SECURITY REPORT 2017

mnemonic
- securing your business

GLIMMERS OF HOPE

In the world of cybersecurity, 2016 was one for the history books.

We saw the largest and second largest data breaches in history (both Yahoo!), record-shattering DDoS attacks and alleged hacking that influenced the outcome of the US presidential election. Cybersecurity has never been so focused in the limelight as it is today. Closer to home, 2017 began with multiple politically-driven cyberattacks towards various facets of the Norwegian government.

It was not all doom and gloom though – 2016 also provided some reassuring glimmers of hope. The EU Parliament approved the EU General Data Privacy Regulation (GDPR) – perhaps this century's most significant advancement in data privacy legislation. There was also the approval of the NIS directive – a commitment to acknowledge the necessity to protect critical infrastructure across Europe from cyberattacks. And in an acknowledgement of the importance of information security, the United States also appointed its first ever Federal CISO.

At mnemonic, we continued to play our part in the collective fight against cybercrime. It is an uphill battle, but through perseverance we continue to gain ground. Highlights from 2016 include:

- Being appointed to Europol EC3's Advisory Board for Internet Security, joining the Global Cyber Alliance and establishing a threat sharing agreement with CERT-EU

- Commencing our collaborative research project - Semi-Automated Cyber Threat Intelligence (ACT) to develop an open platform to solve the challenges in collecting, analysing and sharing threat information, data and intelligence.

- Sponsoring the PhD research project Threat Ontologies for Cybersecurity Analytics (TOCSA), being carried out by our threat intelligence research and incident responder Siri Bromander

- Being recognised as the only European vendor in Gartner's Market Guide for Managed Detection and Response services

Despite the fact that we at times may have differing and even competing goals, we all benefit when joining forces against our common adversary. It is not until we combine our individual assets and expertise that we will become equipped to win this war.

Are you interested in collaborating and joining the collective fight against cybercrime? Send an email to **JoinTheTeam@mnemonic.no** and we'll continue to unify our efforts.

TØNNES INGEBRIGSTEN
CEO, mnemonic

TABLE OF CONTENTS

SECURITY PREDICTIONS: 2017 AND BEYOND

JON RØGEBERG
Head of Threat Intelligence
mnemonic

SECURITY PREDICTIONS: 2017 AND BEYOND

Another year has passed us, and you know what that means: it's time for us to dust off our crystal ball, summon our supernatural abilities and offer our security predictions for 2017. If 2016 was any indication, you'd better strap yourself in – it's going to be a bumpy ride.

Predictions are no easy task. If your prediction is wrong, you can be certain that someone will point it out and you will likely never be allowed to forget it. Meanwhile if your prediction is correct, it will be discredited after-the-fact as something that was obviously going to happen, or simply attributed to a lucky guess. Hindsight is 20/20, and predictions are a lose-lose situation for the predictor.

The reality is, despite how much time we spend analysing past events, the amount of technology we build to crunch the numbers, or the frequency we rub our crystal ball, the simple fact remains that we cannot know for certain what the future has in store for us. However this doesn't mean we can't prepare ourselves.

We do not share these predictions to become the Nostradamus of cybersecurity. These predictions are not the almanac for exactly what will happen this year, or the next, or the one after that. Our intention is to share insights and reflections that our readers can use to gain situational awareness of the state of cybersecurity today, and use these insights to better prepare for what may be ahead of us.

We hope you find the insights useful, and if we're wrong, you allow us to (eventually) forget it.

With that, let's get started.

IoT technology is designed with convenience in mind, often at the expense of security.

CONVENIENCE TRUMPS SECURITY WITH THE INTERNET OF THINGS

For years, the security community has hypothesized, discussed and conducted proof of concepts on the devastation unsecured IoT devices could pose to our digital and physical worlds. In the second half of 2016, the Mirai botnet very publicly demonstrated the implications millions of unsecured devices can have.

By infecting internet-connected devices such as security cameras, printers and home routers, Mirai builds an army of unsecure IoT devices to launch record-shattering DDoS attacks. This botnet was responsible for the 620 Gbps attack that took down Krebs On Security in September. Shortly afterwards, the source-code was released and the botnet was used to launch a reported 1.2 Tbps attack against domain management company Dyn, resulting in popular services such as Netflix, Amazon, Spotify, Twitter, amongst many others being temporarily unavailable across the US and Europe. Mirai was also responsible for taking an entire country offline when Liberia's fibre infrastructure was repeatedly targeted with over the course of a week.

This is merely scratching the surface on the threat billions of internet-connected devices pose if they're not secured properly, and we can only expect the problem to get worse before it gets better. IoT devices will continue to be used in DDoS attacks of increasing volumes, and will be exploited for other types of attacks and objectives as well. Cameras and video solutions in conference rooms can be used for espionage and extortion. WiFi-enabled coffee machines, wall integrated meeting room panels and light switches may be used as the jump-point to infiltrate a corporate network. Data leaks of private data from health related IoT devices. The list is goes on.

IoT technology is designed with convenience in mind - a way to bring us further into the 21st century. Unfortunately this convenience and leap into the future is often at the expense of security, which comes as an afterthought, if it's even thought about at all.

CEO FRAUD – THE DIGITAL BANK HEIST

Organised crime comes in many forms. In Norway, the largest bank heist in history occurred in 2004 when a team of bandits armed with assault rifles robbed a cash depot. Popularly known as the NOKAS heist, the assailants delayed police response by attacking the local police station with tear gas, laying road spikes near the station, and setting a car on fire to block police vehicles from leaving. The robbery resulted in a shootout where one police officer was killed. The bandits made off with €6.3 million, of which €5.6 million is still missing. The heist and subsequent trials (resulting in 13 prosecutions) were popularised in the media and even made into a movie. Sounding like the plot of a Hollywood blockbuster, you would be hard-pressed to find a Norwegian who was not familiar with the burglary.

Meanwhile in January 2016, a fraudster impersonating a top manager in an international company won the trust of employees in a Norwegian subsidiary and instructed them to transfer over €55 million to foreign accounts. The fraud was detected soon afterwards, which allowed much of the money to be recuperated, however more than €11 million is still missing.

Lacking guns, masks and escape vehicles, yet making away with twice the amount as the NOKAS heist, the electronic heist received minimal news coverage and is all but forgotten. Such is the case with much electronic fraud. Commonly known as CEO fraud or Business Email Compromise (BEC), the scam is fairly straightforward. Impersonate a top ranking official and use social engineering techniques to convince employees to make large, fraudulent transactions. No malware, no hacking, just clever deception.

How big of a problem is CEO fraud? Nobody really knows. Between October 2013 to February 2016 the FBI has recorded over 17,000 victims of so-called CEO fraud with loses exceeding $2.3 billion USD. This however merely represents the cases that were reported.

With enormous gains, minimal risk and requiring limited technical knowledge, CEO fraud is a rapidly growing form of organised cybercrime. Expect to see more cases of CEO fraud to be publicised in the media, and even more go unreported.

INVESTIGATIONS BECOME CLOUDY

It's estimated that more than 90% of companies are using cloud technology in some form, and this adoption continues to grow. Trust in cloud platforms continues to rise, and organisations are becoming more comfortable with the amount of sensitive data they store with these providers.

A cornerstone of cloud technology is leveraging shared platforms. When customers share an infrastructure, resource usage becomes more efficient, the service can be scaled more easily, and costs can be minimised. Unfortunately it is this same benefit that introduces serious security challenges for many cloud providers.

In our experience, most cloud environments simply aren't designed to facilitate security auditing, are built with an infrastructure that complicates the ability to detect and respond to security incidents, and are lacking the tools to support adequate investigations. The result is that when a security incident is detected, the lack of auditing at an organisational level or support for investigative tools means the scope of the incident is rarely understood, and it is near impossible to isolate the data that has been compromised. The consequence is individual organisations have little chance of knowing if *their* data was compromised should their provider be suffer a data leakage incident.

Now let's say that the data that was potentially compromised is customer data. Your organisation is now responsible to inform your customers that their data *may* have been leaked. You can see how the domino-effect begins here.

This is not to say organisations shouldn't use cloud technology or that all cloud providers lack essential security. The underlying point is that due diligence is critical when it comes to choosing cloud providers. It is important to assess the security practices and information security policy of the provider, and understand their ability to detect, respond and investigate security incidents that involve *your* data. And in cases where you as a customer may be responsible for carrying out the investigation yourself, it is essential that the cloud platform you are using supports extracting the necessary data.

Expect to see a steady increase in the number of high-profile breaches of cloud providers and an inability to determine what data was compromised, how it happened, and how to prevent it from happening again.

The sinister imagination and creativity of financially motivated cybercriminals should not be underestimated.

RANSOMWARE EVOLVES DIGITAL EXTORTION

Ransomware is broadly defined as malicious software that adversely affects a victim's computer and demands a ransom payment. For modern variants, this has generally been in a crypto-form where the victim's files are encrypted and a ransom payment required to obtain the decryption key.

The concept of ransomware is however not new. It was first observed back in 1989 and spread using floppy disks, but the concept has not changed – pay up or else. Ransomware as we know it today is believed to have been first observed in 2005. Primarily found in Russia and surrounding countries, this ransomware would move specific file types, like all JPEGs for example, to an encrypted and password protected ZIP archive. Pay a ransom and you get the password to retrieve your files.

From here, the scams evolved into two primary variants. The first would use various methods to prohibit users from accessing their machines without intentionally destroying data. The second would utilise a scare-tactic that displayed a message, often from a law enforcement agency local to the victim, that their computer was used to conduct illegal activities, such as downloading pirated media or accessing child pornography.

The tactics ransomware uses to coerce victims to pay ransoms continue to evolve. Towards the end of 2015, Chimera was the first reported ransomware variant that not only encrypted a victim's files, but also threatened to publish the stolen files on the Internet if the ransom is not paid – a practice known as doxing (the act of leaking private data or information about an individual to the public). While Chimera only threatened to dox a victim's data without actually having the capability to do it, it set a precedent for new ransomware variants to follow.

Jigsaw, named after the iconic character from the Saw movie franchise, is similar to other crypto-ransomware variants, except it begins to systematically delete files based on a timed countdown. Jigsaw also reported added live support to help victims make Bitcoin transactions to pay their ransoms. Later variants of Jigsaw also threatened to leak personal data to stored contacts.

Popcorn Time (not affiliated with the video streaming software) was observed towards the end of 2016. This ransomware operates similar to existing crypto-variants – install, encrypt files, demand ransom - except it has a small twist. Victims are offered a second option to decrypt their files – infect two friends with the ransomware and if they pay, your decryption key is free. While when released Popcorn Time was riddled with workarounds and technical issues, conceptually it represents a new tactic for fraudsters to infect more victims and puts those victims into an ethical dilemma.

The sinister imagination and creativity of financially motivated cybercriminals should not be underestimated. Expect this to be demonstrated throughout 2017 and onwards as new extortion techniques are developed.

MACHINE LEARNING AS AN ATTACKER'S TOOL

We often regard technological advancements from the perspective of how they will improve our daily lives, benefit society as a whole and how they will advance humankind. One such advancement receiving significant attention in recent years is machine learning.

Applications for machine learning are already all around us. Netflix uses it to analyse your viewing habits and recommend new content, PayPal and other financial institutions use it to detect fraudulent activity, and Siri uses it for natural language processing that allows us to speak naturally to our favourite digital assistant instead of only understanding a predetermined list of commands. We're beginning to see it used in healthcare to scan images for early disease detection, in driverless cars, and of course in security as well.

In IT security, machine learning helps us to better protect our networks, data, and businesses. We use it to detect malware more accurately, to detect anomalies in user behaviour, and to analyse a seemingly insurmountable volume of data. In our Argus platform, we use machine learning as a mechanism to validate the decisions our security analysts make when assessing potential security incidents (amongst other applications). The application of the technology is seemingly endless, which goes for attackers as well.

While there are few observations of the technology currently being used in offensive attacks, it is only a matter of time. Our hypothesis is some attackers, particularly well-funded nation states, are already using machine learning and natural language processing for the identification, analysis and classification of the data they are acquiring through their espionage efforts.

Expect machine learning to lead to more targeted, believable and automated social engineering attempts. Before sending a phishing email, machine learning can be used to automatically scour news sites and social media to craft a personalised message that is more believable and likely to fool the target by highlighting current affairs.

The machine learning that security vendors are relying upon also has the potential to be manipulated by attackers. For technology that relies on detecting anomalies, attackers can systematically (even using machine learning themselves) use junk data to poison and alter what the machine learning model understands as normal behaviour to the point where their malicious activity is no longer perceived as an anomaly.

While we don't expect attackers to be at the forefront of machine learning research, we can guarantee that they will find imaginative ways to leverage the technology as it continues to become commoditised.

Norway

KRAFTCERT
NORWEGIAN ENERGY SECTOR CERT

Margrete Raaum

General Manager

KraftCERT assists the energy sector in handling digital security incidents and participates in the national emergency response organization. KraftCERT aims to provide reliable, secure and efficient information sharing between companies and organizations, both nationally and internationally.

WHAT IS YOUR BIGGEST SECURITY CONCERN?

One of my biggest concerns is related to attribution. Companies and countries are more eager to attribute an incident to a certain adversary than to look at their own vulnerabilities. I see the importance of threat intelligence, but when your baseline security is so low that old and easily-exploitable vulnerabilities still are attack vectors, you should worry about other things than attribution. The fact that some nation states say that they have, or are building, cyber capabilities to attack national infrastructure should not go unnoticed, but should make the companies take a good look at themselves and be honest about the reality of their security levels.

IN WHAT AREAS OF SECURITY DO YOU THINK WE'RE FALLING BEHIND?

We're still talking too much about malware detection and prevention, and too little about baselining, both traffic and systems. Expected traffic patterns will be highly predictable in many critical networks. Effective baselining first requires an updated inventory, something that many companies may be missing - you don't know what to expect if you don't know what you have. And remember when baselining/white listing systems: the adversary is on to us, utilizing names and directories that are normally white listed. What looks like a duck does not necessarily quack like a duck.

WHAT GIVES YOU HOPE FOR THE FUTURE OF SECURITY?

I hear less and less that "the user is our weakest link". Unless one plans to replace humans in the workplace, it is time to look at how the user can be better protected. It is not possible to get everybody up to the high level of skepticism that some seem to think we can aim for. There are plenty of tools and tricks to make the workday less difficult for the average human, while still improving security.

THE CLASH OF INTERNET MORALITY: CRYPTOGRAPHY VS. SURVEILLANCE

TOR E. BJØRSTAD, PH.D.
Application Security Lead
mnemonic

In the modern days' Internet, an individual's privacy rights inversely conflict with the authorities' need to protect society from cybercrime and defend national interests. While cryptography aims to protect the confidentiality of our information, surveillance sets out to protect our society. In a longstanding clash that challenges the essence upon which the Internet was born, has this duality reached a breaking point?

Cryptography (from Greek: *kryptós,* hidden, secret; *graphein,* writing) is the theory and practice of techniques to protect information in an adversarial setting. Since antiquity, the study of encryption, which are techniques that transform messages into incomprehensible gibberish, have been used to prevent military secrets from falling into enemy hands. Only somebody possessing the correct encryption key, should be able to unlock an encrypted message and access the information thus protected.

Throughout history, there has been an ongoing cat-and-mouse game between those cryptographers attempting to devise better encryption, and those attempting to find weaknesses and break enemy encryption [1].

MODERN CRYPTO HISTORY

Widespread civilian use of cryptography began concurrently with the Internet era. First steps were taken in the mid 1970s, with three ground-breaking publications: the invention of so-called public-key encryption by Diffie and Hellman in 1976, the establishment of the Data Encryption Standard in the United States in 1977, and the invention of the Rivest-Shamir-Adleman (RSA) encryption scheme later that year. Public research into cryptography bloomed after this, and for the last 40 years, a large amount of cryptographic theory and practice has been established in the open.

Between the 1970s and today, there have been major debates about the legal status of encryption technologies. During the Cold War, encryption was initially classified as munitions by the United States, and underlaid strict export controls. At the same time, commercial needs for encryption were gradually increasing, particularly in the financial sector. The publication and distribution of Pretty Good Privacy (PGP) on the Internet in 1991 was another watershed, as it gave individuals world-wide access to military-strength encryption tools.

The policy debate of whether strong encryption technologies should be available to the general public ran hot during the 1990s, and is often referred to as the "crypto wars". However, the cat was largely out of the bag, with the technological know-how being both widely disseminated and geographically dispersed. While there are still certain export regulations in place, particularly for military-grade encryption hardware, the theoretical knowledge of how to build highly secure encryption systems is widely available today, together with well-established technical standards and multiple open-source implementations.

[1] *A particularly famous example of this were the British efforts, code-named 'Ultra', to break the German Enigma encryption as well as other ciphers, during World War 2. At the end of the war, more than 10,000 people were working on Ultra, amongst them Alan Turing. The breaking of Enigma was stated by Eisenhower in 1945 to have been a 'decisive contribution to the Allied war effort', yet it was kept secret until 1974.*

ENTER SNOWDEN

After the Snowden leaks in 2013, the "crypto wars" were somehow reignited. There were multiple revelations regarding attacks on cryptographic technology and use of encryption. Some of the more notorious revelations concerns the systematic attempts by the NSA to weaken public standards, the wiretapping of corporate networks of companies such as Google and Yahoo, and the discovery that US authorities could compel smartphone manufacturers to help them decrypt smartphone storage.

The backlash from the technology industry has been quite strong. Use of mandatory encryption both on the Internet and for internal communications has grown strongly between 2013 and 2016. Around 2014, Apple and Google both redesigned their encryption modules for iPhone and Android to "tie their own hands" – in effect, taking a public stand in favour of consumer privacy, by ensuring that they could no longer cooperate with such requests. Consumer awareness regarding online privacy also appears to have increased.

This all came to a head with the dispute between FBI and Apple in the San Bernardino shooting case, where the FBI wanted Apple to develop custom software to help them unlock the shooter's iPhone 5C, which Apple refused. As the case was withdrawn a day before it would have gone to court, no precedent was established.

