INTERNET SERVICE PROVIDERS AND BOTNET MITIGATION
A Fact-Finding Study on the Dutch Market

Report prepared for the Netherlands Ministry of Economic Affairs, Agriculture and Innovation

This is an anonimized version of the report, meant for public release

January 2011

Michel J.G. van Eeten*
Hadi Asghari*
Johannes M. Bauer**
Shirin Tabatabaie*

* Faculty of Technology, Policy and Management
Delft University of Technology
The Netherlands
** Quello Center for Telecommunication Management & Law
Michigan State University
East Lansing, USA

Table of Contents

1.
Mitigating the Threat of Botnets: From End Users to Internet Service Providers[1]

Background

The internet economy is highly dependent on information and network security. Estimates of the overall damage of internet security incidents vary wildly, but typically range in the tens of billions of US dollars per year for the U.S. alone.[2] While this damage is related to a wide variety of threats, the concurrent rise of malicious software ("malware") and botnets are seen as one of the most urgent security threats.

Malware has become a critical security threat to all users who rely on the Internet for their daily business, whether they are large organizations or residential subscribers. While initially a nuisance more than a threat, viruses, worms and the many other variants of malware have developed into a sophisticated set of tools for criminal activity. Computers around the world, some experts estimate as many as one in ten, are infected with malware, often unknown to the owner of the machine. Many of these infected machines are connected through botnets: networks of computers that operate collectively to provide a platform for criminal purposes. These activities include, but are not limited to, the distribution of spam (the bulk of spam now originates from botnets), hosting fake websites designed to trick visitors into revealing confidential information, attacking and bringing down websites, enabling so-called 'click fraud,' among many other forms of often profit-driven criminal uses. There are also reports that indicate terrorist uses of malware and botnets. This report, however, focuses primarily on botnets as an economic threat.

While originating in criminal behavior, the magnitude and impact of the botnets is also influenced by the decisions and behavior of legitimate market players such as Internet Service Providers (ISPs), software vendors, e-commerce companies, hardware manufacturers, registrars and, last but not least, end users. As security comes at a cost, tolerating some level of insecurity is economically rational. Market players make their decisions based on the perceived costs and benefits of a course of action. In many situations, these private decisions also reflect the resource costs and benefits of a course of action to society at large. However, economic research and policy analysis have also identified situations in which this correspondence is weakened, for example, because players impose costs on or generate benefits for others without a corresponding market transaction, situations for which the term "externalities" is used.

---

[1] The material in this chapter relies heavily on Van Eeten, M., J. Bauer, H. Asghari and S. Tabatabaie (2010). *The Role of Internet Service Providers in Botnet Mitigation: An Empirical Analysis Based on Spam Data. STI Working Paper 2010/5.* OECD. Available online at http://www.oecd.org/officialdocuments/displaydocument/?doclanguage=en&cote=dsti/doc(2010)5.

[2] See US GAO (2007). *Cybercrime: Public and Private Entities Face Challenges in Addressing Cyber Threats*. United States Goverment Accountability Office. Available online at http://www.gao.gov/new.items/d07705.pdf.

Recent economic research has found that the infected machines of end users are a key source of security externalities, most notably home users and small and medium-size enterprise (SME) users.[3] In contrast to larger corporate users, these groups often do not achieve desirable levels of protection. The large number of infected end user machines can be recruited into botnets and abused as a powerful platform for a criminal underground economy.

Measures that address end users directly – most notably awareness raising and information campaigns – are useful, but they have proven to be insufficient to reduce the overall problem. Recent studies have therefore shifted attention to key intermediaries, most notably, the Internet Service Providers (ISPs) that provide access to end users.[4]

Botnet Mitigation by Internet Service Providers

ISPs form, to some extent, a natural control point for the effects of infected machines. Of course, the fact that ISPs can potentially mitigate this threat, does not mean that they *should* mitigate it. They are not the source of the externality but would have to bear substantial direct and indirect costs if they internalize the externalities of their customers. Nevertheless, the leading ISPs in the Netherlands have entered into a covenant that expresses their commitment to mitigate botnet activity in their own networks.

Fourteen Dutch ISPs, representing over 90 percent of the access market, are now collaborating in a Anti-Botnet Working Group. All have agreed to put botnet mitigation practices in place – more precisely, they committed to contacting and in some cases quarantining customers whose machines are infected with malware. There is currently no data available that indicates the scale on which these practices are being carried out. Neither has there been any research into the impact of these practices on the infection levels in the networks of Dutch ISPs.

Scale is critical. There are indications that ISPs only deal with a fraction of the infected machines in their networks. For example, in an earlier study we found that a large European ISP with over four million customers contacted around 1,000 customers per month.[5] Estimates of security researchers put the number of infected machines at around one to five percent of all connected machines at any point in time. This would mean between 40,000 and 200,000 infected machines for this specific ISP, a number that stands in stark contrast to the 1,000 customers that the ISP claimed to be contacting – even when we optimistically assume that all contacted customers are willing and able to clean up their infected machine.

To reiterate: We are not claiming that ISP *should* contact all the owners of infected machines. That is a matter for policy development to consider and will be dependent on the costs and benefits of alternative courses of action. However, we are claiming that policy development in this area urgently needs empirical data on the role of ISPs. The data should inform us about the extent to which ISPs are actually engaged in botnet mitigation, as well as their performance relative to each other, in the Dutch market, as well as internationally.

---

[3] See Van Eeten, M. and J. M. Bauer (2008). *Economics of Malware: Security Decisions, Incentives and Externalities, OECD STI Working Paper 2008/1*. OECD. Available online at http://www.oecd.org/dataoecd/53/17/40722462.pdf.

[4] See for example: Anderson, R., R. Böhme, R. Clayton and T. Moore (2008). *Security Economics and the Internal Market*. ENISA (European Network and Information Security Agency). Available online at http://www.enisa.europa.eu/doc/pdf/report_sec_econ_&_int_mark_20080131.pdf.

[5] See Van Eeten and Bauer (2008, pp. 26-34).

The Dutch Ministry of Economic Affairs, Agriculture and Innovation has commissioned us to conduct a fact-finding study on the problem of botnet infections in the Netherlands and the role of ISPs in mitigating this problem. Identifying recommendations for government policies or ISP practices are outside the scope of this study. Though we briefly reflect on the policy implications of our findings in the conclusion of the report, the task to identify the strategy to combat botnets is part of the ongoing collaboration among government and industry in the Platform Internet Security. We aim to facilitate this collaboration by providing new empirical evidence on the state of infected machines and botnet mitigation in the Netherlands.

Our research has been executed independently, but it has benefited greatly from feedback from the ISPs participating in the Dutch Botnet Mitigation Working Group. Our methodology and findings have been discussed in depth during a workshop with the Working Group. We reflect on that workshop in Chapter 4 of this report. Part of the collaboration with the Working Group was the agreement that ISPs would provide us with confidential information on the number of customers that they had during the period under study. This allowed us to corroborate our own data. To protect the confidentiality of the data that was provided to us, we agreed to remove from the public version of the report all ISP names in relation to specific numbers of infected machines in their networks. We have replaced each name consistently with a generic label, such as NL01. Our client, the Ministry of Economic Affairs, Agriculture and Innovation, received a confidential version of the report from which the ISP names were not removed.

Research Objectives

Our report addresses the need for data on botnet infections and mitigation efforts in the Netherlands. Its main goal is to produce a robust fact-finding report on the number of infected machines located in ISP networks in the Netherlands. To this end, we have identified the following objectives:

1. Collect data from different sources to assess the number of infected machines in the Netherlands from January 2009 to June 2010 and benchmark the findings against other countries – e.g., on a per capita basis;
2. Establish what percentage of infected machines in the Netherlands are located in the networks of Dutch ISPs – in other words, the extent in which Dutch ISPs are indeed control points for botnet mitigation;
3. Collect data from different sources to assess the number of infected machines within the networks of Dutch ISPs;
4. Develop preliminary benchmarks that rank the rate of infection of Dutch ISPs against each other and against ISPs in other countries;
5. Discuss the methodology and findings during a workshop with the ISPs participating in the Anti-Botnet Working Group.

In the Dutch context, this research project is directly relevant to several ongoing initiatives, most notably the Anti-Botnet Working Group. But there are other collaborative efforts of ISPs and the Dutch government. We mention the Platform Internet Security (“Platform Internetveiligheid”) and project Taurus of the Dutch national police agency, who have made great progress in investigating and combating the command and control infrastructure of botnets.

Our study is not only relevant for the Netherlands. In a variety of other countries, ISPs have also indicated that they are willing to share responsibility for botnet mitigation. One such example is the recently signed code of conduct of the Australian Internet Industry Association (IIA) that suggests ISPs should contact, and in some cases disconnect, customers that have malware-infected

computers. Within the OECD, other countries have indicated they are pursuing similar initiatives. Public-private initiatives in Japan, Korea and Germany include the distribution of malware removal tools to infected users and the establishment of government-funded call centers to which ISPs can direct customers in need of support to disinfect their machines. Our report develops an empirical approach that may be helpful to increase the understanding of botnets and botnet mitigation beyond the Netherlands.

Who Is Included?

Our study has collected data on infection levels for all participants of the Anti-Botnet Working Group: Bbned, KPN, Luna, Online, Scarlet, Solcon, Tele2, UPC, Xenosite, XS4All, Ziggo. Together, these ISPs have an aggregate market share of well over 90 percent of the broadband market.

In the emprical part of the report, we report on all participants, except the three smallest providers: Xenosite, Luna, Scarlet. They harbor less than 0.5 percent off all infected machines in the Netherlands. This cut-off point has been used to keep the descriptions and graphics readable. Also, we have treated Telfort as part of its parent company KPN, because we lacked separate customer data for the period under study.

There are two ISPs in the Netherlands that are not part of the Anti-Botnet Working Group, even though they are comparable in size to some of the ISPs that do particpate: Zeelandnet and CAIWAY. We have added Zeelandnet to the study. We also have data on CAIWAY, but opted not to include them, because of measurement issues around use of dynamic IP addresses with very short lease times. The number of infections in their networks is hard to assess, but certainly large enough to merit inclusion in the Working Group. This may be an option for the Working Group to consider.

Report Outline

In the next chapter, we first outline the methodology we developed to identify the location of infected machines, as well as the data sets on which this methodology was applied. Then, in Chapter 3, we turn to the actual findings. We assess infection levels of the Netherlands as a whole, compared to other countries. Then we investigate the extent to which these infected machines are located in the networks of the Dutch ISPs. We also address the question of how infection levels have fluctuated during the period under study (January 2009 to June 2010). Last, we discuss how ISPs perform relative to each other. In Chapter 4, we report on a workshop with the ISPs where we received feedback on our methodology and explored the implications of our findings. Finally, in Chapter 5, we summarize the main findings and reflect on their implications for ISPs.

This is the public version of the report, from which we have removed all ISP names. We replaced each name consistently with a generic label, ranging from NL01 to NL14. Not all numbers in this range are used in the report as they refer to ISPs which are not included in the anlysis, for reasons outlined above. The removal of ISP names was part of the collaborative agreement with the ISPs in the Botnet Mitigation Working Group. They supplied us with confidential data on the number of customers over time. To protect that confidentiality, all names were removed in relation to specific results.

2.
Research Approach

Data on Infected Machines

There is currently no authoritative data source to identify the overal population of infected machines around the world. Commercial security providers typically use proprietary data and shield their measurement methods from public scrutiny. This makes it all but impossible to correctly interpret the figures they report and to assess their validity.

The publicly accessible research in this area relies on two types of data sources:

- *Data collected external to botnets*. This data identifies infected machines by their telltale behavior, such as sending spam or participating in distributed denial of service attacks;
- *Data collected internal to botnets.* Here, infected machines are identified by intercepting communications within the botnet itself, for example by infiltrating the command and control infrastructure through which the infected machines get their instructions.

Each type of source has its own strengths and weaknesses. The first type typically uses techniques such as honey pots, intrusion detection systems and spam traps. It has the advantage that it is not limited to machines in a single botnet, but can identify machines across a wide range of botnets that all participate in the same behavior, such as the distribution of spam. The drawback is that there are potentially issues with false positives. The second type typically intercepts botnet communications by techniques such as redirecting traffic or infiltrating IRC channel communication. The advantage of this approach is accuracy: bots connecting to the command and control server are really infected with the specific type of malware that underlies that specific botnet. The downside is that measurement only captures infected machines within a single botnet. Given the fact that the number of botnets is estimated to be in the hundreds (Zhuang et al. 2008), such data is probably not representative of the overal population of infected machines.

Neither type of data sources sees all infected machines, they only see certain subsets, depending on the specific data source. In general, one could summarize the difference between the first and the second source as a tradeoff between represenativeness versus accuracy. The first type captures a more representative slice of the problem, but will also include false positives. The second type accurately identifies infected machines, but only for a specific botnet, which implies that it cannot paint a representative picture.

This study draws upon three data sources: two of the first type (spam data and Dshield data) and one of the second type (Conficker sinkhole). We have access to three large, independent datasets that allow us to identify infected machines: (1) a spam trap collecting the IP addresses of spam-sending machines, which are typically infected machines; (2) data on global security incidents from the SANS Institute; and (3), a sinkhole for the Conficker botnet that logs which infected machines “call in” to receive instructions from the botnet command and control structure.

Spam Dataset

The spam data is drawn from a spam trap – an Internet domain set up specifically to capture spam, whose email addresses have never been published or used to send or receive legimitate email traffic.[6] There is no legitimate way to deliver email to the domain. All the email it receives is indeed spam – as confirmed by logging the content of the messages. In the period of 2005-2009, the trap has received 109 billion spam messages from about 170 million unique IP adresses worldwide.[7]

Spammers use thousands or even millions of infected machines in a botnet to send out spam. Of the total volume of spam messages that are being sent out everyday, the overwhelming majority is sent through an infected machine. A variety of studies published during the period under study (2005-2010) found consistently that around 80 to 90 percent of the total amount of spam comes from botnets.[8] The IP address of the machine that delivered the spam message, therefore, very likely indicates the presence of an infected machine. Previous studies have also employed the origins of spam messages as proxy data to identify infected machines.[9]

DShield Dataset

The DShield data is collected from a global network of sensors run by volunteers. This data, which has been provided to us by the SANS institute, is the same data that is used by the Internet Storm Center (ISC) for monitoring levels of malicious activity on the Internet.[10] These sensors include firewalls, intrusion detection systems, and home broadband devices that log and report 'unwanted' network traffic to the DShield database. The definition of unwanted network traffic is not that clear cut, but a simple form of it is an attempt to connect to a network from the outside that has not been sanctioned by the network administrator – such as what an Internet worm trying to propagate would do. A more advanced form is the detection of actual attacks, such as a SQL injection or Denial of Service attempts to hosts that are authorized to respond to external requests. The logs from the sensors are aggregated in the DShield database. The result is a list of 'offending' IP address each day, with the number of times each source has been reported (by the sensors), and the number of network hosts and ports it has attempted to target.

The IP addresses in the DShield dataset typically point towards infected machines – be they bots trying to propagate, or being used as attack vehicles. The DShield dataset has its limitations as well. One is the presence of false positives within this dataset, caused by what a firewall device would interpret as an attack. We have taken the step of removing from the dataset all IP addresses that are only logged once or having been reported by only one target. These logged "attacks" can also be caused by certain DNS errors or mistyped URLs, which are not indicative of a bot. This simple step has reduced the size of this dataset by around half. The other problem with the DShield dataset is that it is based on IP addresses and hence distorted by dynamic IP addressing and NATs, just like the

---

[6] This spam trap is operated by Dave Rand, CTO of TrendMicro, and was generously made available to us for several research projects, among which is this study.

[7] We have conducted extensive triangulation efforts to compare our data to the publicly available reports of commercial security providers. Most of the public data relates to the relative spam volume of countries. It turned out that the commercial reports present different numbers, sometimes quite substantially different numbers. Most of our findings are located within the range reported by the commercial security providers. This has led us to conclude that our data set provides a valid basis for the analysis, taking into account the objectives and constraints of the project. For more details, see Appendix 1 of Van Eeten et al. (2010).

[8] See Van Eeten et al. (2010, p. 20).

[9] For example, see Zhuang, L., J. Dunagan, D. R. Simon, H. J. Wang, I. Osipkov, G. Hulten and J. D. Tygar (2008). *Characterizing Botnets from Email Spam Records*. LEET '08. First Usenix Workshop on Large-Scale Exploits and Emergent Threats, San Francisco. Available online at http://www.usenix.org/event/leet08/tech/full_papers/zhuang/zhuang.pdf.

[10] See http://dshield.org/about.html for more information about the Internet Storm Center.

spam-source metrics. It is valuable, however, for capturing an different part of the botnet activities, and also for not being affected by port 25 blocking.

The DShield dataset provides an alternative perspective into the botnet universe, as many of these bots might not be used for sending out spam, and would hence not show up in the spam database. In this regards, the datasets are complimentary. For the purpose of this study, we have used the DShield data for 2009 and the first half of 2010.

