mnemonic
SECURING YOUR BUSINESS

SECURITY 2018 REPORT

...2018

SECURITY COOPERATION WITHOUT BORDERS

In mnemonic, we have long preached the value and need for cooperation between security companies, governmental organisations, law enforcement and academia.

Throughout 2017 we saw the security community come together like never before. A range of sector and industry dividing lines were crossed (if not trampled) on national, European and global levels.

Last year, mnemonic was pleased to join several initiatives organised both at home and internationally. I would like to take this opportunity to highlight a few of the great initiatives out there.

In **Norway** there were several positive initiatives stemming from both the private and public sector. One in particular is the Norwegian National Security Authority's (NSM) quality scheme for Incident Response. The initiative intends to help organisations and companies find security providers that meet strict requirements for incident response. By pre-vetting incident response teams, organisations have access to a list of trusted providers who can assist them when responding to incidents – a valuable service where time is often of the essence. We are proud to add this to our growing list of incident response accreditations and further strengthen our collaboration with NSM.

All over **Europe** there are initiatives leveraging security organisations' collective knowledge and experience. One of them is the STOP-IT project (Strategic, Tactical, Operational Protection of water Infrastructure against cyber-physical Threats). By gathering major water utility providers, industrial technology developers, high tech SMEs and top European R&D, the collaborative project aims to find solutions to protect critical water infrastructure throughout Europe.

Last year, mnemonic also joined collaborative initiatives on a **global** scale. One such project is No More Ransom - a global initiative between law enforcement and security companies that aims to disrupt cybercriminal businesses with ransomware connections, and help victims of ransomware retrieve their encrypted data without having to pay the criminals.

mnemonic's list of international threat intelligence partners now includes over 200 collaborators. Together, security providers from all over the world are looking beyond competing goals, and arming the community with the ability to collectively see the bigger picture.

All of these security initiatives are examples of projects driving security innovation and solving real world issues. They are proof of the positive effects of cooperation within the security community.

Are you interested in joining the collective fight against cybercrime? Send an email to JoinTheTeam@mnemonic.no and we will be in touch to explore how to best unify our efforts.

TØNNES INGEBRIGTSEN
CEO, mnemonic

ARTICLE

INTERVIEWS | STATISTICS

ARTICLES

ARTICLES

Security Predictions 2018

JON RØGEBERG
Head of Threat Intelligence
mnemonic

Well here we are again. Our crystal ball has been freshly polished, we've observed the stars and our hallucinogens have finally worn off. We're now ready to offer our security predictions for 2018. We predicted 2017 would be a bumpy ride, which turned out to be the understatement of the year. Between The Shadow Brokers leak, WannaCry, and Equifax, 2017 wasn't a bumpy ride – it was like off-roading in a 3-wheeled Ford Model T.

The art of prediction is one of managing expectations. Let's take the TV weatherperson as an example. Their job is straightforward enough - predict the weather. Now as we all know (and have likely experienced first-hand getting soaked by rain on a 'sunny' day), TV weatherpersons are notorious for being inaccurate in their predictions. But they are masters of expectation management. From years of repetitively being wrong, when they by chance accurately predict the weather, they receive incredible praise for simply doing their job. Like I said, it's an art form and they have mastered it beautifully.

Our predictions come in a different form. We do not deal in absolutes or the definitive. It would be naïve for us to think we can predict the future (despite our predictions going 5 for 5 last year, but who's counting). The expectations we set are that we share insights and reflections that, we hope, will add to your cybersecurity situational awareness and help you prepare for what may be ahead of us.

The reality is we, or anyone, simply don't know what the future holds, but that doesn't mean we can't prepare for it. Who would have predicted that in 2017, the NSA would have a collection of their own hacking tools leaked, which North Korea would then use to launch a global ransomware attack? No amount of hallucinogens would lead to that kind of prediction.

2017, you certainly didn't disappoint. Now let's have a look at 2018.

THE LONG OVERDUE DEATH OF SINGLE-FACTOR PASSWORDS FOR AUTHENTICATION

Passwords are the bane of our modern world. We use them tens, if not hundreds of times a day. Unlocking your phone? Put in your PIN. Logging onto your laptop? Let's have that password. Want to order a pizza online? It looks like you've ordered from us before, so we'll be having that password please. It's estimated that the average person has upwards of 30 online identities to manage, whereas they cycle a pool of 6 passwords for all of these identities. Looking at it another way, almost 75% of accounts are guarded by duplicate passwords.

Passwords alone are not enough. This has been the mantra of the security community since the turn of the century. We re-use usernames, we re-use passwords, and as people we're fallible to phishing scams that steal this information. Combine this with even more services and data being pushed to the cloud that enables would-be cybercriminals to conveniently access our data with 99.99999% availability, and we are presented with a recipe for disaster.

As society becomes more aware of their data, their rights, and the consequence of their data being leaked, there is a slow but distinct shift towards the general population acknowledging the necessity for security controls, even if it means the slightest of inconveniences in their user experience. Lucky for us all however, there are some technological advances that are pushing us in the right direction.

The proliferation of smartphones, complete with facial recognition, fingerprint scanners, voice recognition and others, puts biometric authentication in most people's pockets, along with traditional PIN authentication. This certainly helps with the adoption of multi-factor authentication. Rather than having to carry multiple tokens around that generate one-time passwords, users can perform a simple task like scanning their fingerprint to serve as a secondary authentication method. That seems like a reasonable 'inconvenience' to protect one's entire online identity.

Multi-factor authentication can also be in the form of risk-based authentication, which is increasing in both popularity and effectiveness. Also known as adaptive authentication, users may be challenged with an additional level of authentication based upon various risk-driven scenarios in either their behaviour or the task they're performing. Is the user logging in from a new location or from a new device? Is their behaviour on the login page drastically different from previous visits or that of other users? Does the activity they are performing, such as transferring money, warrant an additional level of authentication? Adaptive authentication attempts to identify and prevent abnormal behaviour while simultaneously improving the user experience by dynamically reducing authentication requirements.

As biometric authentication becomes more adopted in mobile phones, and users become more aware of the consequence of their online identity being compromised, we expect to see the password as a sole method of authentication slowly dwindle away and die a slow, long overdue death. Will this happen in 2018? Unlikely, but the writing is on the wall, and we do expect to see more users generally accepting the need for stronger authentication mechanisms, and demanding these from their service providers.

Ransomware, malware, and DDoS attacks are amongst the services available to enable anyone to conduct cyberattacks.

ENABLING THE LESS SKILLED THREAT ACTOR

Complex techniques trickling down in a usable form to the laymen has occurred throughout history. It's a natural occurrence through the industrialisation of any new technological advancement. Look at production. Only a decade ago, the creation of virtually any moulded object was restricted to specialised manufacturing plants. Today, the advances in additive manufacturing, commonly known as 3D printing, enables anyone to print custom objects right at home. Wordpress and other publishing platforms enable individuals with no technological background to create a website in minutes.

The same trickle down effect is happening with attackers as well.

This is partially driven by the commoditisation of cyberattacks. Where everything is available as a service these days, cybercrime is no exception. Ransomware, malware, and DDoS attacks are amongst the services available to enable anyone to conduct cyberattacks.

The Shadow Brokers leak is also an example of the devastating, yet quite predictable aftermath of the public release of sophisticated attack tools. Leaked to the general public throughout 2016 and 2017, The Shadow Brokers released a collection of exploits and tools allegedly stolen from The Equation Group, the code name for Tailored Access Operations (TAO) - the cyber intelligence gathering unit of the United States National Security Agency (NSA). Within the first weeks after the leak was posted on April 14th, 2017, more than 200 000 machines globally were infected with tools from the leak. It was also these tools that were used in the WannaCry, NotPetya, and Bad Rabbit attacks that dominated headlines through the second half of 2017.

In late 2016, the Mirai botnet was infecting internet-connected devices such as security cameras, printers and home routers and enrolling them into a botnet measured in the hundreds of thousands. The purpose was to build an army of unsecure IoT devices to be used to launch DDoS attacks, and it did so quite successfully. When the source code was released, attackers were quick to use the code as it was, and further develop the malware. This led to the botnet being 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 over the course of a week.

Such leaks provide insight into the advanced tactics, techniques, and procedures from nation states and other sophisticated threat actors. When these tools are packaged, with functional exploits and come complete with operative guides, we cannot be surprised that less sophisticated threat actors will jump at the chance to leverage these tools.

Nor can we be surprised that opposing nation states, whom themselves are advanced in their own right, will use tools leaked from their adversaries in their own attacks. This was observed months after The Shadow Brokers' leak, with North Korea using the NSA's tools to launch WannaCry, and Russia's subsequent NotPetya attack.

This is the natural evolution as cyberwarfare becomes commoditised. The inventors, the innovators, the boundary pushers will continue to stay in front, all the while enabling the capabilities of those with less knowhow, less skills and less resources, but just as much motivation to join the race. Keep your eyes peeled in 2018 for increased capabilities from less sophisticated threat actors, including a growing presence of industrialised nations as 'new' entrants to the offensive cybersecurity game.

SUPPLY CHAIN ATTACKS PROVE WE'RE NO STRONGER THAN OUR WEAKEST LINK

While not a new concept, recent years have seen supply chain attacks becoming more publicised. These attacks are predominantly found in advanced, targeted attacks, and often with effective and devastating results.

The premise of supply chain attacks is fairly straight-forward. Rather than attempting to thwart your target's defences by directly breaching their fortified perimeter – the avenue your target is most likely expecting you to take – supply chain attacks focus on a target's trusted third party suppliers. So instead of attempting to breach the front walls of the castle, an attacker will compromise the local wheat farmer that makes daily deliveries to the castle, and is trusted to make their way through the side entrance without so much as a second glance from the guards. While the castle has mature, well-developed and properly funded defences, the wheat farmer does not, and is a prime target.

> Supply chain attacks are notoriously difficult, if not impossible, to prevent.

This was the case in the summer of 2017 with NotPetya, where the software update mechanism in M.E.Doc, a popular accounting software in the Ukraine, was compromised, weaponised, and used to deliver the pseudo-ransomware to unsuspecting victims. The United States CIA has since attributed the attack to Russia's foreign military intelligence agency, GRU, and it is widely accepted as a politically motivated attack by Russia against the Ukraine. The collateral damage of the attack impacts a list of global organisations. Most notably was Maersk Line, which has documented a loss of USD $300 Million as a result of the attack.

A supply chain attack was also used in the 2013 attack against Target. Credentials of an HVAC supplier were compromised, which led to the breach of Target's payment systems and 41 million customer payment card accounts being stolen.

In 2017, a version of CCleaner, the popular tool used to optimise the performance of your computer, tablet or mobile, was compromised and used to infect more than 2.2 million Windows machines. While appearing initially to be an attempt to arbitrarily infect as many clients as possible, it was later discovered that this was a highly targeted attack against global technology companies such as Samsung, Intel, HTC, Sony, Google, amongst others. The 2.2 million infections were filtered for machines that belonged to the targeted list, of which 40 victims from global technology companies received a second-stage payload that enabled a persistent presence on the devices.

Supply chain attacks are notoriously difficult, if not impossible, to prevent. Despite an increase in organisations evaluating the security and risk profile of their suppliers, at some point there is an inherent trust that must exist with suppliers. Whether they are maintaining one of your systems, your donut supplier or a tech giant like Google, there is a level of necessary trust that suppliers are taking reasonable steps to protect themselves and by proxy, their customers. This can (and should be) audited, regulated and enforced, but the truth remains that there is a level of risk that must be accepted in order to do business in the modern era.

If anything, the recent publicity of supply chain attacks is encouraging organisations to be more aware of the information they are making available and sharing with their suppliers, who their suppliers actually are, and evaluating how much trust they should be given. This is a positive trend that we hope continues. However, while organisations are increasingly focused on the risk of their suppliers being compromised, we fear they will continue to overlook that they as a supplier themselves will be targeted and used to attack their customers.

EXTORTION WITH GDPR

For our European readers, or those conducting any business in Europe, GDPR (General Data Protection Regulation) is a hot topic that will only increase in temperature throughout 2018. Perhaps the most talked about concept that GDPR introduces – at least in corporate boardrooms and in solution vendors' marketing material – are the potential fines and penalties for breach of the regulations. Namely, this is the maximum fine of 4% of annual global turnover or €20 Million, whichever is greater. Seems like a nice payday for regulators and cyber-criminals alike, but how?

Cybercriminals are canny, ruthless and abide by their own moral code. These criminals are not above holding your personal data hostage, and demanding a ransom be paid to the attacker if you ever want to see your data again. This is the very premise of ransomware – a word that unfortunately has such a prominent presence in our society that it was officially added to the Merriam-Webster dictionary in September 2017 (along with 'Internet of Things', 'troll', and the unrelated but delicious 'froyo').

Likewise, cybercriminals are not above other methods of extortion to gain a buck. Doxing is the act of publishing private information about someone onto the public Internet as a form of punishment, and has been leveraged for extortion in recent years (and also a word added to Merriam-Webster in April 2016, along with 'Bitcoin', 'wacky tobacky', and 'nomophobia' – the fear of being without access to a working mobile phone).

DDoS (Distributed Denial of Service) attacks are used as an extortion method: pay our demands, or we will regularly attack your services and make them frustratingly slow or unavailable for your customers. There are also well documented (and far more undocumented) cases of criminals stealing private company data and threatening to sell it, or simply release it publicly unless a ransom is paid. The list goes on, and this is nothing new.

A challenge with these types of extortion threats is that it is difficult to put a price on their value. What is the perceived cost to the organisation if the personal data of their customers is leaked? This is difficult for both organisations and cyber-criminals to put a price on – until now. With the regulatory fines that GDPR enforces, organisations and cybercriminals now have a common perceived value of what a considerable personal data leak is worth. Granted the GDPR fines represent the upper limit, it nonetheless puts a value that is no longer arbitrary.

In 2018, we expect to see cybercriminals threaten to release personal data that an organisation is responsible for, and use GDPR fines as their bargaining chip to both convince organisations to pay, and increase the amount they're paid.

WHAT IS THE (REAL) VALUE OF INFORMATION SECURITY?

MARK TOTTON
Senior Consultant
mnemonic

After reading this article, you will:

- Know what questions need to be answered in order to properly protect your organisation's information
- Understand who in your organisation needs to play an active part in information security
- Appreciate how GDPR helps us kill two birds with one stone

I would like to begin with some general observations.

I started in the IT business in 1974 as a computer operator. There were no passwords to the machines, you just pressed the "Interrupt" button and a prompt came up on the teletype console. Data transmission went over leased lines from point to point, with no encryption, sometimes using acoustic modems. Development, testing and production ran on the same machines. The list goes on.

This does not mean there was no risk; it just means we had yet to understand the risks.

Today, we have passwords, encryption, separate environments for development and production, segmented networks, monitoring, log analysis and more. So why, after 43 years, do we still have data breaches caused by poor configuration and human error, not to mention criminality and espionage?

WHAT IS INFORMATION SECURITY?

There are many information security companies out there. They have exciting, cutting (if not bleeding) edge solutions using learning AI, advanced algorithms, boxes that are smarter than we are, and employing whatever buzzwords containing "cyber-" we can think of. Each of these solutions will obviously solve any and all of the security concerns we have, whatever they may be. From DDOS attacks to APTs to phishing, e-phishing, spear phishing, trout fishing, deep sea fishing, and so on, our prayers have been answered. Even better, IT and security teams can sort it all out while the rest of us get on with running the organisation. Hallelujah!

Maybe we should take a short pause before we buy more boxes, and ask ourselves *what is information security anyway and why do we need it?*

Information security is simply our response to a threat to our information. We have no doubt all read the frightening news articles about organisations losing control of their email servers, or customer information, or massive data breaches measured in the tens and hundreds of millions of stolen personal records (see: Equifax, Yahoo!, eBay, Target, JP Morgan Chase, The Home Depot, etc.). Therefore - we have to buy stuff to protect ourselves, right? Clearly, the fancier the solution, the better we control our risks, and *then* we can let IT and security take care of it while we get on with running the organisation?