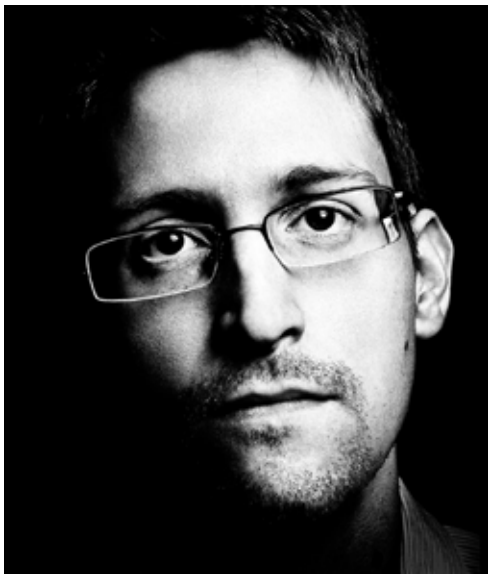

Credit: wired.com

ENCRYPTION VS. LAW ENFORCEMENT

In stable and peaceful Western democracies, there are times where surveillance and communications control are considered proportionate and legitimate law enforcement tactic, as long as it is subject to the appropriate checks and balances. The advent of strong encryption apps in the hands of individuals, in combination with the rise of the smartphone as the main communication platform, disrupts this, and there is a growing concern in the law enforcement community that criminal net-

A BRIEF HISTORY OF ENCRYPTION

Ca 1500 BC

Craftsmen in Mesopotamia use rudimentary encryption techniques to protect recipes, presumably of commercial value.

7th – 3rd century BC

The Spartan military employ a simple encryption device, the *scytale,* for important messages.

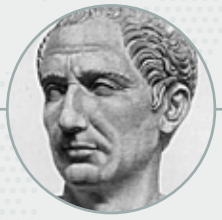

1st century BC

Julius Caesar uses the eponymous "*Caesar cipher*", which shifts the alphabet by 3 positions, for important correspondence.

works are "going dark". Similarly, intelligence agencies would for obvious reasons strongly prefer that everybody else used encryption mechanisms that are breakable (only) by them.

Until Snowden, some stated that society was in a period of "peak surveillance". On one hand, the use of modern communication tools had become pervasive, on the other hand, cryptography was often not used, or used incorrectly, and awareness of issues relating to operational security was low. Thus digital surveillance became an extremely effective tool to track and surveil. After Snowden, this state of affairs is rapidly changing.

> **Activists, journalists, whistleblowers, law enforcement themselves, and many others, have legitimate and strong needs to protect the privacy of their communications.**

At present, encryption is an essential tool in the toolbox used by both government, the private sector, and individuals, to protect personally identifiable information, commercially sensitive information, credit card data, financial transactions, login credentials and passwords, and everything in between. In repressive regimes, having access to secure (i.e. encrypted) communications is a key enabler for opposition groups to organize and communicate. Activists, journalists, whistleblowers, law enforcement themselves, and many others, have legitimate and strong needs to protect the privacy of their communications.

On the flip side, criminals are no less technology-savvy than the general public, and are utilizing the exact same encryption technologies to protect their communications infrastructure and offensive operations, as is used for cyber defense. We are also witnessing the rise of malicious cryptography, sometimes known as cryptovirology. In 1996, researchers Moti Yung and Adam Young predicted the use of public-key cryptography to create ransomware. Around 2005, the first extortionate ransomware was seen in the wild, and for the last 3-4 years ransomware such as Cryptolocker has become a favourite tactic of cybercriminals.

There is also the intelligence dimension. In Norway, a recent Norwegian Official Report (NOU) was made on the topic of creating a digital border defense ("Digitalt Grenseforsvar"). The report advocates monitoring all data traffic crossing the Norwegian border, by introducing legislation modeled on the existing Swedish FRA law. However, it also recognises "that the increasing use of strong encryption will continue, and that this will affect the value of surveillance activity". To counteract this trend, the report suggests compelling technology companies to allow traffic decryption on the data link layer. The proposals have been condemned by privacy advocates in Norway, and are currently under review.

Ca 850 AD

The Arab philosopher al-Kindi writes a book describing how to break substitution ciphers by frequency analysis; probably the earliest use of statistical methods in cryptanalysis.

1553

The Vigènere cipher was invented by Giovan Battista Bellaso (though attributed to and named for Blaise de Vigènere); this system remained unbroken for 300 years.

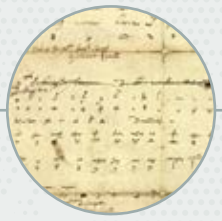

1586

Queen Elizabeth I's codebreaker deciphered correspondence between Mary, Queen of Scots and Anthony Babington, plotting to kill Elizabeth. This is known as 'the Babington Plot'. Both were executed.

Existing Swedish surveillance practices come with their own controversy. Former Swedish foreign minister Carl Bildt has publicly defended the law, stating that "there is a difference between good states and somewhat less good states". The Swedes are known for their advanced signals intelligence capacities; with a geographic location that is strategically situated to monitor a significant percentage of all fiber optic communications from Russia, but also from Norway.

THE CALL FOR COMPROMISED CRYPTO

Because of this, and commonly with reference to the fight against worthy foes such as child pornographers, organized crime, and terrorists, there are periodic calls to regulate civilian use of encryption to enable government access, typically through technological "back doors". Similarly, there have been calls to weaken or ban cryptographic anonymity tools such as TOR (The Onion Router).

Such proposals may seem benign and in the common interest, but there are several important reasons why the approach is in fact counterproductive.

TECHNOLOGICAL PERSPECTIVE

From a technological perspective, it is not technologically feasible to build back-doors in encryption systems that only provide access to the "right" people. One will introduce a significant risk that the back-door may be abused or fall into the wrong hands, which would lead to significant vulnerabilities for everybody using the system;

ETHICAL PERSPECTIVE

From an ethical perspective, mandated encryption back doors and government access is just as desirable for use by regimes that would use it to systematically violate human rights; because of this, one should not erect a framework to enable such activities;

PRACTICAL PERSPECTIVE

From a practical perspective, the deployment of "nationally surveillable" encryption technologies breaks interoperability and standardization on the Internet, since each nation-state will want to have control over their own networks;

PRAGMATIC PERSPECTIVE

From a pragmatic perspective, the know-how and relevant technology for encryption is already out there, and easily adaptable for use by putative lawbreakers; such use would most likely not be readily detectable outside a complete surveillance state.

1791-1792

Marie Antoinette sent multiple encrypted letters to Count Axel von Fersen while imprisoned during the French Revolution. (See https://eprint.iacr.org/2009/166 for a modern analysis of these letters.)

1863

Prussian military officer Friedrich Kasiski publishes the first general method for breaking the Vigènere cipher, known as Kasiski examination. A similar method may have been known to Charles Babbage as early as 1846.

1883

Dutch linguist Auguste Kerckhoffs publishes two articles outlining principles for practical military encryption design. His second rule remains fundamental: "A cryptosystem should be secure even if everything about the system, except the key, is public knowledge."

It is difficult to find "middle ground" here –cryptographic technology will either be practically unbreakable, or possible for any well-funded and capable adversary to defeat. If strong encryption becomes strictly regulated (again), the negative consequences will thus be significant and immediate.

Looking back to the export control era of the 1990s, we also see that the weakened "export-grade" encryption mandated then, is causing very real security problems and damage today. Examples of this include the FREAK, DROWN and LOGJAM attacks on Transport Layer Security.

It is thus difficult to see how restricting the use of encryption technologies will yield a significant benefit for society, beyond perhaps (in an initial phase) catching the "stupid" criminals who do not adapt to changing circumstances. On the other hand, the negative externalities that would be imposed by such regulation are significant and clear.

As the joke goes:

When encryption is outlawed, bayl bhgynjf jvyy unir rapelcgvba [2]

[2] "... only outlaws will have encryption"

An ongoing challenge for the technology community is to communicate these perspectives clearly to policymakers and law enforcement, and help look for constructive alternate approaches.

1917

The "Zimmerman telegram", proposing a military alliance between Germany and Mexico against the USA, was intercepted and decrypted by British Intelligence. This was before the US entered World War 1, and revelation of the contents enraged the public.

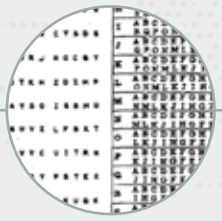

1917

American Gilbert Vernam proposes the "one-time pad", a theoretically unbreakable (but extremely impractical, and therefore rarely used) encryption method.

1938 – 1945

Polish efforts to break Germany's Enigma encryption were continued by Britain at Bletchley Park. Operation ULTRA was a major Allied intelligence breakthrough. It involved as many as 10 000 personnel, but remained a secret until 1974.

SOME ALTERNATE APPROACHES

Unfortunately, several of the alternative approaches that are being discussed are also quite controversial, as they too prove problematic from a privacy perspective. Yet at the very least, they may allow some middle ground to be found.

To some extent, the problem of criminals' communications "going dark" can probably be mitigated by changing how law enforcement agencies are organized and staffed. Although there certainly exist units who possess advanced technological capabilities and skills, it seems evident that much of law enforcement is structured and staffed to fight traditional crime, and poorly adapted to fight crimes conducted or facilitated via the cyber domain. Consequently there should be a significant potential in developing new investigative techniques and using new technology more effectively, which could compensate for the loss of some existing surveillance capabilities.

> "
>
> Access to metadata can be just as useful as access to the encrypted data itself.

As signals intelligence agencies (and managed security service providers) know well, access to metadata, such as information about who are communicating and how often, can be just as useful as access to the encrypted data itself. Gathering such metadata will usually not require breaking any content encryption, though it may be made difficult by anonymization technologies such as TOR. Large-scale metadata gathering, as mandated by the now-defunct EU Data Retention Directive, was found to be in violation of fundamental citizen's rights. However, targeted metadata gathering and analysis should, at least to some extent, be a viable tool.

An alternate hypothesis which has been discussed, is that the spread of the Internet of Things (IoT), combined with widely-deployed insecure devices, will provide so much accessible information that law enforcement will not need to break encryption. Yet the practical consequences of widely-deployed insecure devices also seems like a large negative, and one that the technology community should try to prevent for the good of the Internet.

In some jurisdictions, including the UK and Australia, law enforcement may compel individuals to surrender their cryptographic keys or provide access to materials in decrypted form. However, this can easily clash with the principle of protection against self-incrimination. Many encryption products also implement technologies to provide plausible deniability – for example by enabling a cryptographic storage volume to decrypt in two different ways, with two different keys.

1949

Claude Shannon creates the theoretical foundations of modern cryptography with the paper "Communication Theory of Secrecy Systems", based on classified work he did during World War 2.

1976 - 1977

In a groundbreaking discovery, Whitfield Diffie and Martin Hellman invent public-key cryptography and the Diffie-Hellman key exchange in 1976. At MIT, Rivest, Shamir and Adleman publish their eponymous RSA encryption algorithm the subsequent year.

1991

Phil Zimmerman publishes the software Pretty Good Privacy (PGP) on the Internet.

Another highly controversial topic has been whether law enforcement should be able to actively perform hacking activities and/or use "government malware" to gain control of a suspect's devices and computers. The advantage of gaining device access is indeed that it may make it possible to monitor activity at the point where encryption is added or removed from communications. Indeed, for (defensive) security monitoring, an ongoing trend is the gradual move from purely network-based monitoring to more endpoint-centric solutions, in part due to the increasing difficulty of monitoring encrypted network traffic.

Compromising endpoints also circumvents many anonymization technologies. In 2013, the FBI used a TOR browser exploit to take down a child pornography ring. On the other hand, such government-supported hacking is not a reliable way of obtaining access, as vulnerabilities are patched and suspects grow wise. Hoarding vulnerabilities, weakening ongoing software-security initiatives, and introducing deliberate back-doors in software are all tactics that open up large ethical cans of worms, and neither provide stable long-term solutions.

CONCLUSION

At a meeting in May 2016 at Europol's headquarters in The Hague, on privacy in the age of surveillance, a joint statement was given by Europol and ENISA. The statement stresses that the use of intrusive investigative tools must be proportionate to the crimes being committed, and notes that the introduction of mandatory back-doors or key escrow would be likely to have limited practical effect.

> **The use of intrusive investigative tools must be proportionate to the crimes being committed.**

This appears to be a significant step forward from the suggestions of banning strong cryptography, as for example advocated by David Cameron in 2015. Yet the UK (and US) have strong influences, and post-Brexit it is unclear what the future relationship between the UK and Europol will be. Moving into 2017, it will be interesting to see what policies will be advocated by the new governments in the UK and US.

There are no easy solutions, and even if the "crypto wars" will be put to rest once again, the debate on the balance between individual privacy rights and law enforcement capabilities is likely to continue for the foreseeable future.

1999

Transport Layer Security (TLS) is standardized in 1999, building on work carried out at Netscape to develop a secure communications protocol for use with e-commerce (SSL) since 1994.

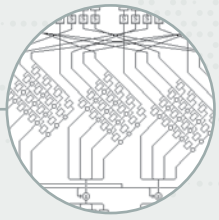

2001

Advanced Encryption Standard (AES) is published, after a five-year standardization process. The standard is based on the algorithm 'Rijndael' by Belgian cryptographers Vincent Rijmen and Joan Daemen.

The canonical reference for the history of Cryptography is The Codebreakers – The Story of Secret Writing (ISBN 0-684-83130-9) by David Kahn. It can be used as a reference and further reading.

2016: A VIEW FROM MNEMONIC'S SECURITY OPERATIONS CENTER

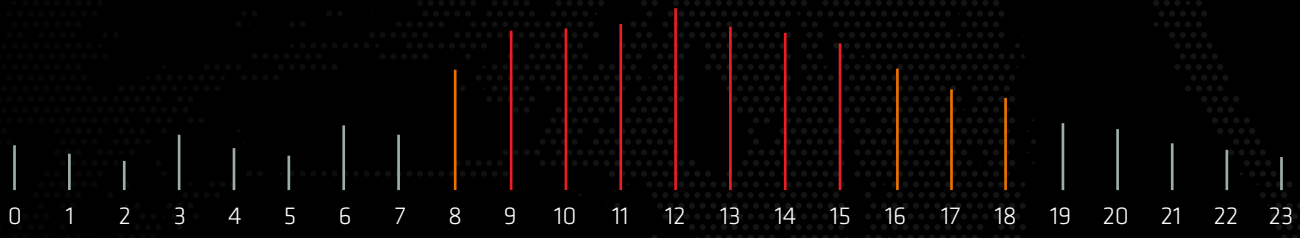

WHEN ARE SECURITY INCIDENTS OCCURRING?

Unsurprisingly, security incidents are occurring at all hours of the day. There are gradual jumps during office hours, and the number of incidents peaks between 12:00 – 13:00 CET. This is likely related to users browsing the internet and performing more personal related activities during their lunch period

30% of security incidents occurred outside the daytime office hours of 07:00 – 18:00 CET.

There is a 57% increase in users being exposed to malware during the lunch hour of 11:00 – 13:00 CET when compared to the 09:00 – 11:00 timeframe. Interestingly, users are 30% more likely to be successfully infected by malware between 9 - 11. This increase could be related to users becoming infected while off the security controlled zones of the corporate network – at home for example if all communication is not forced through a VPN – and the infection being detected once the laptop is re-connected to the network.

ATTACKS AGAINST USERS VS ATTACKS AGAINST INFRASTRUCTURE

Attacks against users commonly come in the form of phishing and the distribution of malware attacks. The target is the user (and their computer) and typically require the user to perform some action to facilitate the infection, such as clicking a link, opening an attachment or executing a downloaded file.

Attacks against infrastructure are a different story. Reconnaissance, brute force attempts, cross-site scripting, SQL injections, vulnerability exploits, and so on do not require user interaction and are targeted against an organisation's infrastructure.

14%

of reconnaissance activity against company infrastructure is occurring on weekends. This is higher than what is observed on Fridays.

33%

of attacks against infrastructure occurred outside the office hours of 07:00 – 18:00 CET.

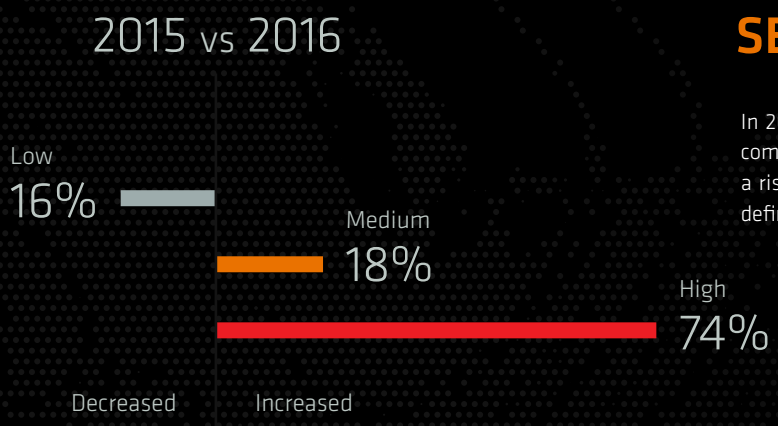

SEVERITY OF INCIDENTS

In 2016 we saw a 74% rise in high severity incidents when compared to 2015. This is largely due to the combination of a rise in ransomware infections, and high-value assets (as defined by our customers) being involved in security incidents.

WHEN ARE USERS BEING INFECTED?

Users are statistically most likely to be involved in a malware related incident on a Tuesday or Wednesday. Exposure to malicious code was most frequent on Tuesdays, while actual infections were more prominent on Wednesdays, these numbers are too close to have statistical relevance.