Conficker Dataset

The Conficker dataset is based on log-files provided to us by the Conficker Working Group.[11] Several members of the working group run sinkholes that continuously log the IP addresses of Conficker bots. The sinkholes work in this fashion: computers infected with Conficker frequently attempt to connect to command and control servers to receive new payloads (i.e., instructions). In order to protect the botnet from being shut down, Conficker attempts to connect to different C&C domains every day. The working group has succeeded in registering some of these domain names and logging all connections made to them. Since these domains do not host any content, all these connections are initiated by bots. Therefore, we can reliably identify the IP addresses of the Conficker bots.

The Conficker dataset is unique in several ways. First of all, unlike the other two datasets, it is not a small sample of a much larger population, but rather captures the universe of its kin. This is because of the way the bot works – most of them will eventually contact one of the sinkholes. Second, this dataset is free from false positives, as, apart from bots, no other machine contacts the sinkholes. These features make the dataset more reliable than the spam or DShield datasets. The difference, however, is that the dataset is only indicative of the patterns applicable to one specific botnet, namely Conficker. Although Conficker has managed to replicate very successfully, with around several million active bots at any given moment, it has not been used for any large-scale malicious purposes – or at least no such uses have been detected yet. This means ISPs and other market players may have less powerful incentives to mitigate these infections, different from spam bots, for example. These differences make the Conficker dataset complementary to the two other sets.

Overall, the Conficker dataset adds a fresh, robust and complimentary perspective to our other two datasets and brings more insight into the population of infected machines worldwide.

Relation among the Datasets

One finding that came as a surprise to us was the low percentage of overlap among the three datasets. This is shown for January 2010 in Figure 1. We looked at the list of IP addresses in each of the spam, DShield and Conficker datasets, and counted the number of addresses that were present in more than one dataset. For the Netherlands, the overlap among the datasets is less than 10 percent and only 200 addresses are present in all three. This pattern is consistent for other months as well. For the global dataset, around 12 percent of IP addresses were part of more than one dataset in January 2010.

This figure is surprisingly low, given that it is generally assumed that bots are used for multiple purposes over time – i.e., a bot is used to send out spam at a certain point in time and to perform a network attack at another. Furthermore, it is possible for infected machines to be infected with more than one strain of malware at the same time. Putting these assumptions next to the relatively large

[11] The Conficker working group is an industry consortium to combat the effects of Conficker. See http://www.confickerworkinggroup.org/wiki/ for more information.

size of our samples, and we would expect to see a much larger number of machines to appear in multiple datasets. Two reasons come to mind for the low overlap.

One reason could be that the size of the botnet population is much larger than our samples. In January 2010, we observe around a hundred thousand bots in the Netherlands, by just looking at three datasets. What would happen if we add a fourth dataset of malicious activity, from another sinkhole, spam-trap or honeypot? Judging by the low overlap between the current data, we should expect the overall count to rise higher. And even higher with a fifth dataset, and so on. All of this strongly suggests that the estimates we are going to derive from our three datasets are conservative and significantly undercount the number of infected machines.

Another reasons could be that there is specialization among bots: a machine that is being used to send out spam is not used to perform network attacks, and vice-versa. If this were the case, out data is not a purely random selection from all infected machines, but rather has some form of systematic bias where membership in one dataset excludes membership in another. This would have interesting implications, but as far as we know such a phenomenon has never been report before. It is outside the scope of this study to test this hypothesis.

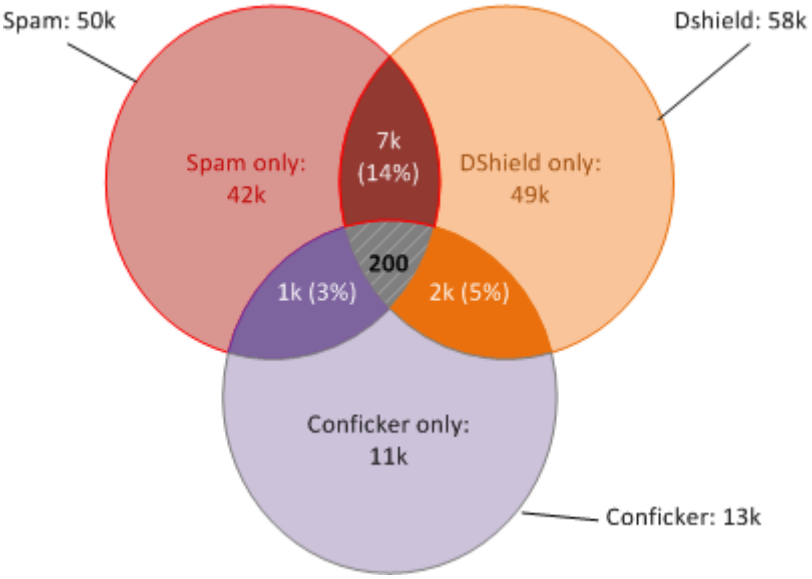

Figure 1 - Overlap of IP addresses among the three datasets for (January 2010, Netherlands).

Identifying the Location of Infected Machines

For each unique IP address that was logged in one of our data sources, we looked up the Autonomous System Number (ASN) and the country where it was located. The ASN is relevant, because it allows us to identify what entity connects the IP address to the wider Internet – and whether that entity is an ISP or not. We looked up the country of an IP address by using so-called geo-IP data, which associates IP addresses with geographical locations – in this case, we used the MaxMind geoIP database.

As both ASN and geoIP information change over time, we used historical records to establish the orgin for the specific moment in time when an IP address was logged in one of our data sources (e.g., the moment when a spam message was received or network attack was detected). This effort

resulted in time series for all the variables in the datasets, both at an ASN level and at a country level. The different variables are useful to balance some of the shortcomings of each – a point to which we will return in a moment.

We then set out to identify which of the ASNs from which the trap received spam belonged to ISPs. To the best of our knowledge, there is no existing database that maps ASNs onto ISPs. This is not surprising. Estimates of the number of ISPs vary from around 4,000 – based on the number of ASNs that provide transit services – to as many as 100,000 companies that self-identify as ISPs – many of whom are virtual ISPs or resellers of other ISPs' capacity.

So we adopted a variety of strategies to connect ASNs to ISPs. First, we used historical market data on ISPs – wireline, wireless and broadband – from TeleGeography's GlobalComms database. We extracted the data on all ISPs in the database listed as operating in a set of 40 countries, namely all 34 members of the Organisation for Economic Co-operation and Development (OECD), plus one "accession candidate" and five so-called "enhanced-engagement" countries.

This resulted in data on just over 200 ISPs (see Appendix 1). Together, these ISPs control the bulk of the market share in the 40 countries. To cross-check the completeness of our market data (as drawn from the Telegeography GlobalComms database), we compared it to the publicly available data on the total number of Internet subscriptions in each country. These public sources of data have their own shortcomings, as they rely on reports by the countries themselves, not on direct measurements. Still, if we use the 2009 OECD broadband statistics as a base of comparison, then the ISPs in our analysis account for 89 percent of the total market in the OECD.

The process of mapping ASNs to ISPs was done manually. First, using the GeoIP data, we could identify which ASNs were located in each of the 40 countries. ASNs with one percent of their IP addresses mapped to one of the 40 countries were included in our analysis. For each of these countries, we listed all ASNs that were above a threshold of 0.5 percent of total spam volume for that country.

We used historical WHOIS records to lookup the name of the entity that administers each ASN in a country. We then consulted a variety of sources – such as industry reports, market analyses and news media – to see which, if any, of the ISPs in the country it matches. In many cases, the mapping was straightforward. In other cases, additional information was needed – for example, in case of ASNs named after an ISP that had since been acquired by another ISP. In those cases, we mapped the ASN to its current owner.

Compensating for Known Limitations in Internet Measurements

Our approach allows us to robustly estimate the relative degree in which ISP networks harbor infected machines. It has certain limitations, however, that need to be compensated for. The effects of three technical issues need to be taken into account when interpreting the data: the use of Network Address Translation (NAT), the use of dynamic IP addresses with short lease times, and the use of port 25 blocking. The key issue is to understand how these technical practices affect the number of machines that are represented by a single unique IP address.

NAT means sharing a single IP address among a number of machines. Home broadband routers often use NAT, as do certain other networks. This potentially underrepresents the number of infected machines, as multiple machines show up as a single address. Dynamic IP addresses with short lease times imply that a single machine will be assigned multiple IP addresses over time. This means a single infected machine can show up under multiple IP addresses. As such, it over-represents the

number of infected machines. Both of these practices counteract each other, to some extent. This limits the bias each of them introduces in the data, but this does not happen in a consistent way across different networks.

This is a classic problem in the field of Internet measurement: how many machines are represented by a single IP address? Ideally, one IP address would indicate one machine. But reality is more complicated. Over an extended time period, a single address sometimes indicates less than one machine, sometimes more than one. This varies across ISPs and countries. Earlier research by Stone-Gross et al. (2009) has demonstrated that in different countries, there are different ratios of unique IP addresses to infected machines – referred to as "churn rates".

We have two ways to robustly control for the potential bias that these churn rates introduce in our data. First, for the spam dataset, we look at the volume of spam in addition to the number of unique sources. If there are several machines behind a single IP address, the spam volume is also several times higher that of a single machine. If there is one machine behind many IP addresses, the spam volume is proportionally lower for each address. We have calculated the ratio of unique sources to spam volume in our data. The Spearman correlation between the churn rates reported by Stone-Gross et al. (2009) and ratios we calculated is very high, namely 0.88. This resemblance suggests that spam volume can indeed control for churn.

A second way to control for the bias caused by churn rates can be applied to all three datasets; namely, to use shorter time scales when counting the number of unique IP addresses in a network. On shorter time scales, the potential impact of churn is very limited. Earlier research found that churn starts to affect the accuracy of IP addresses as a proxy for machines on timescales longer than 24 hours.[12] We therefore worked with a time period of 24 hours. All our comparative analyses are based on the daily average number of IP addresses from an ISP network. This compensates for churn, but has a downside: in these estimates, the number of infected machines is now grossly undercounted. This is because our spam and DShield datasets are samples. They obviously do not capture all spam or attacks worldwide. Since they capture a sample, it means that measuring over short time periods dramatically increases the odds of missing sources of spam and attacks. While the number of bots measured in a 24 hour period is the most reliable for comparisons across networks, it cannot indicate the actual infection rate of a network in absolute terms. For absolute estimates – in other words, of the actual number of infected machines – we use larger time periods, depending on the situation: months, quarters or even the whole 18-month measurement period.

Another limitation, relevant only for the spam dataset, is the use of port 25 blocking by ISPs. The effect of port blocking is that infected machines can no longer directly send email to the wider internet, but have to go through the ISP's outgoing email servers. This affects both the number of sources as well as the spam volume. The ISP's network may harbor thousands of infected machines, but they can no longer reach the spam trap directly and thus do not reveal their IP address through spam distribution. There is one important way in which the attackers themselves compensate for this problem: when the bots notice they cannot connect anymore via port 25, they start to redirect their spam through the ISP's official outgoing email servers. In various cases where port blocking was introduced, we saw that it led to a brief reduction of outgoing spam, only to return to the previous spam volume within about a month. It is difficult for the ISP to prevent this from happening, as each bot sends out a relatively low level of spam, and thus rate limits and similar controls often do not pick up on it. This adaptation of the spam bots allows us to use spam volume to cross check our

[12] See Moore et al. Moore, D., C. Shannon and J. Brown (2002). *Code-Red: a case study on the spread and victims of an Internet worm*. Proceedings of the 2nd ACM SIGCOMM Workshop on Internet measurment. Available online at http://portal.acm.org/citation.cfm?id=637244.

findings. It is not perfect, however. Port blocking is an unavoidable limitation to our data. If the spam volume remains consistently lower, port blocking obscures the presence of infected machines. That being said, the effect of the bias is not wholly unreasonable. The ISPs that adopt port blocking improve their ranking in terms of botnet activity compared to those that don't – which is not without merit, given that the measure of port blocking is part of many guidelines on best security practices for ISPs and that it cuts into the criminal business model of spammers, one of the revenue streams for botnets.

For all the analyses we discuss in this paper, we have always checked whether the pattern we found persisted across all different metrics: spam sources, spam messages, DShield sources, and Conficker sources. For sources we have checked both the daily average number of unique IP addresses, and the total number of unique IP addresses for that particular metric. That way, we can compensate for the various measurement issues. Patterns that hold across these different measurements can be said to be robust and valid. These measurement issues are revisited in more detail in each section of the findings chapter.

The result of this approach is time series data on the number and the location of infected machines across countries and ISPs. We have paid special attention to whether these machines are located in the networks of the main ISPs in the wider OECD. With this data in hand, we can now turn to answering the research questions.

3.
Findings

Infected Machines in the Netherlands

We shall start off by benchmarking the Netherlands – in terms of infected machines – to other countries. This will be done using metrics based on each of our datasets.

Figure 2 shows a scatter plot based on the spam dataset. Each point in the graph represents a country, with the Y-axis showing the number of IP addresses seen as sending out spam at any point during 2009 (see Appendix 1 for the list of countries included in the analysis and the country codes). Each of these IP addresses represents, with about 90 percent likelihood, an infected machine.[13] This is plotted against the number of Internet users in that country, on the X-axis. Both axes are on a logarithmic scale.

An initial observation is that the points in this scatter-plot reveals a more or less linear relation: countries with more Internet users have more infected machines. Although this may be unsurprising, it highlights the point that the number of users going online is the driving force behind the increase in the number of infected machine. On the other hand, we also see significant variantion among countries with the same number of Internet users. As an example, in the case of Finland (FI) and Chile (CL), there is an order of two magnitudes or a hundred-fold difference in the number of spam-sending bots. If you imagine a diagonal line going from the lower left to the upper right, countries below this line are doing better than average and countries above the line, worse than average. Based on the spam-sources metric the Netherlands falls closely on this line, indicating it to be in the middle of the countries covered in the study – that is, the Netherlands has an average number of infected machines.

There are several notes to keep in mind when interpreting this scatter-plot: foremost, it shows all spam-sources within the countries, be they in ISP-networks or other networks – e.g., large enterprises, university networks and hosting providers. That is why on the X-axis we have used the total number of Internet users in that country, not just the number of broadband subscribers.[14] Measurement limitations are another factor to bear in mind, as discussed in the previous chapter. These shortcomings can be compansated for, to some extent, by looking at another metric: spam volume.

Figure 3 shows the spam volume – that is, the number of spam messages sent out from each country plotted against the number of Internet users in the same country, both on a logged scale. This metric has a direct relationship with the number of infected machines, and their level of activity, and is hence relevant for our comparison. It compensates for port 25 blocking to some degree, as most spam-bots reroute their spam messages through the ISPs own email infrastructure and relay servers,

[13] For a more detailed discussion about the reliability of using spam sources to identify infected machines, see Chapter 2 and also Van Eeten et al. (2010, pp. 19-23).
[14] Numbers of Internet users are based on reports by the International Telecommunication Union for 2009.

and although their IP address is concealed, their presence can be detected via the volume of messages. Another opportune feature is that the metric is not influenced by dynamic IP addressing: the changing of the address of an infected machine will not affect the amount of messages sent. And it even includes to a certain extent in itself the number of infected machines behind one shared connection - as the total amount of spam sent out will increase if there is more than one infected machine.

