



European Cyber Security Perspectives 2015



National Cyber Security Centre
Ministry of Security and Justice



Login

PASSWORD



Preface

Dear reader,

Following the success of last year's publication, we are proud to present the second edition of our European Cyber Security Perspectives report. Through this collection of articles, we aim to share our different perspectives and insights, the latest developments and achievements in the field of cyber security, cybercrime investigations and cyber resilience.

By uniting the expertise of four parties with diverse roles in the cyber security domain, we hope to offer some fresh perspectives on issues and developments that we believe to be relevant for society.

Rather than presenting you with bare facts and figures, we describe real-life cases and experiences from our professional practice. Topics include the trends to watch in 2015, responses to high-profile vulnerabilities and advances in the detection and investigation of targeted cyber attacks. Each article in the report can be read independently, allowing you to focus on the topics that interest you most.

A central theme in this year's report is cooperation in cyber security. We strongly believe in the value of uniting cyber security capabilities across public and private organisations and even national borders. In fact, you will see that many of the articles in this publication have been co-authored by specialists from different parties. This reflects the growth in our collaborative projects when addressing the latest cyber security challenges. We aim to foster more of those partnerships in the coming year.

We encourage you to build on our work and experiences and hope that this report will inspire further enhancements in cyber security, cybercrime investigations and cyber resilience, both in the Netherlands and abroad.

Enjoy and be safe!

The European Cyber Perspectives Team.



Quotes contributing partners



Jaya Baloo
CISO, KPN

As more and more countries work to develop offensive 'cyber attack' capabilities, the market for security vulnerabilities increases. The productizing of these vulnerabilities will be offered to anyone willing to pay the price, not just state actors. As the attack surface widens, the recourse for companies and individuals is still to follow best current practices to keep software up to date, secure and backup data, use strong authentication and encryption, and above all have a healthy distrust when online and sharing information. We are fighting an information security arms race, we need to win it incrementally.

Henk Geveke

Managing director defence, safety and security, TNO

Coming together is a beginning, keeping together is progress and working together is success." It's almost as if Henry Ford anticipated today's challenges in cyber security and cyber resilience. Collaboration is indeed a key requisite for establishing a secure and reliable cyberspace, particularly because our adversaries are actively joining forces as well. TNO has built strong alliances and partnerships in its pursuit of cyber security innovations and this report reveals some of the tangible achievements emerging from these collaborative efforts. We hope to intensify this successful approach even further in years to come.



Wilbert Paulissen

Head of Central Criminal Investigations Division, National Police



In 2014, only one year after the merger of 26 police forces, the National Police consolidated its internal co-operation, enabling us to work more efficiently and jointly together, within a European framework through Europol. We improved our results through public-private partnerships: enhancing cyber security with KPN, building on factual knowledge with TNO, and establishing communities with NCSC. 2015 will bring even more police co-ordination. Law enforcement worldwide will be permanently supported in the fight against high tech crime, through Interpol. I look forward to creating more, larger and longer lasting successes by working together with public, private, and police partners.

Hans de Vries

Head of National Cyber Security Centre

Innovative new digital products and services are created at a breathtaking pace, enriching individual lives and redefining the way we work and communicate. The associated freedom and opportunities for growth can only fully flourish if the underlying technology functions in a safe and secure fashion. Considering the growing interdependence of organizations due to the interconnected character of the digital world, it is paramount that ever closer cooperation on cyber security is called for. This publication offers an excellent demonstration of this resolution by exploring this important subject from a multitude of perspectives.



Contents



Trends to watch in 2015	6
Pest control	11
Heartbleed: Lessons learned from a broken heart	14
Yo, Trust me	19
Did you floss today?	22
CyberDawn, training for a cybersecurity crisis	26
On Her Majesty's Secret Service GRX & A spy agency	29
Spotting the prowler	35
Shedding light on the Dark Web	40
Ahead of the Threat	44
Honeypots at KPN	49
Help, the app ate my password!	52



Trends to watch in 2015

Frank Fransen, Richard Kerkdijk, Reinder Wolthuis (TNO)

The year 2014 brought us many interesting events and developments in the field of cyber security and cyber resilience. The commotion surrounding the Heartbleed vulnerability was certainly hard to miss, even for those not particularly interested in security. More recently we saw the unveiling of Regin, a piece of highly sophisticated – and allegedly state sponsored – cyber espionage malware. Meanwhile there were also successes in the pursuit of cyber criminals and advances in (cyber) security solutions and collaborations. All in all a great wealth of events to take note of, as illustrated by the timeline that runs through this report.

This article will highlight the most prominent new developments that we believe will continue to be of interest in years to come. Here we note that the trends described in last year's edition are all still ongoing and significant as well.

Trend #1: High profile vulnerabilities preoccupy security practitioners

We have seen a variety of high profile “zero-day” vulnerabilities¹ come to light. Prominent examples are the security bugs now known as Heartbleed², Shellshock³ and Poodle⁴. Interestingly, some of these vulnerabilities stem from flaws in widespread open source software that took several years to discover. The Shellshock vulnerability, for instance, has reportedly existed since 1989 and was not discovered until this year. Attackers might therefore have exploited these vulnerabilities well before their public disclosure, although there have not been any public reports of actual incidents.



‘Heartbleed’ logo

Particularly striking about the aforementioned vulnerabilities is the hype that surrounded them. All came with appealing names and fancy logos (see figure) and all received unusually strong media attention. With all the commotion raised, security practitioners had little choice but to respond to these vulnerabilities with full priority. Some have undoubtedly felt pressure from their organisation’s leading executives, who were also alerted by the prominent headlines.

The attention surrounding high profile vulnerabilities has undisputedly had some positive effects. If anything, it served as a forceful call to action that compelled many organisations to remediate these issues quickly. However, there are also some disconcerting aspects to all this hype.

Firstly, we observe that risk considerations were of minor importance in the above cases, particularly at the level of individual organisations. Priorities were mostly dictated by the visibility and attention surrounding these issues.



‘Shellshock’ logo

If we take Heartbleed as an example, the effort spent on remediation was by all means justifiable since, arguably, this was one of the worst security bugs seen in the past years. If we take a broader perspective, however, we see that Heartbleed is not the vulnerability most exploited by actual attackers. Rather, such attackers seem to favour fairly common vulnerabilities such as those present in the well-known OWASP Top 10⁵. Arguably, many organisations would have reduced their exposure to attacks more effectively by resolving these less prominent vulnerabilities first.

The greatest concern enclosed in this development is that it could promote a culture in which the priorities of security practitioners are determined by publicity rather than risk. Anyone discovering a security vulnerability may develop a website and logo to make a name for himself. If marketed well, this might (again) invoke coordinated effort from the security community to resolve the issue. If the actual risk is relatively minor then all this achieves is distraction of valuable resources.

As more high profile vulnerabilities will probably appear in the near future, security practitioners will hopefully keep a keen eye on the actual risk that each threat represents to their specific organisation and business.

⁽¹⁾ A zero-day vulnerability is a security deficiency in software for which a patch or fix is not yet available.

⁽²⁾ CVE-2014-0160, a security flaw in the OpenSSL cryptography library, see <http://en.wikipedia.org/wiki/Heartbleed>

⁽³⁾ a family of security flaws in the Unix Bash shell, see [http://en.wikipedia.org/wiki/Shellshock_\(software_bug\)](http://en.wikipedia.org/wiki/Shellshock_(software_bug))

⁽⁴⁾ CVE-2014-3566, a vulnerability in the security protocol SSL 3.0, see <http://en.wikipedia.org/wiki/POODLE>

⁽⁵⁾ https://www.owasp.org/index.php/Top_10_2013-Top_10

Trend #2 Collaboration in cyber security successfully materialises

Experts have been pointing out for a long time that collaboration is a critical ingredient for establishing effective cyber security and cyber resilience strategies. The Dutch National Cyber Security Strategy¹, for instance, emphasises the need for public and private parties to work together in this field. This year we have truly seen such collaboration materialise into tangible developments and effects.

A visible manifestation of this development is that law enforcement agencies successfully collaborate across national borders in the fight against cybercrime. A prominent case took place in May 2014, when buyers of the “Blackshades” malware were arrested through an internationally coordinated action². This particular operation involved over 300 police raids in 16 countries, including The Netherlands. More recently, collaboration amongst law enforcement agencies led to the simultaneous take down of several illegal online markets, including the infamous Silk Road 2.0³. What these cases reveal, is that long-term efforts to establish cross border coordination and collaboration in law enforcement are truly paying off. In fact, the successes achieved suggest that the international nature of cybercrime is no longer a fundamental obstacle for the arrest and prosecution of perpetrators. For the recently formed Dutch National Police, these cases were also great examples of what has now become possible through co-operation between regional units, supported by the central high tech crime unit.

Collaboration in cyber security has also manifested itself through various threat intelligence sharing initiatives. The concept of threat intelligence communities, i.e. networks of organisations that actively exchange threat intelligence amongst each other, is rapidly dispersing and professionalising. To keep pace with the continuous surge of threats and vulnerabilities such communities now deploy dedicated automation solutions⁴. These tools and platforms not only facilitate an efficient exchange of threat-related information (e.g. so-called Indicators of Compromise⁵), but also optimise their utilisation in company internal security processes (e.g. patch management and incident monitoring). The latter is fortified by the advent (and industry adoption) of standardised technology for threat

intelligence sharing, such as the framework of protocols developed by the MITRE Corporation⁶.

These developments clearly demonstrate the value of collaboration in cyber security, whether it be among law enforcement agencies, through public-private partnerships or within specific private sectors. It is likely that its importance will increase even further in the coming years.

Trend #3 Responsible disclosure is the norm and bug bounties continue to increase

How to disclose information on newly discovered vulnerabilities has been the subject of debate among security experts and hackers for years. Some believe that vulnerabilities should not be disclosed until an effective solution is available, whilst others are of the opinion that vulnerabilities should be published via open media as early as possible and without any restriction. The latter is partly driven by the experience that many companies (e.g. hardware or software manufacturers) are slow in their response to a discrete vulnerability disclosure and might even try to silence or sue the person reporting the problem.

We now see that the concept of responsible disclosure has become common, if not the norm, for any ICT-intensive company with a reasonable degree of security maturity. Responsible disclosure basically means that the owner of a vulnerable product or system is allowed some time to fix the problem before it is disclosed to the public. In turn, the person reporting the vulnerability might get certain protective guarantees, for instance that he or she will not be facing legal repercussions as long as the organisation’s rules of play are properly followed. Organisations that embrace this concept, usually publish a responsible disclosure policy and corresponding procedures to clarify their ground rules and guide the process of actually reporting and disclosing a vulnerability. In The Netherlands, the responsible disclosure approach was heavily stimulated by the Minister of Security and Justice, who commissioned the NCSC to develop a Guideline for arriving at a practice for responsible disclosure⁷.

In addition to the above, more and more companies maintain so called bug bounty programs that involve financial rewards for reporting vulnerabilities and other security problems. Whilst up until recently such rewards were only issued by major players such as Google, Yahoo, Twitter, Facebook and Microsoft, we now see them at smaller companies as well. Examples in The Netherlands include the on-line market marktplaats.nl

⁽¹⁾ see <http://www.rijksverheid.nl/documenten-en-publicaties/rapporten/2013/10/28/nationale-cyber-security-strategie-2.html>

⁽²⁾ see article “Pest Control”


⁽³⁾ see article “Shedding light on the Dark Web”

⁽⁴⁾ see article “Ahead of the threat – Enhancing cyber intelligence communities”

⁽⁵⁾ artefacts observed on a network or computer system that indicate (potential) intrusions

⁽⁶⁾ STIX, TAXII and CybOX, see <http://www.mitre.org>

⁽⁷⁾ see <https://www.ncsc.nl/actueel/nieuwsberichten/leidraad-responsible-disclosure.html>



Screenshot of Twitter page showing continuous feed of rewarded bugs.

and the relatively small SNS Bank. The size of the reward usually depends on the severity of the problem that was reported. Interestingly, the rewards have substantially increased in the past years. At Google, for instance, today's awards are five times greater than in 2010 and this year's Pwn2Own computer hacking contest - a recurring event at the CanSecWest hacking conference - featured the highest amount of prize money since the first competition in 2007⁸.

Notably, not all companies combine their responsible disclosure programme with a financial reward system and this does not necessarily imply that they are ineffective. For many ethical (or "white hat") hackers, the primary motivator for reporting a vulnerability is still the glory and in such cases a T-shirt stating "I hacked < company > and all I got was this lousy t-shirt" can already be quite rewarding. In fact, many young hackers collect such T-shirts to build on their resumes. Thus there are several ways of motivating responsible disclosure of vulnerabilities. However, having a proper policy and procedure in place is a key requisite.

Trend #4 Widespread use of smartphones stimulates two-factor authentication

Traditionally, the combination of a username and a password has been the primary means for a user to authenticate himself for the use of services over the internet. The username/password mechanism is easy to implement and relatively easy to use. The weaknesses of password-based access control are, however, becoming more and more apparent and alarming. There have been various reports of large files with stolen username/password combinations being leaked on the Internet. We also see that many users still apply the same password for different services and in many cases their e-mail address serves as the username. All of this is of course very convenient for hackers. Passwords are becoming a millstone around the neck of access control on the internet. A password only represents knowledge of the user. Once this knowledge factor has been compromised, the authentication scheme is breached. The quality of authentication credentials can be improved by combining at least two factors from the well-known triplet: something a user knows, something a user possesses, something a user is. The problem has always been that it is quite expensive to provide users with tokens or cards (possession) or register their personal properties (e.g. through biometrics). Traditionally, only financial institutions had a clear business case to enter this arena, since the magnitude of losses due to fraud began to rise to worrying numbers.