It is only when a risk is felt to be real and relevant that we address it.

Well...no. You see, information security has no value in itself. It makes no improvement to your bottom line; it does not save you money and does not sell more products or services. This, by the way, is one of the main reasons many businesses consider information security "a necessary evil" (and not that necessary either).

WHO IDENTIFIES RISKS?

At its core, risk is a straightforward concept – roughly it can be defined as the consequences of something bad happening, multiplied by the likelihood of that something happening. We could add some threat agents and other factors, but that is the basic idea. Information security has only one purpose - to reduce the likelihood that a perceived risk will occur.

What does that mean? It means that if you do not feel there is a risk, you have no reason to do anything about it. It is only when a risk is felt to be real and relevant that we address it. So who identifies the risks? Who is responsible for deciding whether a risk is acceptable, or whether it requires action?

All too often this task falls to IT and security teams. Do you believe your IT department or your security group have the necessary understanding to decide which business risks you face and how best to address them? Business risk is the key-word here. In my experience, IT people join companies in order to work with IT, and security people to work with security. It is usually of far less importance to them whether the company deals in shipping, banking, or chicken as long as the IT/security environment is interesting.

The only people with the understanding and knowledge necessary to evaluate business risk are leaders. They are also the group who are ultimately responsible for identifying and handling risk, and likewise those who suffer the consequence when a risk is realised. The failure to fulfil this responsibility is a strong contributing factor to the vulnerabilities that lead to the various breaches we have read about recently. Clearly, one cannot easily avoid all breaches; state sponsored hackers are hard to stop, but statistics show that the largest cause of incidents is employee error.

Then there is the third option, risk managers and auditors. They are usually very good at identifying the probability of something bad happening. They use threat scenarios, they do testing, they interview and they read documents. And at the end of this, they produce a report giving a clear view of all the vulnerabilities in whatever they are assessing; what those vulnerabilities can lead to; how likely it is that a threat agent will exploit those vulnerabilities and recommendations for reducing the likelihood of those vulnerabilities being exploited.

There is another step before we can really apply information security well; we need to know how much risk is too much. An organisation's risk appetite expresses its willingness to accept risk in order to achieve its goals. How much money can we accept losing, how important is our reputation? By asking and answering these questions and others, we can identify acceptable levels of risk, and ensure that we address all risks over this level with suitable security measures.

WHAT ARE YOU PROTECTING?

Answering the questions above gives us a lot of valuable information, but there is a major factor missing – what are you protecting and why? There is a clue in the title of this article - we want to protect information. All organisations, private or public, process, store and transport information. Information is essential to all organisations, whether a global manufacturing company, or a service department in a county council. Without access to information everything stops.

Many companies are getting into (or are well into) panic mode over the recent spate of information thefts, and even more frightening, the impending GDPR. Whilst we will all need to do more work to comply with GDPR, it makes sense to work effectively and find out what information we need to protect and where we store, transport and process it. This will help both with information security and with GDPR compliance.

When we know what the information we need to protect is, we can classify it. No, not in terms of Top Secret (though that may come into it), but in terms of its value to the organisation and the potential consequences of breaches to confidentiality,

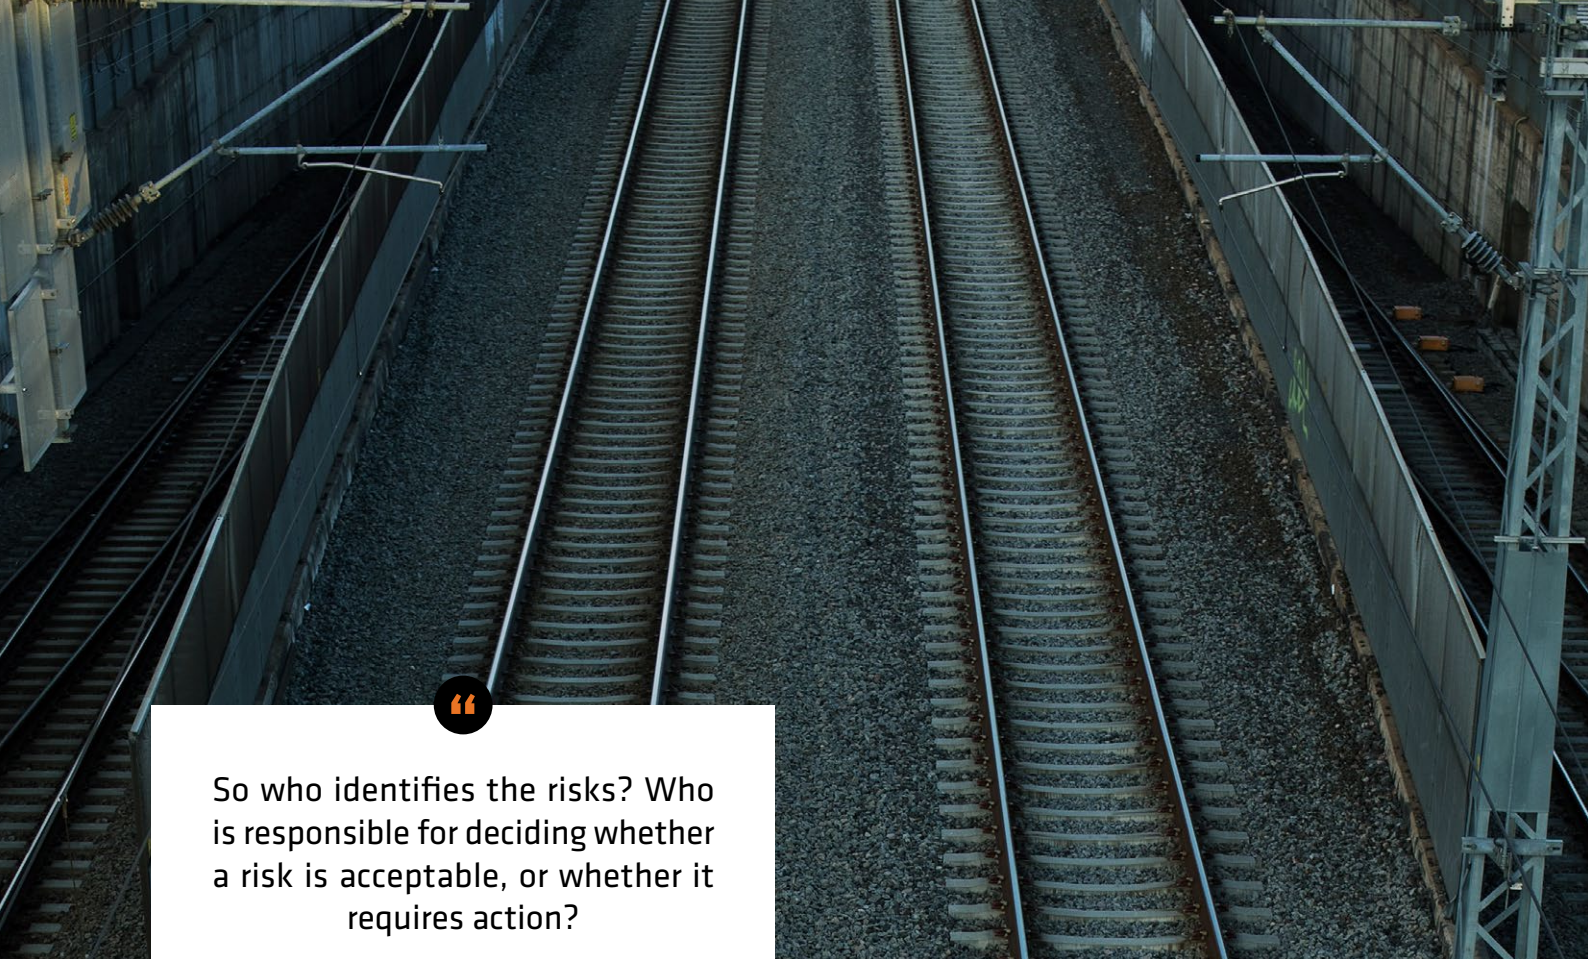

> So who identifies the risks? Who is responsible for deciding whether a risk is acceptable, or whether it requires action?

integrity or availability. When we know where the information is, we can apply information security in a strategic, cost-effective manner.

WHY CLASSIFY INFORMATION?

Information classification is really the starting point for information risk management and gives tremendous value to the risk management process. It allow organisations to set a value on their information, and make a proper evaluation of the investment they are willing to make to protect it.

Classification is a job that only management can do. They can evaluate the possible consequences of various scenarios of compromise to the different types of information the organisation uses (intellectual property, market research, strategic plans, pricing information and many more). They can rate the consequences for different aspects such as financial loss, life and health, production, reputation, compliance with laws and regulations, etc. This process identifies critical information, and since we know now where this information is, we can apply the same priorities to the systems and processes that use, store, and transport the information.

LEAVE IT TO IT AND SECURITY

So, where are we now? We know what information we have and where it is, we know how critical the information is, and therefore how critical our systems and processes are. We know what risks we are willing to accept, so *now* can we buy some more boxes and leave it all to IT and security? Not quite yet.

Before IT and security can go off and identify the best security solutions, they need to know what vulnerabilities exist in our existing systems and processes. Boxes cannot reduce all risks; changes in routines and procedures or training and increased security awareness can be a better approach for many risks.

In order to identify vulnerabilities we perform risk assessments. All the work we have done up to now allows us to focus on critical systems, as we know how the information moves through our processes. With our overview of information, systems and processes, we can design cost effective and efficient security strategies. We can implement solutions that reduce a whole range of risks, and place them where they will do the most good.

We can then even go to the security vendors and tell them what we need, and let them try to convince us that they can deliver.

Finally, at last, we can let IT and security take care of it while we get on with running the organisation. That way, perhaps we can stop making the same mistakes from the last 43 years.

THE FIREWALL IN THE NEW WORLD OF IT

NASA Image ISS044E18893: the cays of the Bahamas, one of the most recognizable points on the planet for astronauts

TORMOD EMSELL LARSEN
Senior Consultant
mnemonic

HARALD HANSEN
Senior Consultant
mnemonic

After reading this article, you will:

- Have a general overview of the major changes the firewall has been through the last 30 years
- Understand how new trends and developments challenge the value of the firewall
- See why we still need the firewall and in what contexts it has future relevance

The IT infrastructure is changing. The use of cloud services is increasing and users can access their corporate resources from anywhere in the world. More traffic is encrypted. Boundaries are becoming unclear. Virtualisation technology is taking over the data centre. So where does the firewall fit in the new world of IT?

THE FIREWALL'S EVOLUTION

From the first packet filters to the firewalls of today, firewall technology has gone through several evolutionary steps. As the threat landscape changes, so does the firewall. The timeline on the next page describes the firewall's main evolutionary phases over the past 30 years.

> Where does the firewall fit in the new world of IT?

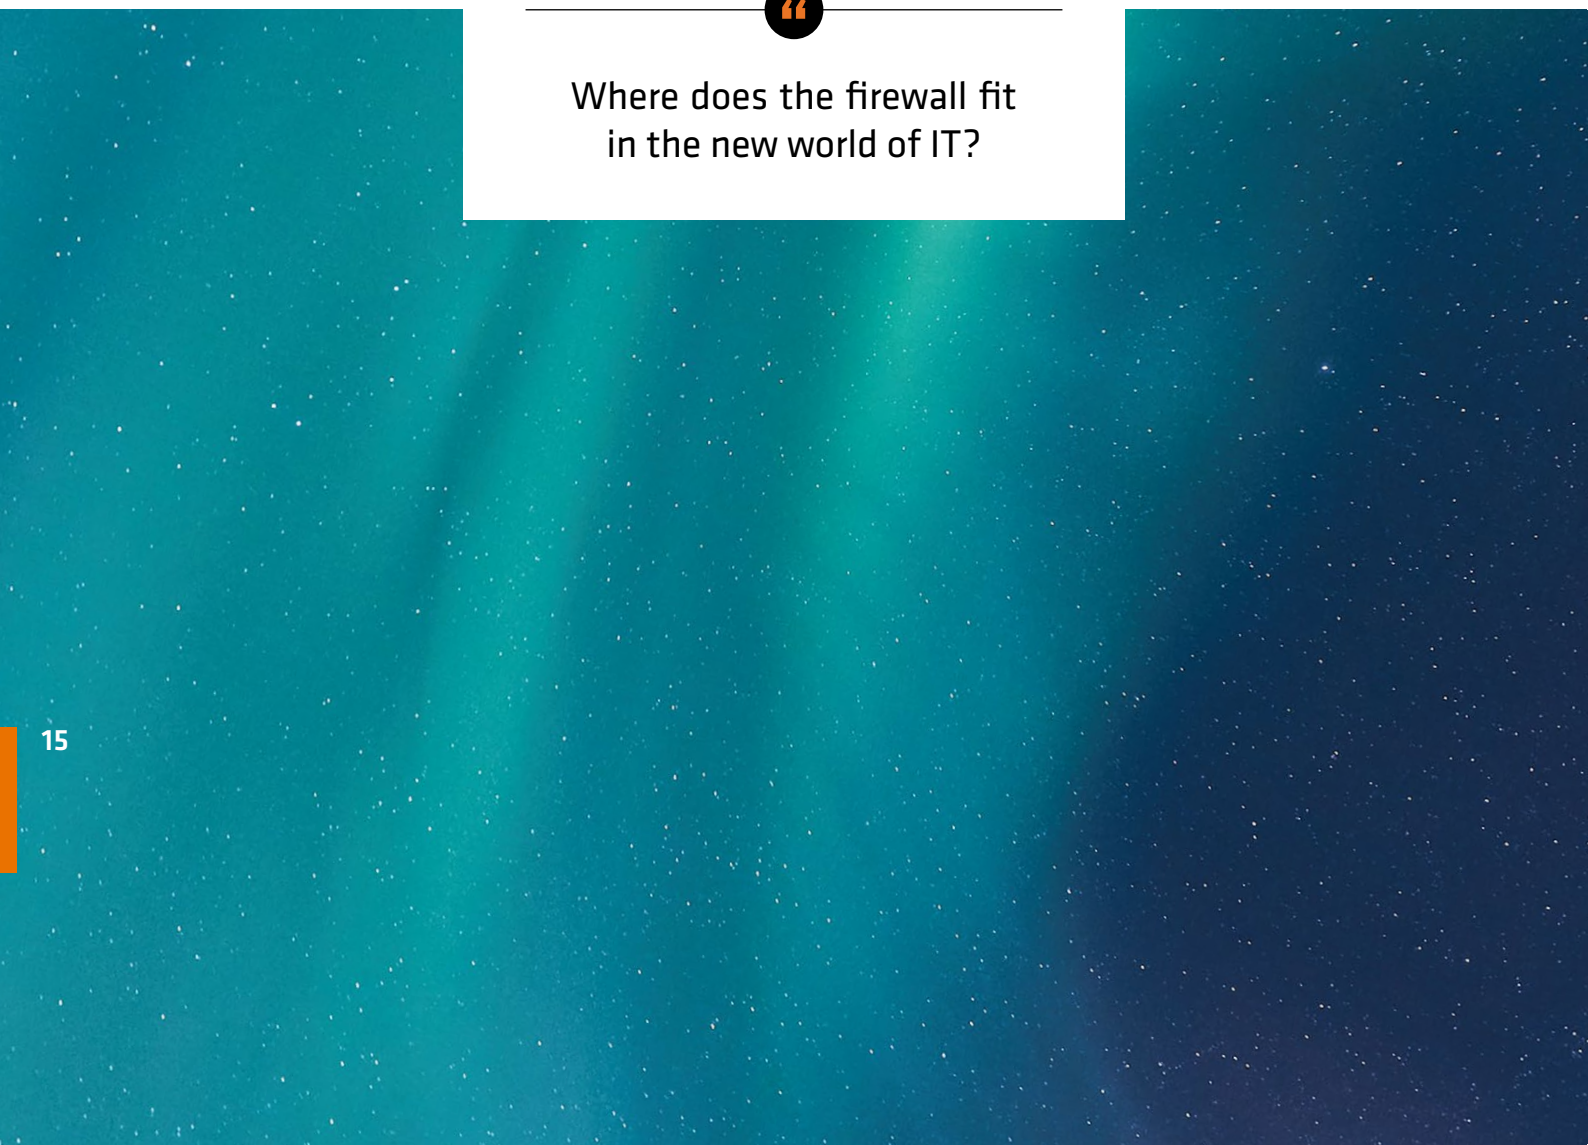

80s

It started with simple stateless packet filters in the late 1980s, based on IP-addresses and ports. Each packet was evaluated without relation to other packets. These packet filters were vulnerable to spoofing and were not able to determine if a return packet belonged to a legitimate connection.