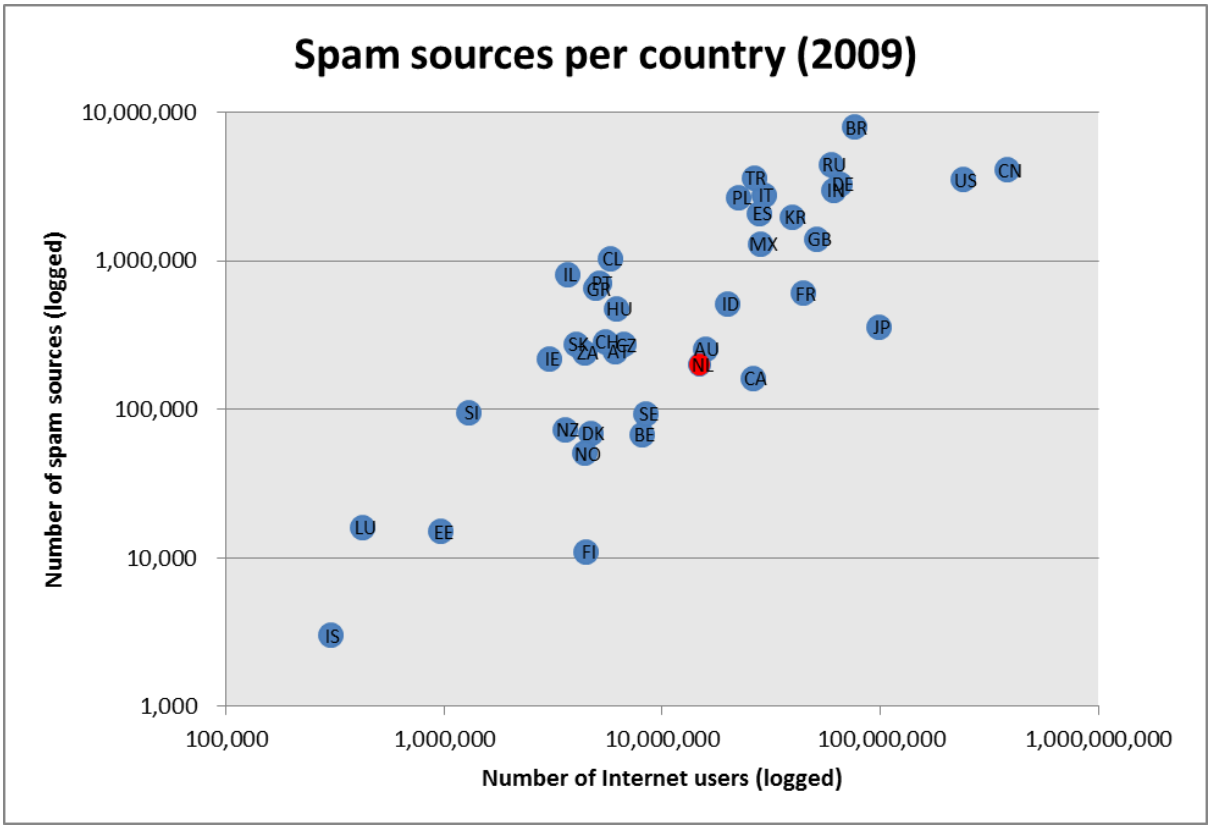

Figure 2 - Number of spam sources versus Internet users per country (logged-scale)

Given the differences between the two metrics, the overall pattern maintains a remarkable similarity. A linear relation can again be observed between the numbers of emitted spam messages and Internet users. And so can the near hundred-fold difference among countries of similar size. Imagining the diagonal regression line again, some countries such as Slovenia (SI) and Germany (DE) shift position, in this case to the lower triangle that indicates a better than average performance (i.e., a lower number of spam-messages than the average for their size). Interestingly enough, the Netherlands stays in the same position in terms of this performance metric, which is in the middle of the pack.

It should be noted that the spam messages metric has its own shortcomings. It is affected by the average Internet speed available in each country, and the amount of time people stay online, with countries having lower Internet speeds or where people spending less time online emitting less spam messages compared to a counterpart with the same number of bot infections. This would give the impression that these countries have a fewer infected machines than they actually have.

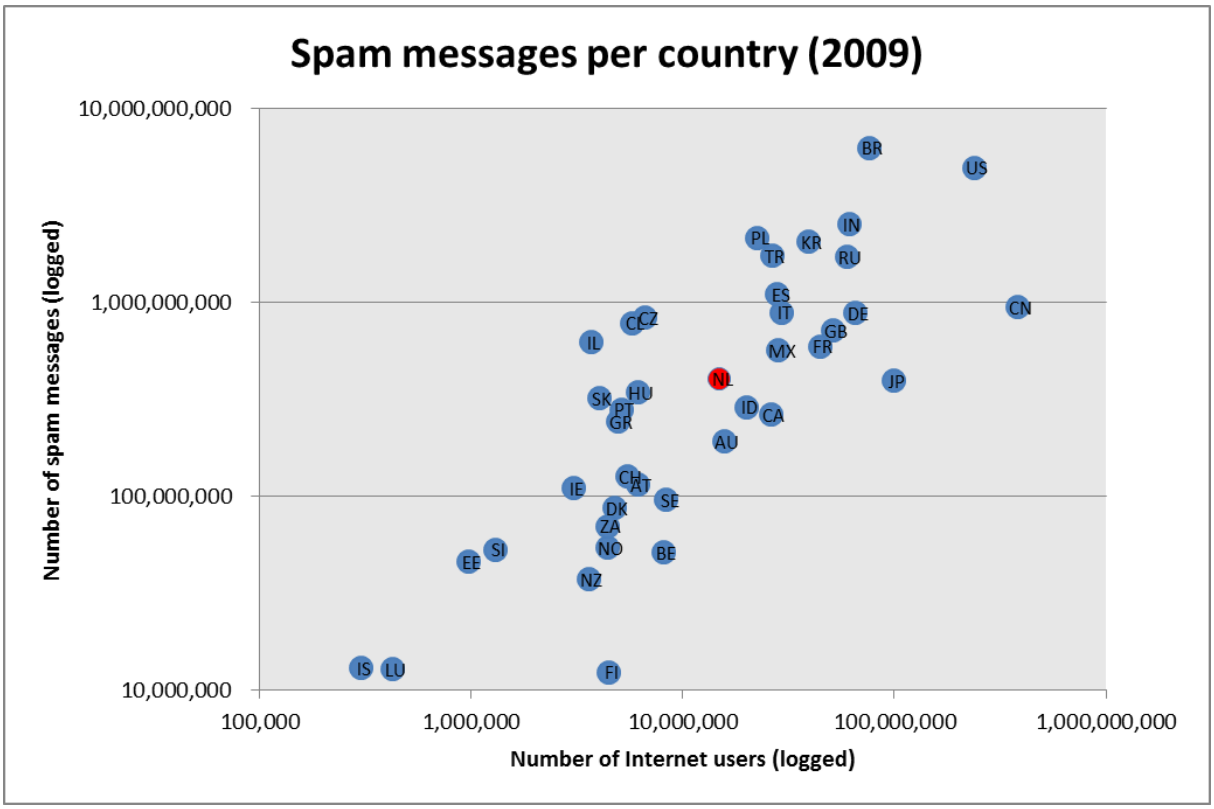

Figure 3 - Number of spam messages sent versus Internet users per country (logged-scale)

The third metric that we use is presented in Figure 4. In this graph, similar to the previous two, each point represents a country with the X-axis indicating its count of Internet users. On the Y-axis this time we have the number of IP addresses seen for that country in the DShield dataset. Both axes are again logged. To recap, each IP address in the DShield dataset is indicative of a network attack traced back to that particular IP address. These network attacks can have different forms, including denial of service attacks, port scans in attempts to infect other hosts, etc, and have been logged by various entities worldwide. Since most of these attacks would logically originate from within botnets – so that the main attacker cannot be identified – we can associate each DShield source with an infected machine.

The overall trend here basically reiterates what we have already seen: a linear relationship exists between the number of DShield sources and the number of Internet users in any country, pointing to the fact that the number of users going online is the major driving force behind increased bot infections worldwide. Again, we also see the spread among countries and the orders of magnitude difference in performance between countries with a similar size of the online population. In this graph we also observe that Netherlands lies in the middle of the pack – on the diagonal line that would split the sample into above and below average performers.

The final country-level comparison is presented in Figure 5. This is based on the Conficker sinkhole data, with the Y-axis showing the number of unique IP addresses that have logged into the sinkhole from each country. Conficker bots contact the sinkholes for instructions, believing that they are C&C servers. The servers log these access attempts. As in the previous cases, the axes are on a logged scale, and the X axis is the number of Internet users in each country. Strikingly enough the overall pattern is again the same: a linear relation between the number of bots and the number of Internet users, but spread over a spectrum.

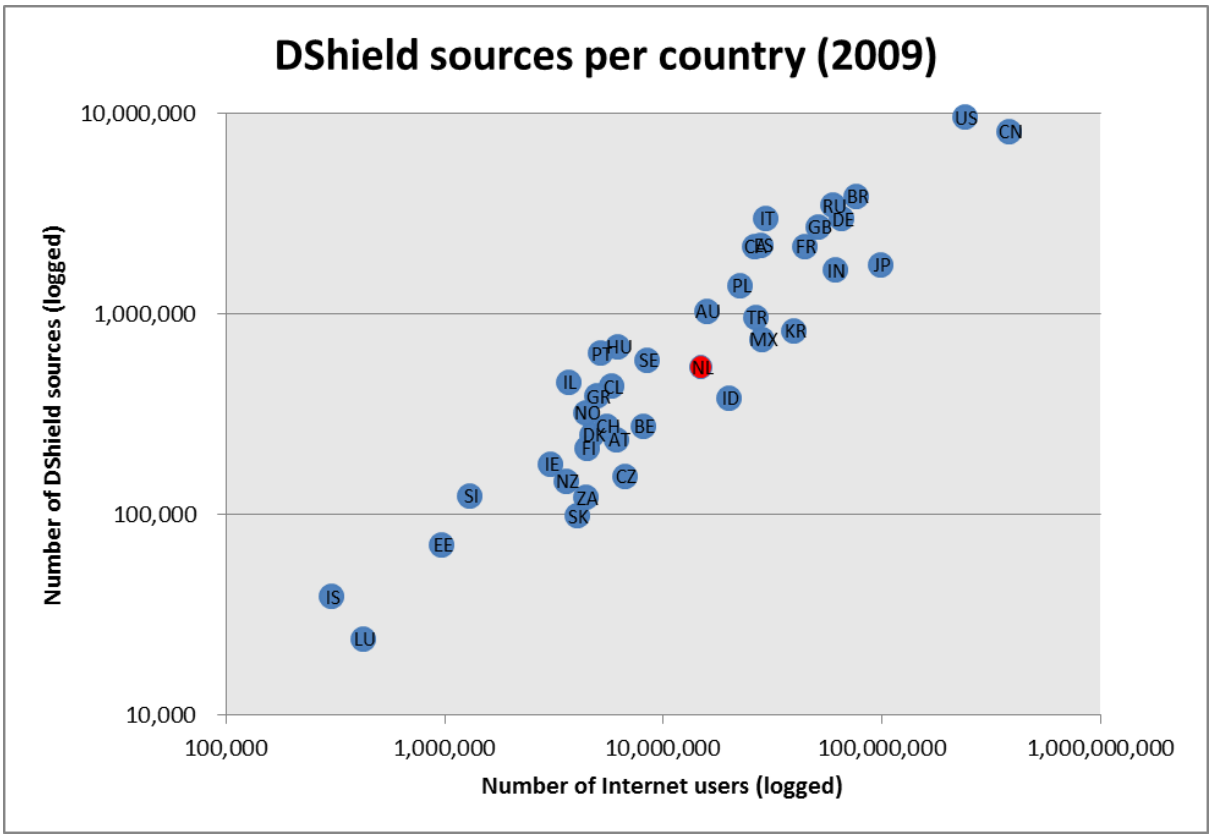

Figure 4 - Number of DShield sources versus Internet users per country (logged-scale)

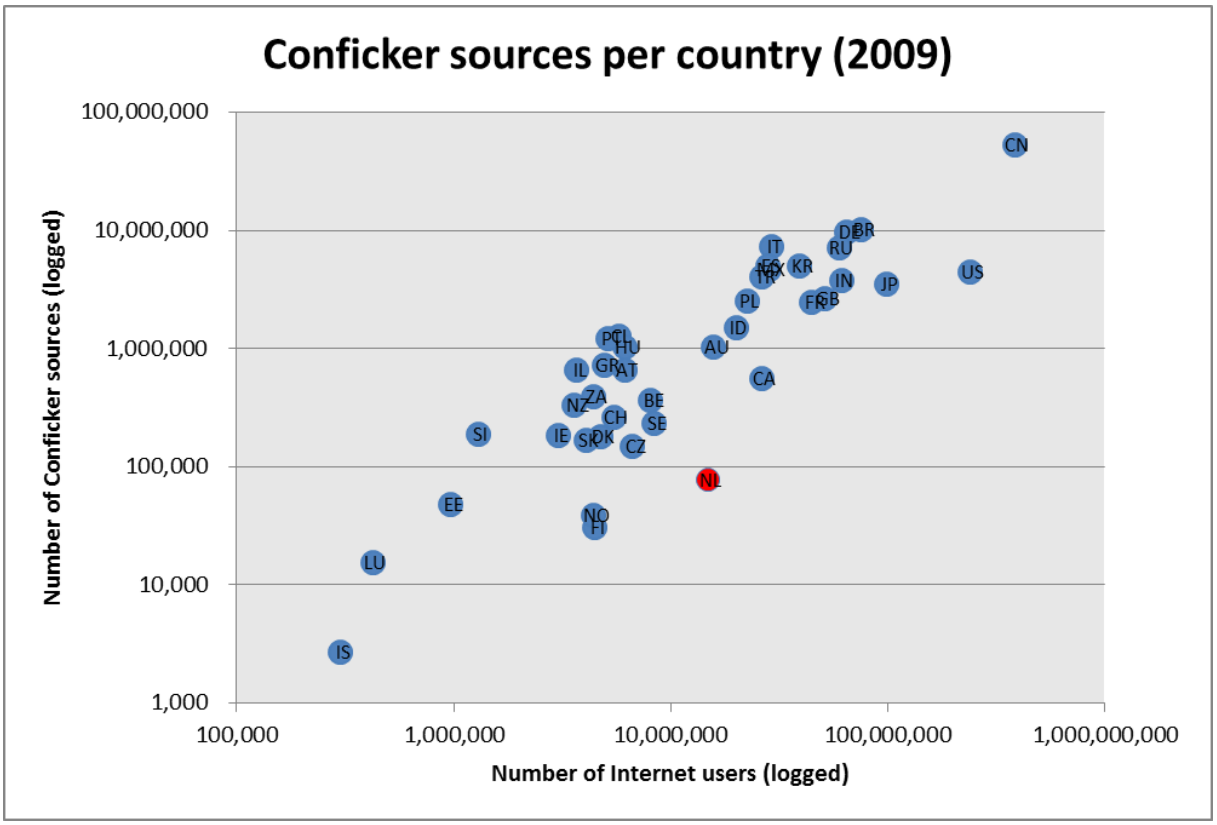

Figure 5 - Number of Conficker sources versus Internet users per country (logged-scale)

As mentioned in chapter 2, this dataset has the advantage of being free from false positives. This metric also suffers from the limitations of correlating IP addresses with infected machines (e.g., dynamic IP addressing and NAT). It is nevertheless an interesting dataset and as it captures an "internal" perspective of a botnet.

The ranking of the Netherlands is similar in the spam and Dshield datasets, but very different for the Conficker data. Considering the diagonal line that splits the points into the lower-right and upper-left triangles, the Netherlands falls clearly into the lower triangle. In other words, it has a much lower number of Conficker infections than expected for a country of its size. Based on this metric, the Netherlands is doing pretty well. This begs the question of why it performs so much better for Conficker. We are not sure why this is the case. There are two plausible hypotheses. First, Dutch ISPs and other network operators have taken Conficker more seriously than many of their peers in the rest of the world. We know, for example, that one of the largest Dutch ISPs actually uses the sinkhole data to contact infected customers and, when needed, quarantine them. This is unusual for consumer market ISPs. The second hypothesis is that the Netherlands has relatively high patching levels for the Microsoft Windows platform. Microsoft has distributed a patch for the vulnerability that is exploited by Conficker since October 2008. We suspect both of these hypotheses combine to produce the remarkably low infection level of the Netherlands.

In summary, the four country-level metrics – *spam sources, spam messages, DShield sources,* and *Conficker sources*, rank the Netherlands as average among the 40 countries studied, or doing better. The actual ranks will be given in the next section.

Ranking countries

In Figure 6 and Figure 7, we use the data presented above but represent it in a different format. The metrics are presented in bar charts, ranked by the relative number of infections in each country. Relative metrics, which enable ranking, are calculated by dividing the absolute metrics by the number of Internet users. Additionally, to correct for the over-counting caused by the use of dynamic IP addressing in some countries, we use daily averages. That is, we look at the average value of each metric (e.g., number of spam sources) per day over the whole period, instead of the total value over the period – see Chapter 2 for more details.

It goes without saying that the daily averages are much lower than the totals. Over the course of one year, some infected machines get cleaned, while others become infected. A second relevant factor reducing the count is the effect of sampling. Most measurement methods capture only a sample of the active machines. The shorter the time frame of the measurement, the lower the odds are that an active machine shows up in the sample. As an example, daily average number of spam sources in the Netherlands for 2009 is around 4,000 and the total number is around 200,000 for the same period. Using daily averages compensates for differences in the dynamic IP addressing assignment policies across countries, creating a more meaningful comparison. However the other limitations inherent in each dataset that were mentioned in the previous section still remain.

It should also be noted that a there is a slight mismatch between the numerator and denominator of the relative metrics. The numerator in cases where it is a count of IP addresses is indicative of Internet subscribers while the denominator is based on Internet users. Multiple Internet users often share one subscription and IP address in households or small companies. In countries where such sharing is higher on average, the relative metrics would have a lower value than their counterparts with the same number of infected bots. One remedy would be to use the number of Internet subscribers in each country as the denominator, but such a figure is not available reliably at the country level.