⁽⁸⁾ <http://www.cnet.com/news/all-hacking-eyes-on-the-prize-money-at-cansecwest/>

In 2011, some online service providers introduced two-factor authentication, also called two-step authentication or 2FA. We now see that 2FA mechanisms increasingly involve the end user's smartphone. On top of the normal username/password, this setup involves a separate code that is acquired by the user by SMS or generated by a personalized app on their mobile phone. An example of the latter is the Google Authenticator¹. It can generate one-time passwords, and is implemented using open standards from the Initiative for Open Authentication (OATH)². The large number of users that possess a smartphone has greatly contributed to the relatively easy implementation of 2FA. Notably, most service providers offer 2FA as an opt-in, it's use is not mandatory. Since 2013, we have seen a rapid growth in the service providers that offer two-factor authentication³. We expect this trend to continue, especially given the fact that several smartphones have been released in 2013 and 2014 which include built-in biometrics (the third authentication factor), such as the Apple iPhone 5S/6 with its touch ID and the Samsung Galaxy S5 with a similar fingerprint scanner.

Trend #5 Technology giants publicly oppose governmental data collection

To protect their country's interests and its civilians, governments always have had a need to collect information. The Echelon interception system that supposedly has been in place since the early 70s of the last century is a good example of this. Since 9/11 in 2001, we have seen a steep rise in the collection of electronic information with the intent to prevent terrorist attacks. This was by many considered as acceptable, until 2013, when Edward Snowden suggested that this data collection was lacking the checks and balances that some would expect in an open and democratic society.

Although publicly denied, there were allegations that in some countries the information was not always collected in compliance with applicable regulations. The large service providers were the primary target for these 'surveillance tasks', supposedly conducted by some governments, because a lot of electronic information is transported or stored by service providers such as Google, Apple, Microsoft and Facebook. These companies found themselves in an awkward position. On the one hand, they have to comply with regulations and

cooperate with authorities, on the other hand they need to earn the trust of their users, not all of whom were amused by the fact that their private data seemed to have been easily given away to one or more authorities. Nine large US based service providers joined forces and called upon the US Senate in a letter signed by their CEOs⁴ to adapt the USA FREEDOM ACT (H.R.3361) in such a way that it 'would help to restore the confidence of Internet users here and around the world, while keeping citizens safe'. Also they published 5 principles⁵ concerning the attitude towards information collection requests from authorities. At the same time, these companies adopted or will adopt these principles in their own privacy statements.

Most of these service providers have announced to be more transparent and will from now on publish the amount of data collection requests they have received, what the nature of these requests was and how they have responded⁶. Also, most service providers have started to offer new levels of privacy and security to their customers, e.g. default encryption of user data that – allegedly – cannot be circumvented by the service provider itself. An information request of the authorities would therefore be useless (unless the data can be somehow decrypted). On top of that, new security-specific services and devices are becoming available to users (e.g. the Blackphone and nation specific cloud services). We can conclude that this is a development with a more proactive role for the service providers. The world is searching for a new balance where on the one side governments have the right tools to prevent and prosecute criminal and terrorist acts and on the other side civilians and organisations can rely on the fact that no unnecessary personal or business information is collected.

We do have to keep in mind however, that the same service providers that are now on the barricades to fight for user rights are also collecting and storing enormous amounts of data on their users. And they continuously launch new services that potentially invade in the privacy and personal lives of users (e.g. Google Glass). Let's see where the new balance will bring us...

⁽¹⁾ https://play.google.com/store/apps/details?id=com.google.android.apps.authenticator2&hl=nl_NL

⁽²⁾ <http://www.openauthentication.org/>

⁽³⁾ <https://twofactorauth.org/>

⁽⁴⁾ <https://www.reformgovernmentsurveillance.com/USAFreedomAct>

⁽⁵⁾ <https://www.reformgovernmentsurveillance.com/>

⁽⁶⁾ <https://www.apple.com/privacy/docs/government-information-requests-20131105.pdf>



Pest control

Peter Zinn and
Roeland van Zeijst (National Police)


Though the Chinese zodiac claims differently, for law enforcement, 2014 has been the Year of the RAT. This sneaky piece of software was abused by hundreds of Dutch cyber criminals alone. In an international joint action day, Dutch regional police units paid a visit to the most notorious users.

RATs (Remote Administration Tools, or Remote Access Trojans, depending on how you look at them) have been around for a while. But this year it became painfully clear how their ease of use keeps increasing and their ubiquity makes them the tool of choice for less advanced hackers. The RAT, when installed on someone else's computer, will give the operator full access to all system resources. Its plug-and-play functionality offers the hacker any functionality (s)he can think of and its low price (of usually less than a hundred euros) takes away another hurdle to venture into cybercrime. Not all RATs have to be criminal tools per se, as they might be used for legitimate purposes. However, that always requires permission from the owner of the computer on which the RAT is installed. Many RATs come with built-in functionality to infect victims' systems or damage their data on purpose

and should therefore be categorised under hacker tools. Owning such tools, with the intent to abuse them, is illegal in the Netherlands.

Blackshades

An example of a RAT that's been in the news greatly throughout 2014 is Blackshades. It was advertised as a so-called 'pentesting tool' (for testing your own system's vulnerabilities) but actually offered ransomware-, botnet- and other functionalities (see also the article 'Did you floss today?'). It contained a USB infector and other methods to intrude on victims' computers, such as spear-phishing.



Blackshades comes with built-in infection and spreading tools

Blackshades was sold publicly through a site on the regular web. Its reseller offered customer service, advertisements and user manuals on YouTube, and the owners also maintained a professional customer database. Moreover, they offered customers to use the NO-IP service. This service translates a fluctuating home IP address to a static URL. Blackshades users could use this service to use their own home computer as a command & control server for small botnets of infected PCs.




In 2012 it was discovered that the Blackshades backdoor had been spread to Syrian activists through fake anti-government video files. These movies required an update to Adobe Flash Player. However, the actual installer for this update contained the malware.

Rotterdam

In 2013, a 20-year old man was arrested in the United States in what is called the Miss Teen USA Sextortion case. He had been able to hijack the devices of Miss Teen USA 2013, Cassidy Wolf, and several other girls, using Blackshades, taking control of their webcams and snapping videos and photos of them while they undressed. The man was sentenced to 18 months in prison.



Meanwhile, in the Netherlands another Blackshades user was behaving creepily. In late 2013, the Rotterdam regional police unit stumbled upon Blackshades when a girl reported that nude pictures of hers were placed on her own Facebook account. Her computer turned out to have been infected with Blackshades. Police soon found out she wasn't the only victim. A man from Rotterdam had targeted over a thousand victims, taken over their computers using Blackshades, and was 'playing around' with all system resources. From webcam captures, he selected many obscene pictures and videos, some of minors, and even tried to blackmail some of his victims.



Copies of Dutch school exams were leaked by the Rotterdam hacker using Blackshades

He displayed random sadistic behaviour, for example deleting someone's master thesis just before it was finished. Only after his arrest it became clear that he had also played a role in the Ibn Ghaldoun fraud. Ibn Ghaldoun was the school where in 2013 exam papers had been stolen. This came to light when one of the exams was published online. The Rotterdam suspect had found the stolen papers on one of the computers he had hacked into using Blackshades. He had used yet another hacked computer to publish the materials, afterwards destroying all evidence of Blackshades on that computer, making someone else get the blame (at first). After his arrest, the Rotterdam police were able to tie all these cases together. They also found out that the hacker had actually used a non-paid-for copy of Blackshades, proving that it's not only legitimate software some people prefer to download illegally.

Operation 'Partition'

By early 2014, the Blackshades phenomenon had grown into a world-wide menace. The United States facilitated an international joint action against Blackshades users. Sixteen countries participated in an action coordinated by Eurojust, the European Union's Judicial Cooperation Unit. This is a new form of cooperation, and offers a way forward when we want to be able to tackle organised cybercrime gangs operating in several different countries. Although the Blackshades case concerned individual buyers rather than organised crime, the joint action was a good proof-of-concept. In May 2014 a joint action day was planned, on which law enforcement from all of these countries performed 359 house searches and 81 people were arrested. More than 1100 data carriers were confiscated, including computers, hard drives and USB sticks. In some cases sums of money, illegal arms and drugs were found and confiscated. Within the Netherlands, all 10 regional units of the newly reorganised Dutch National Police took part, plus the National High Tech Crime Unit (THTC). The latter also coordinated the action, which was dubbed Operation 'Partition'. A cooperation like this, horizontal rather than hierarchical, shows how co-operation will shape the future way of tackling cybercrime. The regional units were each in the lead of their own cases, but a good communication between the units assured mutual assistance if needed. In order to tackle the most likely and relevant suspects, triage was put into place. Out of 284 unique Dutch users, 34 suspects were chosen, based on live information about their recent use of Blackshades. This included activation, the use of NO-IP, and the number of current victims. So, for example, a cyber security professional whose behaviour showed only an examining of the malware was excluded from the hitlist. On average every regional unit would visit three suspects for a house search. This resulted in the

confiscation of 120 data carriers for further investigation. No arrests took place in the Netherlands to date, though several serious suspects will find themselves prosecuted, fined and/or punished in other ways.

Not tole-rat-ed

Cybercrime is one of the main focus areas for the Dutch police. In the coming years, the regional units will take up an agreed, and ever increasing, minimum number of cybercrime cases, including more advanced cybercrime. The National High Tech Crime Unit will co-operate with the regional units to help them build up their capacity and know-how if needed. Most regions already have the technical capability to run cybercrime cases. The agreement to run a specified number of investigations per unit is expected to improve the intake process, might increase operational capacity and will help balance internal priorities. Eventually, this might lead to designated cybercrime teams in the regional units, able to run relevant cybercrime cases from start to end, such as investigations into pesky RATs. Beware: with the wrong intentions, this ease of control over another system's resources opens Pandora's box for many types of victims, ranging from teenage models to large companies. Blackshades and the like significantly lower the threshold for criminal behaviour, creating new types of (non-tech savvy) local criminals as well as new types of victims. Society's message should be clear: the illegal use of RATs will not be tolerated. Law enforcement will continue to play its part in cracking down on RAT-enabled cybercrime.





Heartbleed: Lessons learned from a broken heart

Pieter Rogaar (National Cyber Security Centre)

Heartbleed: a poetic name for a serious vulnerability. The programming error in OpenSSL has caused quite a stir, inside and outside¹ the tech-oriented media. The servers have since been patched, certificates have been replaced and those concerned have been informed. The successors to Heartbleed have already grabbed the public's attention. But where do we go from here? Will vulnerabilities such as Heartbleed be the new 'business as usual', or can something be done? A look back and forward from the National Cyber Security Centre.

⁽¹⁾ CNN on Heartbleed: <http://edition.cnn.com/2014/04/08/tech/web/heartbleed-openssl/>.

Heartbleed in brief

On 7 April 2014, the information security community was alarmed by the disclosure of Heartbleed², a serious vulnerability in OpenSSL that had been in the software for two years. The vulnerability allowed attackers to extract information from the internal memory of a vulnerable computer system. All kinds of information are stored in this internal memory: secret keys of certificates, passwords, customer information, source code of web applications, et cetera.

Using the acquired information, an attacker may execute several types of attacks. With a password, he can penetrate computer systems. With secret keys, he can break the encryption on secure communications. With customer information, many kinds of attack are possible: identity fraud, credit card fraud or sending spam.

Two thirds of the web runs OpenSSL

OpenSSL is open source software: it is the most popular programming library for establishing a secure connection based on the SSL/TLS protocol (the most frequently used protocol for this purpose). Much other software uses OpenSSL: Linux distributions, Android, web servers and firmware of networking appliances. In April 2014, two thirds of all active websites used web server software that depended on OpenSSL³.

Updating and searching for possibly leaked data

Repairing the Heartbleed vulnerability consists of updating to a version of OpenSSL that is not vulnerable. After that, cleanup begins: all secret values that may have been compromised (passwords and secret keys of certificates) should be replaced. If the vulnerable system contains customer passwords, the organization should instruct customers to replace their password as well.

Response: what *do* we know for sure?

At NCSC, we soon decided to write not only an advisory⁴ but also a factsheet⁵ about Heartbleed. An advisory is a piece of technical security advice. A factsheet provides more background. We chose to write a factsheet because of the seriousness of Heartbleed, the complexity of

mitigation measures and the expected attention for the vulnerability within and outside of the information security community. In short: there would be a large need for information about Heartbleed.

On 8 April, information security professionals all started looking for OpenSSL installations within their organization. These turned out to be all around: Linux servers, Linux workstations and Android smartphones were prime suspects. And then there were those appliances and networking equipment: possibly these also used OpenSSL, but who really knew for sure? Could these just be upgraded? Ironically, on the day Windows XP became end-of-life, it was the Linux systems that gave their administrators a headache.

The rumor mill started running soon as well: had the NSA already known about this vulnerability for two years and had they possibly even abused it? Had they maybe even added it to the OpenSSL code themselves? While there is no evidence to support such a claim, it did make sure that organizations energetically started applying their mitigating measures.

Mitigating, such a hassle!

At NCSC, we received a lot of calls and e-mails asking for advice. The question most frequently asked was whether it was possible to detect an attack on Heartbleed from the server logs. The answer to that question is, unfortunately, 'no'. Only by carefully studying network traffic can it be determined that an attack is taking or has taken place. Only an organization that stores its network traffic for an extended period of time can retrospectively determine if it was attacked.

Because it is so hard to detect an attack, the next question was to be expected: is it really necessary to replace all our certificates and passwords? In other words, how likely is the occurrence of such an attack? Not much can be said about the period before 7 April, but after that no doubt could exist. Enthusiasts from around the world wrote their own Heartbleed scanner and started looking for interesting targets. People shamelessly published login credentials of online services like Yahoo to demonstrate that these were vulnerable to Heartbleed.

And our customer data?

The question that most sharply divided information security professionals was the one about customer data. Should organizations instruct their clients to change their passwords, and perhaps also tell them that their personal data could have been compromised? Every organization has made its own decision in this matter, based on its risk profile and the nature of the processed personal data. A

⁽²⁾ See https://www.openssl.org/news/secadv_20140407.txt for the public announcement of the vulnerability. See <http://heartbleed.com> for background information.

⁽³⁾ According to the Netcraft Web Server Survey of April 2014, over 66% of the active websites uses Apache or nginx: <http://news.netcraft.com/archives/2014/04/02/april-2014-web-server-survey.html>.

⁽⁴⁾ At the time of writing, <https://www.ncsc.nl/dienstverlening/response-op-dreigingen-en-incidenten/beveiligingsadviezen/NCSC-2014-0215+1.26+Zeer+ernstige+kwetsbaarheid+gevonden+in+OpenSSL.html> was the most recent version of the Heartbleed advisory.

