5 Time based fingerprinting

While the port state based fingerprinting took into account only the different numbers of ports in particular states, the time based method observes the time which Nmap required to complete the scan. Scanning 1000 default ports on 12 techniques could range anywhere between 1.38 and 3769.63 seconds. Again, every technique was used on both IPv4 and IPv6.

Many firewalls had significantly different behavior based on the IP version used. Refer to *Table 4* and *Table 6* for the result both port scans.

Notice the last 3 columns used in each of these tables. These values represent the number of high deviations from ideal values. Ideal values are listed in the *Table 6: Ideal port scanning times (in seconds)*. For example, the value 12 with the Kaspersky in “5 %” column shows that all 12 scanning techniques had their difference from the ideal value higher than 5 % of the ideal value of every technique. The formula used in Excel sheet was “=IF(ABS((observed_value-ideal_value))<=percentage* observed_falue; TRUE; FALSE)”. Then the number of “TRUE” values was counted and inserted into Tables 4 and 6.

The higher the number in the last 3 columns is, the better and more certainly can we fingerprint these firewalls. If the numbers are low, firewalls are hardly distinguishable (at least for the IPv4 protocol). Number 7 with Emsisoft on IPv4 in “200 %” column says that 7 scanning techniques had deviation within 200 % of the ideal value which makes the most easily distinguishable firewall using time based fingerprinting technique. As we can see in the Table 6, there are far greater numbers in every column which makes IPv6 far more easy to use for fingerprinting purposes.

For the TCP SYN scan, the whole scanning took somewhere between 1.45 seconds (Panda on IPv6) and 1243.67 seconds (Avast! on IPv4). If we take into account global extremes, scanning 1000 ports took somewhere from 1.38 seconds (Panda on IPv6 under the TCP ACK and TCP NULL scans) to 3769.63 seconds (Panda on IPv4 under UDP scan). Panda was therefore the shortest and longest firewall to be scanned.

In almost every case, using IPv6 instead of IPv4 resulted in much faster scan. The most significant difference was observed with Panda – using UDP scan on IPv6 was faster by 2698.11 seconds compared to IPv4 scan. On the other hand, Emsisoft UDP scan on IPv6 was slower by 1137.09 seconds compared to the IPv4 scan. Another example was TCP SYN scan on Avast!, which was faster by 1220.48 seconds on IPv6 compared to IPv4.

Company	-sS	-sT	-sA	-sW	-sM	-sU	-sN	-sF	-sX	-sY	-sZ	-sO	5%	75%	200%
Agnitum	22.91	93.69	22.91	4.34	4.76	22.91	22.91	1236.80	1243.58	2.34	2.34	3.95	6	4	2
Avast!	1243.67	2299.29	23.22	22.91	22.91	22.91	22.91	22.91	22.91	2.34	2.34	3.02	3	3	0
AVG	25.19	45.68	22.91	22.91	22.91	22.91	22.91	22.91	23.64	2.34	2.34	6.72	1	0	0
Avira	22.91	45.57	22.91	22.91	22.91	22.91	22.91	22.91	27.14	2.34	2.34	6.72	1	0	0
Bitdefender	22.91	45.56	22.91	22.91	22.91	22.92	22.91	22.91	22.91	2.34	2.34	6.72	0	0	0
COMODO	22.91	46.56	22.97	22.91	22.91	22.91	22.91	22.91	22.91	2.34	2.34	226.84	1	1	0
Emsisoft	3.77	46.79	2.38	3.59	3.52	22.91	3.41	5.02	3.36	2.34	2.34	2.58	8	8	7
ESET	22.91	45.71	22.91	22.91	22.91	22.91	22.91	22.91	22.91	2.34	2.34	5.97	1	0	0
F-Secure	22.91	46.36	23.19	23.03	23.25	23.19	23.22	23.13	23.25	2.34	2.34	6.72	0	0	0
Gdata	5.00	45.57	22.91	22.91	22.92	7.83	22.91	22.91	22.91	2.34	2.34	2.91	3	3	1
Kaspersky	108.39	124.91	2.80	2.78	1.44	1111.39	2.81	1.56	1.44	2.13	2.13	300.11	12	9	6
McAfee	5.20	45.76	22.91	22.91	22.91	22.91	22.92	22.91	22.91	2.34	2.61	6.72	2	1	1
Microsoft	22.92	45.82	22.91	22.91	22.92	22.91	22.91	22.91	22.91	2.34	2.34	6.72	0	0	0
Norton	22.92	46.51	23.36	22.91	23.45	22.91	23.17	23.02	23.16	2.34	2.34	6.72	0	0	0
Panda	1135.03	2669.77	1.58	1.47	1.48	3769.63	22.92	35.98	22.91	2.13	2.13	312.08	10	7	3
Quick Heal	23.19	45.88	22.91	22.91	22.91	22.91	22.92	22.92	22.92	2.34	2.34	6.72	0	0	0
TrustPort	13.00	45.65	22.91	22.91	22.91	15.28	22.91	22.91	22.91	2.34	2.34	6.72	2	1	0
ZoneAlarm	22.91	45.80	22.91	22.91	22.91	22.91	22.91	22.91	22.92	2.34	2.34	6.72	0	0	0