The results of the rankings are as follows: the Netherlands ranks 15th out of 40 based on the metric spam sources per Internet user; 20th out of 40 based on the metric spam messages per Internet user. It ranks 18th based on the metric DShield sources per Internet user, and finally, its rank is 2nd based on the Conficker source per Internet user metric. In all these metrics, lower ranks are better from a security perspective. The overall conclusion is that the Netherlands ranks slightly above better than average in the group of 40 countries with regard to the number of infected machines located both within ISP and non-ISP networks. That said, a variety of countries do better than the Netherlands, except for Conficker, where only Finland performs better.

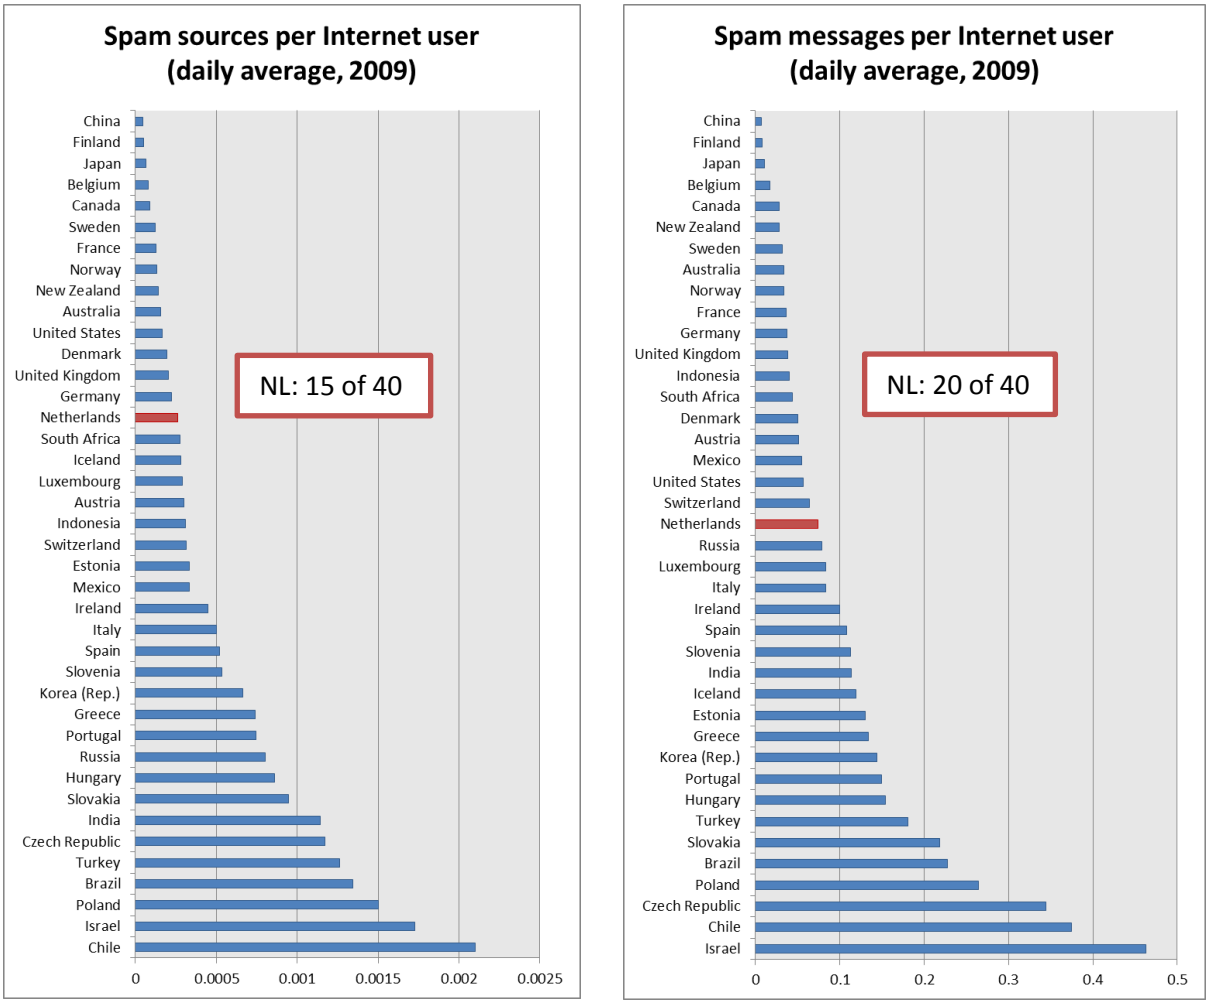

Figure 6 – Spam messages and sources per Internet user by country (daily average 2009)

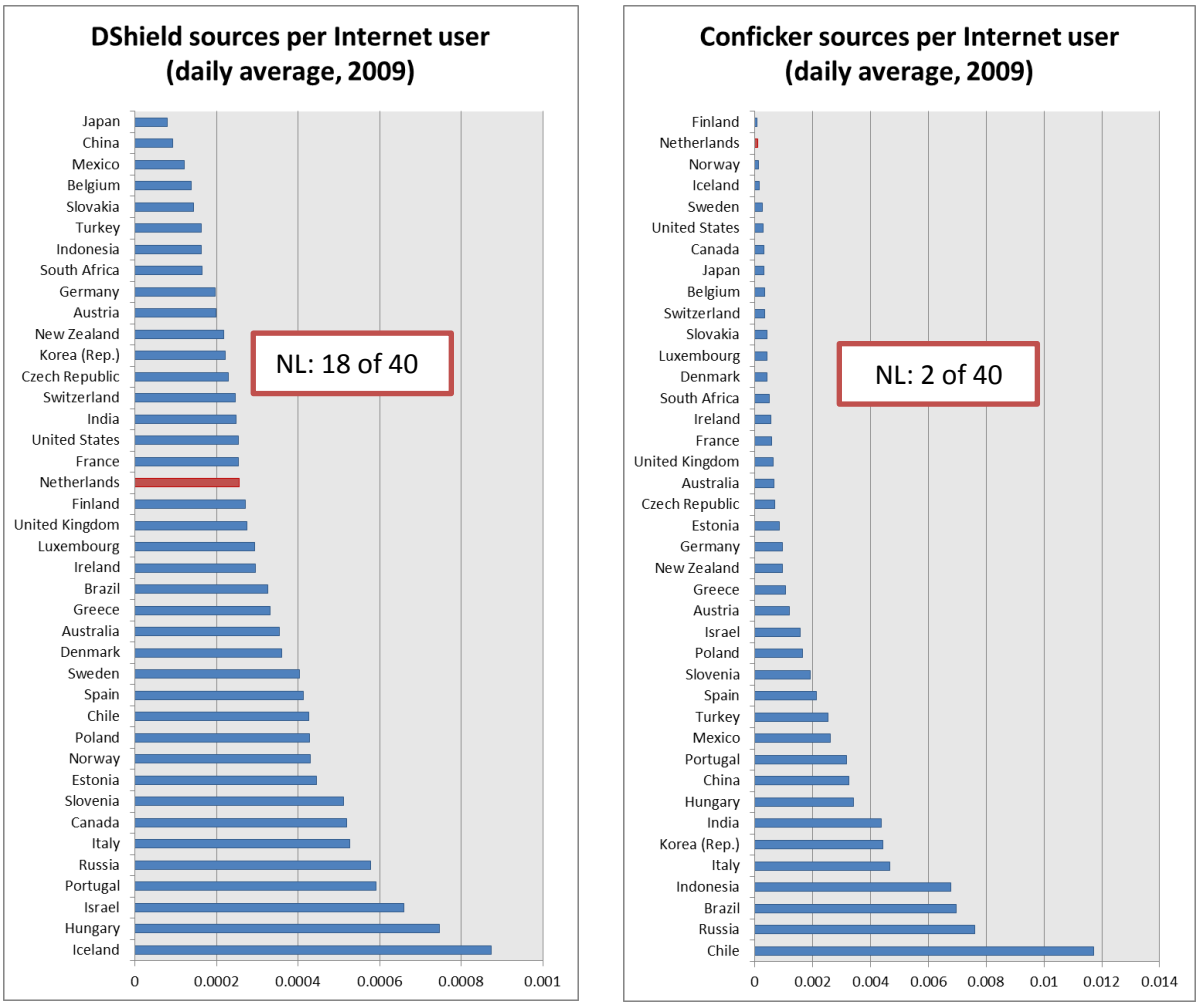

Figure 7 - DShield and Conficker sources per Internet user (daily average 2009)

Infected Machines in ISP Networks in the Netherlands

Are ISPs control points for infected machines?

To answer the question of whether or not ISPs are important control points for infected machines, we look at the portion of all infected machines that are located in ISP networks. For each IP address associated with an infected machine, we look at the type of organization administrating the 'autonomous system' that the IP address of this source belongs to – and whether it is an ISP or not. An ISP in this context is an entity providing Internet access to third parties. Non-ISP entities include enterprises, hosting providers, universities, governments, etc. We count the number of sources within each of these two groups to determine the extent to which these infected sources are located inside ISP networks.[15]

The results are presented in the following figures: Figure 8 shows that on average, across the 40 countries, 80 percent of infected sources are located in networks administrated by well-known ISPs.

[15] The process of determining the type of organisation that administrates an autonomous system needs to be conducted manually as this data doesn't exist in any Internet registry. The manual process is explained in the section "identifying the location of infected machines" of the Research Approach chapter.

This ratio holds in all three datasets that we have investigated. At a country level, these figures will vary based on country specific attributes.

Figure 9 presents these ratios for the Netherlands: 76 percent of spam sources, 84 percent of DShield sources, and 68 percent of Conficker sources are located within the networks administered by the major Dutch ISPs surveyed in this study. Based on these percentages we can conclude that in the Netherlands ISPs are indeed potential control points for infected machines.

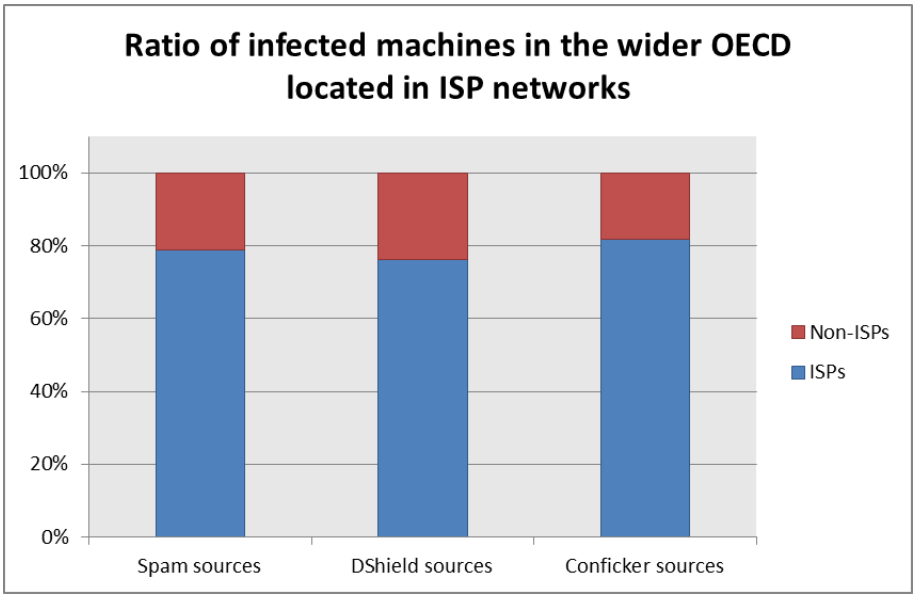

Figure 8 - Ratio of infected machines in the wider OECD located in ISP networks

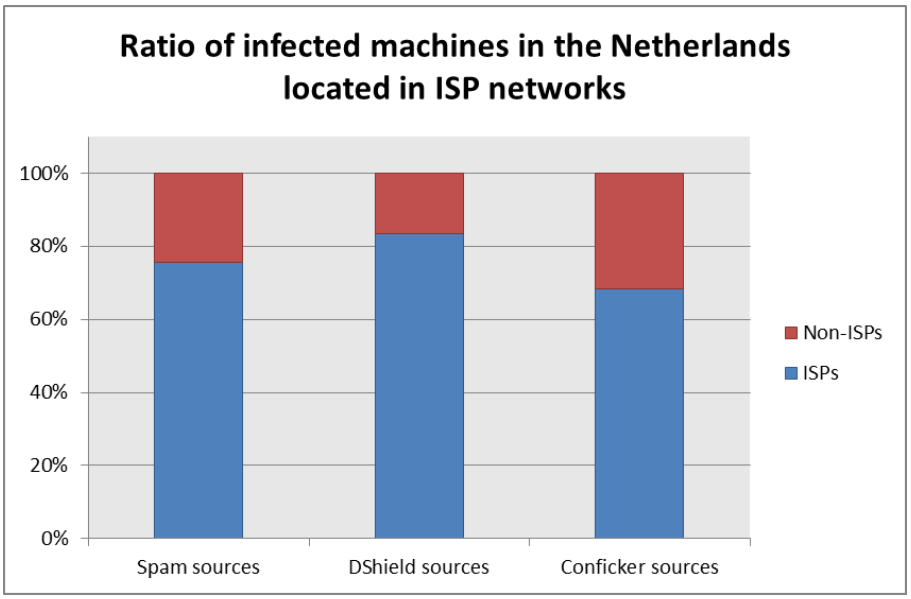

Figure 9 – Ratio of infected machines in the Netherlands located in ISP networks.

Absolute number of infected machines

We estimate that between January 2009 and June 2010, around 450-900,000 infected machines resided in the networks of the Dutch ISPs surveyed for this report. Another way to state this finding, is to say that between 5-10 percent of all Dutch broadband subscribers have suffered an infection in

2009 – and the same will probably hold for 2010. Because of the steps undertaken to derive the estimate, this figure is conservative and represents a lower-bound estimate.

We started by creating a list of all the distinct IP sources present in the spam, DShield and Conficker datasets during the 18-month period. The count of IP address in this list is around 1.1 million, all of which are unique. Addresses that where present in more than one dataset are counted only once. Of this total, as shown in Figure 10, 891,192 IP addresses are located in the networks of the Dutch ISPs.

During our discussions with the Dutch ISPs, it was stated that dynamic IP addressing is implemented in their networks in such a way that subscribers usually maintain a single IP address for very long time periods. In other words, DHCP churn is not a significant factor, according to the ISPs. The approximately 891,000 IP addresses can thus been treated as directly indicating the number of infected subscribers.

However, since our goal was to obtain a lower bound estimate with a high degree of confidence, we also wanted to include a more conservative estimate. For this reason we divided the total number of IP addresses by two, assuming a "churn rate" of 2.[16] In other words, each subscriber on average changes his IP address once during this 18 month period, and hence if a subscriber will be infected twice or will remain infected for more than a year, he will have been counted twice for the two different logged addresses. This leaves us with a figure of approximately 446,000 infected subscribers. Each infected subscriber might have several infected machines at his premises, but as we are interested in the lower bound estimate, we assume that there is only one. After rounding the number, we are left with the estimate of 450-900,000 infected machines in the Netherlands for the 18-month period.

There are certain limitations to the accuracy of this estimate. On the one hand, the count might be too high, considering the number of false positives in the Spam and DShield datasets, i.e., non-botnet spam sources and non-botnet attacks. However, this over-counting is easily offset by the fact that we are only counting the IP addresses of bots in three samples. As we mentioned in chapter 2, we expect the overall count of bots to rise much higher if we add other datasets of malicious activity. Hence we are undercounting rather than overcounting. All in all, we can assert that our figures presents conservative, lower-bound estimates of the number of infected machines.

One important issue to keep in mind is that this estimate is the total number of infected machines for this period. Not all these machines will have been infected at during the same time period. For a specific day or month, there will be fewer infected machines and the measurements will also capture fewer of those that are active. Rather, many machines are disinfected after a certain amount of time, and new ones become infected during the same period. As an example, the estimated number of infected machines located in Dutch ISP networks during the month of June 2010 was around 68,000.

[16] This assumption regarding churn is based on a paper about the Torpig botnet, that provided a unique insight into the relationship between IP addresses and actual machines. That study found that the Netherlands, as a whole, had a churn factor of 1.7. The feedback we got from ISPs suggest that this is not representative for their networks, but we felt that it still provided a basis for a more conservative, lower-bound estimate. Since we are measuring over longer time periods that the Torpig study, we have increased the churn factor to 2. For more detais see: Stone-Gross, B., M. Cova, L. Cavallaro, B. Gilbert, M. Szydlowski, R. Kemmerer, C. Kruegel and G. Vigna (2009). *Your Botnet is My Botnet: Analysis of a Botnet Takeover*. 16th ACM Conference on Computer and Communications Security, November 9–13, 2009, Chicago, Illinois. Available online at http://www.cs.ucsb.edu/~seclab/projects/torpig/torpig.pdf.