90s

In the '90s the stateless packet filters were improved and replaced by the stateful firewall, which kept track of connection states. It handled traffic on layer 3 and layer 4 of the OSI model, while inspection at higher layers were left to other security products.

In the same decade application level firewalls added application level intelligence for certain services, and were able to detect protocol misuse.

2000s

In the first half of the 2000s, a new type of firewall - Unified Threat Management (UTM) - introduced deeper content inspection using technologies such as Intrusion Prevention Systems (IPS), Anti-virus (AV) and Web Filtering. This consolidation of products reduced costs and was widely adopted in the SMB market.

In the second half of the 2000s, web services became more popular, and access policies based on IP addresses and ports were no longer sufficient.

Yet another type of firewall - the Next-Generation Firewall (NGFW) – was introduced. The NGFW enabled access policies based on traffic content, which was able to identify applications by looking for characteristics in the data stream content. NGFWs were designed with better scalability and performance, and were considered a product suitable for large enterprises as well. It introduced new features such as user identification, sandboxing technology, HTTPS inspection and threat feeds.

Today

Today, most firewall vendors offer a broad spectrum of products in addition to the firewall platform, such as end-point protection and cloud security. In addition, they typically put increased efforts into threat intelligence research and develop ecosystems providing threat information sharing amongst their product line.

The firewall has changed remarkably since its start. In the rest of this article, we will discuss aspects that we think will influence the future evolution of firewall technology and usage.

HTTPS INSPECTION

HTTPS ON THE INTERNET IN 2017

There is a significant increase of HTTPS usage in today's network traffic, both internally and externally. By the end of 2017, almost 30% of the top 1 000 000 websites now use HTTPS by default[1] and the trend is increasing by 1-2% every month. The number of page loads measured by Firefox telemetry show that about 65% of web traffic is encrypted[2]. Similar data from Google Chrome users show an encryption load of more than 81%[3]. In short, HTTPS is well adopted.

One of the drivers for this increase in HTTPS usage is the pressure from the browser giants, namely Firefox and Chrome, and search engines, all of whom demote HTTP sites in favour of HTTPS sites. In the future, a red URL-bar is planned for HTTP and insecure HTTPS connections.

Also, new and free domain validated (DV) certificate authorities, like Let's Encrypt[4], help smaller sites adopt to HTTPS, and most web hosting businesses provide basic DV-certificates complimentary with their services.

Lastly, a wider understanding of the security concerns regarding the transport of unencrypted and unverified data over the Internet has hit the broader population's awareness. The simplification of deployment and handling of certificates, full support for Server Name Indication (SNI) in all modern browsers and the general demand for HTTPS from users effectively make HTTPS a requirement for all current websites.

VISIBILITY AND HTTPS INSPECTION

While HTTPS and TLS encryption is securing the Internet for the user, the need to prevent both attacks and data exfiltration still persists. Attackers have seen the advent of free DV certificate services as a possibility to hide their malicious payload with little financial cost or effort[5]. Both the turnaround of these new services (a DV certificate can be issued in minutes with API automation) and the fact that they are free, is moving both exploit and payload delivery to HTTPS.

To gain visibility in an HTTPS connection, it is necessary to perform what in practice is a man in the middle attack (MITM) on the client with a device that is broadly termed an 'SSL Interception Proxy' (SSL-proxy). By replacing the certificate of the web server and acting as the encryption partner, the client believes it is communicating with the server, even though the SSL-proxy is the real partner. This requires the client to trust the SSL-proxy, and requires trusted certificates installed on both devices.

HTTPS INSPECTION IN A NUTSHELL

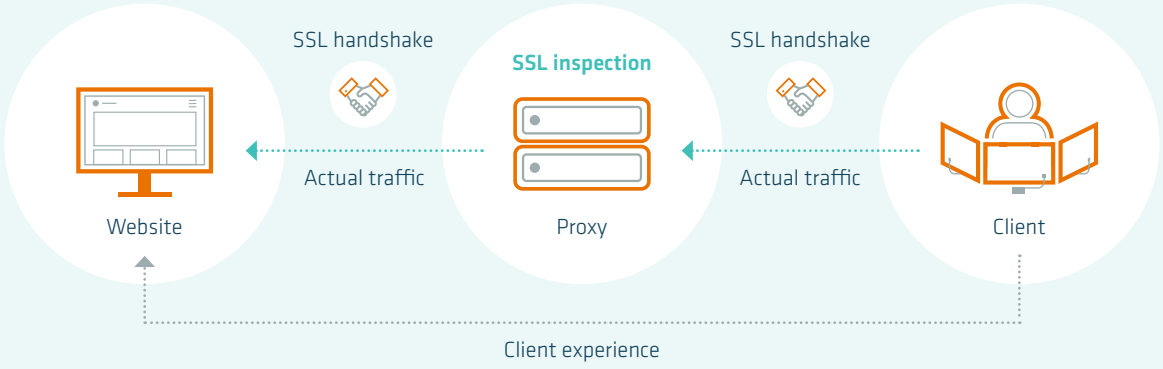

[1][2][3][4][5] *Find the references at www.mnemonic.no/references-2018*

CHALLENGES WITH HTTPS INSPECTION

The challenges IT security is facing with HTTPS inspection can be divided into three issues: legal, design and technical.

Legal issues

The legal side is non-technical, and not part of this article, though worth mentioning as a hurdle when deciding what data to decrypt and whom has access to it. For example, should traffic towards personal banking be inspected? Social media? Government services websites? What about foreign government sites? Governance and legal departments will have to assist in forming the security policy of the organisation.

Design issues

The fact that one has to be between the client and server when doing HTTPS inspection is a limiting factor. This means that either the client traffic has to pass the firewall at all times or the client has to use a proxy service, either on premise or in the cloud (Infrastructure-as-a-Service (IaaS) or Platform-as-a-Service (PaaS)).

Also, the need to issue trusted certificates will increase the risk to businesses if keys are lost. A well-designed Enterprise Certification Authority (CA) is a minimum requirement.

The movement to cloud-based services, especially in the Software-as-a-Service (SaaS) range, increases the load on both the Internet connection and the SSL-proxy. This has to be taken into account, as security measures that were a part of on-premise devices, like AV and (H)IPS, now must be provided by third-party services or in the perimeter firewall/proxy service.

Technical issues

As the percentage of HTTPS traffic increases, the load on NGFW and proxies will increase as well. Before HTTPS inspection, the content of the encrypted traffic was exempt from inspection. This means that not only will the load increase because of decryption/re-encryption, but also because a higher percentage of the data is passing the firewall inspection filters in a readable state. While the firewall was sized to handle the HTTP load, combining both content and HTTPS inspection can be a limiting factor.

In the following sections we will explore some of these technical issues.

Certificate pinning: a boon and a curse

HTTP Public Key Pinning (HPKP), or certificate pinning, is a mechanism to resist impersonation attacks using other certificates - mis-issued or fraudulent - than the certificate originally intended for the service[6]. Because an SSL-proxy's MITM mechanism relies on replacing the certificate, the process will fail.

For end-user security, especially in mobile devices where it is harder for the user to verify certificates, this is a boon. For IT operations however, it is a curse as all pinned sites have to be excluded from inspection to function.

Google Chrome respects the HPKP policy and has built-in pinning of Google services. If an enterprise CA is found in the certificate store, it will ignore pinning when the replaced certificates are issued from this CA.

Upcoming versions of Chrome are going to replace HPKP with Expect-CT headers[7]. How this latest development is going to impact HTTPS inspection is still too early to determine, as the release is scheduled for Q2 2018.

Certificate Stores outside enterprise control

As with certificate pinning, some applications, IoT, embedded and mobile devices and BYOD will not trust the enterprise CA. These devices will always fail on connection through an SSL-proxy, thus needing exceptions in the decryption policy.

New transport protocols

With the advent of SPDY, now HTTP/2[8], the way clients communicate with web servers has changed from a simple, text-based, linear protocol to a binary, multiplexed and compressed protocol. The standard for instance allows the server to proactively transmit information that has not been explicitly requested by the client, and the main browser vendors have decided to only support HTTP/2 over TLS. Thus, the firewall cannot perform passive pattern recognition of the data stream. It has to reassemble, partition and decrypt multiple streams within the same TCP session and has to keep a more complex state machine for each session in memory. Few, if any, inspection solutions fully support HTTP/2 yet, so the real performance impact is yet to be seen.

There still are some advantages when using binary protocols, as there is no need to convert text into data. Semantic inconsistency can lead to severe security consequences such as cache poisoning and filtering bypass, the ambiguous conditions caused by parsing text is therefore reduced significantly.

There are also suggestions for other data transmission protocols to reduce latency. An example of this is the UDP-based protocol QUIC[9] (Quick UDP Internet Connections), which combines elements from TCP, TLS and HTTP/2 into one protocol. For now, no mainstream NGFW supports inspecting these protocols, but they are recognizable and can therefore be blocked.

Handling decryption failure

Decryption failures are handled in various ways depending on the product and scenario - from fail close (if fail then block), friendly error (if fail give an error but permit traffic) to full pass-through (if fail then let traffic through). Most will fail close when the HTTPS inspection device cannot handle the connection due to policy restraints, certificate failure or other issues. Traditionally the end user would receive a security prompt asking them to continue. However, since the firewall is now making this decision on behalf of the user, the failure mode has to be decided by policy. A pass-through policy is more user friendly, while the fail closed policy is more secure.

Encryption ciphers

As new encryption ciphers and hashing techniques are adapted, the HTTPS inspection device has to adopt support for these. This requires a continuous development on the manufacturer's side and keeping up with updates on the customer end. As of January 2018, more than 60% of websites surveyed by ssllabs.com have Grade A or better implementation[10].

[6 7 8 9 10] *Find the references at www.mnemonic.no/references-2018*

HTTPS INSPECTION: A BALANCING ACT

Taking all the discussed challenges related to HTTPS inspection into account, the bottom line is that the challenges are many and inspection of HTTPS is affecting both the firewall load and IT operations. Combining the need for visibility and the users' expectation of problem free access to web and cloud services is a balancing act, and requires new design patterns, routines and thoughts regarding security operations in the enterprise.

PROTECTING THE ENDPOINTS

The firewall is *one* of the network-based solutions that has a role in protecting the endpoints.

Traditionally, network segmentation and network access control have been basic measures protecting enterprise resources, while content inspection has been more effective for detecting threats from the Internet towards the endpoints. However, inspecting the traffic content is not that trivial anymore, because of the HTTPS challenges described above.

The mobility aspect is also challenging. Users have become more mobile, and there are expectations and demands of being able to work from any location. In order to offer this without compromising security, many companies use "always-on" VPN solutions that backhaul all traffic, including the Internet destined traffic, through the company data centre. That way, all client traffic is still protected by the central network-based security solutions, regardless of how mobile a user is.

The drawback with this design however is that it will impact the user experience by introducing latency, and will influence location-based services as users may appear to be in a different country than they are located. Likewise it is not well suited for cloud services that are, by design, accessible from any device and from any location.

To secure the web-based traffic regardless of the endpoint location, we see that many enterprises are now moving to cloud-based secure web gateways for their managed endpoints. Firewall-as-a-service (FWaaS) and software defined WANs (SD-WAN) with built in threat protection are also interesting cloud-based alternatives. FWaaS and SD-WAN can be offered in an integrated solution, providing secure transport channels between Internet, corporate offices and roaming users. All traffic (and not only web-based traffic) can be inspected by NGFW technology in the FWaaS. This field is still maturing and there is limited vendor choice[11].

With the rise of cloud services, there has also been increased focus on identity. Access to cloud resources are mainly provisioned based on the authenticated user identity rather than the origin of the client endpoint. This fits well in the new world of IT.

With the rise of cloud services, there has also been increased focus on identity. Access to cloud resources are mainly provisioned based on the authenticated user identity rather than the origin of the client endpoint. This fits well in the new world of IT.

[11] *Find the references at www.mnemonic.no/references-2018*

PROTECTING CLOUD SERVICES

As enterprises move into the cloud, they also need to think about what security measures to implement. Depending on the types of cloud services, there are different options. Identity Management Services and Cloud Access Security Brokers (CASB) are examples of solutions securing many SaaS applications. CASB are API-based or proxy-based policy enforcement points mainly for SaaS applications. Some of the provided features are authentication, single sign-on, data loss prevention, data protection, malware detection and compromised account detection.

However, when considering the role of the firewall itself in the world of cloud services, the most obvious place is probably in an Infrastructure-as-a-Service (IaaS) environment. The IaaS providers offer the infrastructure, which includes computing power, storage and network. The customer deploys their servers and applications on the top, and is responsible for security measures in the areas of network, operating system, application, data, identity and access management.

Compared to the traditional physical, on-premise environment, the mechanisms for protecting corporate IT resources in IaaS are more or less the same.

An IaaS network is designed using familiar concepts such as network segmentation and access control. The IaaS providers themselves offer basic access control on layer 3 and 4 in the OSI model, similar to what the firewalls offered 20 years ago.

An alternative to solely relying on the basic access control from the IaaS providers is to deploy a virtual firewall/NGFW from third party vendors. Some additional advantages with this approach may include:

Access policies based on	
User Identities	URL categories
Applications	Data types

Threat content inspection	
IPS	Sandboxing
AV	Threat Feeds

VPN
Unified Management for the entire firewall environment
Detailed logging capabilities
Network-based troubleshooting capabilities

These virtual firewall solutions are available not only to public IaaS clouds, but also to private clouds and software defined data centres. The need for traffic inspection and control is the same.

Often, using a third party NGFW is preferable, but it depends on the overall network design, type of resources and data to protect, and so on. For example, to protect cloud-based, Internet exposed web services, other solutions may be more suitable or give more value, such as DDOS Protection or Web Application Firewalls (WAF).

DYNAMIC POLICIES AND AUTOMATION

A challenge in managing firewall solutions is maintaining the policies. Deploying new services and applications in the enterprise is easier, and in more demand than ever before. With such a high pace of changes, it is becoming more difficult to keep firewall policies updated with a clear structure.

However, the firewall products themselves have introduced features to ease the task of policy maintenance. One such feature is the firewall's ability to integrate with other systems, thereby making the firewall policy more dynamic. This can be achieved using APIs or the built-in integrations offered from the leading firewall vendors. Examples are:

- Object synchronization with Configuration Management Database (CMDB) systems
- Dynamic group objects automatically populated and updated from other systems (e.g. based on tags from VMware NSX or endpoint security groups from Cisco ACI)
- Triggering actions in other systems (e.g. if the antivirus solution detects a machine infected with malware, it can notify VMware NSX, which then instructs the firewall to isolate the machine)
- The ability to delegate sections of the policy to system owners without compromising the security policy as a whole

With the increasing support and usage of APIs, the use cases are many.