Table 4: IPv4 port scanning results in seconds with number of deviations from the ideal time.

Company	-sS	-sT	-sA	-sW	-sM	-sU	-sN	-sF	-sX	-sY	-sZ	-sO	5%	75%	200
Agnitum	4.33	45.50	4.55	5.20	5.20	22.91	22.91	4.67	5.20	2.34	2.34	6.72	6	6	6
Avast!	23.19	1257.20	22.91	22.91	22.94	22.91	22.91	22.91	22.91	2.34	2.34	3.88	2	1	0
AVG	8.92	45.81	22.91	22.91	22.91	22.91	22.91	22.91	22.91	2.13	2.13	335.77	4	2	0
Avira	22.91	45.80	22.91	23.41	22.91	22.91	22.91	22.91	22.91	2.34	2.36	6.72	0	0	0
Bitdefender	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
COMODO	22.91	46.84	22.91	22.91	22.91	22.91	22.91	22.91	22.91	2.34	2.34	6.72	0	0	0
Emsisoft	3.80	50.90	3.45	3.67	3.61	1160.00	3.83	5.42	2.39	2.13	2.13	171.55	12	9	7
ESET	11.42	45.70	23.39	23.17	23.98	22.91	23.20	22.91	22.91	2.34	2.34	3.45	2	2	0
F-Secure	22.91	45.68	23.09	23.36	23.16	23.28	23.22	23.19	23.02	2.34	2.34	5.44	1	0	0
Gdata	4.77	45.66	22.91	22.91	22.91	22.91	22.91	22.91	22.91	2.36	2.34	3.89	2	1	1
Kaspersky	312.86	69.89	1.42	2.52	1.47	1117.58	1.45	1.47	1.58	2.13	2.13	300.80	12	9	6
McAfee	4.98	45.75	22.91	22.91	22.91	22.91	22.91	22.91	22.91	2.34	2.34	6.73	1	1	1
Microsoft	23.34	45.86	22.91	22.91	22.91	22.91	22.91	22.91	22.91	2.34	2.36	6.72	0	0	0
Norton	22.91	45.82	22.91	22.92	22.91	22.91	22.91	22.91	22.91	2.34	2.34	3.45	1	1	0
Panda	1.45	45.37	1.38	1.39	1.47	1071.52	1.38	1.39	1.39	2.13	2.13	279.95	11	9	7
Quick Heal	22.91	45.90	22.91	22.92	22.91	22.91	22.92	22.91	22.91	2.34	2.34	2.25	1	1	0
TrustPort	1.47	45.34	1.67	1.56	1.56	1062.59	1.47	1.69	1.45	2.13	2.13	287.66	11	9	7
ZoneAlarm	-	-	1.47	15.19	1.59	44.22	1.48	1.48	1.51	0.00	0.00	4.72	10	7	7

Table 5: IPv6 port scanning results in seconds with number of deviations from the ideal time.

Thanks to these extremes and many different values in between these intervals, it is possible to differentiate between many firewalls with a sufficiently high confidence. We propose *Figure 3: Time based fingerprinting decision tree* to be one of possible decision trees for time based fingerprinting. Avast! is the only firewall that can detect an attacker who is profiling the firewall with this decision tree procedure. We created a second tree which was aiming at reliability instead of detection avoidance. With using that tree, Avast!, Kaspersky and McAfee detected scannings.

This Nmap call tree is easy to use. In rectangles, there are Namp commands used. Based on the value of how long does it take for scan to finish, user can follow different paths, until he reaches the final node (leaf) and discovers firewall. Using this approach in our scope of 18 firewalls, the particular endpoint protection system can be distinguished with using between 1 (BitDefender and ZoneAlarm) to 6 (Avira, Microsoft) default Nmap scans. In contrast to the port state approach, all firewalls were successfully fingerprinted here.