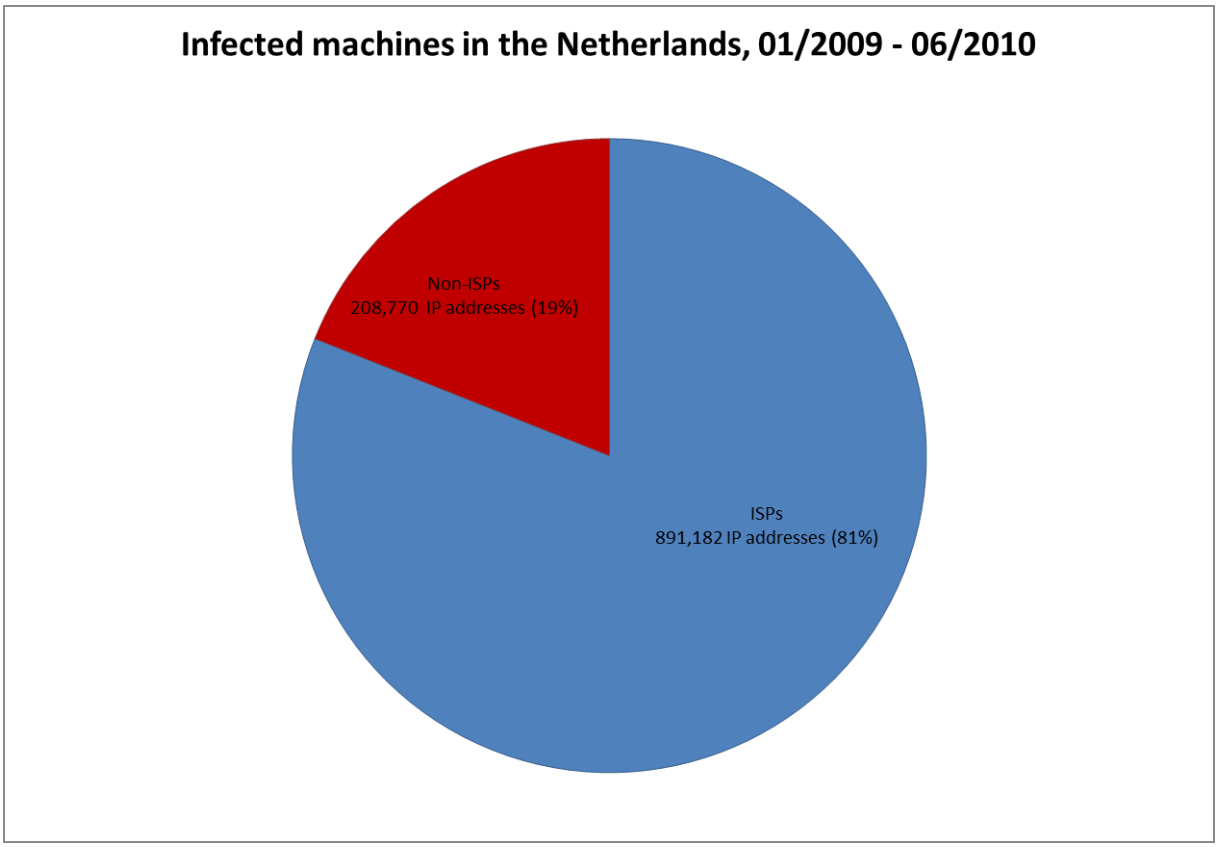

Figure 10 – Number of infected machines in the Netherlands (January 2009 - June 2010)

Our estimate corresponds remarkably well with the figure presented in the 2010 Microsoft Security Intelligence Report (MSIR). The MSIR reported around 200 thousand bots in the Netherlands during the first six months of 2010.[17] Our method yields 260,000 bots for the same period. Microsoft's approach also implies a conservative estimate of the size of the botnet population, as they only include machines that are running Windows updates (which is 600 million machines worldwide) and only check for the presence of the most wide-spread bot infections.

The next figures depict how these numbers are distributed over the various Dutch ISPs. Figure 11 is similar to Figure 10 but the pie chart is broken down by ISP. Again, the numbers are the counts of unique IP addresses across all datasets, corrected for overlaps but not for DHCP churn. Figure 12 presents the number of infected machines per ISP for the month of June 2010. Each bar presents the number of machines in a particular ISP, with the number of IP address counts in each individual dataset shown separately. Please note that neither of these figures should be used as a benchmark between the ISPs, as the numbers are in absolute terms and, as we have seen, larger ISPs have a higher number of infections by virtue of having more customers. Any comparison need to take this into account and use relative infection metrics. We return to this issue later in this chapter.

[17] See p.39 of Microsoft (2010). *Microsoft Security Intelligence Report, Volume 9*. Microsoft. Available online at http://www.microsoft.com/security/sir/.

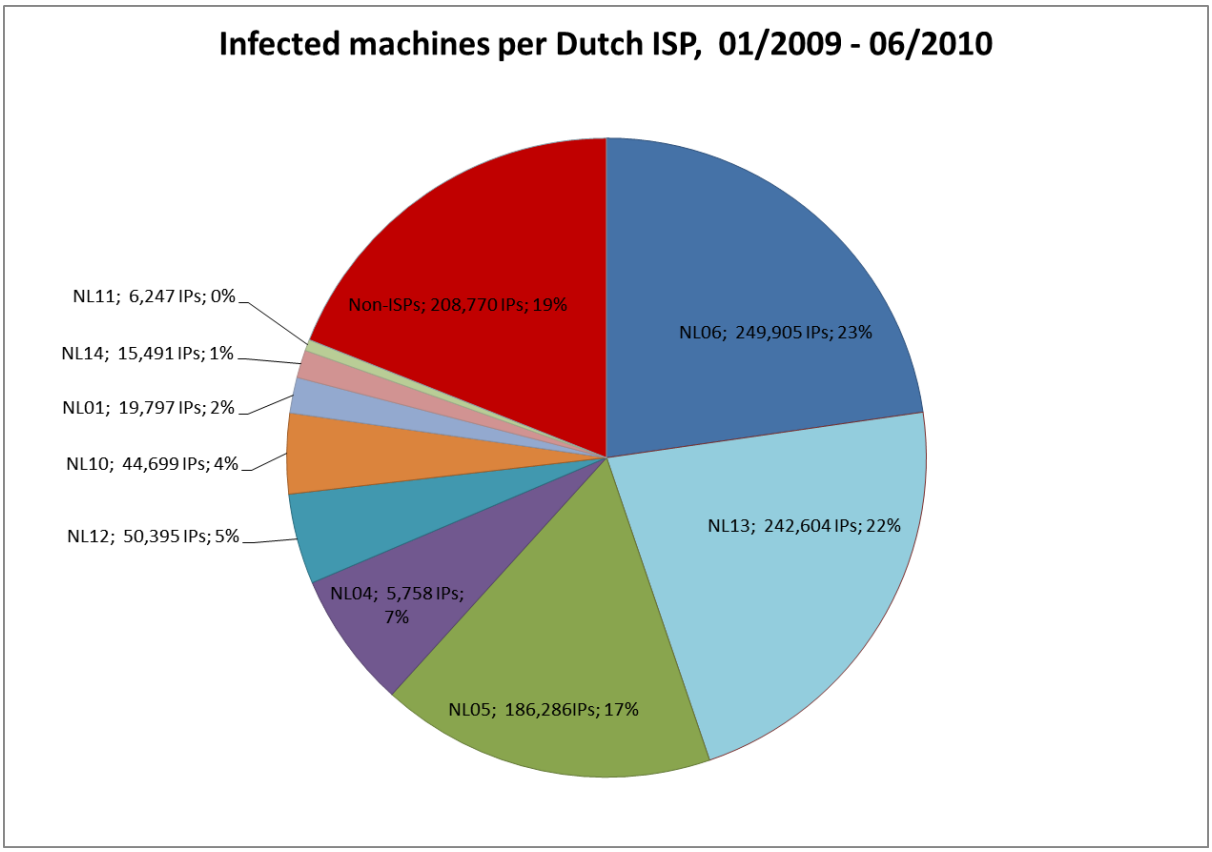

Figure 11 – Number of infected machines in the Netherlands - IP address count per ISP (January 2009 - June 2010)

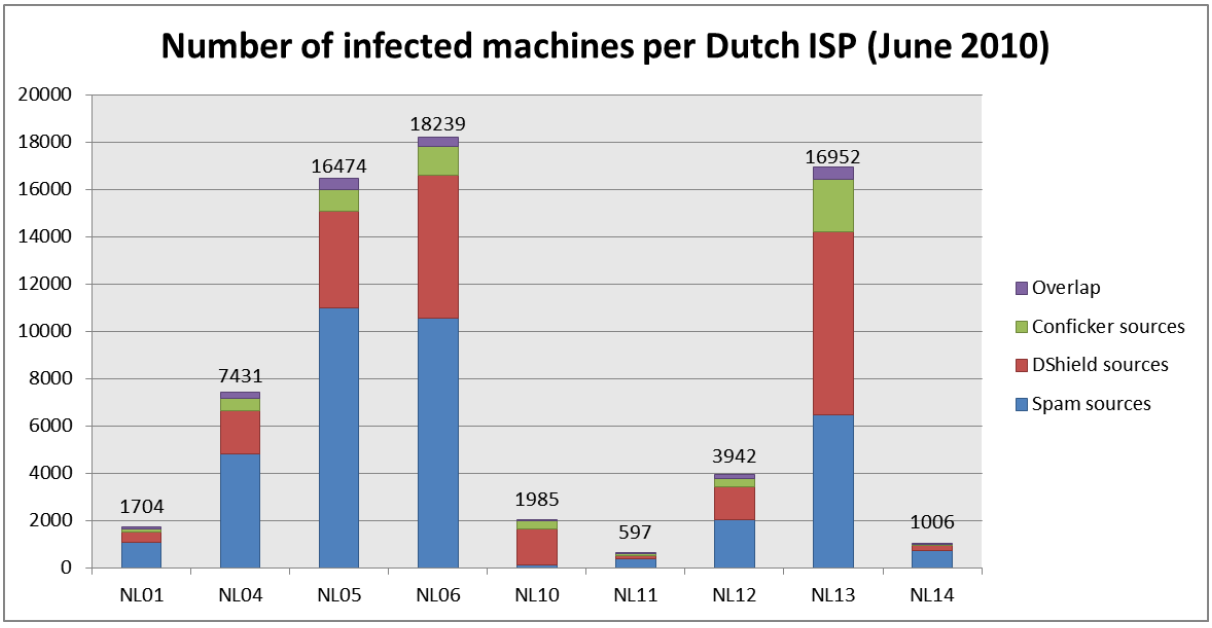

Figure 12 – Number of infected machines in the Netherlands, per ISP (June 2010)

Time Trends

Figure 13 shows how the number of infected machines in each of the ISPs' networks has changed over time. Such a perspective provides insights into whether the situation regarding botnets in the Netherlands is improving or deteriorating. It also allows a first assessment of the effects of events such as the signing of the anti-botnet covenant among the ISPs on the overall pattern. The observations reveal that the overall pattern fluctuates but without any clear upward or downward trend, even after the signing of the covenant in September 2009. What may perhaps be surprising is that the fluctuations in all the ISPs are happening in a more or less synchronized manner.

How can we explain the similarity in the time trends of all the ISPs? The most obvious answer is that the trends are driven by attackers. This is corroborated by the fact that, as we have found through the workshops, the ISPs have not changed their security policies drastically, or have had any significant shift in their subscriber base during the plotted 18-month period. We would hence not expect the ISPs' actions to produce the kind of fluctuations visible in the time trends, let alone in such a synchronized manner. Since the overall pattern shows considerable fluctuations while these internal factors remain more or less stable, the changes must be driven externally, i.e., by attackers, especially when the synchronised manner of the fluctuations across all the ISPs is taken into account. This hypothesis is supported by the trends that Microsoft recently reported, with a steep increase in the first quarter in 2010, followed by a drop in the second quarter, but not back to level of the second half of 2009.[18]

The attackers' actions includes actually infecting new machines, but also 'activating' sleeper zombies? In the latter case, then we are actually not seeing fluctuations in the number of infected machines, but in the number of them that are active. This would mean that the number of bots is actually closer to the maximum of this graph. On the other hand, if we instead accepted the first case, the fact that the percentage of new infections is so similar in all the ISPs would imply that there is a fixed percentage of the population that is always 'at risk', whatever the malware strain. This would make sense if this part of the population consistently employs sloppy online habits, such as not patching, not running anti-virus software, etc. These are all far-reaching conclusions based on the time-trend graphs, and further research to verify these conclusions is required and recommended.

To further understand developments over time, we have broken down the trends for the individual datasets in Figure 14. What can be seen is that during the surveyed period, the number of Conficker sources (per quarter) has remained more or less stable. The number of spam sources was steady for the first three quarters of 2009 but increased from the fourth quarter onwards. The number of DShield sources had the most variability and the overall trend we see is strongly influenced by variations in the DShield data.

[18] See the global figures at: http://www.microsoft.com/security/sir/threat/default.aspx#botnetcmm

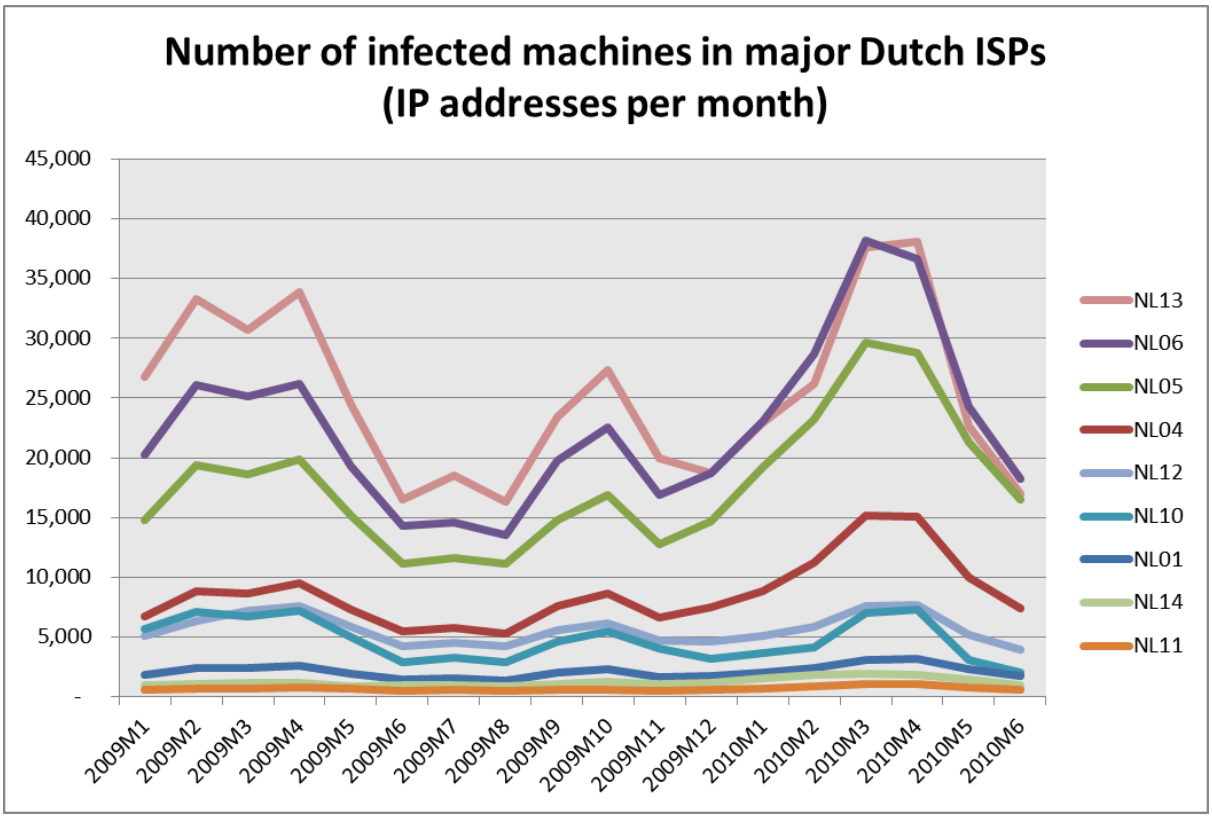

Figure 13 – Time trend of the number of infected machines in Dutch ISPs (IP addresses per month)