We believe making use of more integration and automation on the firewalls will benefit the administration tasks of maintaining the policy and improve security functions. However, one must be careful and aware of where the security controls are, as these in practise can be moved away from the firewall to other systems. For instance, when using dynamic groups learned from tags in VMware, the ones applying the tags in VMware should be aware of how this affects the network-based access policy. Hence, when introducing such changes to the firewall, one should also revise other elements such as roles, responsibilities and routines.

THE FIREWALL IS FAR FROM DEAD

The firewall is far from dead, but some areas of its functionality are certainly threatened.

With the introduction of NGFWs, we gained greater visibility into traffic. However, with the rapid increase in encrypted traffic and the use of HTTPS, we were suddenly losing visibility. Therefore, we started decrypting the traffic, and once again we regained visibility. Things looked promising, but it turned out to be complicated. Decryption is becoming harder to do, and once again we are gradually losing visibility. The situation is still manageable, but as the trend continues, perhaps we are forced to rely more on content inspection at the endpoints instead of the network.

The clear boundary between the untrusted external network and trusted internal network is not that clear anymore. Cloud services and mobility requirements are blurring the boundaries. In situations where we do not control the endpoint or the network, which could be the case with SaaS, other security mechanisms than firewall technology need to be considered. An example is CASB solutions.

For managed endpoints in enterprises with high mobility requirements, we see cloud-based secure web gateways as more appropriate for controlling web traffic. In addition, access control based on identity is becoming more important, while the endpoint location is of less importance.

For securing data centres, firewall technology is still inevitable. This applies for both cloud and on-premise data centres, and includes traditional network segmentation, modern access control policies and threat inspection. The trend of more automation and integrations will probably continue.

Though many have described its demise, the firewall still has a valuable place in 2018, and beyond. However, as the IT infrastructure continues to evolve, the firewall's future existence hangs on its ability to adapt to its new surroundings.

Global

WORD ON THE STREET

MICROSOFT

ENTERPRISE CYBERSECURITY GROUP

MANDANA JAVAHERI

Director, WW Business Development, Cybersecurity

Microsoft (Nasdaq "MSFT" @microsoft) is the leading platform and productivity company for the mobile-first, cloud-first world, and its mission is to empower every person and every organization on the planet to achieve more.

WHAT IS YOUR BIGGEST SECURITY CONCERN?

Despite the industry's investments and efforts, cyber-attacks and breaches are still happening daily. Half of the breaches result from criminal intent, which are very targeted and impactful. Hackers are advancing their skills and methods and hacking tools are readily available in the black market, which makes it easy for hackers to make money and engage in cybercrimes without a lot of effort. A shortage of cybersecurity professionals and expertise is also a global concern, which adds to the complexity of cybersecurity. Cybersecurity is a business imperative and breaches can have significant negative impact on the business and its reputation, and in some cases, can cause the business to fail completely. My main concern is that organizations are not equipped with the right strategy, technology, and resources to survive costly cybercrimes.

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

There are two areas that we have been talking about for a while, without a lot of progress: Consolidation and Automation. Looking at the security landscape, our clients have too many security vendors and products in their environment, which makes it very difficult to manage. Additionally, there is no real integration and collaboration between these products. We keep adding noise to an already noisy environment. Our customers are also asking for orchestration and automation of their security operations and processes. One of our goals at Microsoft Enterprise Cybersecurity Group is to accelerate these efforts through our strategic collaboration and partnerships and to help secure our customers' digital transformation journey.

WHAT GIVES YOU HOPE FOR THE FUTURE OF SECURITY?

Advancement in technology, including the use of AI and Microsoft Intelligent Security Graph to predict, prevent and detect threats before they become an incident even quicker; investments being made in innovation and R&D, for example, Microsoft is spending $1B a year in security R&D; and the industry's commitment to collaborate with the goal of making the world a safer and more secure place for all of us.

Norway

WORD ON THE STREET

INTILITY AS

ANDREAS HISDAL
CEO

Intility is an enterprise grade technology platform that is currently used by 600 companies across more than 1500 locations in Scandinavia and worldwide. The platform includes a complete Workplace-as-a-Service solution and digital business platform for digitalization of core business. Intility is an enabler for companies to utilize information technology more efficiently, increasing their productivity and competitive edge. The platform is continually upgraded with new functionality and supports an increasing number of integrated cloud services such as Microsoft Azure, Office 365, Amazon Web Services and Salesforce.

WHAT IS YOUR BIGGEST SECURITY CONCERN?

As businesses are increasingly focusing on digitalization and automation, we see systems, applications and cloud services becoming increasingly integrated and dependent on each other. This increases the possible attack surfaces and subsequently the risk for security breaches. This trend requires a high level of security competencies, both from the end-user and the system administrator perspectives. As information security expertise and knowledge is a rare commodity these days, we believe that the lack of such skills will impose even greater risks in the years to come.

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

In the age of digitalization, an increasing number of services and devices (IoT) are made available online. This greatly increases business opportunities and can provide businesses with a competitive edge. However, businesses need to be aware of the risks they impose, as this can introduce new risks and create attack surfaces inside an existing security perimeter.

WHAT GIVES YOU HOPE FOR THE FUTURE OF SECURITY?

We have seen increased personal awareness with regards to information security. This is of key importance, regardless of any technical security features or implementations, in order to keep data and information secure. This development is fueled by regulatory drivers such as GDPR, that put security on the agenda across all industries. As public interest and awareness of IT security increases, vendors and service providers are forced to provide and create even more secure systems and services. Our hope is that this will lead to fewer security incidents and breaches, and greater incentives to develop and implement new security features.

ESPEN HOVIND
Senior Consultant
mnemonic

SECURING YOUR WEB APPLICATIONS IN THE CLOUD

A DOWN TO EARTH APPROACH

After reading this article, you will:

- Understand some of the security concerns related to network, access control, encryption and the web application when using Infrastructure-as-a-Service (IaaS) or Platform-as-a-Service (PaaS) cloud models from major vendors like Amazon Web Services (AWS) or Microsoft Azure (Azure)
- See that even though you face several of the same challenges in the cloud as you do on-premise, the solutions may be different
- Walk away with some useful advice on how to secure your web applications in the cloud

Attacks on web applications only account for 8% of reported security incidents. Such attacks, however, are responsible for over 40% of the incidents that result in a data breach.*

If your company, like many others, has decided to utilize cloud services to develop and deploy your own web applications, it is likely that this is because of its many advantages. Reduced capital expenses, short time to market, agility and elasticity are tempting benefits to be gained from utilizing cloud services. But how do we ensure the ability to reap the advantages, while still securing our web applications?

IAAS VS PAAS

The IaaS model implies that you are basically deploying virtual machines and networks in the cloud. The cloud provider is responsible for underlying cloud infrastructure components like the data centre and virtualization platform, and provides you with processing, storage and network resources. You are responsible for the deployed operating systems, application servers, and applications.

The PaaS model delegates operating system and application server maintenance to the cloud provider, giving you the capability to create environments where you can deploy your applications, and permits access to built-in services like a relational database.

IaaS and virtualized compute like AWS Elastic Compute Cloud (EC2) and Azure Virtual Machines can be the best choice if you require full flexibility, but also introduce many of the same security challenges as on-premise. You will have to install, configure, patch and maintain your platform. PaaS services like AWS Elastic Beanstalk and Azure Web Apps trade some of the flexibility for reduced maintenance costs. In this instance, the cloud vendor will configure and patch your infrastructure.

It is important, however, to be aware that in both models you develop and maintain the web application source code and data, and herein lies some of the major challenges.

WEB APPLICATIONS AND THE NEVER-ENDING SECURITY CHALLENGES

From the beginning of web application development, the core security problem has been that the user can provide arbitrary input to the web application. Therefore, all input must be treated as potentially malicious. This fact, combined with the growing complexity of web applications, the increased demands for functionality, and strict resource and time constraints, illustrates the considerable task at hand when trying to secure your web applications.

By looking at *The Open Web Application Security Project (OWASP) Top 10 Application Security Risks* from 2003 until today, we recognize several of the same vulnerabilities always appearing with a high ranking:

- Injection flaws
- Cross-Site Scripting (XSS)
- Broken authentication
- Session management
- Misconfiguration and vulnerabilities in operating systems and application servers

Several of these risks relate to the application code itself, and is a relevant security challenge when utilizing both the IaaS and PaaS models. Securing your web applications is hard, this fact does not change if you move to the cloud.

DELEGATION OF RESPONSIBILITY IN IAAS AND PAAS

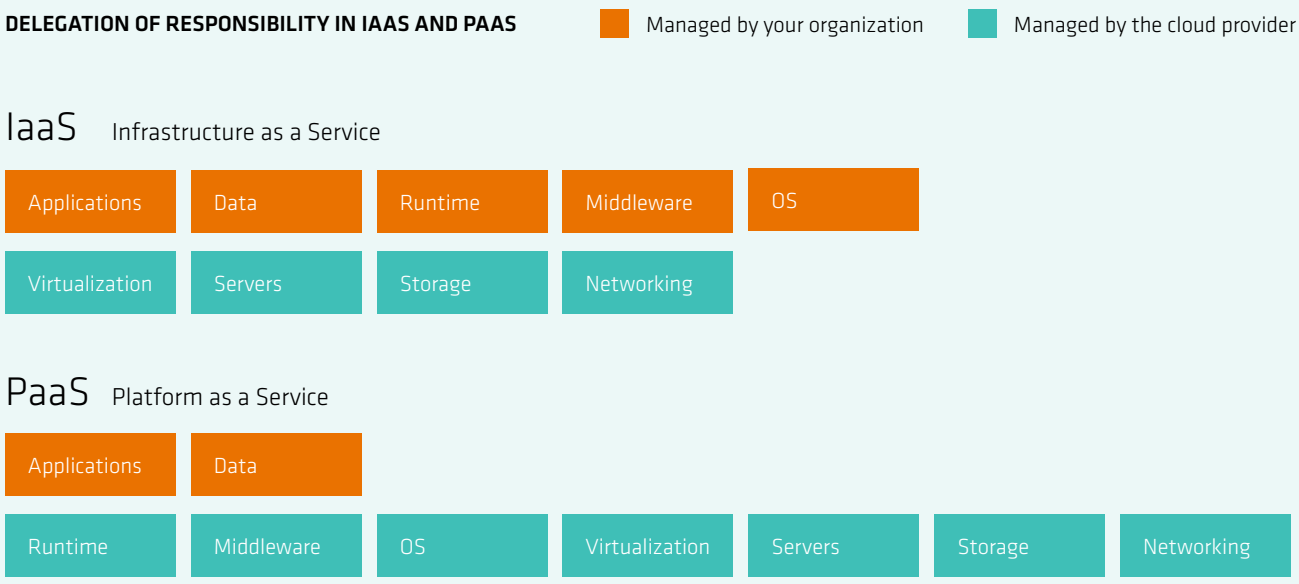

* *According to Verizon's Data Breach Investigation Report (DBIR) based on investigations and reports of over 100,000 security incidents*

NETWORKING

The network layer has traditionally been the focus of information security. An integral part of your on-premise security model is most likely network security controls, and this also needs to be a point of focus in the cloud.

AWS and Azure acknowledge this fact, and offer network features that map closely to their on-premise counterparts, like subnets, routing tables, access control lists, stateful firewalls, and load balancers. Both vendors describe how to deploy these mechanisms to create a robust security policy in their reference designs.

Segment your network into a multitier architecture, control routing, and use the AWS Security Groups' and Azure Network Security Groups' stateful firewalls to restrict access to your subnets. The load balancers AWS Elastic Load Balancer and Azure Application Gateway can act as an entry point to your applications, providing auto-scaling and SSL/TLS termination. This design applies to IaaS, but you should also apply network control services like subnets and load balancers to your PaaS resources.

Remember to be aware of the defaults. In AWS, every computer inside an AWS Virtual Private Cloud (VPC) can potentially reach each other on all network ports, even if they are located in different subnets. In addition, AWS Security Groups allow all outbound traffic by default.

To manage your servers, you should use a Virtual Private Network (VPN) or deploy bastion servers with exclusive access to your server's management ports.

IDENTITY, AUTHENTICATION AND ACCESS CONTROL

Whether you choose AWS or Azure, IaaS or PaaS, the foundation of your cloud infrastructure security is based on identity and access management.

In the cloud, your servers and networks are virtual, and the built-in services are Application Programming Interface (API) endpoints. Every operation on your cloud resources is an API call. While this gives near limitless flexibility and encourages the automation of tasks, the consequences of security misconfiguration is also increased.

AWS Identity and Access Management (IAM) and Azure Active Directory (AD) should be used to control access to your cloud resources. Features like roles, granular permissions, multi-factor authentication, temporary credentials, and identity federation will be useful mechanisms to protect your cloud infrastructure.

The complexity of identity management in the cloud should not be underestimated. It is the responsibility of the customer, and you should invest the necessary time to understand and maintain your configuration. After all, an all-time top 5 question on the AWS Security Blog is "Where's My Secret Access Key?"

Where's My Secret Access Key?

AWS Security Blog

AWS Security Blog's most viewed blog post in 2017

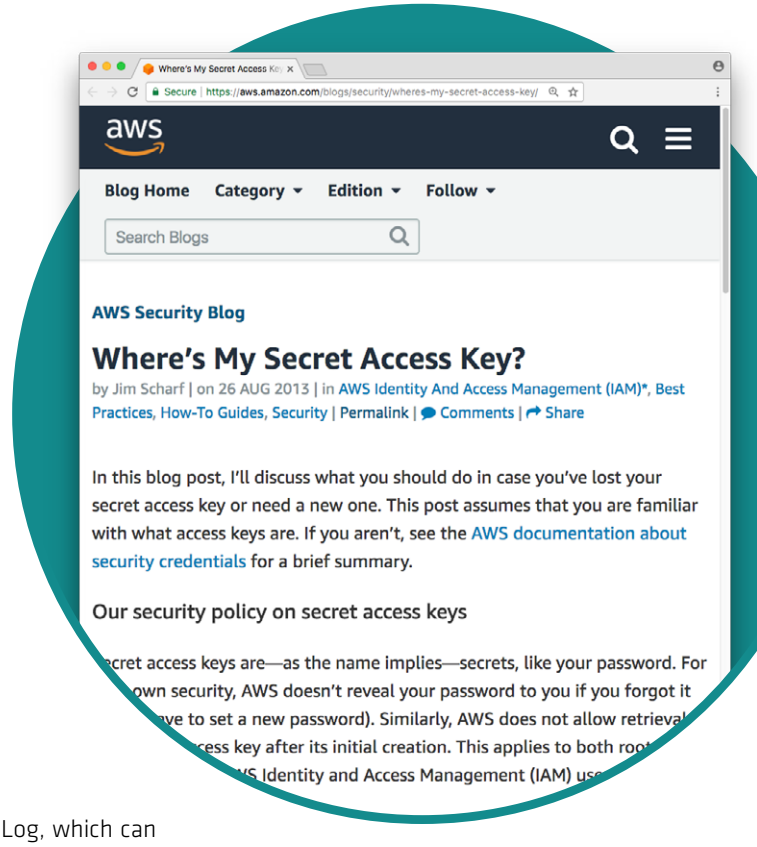

All API calls are monitored in AWS CloudTrail and Azure Activity Log, which can be integrated with built-in push-notification services, or on-premise Security Information and Event Management (SIEM) solutions. These audit logs can be valuable assets during operational and security incidents.

CLOUD AND CRYPTOGRAPHY

The Internet is steadily progressing towards encryption of all web traffic using SSL/TLS. There is a strong chance you may already be encrypting traffic destined for your on-premise web applications, as well as traffic between your application and database servers. Cryptography can also protect your web applications in the cloud. While typical challenges are solved by AWS and Azure, new challenges are also introduced.

CIPHER SUITES AND LOAD BALANCERS

Both AWS and Azure maintain cipher suite profiles that can be applied to their load balancers. A cipher suite profile consists of a set of supported cryptographic algorithms for key exchange, bulk encryption and message authentication. When using protocols like HTTPS, the choice of cipher suite is negotiated between a client and the SSL/TLS termination point.