⁽⁵⁾ See <https://www.ncsc.nl/english/services/expertise-advice/knowledge-sharing/factsheets/factsheet-heartbleed-serious-vulnerability-in-openssl.html> for the NCSC factsheet on Heartbleed.



broad warning for possible identity fraud did not fit the nature of the threat, we decided. Also, such a warning could trigger a panicked reaction.

The months after Heartbleed

A few months have passed, and we are now able to look back and see whether the situation around Heartbleed has improved. To start with good news: many servers have been patched and are no longer vulnerable to Heartbleed. By the end of May, only 1.9% of all internet-reachable servers that offer HTTPS were vulnerable¹. In other words: 98.1% was not or no longer vulnerable. On the other side, even 1.9% of all HTTPS-enabled servers is a lot of servers. It seems like this patch level will not be improved upon much. The graph levels off at 1.5%. The remainder could consist of systems that are not easily updateable, such as embedded systems.

The first major data leaks as a consequence of Heartbleed have since also been reported. Several websites have successfully been attacked. Reportedly, the largest data leak is an attack on a hospital website that netted the attackers 4.5 million patient records².

Shellshock looked like the next Heartbleed (for a minute) A few months after Heartbleed, alarming stories were published about a new vulnerability with a name almost

as beautiful: Shellshock. Bash, a shell that is used in many Linux distributions, contained a programming error in the handling of environment variables. The parallel with Heartbleed was drawn quickly, but was that deserved? On the one hand, Shellshock looked a lot like Heartbleed. It concerned software that was perhaps used in even more computer systems than OpenSSL. A further parallel was the open source nature of both OpenSSL and bash. On the other hand, Shellshock was probably a much less serious vulnerability than Heartbleed. It was harder to exploit Shellshock than Heartbleed, and mitigation of Shellshock was easier.

Shellshock drew a lot of attention, which was surely because Heartbleed was still fresh in our collective memory. For the non-technical media, the difference between the two was hard to understand. The attention for Shellshock undoubtedly led to more prompt and thorough remediation. Still, we see many systems that remain vulnerable to Shellshock. We don't seem to have learned much from Heartbleed.

Three lessons from Heartbleed

The time has come to reflect. Did the Heartbleed vulnerability stand on its own, or is this type of vulnerability the new 'business as usual'? Can we, the information security community, learn something about the way we interact with software, and how we react to the disclosure of such vulnerabilities?

One could argue that the commotion around Heartbleed was all just hype. After all, we have seen such technical vulnerabilities before. However, it would be a mistake to

⁽¹⁾ An analysis of the follow-up after Heartbleed: <https://jhalderm.com/pub/papers/heartbleed-imc14.pdf>.

⁽²⁾ Bloomberg reports this event on the basis of the words of David Kennedy, founder of TrustedSec LLC: <http://www.bloomberg.com/news/2014-08-20/heartbleed-flaw-said-used-by-chinese-in-hospital-hacking.html>.

Stopping the heart bleeding, a practical experience.

Rob Kuiters (KPN CISO)

As an incident response handler you're faced with a challenge when you hear that there's something wrong with the number one encryption implementation. So when the Heartbleed information was published it was all hands on deck.

Let's view KPN as a car, shiny on the outside but the engine under the hood is complex and large. This complexity is something you don't notice from the outside. When you need something of us you go directly to the most evident part of the car, the publicly accessible parts of KPN, one of the brand portals for customers. So the first step we took in mitigating Heartbleed was scanning these portals to see if they were vulnerable.

Once you know what's affected you can determine what actions need to be taken. Technically the solution to fix Heartbleed was rather simple: request a re-issue of your certificate, upgrade the OpenSSL module, install the replacement certificate and inform your users to change their passwords. A single sentence, just a few seconds to put on paper, but a very big operation under the hood of the car.

You can ask yourself why it's such a complex situation then, and for someone who only sees the outside of the car it's a valid question. From a technical point of view it's not really complex but the difficulty is mainly caused by tuning operations and the communication towards our customers. Most of our brand portals are a first point of contact for customers and a vital part of our business, this means downtime has to be limited to a minimum. Still downtime will occur when upgrading OpenSSL to the correct version. Careful and correct communication to our customers about this issue is key in planning mitigations.

The above just pictures what we did with the public customer portals, but of course all of our other services (such as e-mail and VPN services) were inspected and fixed too.

And then there's still the massive internal networks that need attention. A lot of system management is done using a graphical user interface. These user interfaces all have their own web server and were also vulnerable for Heartbleed.

In its entirety the Heartbleed operation took us a month. It was a large operation and it took a lot of effort from the different departments throughout the organization. And what we didn't know back then... It wouldn't be the last time that year...

view Heartbleed as just evidence that programmers make mistakes. Important things can be learned by viewing Heartbleed as the product of an open source community that lacks funding, or by looking at the haphazard way in which organizations tried to remedy the vulnerability. While the technologies were the same, the impact was much different.

Lesson 1: Organizations still hardly have a hold on which software they use

Knowing which software runs where is crucial for an organization. The response to Heartbleed emphasizes that point. Organizations should store this information in a Configuration Management Database (CMDB). When a vulnerability like Heartbleed is disclosed, one only needs to consult the CMDB to know where the vulnerable

software (in this case, OpenSSL) runs, and repair the vulnerability there.

Time and again, this theoretical approach shows to be lacking in practice. CMDBs by no means contain records of all software. What about firmware of devices? The packages used on a Linux server? Programming libraries that are provided with the software that uses them? The software running on an appliance? And what about devices that run software, but that are themselves not even recorded in the CMDB?

The lesson to be drawn is twofold. First: yes, organizations should administer their CMDBs more accurately. They form the basis of any information security policy. Second: it is not realistic to think any CMDB is complete and perfectly accurate. Actively searching for vulnerable software, for example with a scanner, is an important additional security measure.

Lesson 2: The information security community hardly knows which code is critical

Heartbleed originated through a programming error in the heartbeat functionality of OpenSSL. The heartbeat functionality is not complex: it is the SSL/TLS equivalent of 'ping'. It maintains the connection in absence of other traffic. Still, it was possible for this small mistake to have an enormous impact. Had not enough people perused the code? Probably not.

Large differences exist between the criticality of code for the security of systems¹. Most code does not play a role in security and is not reachable to attackers. However, apparently not all critical code receives the attention it deserves. Yes, if you try to change the code that implements AES in OpenSSL, you will be noticed. Much other code does not receive this attention, while programming libraries such as libcurl, libxml or libjpeg can be very reachable for attackers as well. Recognizing such critical code, and acknowledging its role in securing the internet, are important steps towards reinforcing the foundation of security.

Lesson 3: Donations currently do not work to fund open source software

Like many other open source software projects, OpenSSL depends on donations. With these, the team can further develop the software. Until recently OpenSSL received about US\$ 2000 per year in donations². Furthermore, organizations paid the OpenSSL team about one million US dollars per year for programming work, mostly custom implementations for these organizations. In other words, this software plays an important role in securing networks and the internet, but hardly receives any funding to fulfill this task.

Several technology companies have since pledged a few hundred thousand dollars in extra donations for the coming three years³, but no sustainable solution is in sight. The open nature and free availability of the software make it hard to compel organizations to contribute financially. Nonetheless, such contributions seem necessary to ensure software quality in the long run.

And now? Onward.

Heartbleed has occupied our minds for quite some time. The vulnerability has drawn an unprecedented amount of attention in technical and non-technical media outlets. How will the information security community improve the situation in the coming years? By its actions, it influences whether vulnerabilities like Heartbleed become 'business as usual'. If the community is able to learn lessons from Heartbleed, we may yet make the internet and our networks more secure.

About the author:

Pieter Rogaar works as a senior advisor at the Nation Cyber Security Centre. He is the author of the factsheet on Heartbleed. His specialization is cryptography, and his interests include privacy and IT law. As an advisor he writes publications for knowledge sharing such as fact-sheets, whitepapers and the Cyber Security Assessment Netherlands. Pieter gladly pushes buttons without a label, he is politically active and talks much and often about the future of the internet.

⁽¹⁾ I have taken this idea from Dan Kaminsky's blog: <http://dankaminsky.com/2014/04/10/heartbleed/>. I hope he will excuse me for borrowing a good idea from him.

⁽²⁾ Steve Marquess, president of the OpenSSL Foundation, has described the situation in a blog: <http://veridicalsystems.com/blog/of-money-responsibility-and-pride/>.

⁽³⁾ Source: <http://arstechnica.com/information-technology/2014/04/tech-giants-chastened-by-heartbleed-finally-agree-to-fund-openssl/>.



Yo, Trust me

Jeroen van Duuren (KPN CISO)

In this day and age people are regularly confronted by con men. In cyberspace this often takes the form of phishing. Phishing is a methodology used by cyber criminals whereby they send you an e-mail or text message requesting you to either answer a few questions and return the answers by e-mail, or visit a web page where they expect you to submit information in a form.

These questions usually relate to your personal details, credit card information or account login details (such as webmail or social media websites). Most companies communicate with their (potential) customers using a variety of different media and in a non-consistent manner. How can a person distinguish legitimate communication from imposters? In this article I'll try to sketch the issue in an understandable manner, explain the dangers and suggest solutions (and yes, also technical ones).

The problem

Marketing and sales departments in companies can get very creative and want to get things going as fast as possible. That's why they often launch websites independently from the main company website, or send a mass mailing using a distributor of their own choosing which sends e-mail from a separate domain name (this is especially a big issue in large companies with multiple

marketing departments). How is a consumer to know if the message really originates from the company that claims to be the sender?

I'm a customer at the company where I'm also an employee, and even I can't tell if certain communication is real or fake! The fake ones are so good that there's no way for me to verify this, so I'll just have to call the helpdesk.

Example: The main website for Example, inc. is example.com, but sometimes people are redirected to websites such as examplediscount.nl, custhelp.com, or example.net. Or they receive an e-mail from example-mailing.nl. It's not clear for the customer what's really Example, inc. and what's not.

The dangers

So, what do these fraudsters want with your information? That depends on the kind of information they're after. Basically you can divide this into three categories: identity information, payment information and online credentials.

Identity information

Using personally identifiable information a fraudster can impersonate a victim and perform transactions on his or her behalf. E.g. they can order stuff and don't pay the merchant, get a credit loan, open a bank account, start a business, etc. all in the victim's name. This leaves the victim with the problem of proving that he or she wasn't the one that actually performed these actions.

The latest public information about identity theft stems from 2012. In this year there were an estimated 612.000 cases of identity theft in The Netherlands. The estimated total amount of damage as a result of identity theft between 2007 and 2012 was € 280.000.000 (that averages on € 56.000.000 a year).

Source: "2013-update onderzoek 'Omvang van identiteitsfraude & maatschappelijke schade in Nederland'" – PWC, May 2013.

Payment information

This is rather stating the obvious, but criminals are always interested in payment information. For example, they steal bank account information and use this to transfer money they stole from other compromised accounts to (offshore) bank accounts or online payment channels such as PayPal.

Online credentials

Cyber criminals are often trying to steal your username and password for different websites. They use this information mostly to send spam and malware via e.g. your e-mail account or via your social media channels such as Facebook and Twitter. Think of it, who are you more likely to trust when clicking on a link; an e-mail from a random organization or a message on Facebook from a friend?

According to statistics from the social networking website Facebook an estimate of 600.000 accounts are compromised every day by cyber criminals. These accounts are mostly used to send spam and malware.

Source: "600,000+ compromised account logins every day on Facebook, official figures reveal"
– Graham Cluley (Sophos)

The solution!

It's important that companies take their responsibility in enabling people to verify if digital communication is trustworthy. The solution is rather easy to explain but hard to implement for larger companies. I've split this technical solution up into two parts: websites and e-mail.

Websites:

This one is rather easy and logical. How can a (potential) customer easily identify if a website is from a certain company? Help them by using the main domain name in the website. E.g. examplecampaign.nl redirects to campaign.example.com or example.com/campaign. One other thing to add to this is always use TLS with company verification or extended validation. This way someone can visually verify that a website belongs to the expected organization (in the address bar of their browser, as shown below).



The big challenge with this is the culture of marketing organizations. These are often fixated on time-to-market and don't always have time to think about the consequences it may have if they make a quick, insecure and non-verifiable website (e.g. reputational damage to the brand or financial damage for customers if they're scammed).

It's important that marketing and sales departments become aware of the marketing consequences that untrustworthy communication can have on the entire

organization and its customers. This can often be accomplished by organizing awareness campaigns and including specific rules and guidelines in the brand communication handbook. Another option is enforcing proper behaviour by implementing a very strict policy that states that all websites that display the logo of the company and are not part of the main domain name will be taken down using the legal framework of copyright infringement.

E-mail:

The way most marketing and sales departments directly communicate with (potential) customers is by using e-mail. How can someone verify that an e-mail is trustworthy and not a scam without calling the helpdesk? Again, always use the main domain name in the sender address. But that doesn't quite cut it with e-mail, because SMTP, the computer protocol used to send e-mail, is inherently insecure by default and easy to manipulate. This means you should implement extra technical measures that secure the entire verification chain.

This chain starts with DMARC (Domain-based Message Authentication, Reporting & Conformance). DMARC is a method that expands on two other mechanisms; Sender Policy Framework (SPF) and DomainKeys Identified Mail (DKIM). Both of these mechanisms are designed to combat e-mail abuse. With DMARC you basically define a policy for mail servers so that they know what to do with e-mail that fails SPF and DKIM checks. DMARC relies heavily on the Domain Name System (DNS) to function. Because of that it's vital that the DNS infrastructure is also protected against tampering and fraud. The way to do that is to implement DNSSEC (a cryptographic security extension for DNS) and thus completing the chain of verifiable trust.

The ING bank implemented DMARC in 2012. They reported results rather quick. There was a decrease of phishing mail pretending to be from the ing.nl domain of 71% within the first six weeks after implementation.