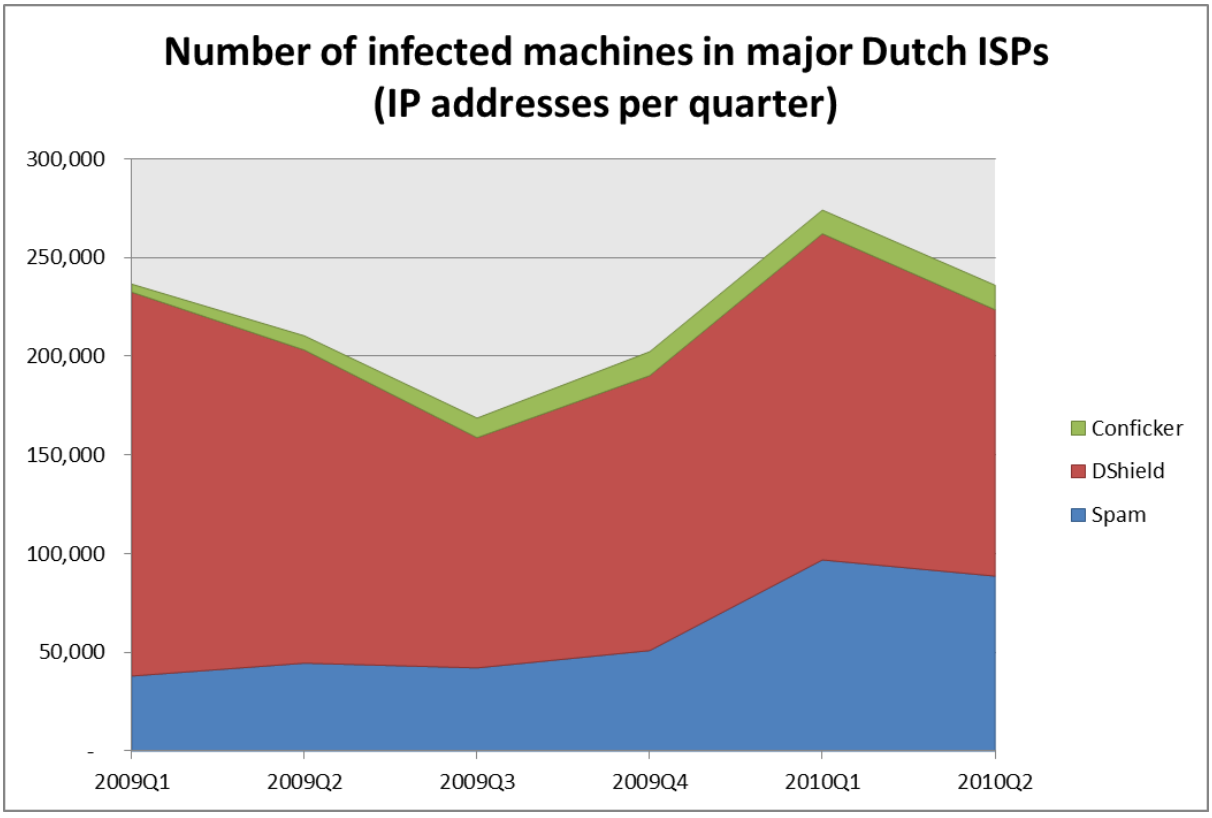

Figure 14 – Number of infected machines in major Dutch ISPs by dataset (IP addresses per quarter)

Both of the above graphs were based on total unique IP addresses seen per month or quarter. It is interesting to also plot the same graphs based on the daily average number of sources seen over the intervals, and compare the trends. Figure 15 shows the number of infected machines active on average each day, in the major Dutch ISPs and during the 18-month period. We see some similarities to Figure 13, such as synchronised fluctuations in the number of bots across all the ISPs.

There is however one key difference between the two patterns – in the daily average graph, a clear and steep rise in the numbers of bots occurs in 2010. But how can this difference be explained? Let's start by reviewing the metrics: the unique IP addresses are the number of bots seen at any point over the time period. Whether a bot is seen once in a month, or every day, it generates the same +1 increase on the metric. [19] In the daily averages on the other hand, the amount of activity of a single bot has a much larger effect than for the monthly or quartely counts. Thus, the clear upswing is probably not only caused by a growing number of bots, but also by these bots becoming much more active. The breakdown of the daily average trend by dataset (in Figure 16) shows that spam bots are the main source of this increased activity.

When comparing Figure 16 and Figure 14, we realize that the ratio of the DShield sources to the other sources is larger in the monthly graph than in the daily average graph. This means that the activities of DShield sources are much sparser over each period. In other words, the activities of the bots being used for network attacks are witnessed less frequently than those of the spam-bots. It would be interesting to investigate the reason of this pattern.

We can summarize this section by stating that: (a) the overall botnet problem in the Netherlands has worsened in 2010; and (b) this change appears to be driven by attackers.

[19] Please note that this explanation only holds in ISPs where DHCP churn is very low. This is the case in all major Dutch ISPs.

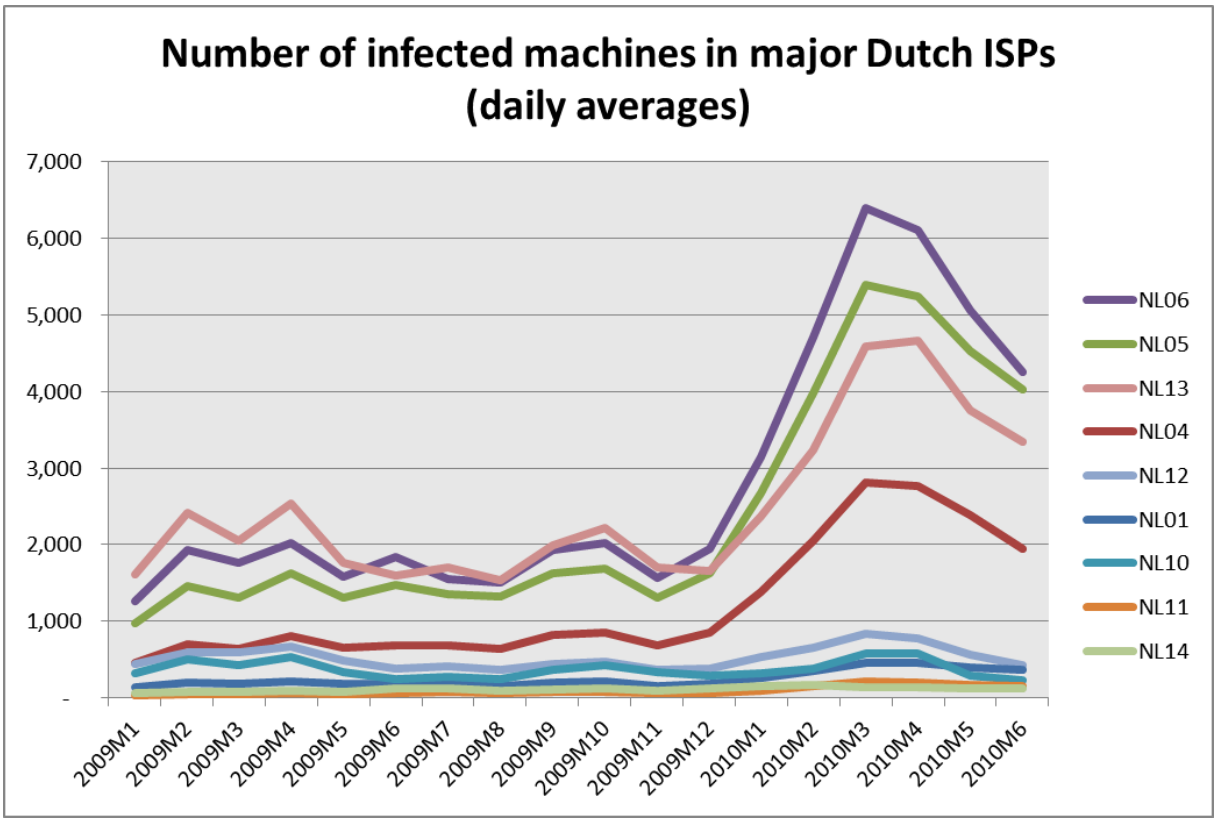

Figure 15 – Time trend of the number of infected machines in Dutch ISPs (daily average over each month)

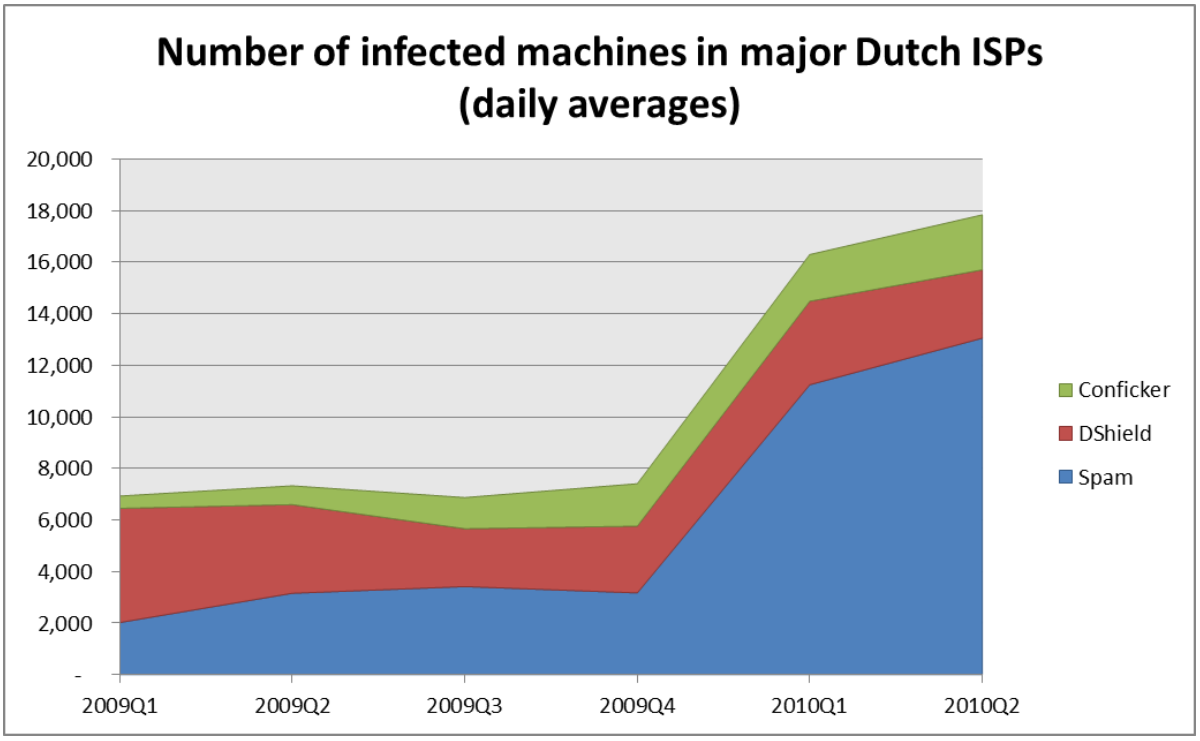

Figure 16 - Number of infected machines in major Dutch ISPs by dataset (daily average over each quarter)

Preliminary Metrics for Relative Infection Rates of Dutch ISPs

In this section we benchmark the Dutch ISPs against each other. Figure 17 and Figure 18 show data similar to the previous time trend graphs except that they have been normalized for size – the y-axes on the graphs are the total number of infected machines divided by the number of subscribers of each ISP in that quarter.

Figure 17 and Figure 18 reveal variation among the ISPs. We also see the fluctuations, which are mostly synchronized across the ISPs, indicating as before that the performance of the ISPs remain relatively stable compared to each other, although some minor position changes do occur. In another study, we found a rather strong intertia effect, where ISPs perform relatively stable over time, compared to each other.[20] And finally, we see that the number of infected machines per quarter varies without a clear trend up or down, while the daily averages clearly rise. The most probable explanation is again that the bots have become more active in 2010. This upward swing is noticeably larger for two of the ISPs in the group.

General observations aside, caution must be exercised regarding the benchmarks, which should be considered preliminary. The main reason is that although our previous studies have indicated that our methodology and findings are robust at the level of the population as a whole, sampling issues and certain measurement limitations could cause rankings to change at the level of individual ISPs. For example, at least two of the major Dutch ISPs have port 25 blocking in place, which would artificially improves their rank in these metrics.[21] We have however performed some extra checks for the Dutch ISPs to increase the robustness of these metrics (see also Chapter 4).

To reiterate the previous sections: the relatively stable rankings indicate that the security policies of ISPs and their subscriber bases have not changed much during the 18-month period, a point confirmed by the ISPs during the workshops. These performance differences might be a result of differences in the subscriber base of the ISPs, not their policies per se – e.g., if one ISP has users that are more at risk of getting infected with malware. But during the workshops, none of the ISPs pointed to any measurable evidence in this regard, and according to the data provided by the ISPs, all of them have a mix of business customers and residential subscribers, with the latter forming the bulk of the customer base.

It is essential to note that although the performance spread among the Dutch ISPs appears high, when viewed in the global context, they are actually performing relatively close to each other. This can be seen in Figure 19, Figure 20 and Figure 21. In these graphs, each dot presents an ISP, grouped in columns by their respective country. The Y-axes of these graphs measure the number of infected machines per subscriber, as counted in the three datasets. Since these rankings are based on the daily averages, the numbers on the Y-axis should not be read as a percentage that accurately represents the infection rates. Also, the differences among the scores on the Y-axes of the three figures are not relevant. Since the size of the data sets is different, these cannot be compared in a meaningful way.

[20] See Van Eeten et al. (2010, pp. 28-32)

[21] These are Online and KPN. In the case of KPN though Port 25 blocking is only performed on one of their ASNs.

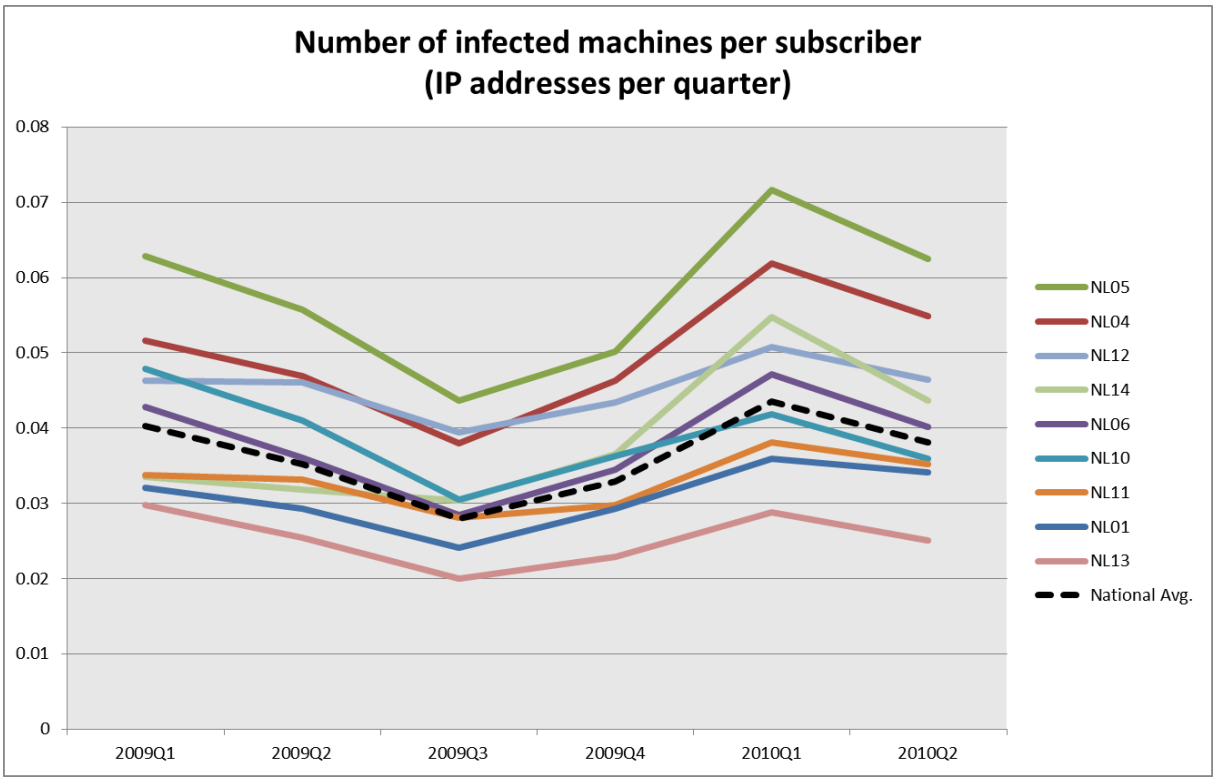

Figure 17 – Number of infected machines per subscriber in the major Dutch ISPs (IP addresses per quarter)