For alleviating the risk of time based fingerprinting, we propose the ideal port scanning times which are stated in *Table 6: Ideal port scanning times (in seconds)*. These are the times that were observed as the most occurring ones on both IPv4 and IPv6 across all firewalls. Provided there would be none or minimal differences between different brands, the time based fingerprinting would be effectively countered. Only BitDefender and Microsoft came close to these numbers in all scans on both IPv4 and IPv6. All other firewalls had anomalies which would allow them being fingerprinted. If we would take into account only IPv4 and less than 5 % deviation from the ideal scanning times across all Nmap scanning techniques, following firewalls would be almost indistinguishable: BitDefender, F-Secure, Microsoft, Norton, Quick Heal and ZoneAlarm. On the IPv6 alone it would be Avira, BitDefender, COMODO and Microsoft.

Port scanning technique	Ideal time
TCP SYN, TCP FIN, TCP Xmas, TCP Null, TCP ACK, TCP Window, TCP Maimon, UDP	22.91
TCP Connect()	45.56
SCTP Init, SCTP cookie echo	2.34
IP protocol scan	6.72

Table 6: Ideal port scanning times (in seconds).

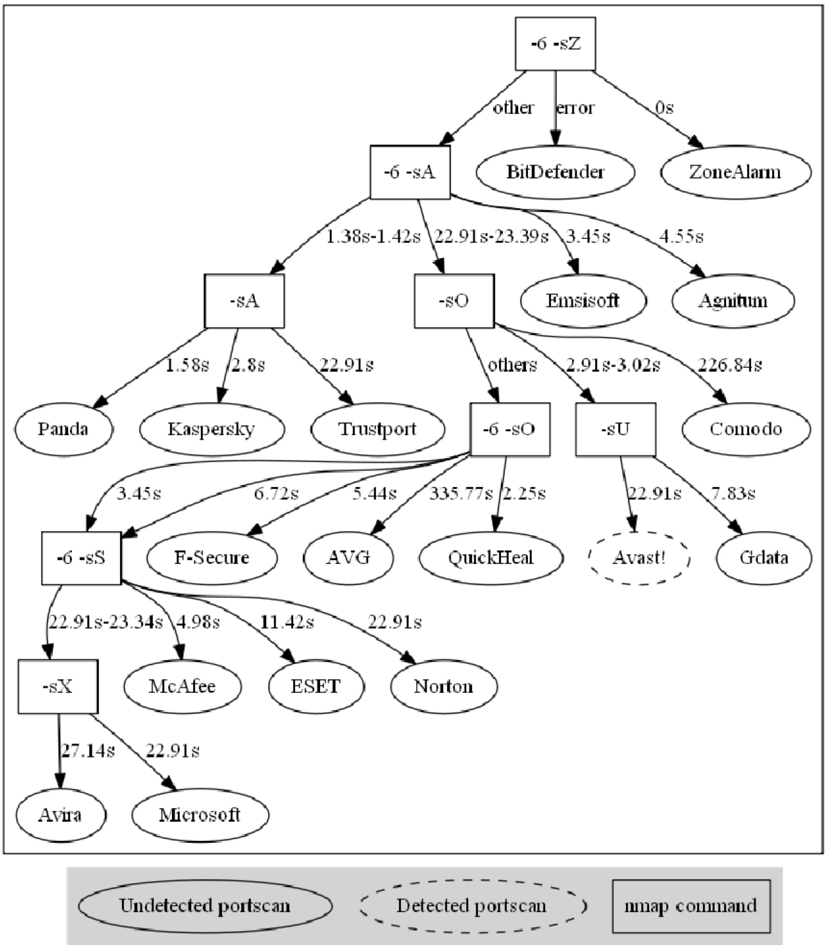

Figure 3: Time based fingerprinting decision tree

6 Conclusions

If the attacker already has an access to the local network and he possesses the port scanning capabilities, he can easily find out which firewalls are installed on the computers inside of the network. This is possible because all firewalls share significant differences – may it be the time consumption of selected scanning techniques, or the differences in port states. Many firewalls are inconsistent in facing IPv4 and IPv6 port scans, which only adds benefits to fingerprinting. Out of 18 firewalls in our testbed, we successfully created two decision trees (*Figure 2: Port state based fingerprinting decision* tree and *Figure 3: Time based fingerprinting decision tree*). Although there are other ways of how to create such Nmap calling trees, these two have proven to be successful.