Since a cipher suite depends on the security of its algorithms, and history has shown that cryptographic algorithms can be broken (DES, RC4, MD5, ...), the set of cipher suites should be flexible. Cipher suites which are found vulnerable must be removed by the cloud vendor.

When trusting the cloud vendor with the cipher suite configuration, it is important to be aware of the fact that they maintain a default configuration that must be suitable for all their customers. The required testing, implementation and communication processes might therefore delay the cipher suite modifications considerably. For example, in February 2015 the RC4 cipher was found vulnerable and its usage was prohibited by the Internet Engineering Task Force (IETF). That same month Amazon issued a statement that RC4 support would be removed from all new load balancer instances using the default cipher suites. Azure, however, did not remove RC4 support for the PaaS Web Apps service until several months later. The customer, if they knew why and how, could manually disable RC4 support at any time.

The cipher suite profiles maintained by AWS and Azure should be adequate in most scenarios, but there are exceptions. Some customers may have stricter requirements than the defaults provided (e.g. in Payment Card Industry (PCI) environments). If so, the customer must evaluate the default configuration, and eventually configure and maintain its own cipher suite configuration. Both AWS and Azure offer this possibility for IaaS and PaaS deployments.

Additionally, on-premise deployments often terminate SSL/TLS on the load-balancer and route traffic in plain-text to the back-end servers. On-premise, this may be acceptable and convenient, since you control and trust your data centre. In the cloud, traffic should be re-encrypted to ensure end-to-end security between the users and your servers. The AWS Elastic Load Balancer and Azure Application Gateway offer such functionality.

The gradual deprecation of the RC4 cipher

2015

- **February:** IETF prohibits the RC4 cipher suites in RFC7465
- **February:** AWS Elastic Load Balancer releases security update to disable RC4
- **April:** AWS CloudFront removes RC4 from list of supported ciphers
- **July:** Microsoft removes RC4 support from Azure Web Apps (Formerly Azure Websites)

2016

- **April:** Microsoft removes RC4 from the supported list of negotiable ciphers on service endpoints

DON'T EXPOSE OR LOSE YOUR SECRETS

Public Key Infrastructure (PKI) and symmetrical encryption require key management. Where should we store these secret keys and how can we control and audit access?

Managing your private key in a PKI is crucial. The private key is used by the key exchange process between the client and your SSL/TLS termination point to protect the unique session keys for this session. For the traditional RSA key exchange, this implies that whoever controls the private key can decrypt the encrypted data. This includes current and past (recorded) sessions. To reduce the impact of a leaked private key, the current recommendation is to instead use (Perfect) Forward Secrecy ciphers like Elliptic Curve Diffie-Hellman Exchange (ECDHE). ECDHE generates transient session keys which only exists during the session's lifetime. The private key, however, is still important because it is used to digitally sign the parameters used during the key exchange. In other words, the private key is used for authentication and "ties" the session keys to your website.

> Protecting your web application's user passwords is relatively easy, but often forgotten.

When protecting your web applications' data at-rest you can utilize symmetrical encryption ciphers like Advanced Encryption Standard (AES). A single secret key is used for encryption and decryption of data; thus it must be restricted to authorized usage only.

The problems mentioned above are well-known from on-premise hosting, and AWS and Azure provide services which can help solve these challenges. AWS Key Management Service, AWS Certificate Manager and Azure Key Vault can generate, store and control access to the keys used to protect the session keys, encrypt your application data or sign your X.509 certificates and API calls. If required, you can even generate your own keys on-premise and import these into the cloud key stores. Access control is enforced by key policies defined in AWS IAM or Azure AD. All key store operations are logged respectively to AWS CloudTrail and Azure HDInsight, or your own cloud or on-premise SIEM.

Customers with very strict security policies might have issues with a shared key management service. If the customer requires physical isolation and total control over the keys and application software, a dedicated Hardware Security Module (HSM) like AWS CloudHSM could be deployed. This solution will introduce additional management costs, but can provide the strongest protection for your private keys and cryptographic operations, as well as auditing functionality.

Protecting your web applications' user passwords is relatively easy, but often forgotten. Salting and hashing your customers' passwords will make all the difference if your user database is leaked. Your chosen programming language library most likely support several secure alternatives, like *passlib* for Python.

THE WEB APPLICATION

After the underlying infrastructure is secured, you can shift your focus to the web application itself. Even when building on top of a solid foundation, it is necessary with a clear understanding of the risks involved and the tools at your disposal.

SECURE CODING AND THE WEB APPLICATION FIREWALL

The core mechanisms of defending the web application consist of preventing unauthorized access to data and functionality, and handling malformed input. These problems are ideally solved in the web application design and source code. Several secure coding guides address this topic (for example the quick reference guide from OWASP[1]). While such guides serve as an excellent starting point, a glance at the history of web applications demonstrates that it can be unrealistic to base the security of your web applications solely on the notion of *perfectly secure code.*

Amazon and Microsoft acknowledge the necessity to protect the application layer, and announced Web Application Firewall (WAF) services in 2015 and 2017, respectively. A WAF can inspect, control, and modify application traffic between your web application and the users.

Since the WAF needs insight into the plain text application traffic, it is usually located on the load balancers. The built-in WAFs, if enabled and configured, can provide an extra layer of defense for your web applications primarily by enforcing a negative security model. This model can block known bad traffic like cross-site scripting (XSS) by using attack signatures (blacklists), but new attacks (zero-days) cannot be detected until a signature exists. The WAFs can also be used to mitigate typical volumetric DoS attacks.

[1] *Find the references at www.mnemonic.no/references-2018*

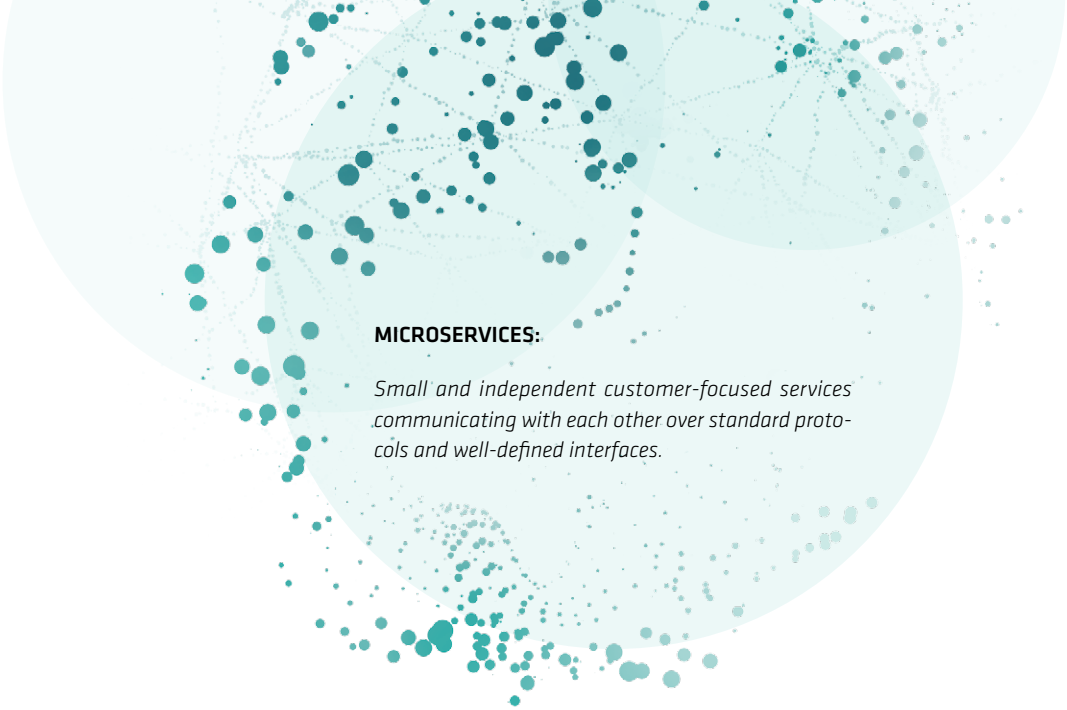

MICROSERVICES:

Small and independent customer-focused services communicating with each other over standard protocols and well-defined interfaces.

If you need advanced features like whitelisting and protection against the more subtle attacks on the application layer (i.e. layer 7) you have to look for alternative solutions. There are several options available from the vendors' marketplace.

A WAF can be a vital layer between your application code and the Internet. In addition, the cloud vendors also encourage the use of vulnerability assessments and penetration tests of web applications. Just remember to ask the vendors for permission first.

HTTP SECURITY HEADERS

After reading *The Tangled Web* from 2011 and learning about the vast array of possible attacks, many of us became somewhat disillusioned about operating securely in the Web. However, in the later chapters the author highlights examples of future technology that can help protect against some of these attacks. The future is here, in the form of new HTTP response headers like; HTTP Strict Transport Security (HSTS) and Content Security Policy (CSP). These can provide a new layer of protection against session hijacking and code injection attacks.

Whether you are using IaaS or PaaS, you need to configure support for these mechanisms on your web server. When utilizing Amazon Elastic Beanstalk, the default HTTP response headers do not include the above-mentioned headers. Behind the scenes the well-known Apache web server is running, and you can customize the environment using configuration files located in an *.ebextensions/* directory. The same applies for Azure Web Apps. Microsoft handles infrastructure and scales your apps, but the headers must be added manually. One method is to use Web.config to set HTTP response headers for your ASP.NET applications.

MANAGE YOUR MICROSERVICES

The primary reason you are moving your applications to the cloud might very well be that you are migrating to a microservices approach for development and deployment. Cloud services include built-in tools that are well suited to support microservices. Docker images can be deployed directly on your AWS Elastic Compute Cloud (EC2) instances and use Docker Swarm for orchestration, or use the built-in services like AWS EC2 Container Service (ECS) or Azure Service Fabric. Either way, the elasticity of the cloud and modularity of microservices are a good match.

Securing microservices is an entire topic on its own. Some of the risks are container breakouts, authentication and handling secret keys inside the containers.

Docker shares a kernel across containers, which increases the probability of a breakout. The Docker instances each run isolated in their own dedicated namespace, and cannot access the processes or resources in the other instances or on the host itself. Vulnerabilities in the kernel or misconfigurations can lead to instances breaking out of their namespace. Therefore, the isolation of Docker instances should be further improved by adding another layer of sandboxing using Mandatory Access Control like SELinux. Be aware though that SELinux may not be enabled on the OS images used by your deployment. For example, the default go-to image in AWS, Amazon Linux AMI, disables SELinux by default.

Handling authentication and secrets inside containers is complicated because there may be a large number of instances and they may be short lived. To manually distribute access tokens is error prone and introduces security risks.

The AWS API Gateway and Azure API Management services align nicely with the API-driven microservices architecture, and can manage your APIs by creating endpoints, centralize authentication, and perform caching and monitoring. For example, AWS API Gateway can control the edge security of your APIs by utilizing authentication frameworks like OAuth 2.0.

Another layer of protection would be to integrate the API gateway with a WAF. This is a frequently requested feature from AWS, but currently not easily implemented.

Microservices, DevOps and continuous deployment is trending, and for good reasons. The rapid deployment of new code, however, can impact the security of an application. A holistic approach to security is required, from network level controls and container isolation, to security scanning, application layer defense, and secure coding practices.

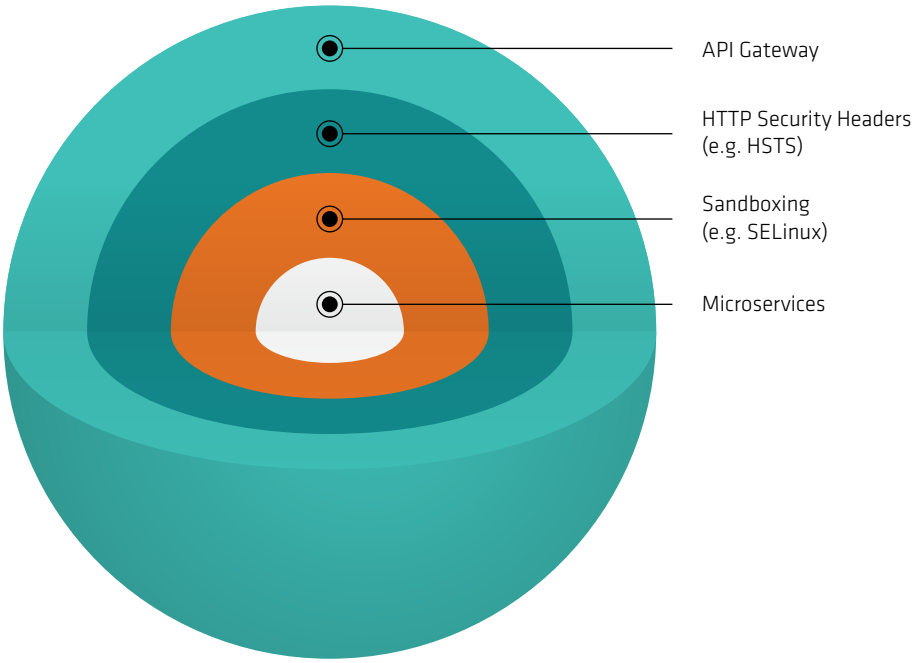

Microservices: a holistic approach

THE REST IS UP TO YOU

This article has highlighted some the challenges related to securely developing and deploying your web applications in the cloud. Essentially, you will face many of the same challenges as you do on-premise, but the solutions may be different. Security will be a joint effort between you and the cloud provider.

Large cloud platforms like AWS and Azure are well-designed and robust. These providers also have a proven history of continuously developing security tools that can be used to secure your applications in the cloud. A cloud deployment can be just as secure, or even more secure than an on-premise deployment. Just keep in mind that sufficient protection is not always achieved out-of-the box. The cloud model still requires that you understand how web applications can be compromised, and how to mitigate these threats. As always, if you deploy vulnerable code you will increase the risk of someone compromising your web application.

An essential concept in information security is *defense-in-depth* by deploying multiple layered and independent security controls. If you break this concept down, you will find that well-known technical prevention security controls like firewalls, cryptography, sandboxing, access control and secure coding still apply in the cloud. The cloud vendors provide the tools; the rest is up to you.

DETECTION

While this article focuses much on preventing data breaches, the core focus of any security team is being prepared. Prevent what you can, and have the capability to detect what you can't. According to a widely read blog post by Ryan McGeehan from 2016[2], practical experience has shown that the presence of detailed audit logs on a centralized log server is a major indicator of whether the root cause and impact can be concluded. The on-premise log integration options in AWS Cloud Trail/Flow Logs and Azure Log Integration should be implemented to securely store your audit, server and application logs.

[2] *Find the references at* ***www.mnemonic.no/references-2018***

Useful online resources:

AWS Security Blog: https://aws.amazon.com/blogs/security/	Microsoft Azure Security: https://azure.microsoft.com/en-us/blog/topics/security/	The Open Web Application Security Project: https://www.owasp.org

You can also find these online resources at www.mnemonic.no/references-2018

TURNING GDPR INTO AN OPPORTUNITY

ARNE CHRISTIAN GJÆRUM
Governance, Risk and Compliance Consultant
mnemonic

After reading this article, you will:

- Know what opportunities lie on the road towards GDPR compliance
- See how GDPR can lead to better internal control, information security and efficiency
- Pick up convincing arguments to make the case for sound GDPR compliance in your organisation

GDPR, a major buzzword in 2017, comes into force this year. This four-letter acronym has struck fear and panic into many organisations, forcing them to review how they approach personal data. Through this process, plenty of skeletons have found their way out of the closet and forced many organisations to face outdated systems and processes.