Source: https://www.ing.nl/nieuws/nieuws_en_persberichten/2012/09/xs4all_en_ing_samen_valse_emails_te_lijf.aspx

The very last step to fortify this chain is to let the mail server of the company communicate securely with the mail server of the client using TLS. This ensures that a personal message to the user isn't intercepted by a third party, but above all makes sure that a message isn't tampered with during transit. To be complete, only SMTP with <<STARTTLS>> is not the answer to this. <<STARTTLS>> is a command that is issued after a connection has been set up. To implement this in a correct way one has to enable a fully TLS secured transmission from the beginning: implement and enforce SMTPS.

One final issue with this, is that to implement DMARC you need all e-mail to go via trusted servers that can sign e-mail with a trusted private key. In big companies however, you often have multiple marketing and sales divisions that use dozens of mailing providers. So, these providers all need to implement the security features or start using a mail proxy that is supplied by the company.





Roeland van Zeijst (National Police)

Did you floss today?

It will end your nightmares

In 2014, the National Police issued several warnings directed at business, government and civilians about the dangers of ransomware and its even more evil twin cryptoware. The solution is easy. But you have to implement it before you become a victim. And you have to do it on a regular basis. Like flossing. Can you guess what it is?

Police Nightmare

Imagine it's early in the evening and you are returning home. A police officer is standing outside, accusing you of a crime you did not commit. The cop (or is she an imposter?) demands that you pay her hundreds of euros before you are allowed to enter your home. And there's another catch - you have to pay her in magic pebbles, which can only be bought in the magic pebble store on the other side of the river, ten miles from where you are. If you don't pay, you will be taken into police custody.

Classic police ransomware blocks access to the boot procedure of a computer and shows police logos, implying this inconvenience is a punishment for misbehaviour online. To transfer and obscurify funds quickly, the victim is told to pay in exotic ways, such as bitcoins or game vouchers:



The first large wave of such ransomware hit the Netherlands in 2013, as we described in the previous Cyber Security Perspectives report. To counter this eerie phenomenon, Dutch police made it one of their ‘high tech crime priorities’ in both 2013 and 2014. Pan-European police investigations have been carried out and gangs were arrested. Also, through large publicity campaigns, citizens were informed about the triple scam behind police ransomware: first, the Dutch police do not monitor innocent people’s online activities; second, they aren’t allowed to block someone’s computer like this; and third, a computer showing police logos rarely is completely blocked.


Most ransomware can quickly be removed with anti-virus software, but some users might need assistance with this. Since ransomware often demands exotic tender, it tells its victims to go outside, to a shop or post office, and buy a special voucher code. Especially for challenged and elderly victims, this was a confusing, surreal, and sometimes scary experience, because they also felt they could not call the police about this. Fortunately, many shops and post offices supported the police campaigns against ransomware scams. Now it seems that less and less people are willing to pay up.

Glass Nightmare

Imagine you wake up in the morning, stepping out of bed and finding that the floors, walls and ceilings of your home are all made of glass. Someone else is on the other side of your house (how long has this guy been there???), going through your things. He's wearing blackshades. When he sees you've woken up, he grabs your favourite wall-art and runs out.

Barefoot you go after the man, but a large piece of paper catches your eye. It contains a pre-printed ransom note, with your name filled in on a dotted line. You have until noon to pay, or the item, your favourite, will be destroyed.

Criminals using the remote access trojan Blackshades can do anything they like to a computer, such as monitor all data traffic, the screen, your microphone and webcam. They can also encrypt files and demand money through a popup screen that actually comes with a pre-filled extortion message:



The article ‘Pest control’ describes how easy it can be for small-time cyber crooks to abuse your system’s resources using, for example, Blackshades. When the extortion message pops up on your screen, you know instantly that you have lost not only your most important files, but also your privacy.

Even worse, victims are being addressed directly by their extortionist, knowing that all of their actions are being monitored by a human who is obviously of bad intent – and whom they can chat with on-screen. On the victim's machine, the pop-up window actually has no 'close' button. It is the perfect start of a horror movie.

Customer-friendly Nightmare

Image you're at home and decide to have a look in your family photo album to relive those wonderful personal moments. When you go to grab the album, you find it has been replaced by a metallic vault. There is no lock, no keyhole, no way to open it. Suddenly, your phone rings. A friendly voice asks how you are doing and if you would like some assistance. You don't really understand at first. Assistance? 'Yes, to help you get to the photo album we took from you.' Then, tauntingly: 'It's now or never!'

Unlike police ransomware, cryptoware does not try to hide its purpose. It states boldly that your personal files have been encrypted and that you have little time to pay up, or your link to them will be severed. Some criminal groups provide excellent 'customer' service. Still, decryption often fails:



Combating crime is an arms race, even when it comes to cybercrime. Perhaps due to the success of anti-ransomware campaigns, professional criminals have weaponized the power of cryptography into fully-automated cryptoware. Scariest of all is the blatant cynicism with which they will actually provide some 'online customer support' to help their victims pay up and (maybe) get their files back.

The Dutch police force advises against paying; giving in to extortion will never make an extortionist stop, and in many cases decryption attempts turn out to be unsuccessful.

Mother of all Nightmares

Imagine sleeping and having your worst nightmare. You know, THAT one. Then, imagine waking up, gasping for air, soaking wet, your heart racing. It's the middle of the night. You try to remember the nightmare, but already can't. It's gone. You switch on a light. Carefully, you step out of bed, convincing yourself that your house is not made of glass. Slowly you move around, checking your wall (wall-art is there), your photo album (no vault, just wonderful pictures), and you peek out of the front door to check that there is no-one standing there. Phew. With a sigh of relief you feel how your body relaxes. You walk back to bed, smiling. When you attempt to climb in, you find that someone is in your spot. It's... it's... you?! You're lying there, sleeping, shaking, sweating. Now you're crying, and screaming in your sleep. Wait – oh wait, you remember now. You haven't woken up at all! Because you can't. You can never wake up from this. You are trapped in your own worst nightmare. Forever.

Many organisations nowadays have smart and continuous backup services, that are easily accessible through their company's network, which is constantly being imaged. Which is wonderful, except during a cryptoware attack. Cryptoware destroys all the data it can find, including on-line backups! Not all convenience is progress. Several companies have run into trouble due to not being able to recover from a cryptoware attack, because the backup had fallen victim to the extortionists as well.

Throughout 2014, Dutch police have noted a further increase in cryptoware attacks and issued several warnings. Unfortunately, during the fall a new strain emerged, dubbed TorrentLocker:

TorrentLocker has been found to harvest e-mail addresses, enabling the criminals behind it to prepare new attacks, perhaps using new false identities. It might





not look as slick as some of its predecessors, but it seems to be operated by clever criminals using vicious spamming techniques to seduce victims into executing their malware.

Many people fell for urgent e-mails 'from a debt collector' or 'from the post office', containing an 'innocent' Word-attachment... Now, you might think that people who fall for such tricks are stupid. Then, to be honest, you should know that some police personnel in training fell victim to this one. You wouldn't call the police stupid, would you? (otherwise we will come and lock your computer)

There is a way to end all these nightmares. It is like flossing. We have to do it on a regular basis. It is somewhat tedious. It is always in danger of being forgotten or skipped due to a busy schedule. But everytime we do it, we feel a bit better. About ourselves. And about our cyber security perspectives. (so it might even be a little bit better than flossing)

Please, from now on, **make external backups of your vital data and then unplug the backup facility.**

This will enable you to do a rollback after an infection has hit your business. And it really can hit you, since no-one can predict how it will present itself next time. Remember: a backup system which is continuously online and updated, is in danger of becoming encrypted as well. Suggestion for smaller users: burning vital backups onto write-only dvds might be a quick fix that is actually not so tedious.

Please, make this your 2015 New Year's resolution and follow through.

You will sleep so much better for it.

Also, don't forget to floss.





CyberDawn, training for a cybersecurity crisis

Nathalie Lokhorst and Oscar Koeroo (KPN CISO)

With everyone hyper-connected, cyber risks are everywhere. Retailers breached. Intellectual property stolen. Systems hacked almost on a daily basis. This makes us dependent on each other. What if a hack would take place which affects us all? How must we deal with this situation? As we can't prevent the possibility of such incidents we will have to prepare ourselves.

On Wednesday October 29th 2014 the biggest Dutch cybersecurity drill was held called "CyberDawn". This cybersecurity drill was initiated by the telecom participants of Nederland ICT, hosted by KPN and facilitated by Deloitte. The purpose of the drill was to understand how to work together optimally within the telecom sector on

a nation-wide crisis which involves the government and other partners in the vital infrastructures.

What is Nederland ICT?

Nederland ICT is an association of more than 500 ICT-companies in The Netherlands. The members



represent a turnover of almost 30 billion euros and have over 250.000 employees combined. Therefore Nederland ICT is the representative of the Dutch ICT sector.

CyberDawn

CyberDawn consisted of two parts: a strategic crisis management drill and a technical drill in the form of a Capture The Flag (CTF) contest.

Strategic crisis management drill

How to deal with a large-scale crisis? That's the starting point for the strategic crisis management drill. During the drill a scenario unfolded which hit the telecom sector first and escalated into a national crisis. This crisis affected organisations in the energy sector and involved several departments and ministries, NSCS and police. The purpose of this is to train the participants in developing their own vantage point by conceiving clear objectives, take their responsibilities and focus on their leadership capabilities.

The plot of the scenario has been engineered in such a way that the best possible outcome can only be achieved when all parties properly work together. It is crucial to understand each other's roles and responsibilities. Roeland van Zeijst (National Police): "I found it very important to understand that lots of telecom-participants had never heard about the ICT Response Board. Generally speaking you can say that the public and private parties have learned more about each other's crisis structures. We will now know how to find each other during a real crisis".

After an intense day the findings were shared privately between the participants before everybody moved towards the CyberSalsa.

Scenario:

A large set of sensitive company data is published in phases on the internet.

Threat profile	
Category	Data loss
Actor	Criminal organisation Hacktivists
Motive	Data theft Cover up
Impact	Data loss Trust in telecommunications sector

The telecom sector receives messages that sensitive data is in the possession of a rogue party. This involves both privacy-related data and company-sensitive data, like infrastructure data. The dataset is made online available piece by piece. It appears that it is not only data from the telecom sector but also from the banking and energy sectors. During the scenario it became apparent that the data had been leaked through a 'zero-day weakness' (zero-day weakness, is a weakness that is discovered today, there was no knowledge about it and there is no available countermeasure) in a network component. This piece of equipment is not only used by the telecom players but also by other vital infrastructure in the Netherlands. Various parties appeared to be actively making use of this weakness. During the course of the scenario the exploitation of the zero-day turns out to be a cover-up for other activities, building up to a true national crisis.

Martijn Ronteltap (Deloitte – facilitator CyberDawn crisis management drill)

The key take-away for the participants is to get to know each other. For example understanding their position, each other's responsibilities in an organisation and with this mutual understanding to solve the crisis. In a crisis situation it is important to know who you need to call. Knowing who to contact makes it possible to operate at maximum efficiency to solve a crisis. But to be clear: one crisis drill is not enough! This exercise shows that both private parties (telecom, energy and banking sector) and the public sector (government, NCSC, police) are taking mutual responsibility for vital infrastructure in The Netherlands. For me it was striking to see that so many different parties where at the table to actively participate in this crisis drill.

Capture The Flag (CTF)

CTF is a training where 'techies' receive several different technical challenges which they need to solve like puzzles. The CyberDawn CTF consisted of offense, defence and cryptographic challenges. The members of CERT teams and Red Teams of the participating companies and organisations were grouped into several teams.

Those teams consisted of maximum 5 people and competed against each other in a virtual environment.

The teams tried, while enjoying Club-mate (a typical long-lasting energy drink that is consumed at hackers parties), to break into poorly protected servers. At the same time they had to defend their own servers from a continuous stream of attacks. Also, they had to solve cryptographic puzzles.

Whenever a challenge was solved, one or more new challenges would appear which had to be solved. For every correct solution the teams gained points which were plotted live on a scoreboards. The number of points you earned was in proportion to the difficulty of the puzzle solved. At the end of the day the team with the most points was declared the winner.

The CTF led to a cosy atmosphere which allowed the participants to share common interests and knowledge.

Jaya Baloo (KPN – Chief Information Security Officer)

There is a proverb, 'praemonitus, praemunitus' which translates roughly to "to be forewarned is to be fore-armed". In relation to cyber security awareness, it means that people will have a natural advantage over cyber criminals simply because they had advanced warning. It should be considered as the beginning of a culture shift to get the balance towards the tipping point for an informed and thereby armed public and corporate defence.

The CyberDawn cyber crisis exercise is an example of a first time ever event in the Netherlands that promises to improve intersectoral cooperation while at the same time allowing opportunities to test and hone hacking skills with a Capture-The-Flag event.

CyberSalsa

How do you move 'techies' after an intense day in sync with the rhythm of a strategic crisis management drill? The answer is CyberSalsa!

A salsa instructor brought his ladies and challenged everybody to join the CyberSalsa. After a hilarious start, several participants joined in and showed they could move more than just their fingertips on a keyboard. The spectators made nice pictures and videos of these moves. After a day of intense concentration, this was a nice way to loosen up and let it go.

An excellent way to close a most excellent day!



Rob Kuiters (KPN CISO)

On Her Majesty's Secret Service GRX & A spy agency

“Have you seen the news on Belgacom today? This looks really nasty!”. The rumours and conspiracy theories flew across the office on the 17th of November 2013 when the news hit the Internet.

As time elapsed more information around the hack was available via the media. Adding to that, new revelations of Edward Snowden gave some more clues on what the stories were all about.



This was not the NSA, as earlier media publications suggested, behind the infiltration of the Belgacom network. The newly revealed documents suggested that it came from about 550 kilometres from Brussels,

Cheltenham, home of the Government Communications Headquarters (GCHQ). This British intelligence and security organisation is responsible for providing signals intelligence and information assurance for the British government.

Also the documents revealed the target of the hack, the GPRS Roaming Exchange network, in short GRX.

“So, why this specific part of the network? What kind of information resides there that could be of interest to intelligence agencies? How vulnerable is the network anyway?”, these were the questions that lead us to investigating one of the largest, if not the largest global mobile IP backbone network which connect all mobile networks together, GRX.

Before diving into our research, let's first take a small tour around mobile data communication, “there is a whole world out there!”