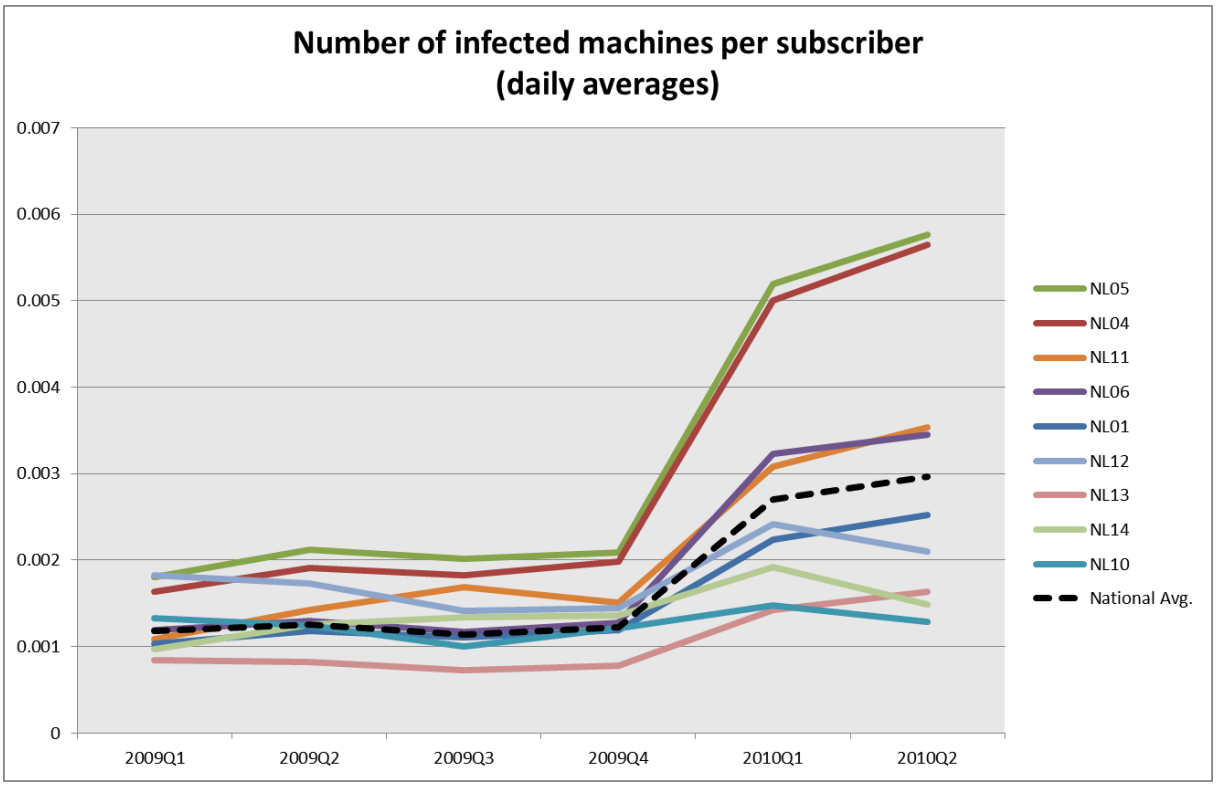

Figure 18 - Number of infected machines per subscriber in the major Dutch ISPs (daily average over each quarter)

It can clearly be seen that (a) the Dutch ISPs have limited variance compared to other countries, especially those on the right side of the graphs; and (b) the average performance of the Dutch ISPs is above the average of the total group of 40 countries, but there a several countries that do better.

Caution must be exercised when interpreting these graphs as country-level benchmarks, as the extra checks that we performed to ensure the accuracy of the data for the Dutch ISPs could not be done for the ISPs in other countries. The graphs show outliers for some of these countries, which means we must keep the limitations of the estimates for those countries in mind when drawing conclusions. The differences between these figures and the country-level benchmarks in Figure 6 and Figure 7 are also noteworthy. They can be explained by the fact that in Figures 19-21 we are only counting infected machines located in ISP networks, where as in Figures 5-6 we looked at the total number of infected machines in ISP and non-ISP networks in each country.

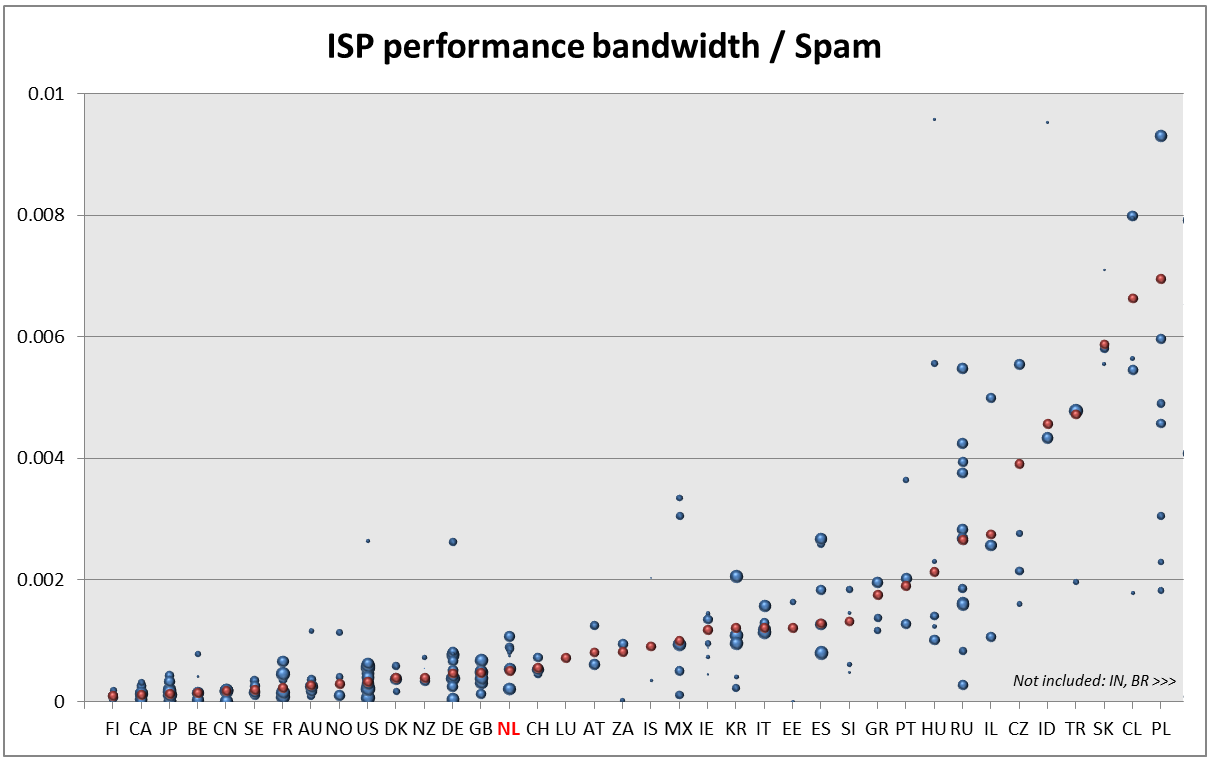

Figure 19 - ISP performance bandwidth across countries, measured by spam sources per subs. (daily average 2009)

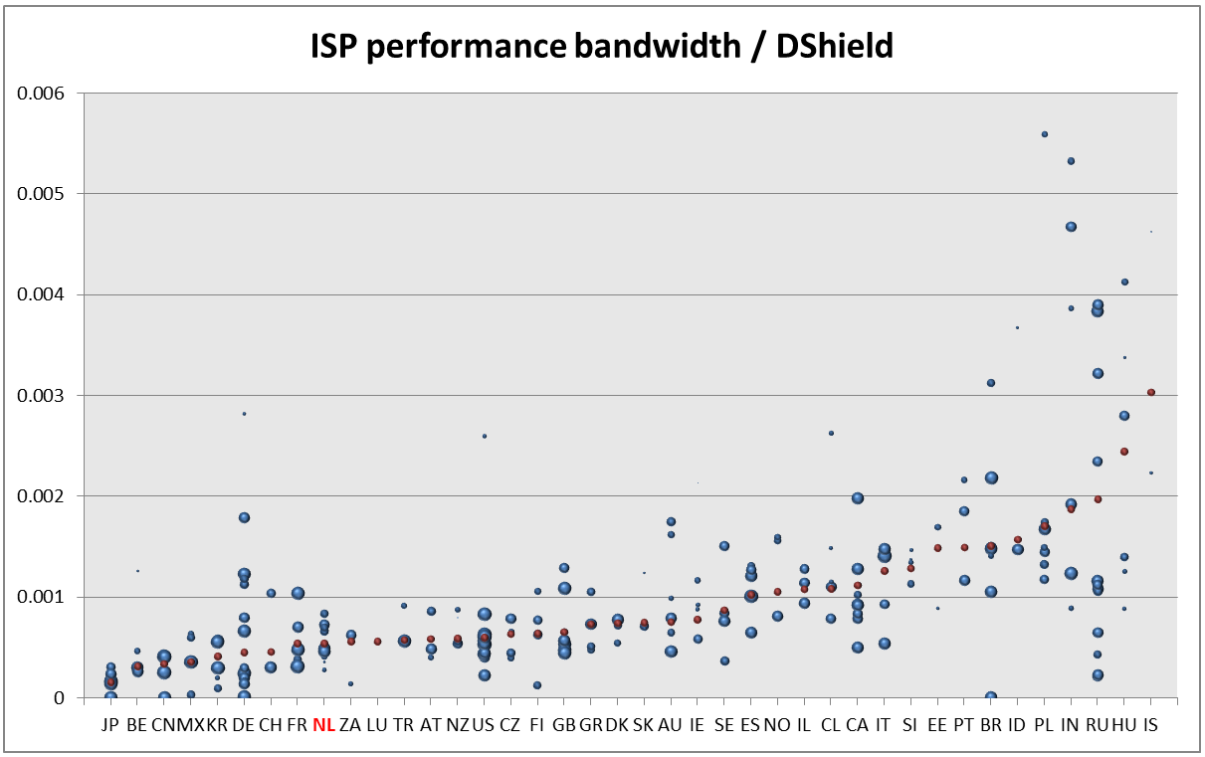

Figure 20 - ISP performance bandwidth across countries, measured by DShield sources per sub. (daily average 2009)

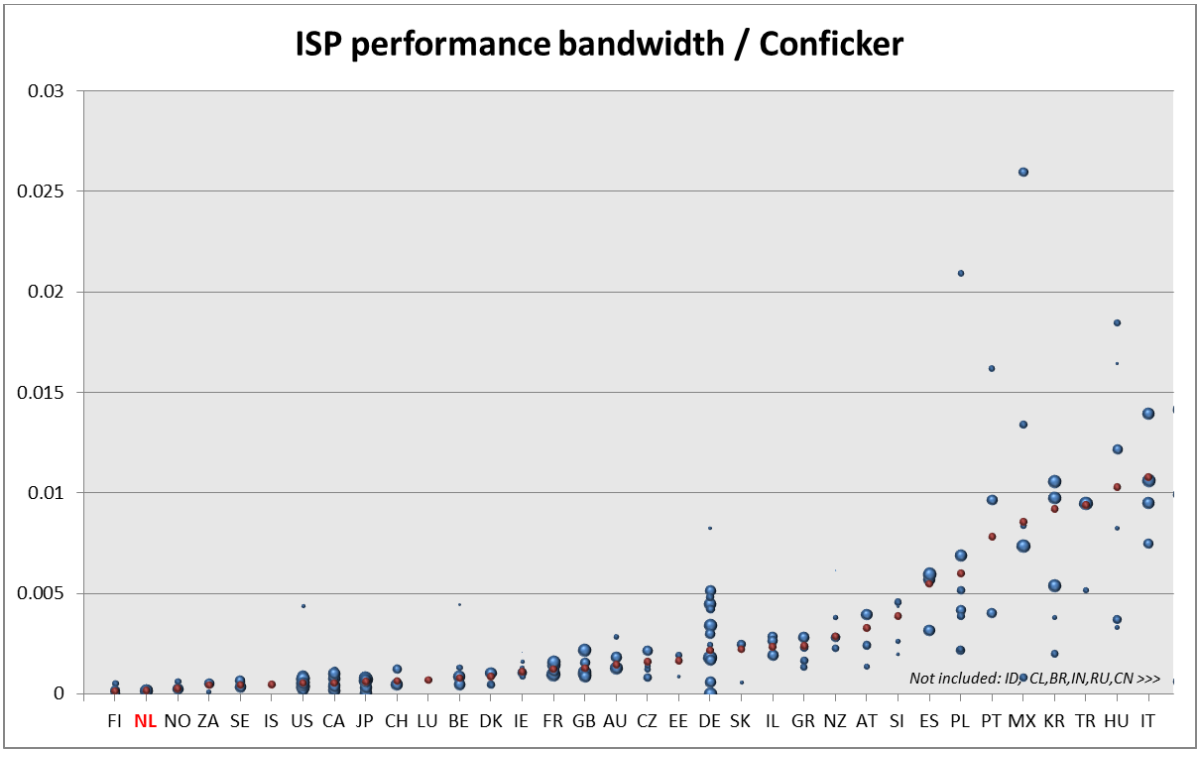

Figure 21 - ISP performance bandwidth across countries, measured by Conficker sources per sub (daily average 2009)

4.
Discussion

Collaboration with the ISPs

The study was designed around publicly available data and was executed independently from government and market players. On three occasions, we sought feedback from the ISPs partipacting in the Anti-Botnet Working Group. At the start of the study, we asked the ISPs to corroborate or correct the information we had on the number of customers over the period under study, as well as the Autonomoums System Numbers of the networks they operated. The information they provided was completely consistent with our own.

Second, halfway during the execution of they study, we presented our methods (Chapter 2) and findings (Chapter 3) during an intensive workshop with experts from all Dutch ISPs participating in the Anti-Botnet Working Group. We discuss this workshop in more detail below. In a nutshell, the feedback indicated that the method was sound and the findings were correct.

Third, after the study had been completed, we presented our conclusions to the Working Group and asked ISPs to comment on how we summarized the main findings.There were no objections to our description of the study's conclusions.

Feedback on the Methodology

On November 3, 2010, we met with multiple experts from the nine largest Dutch ISPs (BBNed, KPN, Online, Solcon, Tele2, Telfort, UPC, XS4All, Ziggo). All of these ISPs were included in our study – with Telfort being treated as part of its parent company KPN, because we lacked separate customer data for the period under study. Together, these ISPs represent over 90 percent of the broadband market. It was a highly constructive workshop that, after discussing the methodology and findings, led to an open-minded discussion on the wider implications of the numbers we had produced.

First, we laid out in detail our methodology. It was judged to be valid by the participants. The main comment we received was that our measurements were probably too conservative – i.e., we underestimated the number of infected machines in the ISP's networks. While our study consciously adopted a conservative approach for dealing with the complexities of internet measurements, the feedback from the ISPs indicated that in one respect we might want to reconsider our design.

This relates to the way we had dealt with DHCP churn – that is, the fact that the same machine may show up under several IP addresses in our data, potentially leading to overcounting the number of infected machines; see Chapters 2 and 3 for more details. The providers all indicated that churn is not a relevant factor in their networks. In a formal sense, many of the IP addresses they assign are dynamic. In practice, however, the bulk of their addresses are assigned for such long periods of time – a year, or longer – that they behave more or less like static IP addresses.

In our preliminary findings, we had corrected for churn. We had first counted the total number of unique IP addresses that harbored an infected machine between January 1, 2009 and July 1, 2010. Then, we had divided that number by 2, as a safe margin to avoid overcounting because of churn. The feedback from ISPs now indicated that this strategy was overly conservative. We have incorporated this feedback by now reporting on the estimate in the form of a bandwidth, where the lower bound corrects for churn and the upper bound does not.[22] In other words, the study reports that between 5-10 percent of all Dutch broadband subscribers have been part of a botnet in the period of January 1, 2009 and July 1, 2010.

Another issue that emerged during the discussion is the different types of users that the ISPs attract as customers. The idea is that some ISP may have more customers with a poor understanding of the risks and limited technical skills. That would be reflected in higher infection rates, even if the ISP put the same amount of effort into mitigation as its peers. This is a valid concern. In another study we did find that characteristics of the user population matter. For example, in countries where users are more likely to use pirated software, infection levels are higher, all other things being equal.[23] We cannot account for this effect in the Dutch market, because there is no publicly available data on the differences in user populations among the ISPs.

That being said, we expect this factor to have a minor impact on the measurements, because there are countervailing forces at work. For example, ISPs that cater to the mass market dominated by users with low awareness and limited competence, are in a better position to implement blanket measures across their whole network. In this market, it is not that contested to block port 25 or use more aggressive quarantining practices. ISPs that cater to more advanced users, sometimes also charging slightly higher rates, have to be much more nuanced in their approach. The users pay a premium precisely because they are not treated as novices and enjoy fewer restrictions on the use of their connection. The measurements in the Dutch markets seem to support this countervailing effect. Of course, the best way to deal with this issue is to get better metrics on the user populations. Hopefully, future research can progress in this direction.

Exploring the Implications of the Infection Rates of Dutch ISPs

In preparation for the workshop, we wanted to assess the scale of the current mitigation efforts. We asked all participating ISPs to answer two questions about their ongoing practices: (1) how often did they contact infected customers per month; and (2) how often did they quarantine or otherwise limit the connections of customers per month?

Most, but not all, ISPs provided answers to these questions. Some wanted these numbers to be treated confidentially. Table 1 compares the answers that we are allowed to report to the number of infected machines we counted in their networks for the month of June 2010.