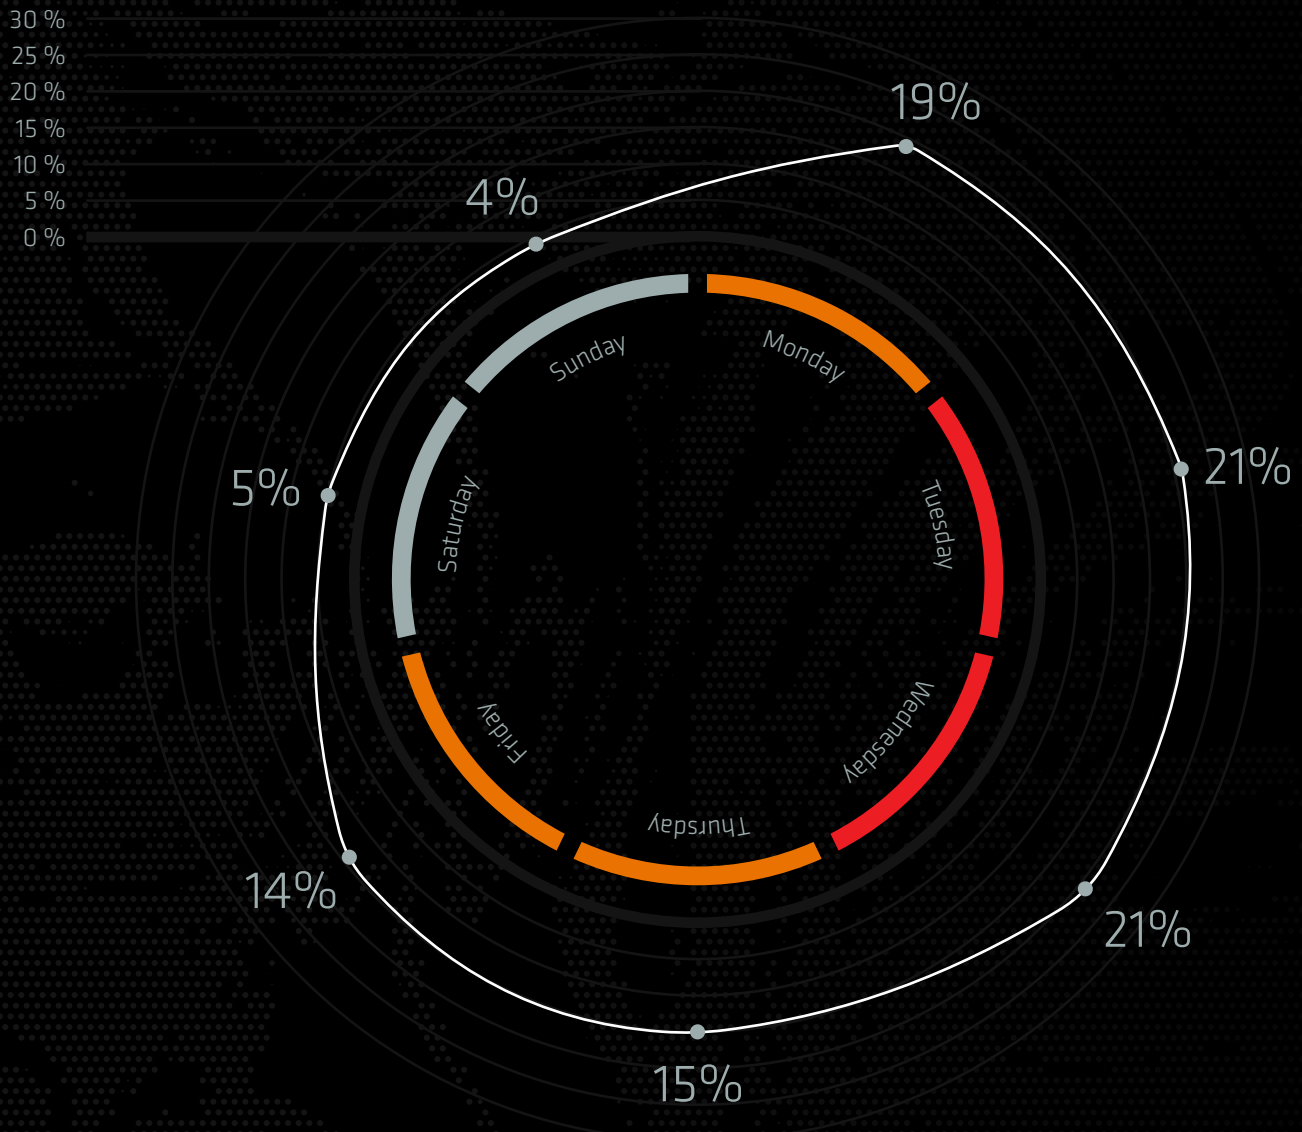

ANGEL ALONSO
Team Leader, Governance, Risk & Compliance (GRC)
mnemonic

THE CISO'S DILEMMA: RESPONSIBILITY CANNOT BE OUTSOURCED

CISOs are entrusted with an essential responsibility: protect an organisation's information while enabling the business. With new risks and threats that arise daily, rapidly evolving technology and emerging compliance and legislation demands, keeping an organisation's information safe is a daunting challenge. Welcome to the CISO's dilemma.

Allow me to introduce you to Mr. Ciso, the Chief Information Security Officer (CISO) at General Global Corp (GGC), a global enterprise. Mr. Ciso is what research shows is your "typical" CISO: male, in his forties, originally educated within business, but with a professional career focused in information technologies and IT security *(The Anatomy of a CISO, Digital Guardian)*. He performs a staff function within the IT department with a solid and effective reporting line to a business-oriented IT Director.

Mr. Ciso's life has become quite complicated in recent times. A distant memory are those comfortable years when his only headache was to occasionally combat worms spreading in his well-defined network and ensure PCI compliance.

TIMES HAVE CHANGED, AND KEEP CHANGING

Now he is not only facing a rapidly evolving threat landscape and new regulations, but also helping (or at least not hindering) the achievement of GGC's core business objectives. His position has become more strategic and requires a new set of skills.

Acting as GGC's second line of defense, he needs to not only ensure that business assets are secure and protected against disruptions, but also manage regulatory obligations and inherited risks. This is all while simultaneously acting as the facilitator between line managers, internal and external auditors, and senior management.

Nevertheless, after years of advocating, he has finally been requested to start reporting to GGC's Board of Directors on cybersecurity and technology risk.

Mr. Ciso is happy with this, since having more reporting lines will help to spread his message. It has been a long, uphill battle for Mr. Ciso to convince management to consider his security strategy as more than a cost center, and to secure the adequate funding required to bring GGC's security program to the necessary maturity level.

By 2020, 100% of large enterprises will be asked to report to their Board of Directors on cybersecurity and technology risk at least annually, which is an increase from today's 40%.

The Comprehensive Guide to Presenting Risk and Information Security to Your Board of Directors, Gartner

SKILLS OF THE CONTEMPORARY CISO

- Leadership
- Relationship Management
- Risk Management
- Business Knowledge
- Communication
- Strategic Planning
- Psychology and Sociology

Develop the Skills of the Contemporary CISO, Gartner

WHAT IS CEO FRAUD?

CEO fraud, also known as Business Email Compromise (BEC), is a financially motivated attack. Fraudsters use social engineering techniques and pose as senior managers to coerce employees to transfer company funds to fraudulent accounts under the premise that it is a legitimate business transaction.

Attackers will typically either gain access to a senior manager's email account through phishing, spoofing a senior manager's email address, or by emailing employees from a domain they register that is very similar to the target's domain name.

According to the FBI, victims reported losses of $2.3 Billion USD due to CEO fraud between October 2013 and February 2016. This does not include the substantial number of un-reported cases.

WHICH SECTORS DOES THE NIS DIRECTIVE APPLY TO?

- Energy (Electricity, Oil, Gas)
- Transport (air, rail, water, road)
- Banking
- Health
- Drinking water supply and distribution
- Digital infrastructure (internet exchange points, domain name service providers, top level domain name registries)

In addition, digital services seen to have general importance with regards to cyber-security, including online marketplaces, online search engines and cloud computing services.

(Un)fortunately for Mr. Ciso, GGC was the victim of a CEO fraud incident in 2016, where the company was swindled out of nearly $100,000 USD. While such an incident is obviously something they want to avoid, it was a blessing in disguise for Mr. Ciso's security program.

CYBERSECURITY IS A BOARDROOM DISCUSSION

It can be challenging to convey to management an organisation's true risks and exposures based on hypothetical consequences and probabilities, especially when dealing with the complex and often misunderstood topic of cybersecurity. This message becomes loud and clear though when there is cold hard cash missing.

But it is not only the CEO fraud incident that has made cybersecurity a topic at the management table. New directives and regulations are on the rise and making cybersecurity a discussion in boardrooms globally.

In the USA, President Barack Obama emphasized the necessity to protect national critical infrastructure from cyberattacks by signing the Executive Order 13636 for Improving Critical Infrastructure Cybersecurity in 2013. As a result, the following year the NIST Cybersecurity Framework (CSF) was published to provide organisations with a set of industry standards and best practices to manage cybersecurity risks. By the end of 2015, more than 30% of US companies have adopted the framework, with this number expected to grow to 50% by 2020. Further emphasizing the importance of cybersecurity, in September 2016, the US appointed their first Federal CISO, responsible for driving cybersecurity policy, planning and implementation across the US Federal Government.

Meanwhile the European Parliament passed two significant pieces of legislation in 2016 to address cybersecurity challenges.

General Data Protection Regulation (GDPR) is a regulation intended to strengthen and unify data protection for individuals within the European Union (EU) and European Economic Area (EEA), while at the same time addressing the export of personal data outside the EU/EEA. The regulation was adopted on 27 April 2016 and, after a two-year transition period, will enter into effect on 25 May 2018. A major discussion point in many boardrooms is the potential penalties for non-compliance with the GDPR, which sees a maximum penalty of up to €20 Million or 4% of annual worldwide turnover, whichever is greater.

The Directive on security of network and information systems (the NIS Directive) is intended to ensure a high, common level of network and information security across the European Union, raise Europe's preparedness to ward off cyber incidents and encourage cooperation among all Member States. The NIS Directive applies to organisations that provide elements of a country's critical national infrastructure and main digital service providers. The directive was adopted by the European Parliament on 6 July 2016, and entered into force in August 2016. EEA member states will have 21 months to transpose the Directive into their national laws and a further 6 months to identify operators of essential services.

BUILDING A MODEL TO SIMPLIFY THE CHALLENGE

Taking into consideration the complexity of the modern cybersecurity landscape, Mr. Ciso set out to build a model to bring order and structure to his information security issues and, with any degree of success, eventually be used as the basis for GGC's future security strategy. This is the model he developed:

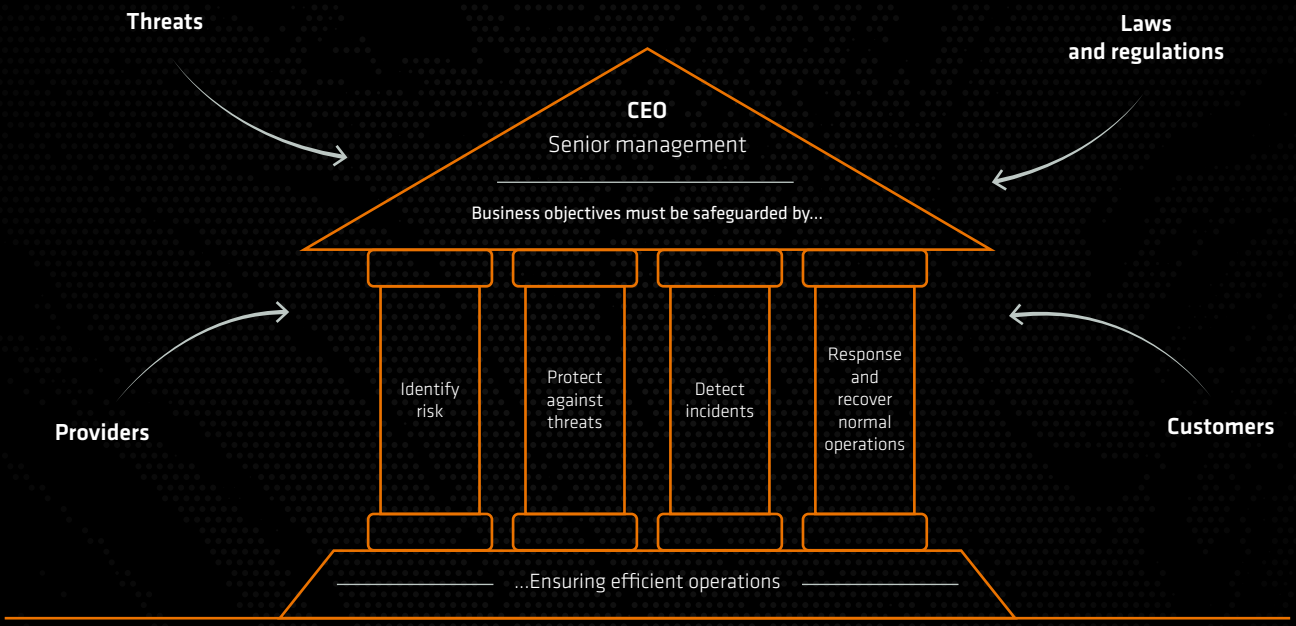

GGC's business objectives must be safeguarded by identifying the risks that are actually applicable for GGC. This includes identifying potential threat actors whom may target GGC, inherited risk introduced by third parties (including partners and service providers), along with the risk produced by non-compliance with laws and regulations in the various countries GGC operates in.

Once these risks have been identified, an informed decision can be made to accept the risk, or institute protection mechanisms to mitigate them within an acceptable cost. For threats that cannot be prevented, it is vital to have the ability to detect when an incident occurs, and effectively respond to the threat to contain the damage and recover to normal operations.

"

For threats that cannot be prevented, it is vital to have the ability to detect when an incident occurs, and effectively respond to the threat to contain the damage and recover to normal operations.

MAKING THE CASE

So the day of his first meeting with the Board has finally arrived, and Mr. Ciso feels confident with his security strategy model. Despite only receiving 15 minutes in the meeting, by using his model he was able to communicate with management in a way that they understood and could relate to the cybersecurity challenges that GGC is facing. Some Board Members even left the briefing accepting that cybersecurity is a common organisational responsibility and that the strategy must be executed from the top.

Mr. Ciso was excited with this newly acquired support in the organisation, and relieved that his model could effectively convey the complexity of GGC's cybersecurity needs in a concise and understandable way. His primary objective was successfully completed, and he finally received the funding commitment needed to implement an effective security program at GGC.

Now with his funding secured, his first priority is to establish the resources needed to begin executing GGC's new security strategy. As Mr. Ciso soon discovered, this is a task easier said than done.

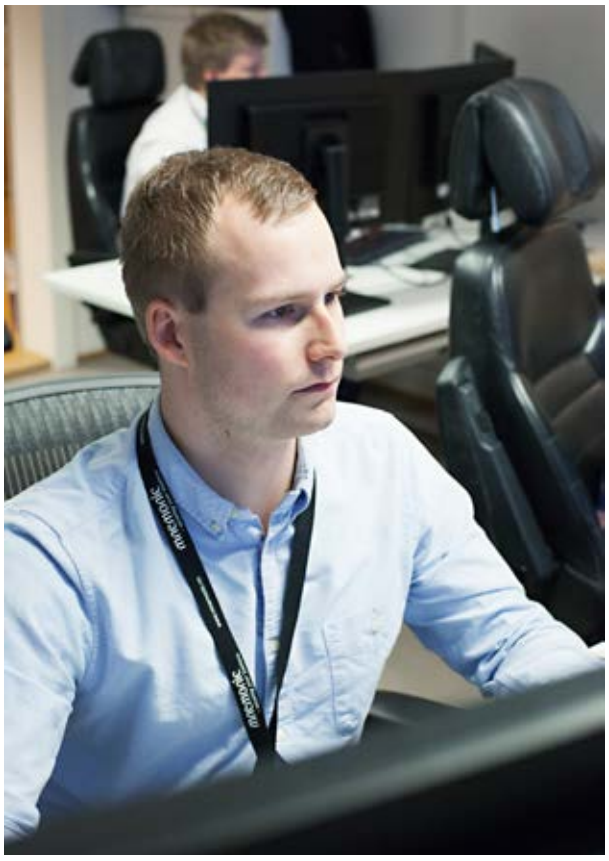

THE MAKE OR BUY DECISION

Make

The intuitive approach was to reinforce GGC's existing security organisation by hiring more internal resources and building the competence in-house. As Mr. Ciso explored this option, he was soon presented with several challenges:

THERE IS A SHORTAGE OF IT AND INFORMATION SECURITY PROFESSIONALS

To say it is difficult to fill positions in cybersecurity is an understatement. With an estimated global shortage of 1 million cybersecurity professionals, organisations looking to fill positions in-house must be ready to compete for cybersecurity talent, and accept that it may take time before positions are filled, if at all.

CYBERSECURITY IS A FULL-TIME JOB

The cybersecurity landscape continuously evolves and changes rapidly, resulting in much of a security professional's role dictated by externalities. Staying up-to-date and relevant requires dedicating a significant portion of time for knowledge acquisition, commonly in the form of attending security conferences, taking courses, and completing certifications to ensure proper professional competence. Aside from the financial investment, these activities also consume a significant portion of a security professional's most valuable and in-demand resource – their time. This is compounded with the daily necessity to stay abreast of changes in the threat landscape, security announcements, new vulnerabilities, product developments, and so on. All this while performing their sanctioned task of protecting their organisation is a full-time job.

A SECURITY STRATEGY IS BUILT UPON A COLLECTION OF NICHE ROLES

Cybersecurity involves a wide array of related, but distinct specialties and disciplines. Risk management, solution specialists, digital forensics, incident response, application security, threat intelligence, and security analysts are merely a subset of the defined roles and skillsets commonly required for any security program. Each of these roles are highly specialised, and requires distinct set of skills, training and experience. While any one person may be qualified in multiple disciplines, the individual roles themselves are often too demanding for even the most adept people to fulfil more than a couple.

Buy

In light of these challenges, Mr. Ciso has now begun evaluating the outsourcing possibilities to fulfil some of the functions outlined in GGC's new security strategy. Mr. Ciso is faced with several key considerations:

FUNCTIONS TO OUTSOURCE

What functions are available to be outsourced? Is it at a strategic level? Operational level? In the form of technology experts? Advisors? Managed services? Will it be on-demand or constant? Temporary or permanent?

SECURITY OF OUTSOURCING PARTNER

What physical, logical and procedural safeguards does the outsourcing partner have in place? Does the partner have any attestations and certifications to back these up? How are personnel screened? Who will have access to GGC's data? Is GGC permitted to audit the partner?

RELATIVE SIZE

How big (or small) is the outsourcing partner? Will GGC represent a large customer for the partner, thereby potentially receiving more attention, have more bargaining power and a stronger influence to shape service development? Or will GGC be one of many similarly sized or non-strategic customers, exerting little influence on the partner or potentially seen as less important?

FLEXIBILITY

How flexible is the partner? Will the partner match the service delivery to GGC's requirements, or will GGC be required to adjust their requirements to match the service? Does GGC need the service tailored to them, or is a standardized service more suitable?

ABILITY TO SCALE UP AND DOWN

Are GGC's outsourcing resource demands for each function constant over time, or do the demands fluctuate throughout the year? Are resource requirements predictable and therefore can be planned for in advance, or will demands arise unexpectedly, such as during a security incident? Can the services provided by the outsourcing partner scale up and down to match these resource demands as GGC's needs change over time?

LOCATION

Where is the outsourcing partner located? Where are the personnel who will actually be delivering services to GGC located? Are resources required to be on-site on occasion or even permanently? Are there any cultural considerations? Where is data stored? Is the partner compliant with the rules and regulations for my industry and country? Does the partner operate in a compatible time zone? Is GGC's intellectual property protected by the trademark, patent and copyright laws in the country where the outsourcing partner operates? Are these enforceable in that country?

COMPETENCE AND SIZE OF RESOURCE POOL

How many people are qualified to meet the requirements? Are these resources actually available? What experience do they have with similar requirements? Is it important to have consistency with the same people delivering the services? What is the employee turnover rate?

GROWTH SUSTAINABILITY

Is the partner capable of supporting GGC's growth, whether this is in the size of the company, expansion to new markets, or entering new industries? Can the partner provide complementary services or take responsibility for additional functions in the future?

COMMUNICATION

How well does the outsourcing partner understand Mr. Ciso's requirements? Can the partner communicate effectively with GGC's team? Does the partner clearly outline the expectations they have of GGC? Are roles and responsibilities clearly defined?