Mobile Data Network 101

In the late 90's mobile networks around the globe were enhanced with a packet switched feature add-on. Circuit switched data was already possible but couldn't exceed the 9600 baud limit. By adding a specific part on the radio link and adding three more core elements the mobile network was able to carry an IP stack across. General Packet Radio Service, GPRS, was an enhancement on the second generation mobile network GSM and could reach an ultimate speed of 40kb/s. Besides the complete new radio specifications a new mobile core protocol was introduced, the GPRS Tunnelling Protocol (GTP).

A connection from a mobile device to its destination is completely tunnelled throughout the mobile network. In this manner mobility management for this specific connection can be achieved. GTP is used to tunnel the connection across the mobile core network. There are a few identifiers in the network that make this possible. From a mobile device point of view the first thing we encounter is the mobile base station, BTS, or nodeB for 3G. Base stations are connected to a radio controller station, a base station controller (BSC) for 2G or a radio network controller (RNC) for 3G. The controllers are the boundary of the radio domain within mobile networks. The setup for 4G is slightly different, the base stations are called e-nodeB's and the controller functions are incorporated in the base station.

What kinds of identifiers are important to the network for setting up a connection? First of all your mobile device or actually the Subscriber Identification Module (SIM).

This module resides on a Universal Integrated Circuit Card. The way this card is identified by the network is by the International Mobile Subscriber Identity (IMSI). This is a worldwide-standardized fifteen digits long unique number. With the first three digits the country is identified. For example 204 is used for the Netherlands and 234 is used in the United Kingdom. This part of the IMSI is known as the Mobile Country Code (MCC).

The next two or three digits is the Mobile Network Code (MNC). Only North America uses a three digit MNC.

Within a country there are likely more than one mobile network and different networks will be distinguished by the MNC. The KPN network has MNC 08, T-Mobile in the Netherlands has MNC number 16. The remaining part of the IMSI is known as the Mobile Subscription Identification Number (MSIN). All IMSI's for a particular operator are stored in a big database also known as the Home Location Register (HLR). This mobile network element provides authorisation and authentication. It also holds the different mobile services per IMSI. So if your subscription would hold voice and data you would be allowed to roam then the services would be attached to your IMSI in the HLR.

The base station is the access to the mobile network. Base stations are numbered uniquely through the network. In a certain area those base stations are grouped together. This is known as a routing area. A group of base stations is controlled by a Base Station Controller (BSC). This is the element name for 2G. In 3G they are known as Radio Network Controllers (RNC).

A group of BSCs or RNCs are connected to the core network. In the core network three elements take care of the mobile data connection. The Serving GPRS Support Node (SGSN) is the 'link' between the radio domain and core domain. The Gateway GPRS Support Node (GGSN) is the gateway towards any kind of packet data network. A SGSN handles the traffic of a certain group of base station controllers. SGSN's and GGSN's are connected to the operators mobile IP backbone network.




Figure 1 – High level mobile data network architecture

Tracking through the network

How do we keep a mobile data connection connected through the network? To keep track of a mobile device in the network, each mobile base station has a unique identifier on the network. A group of base stations together is called a routing area. The whole network is built-up of these routing area's which are also known in the core DNS of the mobile data network. Each routing area has its own IP address. When a mobile device roams through the mobile network, handovers are performed from base station to base station within a routing area. When a routing area is crossed, the network will ask the DNS which IP address goes along with the new routing area so the connection can be continued.

To setup a mobile data connection, also known in mobile terminology as a 'Create PDP context request,' your device has to have an Access Point Name (APN) configured in the mobile device settings. Access Point Names are also known in the DNS. The IP address going along with the APN is the IP address of the GGSN where the APN is configured. APN's can be configured on multiple GGSN for redundancy reasons.

When a mobile data connection is started, the SGSN will resolve the APN at the DNS and will setup a tunnel towards the destination GGSN using the GPRS Tunnelling Protocol. In the IP stack GTP resides between the transport layer and application layer and is based on UDP.

Currently GTP has three versions. The first version, GTPv0, is used in the beginning of mobile data. GTP version 0 uses port 3386. The control traffic, traffic which handles the connection, as well as the user data is not separated in this version. The succeeding version, GTP

version 1, splits the control and user traffic. It uses two different UDP ports for this purpose. The control plane uses port 2123 and all user data in the user plane is carried across port 2152. With the arrival of 4G only the control plane got a new version, GTP version 2. Messages are sent between the SGSN and GGSN to control the mobile data connection. Those messages are defined in the 3GPP specifications. Create PDP context request is one of those messages. In total there are around 20 of those kinds of messages. Each message has specific Information Elements, which are needed to manage the connection across the mobile data network.




Figure 2 – PDP Context Setup and Delete

When you use your mobile device outside of your country to check your mail or to use a messenger services the mobile data connection will make use of a roaming partner of your own provider. It uses the radio network and the SGSN of the visited mobile network. The outgoing connection, the GGSN towards the external packet data network, will still be in your home network back home. But there seems to be an 'air gap' between the visiting SGSN and the home GGSN. To connect all mobile networks worldwide the GSM Association introduced a global IP backbone network for mobile network use only, the GPRS Roaming Exchange (GRX). The GRX network is an extension on the local mobile IP backbone and makes it possible to connect all the mobile data network elements globally. There are just over a few dozen global GRX providers mostly tier one carriers. It uses the same technology that the internet is built upon. The GRX network is separate from the internet and until now only allows access to mobile operators. As a rule mobile network can be no more than three hops away from another operator; GRX providers don't transit traffic from one another.

This roughly defines the mobile landscape we are looking at, a global IP network connecting mobile operators around the globe to make mobile data roaming possible.

The spy agency's perspective – a honeypot

So if you were to gain access to this network, what could you possibly see here? Capturing a full session in a pcap like fashion would give you a full GTP flow. With network forensic tooling like Networkminer or Explico you will see the frames but not the direct content data within the GTP user plane. But it is not difficult to strip off the GTP layer, with a little bit of scripting you will end up with a regular IP session, which can be handled by standard network forensic tooling. Like in any other plain IP session you will be able to see all of the content visited by the mobile end user. Every user session with apps which don't use any kind of encryption will give you visited websites, plain text username and password combination for your mail server, ftp sites or logon to sites. And without strong encryption on user plane level all content would be available during offline analysis. All this information resides in the GTP user plane.




Figure 3 – Wireshark capture Create PDP Context

Ever since mobile operators started with mobile data roaming, the volume of traffic increased every year. It would not be efficient even for an intelligence agency, to capture all traffic streams. But there are ways to pin point the traffic more specifically. For example by capturing the GTP control channel. In every create PDP context request end user information is available.

As stated earlier, the GTP created PDP context request holds a few interesting end user information elements. In this screenshot of the expanded packet view in Wireshark you can see the IMSI, the unique identifier for the mobile end-user and the location information written down as Mobile Country Code, Mobile Network Code and more specific the radio location. Also the IP address of the GGSN, the MSISDN number and the IMEI of your mobile device. An IMEI stands for 'International Mobile Equipment Identity'. It is like the Mac address for the mobile device interface. This number identifies your mobile device. With some open source information tools you can make some more sense of these information elements. For example, making use of a location service as OpenCellid.org or underwirelabs.com will give you the location of the base station from which the connection was setup making use of the routing area code and Cell-ID. Moving around in the mobile network will result in a Update PDP Context Request for handing over to a new base station. Enough information to track someones movements across a cellular network and no fancy LI equipment needed!

Another element useful to an intelligence agency is the IMEI. In this case the IMEI starts with 01147200, this is the part which tells you what kind of phone is used here. Again, lots of open information sources will give you the information that goes along with this number. In this case it would be an Apple iPhone. Information useful for a targeted device attack against an iPhone. All information to see what a user is doing with a mobile device and also where the mobile device is in the network can be extracted from the GTP protocol. This is probably interesting for any intelligence agency.

Hacker tactics – grabbing the honey


What tactics would intelligence agencies use to acquire this information? To get to the 'honey' some 'bears' have to be defeated. We used the kill-chain methodology to plot our 'honey' stealing plan. The kill chain is originally a method used in the military. It is a tactic used to engage the enemy, Find, Fix, Track, Target, Engage and Assess. Lockheed Martin's incident response team used this tactic to proactively remediate and mitigate internet threats. There are different stages how a hacker approaches an attacks. Reconnaissance, Weaponization, Delivery,



Figure 4 – Plotted location based upon mobile identifiers

Exploitation, Installation, Command and Control, Actions and Objectives also known as Exfiltration.

In the reconnaissance phase Research, identification, and selection of targets are the objective. For example, crawling Internet websites for e-mail addresses, social relationships, or information on specific technologies. The weaponization phase is used to create a workable exploit by combining a trojan (get past defences) with



a malware payload constructed to accomplish the attacker's goals. With the delivery the transmission of the weapon to the target is done. Popular attack delivery vehicles are e-mail attachments, websites, and USB removable media. The next phase, exploitation is where the malware weapon is delivered and the payload activates exploiting a vulnerable program or system. Installation will take care of a backdoor on the victimized system allows the adversary to maintain contact. Attackers typically require manual intervention to explore the victim's network. This is accomplished by the malware contacting a remote command and control server. This is known as the command and control phase. The final phase, if everything goes according to plan, the attackers now pursue the reason for the intrusion, possibly compromising additional servers or exfiltration of data.

For our research we limited ourselves to the first stage. Only reconnaissance, just identification of hosts and enumerate the services. The GRX BGP routing table holds 48.000 subnets, roughly 320.000 IP addresses. To have results quickly we used the MasScan tool to discover live hosts.

Out of the 320.000 IP addresses we found 42.000 live hosts who replied on the scan.

To determine if they are really mobile elements we used zmap to find the open GTP UDP ports. To be confident that they are actually running a GTP daemon behind it we took the SGSNEMU tool, a SGSN Emulator, which is a part of the OpenGGSN project. We used the open ports found with Zmap. With the SGSNEMU we sent out a GTP ECHO REQUEST message, a ICMP-like control message to verify the connection between SGSN and GGSN. Live GTP daemons will respond with an ECHO RESPONSE message. As a result of this scan we found 770 GTP Control Plane ports (UDP port 2123) and 1042 GTP User Plane ports (UDP port 2152).

'Well, that is obvious!' you would argue. Those ports are meant to be found from within the GRX network. But the stunning fact is that they were found from the internet. Those ports should never be reachable outside from the GRX network! Besides these ports, we found lots of other stuff reachable from the internet. Some with severe known vulnerabilities. We reached out to the global mobile community and informed them about our findings. We identified and noticed 15 mobile operators who had open GTP ports. Only 5 of those operators acknowledge our findings, this is sad but true.

Can this be prevented? The answer is simple, yes! It is very obvious these IP ranges should not be announced via the BGP advertisements on the internet. Keep your networks separated and patch your systems!

Also the GSMA is working on improving the guidelines for setting up GRX / IPX connections. They recently released an update of the Inter-Operator IP Backbone Security Requirements For Service Providers and Inter-operator IP backbone Providers, that includes best practices. We think lots more work needs to be done in this area. For example, how a mobile operator can test the security posture of their GRX / IPX by themselves.

Closing words

For us this was a deep dive into mobile core networks. But our work is far from finished. We will pursue our investigations and research further to ensure a more secure mobile world.

Tools used in this research:

OpenGGSN project

<http://sourceforge.net/projects/ggsn/>

Location API

<http://unwiredlabs.com/api>

<http://opencellid.org/>

Numbering plans

<http://www.numberingplans.com/>

<http://www.mcc-mnc.com/>

Network Miner

<http://www.netresec.com/?page=NetworkMiner>

Enum4linux

<http://labs.portcullis.co.uk/application/enum4linux/>

ZMAP

<https://zmap.io>

Masscan

<https://github.com/robertdavidgraham/masscan>

GTP specification

<http://www.3gpp.org/DynaReport/29060>



Spotting the prowler

Early detection and rapid response for targeted cyber attacks

Sander Degen (TNO), Bert Jan te Paske (TNO)
and Rob Vercouteren (KPN CISO)


In 2014 numerous high-profile cyber attacks again caught the public eye: ransomware campaigns, compromised ATM machines, and disclosure of celebrity nude pictures acquired from hacked cloud accounts. Less noticed, but possibly more worrisome to governments and private companies, is the continuing threat of targeted attacks (also known as APTs - Advanced Persistent Threats) against specific organizations. Examples are the Russian hacker group dubbed Sandworm that recently infected NATO, and the widespread use of the tool Blackshades to take over computers - read more about that in article 'Pest Control'.

Introduction to targeted attacks

Contrary to conventional computer security issues such as worms, viruses, ransomware and botnets, targeted cyber attacks are aimed at specifically selected individuals or organisations. Typically the goal is to access private and confidential information. (Industrial) espionage, credit card fraud, monetary theft and hacktivism are the primary drivers. As the target is so specific, the attackers can focus all their efforts on trying to breach the security measures that are put in place and (if they are successful) in maintaining access for as long as possible in order to capture more information.

In this article we will focus on those targeted attacks that aren't easily mitigated by conventional means. What are the unconventional means that can keep us safe? First we need to understand how these attacks take place. A typical targeted attack consists of several steps (experts call this the 'cyber kill chain') that are depicted in their chronological order in the following figure.


Chronological order:



The steps can be explained with the following example: A group of attackers use search engines, social media, and web forums to gather personal information of a person working for the intended target organisation (Reconnaissance). They use this information to create an

email that seems authentic enough for the victim that he or she will open the attachment. This attachment itself is malicious: it contains a hidden program that installs a remote access tool on the victim's computer, and special care was taken to make it undetectable by current - up to date - virus scanners (Weaponization). The e-mail is sent to the victim (Delivery). Sure enough, the victim opens the attachment, the malicious code gets executed (Exploitation) and the remote access tool gets installed (Installation). The remote access tool connects to the Command & Control server that the attackers have set up (Command & Control). They can now monitor the victim's every action on that computer and gain an understanding about the programs that are in use and where information is stored. With this, they can begin to search for confidential information, export it out of the organisation's network and into their greedy fingers (Action on Objectives). This last step may require taking control over other computer systems in the same network, e.g. a file server or mail server, when the desired information is not

directly available to the victim's computer. This is called *lateral movement*. Keep in mind that other scenarios are possible, e.g. attackers gain access to a financial server/application



and change IBAN-numbers in an address book or transfer money directly to banking accounts under their control.