Despite the fact that privacy rights and laws have existed throughout Europe prior to GDPR, they have not necessarily been followed nor enforced. Otherwise professional companies have demanded, processed and stored personal data with minimum control and attention to privacy. There have been a multitude of cases where unauthorised personnel have been able to inadvertently (or intentionally) access sensitive information – typically as a result of poor technical design, access control, data management, or processes that prioritised business operations and efficiencies while lacking consideration of the implications on personal data.

There are several reasons to help explain how we got to this point. One is related to enforcement, or the lack thereof, and the absence of strong penalisation or consequences. Another contributing factor is the increase and the relative ease of outsourcing IT operations and development – unfortunately often without proper due diligence of the outsourcing partner's routines and control mechanisms. Unstructured organic organisational growth, both in size and operations, is also one of many contributing factors.

So, how many organisations will be able to comply with this new regulation? At the beginning of 2018, almost none. Studies from AvePoint, Forrester and Gartner predict that 50% - 70% of organisations will not be compliant by the time GDPR comes into force.

INTERNAL CONTROL

No matter how you approach GDPR, it calls for some common requirements for internal controls in the form of audits, procedures, and policies regarding the handling of personal information. This includes knowing where your data is, how it is kept safe, why you have it and whom has access to it.

If the desire for GDPR compliance is driven by the fear of fines alone, it can be tempting to merely perform the bare minimum required to check the appropriate boxes. This is unfortunate, as you risk missing out on a positive ripple effect of GDPR compliance – namely establishing real internal control.

GDPR presents a natural and valuable opportunity to review and/or establish internal control throughout the whole organisation, not only in regards to privacy. This is a healthy activity that will mature the entire organisation and improve its robustness. For the privacy advocates out there, GDPR has provided an excellent opportunity to get more focus on internal control from the CXO and executives.

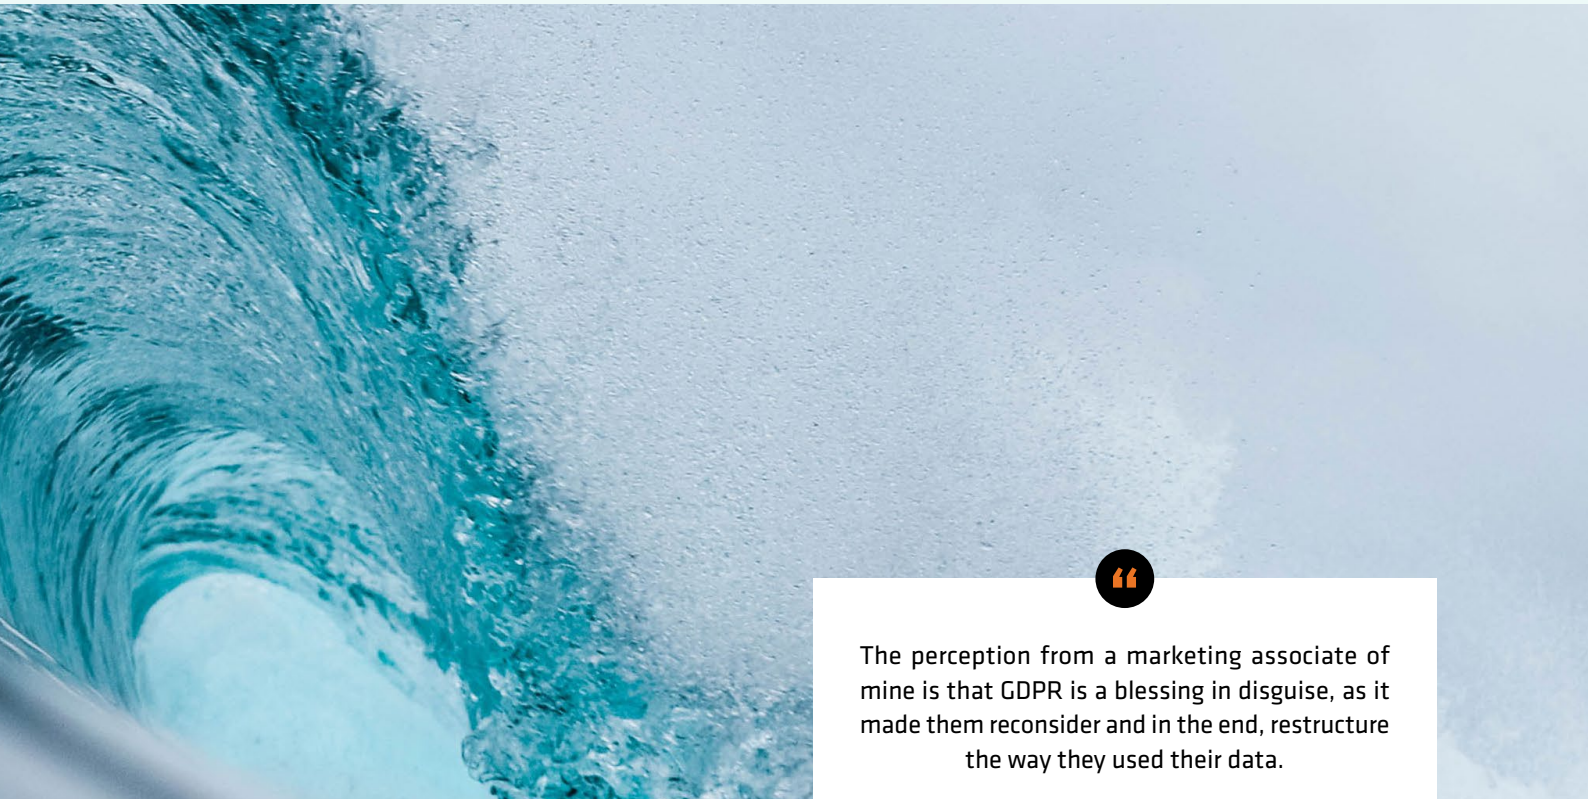

> The perception from a marketing associate of mine is that GDPR is a blessing in disguise, as it made them reconsider and in the end, restructure the way they used their data.

INFORMATION SECURITY

Article 32 of the GDPR states:

"[...] the controller and processor (editor's remark: companies handling personal information) shall implement appropriate technical and organizational measures to ensure a level of security appropriate to the risk".

This excerpt calls for sound information security and controls – an important and quite relevant provision. Unfortunately, the interpretation to what is defined as "appropriate technical and organizational measures" is subjective and will vary drastically from organisation to organisation, and person to person.

Information security has been in the proverbial backseat for far too many years. Fortunately though, information security has received a significant rise in attention lately, in part through the rise in public awareness of cybercriminal activity, such as ransomware and targeted attacks, and now through the GDPR. The fear of fines reaching up to 4% of global revenues in the event of a breach of the regulation certainly helps in raising awareness as well.

As with internal control, some organisations perform the bare minimum concerning information security by "only" safeguarding personal information, and do not understand the value of their internal business information and systems. However, other organisations use the process of getting GDPR compliant as an opportunity to implement measures that boost their overall information security posture. Because of Article 32's emphasis on security, there is reason to believe that organisations will overall become better equipped to deal with modern digital threats.

EFFICIENCY

The GDPR requires organisations to delete/anonymise all personal data it does not need or have a legal basis to process. This provides an opportunity to rethink how your organisation uses information, how to use it smarter, and how to get rid of vast amounts of data to reduce complexity and free up space, thus limiting maintenance, and licencing fees in the cloud environment.

To be able to comply with the new privacy rights given to European citizens, organisations will be forced to evaluate their existing systems for processing personal data. Should they find that their existing systems are simply not secure enough, do not provide proper overview, lack the ability to delete and export individuals' data, or are missing the capability to differentiate in data processing, organisations may quickly realise that their existing systems are not up to date, and new investments are required.

A positive ripple effect of replacing these systems in order to become GDPR compliant is that organisations will be able to handle data more efficiently. For example, the perception from a marketing associate of mine is that GDPR is a blessing in disguise, as it made them reconsider and in the end, restructure the way they used their data. As a result, it enabled them to make better use of their data, and the process of becoming GDPR compliant gave them a competitive advantage, by splitting up consents, using data in a smarter way and producing higher value to targeted customers.

CONSUMER PORTABILITY

The GDPR also introduces the right to data portability for consumers. This empowers consumers with the right to have their personal data exported in a common format. This provision may lead to lower levels of consumer loyalty, as individuals no longer risk losing their historic data when switching from one service provider to another.

Data portability may also give consumers more bargaining power when choosing service providers. By bringing their historic data with them, and hence proving to be a valuable customer, consumers can use this as leverage to negotiate better prices or terms. Therefore, organisations that are prepared to offer new customers a smooth transition when moving from another provider will have an advantage. For example, for industries such as banking, finance and insurance, there is potential to leverage portability to differentiate oneself in what can be seen as an otherwise rather homogenous industry.

GDPR is not simply a project with a set completion date. It will not be enough to achieve compliance on May 25th and forget about the requirements after this date.

BEYOND MAY 25TH

As we prepare for the GDPR to come into effect on May 25th, it is wise to remember that this is not the finish line. May 25th can be considered the official date when data privacy becomes a driving force. It is important to remember that GDPR is not simply a project with a set completion date. It will not be enough to achieve compliance on May 25th and forget about the requirements after this date. GDPR is a regulation, and consumers' perceptions of their data privacy rights are changing – neither of which will disappear.

The, for many, new way of thinking about privacy, and having control over data, policies, routines and documentation, has to live on. As people progressively share more personal data with organisations, privacy and consumer rights will demand larger public awareness. Europe is an attractive market for many, and as all companies handling personal data about EU/EEA citizens are bound by the GDPR, there is reason to believe that other countries and regions will follow the EU's initiative.

There are those hoping for a "soft implementation" of the GDPR in their own country, arguing that some governments have been lagging behind their European counterparts, and being late with guidelines, translation, hearings and adoption of the regulation. However, the GDPR aligns all the Data Protection Authorities in the EU/EEA so that no single country will be "softer" than others in their interpretation of the regulation, as was the case with the Data Protection Directive that preceded GDPR. Our expectation is that within a reasonably short period after May 25th, one or several companies in clear and severe violation of the regulation will receive significant fines and be made an example of.

As for operational business effects of the GDPR, it is hard to predict the future. However, one can be certain that organisations grabbing the opportunities found on the road towards GDPR compliance will gain an advantage over those who do not. The goal of this article has been to show that there is no need to fear the GDPR, as long as you properly prepare. By spending appropriate time and efforts on planning, implementation and follow up, organisations will be able to utilise the comparative advantages of the process – and not be left behind by their more prepared competitors.

WHAT IS GDPR?

The General Data Protection Regulation (GDPR) is a privacy regulation coming into force on May 25th 2018, and aims to give consumers control over their own personal data. It is a regulation that encompasses all of the EU/EEA, and aligns privacy laws and all Data Protection Authorities in the EU/EEA.

THE REGULATION IS BASED ON THE FOLLOWING PRINCIPLES:

- The processing of personal data shall be **lawful, fair and transparent**
- Personal data shall be collected for a **specific, explicit purpose,** and not used for other purposes
- The collection, storage and use of personal data shall be **minimised and limited to what is necessary**
- Personal data shall be **accurate**
- Personal data shall **only be kept for the time period for which it is necessary** for the purpose of the processing
- The **integrity and confidentiality** of personal data shall be ensured using appropriate technical and organisational measures
- The company that collects data (controller) shall be **responsible and able to demonstrate compliance**

BASED ON THESE PRINCIPLES, EU/EEA CITIZENS HAVE THE FOLLOWING RIGHTS:

- **Right of information:** One has the right to know why the data is needed, what is done to it, how long it is stored, where it is collected, with whom it is shared and how to complain
- **Right of access:** The right to know if an organisation has data concerning you, and if so, to access and see all of this information
- **Right to rectification:** The right to have one's data rectified and up to date
- **Right to erasure:** The right to have personal data deleted
- **Right to restriction of processing:** The right to restrict processing if the processing is unlawful or if the accuracy of the data is contested
- **Right to be notified:** The right to be notified when personal data is rectified, erased or processing is restricted
- **Right to data portability:** The right to have personal data exported in a common, structured and machine-readable format without hindrance
- **Right to object:** The right to object to processing and automated decision taking with legal or similarly significant consequences

The maximum penalty for a non-compliant organisation is a fine of 4% of annual global revenue or €20 000 000, whichever is higher.

2017: A VIEW FROM MNEMONIC'S SECURITY OPERATIONS CENTER

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

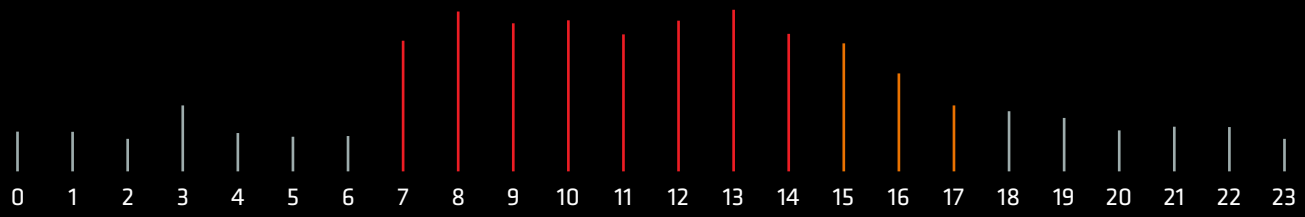

WHEN ARE SECURITY INCIDENTS HAPPENING?

Security incidents continue to occur 24-hours a day – this is no surprise. There is a noticeable increase during office hours, which continues to support the established truth that more user activity tends to lead to more security incidents – or in other words, users cause security incidents. Last year we observed a distinct peak of incidents between 12 – 13; a time when many users were presumably on or returning from lunch. Interestingly, in 2017 the volume of incidents are spread move evenly throughout the most common working hours of the day.

67%

Consistent with our observations in previous years, 67% of security incidents across all severity levels occur during the office hours of 07 – 18 on weekdays.

76%

76% of high and critical severity incidents occur during office hours. So not only are employees causing security incidents, they are responsible for causing more severe security incidents as well.

THE TALE OF TARGETED ATTACKS

Broadly speaking, there are two genres of attacks:

OPPORTUNISTIC

With opportunistic attacks, the victim is arbitrary, and exploited because they happen to be vulnerable. Call it being in the wrong place at the wrong time (and with the wrong vulnerability). This scenario is commonplace with ransomware infections.

TARGETED

Targeted attacks on the other hand are focused on exploiting a specific and chosen victim – be it a user or an organisation – to achieve a predetermined goal. Typically, the endgame here involves espionage, stealing corporate information, or scamming money through CEO fraud / Business Email Compromise. Targeted attacks are also the trademark of nation states, advanced threat actors and those engaged in offensive cyberwarfare.

33%

33% of all targeted attacks occurred during the lunch period of 11 – 13. Considering only 14% of all security incidents occur during the same timeframe, there is a notable concentration of targeted attacks at this time. A possible explanation may be that users are likely to be performing more personal activities during their lunch break and may have their guard down, or that attackers themselves believe that users will have their guard down and choose to attack at this time. Most likely though, it's a combination of the two.

96%

96% of the targeted attacks we observed occurred during the regular working hours of 07 – 18 on weekdays.

THE FALL OF RANSOMWARE?

Proportionate to the total security incidents reported, we observed a 61% decrease in ransomware cases from 2016 to 2017. However, this certainly does not insinuate that ransomware is on the decline. In 2016, there were a few particularly successful ransomware campaigns that led to a somewhat anomalous 431% increase from 2015. If we observe a longer trend, we see that in 2017 there was more than twice as many ransomware incidents than in 2015. The conclusion is that ransomware is on the rise, continues to be a nuisance and is a real security threat to organisations globally.

WHEN ARE USERS BEING INFECTED?

Users are most likely to be involved in a malware related incident earlier in the workweek. Our observations also show that statistically users were more frequently exposed to malicious code on Mondays, but successful malware infections were more prominent on Wednesdays. However, the difference is so marginal that it is likely more related to chance than any meaningful reasoning.