However, fundamental to selecting a security partner is accepting that governing bodies (Board of Directors or equivalent) and senior management are ultimately responsible for protecting the organization's information. Senior management and governing bodies collectively have the responsibility and accountability for setting the organization's objectives, defining strategies to achieve those objectives, and establishing governance structures and processes to best manage the risks in accomplishing those objectives.

This brings us to today, where Mr. Ciso is evaluating all of the considerations above to implement his security strategy. And while there are many functions that Mr. Ciso can outsource, responsibility is not one of them.

The industrialisation of cyberattacks has evolved to the point where if you strip away the malicious intent of the objectives, it is difficult to differentiate the operations of an attack group from that of a normal organisation

Global

GLOBAL CYBER ALLIANCE

Philip R. Reitinger

President & CEO

The Global Cyber Alliance is an international, cross-sector, non-profit dedicated to confronting cyber risk and improving our connected world. GCA's mission is to identify systemic cyber risks and bring together the people and resources to identify and implement a solution and to measure the effect. "Do Something. Measure It."

WHAT IS YOUR BIGGEST SECURITY CONCERN?

My biggest security concern is the Internet of Things. The combination of network complexity orders of magnitude greater than now along with connection to the physical world and weak on-device security capabilities will increase risk exponentially.

IN WHAT AREAS OF SECURITY DO YOU THINK WE'RE FALLING BEHIND?

Many areas, but I'll focus on security at scale. Our systems are not dealing well with the greater risk that stems from more systems, more code, and more vulnerabilities. We need more significant progress in automation and usability.

WHAT GIVES YOU HOPE FOR THE FUTURE OF SECURITY?

Our best hope for the future is to build an infrastructure that defends itself at Internet speed - an Internet immune system. To enable this, we need better automation, security interoperability among services, devices and agents, and strong authentication.

SEMI-AUTOMATED CYBER THREAT INTELLIGENCE (ACT)

MARTIN EIAN, PH.D.
Senior Threat Researcher
mnemonic

Accurate, timely and relevant threat intelligence is a necessity to combat modern threats. However the way we as a community collect, analyse and share threat intelligence today is inefficient, restrictive and needs to be improved. So what can we do about it? Industry, government and academia have joined forces to find out.

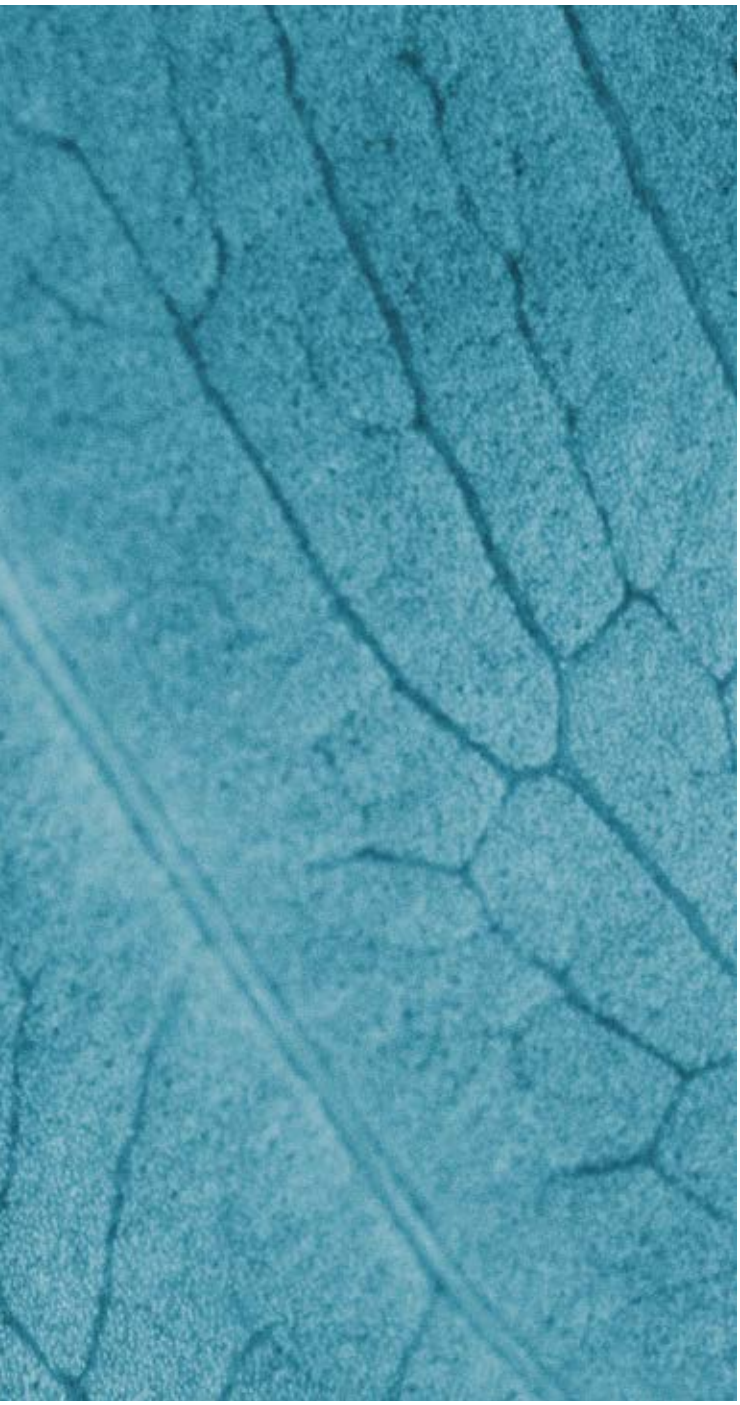

In 2016, mnemonic launched the research project "Semi-Automated Cyber Threat Intelligence (ACT)". The project is funded by the Research Council of Norway, and the project partners are mnemonic, the University of Oslo (UiO), the Norwegian University of Science and Technology (NTNU), the Norwegian National Security Authority (NSM), FinansCERT Norge AS and KraftCERT AS.

The main objective of the ACT project is to develop a platform for cyber threat intelligence to uncover cyberattacks, cyber espionage and sabotage. The project will result in new methods for data enrichment and data analysis to enable identification of threat agents, their motives, resources and attack methodologies. In addition, the project will develop new methods, work processes and mechanisms for creating and distributing threat intelligence and countermeasures to stop ongoing and prevent future attacks. The project duration is three years, and the platform will be made available under an Open Source license.

In order to understand the project background and results, a cursory understanding of cyber threat intelligence is required. Readers familiar with the core concepts of threat intelligence can skip the following section.

The platform will be made available under an Open Source license.

CYBER THREAT INTELLIGENCE

Gartner defines threat intelligence as follows:

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

- Gartner (2013)

In short, threat intelligence is knowledge about threats. This knowledge must be based on evidence, and it must be actionable. Note that knowledge is different from data and information. Knowledge resides in the human brain, while data and information can be stored and processed externally.

Information is data put into context, and the process of extracting information from data can be automated. Intelligence, however, requires explanation, and at the highest ambition level, prediction. Extracting intelligence from information is currently a manual process. We do not believe that it is feasible to fully automate this process without a revolutionary breakthrough in artificial intelligence, hence we use the term "semi-automated" in the project name.

In threat intelligence, the terms threat agent and indicator of compromise are frequently used. A threat agent or threat actor (TA) is an individual or group that poses a threat to a victim. The threat agent might be referred to using its true identity if known, or an alias if unknown. A threat agent has associated capabilities and intents. A threat agent also has a history of

In mnemonic, we use the following model to differentiate between information products and intelligence products:

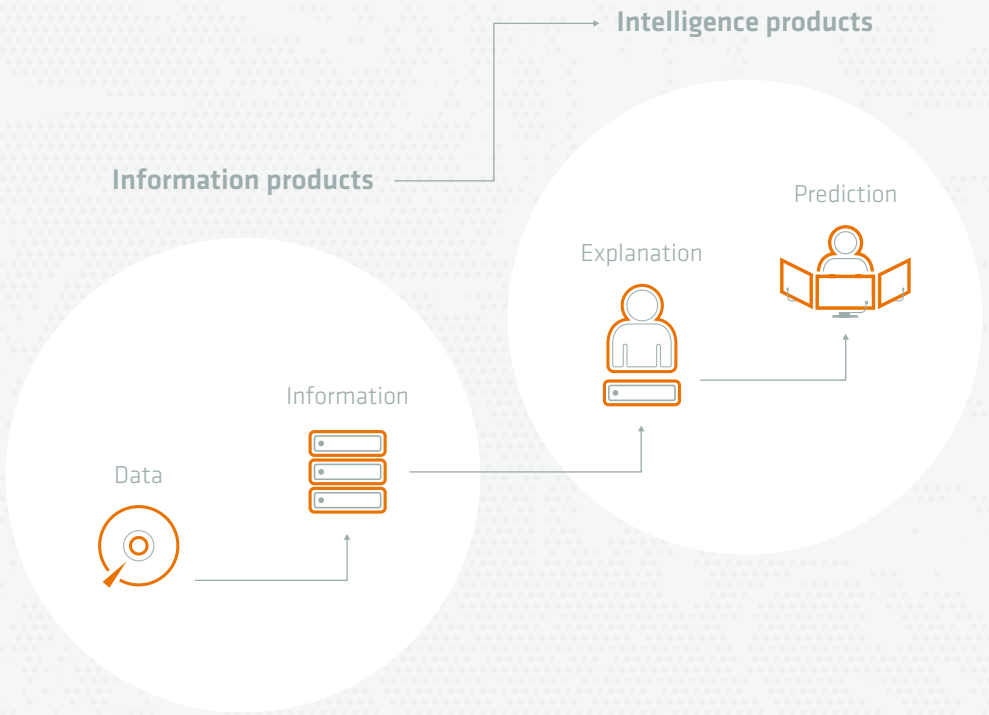

past activity. The analysis of this history forms our knowledge of the threat agent, including its capabilities and intents, and this knowledge can help us predict and detect future attacks.

An Indicator of Compromise (IOC) is defined by RSA as follows:

"An Indicator of Compromise (IOC) is a forensic artifact or remnant of an intrusion that can be identified on a host or network."

- RSA (2012)

An IOC might be an IP address or domain known to be used for command and control purposes by a threat agent. It might be a cryptographic hash of a file identified as malware. In short, it consists of an atomic observable (e.g. IP address) combined with metadata indicating why this observable is malicious.

Threat intelligence is commonly divided into four primary areas or levels: strategic, operational, tactical and technical.

STRATEGIC THREAT INTELLIGENCE

Strategic threat intelligence covers threat agent intents and long term trends to serve as strategic decision support to C-level executives.

OPERATIONAL THREAT INTELLIGENCE

Operational threat intelligence aims to uncover threat agent campaigns or operations and their victims to put several different attacks into context.

TACTICAL THREAT INTELLIGENCE

Tactical threat intelligence studies the threat agents' tactics, techniques and procedures (i.e. how they operate).

TECHNICAL THREAT INTELLIGENCE

Technical threat intelligence studies the threat agents' tools and the observable traces of their attacks, to support detection of the threat agents' activity at a technical level.

PROJECT BACKGROUND

Our primary motives for launching the project were to provide a holistic workspace for analysts, automate repetitive tasks, facilitate advanced automated analysis, improve our knowledge of threat agents, facilitate efficient and accurate manual analysis, automate sharing of threat information and countermeasures, and automate the processing of unstructured data.

The status quo is that threat intelligence analysts use numerous different systems for their daily tasks. Their workspace is fragmented. They copy and paste data from system to system, then manually try to collate the results. As an example, analysts might want to find information related to an IP address. They might issue queries to WHOIS, PassiveDNS, geolocation services, malware databases and reputation lists, then collate the results. The process is then repeated the next time an analyst investigates a new IP address. The platform aims to automate such processes, to provide a holistic view of the collated information, and to retain the information for future use.

An important part of the ACT project is to facilitate sophisticated enrichment of data and the application of artificial intelligence techniques for automated analysis of data and information. These two research areas are the main responsibility of the universities participating in the project.

In order to better understand threat agents, we need to develop formal knowledge representations covering strategic, operational, tactical and technical aspects. We also need to develop methods for making these knowledge representations easily understandable for analysts.

Automated threat information sharing and automated countermeasures can significantly improve detection and prevention capabilities. An important part of the ACT project is to review existing standards and protocols for information sharing and countermeasures, to support the relevant standards, and to develop new standards and protocols if needed. Existing standards include CybOX, IODEF, MAEC, OpenDXL, OpenIOC, RID, STIX, TAXII, and VERIS.

Finally, masses of data relevant to threat intelligence are available in unstructured formats. Examples include threat reports, academic papers, news articles, blogs, e-mail lists, and wiki pages. Unstructured data cannot readily be included in a formal knowledge representation. To overcome the issue of unstructured data, we investigate techniques for extracting structured data from unstructured data.

Many threat intelligence platforms are already available, but none of them meet all of the requirements outlined above. One common deficiency is the ability to model tactical knowledge about a threat agent, linking the strategic and operational levels to the technical level. Another common issue is that several of the platforms are malware centric, with little support for modelling the complete knowledge of a threat agent.

THE DATA MODEL IN ACTION

Objects represent any real world entity, such as an IP address, a person, or a file.

Facts represent relations between objects or additional information about a single object.

EXAMPLE SCENARIO

A security incident is detected – Attack 1. BadDomain.com is found to be used in the Attack 1, and this domain resolves to IP address 1.2.3.4. mnemonic and Europol both attribute 1.2.3.4 to being involved in Attack 1. mnemonic investigates Attack 1 and attributes this to Threat Agent Z.

Attack 2 occurs. Trusted Intel Partner confirms that WorseDomain.com is used in Attack 2. WorseDomain.com resolves to 1.2.3.4 – the same IP address used in Attack 1.

By following the facts through the data model, it is possible to infer that Threat Agent Z was also responsible for Attack 2.

The certainty in your conclusion is dependent on many factors, including the confidence in the attributed facts, the reliability of the source, or if there are conflicting facts, amongst others.

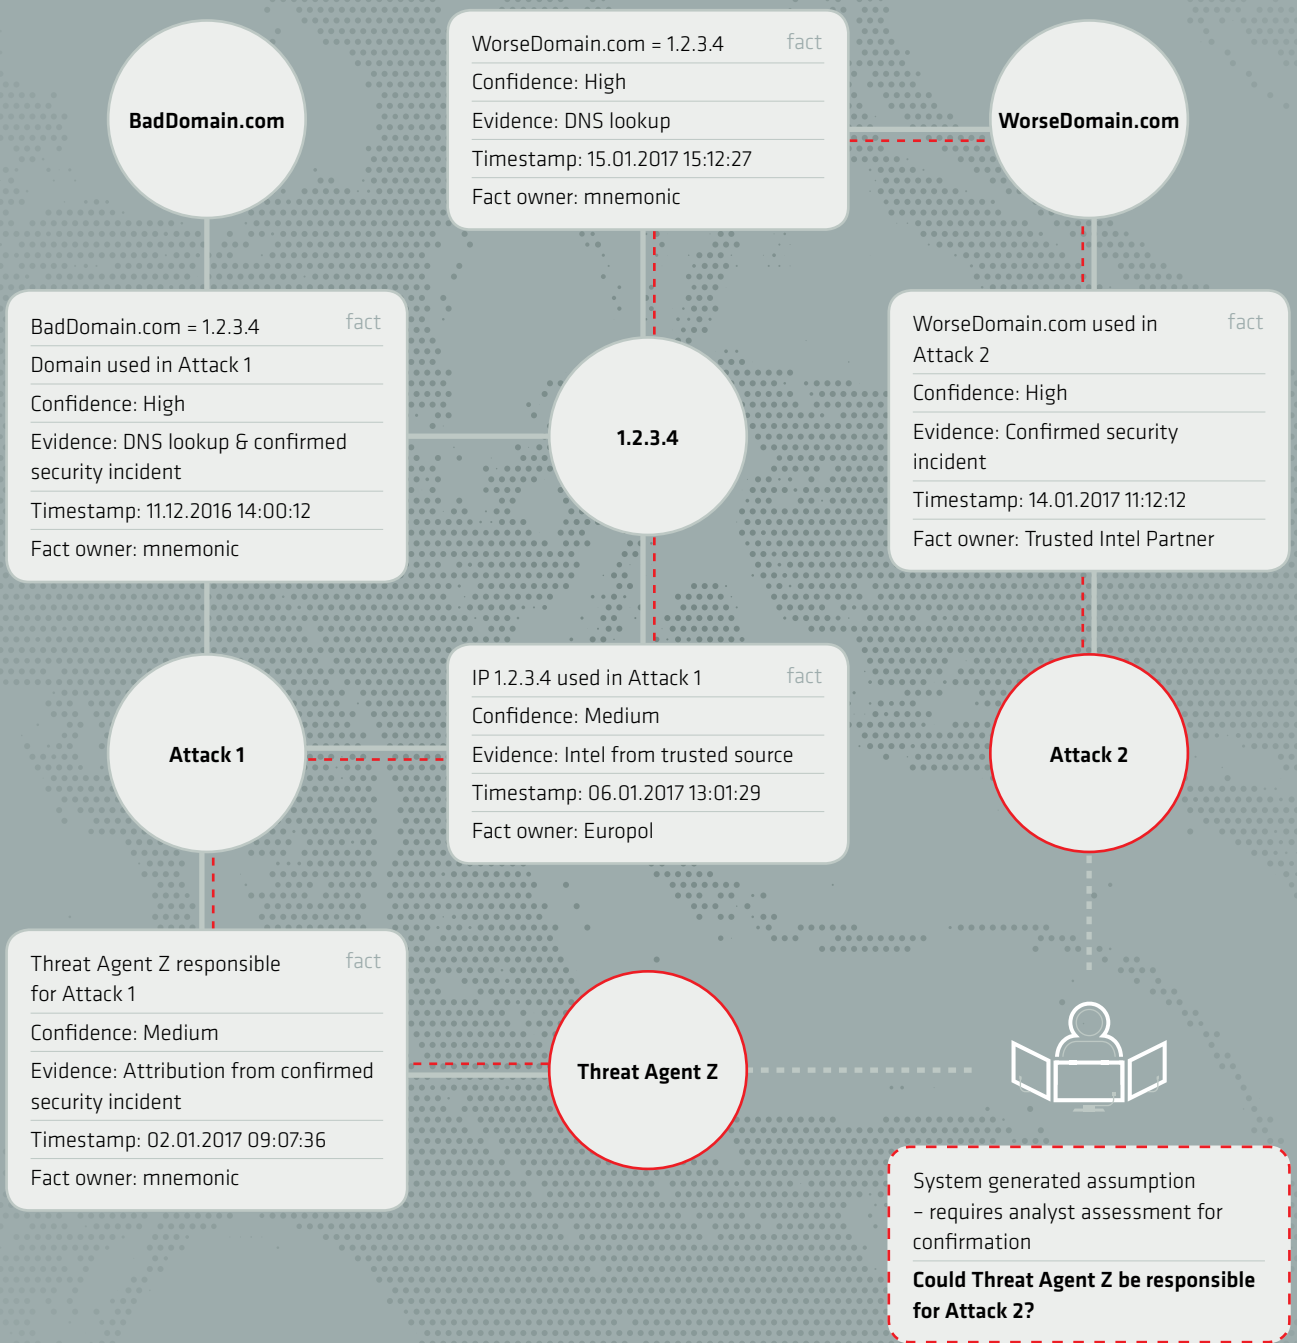

PROJECT RESULTS

During the first six months, we have developed a low-level data model, the platform architecture, a workflow model, and an application programming interface (API).