State-of-the-art detection methods

So how can a targeted cyber attack be detected? The aim should obviously be to note the early attack stages and take responsive action before the attackers have accessed and exported the information they are looking for.

Current detection solutions use different approaches to meet this challenge in terms of what is being monitored and how suspicious behaviour is defined.

A Host-based Intrusion Detection System (HIDS) is a combination of tools that monitor a computer system: are unexpected programs running in the background, are there numerous failed attempts to log in to an application, does an application try to access files it shouldn't?

A Network-based IDS (NIDS) is a system that analyses network traffic. It can detect scans that try to reveal the active programs on a system. Another example is repeated attempts to log in to such a program. Detection can be based on a superficial level, where only IP addresses, communication ports, number of packets and the number of bytes are given, or through inspection of the individual packets.

In both types of IDS, different approaches can be taken to detect suspicious behaviour: signature based (matching a previously created signature) and anomaly based (deviation from the norm).

Examples of defensive tools that employ these techniques are antivirus software, log analysers and Security Information and Event Management (SIEM) tools.

Detection challenges

In combating targeted attacks we often deal with (groups of) skilful attackers well able to mislead regular detection products. Take for example antivirus products, which are commonly trusted to provide sufficient protection against malware infections. These products primarily work signature based, meaning they will only raise alerts when they encounter a malware payload that was earlier added to their virus database. In recent years it has become easier, even for less sophisticated attackers, to disguise known malware samples using packers and other tooling to modify their appearance while preserving the hideous effects. This has led Symantec's senior vice president to state in 2014 that 'antivirus is dead'. Even if this is debatable in general, as it is still effective in blocking older viruses, it is probably true for the case of detecting high-end targeted attacks.

Anomaly-based detection to the rescue? The idea of looking for suspicious activity that stands out from normal activity is appealing. Unfortunately, anomaly-based detection has its own limitations. In particular: it only works when a baseline of what is considered normal can be accurately defined. This may be feasible for specific protocols such as online banking transactions, SCADA communications or DNS traffic. However, defining a workable generic baseline for office IT networks or Internet traffic is an illusion as these are highly diverse with even regular users generating large amounts of anomalies.

In general, we believe it will prove very difficult to prevent the delivery and initial exploitation steps.

We are convinced that determined and resourceful attackers will eventually succeed in taking control of a workstation inside the network perimeter.

Once this has been achieved, the attackers can pose as a regular system user and mask their actions as normal system and network behaviour. A detection system is then faced with the challenge of spotting subtly hidden malicious activity in a huge volume of system and traffic data generated by genuine employees and their systems. On one hand, this involves a big data challenge: how to process and store all relevant logs and network traces. On the other hand, as the KPN-CERT team experiences in practice, the detection system must be continuously tuned to balance a good detection rate (not missing out on signs of an actual attack) and a low false positive rate (not flooding a human analyst with alerts for benign activity). In typical office IT environments with few restrictions for its users this is more difficult to achieve than in military networks with hardened workstations, few applications, and tightly controlled network boundaries with proxying and filtering policies in place.

In common IT networks, attackers have numerous ways to hide their outbound Command & Control and data exfiltration traffic in innocuous-looking outbound connections such as e-mail or web traffic. With encrypted (https) web traffic becoming commonplace, analysing traffic payload based on deep packet inspection is made increasingly difficult.

While a large research effort is being made to develop technology for detecting Command & Control and exfiltration traffic, we doubt that focusing on these attack steps individually will be effective.

New directions

So things are looking grim. Neither signature-based detection nor anomaly-based detection can prevent

these tenacious opponents from getting a foothold in our network. They will get in and they have ample opportunity to export valuable information once they succeed in finding it.

While this illustrates why targeted attacks constitute such a complex challenge, we believe it is still possible to put up a defence, holding on to the aim of early detection and response before the attackers have taken their 'action on objectives'.

Instead of focusing on the first attack stages of delivery, exploitation and installation, we believe a stronger effort should be made to detect the lateral movement and data collection attack steps. After compromising one host system, the attackers will want to prowl the network, and gain access to other systems looking for information of interest. This generates network traffic and log events that are atypical for the compromised host system (or for the active user account, in case such a correlation can be made). Anomaly based detection is more effective here, especially since the network protocols used to explore a network and communicate with servers are well-defined and can be analysed for protocol anomalies.

Examples are SQL database communication and access control handshakes.

The detection of lateral movement can be aided by employing honeypot systems, (virtual) computers configured to look like operational servers or file systems that are often made easy to break into. When performing network scans attackers will note such a tempting target and try to access and explore it. These actions are immediately suspicious as no authentic user would have a reason to communicate with the honeypot system or even know about its existence.

In addition, improving the detection rate by combining different detection sources and techniques - not just aggregation but correlation beyond the level of current SIEMs - can reveal traces of attacks that would go undetected by individual detection tools. Together with rapid response, this has the potential to thwart targeted attacks.


Automated forensics

Once an attack is detected, rapid responsive action may prevent the attackers from achieving their goals. One approach is to streamline the forensic analysis process,



At the Black Hat security conference in Las Vegas, two researchers reveal that Open Mobile Alliance Device Management (OMA-DM), supported by many **smartphones**, can be exploited over the air using a false base station.

Hold Security states that it has obtained over 1,2 billion passwords, stolen from over 420,000 websites. Much stays unclear about the origins of the data set, as well as Hold's online service where people may enter usernames and passwords (!) for a check.



thus reducing the time required by an analyst to investigate a malware infection. In the past few years we have seen advances in the forensic investigation of memory dumps.

The volatile (RAM) memory of a computer contains the active applications (and related information such as ownership, contents, and permissions), current Windows registry settings, and active user profiles. By ‘dumping’ the current snapshot of the memory into a file, an archive is created that can be analysed offline. Analysis involves identifying anomalies, usually within a single dump. State-of-the-art techniques increase the amount of information that can be obtained and gives a better chance at detecting the presence of targeted malware. One example is the automatic creation of memory dumps - e.g. when a SIEM notices something might be wrong. This may save hours in precious response time. Additionally, the forensic analysis can be improved by comparing a memory dump with previous memory dumps from the

same system or similar systems. By highlighting the differences, it should be much easier to pinpoint the malware, allowing for a faster and more accurate reaction.

Perspective on 2015

Any organisation working with valuable and sensitive information should consider themselves prey for the attackers discussed in this article. Furthermore, with the current state-of-the-art in detection products no organisation can be sure they haven’t already been targeted and compromised. In 2015, TNO and KPN will continue to develop and employ new approaches to early detection and rapid response. In the meantime, each organisation should make their own effort to build an accurate picture of what is happening on their networks and systems. Good logging and live forensics not only support adequate detection and response but can also aid police in investigating the threat against your organisation and, even internationally, pursue the perpetrators.

Shedding light on the Dark Web*

Mark van Staalduinen (TNO) and Roeland van Zeijst (National Police)

In cyber security, co-operation is essential, we are all well aware. But what about our adversaries: where do cyber criminals meet? And why are they avoiding the Netherlands more and more?

In November 2014, the illegal TOR market place Silk Road 2.0 was taken down, as well as its alleged owner, 'Defcon' a.k.a. Blake Benthall (26). After being in business for exactly 1 year, Silk Road 2.0 approximately had 150,000 active users and generated millions of dollars of revenue per month. Benthall by then drove a \$127,000 Tesla car that he had paid for in bitcoins. On the last morning he drove it, he didn't get out of his driveway. 'Defcon' found himself surrounded by




^(*) This article discusses several suspects in current criminal court cases, who of course are legally considered to be innocent until proven otherwise.

20 armed FBI agents. Immediately after his arrest, the digital drug lord confessed to being the owner and operator of Silk Road 2.0. The alleged kingpin might end up in prison for life. An evil empire collapsed, partly because the owner had made the mistake of hosting it, for some time, in the Netherlands.¹



Alleged digital druglord 'Defcon' is facing life in prison because he chose the wrong hosting country

But what does Benthall's arrest have to do with cyber security? Well, Silk Road 2.0 was one of the so-called darknet markets within the TOR hidden services, where users tend to think they are completely untraceable and anonymous. Darknet markets mainly offer drugs and weapons. Some marketplaces even provide human trafficking, hitmen and child abuse. Less well-known is the trade in illegal cyber products and cyber services, e.g. DDoS attacks, espionage services, ID fraud, botnets, etc. This is one of the ways cyber criminals meet, team up and exchange tricks, tools and bitcoins.



DDoS Attack Service offered by 'Hackyboy' on anonymous marketplace Evolution



Darkweb hackyboy ddos attack

Taking down Silk Road 2.0 (and a number of even worse illegal marketplaces simultaneously) was a huge hit to criminals operating on TOR. Operation Onymous² is not the first action of Law Enforcement against the darknet markets, but by far the biggest and best coordinated one until now. Over 404 'onions' were taken down and illegal traders were arrested worldwide. There were many questions about how Law Enforcement was able to track down so many illegal marketplaces and their (ab)users.³

Silk Road 2.0 was called '2.0' because until late 2013 there had been a previous marketplace called Silk Road, named after the legendary trade route in ancient Asia. It was (allegedly) operated by Ross Ulbricht,⁴ who is standing trial for, amongst committing other crimes, having paid a hitman \$500,000 to murder six people.⁵ Shortly thereafter, darknet markets Black Market Reloaded and Utopia were seized, the latter as part of Operation Commodore by the Dutch police. A total of five people were arrested and half a million euros worth of bitcoins was seized in the Netherlands and Germany.

As a consequence of this sequence of takedowns and arrests, the traders spread their risks and started to act on multiple marketplaces at the same time. In 2013 less than ten marketplaces were active, while in November 2014 approximately twenty marketplaces⁶ were in the air.

⁽¹⁾ Dutch police investigated Silk Road 2.0 when its servers were temporarily located in the Netherlands, and were able to confirm Benthall's identity. They also were able to offer proof to the FBI that the webserver under investigation, and operated by Benthall, was running the criminal TOR hidden service Silk Road 2.0.

⁽²⁾ literally; 'not anonymous'

⁽³⁾ We will explain this another time.

⁽⁴⁾ The FBI claims that Blake Benthall was his second-in-command.

⁽⁵⁾ It seems that none of the murders ordered have taken place. Rip-offs are another big risk of darknet markets.


⁽⁶⁾ <http://www.dnstats.net>; before the mass takedown, the number of active marketplaces was well over thirty.



Message of the FBI and European law enforcement agencies after taking down Silk Road 2.0

Darknet market traders are getting more and more professional. Some traders, we estimate, have a turnover of millions of euros per year. Access to these marketplaces is simple and traders reach a global scale immediately, their effectiveness augmented by the provision of secure payments (bitcoin) and international escrow services. Discussions in the forums illustrate their increasing professionalization. Traders only discuss business and take good care of their reputation.

As we explained before, darknet markets tend to offer not only goods, but also criminal services. This is one of the places where cyber criminals meet and team up. The now defunct Rent-a-Hacker webshop had a very user-friendly shopping cart system that even the most computer-illiterate buyer could use:



Service catalogue of Rent-a-Hacker


Dealing with such marketplaces poses fundamental problems. Forbidding networks like TOR, Freenet or I2P is not a viable solution, because this would negate Dutch net neutrality and moreover it would eliminate the positive side of such networks from a worldwide human rights perspective.



Logos of several anonymization networks

But still, given the fact that traders have turnovers of millions of euros per years, whilst people (and perhaps your networks or computers) are falling victim to crime, we need to act. Therefore, we are gathering information to better understand what is happening on the darknet markets. For example, is this a typical Dutch problem? What geographical attributes can be identified?


Based on online discussion topics on darknet markets, TNO technology provided data, which was then analysed in co-operation with the Dutch Prosecutors' Office. This resulted in a picture of European countries that were mentioned on darknet markets more often (red) than others (green).



Geographical analysis of mentions of countries and capitals on darknet markets

For the emerging pattern, one explanation might be that the 'red' countries have the facilities and infrastructures to enable many people to mine for bitcoins. This is the most important way of payment on the marketplaces.¹ Another reason might be that attractive cybercrime victims can be found in the 'red' countries, for which specific online banking malware can be traded.

Finally, some countries are mentioned a lot because criminals are constantly debating their modus operandi. Dutch police have found several examples of criminals stressing the need to avoid the Netherlands, or even explaining to each other how to avoid the Dutch police.² Of course, most of the darknet criminals are not physically acting from within the Netherlands. But when they are abusing Dutch infrastructure, the National High Tech Crime Unit (THTC) might choose to track them down.




Cyber criminals discussing their modus operandi, trying (idly) to avoid law enforcement

From a cyber security perspective, our common challenge is to prevent and disrupt criminal business cases. To develop and implement this strategy is not only a government task: the responsibility is shared by companies and scientific institutions alike. So, also your organization can join in!

There is good news here. Unlike with zero-days, it is possible to know what is for sale, due to the fact that products and services are offered in open forums. Next step will be to enhance the information into a global picture, so it will be possible to judge the potential of new criminal products or services.

Invitation

This innovative approach requires deep understanding of the newest malware, the cyber underground, and will also apply big data analytics to process the enormous amount of information into actionable intelligence. TNO, National Police and partners are working to make the Netherlands more resilient, and to strengthen cyber security as a whole. Therefore, we will be shedding more light on darknet markets and surrounding phenomena. This can be achieved by public, private and science partners cooperating closely. It is a challenge we are happy to embrace together with you.



⁽¹⁾ It is good to know that bitcoin transactions are less anonymous than is often presented, partly due to the transparency of the transaction table (blockchain).

⁽²⁾ Some of these criminals are, however, now awaiting their trials in the Netherlands. Others will too, in 2015.



Ahead of the Threat

Enhancing Cyber Intelligence Communities

Richard Kerkdijk (TNO) and Michael Meijerink (National Cyber Security Centre)

There has been much roar about “threat intelligence” and its purpose in the ever changing security landscape. In this article we discuss the need to collect and manage such threat intelligence, the limitations of current intelligence sources and our initiatives with respect to threat intelligence sharing.