Firewall fingerprinting is largely possible because the majority of firewalls are not sufficiently protected against port scanning attacks. We proposed *Table 6: Ideal port scanning times (in seconds)* to be the “holy grail” of countermeasure against time based fingerprinting approach. As for the port state based approach, firewalls would not leak any information about the port states and hence they should respond to port scanning probes in the same way no matter the IP protocol, nor different Nmap techniques used.

7 Discussion

There are plenty of improvements of this research which can be done to push the boundaries of personal firewall fingerprinting even further. Testing more versions of the same endpoint security protection system could reveal if there are any interesting variations across major releases of these products. Individual responses to malformed packets could be pursued [7] (e.g. invalid checksums or unexpectedly big data parts of port scanning packets) for more different ways of fingerprinting. Use of fragmentation or permutations of flags which are not used by Nmap scanner can also give interesting results. Scanning all 65 536 ports with every technique can give more interesting results to which ports are behaving out of the ordinary. For example, if port X is not part of top 1000 scanned ports in Nmap and it would be open only for one firewall, scanning this particular port while fingerprinting would be very effective. Noting particular ports with their states as a base for new approach is also feasible.

Therefore, devising a strategy of for example scanning port A on TCP ACK on IPv6 and scanning port B on UDP on IPv4 could fingerprint firewalls with the least amount of scans under few milliseconds. Observing the exact responses to scanning probes (e.g. which ICMP messages are generated) in Wireshark could also shed some light on improving the fingerprinting process. Testing the same firewalls across different platforms (Linux, Mac), if available, could point to Windows-specific behavior. Different Windows versions could also be scanned to discover specific behavior of Windows 8.1. Trying different profiles (public/private network settings, or interactive/learning modes), along with tweaking settings of every firewall to observe what can be changed could elaborate on invariable conditions, effectively removing behavior which depends on something that could be changed.

To counter the possibility of fingerprinting, there are 3 main areas where the vendors of these firewalls should improve:

1. No leakage of port states. Where possible, Nmap should say “open|filtered” with all ports scanned. Where this option is unavailable, “filtered” state should be used.

2. Unified time consumption of port scans. Every firewall should have similar or identical results.
3. No differences between behavior on IPv4 and IPv6. This is now perhaps one of the greatest weaknesses among most of the firewalls.

References

[1] Khakpour, Amir R., Joshua W. Hulst, Zihui Ge, Alex X. Liu, Dan Pei, and Jia Wang.: Firewall fingerprinting, in *INFOCOM, 2012 Proceedings IEEE*, pp. 1728-1736. IEEE, 2012.

[2] Lyon, Gordon F.: *Nmap Network Scanning: The Official Nmap Project Guide to Network Discovery and Security Scanning*. Insecure, 2009.

[3] Yarochkin, Fedor V., Ofir Arkin, Meder Kydyraliev, Shih-Yao Dai, Yennun Huang, and Sy-Yen Kuo.: Xprobe2++: Low volume remote network information gathering tool, in *Dependable Systems & Networks, 2009. DSN'09. IEEE/IFIP International Conference on*, pp. 205-210. IEEE, 2009.

[4] Zalewski, M.: p0f, http://lcamtuf.coredump.cx/p0f3/

[5] Šimko, Andrej. Comparative Analysis of Personal Firewalls [online]. 2015. Master thesis. Masaryk University, Faculty of informatics. Available from: http://is.muni.cz/th/359952/fi_m/.

[6] Touch J., Lear E., Mankin A.: Service Name and Transport Protocol Port Number Registry, http://www.iana.org/assignments/service-names-port-numbers/service-names-port-numbers.xhtml, 2015/03/16.

[7] Ptacek, Thomas H., and Timothy N. Newsham. *Insertion, evasion, and denial of service: Eluding network intrusion detection*. SECURE NETWORKS INC CALGARY ALBERTA, 1998.

[8] Oberheide, Jon, Evan Cooke, and Farnam Jahanian. "CloudAV: N-Version Antivirus in the Network Cloud." In *USENIX Security Symposium*, pp. 91-106. 2008.

[9] Durumeric, Zakir, Michael Bailey, and J. Alex Halderman. "An Internet-wide view of Internet-wide scanning." *USENIX Security Symposium*. 2014.

[10] Julian Kirsch, Christian Grothoff, Monika Ermert, Jacob Appelbaum, Laura Poitras, Henrik Moltke, NSA/GCHQ: The HACIENDA Program for Internet Colonization, http://heise.de/-2292681, 2014/08/15.