The data model supports a large set of objects, on which the system can attribute facts collected from analysis and intelligence sources. Each object represents a real world external entity, which exists independently of any other information, such as an IP address, a person, or a file. The facts represent relations between objects or additional information about a single object. Such facts may be uncertain, may be sensitive, and even contradicting or overlapping, and need to be traceable and backed by evidence and assessments.

A key property of the data model is that facts are immutable, meaning once they have been added to the system, they cannot be modified or deleted. Furthermore, all facts are timestamped. The combination of these two properties makes it possible to "rewind time", to review the knowledge representation that existed at any point in time. Having this ability can be very useful – for example to determine the basis for a decision that was made several months ago.

We have investigated several models for representing knowledge about threat agents: the Detection Maturity Level (DML) model, the Pyramid of Pain, the Cyber Kill Chain, and the Diamond model. We have selected the DML model as a possible starting point for developing a threat ontology. The DML model was proposed by Ryan Stillions in 2014. Our slightly extended version of this model is shown below:

THE DETECTION MATURITY LEVEL (DML) MODEL

In collaboration with two other research projects, TOCSA and Oslo Analytics, we presented an academic paper at the 11th International Conference on Semantic Technology for Intelligence, Defense, and Security (STIDS 2016) where we explain how the DML model can be used to develop a threat ontology. One major obstacle to overcome is to develop formal definitions of goals, strategy, tactics, techniques, and procedures. We have reviewed the MITRE CAPEC and ATT&CK taxonomies in order to determine if they can be used to describe tactics and techniques in the DML model. Our findings so far indicate inconsistencies that make it difficult to use them directly as formal definitions.

In parallel with the above work we started developing an ontology for DML-1, DML-2 and DML-3. This ontology was validated during incident response, where we used a graph database to implement a knowledge graph of the incident. The knowledge graph quickly proved to be useful, enabling us to compare evidence from the incident to known threat agents, quickly assess what happened during a time period, and to quickly get an overview of different types of threat agent activities.

We have also researched methods for extracting structured data from unstructured data. Candidate approaches to this goal are natural language processing (NLP), natural language learning, and neural networks with word embeddings. The most significant finding so far is that there is no readily available NLP corpus for the cybersecurity domain, and that a domain specific corpus has to be developed in order to apply NLP.

CONCLUSIONS

The ACT project is still in an early phase, but we have already produced some interesting and useful results. We will in the near future set up a Github repository for the project, where we will publish platform documentation and code. The code will be available under an Open Source license.

Based on observations from mnemonic's Security Operations Center, this year you should expect one confirmed security incident for every four users in your organisation

Europe

EUROPOL EUROPEAN CYBERCRIME CENTRE (EC3)

Steven Wilson

Head of European Cybercrime Centre

Europol set up the European Cybercrime Centre (EC3) in 2013 to strengthen the law enforcement response to cybercrime in the EU and thus to help protect European citizens, businesses and governments from online crime. EC3 delivers technical, analytical and digital forensic support to cybercrime investigations in the three main crime areas of: online child sexual exploitation, transnational payment fraud and online fraud and high-tech crime.

WHAT IS YOUR BIGGEST SECURITY CONCERN?

The development of IoT based attacks is a major concern. This past year we have seen a range of DDoS attacks from IoT connected devices at a magnitude never experienced before. The scale, impact and consequences of these attacks will only continue to grow as more vulnerable IoT devices come online. The sheer volume and size of these attacks have the ability to knock out significant organisations, critical national infrastructure or even countries.

IN WHAT AREAS OF SECURITY DO YOU THINK WE'RE FALLING BEHIND?

The lack of basic security on IOT devices has created a significant vulnerability that can be exploited by a wide range of actors. Basic security principles should be applied to all devices capable of being attached to the internet or we will have an uncontrollable problem as these devices proliferate.

WHAT GIVES YOU HOPE FOR THE FUTURE OF SECURITY?

The increased cooperation between the technical security sector and law enforcement gives me confidence for our future. I see significant examples of Corporate Social Responsibility where companies are working to assist law enforcement in investigations or research to solve problems for no reward other than making the internet a safer place for society.

THE HUMAN ELEMENT OF CYBER ATTACKERS

#FinallyTimeOffWork #FamilyVacationTime #FullyFunctional

ERIK ALEXANDER LØKKEN
Manager of Managed Security Services
mnemonic

Attackers are people too. And as people they have the same basic and fundamental needs as the rest of us. Shelter, food, water, family, health, financial security and so on. Despite the fact that we're on opposite sides of the battlefront, are we really all that different?

Let's start with a familiar story: a person wakes up to their alarm. This is the same alarm, set at the same early hour of the day that has woken them up every morning for the past 10 years, and kicks off their daily routine: wake the kids up, make them breakfast and try to get them out the door in time to catch the school bus. Jump in the car, get caught in highway traffic and be amazed that despite the number of times they change lanes, the other lane always seems to be going faster. Arrive at the office, pass through security and sit at their desk adorned with artwork from their kids and pictures from their most recent family vacation.

This story could be used to describe a large portion of the global population, except for a simple, albeit critical difference: this person is a cyber attacker and their job is to breach your network, digitally extort you, steal confidential secrets and conduct cyber espionage.

MALWARE DOESN'T HACK PEOPLE – PEOPLE HACK PEOPLE

Unlike the natural world, where viruses and bacteria are a naturally occurring, biological piece of the evolutionary puzzle, there is nothing natural about computer viruses and malware in general. All malware is created by people. All of it. And in similar fashion to the adage "guns don't kill people, people kill people", all malware is consciously executed by a person.

Malware and attackers target vulnerabilities in software and in people (e.g. social engineering). However, vulnerabilities themselves are not dangerous. They are merely a susceptibility or flaw that represent a potential to be exploited, and this exploitation can only occur if two other conditions are true: a person (attacker) has access to the vulnerability, and a person (attacker) has the capability to exploit the vulnerability. People are the real threat, not vulnerabilities.

People are the real threat,
not vulnerabilities.

JUST ANOTHER DAY AT THE OFFICE

The industrialisation of cyberattacks has evolved to the point where if you strip away the malicious intent of the objectives, it is difficult to differentiate the operations of an attack group from that of a normal organisation.

These attack groups – particularly the well-funded, operationalised nation states – have regular employees who show up on a daily basis, working the traditional 9-to-5 job with a 30-minute lunch break. Once 5 o'clock hits, a shift change will occur, where the night crew take over. When times get busy, threat actors will also scale up their operations by increasing the number of operators and moving to 24-hour operations.

The financially motivated attackers targeting netbanks are no different. Mechanisms like two-factor authentication make it more difficult for attackers to operate fully automated attacks. Attackers countered this and now rely on human operators to work in real-time to phish the one-time passwords generated by two-factor authentication and empty your bank account. The hours we see these operators work? Consistently 8-to-4 in their local time zone. Attackers also need to take their kids to soccer practice in the evenings.

As with any service operation, proper IT service management is important to reduce downtime and improve service quality. Threat actors are no different. They have operating procedures and routines to perform their duties, defined roles for service delivery and paths for escalation – yes even the bad guys need 2nd and 3rd line support.

[Cathy] Hi, I'm Cathy, and I'll be your malware support specialist today. How can I help you?

[me] I have successfully infected a group of clients, but the keylogger does not seem to be working.

[Cathy] This is a known problem with version 1.2.4. Upgrading your toolkit to 1.2.5 should resolve this issue.

AN UNDERGROUND FREE MARKET

Whereas nation states may have the financial and technical resources to produce custom malware in-house to fulfil their espionage operations, financially motivated and opportunistic threat actors – such as those engaged in netbank fraud or ransomware – may not have this luxury. But attackers need not fear – the entrepreneurial spirit is well represented in the cyber underworld, which has led to a mature, free market economy being established.

> **The entrepreneurial spirit is well represented in the cyber underworld.**

There exists an entire supply chain ecosystem for the production, management, operation and maintenance of malware. This is complete with software maintenance agreements, multi-tiered support agreements (yes, this even means silver, gold and platinum packages with 8x5 or 24x7 phone and email support options), tutorials, regular software upgrades, and even Service Level Agreements (SLAs). We have even observed some lucrative SLA agreements, complete with money-back guarantees.

Not looking to make a capital expenditure on malware? Fear not, you can claim your next attack on your operational budget by renting a DDoS attack by the hour.

Whether performing espionage, sabotage, supporting political motivations or simply looking to make a quick buck, one common truth remains – malware is big business. The production of malware to support this business follows standard software development lifecycles, and similar to non-malicious software development, malware is prone to errors, mistakes and vulnerabilities – after all attackers are people as well.

The industrialisation of cyberattacks has evolved to the point where if you strip away the malicious intent of the objectives, it is difficult to differentiate the operations of an attack group from that of a normal organisation.

ATTACKERS ARE PRONE TO MISTAKES

Attackers are no different when it comes to humanity's basic needs – namely food and sleep. As humans, we're prone to make mistakes and have reduced productivity when we're deprived of either of these basic needs.

For example our observations have shown that attackers are more prone to make errors before their lunch break. Reflecting on your own workday, there are several plausible explanations for this. It can be that their glucose levels are too low, which affects concentration and promotes inaccuracies. Alternatively, or most likely in combination, just before lunch the attackers may be rushing to quickly finish the task they are performing so they can take a break, have some food and replenish those sinking glucose levels.

> **Moral of the story: don't get attacked on Tuesdays.**

Similarly, we have found that attackers are more likely to make mistakes on a Friday than on any other day of the week. Tuesdays on the other hand are when attackers are at the top of their game and proportionately make the fewest number of mistakes. Moral of the story: don't get attacked on Tuesdays.

But what kind of mistakes are we talking about?

FAT-FINGER SYNDROME AFFECTS US ALL

Attackers are just as susceptible to fat-finger syndrome as the rest of us, where mis-hitting the keyboard and striking adjacent keys results in tuyping ewrrors. Attackers are also concerned with speed and efficiency, which as anyone who has hastily sent off an email before proofreading it can attest to, results in minor, often preventable mistakes. These minor mistakes however can have major consequences.

Traces and evidence of these errors can be found throughout a victim's network. For example, in DNS logs we can observe the effects of when an attacker has mistyped a domain name in their malware configuration, leading to their infected client not being able to connect to the command & control infrastructure (and likely forcing the attacker to contact their 2nd level support).

We also see evidence in configuration files – dynamic files used by threat actors to remotely configure malware on infected clients – where attackers make mistakes that should be caught in their QA process. A more common example is when the configuration is explicitly referencing file locations on the machine the attacker used to compile the malware rather than using environmental variables to make their configuration compatible with the victim's machine.

Live operators are in real time remotely connected to infected clients and are issuing hidden commands to infect other machines, steal passwords, and exfiltrate data - all unbeknownst to the victim.

A BRIEF EXPLANATION OF DNS

The DNS protocol resolves human readable domain names to machine friendly IP addresses so that you can type www.google.com into your browser instead of 216.58.217.142. It is important for DNS requests to be communicated quickly (who wants to wait for a webpage to load?), so to encourage speed in message delivery, there is a sacrifice made in communication reliability.

DNS operates using the UDP protocol, one of the two core protocols that make up the Internet protocol suite. This is a connectionless protocol that operates on a best-effort basis, meaning its messages are sent without confirmation that they have been delivered to the recipient. The messages themselves may take different paths on the Internet to get to the recipient, resulting in some messages being received out of order, while others can be delayed or even lost along the way.

NB: DNS can also operate on TCP under certain circumstances, but UDP is the default.

REAL-TIME MISTAKES

These errors tell us that there are humans behind the production, configuration and distribution of the malware. The nature of these errors indicates the mistakes were likely made at some point in the past, occurring as part of the development and configuration process.

However, other errors that prove to us that in many cases there are live operators on the attackers' side. These live operators are in real time remotely connected to infected clients and are issuing hidden commands to infect other machines, steal passwords, and exfiltrate data - all unbeknownst to the victim.

To operate undetected and under the radar in well-defended networks, attackers – particularly those conducting extended espionage campaigns – must employ creative and cunning methods to maintain open communication channels with infected clients. One such example is tunnelling and concealing their communication (and thereby issuing their commands) in the underlying protocols we rely upon for the Internet to function, such as DNS.

As a connectionless protocol, DNS messages are susceptible to being delayed, delivered out of order, or simply lost without a mechanism to correct errors or re-request missing messages. For legitimate DNS messages, this functions well, however when attackers manipulate DNS messages to contain their malicious commands, they will receive mixed results.

When an infrastructure is built upon unreliable protocols, it can only be expected to receive unreliable results. For example, it is common to see typing errors that are presumably caused by the delay between what the operator is typing and what is being received by the infected client. Think about when your computer is running slowly and it takes a few second for your keystrokes to show up on the screen – what do you do? Most of us will hit the same key a second, third or fourth time to

see if that will help (hint: it never does). Operators are no different, and we see this in the commands they issue with repeeeeeated characters.

Observing this type of behaviour is usually a strong indicator that there is a human operator on the other side.

DEFEND AGAINST PEOPLE, NOT MACHINES

Humans are an adaptive species. In biological terms, plasticity represents an organism's ability to adapt to changes in its environment. It is this plasticity that has enabled humans evolve to where we are today. This same plasticity enables attackers to adapt to your defences, making calculated decisions and adjusting their behaviour as a matter of survival within your environment. When your defences hinder their attack, they will adapt and counter in a calculated and often unpredictable manner that is uniquely human.

Malware itself is not the threat. Malware is merely the tool that enables attackers to execute their objectives. While technology plays a large and important role in enabling organisations to detect the tools attackers use, do not overlook the human element when building your cybersecurity strategy.

Because people are the real threat here. People who are going to work to put food on the table, a roof over their head and support the basic needs of their families. And people who are as equally motivated to succeed in their attack as you are in defending against them.

mnemonic's Security Operations Center saw a
74% increase in high severity incidents in 2016

Global

GLOBAL OIL AND GAS COMPANY

Anonymous

Chief Information Security Officer

WHAT IS YOUR BIGGEST SECURITY CONCERN?

Bureaucracy which lead to (security and infrastructure) projects being crippled before deployment or worst case terminated.

IN WHAT AREAS OF SECURITY DO YOU THINK WE'RE FALLING BEHIND?

Many, but to name one: attracting the talent. The dark side has way more sex appeal in terms of money, challenge and freedom from Bureaucracy.

WHAT GIVES YOU HOPE FOR THE FUTURE OF SECURITY?

Not much to be honest… It is true that security is gaining traction in the corporate world. That said we are surely on the losing side of things and if we don't start focusing on the offensive aspect of security we will not only lose most battles but eventually also lose the war.

FRODE HOMMEDAL
Head of Incident Response and Security Analytics
Telenor

MAKING YOUR MOVE: BOOTING A PERSISTENT ADVERSARY OFF YOUR NETWORK

Imagine that you're riding the bus. Seated across from you is a person you find cute and attractive. Your mind wanders, as the mind tends to do in situations like these. Maybe you even imagine what it would be like to kiss this person. Then the person looks at you and smiles. What do you do?

This is perhaps not the question you expected in an essay about kicking spies and criminals off your network, but please bear with me.

The easiest thing to do is pretend you didn't see anything and just look away. You can of course also smile back, and perhaps maintain eye contact. If you're good at breaking the ice, you can even try to chat this person up. If successful, you may even agree to meet for a cup of coffee later on and take it from there.

One thing you most likely *won't* do, regardless of how attracted you are to this person, is lean across the aisle and start kissing them. No, *making a move* requires a lot more finesse than that, a lot more reading of subtle signals and figuring out chemistry and dynamics.

You don't just kiss strangers. It is counterproductive on all but the rarest occasions.

RUSHING TO ACTION

It's not just between smiling strangers on the bus that prematurely rushing to action will be counterproductive. Truth be told, this is probably more of a rule than it is an exception. It is particularly true in an unexpected high risk situation – which is the topic of this essay. Still, we often rush to action anyway. Our sense of urgency pushes us to do at least *something,* and often enough this *something* turns out to be counterproductive.

If you have ever worked for a company that was hit by a targeted cyberattack – such as an espionage campaign – for the first time, and fought against a determined and well-resourced adversary, I'm pretty sure you share my experience on at least two accounts:

- A lot of people wanted to pretend it wasn't happening. (Who would spy on us?)
- A lot of people wanted to rush to action. (We need to block this in the firewall!)

Or to say it with the strangers-on-the-bus analogy: a lot of people wanted to turn away and look out the window, and a lot of people wanted to rush over and start kissing the stranger.

During a time of crisis, such as discovering a targeted attack, relatively few people instinctively go for the third and more sensible alternative: stopping and thinking everything through before acting cautiously with a clear sense of what the best course of action would be.

If you want to stand a chance when faced with determined spies and criminals, this ability must be deliberately developed.

GETTING A READ ON THE SITUATION

Back to our stranger on the bus. You know it all ends if you just turn away. You know it also ends, and possibly horribly so, if you immediately try to kiss the stranger. You know that for any chance of a future romance you need to start talking, arranging to meet and generally play your cards just right through several phases of intricate interaction.

Very few people have the same kind of instinctive read on the situation when it comes to facing a determined adversary with a foothold in your network. The reason is simple: normally we aren't faced with adversaries who are actively and deliberately exploiting our weaknesses and turning our IT infrastructure against us in ways most people are not aware is even possible.

This is what's happening when you're faced with espionage from a well-resourced and determined adversary, and most people and organizations aren't prepared to effectively deal with it.

If you want to stand a chance when faced with determined spies and criminals, this ability must be deliberately developed.

THE SPY WHO LOVED ME

It is no secret that the large, national intelligence agencies of the world have been conducting an ever-growing amount of highly successful cyberespionage campaigns for more than 20 years now. For the last ten years or so, many of the smaller nations have followed the larger nations' lead.

This is not just strategic military and political espionage we're talking about. There has been just as much economic espionage and theft of intellectual property.

Alongside this development, an economy of government contractors has developed, and some of these contractors have also chosen to operate in the…shall we say grey area outside of their government contracts.