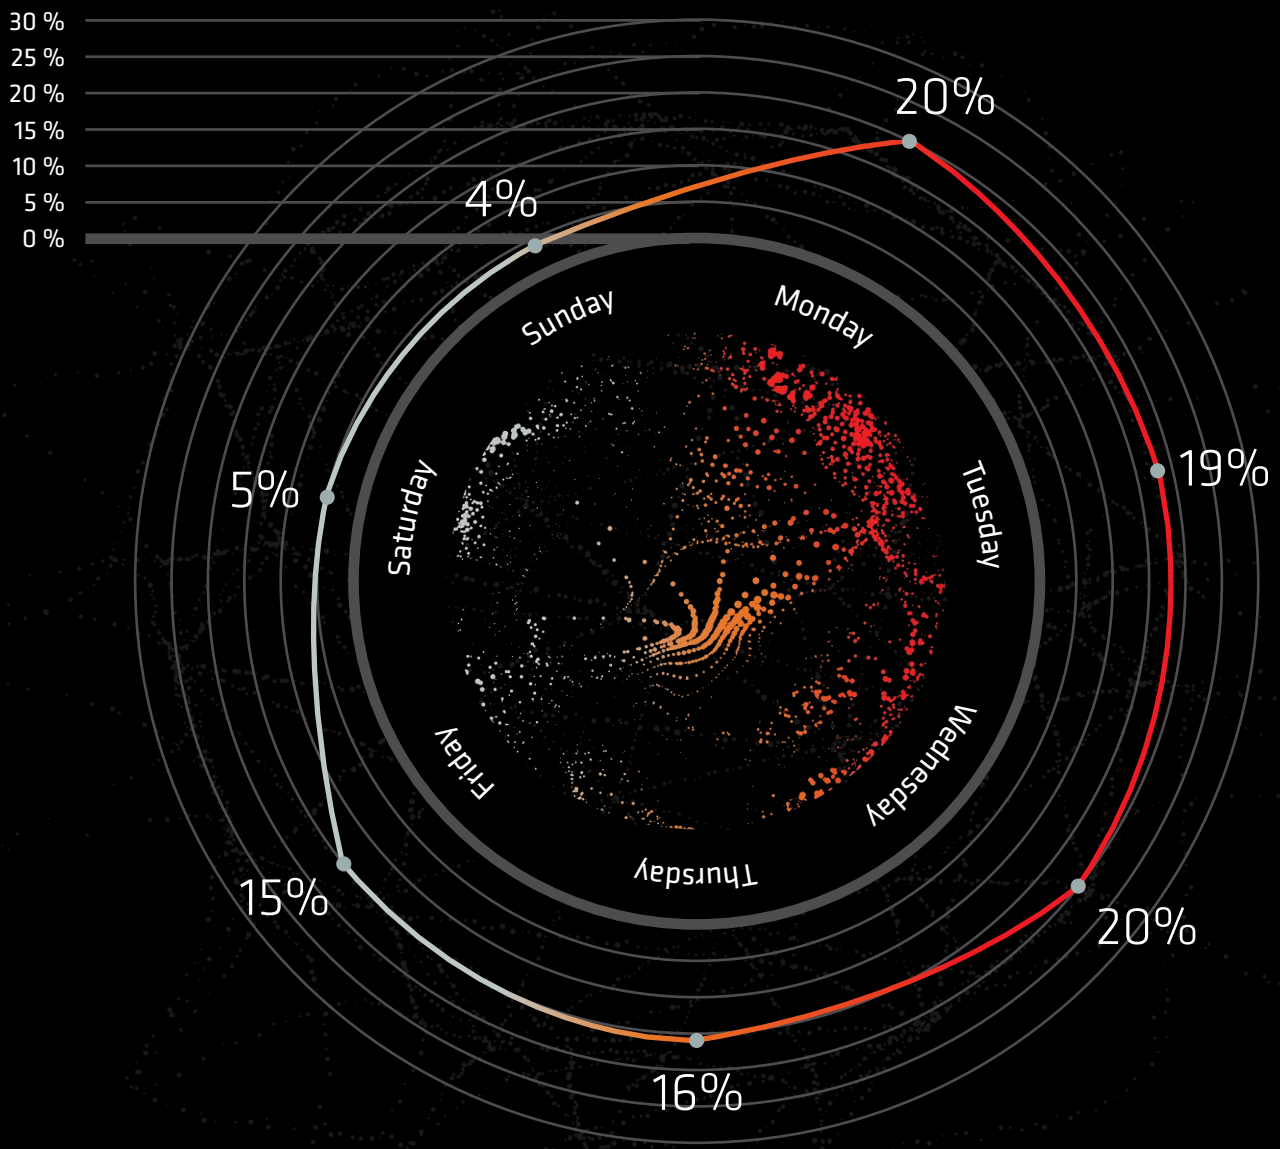

EMAIL FRAUD

HOW CAN WE PROTECT OURSELVES?

OLE KRISTIAN ROSVOLD
Security Infrastructure Consultant
mnemonic

JON-FINNGARD MOE
Department Manager Security Infrastructure
mnemonic

After reading this article, you will:

- Understand the basics behind common email fraud techniques
- See that there actually exist effective techniques against email fraud that go beyond awareness training
- Know how to implement protection techniques that work

We can love or hate email, but the facts are unambiguous. Even if email as a technology is antiquated, and based on a protocol defined in 1977, its usage continues to increase every year at a steady pace, with no signs of it slowing down[1].

Today, 90% of all sophisticated cyberattacks target people, not machines. The initial phase of these attacks commonly involves some kind of phishing and fraud attempt, and most often this is via email. While attacks often attempt to deliver malicious code, even more often they do not. With increasing usage of cloud-based services and social media, the attacks targeting people will remain the most popular way of compromising systems and corporations. The reason for this is simple - in an enterprise environment email is by far the most popular messaging platform, and people are the weakest link in the security chain.

Protection techniques and technologies to protect against email-based fraud are quite straightforward to use, and relatively cheap to implement. Still, enterprises are not using the most modern tools and updated techniques to protect email as a communication platform. Despite being the most common infection vector, only 8% of security budgets are spent on securing email. Meanwhile over 50% of budgets are used for traditional network protection, such as firewalls, IPSs, sandboxes, and so on. As a consequence, most enterprises are lagging behind both in protection technology and competence, and phishing and fraud through email remains effective.

In recent years organisations, CERTs and security vendors have increased their focus on email-based fraud. However, the focus has mostly been on awareness and user training. As an industry we have been trying to educate our users not to trust email as a platform for years, if not decades. We are repeating the mantras: "Be sceptic to the emails you receive, AND do not EVER open attachments". "Do not click on links!" "If your boss tells you to do something in an email, do not do it! Call your boss to double check". Awareness and training is a necessary and valuable tool, but never enough on its own. Regardless, fearmongering around email as a communication platform is not productive.

COMMONLY USED PHISHING TECHNIQUES (AND HOW TO MITIGATE THEM)

In our daily work, we observe a number of techniques to phish credentials, information or money transfers from victims through email fraud. Email-based attacks without sending malicious files or links to the victim are commonly referred to as "Zero payload" attacks. Threat actors often target financial departments in large organisations to perform a money transaction. The reconnaissance phase of the attack relies heavily on social engineering and exploiting people's trust and respect for authorities.

At first glance, technical mitigation measures seem futile. Hence, many organisations have instead invested in awareness training for their employees. While awareness is a productive and positive measure, it is not enough by itself. The pros and cons of awareness training are another topic to be discussed on another day, but for this article we will work on the assumption that technical controls are more accurate and harder to fool than the fallible human eye and mind. The following sections will illustrate that there actually is a lot that can be achieved by using simple rules and open technology.

[1] *Find the references at www.mnemonic.no/references-2018*

"PHONE BOOK" DICTIONARY OF KEY PERSONNEL

A common technique used by attackers is to spoof or namedrop executives in the organisation to exploit the business' processes. Many organisations use email to exchange payment requests and details internally on a daily basis. As an employee in the finance department, when your executive instructs you to perform an action, generally you do as requested.

Technically is it simple to create a dictionary from the organisation's phone book, and perform a search on the sender's name in every email message to look for a predefined set of key personnel that are more likely to be misused or spoofed in an email fraud attempt. Through this you will be able to identify if the sender's email address does not match known addresses from the dictionary. This is a good approach to sift out which messages to look further into.

Protection mechanism:

- Identify personnel in the organisation that have a higher risk for being spoofed. Dictionary search if name is misused in email

SCENARIO 1 – SPOOFING AN INTERNAL EMPLOYEE AND CHANGING THE "REPLY-TO" HEADER

This is a simple, but effective method. At first glance the email will appear legitimate to the user. The *From:* address is from within the organisation. Email clients like Outlook even provide a picture of the person. But when clicking "reply", another address (the reply-to address) is used for actually sending the email. Thus this technique tricks the user to start communication with a malicious third party.

Protection mechanisms:

- Use a SIEM/log analysis tool to flag emails with non-corresponding reply-to addresses
- Implement rules in the email security gateway to quarantine these emails for review by internal SOC/CSIRT

Example of spoofing an internal employee and changing the "reply-to" header

Header To: recipient@mycompany.com

Header From: sender@mycompany.com

Header Reply-To: hacker@fraudster.com

SCENARIO 2 – NAME SPOOFING IN *FROM:* ADDRESS

Another much used technique is to include the name of an internal employee or trusted third party in the *From:* header, but actually providing an address leading to a malicious third party. This is fairly easy to spot by the receiving user, as long as the email client is showing the full From: address. This might not always be the case on mobile devices and smartphone clients, where text-space is limited.

Protection mechanisms:

- Use a SIEM/log analysis tool to flag externally received emails containing names of internal key personnel
- Implement rules in the email security gateway to quarantine these emails for review by internal SOC/CSIRT

Example of name spoofing in From: address

Header To: recipient@partner.com

Header From: CEO in company
hacker@fraudster.com

SCENARIO 3 – DOUBLE ENTRY IN FROM: ADDRESS

This scenario differs slightly from Scenario 2, as the *From:* address contains two different addresses. One of the addresses is a legitimate address of a person the fraudster is trying to impersonate, however, the real from address, which emails are replied to, are leading to a malicious third party. This is even harder to spot for the human eye, as well as less probable to be shown in a mobile email client.

Protection mechanisms:

- Use a SIEM/log analysis tool to flag externally received emails containing names of internal key personnel
- Implement rules in the email security gateway to quarantine these emails for review by internal SOC/CSIRT

Example of double entry in From: address

Header To: recipient@mycompany.com

Header From: CEO in company
ceo@mycompany.com
hacker@fraudster.com

SCENARIO 4 – SPOOFING A PARTNER ORGANISATION AND CHANGING "REPLY-TO" HEADER

This scenario is similar to Scenario 1, but the fraudster is spoofing an address outside your organisation. Most often this is a trusted partner. This scenario is harder to prevent technically than Scenario 1 as prevention requires the partner to have a DMARC "reject"-policy in place.

Protection mechanisms:

- Use a SIEM/log analysis tool to flag emails with non-corresponding reply-to addresses
- Implement rules in the email security gateway to quarantine these emails for review by internal SOC/CSIRT
- Advise the partner organisation to implement email authentication (DMARC). Ensure DMARC validation and enforcement is active on your email security gateway

Example of spoofing a partner organisation and changing "reply-to" header

Header To: recipient@mycompany.com

Header From: sender@partner.com

Header Reply-To: hacker@fraudster.com

SCENARIO 5 – USING A "LOOKALIKE"-DOMAIN

This scenario requires the attacker to register a domain that looks similar to the company's internal domain, which is then used by the attacker to send emails from. This is referred to as *typosquatting*. Users are easily tricked by this as it requires a second look at the "from"-domain to detect the slight mismatch in writing.

Protection mechanisms:

- Create filters in a SIEM/log analysis system to detect phishing attempts
- Use the algorithm to calculate the domain variations most susceptible for phishing and implement manual block-rules for these domains in your email security solution. Using a tool like "dnstwist" might be useful[2]
- Register/buy the domains to prevent third parties from using them for malicious purposes

Example of using a "lookalike" domain

Header To: line@mnemonic.no

Header From: tonnes@mnemOnic.no

[2] *Find the references at* ***www.mnemonic.no/references-2018***

The Damerau-Levenshtein distance

Technical detection of lookalike domains is more accurate than the human eye, as we read known words as pictures rather than individual letters, while a computer sees the different characters as a unique set of bits.

The Damerau-Levenshtein distance is a metric for measuring the difference between two strings in terms of operations required to change one word into the other. Operations consist of insertions, deletions, substitutions and transpositions of single characters to transform the one string to the other. This is a brilliant metric to detect similarities among unique strings like domains in DNS.

Implementing the algorithm and using this technique to detect lookalike domains in Secure Email Gateways could enhance spoofing detection. At this point there are no known email security solutions officially supporting this technique, however it can be found in use by some managed security service providers.

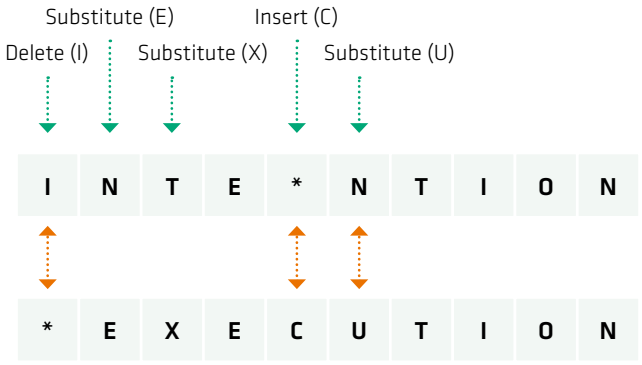

The Damerau-Levenshtein distance

DNS, DNSSEC AND DANE

Domain Name Service (DNS) is vulnerable to forged responses and other attack vectors that exploit lack of authentication and integrity of DNS records. The proposed mitigation is DNS Security Extensions (DNSSEC). By using DNSSEC, every record in each level in DNS or zone (including the root zone ".") is cryptographically signed. This means that a resolver can validate the authenticity of a record in the lookup answer, including the chain to the root zone, which it ultimately trusts.

DNSSEC not only ensures that requested domain names for websites or email servers resolve to the correct IP address, but more importantly, offers the ability to use DNS for trust relationships within security services. The fact that DNS is globally available, as well as commercially and politically independent makes this a potent system for managing and verifying trust relationships on the Internet.

Another initiative is DNS-Based Authentication of Named Entities (DANE). Organised as an IETF working group, the objective is to specify a set of mechanisms and techniques that allow Internet applications to establish cryptographically secured communications by using information made available in DNS. This has led to a standard for authenticating Transport Layer Security (TLS). By anchoring public key crypto certificates to their corresponding domain name in DNS, validation of authenticity and propagation of trust is possible in a timely manner across the world.

An important security measure for enterprise email systems is *enforced TLS*. This mechanism requires TLS secured email delivery between the organisation and its defined partners. To ensure a certain minimum level of security for each partner some parameters are possible to enforce for the encryption and certificate validation of the connection.

Implementing and enforcing DNSSEC, DANE and TLS for email exchange is a huge step in the direction of secure communication between organisations. However, DNSSEC is dependent of support and management from DNS providers and resolvers. Not all provide this by default, which is one reason for DNSSEC's limited deployment. The situation for DANE is even more deceptive. The support in enterprise email solutions does not exist and there is no planned future support in major web browsers.

EMAIL AUTHENTICATION: A KNIGHT IN SHINING ARMOUR?

If we take a look at the scenarios in the last chapter, several could be completely prevented if email authentication had been implemented at both the sending and receiving parties (of course this does not apply if the email systems are compromised, but that is another topic). A perfect implementation of email authentication technologies would ideally protect an organisation from unauthorised third parties sending emails on the organisation's behalf. Another important use case supported by implementing email authentication is to maintain internal policies of authorised third party senders.