The need for threat intelligence

The landscape of cyber threats is rapidly evolving. New vulnerabilities emerge at a tremendous pace and these vulnerabilities are increasingly qualified as severe¹ (see figure). What’s more, present day cyber attacks are more sophisticated than ever. State of the art malware is greatly autonomous and employs specific stealth techniques to avoid detection. High end attacks are persistent and targeted and involve elaborate combinations of methods

⁽¹⁾ Cyber Security Assessment Netherlands 2013 (CSAN-3), <https://www.ncsc.nl/english/services/expertise-advice/knowledge-sharing/trend-reports/cyber-security-assessment-netherlands-2013.html>

and channels, ranging from specific technical exploits to social engineering of critical staff. On top of all this, it is evident that the attackers are increasingly organised by actively collaborating, sharing tools and techniques and offering services to one another.



Five year vulnerability disclosure trend, source: Secunia²

In the midst of these developments, the dependency on ICT and thus the potential impact of any security incident is ever increasing. Due to the dynamics of present day cyber threats, organisations cannot passively rely on traditional (preventive) measures. To avoid unnecessary damage, they must continually stay on top of the latest threats, vulnerabilities, attack methods and attacker campaigns. To this end, organisations are in need of appropriate threat intelligence.

Many definitions exist of the term “threat intelligence”, some accurate and some not so accurate. The authors believe that the definition put forward by Gartner³ covers the essence rather well:

Threat intelligence is evidence-based knowledge, including context, mechanisms, indicators, implications and actionable advice about an existing or emerging menace or hazard to assets that can be used to inform decisions regarding the subject’s response to that menace or hazard.

This definition reveals some important characteristics of proper threat intelligence, among other things that it should combine bare technical indicators (e.g. potentially harmful IP addresses) with contextual information

(e.g. concerning adversaries and their motives) so that organisations can appraise its relevance and (potential) impact. It also needs to be actionable, i.e. facilitate tangible follow-up actions, and timely.

Threat intelligence is not only vital for ICT-intensive organisations seeking to maintain a solid level of cyber resilience, but also for bodies coordinating cyber security on a national level such as the National Cyber Security Centre (NCSC) in The Netherlands. For such entities, sharing threat intelligence can be instrumental for monitoring changes in the threat landscape and predicting major cyber threats that might have a disruptive effect on society. Notably, such “situational awareness” requires specific data and insights not typically needed for increasing the resilience of individual organisations. We will come back to this in the final section of this article.

The value of threat intelligence is ultimately determined by the ability of organisations to assess its relevance and impact and bring the intelligence to action in operational security processes such as firewall management, (security) patching and incident monitoring. Most organisations have yet to become mature on this aspect.

Limitations of common intelligence sources

There exists a great variety of sources to acquire threat intelligence. We distinguish the following categories:

- **Company internal sources.** Technical security solutions in use within an individual company may offer a wealth of insights with respect to emerging threats. Relevant examples include Intrusion Detection Systems (IDSs) and SIEM⁴ solutions. Suppliers of such solutions usually feed them with intelligence of their own to keep them aligned with the latest threat developments. Apart from this, internal staff can also be a valuable source of threat intelligence, especially if the company has invested somewhat in their security awareness.
- **Public sources.** Institutes such as SANS offer threat intelligence through their public websites and many vendors of security solutions publish periodic threat reports. Some national CERTs collect, analyse and report threat intelligence to stimulate the resilience of organisations in their country. In The Netherlands, the NCSC does so through advisories, factsheets and a specific application called Taranis⁵.

⁽²⁾ http://secunia.com/vulnerability-review/vulnerability_update_all.html. Enquiry at Secunia revealed that they expect even higher numbers for 2014 (exact figures still pending at the time of writing).

⁽³⁾ <https://www.gartner.com/doc/2487216/definition-threat-intelligence>

⁽⁴⁾ Security Information and Event Management

⁽⁵⁾ <https://www.ncsc.nl/english/services/incident-response/monitoring/taranis.html>

- **Commercial sources.** Companies such as Mandiant, Secunia and IBM supply threat intelligence on a commercial basis. This usually takes the form of a subscription based service.

This wealth of sources presents security practitioners with an overwhelming amount of data in which it is hard to assess what is truly relevant for their specific organisation and business. Apart from the sheer volume, much of this data is unqualified and lacks proper context (e.g. where has the threat been sighted, what impact is anticipated, who are the attackers and what are their motives, etc.). This greatly complicates the identification of appropriate responsive actions. At the very least, the task of structuring and rationalising this flood of incoherent data would require a substantial investment in specialised staff.

Notably, some of the most interesting threat intelligence can (at least for a certain period) be bound by restrictions. Organisations struck by a severe breach of security will initially be reluctant to share the corresponding intelligence. What's more, information concerning state actors or new exploits is usually reserved for selected parties such as international intelligence agencies and major software companies. Threat intelligence of this nature will appear fairly late – if at all – through the above channels.

All in all we see that common sources of threat intelligence, whilst valuable in their own right, come with some intrinsic limitations..

The promise of intelligence communities

To overcome the aforementioned issues, the authors see great potential in the concept of threat intelligence communities, i.e. networks of organisations that exchange threat intelligence amongst each other. We believe that there is strength in numbers and that communities for threat intelligence sharing offer some appealing benefits:

- **Trust among peers.** Participants in a threat intelligence community can develop trusted relationships, thus stimulating openness in information sharing and offering access to intelligence that might not be shared with public or commercial sources.
- **Leveraging of capabilities.** Through the intelligence community, experts at individual organisations could effectively build on each other's analyses and insights, thus strengthening each other's work rather than duplicating it. As mentioned above, the attackers are collaborating so it would make sense that potential victims do the same.

- **Operational insights.** Within a trusted network of peers, intelligence shared could extend to operational insights such as the impact and effect of response and mitigation strategies. Again, rather than trying the same mitigation approach in parallel, organisations could assess measures proposed by their peers and use their own capabilities to enhance these further.

The concept of a threat intelligence community is not entirely new. Many security teams already take part in a community of some sort through which they gather threat related insights. Prominent examples include the global Forum for Incident Response and Security Teams (FIRST¹), the European Government CERTs (EGC²) group and the European community of telco CERTs organised under the umbrella of ETIS³. On top of this, certain threat intelligence is exchanged in industry oriented Information Sharing and Analysis Centers (ISACs). An example is the FS-ISAC⁴ for sharing critical security threats in the global financial services sector.

What is good about existing intelligence communities is that they all represent a certain level of trust among the participating organisations. At the same time, however, present community efforts are often inefficient. Threat intelligence might for instance be shared through mailing lists and discussed in conference calls or on-line forums. This will ultimately not suffice to keep pace with the continuous stream of threats and vulnerabilities (see above). What's more, the actual intelligence shared is often insufficiently actionable, in part because it is too general in nature and in part because it is not aligned with the operational (security) processes of individual organisations.

Thus, whilst a foundation of threat intelligence communities exist, some shortcomings need to be resolved to exploit their full potential.

The merits of automation

To enhance the speed and efficiency of threat intelligence communities, it would make sense to introduce some standardisation, automation and dedicated tooling. This need is widely recognised and has led to some interesting developments. Firstly, as already indicated in the "Trends to watch" article, various protocols and standards to facilitate the exchange of threat intelligence have emerged. Here, the framework put forward by the

⁽¹⁾ <http://www.first.org/>


⁽²⁾ <http://www.egc-group.org/>

⁽³⁾ http://www.etis.org/?page=CERT_SOC

⁽⁴⁾ <https://www.fsisac.com/>

MITRE corporation⁵ is of particular interest. This package of standards encompasses three distinct elements (STIX⁶, TAXII⁷ and Cybox⁸) that can greatly facilitate automation in the exchange of threat intelligence.


There are some appealing aspects to the MITRE framework that distinguish it from competing standards. Most notably, it has strong momentum. Industry leaders such as Microsoft have committed to employing these standards and vendors of security solutions (e.g. SIEMs) are incorporating them in their products. The aforementioned FS-ISAC also embraced STIX and TAXII for sharing threat information. Such wide adoption can ultimately enable an integrated eco-system in which intelligence shared amongst peering organisations is easily mixed with other sources and easily processed in company internal security systems. Another attractive feature of the MITRE standards is that they are process oriented and that the underlying body of thought not only encompasses collection and sharing of threat intelligence but also the subsequent analysis and effectuation within individual organisations.



Screenshot of the MISP application

On the tooling side, the Malware Information Sharing Platform (MISP) has been receiving quite some attention in Europe. MISP⁹ was developed by CERT-EU, Belgian Defence, NATO and CIRCL to facilitate their own threat intelligence community. It's primary purpose is to store and exchange so called Indicators of Compromise¹⁰ (IoCs, see figure). It also supports correlation of events

and export to popular formats such as Snort and OpenIOC. MISP is available as open source and can thus also be embraced by other communities. Efforts are ongoing to make MISP comply with the aforementioned MITRE standards.



IoC and underlying data model for a malicious URL scenario¹¹

An interesting characteristic of IoCs is that they are dynamic in nature. An IoC for a certain attack scenario might initially be rather rudimentary and be updated and refined as the attack is further analysed. In a community context, IoC updates might be compiled by (security specialists of) different companies, thus effectively combining insights and capabilities across the community.

Upcoming initiatives

Both NCSC and TNO are preparing initiatives with respect to threat intelligence sharing. Whilst these initiatives involve different objectives and target groups, both seek to utilise the potential of (automation in) threat intelligence communities.

In 2012 the NCSC and the Dutch Intelligence and Security Services started the preparations for a threat intelligence sharing initiative among public and private organizations involved in critical infrastructures. These organisations also rely on intelligence that is not publicly available and can only be shared within a highly trusted, public and private, cross-sectorial network. Key attributes of this initiative are transparency in the approach and processes and high demands on privacy measures. In 2013, a pilot at a national data centre proved successful for sharing threat intelligence among individual organisations as well as creating near real time insight into the severity of emerging threats. As of December 2014, the community will be expanded with more private organisations. Information shared will include operational experiences with actual threats and the intelligence compiled will also be used to generate a national operational

⁽⁵⁾ <http://www.mitre.org>

⁽⁶⁾ Structured Threat Information eXpression – a structured language for cyber threat intelligence information.

⁽⁷⁾ Trusted Automated eXchange of Indicator Information – a transport mechanism for cyber threat information exchange

⁽⁸⁾ Cyber Observable eXpression – a structured language for cyber observables, i.e. standardised format for specifying so called Indicators of Compromise (IoCs)

⁽⁹⁾ <https://github.com/MISP/MISP>

⁽¹⁰⁾ Artefacts observed on a network or computer system that indicate intrusions with a high degree of confidence, e.g. a file or memory pattern

⁽¹¹⁾ Source: <http://stixproject.github.io/documentation/idioms/malicious-url/>

picture about threat developments. Notably, the network is not just about sharing indicators, but has the potential of becoming a knowledge base where members actively work together to improve each other's resilience.

In the first half of 2015, TNO will conduct a threat intelligence pilot with a group of European telecoms providers. This project will take place under the umbrella of ETIS and feature 4-6 participating companies. The pilot essentially seeks to professionalise the already existing ETIS community of telco CERTs and CSIRTs (see above). The project will streamline the exchange of threat intelligence among telcos in terms of the exact information that requires sharing and the corresponding conditions and rules of play. Information feeds will be automated by means of the aforementioned MISP platform and participants will jointly assess both the quality of intelligence exchanged and the subsequent effects in their company internal security processes. Ultimately, this small scale pilot aims to lay the foundation for a more elaborate operational setup involving many of the (20+) ETIS member telcos.

There are some interesting parallels between the above initiatives. Both involve organisations that maintain nationally critical infrastructures and both have embraced the same technical platform to facilitate the information exchange. There are also some distinct differences, most importantly:

- Whilst the NCSC pilot is cross-industry on a national level, the telco pilot is industry specific but has a European scale. The authors believe that intelligence communities such as these can coexist and even

mutually interact, thus establishing a network of communities.

- Through its role of national CERT, NCSC can feed its community with specific threat intelligence that individual organisations could not easily acquire.
- Contrary to the telco initiative, the NCSC community also has a "situational awareness" component. Specifically, NCSC wishes to correlate intelligence from participating organisations to identify major cyber threats with a potentially disruptive effect on society.

Notably, both the telco pilot and the NCSC community are based on voluntary participation. Whilst this is rather obvious for a private sector initiative, many governments prefer a supervisory approach. In the Netherlands, however, the larger part of the (critical) IT infrastructure is owned by private organisations and the NCSC has neither the desire nor the mandate to enforce threat intelligence sharing among such parties. As it turns out, collaboration on a trust basis requires a larger investment in establishing and maintaining the network but in turn also facilitates intelligence sharing at greater levels of (operational) detail.

Both initiatives also hope to enhance the maturity of participating organisations in processing the intelligence received and applying it in operational security processes. Supplying threat intelligence in structured formats that align with company internal processes should serve as an effective catalyst to this end.

NCSC and TNO will actively exchange their insights and experiences with respect to intelligence sharing. The intention of this collaboration is to enhance the success of both initiatives.





Wesley Post (KPN CISO)

Honeypots at KPN

During the last year KPN has developed a Honeypot setup for their infrastructure. This article describes the basics of a honeypot and some details about KPN's setup.

Honeypot 101

Honeypots are systems that look like vulnerable systems and are set as a trap for hackers searching for systems to attack. The goal of such a setup is to collect information about what attackers try to do. This information can then be used to set priorities in other processes. For example if you have a vulnerability in OpenSSH and Apache, and you see that the one in OpenSSH is actively exploited, a logical choice would be to give priority to upgrading OpenSSH. Another use of honeypots is to see what attackers actually do, this might give you valuable information on how to mitigate the attack. Theoretically honeypots could have detected Heartbleed attacks, an attack which was hard to detect by regular means. Of course this is theoretical, since at that time no honeypots were available which could have provided the level of logging required to do so. In case of KPN there is another advantage. Of course there is interest in what happens on the network. Simply monitoring isn't an option since

this involves capturing customer's traffic, which is only allowed in very specific cases. Also, when handling encrypted protocols like SSH, monitoring the network won't reveal the actions of the attacker. With a honeypot there is no customer data involved at all, so there are no legal issues to monitor and log every aspect of it. There are different kinds of honeypots, grouped by how well they look like the actual service. These levels are called low, medium and high interaction honeypots. Low interaction honeypots may just give you an open tcp port or the banner of a vulnerable service, medium interaction honeypots go through a significant effort to actually make the attacker believe he is connected to the actual service. High-interaction honeypots are usually the actual service, but configured in a way that it is still vulnerable without causing actual harm.