Then of course, there's crime. The criminal underbelly of the emerging digital society and economy is proven to be substantial, and the groups operating in this area are often performing on a level that is comparable to that of government contractors and nation states.

There are a lot of potential adversaries out there for companies of any size and type, and it has been like this for some time. Historically they were left to operate with impunity – and they did. They happily hacked the world, and we were none the wiser.

They still do of course, but now we know, and it's time we deal with it.

You can change as many locks you want, but if you leave your attacker with just one working key, they can just wait for the dust to settle before sneaking back in.

MASTER OF THE HOUSE

Nowadays there are security vendors and security teams everywhere trying to combat cyberespionage. This was not the case 10 years ago. Go back 20 years and most people hadn't even heard of it.

However, some defenders came along for the ride early. Many were working for government and military agencies and offices that over the years were repeatedly targeted by persistent adversaries. They were given the opportunity to develop skills and knowledge over a decade before anyone else. It is these people that first gave us one of the most important lessons learned from dealing with targeted attacks from advanced, determined and persistent adversaries:

Do not take overt defensive actions that will tip off your attacker before you have thoroughly scoped the breach, and are ready to shut down the attacker's operation abruptly and completely. Even if you are 99% successful in closing down your attacker's access to your network, you will be 0% successful in keeping the attacker out of your network if you "clean up" and move on.

You can change as many locks you want, but if you leave your attacker with just one working key, they can just wait for the dust to settle before sneaking back in, establishing a new and improved foothold and continue their mission as before.

The same goes if you don't cut off an adversary's entire access simultaneously, but instead spend days or weeks doing so. This will give your adversary time to evade your counter-measures and place a well-hidden backdoor they can use to re-enter your network once you're done "cleaning house".

Not evicting an attacker properly is a massive waste of resources.

Your adversaries hide and operate in your knowledge gaps. Fighting your adversary means fighting those gaps.

EVICTION NOTICE

But how can you know you have uncovered everything? How do you distinguish 99% from 100%? The simple answer is that you can't, and even determining if you're closer to 100% than 20% is surprisingly difficult. Let that sink in. It's actually hard to figure out if you have a real chance of success, or if you have no chance at all.

It's a tall order to fight advanced and persistent adversaries, and the you-gotta-be-this-tall-to-ride mark is carved in quite high on the wall for this particular rollercoaster of a ride. It requires a motivated, mature and highly capable security team, and a defendable infrastructure with a well-designed and effectively executed security monitoring and incident response strategy to succeed.

This doesn't mean we should give up, but rather illustrates the need for a structured approach to reduce this uncertainty. We know we will never know for sure, but we need to find a way to gauge whether we have a real chance of success or not.

Your adversaries hide and operate in your knowledge gaps. Fighting your adversary means fighting those gaps.

MIND THE GAP

The number and severity of knowledge gaps you have about your adversary will likely have a strong correlation with your chances of successfully evicting this adversary. If you can identify obvious gaps in your knowledge about how they operate within your network, you're running a substantial risk of failing to evict them successfully. So look for knowledge gaps, and work to close them.

When you start to struggle in finding gaps you can infer that your knowledge gap is closing. You can estimate that you are at least getting closer to "100%", and hence closer to being ready to evict.

Either that, or you have reached the limits of your abilities. Remember, you need to be pretty tall to ride.

It's also worth considering that even if you know enough to show your hand and act, you still need to be able to swiftly transform your knowledge into eviction, prevention and detection to successfully thwart the intrusion and keep your adversary out.

> **Knowledge doesn't save you if you lack the ability to act on it, and all your effort is wasted if your adversary can effortlessly re-enter your network after your eviction attempt.**

Knowledge doesn't save you if you lack the ability to act on it, and all your effort is wasted if your adversary can effortlessly re-enter your network after your eviction attempt.

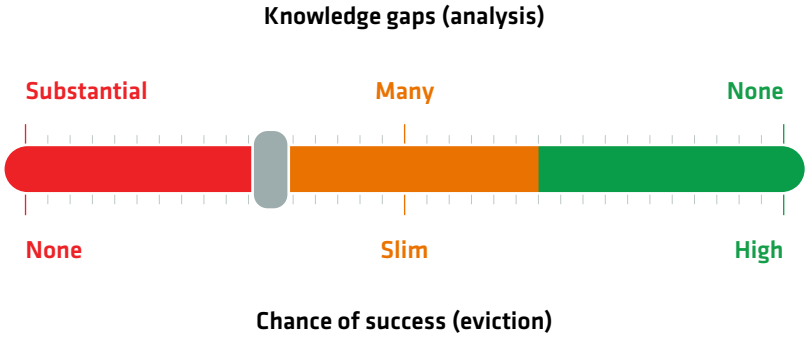

The slider must move to the right to have any chance of success

TAKING THE RED PILL

But how exactly do you use the knowledge you've acquired about an adversary to make an informed decision on when it's time to give them the boot? This is where The Cyber Threat Intelligence Matrix comes into play.

The Cyber Threat Intelligence Matrix (CTIM) was created as the result of a seemingly simple question: how can threat intelligence help improve incident response?

The overarching goal of CTIM is to provide a better foundation for good response strategies and informed decision making during an incident, especially regarding attacker eviction. It aims to provide this through a system for assessing your knowledge about an adversary and their campaigns against you. The core idea of CTIM is actually really simple: encourage structured analytic reasoning regarding targeted intrusions.

THE CYBER THREAT INTELLIGENCE MATRIX

The concept behind the CTIM is to categorize knowledge as it is produced, and map it into several three-by-three matrices representing depth of knowledge along stages of an attack. The main purpose of the mapping is to help communicate your findings and recommendations to management (e.g. the crisis management team, who in all likelihood is making the final decision on when to evict).

I've nicknamed the matrix **FAT PIE** after the first letters of each of the rows.

FOOTHOLD:
Atomic indicators, like IPs, domains, hashes etc.

ARSENAL:
The toolbox the attack has been built from, like family of malware being used.

TRADECRAFT:
The behavior of your adversary's operators, like routines.

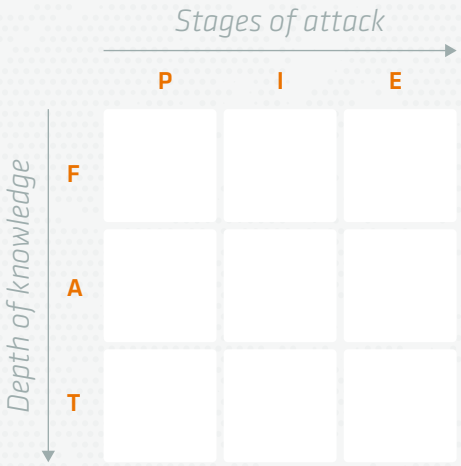

PREPARATION:
The attacker's target development, reconnaissance, preparation and staging.

INTRUSION:
The attacker's penetration of your perimeter, and the establishment of foothold.

EXECUTION:
The attacker's execution on mission objectives within your network.

MATRIX DYNAMICS

Imagine an attacker is tasked to collect specific information from your company. They will start with the preparation, and will be guided by their tradecraft and aided by their arsenal. As defenders, we don't get to observe this directly. Our first observation will probably be footprints left by the attacker operating in the intrusion or execution phase.

What we must do is track those footprints as far back as we can, and try to understand what has been going on. Through analysis we can make inferences about the attacker's preparation, along with their arsenal. Eventually we may also be able to make inferences about their tradecraft.

It's all about catching up with the attacker, keeping up with their operational tempo and start predicting and influencing their decision loops and actions.

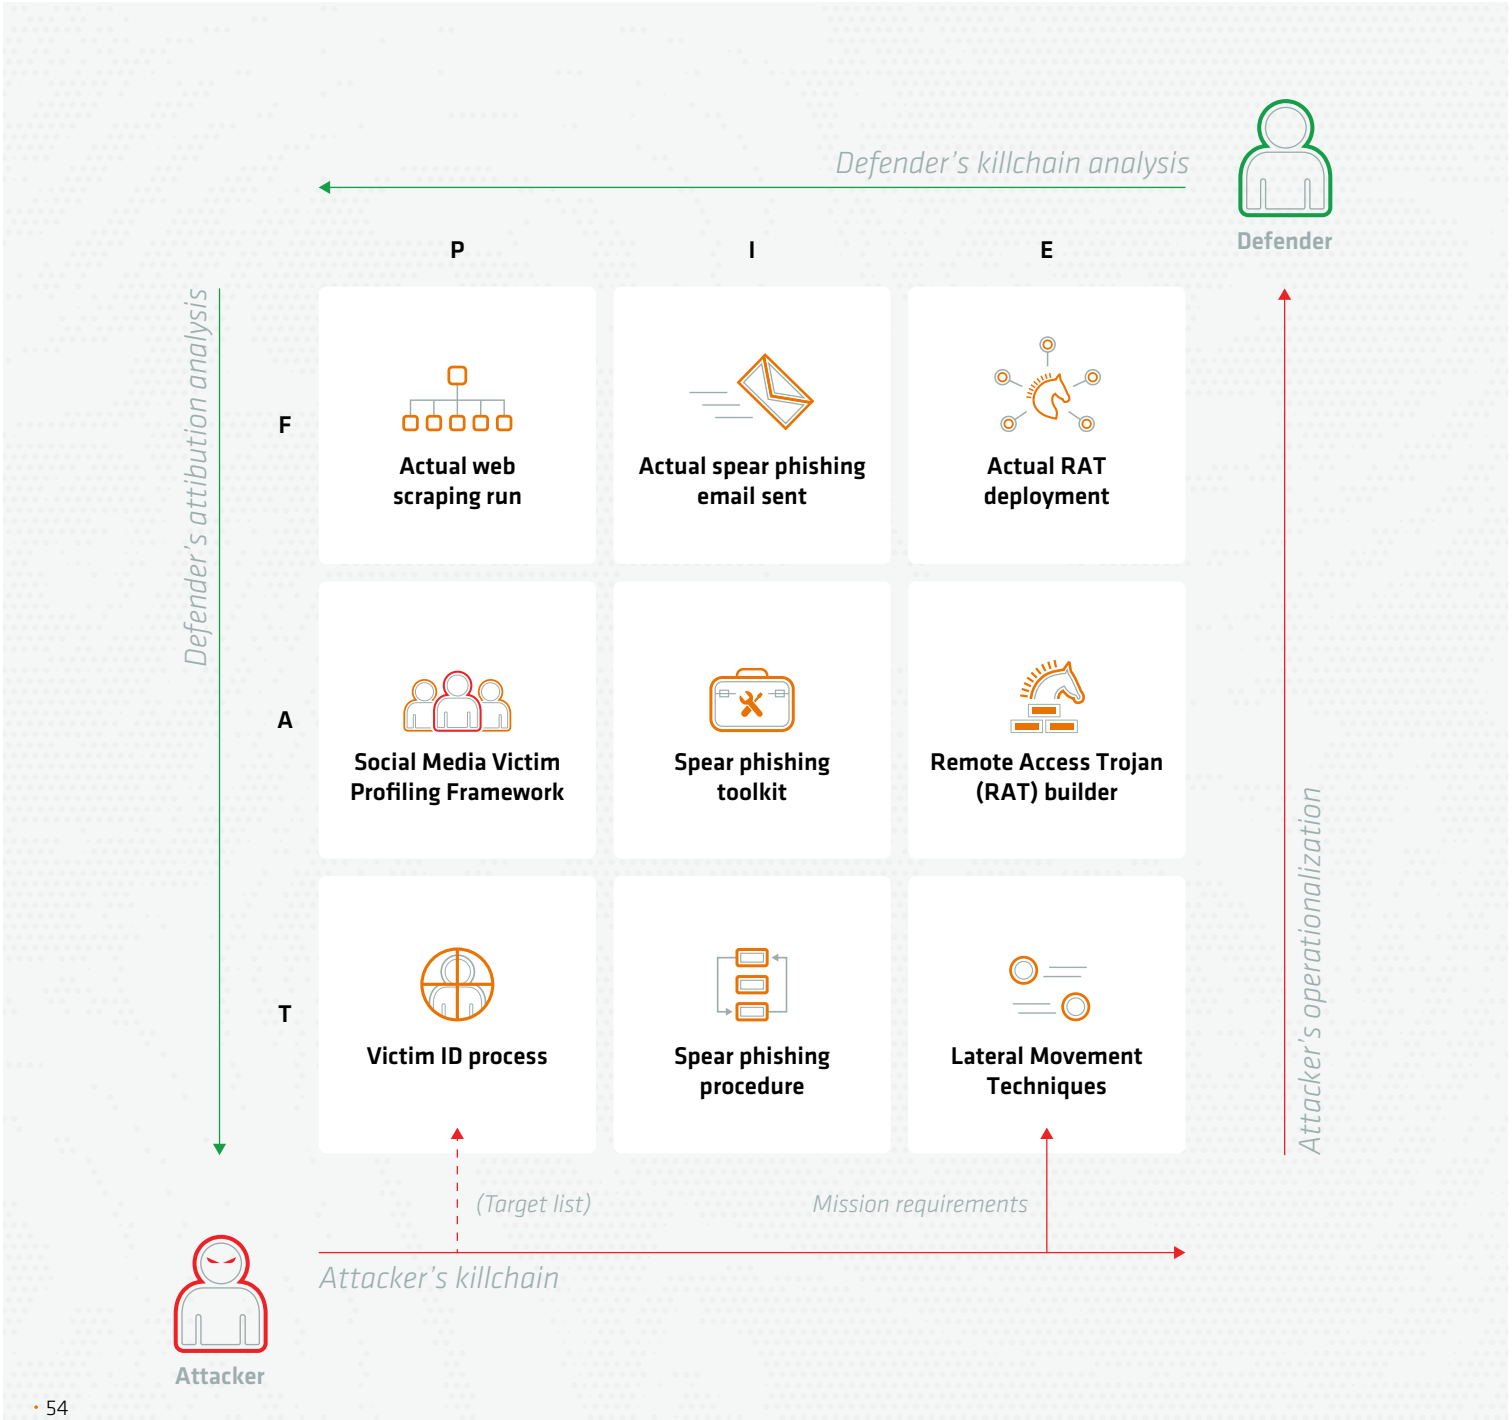

PEDIGREE

The more seasoned incident responders among you will recognize elements from both Ryan Stillions' Detection Maturity Levels, David Bianco's Pyramid of Pain and Mike Hutchins et al.'s Cyber Kill Chain in the matrix. This is not by accident. These are the fundamental building blocks of the matrix's axes – depth of knowledge and stages of attack.

Maybe not so obvious are the ties to the Diamond Model of Intrusion Analysis by Caltagirone, Pendergast and Betz, but this would be the preferred method for actually identifying knowledge. For proper classification of findings you will also need to be intimately familiar with MITRE's ATT&CK nomenclature to be able to categorize your findings, and a solid understanding of Boyd's OODA loop principles will be helpful when assessing your dynamics with the attacker.

USING THE MATRIX

A central idea of the visual design of the matrix is to more effectively communicate to management the risk of prematurely rushing to action. This is achieved by using a scheme that will be familiar to them: a matrix with traffic light colours, in addition to white.

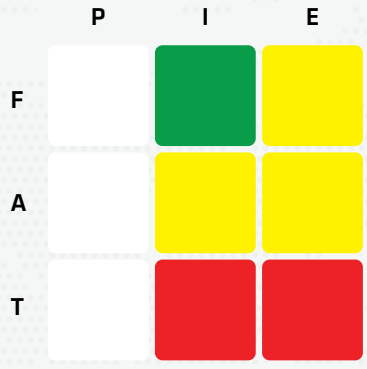

The color of the cells in the matrix are set using these guidelines:

- White: You have no information about this yet. (Constitutes an obvious knowledge gap.)
- Red: You have confirmed important knowledge gaps.
- Yellow: The likelihood of important knowledge gaps is above the accepted threshold.
- Green: The likelihood of important knowledge gaps is below the accepted threshold.

You should map what you have observed in the ongoing incident in one matrix. The gap between this and everything else you know about this particular adversary from older incidents and from threat intelligence sources is mapped in the other.

The first matrix will tell you if you have gaps in your own data and analysis. The second will alert you if there is reason to believe you may be missing an important piece of the puzzle.

Example: You identify outbound command and control traffic from a server, but you are unable to find the malicious binary on the system. Here you have a gap in your own data and analysis. You have also found two different attack tools during your analysis, but from threat intelligence you learn that the suspected adversary group normally uses twelve tools. This tells you there is a chance you are blind to a large part of the ongoing intrusion.

ASSESSING YOUR READINESS

Before deciding to evict, it is also important to map the knowledge you are able to act upon to evict, prevent and detect re-entry attempts after the eviction. This difference can sometimes be unpleasantly large, and you may eventually conclude that you have enough knowledge to start the eviction process, but you lack the ability to execute effectively on it. Should you find yourself in this predicament you will have to establish this ability before you can proceed. In this situation the CTIM becomes your guide for mapping the tools you miss to successfully detect, deny, disrupt, degrade, deceive and finally fully evict your adversary.

A FINAL WORD OF CAUTION

Fighting a resourced, advanced and persistent adversary that has successfully established a foothold within your infrastructure is no easy thing to do. If you are not certain you can manage it alone, do seek help. The complexities in the analysis, the assessments and the actions you need to take are substantial, and it requires a certain level of maturity to even start getting it right.

When I sat out to create the CTIM, my hope was that it could help responders untangle some of those complexities. I can only hope that I have succeeded, and that you will be able to put it to good use.

Good luck fighting the bad guys and good luck determining when it's the right time to make your move.

THE RISE OF RANSOMWARE IN 2016

While the concept of ransomware has been around for 25 years, the crypto-variants – malicious software that encrypts your files and demands you to pay a ransom to unencrypt them – started making more frequent appearances towards the end of 2013. This trend rose well throughout 2015, and in 2016 we saw an explosive growth in successful ransomware infections.

For this study, we looked at the most prevalent ransomware variants observed from our SOC in 2015 and 2016. This includes TorrentLocker, Locky, Cerber, Cryptowall, CryptoLocker, Teslacrypt and CTB Locker.

[All statistics are from real customer cases detected from our Security Operations Center.]

+199%

From 2015 to 2016, we saw a 199% increase in the number of ransomware incidents with our customers.