This comparison suggests that, roughly speaking, ISPs contact around 10 percent of the total number of infected customers. During the workshop, it became clear that the ISPs were surprised by this remarkable discrepancy. The subsequent discussion revealed that an important explanation for the discrepancy was the difference in our datasets and the ones they were using in their mitigation efforts and abuse teams. Our datasets were much larger and revealed many more infections. Their number of infections identified by their own data was much closer to the size of the mitigation effort.

[22] See page 26 for more detail on this issue.
[23] See Van Eeten et al. (2010, pp. 36-45).

Table 1 – Comparing infection rates and mitigation efforts at selected Dutch ISPs

ISP	Number of infected machines in June 2010	Number of customers receiving notifications	Number of customers being quarantined
NL13	16952	800-1000	200-300
NL10	1985	1000	7
NL04	7431	60-80	Not provided
NL12	3942	Not provided	"50 filters per day"*
NL06	18239	900	180

* the total number of customers being quarantined per month was unknown because some customers resolve the problem quickly and have the filters removed within hours, while the connections of others were filtered for days, weeks or even indefinitely.

The ISPs concluded that their data about infected machines is much less comprehensive than they had hitherto presumed. They expressed a need for better data. It was also discussed that this represents an opportunity for collaboration with each other, with researchers and with government. It was deemed inefficient if every ISP on its own had to build the same infrastructure for capturing and parsing the right datasets and feeds. A form of centralized, shared clearinghouse might be an efficient way to drastically improve the intelligence that ISPs are using to protect their networks and customers. The Australian Communications and Media Authority (ACMA) has established such a clearinghouse that aggregates numerous data feeds and transforms them into weekly reports for each Australian ISP.

There is on other factor that explains the discrepancy between the number of infections and the number of customers being contacted. ISPs want to be very careful in weeding out false positives, before they act on the information. One way to do this is to wait for corroborating evidence, for example, two independent reports on the same customer. The customers' IP addresses would first show up in, say, a spam data feed and then later in another feed. Only then would the mitigation procedure be initiated. While this is a very legitimate approach, it greatly reduces the number of customers that meet the threshold of being contacted. As we discussed in Chapter 2, there is very little overlap among the different data sets that we use. Somewhere around 12 percent of the IP addresses in one set, also show up in the other. We suspect that something similar may be the case for other data sets. That means that if an ISP waits for a second confirmation before acting, it would ignore almost 90 percent of all the IP addresses in any given data set. This highlights an important issue for future work: how should ISPs balance the need for due process before a customer is contacted with the need to acts against infections at the appropriate scale?

The discrepancy also raised another discussion: if ISPs would receive more comprehensive intelligence on infected machines, how could they scale up their mitigation efforts to match the now much higher number of cases they can or should act upon? Mitigation procedures are costly, especially in terms of customer support. They cannot be expanded at will. These costs places a premium on finding more efficient ways to respond to infected machines. The less costly it becomes to contact, quarantine and help customers to deal with infections, the more cases the ISPs can take on. Reducing these costs also depends on the efforts of other market players, of government and, last but not least, of the end users themselves. As we have argued before, the responsibility should not be assigned exclusively to ISPs. Other market players can also contribute, and tolerating a certain level of infection is probably economically rational, also from a societal perspective.

5.
Conclusions

Main Findings

This study was designed to fulfill five objectives:

1. Collect data from different sources to assess the number of infected machines in the Netherlands from January 2009 to June 2010 and benchmark the findings against other countries – e.g., on a per capita basis;
2. Establish what percentage of infected machines in the Netherlands are located in the networks of Dutch ISPs – in other words, the extent in which Dutch ISPs are indeed control points for botnet mitigation;
3. Collect data from different sources to assess the number of infected machines within the networks of Dutch ISPs;
4. Develop preliminary benchmarks that rank the rate of infection of Dutch ISPs against each other and against ISPs in other countries;
5. Discuss the methodology and findings during a workshop with the ISPs participating in the Anti-Botnet Working Group.

The study analyzed three large data sets that can identify infected machines worldwide: (1) a spam trap collecting the IP addresses of spam-sending machines, which are typically infected machines; (2) data on global security incidents from the SANS Institute; and (3), a sinkhole for the Conficker botnet that logs which infected machines "call in" to receive instructions from the botnet command and control structure. To indicate the size of these data sets: in 2009, they contained 67 million IP addresses, 130 million IP addresses and 169 million IP addresses, respectively.

We have identified the locations of these IP addresses in terms of the country and the network (Autonomous System) to which they belong. The Autonomous Systems were then checked to see whether they were part of an ISP network or not – where ISPs is defined as a provider of Internet access. For the Dutch market, this means that we have mapped the infected machines in the networks of all ISPs participating in the Anti-Botnet Working Group, who hold an aggregate share of over 90 percent of the broadband market. Our approach to analyze this data consistently sought to generate a robust, conservative estimate. Our numbers should be interpreted as lower-bound estimates of the size of the problem.

We are now in a position to summarize the main findings for each of these objectives:

- At the country level, the Netherlands has an average number of infections per internet user, when compared to a group of 40 countries. The exception is the Conficker botnet, where the infection level is much lower.
- Around 80 percent of all infected machines in the Netherlands are located in ISP networks. The remaining 20 percent are located in the networks of hosting providers and Surfnet, the Dutch academic network. This pattern is very similar to that across the whole group of 40 countries.

- Over 60 percent of all infected machines are located in the networks of the three largest Dutch ISPs, by and large consistent with their share of the broadband market.
- In the period under study, January 2009 to June 2010, we identified around 1.1 million IP addresses that indicated the presence of an infected machine in the Netherlands. Around 900,000 of those were located in the networks of the main Dutch ISPs. This can be conservatively interpreted as 450,000 to 900,000 infected machines.
- To put it differently: During 2009, between 5-10 percent of all Dutch broadband subscribers have suffered an infection that made their machine part of a botnet. The data for the first half of 2010 suggests this pattern will hold or get worse for 2010.
- In reality, the number of infected machines in the Netherlands is probably significantly larger than our estimates suggest. This is because only a fraction of the infected machines we identified show up in more than one dataset. In other words, there is only a small overlap among the datasets. This suggests that if we would include additional datasets, the number of infected machines we would identify in the Netherlands is likely to be substantially higher.
- There is no consistent trend in the number of infected machines in the Netherlands, but the problem has certainly not gotten smaller during the period under study (January 2009 to June 2010). If anything, the data suggest the problem has been getting worse.
- Dutch ISPs perform better than average – that is, they have fewer infections per customer – compared to the total population of ISPs in the wider OECD. Still, ISPs in a variety of countries perform better.
- Dutch ISPs contact around 10 percent of the total number of infected customers in their networks. This low ratio has two main explanations: (1) the ISPs lack intelligence on infected machines – their own data feeds capture much less than the datasets used in the study; (2) ISPs need to be careful to avoid false positives when deciding whether to contact or quarantine a customer, so they cannot act on every single piece of data.
- The extent to which the mitigation process is automated – contacting infected customers and, when needed, limiting or quarantining their connection – has a direct impact on the number of customers that are contacted or quarantined. Automation drives down the costs of mitigation.
- There are significant differences in performance among the Dutch ISPs. Some ISPs suffer three to five times more infections per customer than others. That being said, this variation in performance is not that large compared to the variation in the overall population of ISPs in the 40 countries of the wider OECD. Those differences can span one or even two orders of magnitude – in other words, ten times or one hundred times more infections per customer.

Exploring Next Steps

Our assignment was to generate a fact-finding report for the Dutch market. It was explicitly not intended to draw policy implications. Keeping that in mind, we still feel it is valuable to briefly reflect on our findings, outside the immediate scope of the study.

A number of results merit further reflection on the role of ISPs. We found they form a critical control point, harboring around 80 percent of all infected machines in the Netherlands in their networks. Notwithstanding the efforts of Dutch ISPs collaborating in the Anti-Botnet Convenant, the problem of botnets is substantial and shows no signs of diminshing. This is not meant to say that ISP efforts have no effect, on the contrary. Rather, it implies that the problem seems to be larger than was assumed so far. One of the most striking results is the discrepancy between the current mitigation efforts of ISPs and the size of the problem. It appear that ISPs contact about 10 percent of the customers that suffer from an infection. That is much lower than most people, including the ISPs themselves, had expected.

How should this discrepancy be evaluated? Presenting a number like '10 percent' seems to invite a straightforward conclusion: the ISPs are not doing enough – not until they get closer to mitigating all infected machines in their network.

We would argue that this conclusion is misguided. First of all, as all security comes at a cost, it is economically desirable to tolerate a certain level of infection in our networks. What the appropriate level is, will be a matter of debate. In general, one can argue that the closer one tries to get to infection-free networks, the more rapidly the marginal costs will rise of reducing the infection levels even further.

Second, it is not reasonable to expect ISPs to fully internalize this problem. As stated earlier, these infections originate in criminal behavior. They are not caused by the ISP itself. Even when we only look at the legitimate market players that can influence the magnitude of this problem, we can see a broader set of players that contribute to this problem: hosting providers, software vendors, computer retailers, registrars, e-commerce companies and, last but certainly not least, the end users themselves, be it home or business users. It is their machines that are infected. In addition to the market players, we should also not neglect the role of governments, who can help by raising awareness, providing market oversight and enforcing the law. Governments in Japan, Korea, Australia and Germany have gone even further and now provide direct support to citizens who are struggling with infected machines.

These considerations suggest it would be wrong to treat botnet mitigation exclusively, or even predominantly, as the task of ISPs. That being said, we also have to acknowledge that ISPs can play a crucial role and that the economic incentives under which they operate will make them reluctant to take on that role. As earlier studies reported, most of the damage of botnets are borne by other actors than the owners of the infected machines or the ISPs that connect them to the Internet.[24] Since they do not suffer the full extent of this damage, ISPs and their customers do not have the economic incentives to invest in mitigation at the level that is socially desirable. The cost of mitigation will therefore be a major factor in influencing how much ISPs are willing and able to do in this area.

Keeping these considerations in mind, what options do our findings suggest for the Dutch ISPs? In general terms, we see two ways to improve their current mitigation efforts: (1) improving detection, and; (2) improving mitigation of infected machines.

The first option, improving the detection of infected machines, has already been discussed briefly in the previous chapter. Our study made it clear that the data feeds that the ISPs are currently using, does not give them adequate intelligence on the total number of infected machines in their network. There are additional data sets that ISPs can tap into to improve their intelligence, without having to install intrusive and controversial monitoring technologies in their networks.

In light of the costs of acquiring and processing this data for use in the mitigation efforts of the abuse departments, it seems worthwhile to explore whether this effort could be pursued collectively. It may be possible to achieve economies of scale by building one platform for all ISPs, rather than each ISP building a platform on its own. A centralized, shared clearinghouse might be an efficient way to drastically improve the intelligence that ISPs are using to protect their networks and customers against modest cost. There is an interesting real-world example that could be studied: the Australian Communications and Media Authority (ACMA) has established a clearinghouse that aggregates numerous data feeds and transforms them into weekly reports for each Australian ISP.

[24] See Van Eeten et al. (2008) and Anderson et al. (2008).

Another function of such a clearinghouse could be to provide light-handed oversight on the self-regulation of the Dutch ISPs. By enabling trend analysis and benchmarks to compare Dutch ISPs amongst each other and to peers in other countries, the government can get a sense of the effectiveness of the mitigation efforts – and thus of the Anti-Botnet Covenant. This would be a way to provide oversight, without having to interfere in the freedom of ISPs to pursue different strategies of mitigation and without imposing regulatory costs on them. As an aside: this analytical function is completely absent from the ACMA clearinghouse. For a variety of reasons, it may be easier to achieve such functionality if the clearinghouse is set up as a public-private entity, at arms length from the government.

The second option, improving the mitigation of infected machines, focuses on ways to enable ISPs to better deal with infected customers. Sharing tools and procedures may be helpful here. The critical issue will be to reduce the cost of customer contact and support. The more efficient an ISP can deal with a customer, the more infections it can take action on, within the same amount of resources. The discussion during the workshop with the ISPs confirmed a finding from an earlier study we did on botnet mitigation, namely that automation is a critical factor in increasing the efficiency of dealing with infected machines.[25]

As with improving detection, there may be economies of scale in mitigation as well. One interesting development in this respect are the initiatives in Japan, Korea and Germany. In those countries, the government has stepped in to help citizens that have their machines infected. This help consist of the distribution of a software tool to remove bot infections (Japan, Germany) and the operation of a national call center for infected consumers (Korea, Germany). These initiatives potentially reduce the cost to an ISP of dealing with an infection. It stands to reason that lower cost enables ISPs to contact, quarantine and help more customers, within the same resources. This promises lower infection levels and more protection for citizens in those countries.

Such efforts should of course be part of a more comprehensive strategy to combat botnets. Software vendors, for example, can help to prevent infections and to clean them up more effectively. Hosting providers also harbor infected machines that need to be cleaned up. More importantly, they can collaborate with security experts and law enforcement to take down command and control servers in their networks, as they have done with some success over the past few years. While taking down the command and control infrastructure of botnets can be an important tactic, we also have to realize that this provides no structural solution. When it comes to taking down a botnet, there is no alternative to the laborious process of actually cleaning up large numbers of infected machines.

Whatever options ISPs, other market players and government choose to pursue, it is clear that the fight against botnets has only just begun. The efforts of all these stakeholders will determine whether progress can be made in reducing the societal impact of this kind of cybercrime.

[25] See Van Eeten et al. (2010, pp. 36-45).

Appendix 1

Code	Country Name	Number of ISPs
AT	Austria	3
AU	Australia	6
BE	Belgium	4
BR	Brazil	8
CA	Canada	9
CH	Switzerland	3
CL	Chile	5
CN	China	5
CZ	Czech Republic	4
DE	Germany	13
DK	Denmark	3
EE	Estonia	2
ES	Spain	6
FI	Finland	4
FR	France	5
GB	United Kingdom	8
GR	Greece	3
HU	Hungary	6
ID	Indonesia	2
IE	Ireland	7
IL	Israel	3
IN	India	6
IS	Iceland	2
IT	Italy	4
JP	Japan	6
KR	South Korea	4
LU	Luxembourg	1
MX	Mexico	5
NL	Netherlands	6
NO	Norway	5
NZ	New Zealand	4
PL	Poland	5
PT	Portugal	4
RU	Russia	10
SE	Sweden	4
SI	Slovenia	5
SK	Slovakia	2
TR	Turkey	1
US	United States	15
ZA	South Africa	2
TOTAL	**40**	**200**

References

Anderson, R., R. Böhme, R. Clayton and T. Moore (2008). *Security Economics and the Internal Market*. ENISA (European Network and Information Security Agency). Available online at http://www.enisa.europa.eu/doc/pdf/report_sec_econ_&_int_mark_20080131.pdf.

Microsoft (2010). *Microsoft Security Intelligence Report, Volume 9*. Microsoft. Available online at http://www.microsoft.com/security/sir/.

Moore, D., C. Shannon and J. Brown (2002). *Code-Red: a case study on the spread and victims of an Internet worm*. Proceedings of the 2nd ACM SIGCOMM Workshop on Internet measurment. Available online at http://portal.acm.org/citation.cfm?id=637244.

Stone-Gross, B., M. Cova, L. Cavallaro, B. Gilbert, M. Szydlowski, R. Kemmerer, C. Kruegel and G. Vigna (2009). *Your Botnet is My Botnet: Analysis of a Botnet Takeover*. 16th ACM Conference on Computer and Communications Security, November 9–13, 2009, Chicago, Illinois. Available online at http://www.cs.ucsb.edu/~seclab/projects/torpig/torpig.pdf.

US GAO (2007). *Cybercrime: Public and Private Entities Face Challenges in Addressing Cyber Threats*. United States Goverment Accountability Office. Available online at http://www.gao.gov/new.items/d07705.pdf.

Van Eeten, M., J. Bauer, H. Asghari and S. Tabatabaie (2010). *The Role of Internet Service Providers in Botnet Mitigation: An Empirical Analysis Based on Spam Data. STI Working Paper 2010/5.* OECD. Available online at http://www.oecd.org/officialdocuments/displaydocument/?doclanguage=en&cote=dsti/doc(2010)5.

Van Eeten, M. and J. M. Bauer (2008). *Economics of Malware: Security Decisions, Incentives and Externalities, OECD STI Working Paper 2008/1*. OECD. Available online at http://www.oecd.org/dataoecd/53/17/40722462.pdf.

Zhuang, L., J. Dunagan, D. R. Simon, H. J. Wang, I. Osipkov, G. Hulten and J. D. Tygar (2008). *Characterizing Botnets from Email Spam Records*. LEET '08. First Usenix Workshop on Large-Scale Exploits and Emergent Threats, San Francisco. Available online at http://www.usenix.org/event/leet08/tech/full_papers/zhuang/zhuang.pdf.