Analogy

To understand what honeypots are doing in the digital world, it's good to compare it to something in the real world. In the case of honeypots we can compare it to the police trying to catch bicycle thieves using fake bikes. A low interaction honeypot would be comparable to a cardboard bicycle. From a distance it might look interesting and attract thieves, but once they get close they'll notice it's fake and non-functional. However, you could still register who's interested without any risks. A medium interaction honeypot would be comparable to a real bike, but all parts are welded together. So if you come closer it still looks like a bike, but once you actually try to use it, it turns out to be useless. In this case you could register who's interested, and which techniques were used to open the locks. Of course the risk is also a bit higher, but there is not much to lose. A high interaction honeypot would be comparable to a real bike, and everything works this time. However, the bike is also equipped with a GPS tracker and a microphone. In this case you know who's interested, how he opened the locks, and what he's doing with the bike. Of course this poses the highest risk, but it also gives us the most information with a fair chance to find back the bicycle.

Honeypots

Honeypots are often started as a project and developed as Open Source/Free Software. There are some excellent projects out there but most are old and/or poorly maintained. The standard tools used in these less active projects continue to evolve but the scripts / programs used are not. This will eventually cause compatibility problems when you try to use these old scripts / programs. In that case that project will need modifications to make them work again but some do not work at all. A number of protocols have no honeypots available, for example DNS and NTP. However, having a honeypot for those services

does make sense because they can be abused for DDoS reflection attacks. Instead of writing a honeypot from scratch, it is possible to use real servers and patch them to get the logging you need. This approach was used for the KPN DNS and NTP honeypots.

While the ISC NTP daemon can be configured to log a lot of data, this data is only relevant to it's normal use, timekeeping. For a honeypot it's much more interesting to see what's happening via the network. For this, KPN patched this program to log all commands and opcodes that it receives via the network. This logging allows us to see what techniques attackers are trying to abuse NTP in reflection attacks.

KPN's Setup

When considering a honeypot setup the main question is: "Is a honeypot secure?". While the realistic answer is probably that "It depends ...", the only safe approach is to assume that they are not. This implies that the honeypot process has to be run in a safe, constrained environment where little to no damage can be done whenever the process turns out to be insecure. This is why KPN's setup is done using Linux Containers (LXC, more details on that later). Each process will get it's own container where only a very limited environment is available. Any log data will immediately be transferred via syslog to another container for storage. The setup also encompasses a routing container. This takes care of NAT between the individual honeypot containers and the outside world. Additionally it will continuously run tcpdump to capture any traffic to the honeypot to study what has been done in case the honeypot itself does not provide enough information.

The honeypots currently in use at KPN:

- Kippo (ssh)
- Deception ToolKit (telnet, ftp, pop3, smtp)
- Wordpot (http/wordpress)
- Samba (cifs)
- ISC NTP (ntp)
- ISC BIND (dns)
- xrdp (remote desktop)
- snmpd (snmp)

The picture shows the typical setup within a physical server. You can clearly see the separation done with containers. The connections between the containers (especially the honeypots and the router/collector)

are based on bridged interfaces on the host OS. As an extra security measure the bridges are limited and only forward traffic from/to the collector and router, i.e. communication between two individual honeypot containers is not possible.


Linux Containers

Linux Containers (LXC's for short) are used extensively within KPN's honeypot setup. This technique is not very well known, although it has been available for a few years already. LXC's are easily described a "chroot on steroids". Where chroot limits a process' access to the file system, a LXC applies similar principles on other system resources like memory, CPU time, network interfaces and processes. When applying such limits and combining all those separated bits, you can create the illusion of a virtual machine. Of course it isn't full virtualization. Isolation is a much better term to describe it. One of the major differences is that it doesn't have it's own kernel, the benefit on the other hand is that you don't

need a full isolated system to run an isolated process, usually the files needed to run a process in a chroot jail are enough to run a process in a container. Using LXC's makes the system maintenance much easier since you don't have to manage a lot of virtual machines, instead you can copy the required files from the host machine into the container. With this process automated it's a matter of updating the host system and re-generating the containers and all your LXC's are updated.

Location

When deploying honeypots, much like real-estate, a key factor is: location, location, location. The location in your network that is. It's good to have a clear idea of where you want to place your honeypots before you even start building them, because in the end, the location will greatly influence how useful your honeypot data will be. At KPN we have a few physically separated networks with different functions, so based on the functions we selected the locations and honeypot servers we need to be active.





Jesse Helder (KPN CISO)

Help, the app ate my password!

On a gloomy day in January with the office still buzzing of fresh new year's resolutions we received an e-mail from a journalist¹ containing a link to a website. While expecting an interesting article or a scandalous new story my eyebrows went up half a meter when I looked upon the contents displayed on my screen. Before my eyes were more than a thousand credentials of users of the Telfort website. And even worse, a test of some of these credentials actually proved them correct as well.

Now this is odd... why would a hacker keep stolen credentials on a publicly accessible website indexed by Google and easy to find for anyone? And even more important, how did he manage to get these passwords? In

this article I would like to give you an insight as to how we tried to answer these questions, but also to point out the risk involved in the millions of easy smartphone apps available.

⁽¹⁾ <http://www.destentor.nl/algemeen/binnenland/inloggegevens-klanten-telfort-op-sstraat-door-nep-app-1.4176214>

During its annual congress, political party **Piratenpartij** becomes the victim of a hack. Its entire website is compromised, including the member list. Piratenpartij, aiming for more transparent government, quickly issues a list of compromised database fields.

Stopping the bleeding

Our first concern was of course to solve this operational issue. Since the accounts listed on the website were all compromised we had to reset the password of all accounts involved to prevent active abuse. We also had to warn the involved users about our findings. While we were busy doing this we actually noticed a remarkable fact: during our investigations and password resetting the list had grown by two accounts. This made it clear that however the accounts were compromised the evildoers were still actively harvesting account data. This made it clear we needed to find the source of this information fast!

Finding the culprit

One of the first traces to follow in this case was to find out who actually owned the domain containing the website. Unfortunately this did not lead to a quick resolution as the domain was registered to a domain registrar and did not yield any data on the actual owner of the domain. At the same time it was also investigated where the actual server hosting this content was located. We quickly found out this machine was located in France and was rented from a large hosting provider, so this too did not give a quick result on where to find the culprit who stole this data.

By then we caught a lucky break by studying the other content found on the webserver containing the compromised account data. In one of the script files also publicly available, we found a person's name inside one of the variables used in the code. When we did an online search we quickly found out that someone by this name published apps in the Microsoft appstore for Windows Phone. And one of the apps was called "Telfort Abonnement Status". This app proudly announced itself as the unofficial Telfort App for Windows Phone to give users insight into how many minutes, texts or MBs of their subscription they had already used. Although the app description declared this was an unofficial app, the app itself was riddled with Telfort logos so it looked official enough to convince people to install and use this app.

Using the name and developer ID in the Microsoft app store we were able to track down and contact the actual developer of the app and inform him of possible legal steps concerning this issue.

An investigation proved how this app actually worked. When installed on a Windows Phone the user would supply their credentials for the Telfort website. The app would then log in on the website using these credentials and get the users subscription data to display these on their phone. However we found out that when you first put in your credentials the app would upload those to the file on the mystery webserver we had found earlier. So

we had found the source of the leaked passwords and we also found a stick to beat it out of the Windows App store. Because the developer had used Telfort logos without our consent, he had violated the conditions of the Windows app store and thus a request to take down the app was quickly honoured by Microsoft. This effectively stopped the growth of the compromised account list that we started this incident with.

After we had tracked the person responsible for this app we filed an injunction against him for using trademarks without permission. As a result he reported himself to our front desk and wanted to explain what had happened. According to him he was let down that Telfort had not released a Windows Phone app that enabled him to check his balance. For that reason he decided to write such an app himself by scraping the website of Telfort. Because he ran into problems during the testing phase he had performed logging of data on the webserver we had found. Although we never will be able to tell what his real motivations were, it had already occurred to us that logging credentials on an open website was not a smart plan if he would want to actually use these credentials. Whatever his intentions truly were, we were able to stop the leaking of credentials, additionally we actually eliminated several more fake branded apps from the Windows Phone store.

Wider concerns regarding apps

What is worrying is that if the developer of this app had not been so silly to log the credentials on a publicly available website, it would have taken us a lot longer before we had found out what was happening here and the damage could have been a lot more substantial. This gave us the valuable lesson that it is worth the effort keeping track of apps released with your company name as they might not actually have been released by yourself.

Telfort is not the only brand suffering from users being duped by fake apps. A recent study by Trend Micro² showed that there are over 900.000 apps available that are fake. Fake here means that they mimic an app that already exists. The same study showed that 51% of these fake apps actually have a malignant component and although some of them are just very aggressive adware most of them have more sinister purposes like stealing credit card details.

In most cases these malignant apps can't be downloaded from an official app-store. Google, Apple and Microsoft have a vested interest in keeping their stores clean and

⁽²⁾ <http://www.trendmicro.com/cloud-content/us/pdfs/security-intelligence/white-papers/wp-fake-apps.pdf>

trustworthy to maintain their user base. So most people are duped into installing apps from an unofficial store offering a fake version of a paid premium app for free. Although this requires the user to switch off a security feature on his phone, the lure of a free app seems to be enticing enough to ignore all warnings and go ahead. However as can be seen in our case and also in other cases app removals are often done on base of complaints after the app is published. This means that downloading from an official app-store is not a 100% guarantee of downloading a secure app. Therefore it is always wise to check if the publisher of the app listed in the store is indeed the official publisher of this app. By performing this check one can prevent being duped into downloading a fake app that slid through the official store controls.

There is however another interesting category of fake apps that are available in official app stores. Apps that promise you the world but actually do nothing at all. A

good example in 2014 was the VirusShield app, this was available in Google's play store. This so-called anti-virus app did not do anything at all besides showing you a very comforting green checkmark on the screen indicating your device was completely "protected" and of course charging you 4\$ for the privilege of downloading it. This app was downloaded more than 10.000 times before it was banned from the app store by Google. Minimal investment for maximum profit.

This again shows that cyber criminals are seeing the potential of profit by exploiting mobile devices and that this threat is indeed growing in size and severity. The first botnets based on Android have already been found and used in the wild¹. Although the mobile tsunami of doom and gloom as preached by mobile anti malware vendors has not hit us yet, the waves are certainly picking up. So it is surely wise to be extra vigilant when installing the latest app on your phone. By following a few basic rules one can avoid the most ominous threats:

- Do not install apps from unofficial app stores.
- Install a well reputed anti-virus app on your phone.
- When installing an app, always check if the publisher listed in the store is the same as the publisher listed on the internet.
- Never install apps advertised to you by SMS, WhatsApp or other sources by an unknown party.
- Perform regular software updates of your smartphone as they often include security updates.



⁽¹⁾ <http://www.zdnet.com/first-case-of-android-trojan-spreading-via-mobile-botnets-discovered-7000020292/>
<http://www.zdnet.com/android-botnet-poses-as-google-app-pilfers-email-and-sms-7000024495/>

Overview contributing partners



KPN is the largest telecom and IT service provider in the Netherlands. Our network is Dutch to the core. We have a clear mission – to help the Netherlands move forward through that network. We believe in a society in which communication technology makes life richer, easier and fuller. KPN wants to be the unifier of that society, for people and companies. At home, at work and on the move. We have the resources, and the technology and the reliable fixed and mobile networks.



The National Police is investing in one strong approach towards cyber crime. From 2015 onwards, regional police shall investigate a growing number of such cases, applying expertise gained in recent years by themselves and by the National High Tech Crime Unit (in Dutch: Team High Tech Crime, or THTC). Since 2012, the focus of THTC lies on preventing and investigating high-impact cyber crimes with an innovative and complex nature, undermining society. In 2013, the unit won the Mensa Fonds Award for being the best employer for highly-gifted individuals. Working closely together with both regional and international law enforcement, THTC can be invoked in any high tech crime with a victim, a suspect or abused infrastructure residing in the Netherlands. We encourage you to report cyber crimes. Please make an appointment with your local police station and ask for the presence of a digital expert. If needed, they will invoke THTC. You are also welcome to follow THTC on Twitter: @PolitieTHTC



TNO, The Netherlands Organisation for Applied Scientific Research, is one of Europe's leading independent research and development organisations. TNO is not for profit and operates independently and objectively. Its unique position is attributable to its versatility and the capacity to integrate knowledge across specialist disciplines.

TNO innovates for a secure cyberspace and provides cyber security research, development, engineering and consultancy services to government and industry. Customers include Dutch government departments (e.g. Defence, Economic Affairs and Security & Justice) and private sector companies across Europe, including providers of national critical infrastructure (a.o. in telecoms, finance and energy).

TNO is an active member of numerous cyber security partnerships, including the European Network for Cyber Security (ENCS), the Hague Security Delta (HSD) and the EU NIS platform. TNO was part of the core team that formulated the Dutch National Cyber Security Strategy (NCSS) II and was one of the lead authors of the Dutch National Cyber Security Research Agenda (NCSRA) II.

www.tno.nl



The National Cyber Security Centre (NCSC-NL) is a government organization dedicated to increase the resilience of Dutch society in the digital domain and, by doing so, to help create a safe, open and stable information society. Through its network strategy, it aims to increasingly involve various other organizations in order to fully utilize their combined knowledge and experience. At this stage, the NCSC has established fruitful cooperation with other government bodies, public and private parties, the academic community, and with numerous international partners. One of its main products is the annual Cyber Security Assessment report, in which the NCSC - in cooperation with its partners - describes the current state of affairs of cybercrime and digital security in the Netherlands.