DISTRIBUTION OF RANSOMWARE VARIANTS

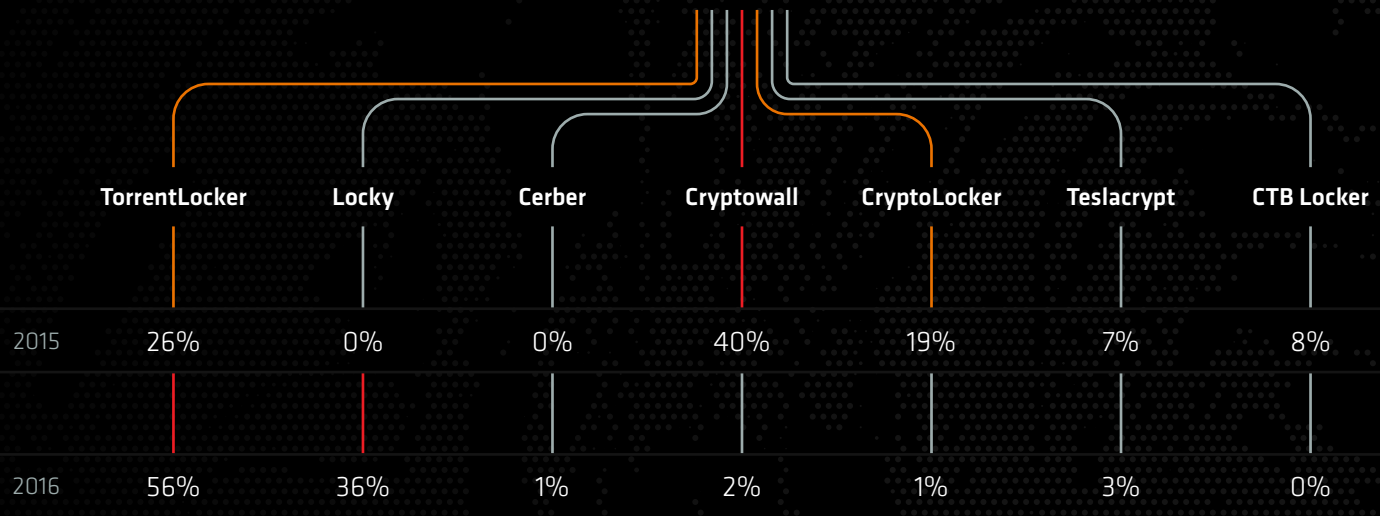

TORRENTLOCKER

TorrentLocker campaign trends in 2016

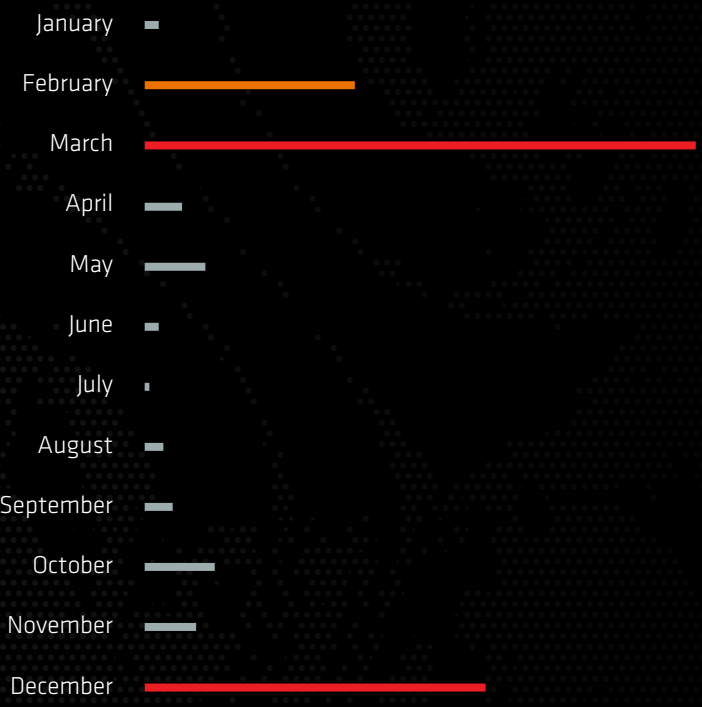

+536%

TorrentLocker was the most prevalent ransomware variant in 2016, leading to a 536% increase in infections from 2015.

2

Two successful TorrentLocker campaigns alone resulted in more incidents than all ransomware combined in 2015. One of these campaigns in March was responsible for 22% of all ransomware incidents we saw in 2016.

LOCKY: THE NEW KID ON THE BLOCK

First detected in February 2016, Locky was consistently distributed throughout the year. Within 8 months, we saw more incidents with Locky than total ransomware incidents in all of 2015.

THE FALL OF CRYPTOLOCKER AND THE COPYCATS

The original CryptoLocker hit its distribution peak in the first months of 2014. However the infrastructure used to distribute CryptoLocker suffered a major disruption as part of Operation Tovar - an international collaborative takedown of the Gameover ZeuS botnet. By the time of the takedown, there was approximately 500,000 victims worldwide who had been infected with CryptoLocker. Of these victims, 1.3% had paid the requested ransom, netting the criminals an estimated $3 million USD in about 9 months.

Since then, several CryptoLocker copycats have emerged – two of which are Cryptowall and a similarly named but unrelated CryptoLocker. Both copycats enjoyed success in 2015. Cryptowall represented 40% of all ransomware incidents we saw, while the copycat CryptoLocker held steady at 19%. However both copycats have all but dropped off the map according to our observations in 2016.

PREVENTING THE INEVITABLE: THE NEED FOR RAPID DETECTION AND RESPONSE

You've fortified your defences. You follow industry best practices. You've purchased the latest and greatest technology. Yet attackers still penetrate your defences. In a world where it's expected that attackers will successfully breach your perimeter, what chance can you possibly have to protect your business?

BJØRNAR PRESTAASEN
Head of Security Operations Center
mnemonic

Attackers will breach your defences. This is not a hypothesis, exaggeration or a fear mongering statement – it is a simple, undisputed fact.

Regardless of the preventative measures we put in place, a determined attacker with the right motivation, financial backing and skillset will evade these measures. In some cases, the attackers do not even need to evade your preventative measures – your employees take care of that task for them.

It should be no surprise that there is a direct correlation between the number of users in an organisation and the number of confirmed security incidents the organisation experiences each year. From our 15 years' experience, we find that for every user an organisation should expect to see 0.2 to 0.3 confirmed security incidents annually. That means that for a company with 1000 users, there are an expected 200 to 300 confirmed incidents each year. This does not speak to the severity of the incidents, but serves as an indicator to the immensity of the task security teams are faced with.

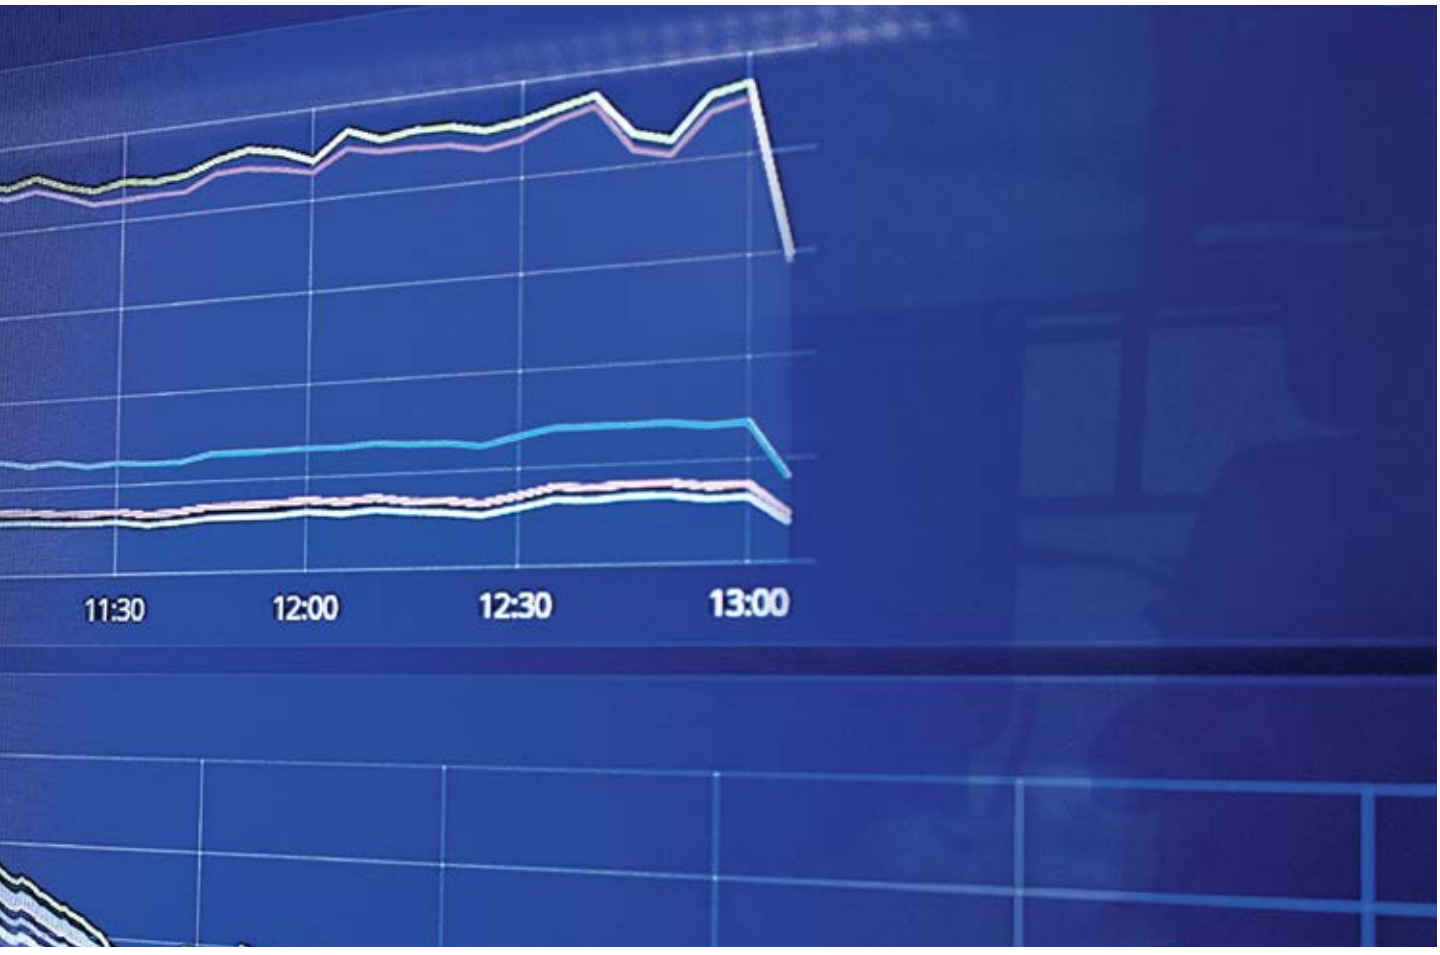

By 2020, 60% of enterprise information security budgets will be allocated for rapid detection and response approaches, which is an increase from less than 30% in 2016.

Special Report: Cybersecurity at the Speed of Digital Business, Gartner

PREVENTION EXAMPLES INCLUDE:

- Vulnerability management
- Password management
- Access control
- Inline security products with blocking capabilities (e.g. firewalls, web/email proxy, anti-virus, endpoint protection)
- User awareness training

"IT risk and security leaders must move from trying to prevent every threat and acknowledge that perfect protection is not achievable. Organizations need to detect and respond to malicious behaviours and incidents, because even the best preventative controls will not prevent all incidents.

Special Report: Cybersecurity at the Speed of Digital Business, Gartner

THE ROLE OF PREVENTION

This is not to say that prevention is not an important component of a well-rounded security strategy – quite the contrary. Prevention is a critical capability that fortifies your cyber defences, and represents best practice for protecting your organisation.

Prevention shapes the attackers path, and makes it more difficult for them to infiltrate a network, move laterally, escalate privileges and steal data. Attackers are only human, and are most likely to pursue the path of least resistance to achieve their goals. If it is dangled in front of their face, attackers will go for the low-hanging fruit. However if you cut the low-hanging fruit, hungry attackers will bring a ladder to reach the higher branches, or a chain-saw to simply cut the tree down.

In opportunistic attacks, where the target is arbitrary, deterring an attacker with enough preventative measures may be enough to cause them to simply move on to another target with lower hanging fruit. Or it might not be. This will vary depending on factors such as the attacker's ambitions, skill level, motivation, and because they are human, their mood.

> **The harder we make life for an attacker, the more likely they are to knock on different doors, generate noise, and trip an alarm.**

On the other hand, a determined attacker with a clear target and goal will relentlessly raise the sophistication of their attack to match the security maturity of their target. However, the harder we make life for an attacker, the more likely they are to knock on different doors, generate noise, and trip an alarm that allows defenders to detect and respond to their presence – provided there is someone listening for the alarm.

IF AN ALARM GOES OFF AND NO ONE IS THERE TO HEAR IT – DOES IT MAKE A SOUND?

There are a wide-range of techniques that can be used to detect suspicious activity in your environment. You can inspect network traffic, install agents on endpoints, collect logs in a SIEM, execute files in a sandbox, monitor user account activity, uncover anomalies in user behaviour, identify spikes in bandwidth usage, amongst a plethora of other techniques. A vast amount of technologies exist supporting these techniques, exclusively designed to detect an attacker's presence.

The challenge is that the security solutions designed to alert you to suspicious and unwanted behaviour do exactly that: they alert you - a lot. In our own security operations center, we see that an organisation with 2,000 employees will, on average, generate over 400,000 security events every day. It is simply not feasible for any organisation, regardless of how well funded they are, to manually assess this volume of alerts.

This challenge is compounded by the fact that the alerts themselves are often low-value and commonly considered unreliable. According to a study by the Ponemon Institute,

> "
>
> **An organisation with 2,000 employees will, on average, generate over 400,000 security events every day.**

81% of malware alerts are considered unreliable by security professionals. It is really no surprise then that only 4% of these alerts are ever investigated.

But let's assume that an alert is investigated – then what? These automatically generated alerts provide marginal value as they have limited information, little-to-no context, and often represent an isolated, technical event from an individual point in a network.

These automated alerts serve the function of bringing the existence of a potential security incident to an operator's attention. It is up to a security team to make sense of these alerts, validate the security incident, assess if they represent a threat to the organisation, and take appropriate action to respond to the potential threat.

Without a plan and capability to respond to a detected threat, you are not much better off than having not known about the threat in the first place.

DETECTION MEANS NOTHING WITHOUT THE APPROPRIATE RESPONSE

The ability to detect malicious and unwanted behaviour in a network only represents half of the battle. Without a plan and capability to respond to a detected threat, you are not much better off than having not known about the threat in the first place.

Responding to a security incident requires that an organisation understands more than just the technical details. An effective response requires that an organisation observes the threat not just as a single malware alarm that has been triggered on some client, but as a security incident that has the potential to (further) impact key business processes if not handled appropriately.

STAGES OF DETECTING AND RESPONDING TO SECURITY INCIDENTS

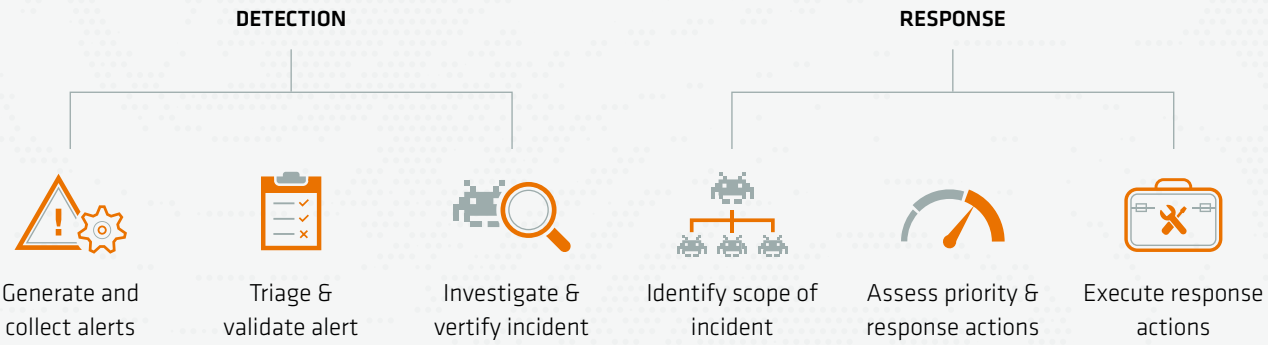

Every organisation will need to go through various stages when detecting and responding to security incidents. These stages are the same, regardless of the threat, industry or technology a company may have.

The detection phase identifies and validates potential threats against your organisation. Much of this phase can be automated and driven by technology, however people will need to be involved at some point to validate the threats.

The response phase is focused on understanding the threat in the context of your business, and taking appropriate actions to remediate the incident. What is the significance of the assets, data and users involved? What services are impacted? How will this affect core business functions? Evaluating the incident in this context enables a response that reflects the severity of the incident based on what it means for your business, rather than solely on the threat itself.

IT'S ALL ABOUT THAT CONTEXT

Context is an important aspect of all decision-making processes. The more information we have surrounding the circumstances of a situation, the more likely we are to make an informed decision on how to proceed. This applies not only to incident response, but every decision we make.

Consider Halloween night as an example. On any other night of the year, if you saw a bloodied person with an axe in their head, you would quickly assume they are severely injured and likely call an ambulance. However the information that it is Halloween night provides context to the situation and will influence how you assess the situation and ultimately, the decision you make. While the decision you make may not change, the extra information adds context and supports a more informed decision to be made.

The same concept applies when responding to security incidents. The more information we understand surrounding the incident, the more informed of a decision we are positioned to make.

Some examples of context that can assist in the decision making process includes:

USER CONTEXT:

Which users are involved? What is their role? Where are they located? What systems and information do they have access to? What influence do they have in the organisation should they be impersonated by an attacker?

INCIDENT CONTEXT:

What systems are affected? Were more than one system involved? Was the threat blocked? Are there alerts from multiple systems? What is the technical scope of the incident?

THREAT CONTEXT:

What type of attack is it? How sophisticated is the attack? Is this a targeted or opportunistic attack? Can the attack be attributed to any individual threat actor? Where is the attack originating from? Have we seen similar attacks in the past? Are other organisations in my industry or region being attacked?

BUSINESS CONTEXT:

Are any of the involved systems or users connected to any critical business processes? What is the potential impact towards these business processes?

How a security incident impacts a business is driven not only by the type of threat, sophistication or attack vector, but the business under attack itself. Capabilities aside, each organisation will have a different set of key business functions and respective priorities for these functions.

For an online retailer, this may mean prioritising that the webstore is available and able to process sales. Meanwhile a law firm may be far less concerned with downtime on their website as they would in protecting their clients' personal and confidential data.

Putting a security incident into the context of a business' core functions enables a response that proportionately reflects the severity of the incident as it relates to how the business is impacted, rather than on the threat itself.

In a world where minutes can be critical,
months are an eternity.

THE IMPORTANCE OF RAPID DETECTION AND RESPONSE

There are an array of regional and global reports that provide insight into how breaches happen, how long attackers go undetected, and how prepared organisations are to respond. One report puts the average time from compromise to detection at around 2.5 months, while another has it at just under 5 months.

The exact figure is less important than the reality of these numbers – we are measuring our detection time in *months*. And bear in mind this is just the time it takes to detect the attacker, not the time it takes to respond. Imagine what you have done in the past 2.5 months, let alone 5 months. In a world where minutes can be critical, months are an eternity.

Despite the statistics in each report varying, the main takeaways are the same:

Takeaway 1: The longer it takes an organisation to detect and respond to a compromise, the more costly it will be.

The direct costs for responding increase as an attacker moves laterally throughout a network and widens the scope of the incident, and increase the chance of direct financial loses as the attacker has more time to execute their actions. Similarly, indirect costs rise as the probability of intellectual property theft (and in all likelihood, the magnitude of this theft) increases the longer an attacker has to operate unimpeded.

Takeaway 2: Security incidents that are discovered by an external party take far longer to respond to than those detected internally.

It makes sense that incidents detected by external parties (e.g. law enforcement, regulatory bodies, customers, etc.) will take longer to respond to simply from an operational perspective to alert the right person in the organisation with the necessary information. External detection is also commonly based on the magnitude or residual effect of a compromise rather than the compromise itself.

However this expected delay does not account for the fact that organisations take 7 to 28 times longer to respond to externally notified incidents than internally notified incidents (internal includes those detected by Managed Security Service Providers). One plausible explanation is that organisations without the capability to detect incidents internally are also less likely to have the plans, processes and routines required to rapidly respond to security incidents in general.

Takeaway 3: The probability for data being stolen rapidly increases when an attacker's presence in an environment moves from days to weeks.

This should not come as any surprise. The longer an attacker has in your network, the more likely they are to succeed in their goals.

Takeaway 4: The longer an attacker is present in a network, the lower an organisation's confidence becomes in their understanding of the full scope of the incident and ability to completely extinguish all effects of the breach.

Time is an invaluable resource for your adversaries. The more time an attacker has in your network, the more complex the incident is likely to become, and the more difficult it will be to understand the full scope of the incident. This leaves an organisation with avoidable gaps in their knowledge, forces them to make less-informed decisions, and a general sense of uncertainty.

Takeaway 5: Technology alone is not enough.

Incident response is a decision-making process. Technology is a tool to support this process, however it is exactly that – a tool. Assessing a security incident, understanding the potential impact it has on a business, and determining how to respond involve a series of complex decisions that only people can make. If you are an online retailer, and you suspect your payment systems have been compromised the day before Black Friday, technology cannot decide whether to take your website offline during the biggest shopping day of the year - only people can make that type of business decision.

OPERATIONALISING INCIDENT RESPONSE

Rapid detection and response operationalises incident response. Handling security incidents become a part of the daily routine and an organisation enters a state of continuous response. Security incidents are no longer seen as an anomaly but an expected occurrence on any given day. From an operational viewpoint, the handling of security incidents is seen no differently than traditional operational activities, such as server maintenance or user management.

The effect of operationalising incident response is that the efficiency of the process dramatically increases. These efficiencies mean that an organisation moves from a purely reactive state of defence to a (near) real-time level of incident response. The resources required to facilitate the response decrease, the elapsed time between initial compromise and

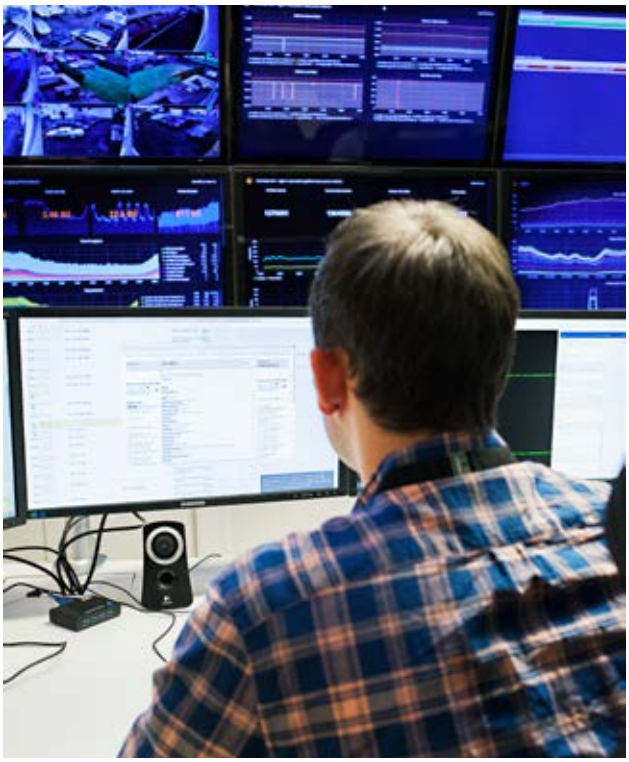

recovery shortens, and the process as a whole becomes predictable and measurable.

In the end, this translates to cost savings, an ability to confidently report on the state of cybersecurity and reduces the potential impact a security incident may have on the organisation.

A FINAL NOTE

Cyberattacks will happen. Some of them will be blocked, others will be successful. Some will be caused by employees, and others will elude even the most fortified defences. While it is not possible to prevent all attacks, it is not realistic or feasible to completely abandon prevention and rely solely on detection and response.

It is not a choice between either prevention or detection and response, but how to best use them in combination.

PERSONAL DATA COMPLIANCE MANAGEMENT UNDER THE NEW GENERAL DATA PROTECTION REGULATION

NILS KRISTIAN EINSTABLAND
Partner/Attorney at law
Advokatfirmaet Selmer

KRISTINE JUELL JOHNSEN
Associate Lawyer
Advokatfirmaet Selmer

Approved by EU Parliament in April 2016, the EU General Data Protection Regulation (GDPR) aims to unify data privacy laws across the EU and EEA, and protect data privacy of European citizens. With enforcement commencing in May 2018, what does this mean for you and your business?

Personal data is business critical information in line with financial data, trade secrets, intellectual property rights (IPR) and technical and strategic information – and is (or should) as such be included in an organisation's information security policy, which again is a part of the enterprise's general compliance strategy. "Compliance" in this context is the enterprise's strategies and guidelines for mitigating business risks and meeting investors and counterparties´ expectations, as well as acting in accordance with regulatory requirements, proper corporate governance, and ethical and integrity guidelines.

The generally increased focus on compliance is an expression of the complex environment in which today's business life operates, with increased concerns related to political risks, corruption and bribery, money laundering, cybercrime and personal data issues, etc.

Protection of personal data is for many organisations particularly challenging due to a comprehensive and complex regulatory framework. It is immensely difficult to derive from the words of the law what it actually means for your business in terms of duties and obligations.

The data protection legislation regulates both what to do to minimise the risk for a breach of information security, and what to do when security incidents inevitably occur. For the arrangement to be effective, it must be properly anchored to senior management and properly implemented so that it becomes part of the organisation's DNA, not only a "paper compliance". As with all compliance issues, it should be integrated into business operations based on an updated risk analysis.

A NEW REGIME TO ADDRESS NEW SECURITY CHALLENGES

The EU's new General Data Protection Regulation (popularly known as the GDPR) is a response to this new, complex environment.

The GDPR does not entail any major changes with regard to the basic principles for the processing of personal data. The categories of data considered as personal data and sensitive personal data mainly remain the same, and all processing of personal data must – as is the rule today – be in accordance with the general principles for processing. An area in which the GDPR however represents some changes are with regard to the security regime required for data protection compliance.

WHAT IS PERSONAL DATA?

Personal data means "any information relating to an identified or identifiable natural person". The GDPR states that an identifiable natural person is one who can be identified, directly or indirectly, in particular by reference to an identifier such as name, identification number, location data, online identifiers or to one or more factors specific to the physical, physiological, genetic, mental, economic, cultural or social identity of that natural person.

Thus, completely anonymized data is not considered as personal data. However, the data controller must be confident that the data can in fact not be connected to any individual, now or in future (within reason). Further, note that the actual anonymization of personal data is to be considered as processing, requiring legal basis, etc.

It is immensely difficult to derive from the words of the law what it actually means for your business in terms of duties and obligations.

THE GDPR AT A GLANCE

Entered into force on 24 May 2016, and will apply in the EU from 25 May 2018.

The GDPR establishes a two-tiered system of administrative fines, giving basis for imposing administrative fines of up to EUR 20 000 000 or 4% of the total worldwide annual turnover of the preceding financial year. In comparison, the current data protection regime in Norway places the fine limit at approximately EUR 100 000.

The new rules are given in the form of a regulation, meaning that they become immediately enforceable in all EU member states simultaneously, and that all member states must comply with the regulation directly. As Norway is not member to the EU, the GDPR will have to be implemented in Norwegian law prior to being given effect. According to the Norwegian Data Protection Authority we can expect the new Data Protection Act to come into effect in or around May 2018.

MAIN CHANGES UNDER THE GDPR

- Abolishes notification requirements to the Data Protection Authorities, but requires that the data controllers and processors keep relatively detailed records of all processing activities.
- Gives the data subjects new substantial rights, including the right to be forgotten, data portability rights and the right to object to certain processing activities.
- Applies to both data controllers and data processors, entailing that data processors will be required to comply with data protection legislation directly and not just via their contractual obligations to data controllers.
- Requires organisations to adopt significant new technical and organisational measures to demonstrate their GDPR compliance, including by adopting certain "data protection by design and default" measures, staff training programs and undertaking audits.
- Introduces a security breach communication framework for all data controllers regardless of the sector in which they operate.

PREPARING FOR THE GDPR

In preparation for the GDPR, the first step should be to identify and address potential gaps in current data protection compliance, in order to identify any gaps that does not necessarily impose a high risk now, but may lead to high exposure risks when the GDPR is implemented. For instance, all businesses should perform risk assessments of all current data processing to identify processing that might require the performance of a detailed privacy impact assessment (DPIA), and in extension, an evaluation by the data protection authorities.

Further, to secure compliance with the new breach notification-regime, data controllers should also develop and implement plans for handling of data breaches. This is particularly important to secure compliance with the strict new notification regime, which requires notification to the data protection authority to be effected within 72 hours of a suspected breach. Both data controllers and processors should also establish routines for documenting all processing of personal data, as this is a requirement under the GDPR.

As a second step, all businesses should evaluate current personal data compliance against the new or more stringent requirements introduced by the GDPR on a more general level, in order to establish which updates of systems and procedures that will be necessary prior to implementation in 2018.

Data controller

A data controller is defined as the natural or legal person, public authority, etc. which, alone or jointly with others, *determines* the *purposes* and *means* of the processing of personal data.

VS

Data processor

A data processor refers to a natural or legal person, public authority, etc. which processes personal data *on behalf of* the controller.

The key distinction is the degree of independence the parties have in determining how and in what manner the data is processed, as well as the degree of control over the content of the personal data processed.

DIRECTIVE ON SECURITY OF NETWORK AND INFORMATION SYSTEMS (THE NIS DIRECTIVE)

The Directive on security of network and information systems (the NIS Directive) will enter into force in May 2018. Norwegian Authorities has signalled that the directive has EEA-relevance.

The NIS Directive requires EU Member States to ensure a certain level of cybersecurity by (i) creating a strategy for safety, (ii) establish a cybersecurity readiness unit (CSIRT) and (iii) to impose upon certain operators of essential services (such as energy, transport, water, banking, financial market infrastructures, healthcare and digital infrastructure) certain cybersecurity requirements, including implementing appropriate and proportionate security measures and an obligation to notify serious incidents to the relevant national authority. Also key digital service providers (search engines, cloud computing services and online marketplaces) will have to comply with the security and notification requirements under the NIS Directive.

COMPARING THE GDPR AND NIS DIRECTIVE

There is significant overlap between the GDPR and the NIS Directive. Both set of rules impose an obligation to implement risk-based security measures and both include notification requirements in the event of an incident. However, they also protect distinct interests and may apply to different types of incidents.

For instance, the rules have different aims. Where the aim of the GDPR is to safeguard personal data, the NIS Directive focus on network security. Consequently, the GDPR requires adoption of measures to secure personal data, while the NIS Directive focuses on network integrity in order to protect the provision of the relevant service. These aims will often overlap, such that in many situations network integrity will be an appropriate prerequisite for personal data protection. However, there might be areas of divergent protection, for example encryption, which might help secure compliance with the GDPR, but will not protect the network in accordance with the NIS Directive.

Further, the GDPR requires breach notification only where a security breach has compromised personal data, while the NIS Directive imposes notification obligations in case of significant disruption to the provision of the relevant service.

Thus, the GDPR protects and requires notification in case specific information – personal data - is compromised, while the NIS Directive targets operators that fail to implement appropriate security measures or fail to notify the competent authorities of an incident in general.

The NIS Directive states that processing of personal data under the directive shall be in accordance with the data protection rules. However, the Directive does not, for example, regulate a situation of conflicting interests, additional liability in case of non-compliance with both set of rules, and whether compliance with one set of rules can be used as a defence against non-compliance under the other set of rules.

NEW CONCEPTS UNDER THE GDPR

The GDPR introduces new technical and organisational measures to demonstrate compliance, including by adopting certain "data protection by design and default" measures, as well as stricter requirements related to risk and impact assessments. Unlike the current regime, the GPDR also places express obligations on the data processor, meaning that many data processors must have an increased focus on the new requirements.

Data protection by design and default requires that businesses must have privacy in mind when designing new projects, processes, products or systems. Privacy by design implies that data controllers must consider data protection at the initial design stages of a project, and not only as an after-thought. This requires compliant policies and procedures, as well as compliant systems, at the outset of any product or process development.

Privacy by default requires that the data controller establishes measures for ensuring that, by default, only personal data which are necessary for each specific purpose of the processing are in fact processed, making "data minimisation" a default setting for processing of personal data. Thus, the strictest privacy settings shall automatically apply when, for example, a customer acquires a new product or service, without any manual change to the privacy settings being required on the part of the user. The "default obligation" applies to both the amount of data collected, to the extent of the processing, the period of storage and the data's accessibility.

An example according to the GDPR of a measure that may help satisfy the requirement of appropriate technical and organisational security measures is the concept of "pseudonymisation". The concept involves processing of personal data in a way that the data can no longer be attributed to a specific person without the use of some other information that is kept separately and is subject to sufficient security measures.

Data protection by design and default shall be implemented and continuously followed-up at all stages of processing – requiring businesses to make necessary improvements along the way, unless the cost/"privacy gain" are in-proportional.

In addition to the risk assessment also known under the current regime (by data controllers at least), the GDPR introduces a detailed privacy impact assessment (DPIA) concept. If the risk assessment shows that the processing might entail a large risk, the controller/processor shall perform and document a DPIA, addressing the impact of the planned processing operations on the protection of personal data, and subsequently seek the data protection authorities' opinion on the proposed measures.

The GDPR introduces an obligation to notify both the data protection authority within 72 hours and affected data subjects if the security incident imposes a risk to a data subject's rights.

WHEN A BREACH OCCURS

When a security incident occurs, your main concern should be to correct it and to implement measures to prevent it from happening again. In addition, the GDPR introduces an obligation to notify both the data protection authority within 72 hours and affected data subjects if the security incident imposes a risk to a data subject's rights.

To manage the situation properly you must have an emergency team ready, escalation procedures and communication plans in place to facilitate damage control from the start, and to get right messages out, both internally and externally.

To investigate the incident and to implement the appropriate preventive measures, you may have to team up with a technical and legal forensic team. If actions shall be taken against those behind the security incident, the investigation must be appropriately conducted to substantiate legal positions and avoid damaging evidences. This is particularly important when the investigation involves your own employees.

Proper handling of security incidents is particularly important when the GDPR is implemented. The new enforcement regime presented in the GDPR can surely make even the toughest business man (or woman) shiver. The GDPR introduces a two-tiered system of administrative fines, giving basis for imposing administrative fines of up to 20 000 000 EUR or 4% of the total worldwide annual turnover of the preceding financial year. This does certainly make the risks related to non-compliance much more severe than under the current regime.

However, it is a misconception that administrative fines will be the result in every case of non-compliance with the GDPR. Administrative fines are just one of the data protection authorities' tools. They may also resort to other administrative sanctions, such as issuing warnings or reprimands, as well as different types of orders (e.g. ordering the rectification or erasure of personal data), and temporary or definitive bans of processing, including suspending overseas data flows.

The purpose of the GDPR is to harmonise the method for determining sanctions, reducing differences with regard to sanction types, levels and frequency across the European Economic Area. How this will turn out in practice remains to be seen.

Netherlands

ASML

Ewoud Smit

Manager Cyber Defence Operations

The opinions expressed are the respondent's own and do not necessarily reflect the views of their employer.

ASML is one of the world's leading manufacturers of chip-making equipment. ASML's guiding principle is continuing Moore's Law towards ever smaller, cheaper, more powerful and energy-efficient semiconductors. ASML is a multinational company with over 70 locations in 16 countries, headquartered in Veldhoven, the Netherlands and employing more than 16,000 people.

WHAT IS YOUR BIGGEST SECURITY CONCERN?

I see a trend of criminal and malicious actors increasing their ability to cause harm through specialisation, cooperation and maturing of an underground marketplace. At the same time I see law makers putting in legislations based on symptoms and incidents without addressing root cause drivers. These laws still are mostly aimed at creating incentive for prevention while the security world has learned detection and response should get more focus. Currently the legal landscape is still favouring the attacker who acts with impunity and burdens the defender, who must expend resources on compliance that could be better spent on true security.

IN WHAT AREAS OF SECURITY DO YOU THINK WE'RE FALLING BEHIND?

I believe as a society we are not able to put in structures to contain and control criminal behaviour, let alone state sponsored malicious activity. Developments in technology and digital information continue to outstrip our collective ability to ensure their use for the collective greater good outweighs the abuse. This is largely due to a lack of principles based law making and international alignment. I hope to see more international collaboration and alignment in law making for the next years.

WHAT GIVES YOU HOPE FOR THE FUTURE OF SECURITY?

I see a lot of operational security intelligence sharing initiatives taking hold ensuring lessons learned on the defence side are being shared and reduce the costly mistakes of first time responders. In the wake of large corporations bolstering their security, the sharing initiatives and shared SOC/CSIRT services being offered at reasonable commercial prices make it possible for small and medium enterprises to enhance their security cost efficiently. I am hopeful this will in turn increase the resiliency of the entire eco-system against malicious adversaries.

For more information about mnemonic, visit www.mnemonic.no

mnemonic AS
Wergelandsveien 25
0167 Oslo
Norway
+47 2320 4700
contact@mnemonic.no
www.mnemonic.no

Publication designed by Bjørnar Løvtangen, Make Noise AS

Photo credits page 14/15, 24, 28, 30, 32/33, 38, 46, 51, 58/59, 61, 65, 66:
Charlotte Sverdrup Photography

The views and opinions expressed in this report are those of the authors and do not necessarily reflect the views of their respective employers.