However, even if your company has implemented the technology to perfection, that is only half the battle and you are still dependent on the third parties you are communicating with for effective enforcement to be in place.

Protocols, standard and tools have been available for some time, but have not reached broad popularity. Due to the massive increase of phishing and message spoofing, the industry as a whole have finally understood the importance of getting email authentication techniques in place.

EMAIL AUTHENTICATION TECHNIQUES – AND THEIR SHORTCOMINGS

Initiatives for providing email authentication techniques have roots as far back as 2000. The work has so far resulted in three standards that have reached a certain level of adoption - SPF, DKIM and DMARC. In addition, a new standard – ARC - is in draft. It is not difficult to implement the technologies, but you should be aware of the most common limitations.

Sender Policy Framework (SPF)

This standard is used to define which IP-addresses are legitimate senders of a domain, and sets a policy for the receiving system on what to do with emails coming from other IP addresses. The implementation is straightforward and results in a simple string published in DNS.

Limitations and considerations:

- SPF only authenticates the "envelope-from" (RFC5321) field in the communication. This is represented as a header in the email. This header is not visible in the email client for the recipient. This makes SPF very easy to circumvent, and provides limited value alone for stopping phishing emails.

The envelope-from address (RFC5321) is the address provided in the envelope of an email message. It is also called the return-path address. Email clients strip the envelope for the recipient, only showing the MailFrom address (RFC5322).

- The only policy that provides any real protection when using SPF is "HardFail". We see great variation in the policies in use, and confusion amongst our customers.

SPF-policies can be either Pass, Neutral, SoftFail or HardFail. Only the HardFail-policy instructs the receiving gateway to discard messages that fail SPF-authentication.

- Many organisations depend on third parties for mass email and marketing services that send email directly from web or SaaS applications. In order to use these systems, organisations need to include the third party's SPF record in their own record. The quality of the third party records will vary, and the SPF-standard limits the number of DNS-lookups to 10 per record. Poor knowledge and practises result in faulty records and syntax errors which again of course, offer no protection, just complexity and problems.
- The configuration requires maintenance. New senders, either internal or third party, need to be added to the SPF-record. SPF provides no features to detect legitimate emails rejected by the receiving party (this is where DMARC comes into play).

Domain Keys Identified Email (DKIM)

DKIM is also a very simple and straightforward technique. The main goal of DKIM is to provide the receivers with emails that have cryptographically signed messages. It leverages PKI and provides the ability to verify signatures by using the sender's public key published in a DNS record.

Limitations and considerations:

- Office365/Exchange Online by default provide a DKIM-key for domains in use by the SaaS-suite. Exchange Online will sign outgoing emails by <domain>.onmicrosoft.com. This is positive, as emails are provided by verifiable signatures. However, the signature does not align with the actual domain we are sending the email domain from. Thus, messages are failing authentication. This is easily solved within Exchange Online, however not everyone is aware of this, or take the time to implement the necessary change.
- Third party services include mass mailing services, SaaS applications and more. These third party providers must be able to sign with DKIM signatures on the domain owner's behalf, however not all providers offer this option. In addition, if they do, the service's public keys are shared between all subscribers of the service, which is not always desirable.
- As an authentication technique alone, DKIM provides limited value. As the messages are prone to arbitrary forwarding and footers from mailing lists and antivirus scanners.

Domain-based Message Authentication, Reporting and conformance (DMARC)

DMARC is currently, when implemented to perfection, the only technique that provides real protection against email spoofing and modern phishing methods.

DMARC can be used to stop emails with spoofed *From:* fields. This is achieved by a so called "alignment" check between the "envelope-from" address (RFC5321.MailFrom) in the email header and the *From:* address (RFC5322.From). Or an alignment-check between the *From:* address and the signing domain used in DKIM. It can also use the authentication results from both SPF and DKIM to determine if an email appears to be spoofed or not. Only SPF *or* DKIM is necessary in order pass DMARC authentication. This is useful to know when using third-party email services that do not support both.

DMARC can also provide reports. These reports contain detailed information and statistics on any malicious third parties that are misusing your domains to send phishing, fraud and malware. The reports also include information about legitimate domains failing authentication. Some forensic reports may even contain examples of the emails failing authentication. As far as we have observed, as of the end of 2017, Outlook.com/Hotmail is the only major service providing these reports.

Limitations and considerations:

- In order for DMARC to reach its potential, it requires either SPF or DKIM to work perfectly for all protected domains in an organisation. Ideally, both SPF and DKIM in combination. To get this alignment right, it is important to get the SPF-definition and policy up to speed.

- DMARC "reject" polices take time to implement, especially for large organisations with many domains and third party senders.

DMARC policies may be either: None, Quarantine or Reject. Both the None and Quarantine policies are used in the initial phases of DMARC implementation and offer limited protection. Reject is the only policy instructing the recipient system to discard messages failing DMARC authentication.

- Errors in DMARC policies may have a large impact, and cause legitimate emails to be discarded.

- DMARC reports (aggregate reports) are currently delivered by the largest email providers on the Internet (e.g. Facebook, Microsoft, Yahoo!, etc.). Corporations and governmental bodies that use their own email security gateways seldom send failure reports to domain owners.

- DMARC has limited support in enterprise email security gateways for granular enforcement policies, and even less support producing DMARC reports.

- Analysis of DMARC reports require third party tools to be efficient, however there are not many effective tools currently available, and those that are can be quite expensive.

IS EMAIL AUTHENTICATION WORTH IT?

Yes.

Although email authentication techniques may appear deceptively simple, establishing an effective and maintainable policy requires competence, time and the necessary tools.

The *garbage in – garbage out* principle applies to email authentication techniques. It is not a fire and forget job. A half-implemented, half-heartedly maintained policy will not work, and in some cases may leave an organisation worse off than not having email authentication at all.

The simplified flowchart to the right shows the steps organisations need to follow to develop a functional DMARC policy.The process may take anywhere from a few days to several months depending on how many domains the organisation manages, the amount of legitimate third-party senders, the DMARC analyser tool, gateway products, outsourcing strategies, etc.

External competence and tools may be required for larger or complex organisations.

EMAIL FRAUD – WE CAN PROTECT OURSELVES

There is no ultimate fix for email security, however effective protection techniques and products that go beyond awareness and user training exist.

By implementing these techniques, email can become a less effective attack vector for phishing and fraud attempts, and continue to have a place in the enterprise environment.

EMAIL AUTHENTICATION
MAIN STEPS TUTORIAL

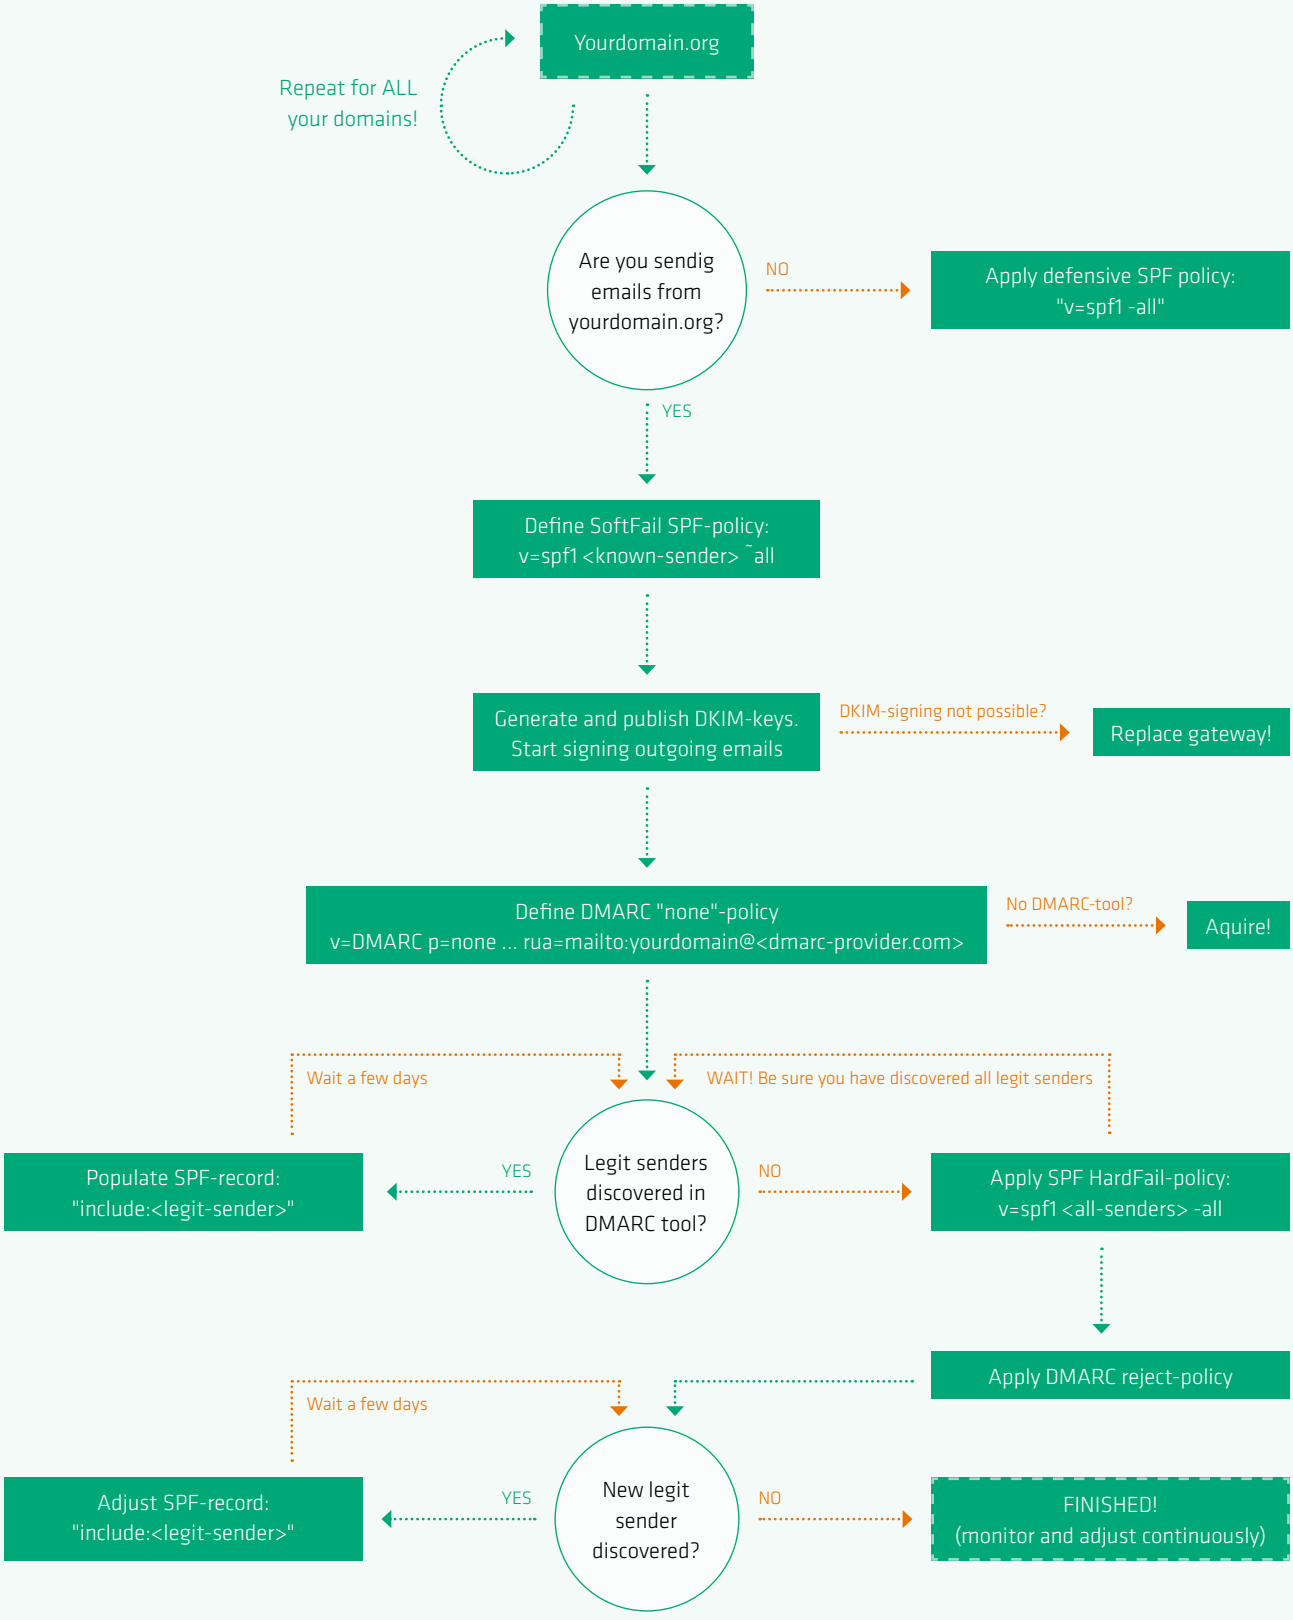

IDENTITY

IS IT THE NEW SECURITY PERIMETER?

SIMEN E. SANDBERG
Senior Consultant
mnemonic

After reading this, you will:

- Understand the basics behind penetration testing
- See how the cloud is complicating, and changing, the traditional risk management matrix
- Know what "old" security measures are still valid in the age of the cloud

I am a penetration tester. As a penetration tester, my job is to look for vulnerabilities in my clients' systems that enable me to perform actions not originally intended. These vulnerabilities may impact the stability and availability of a system, while others may jeopardise the confidentiality and integrity of the data within the system – all of which can pose a risk for the client.

Managing this risk is a process with several steps. During my time as a penetration tester, I have noticed that while the fundamental steps remain the same, some of the technical outcomes of the risk management process are changing. The changes come due to new technologies not only giving us opportunities to manage information in easier and more flexible ways, but also giving threat actors new avenues of attack.

One of the things that has *not* changed is the primitive question "what should be protected?". Commonly this is information that is seen as confidential, such as strategic plans, product designs or financials. Recently, with the GDPR, personal information about customers and employees has moved up to the top of that list.

Penetration testers are often part of brainstorm sessions to imagine and identify the different attack vectors a threat agent may use to gain access to critical information. Listing such scenarios is a step in the risk management process. To measure the risk this unauthorised access to information assets poses, we must assess both the probability of the scenario occurring, along with the technical feasibility of the scenario being successfully executed.

For example, a quite probable scenario is assessing what information an attacker can gain from a web application without a validated user name or password. Likewise, a less probable scenario, but perhaps as equally important, is if an attacker gains physical access to you data centre. Our job is to then play the role of an attacker and test if we are able to fulfil the scenario and gain access to critical information assets. While this may lean towards a more risk-centric penetration test, this process involves the same fundamental premise as all penetration tests.

THE TRADITIONAL PENETRATION TEST

A common scenario for a penetration test is that an employee receives a phishing email, and opens an attachment or clicks a link. In this scenario, the user will often disregard any dialogue boxes, ignore any warnings and simply click "OK" to whatever pops up on their screen. Our experience shows that this action is all too likely in most organisations. Should we succeed, we will now have some or full control over this user's device, account or session. We have now breached the perimeter security, are on the inside of the network and can start attacking internal services. Penetration complete.

We continue to steal passwords and move laterally until we find a computer with access to the information we want, or an administrator account that can give us the rights to grant ourselves the access we need.

Our next step is often to attempt what is known as "lateral movement", which is effectively when an attacker or penetration tester moves from one device to the next in search of their ultimate goal or target. One such technique for lateral movement is to steal, and subsequently use legitimate credentials to move throughout a network. We continue to steal passwords and move laterally until we find a computer with access to the information we want, or an administrator account that can give us the rights to grant ourselves the access we need.

The outcome of a penetration test like this is often general recommendations, such as:

- Make direct attacks on internal systems harder with patch management, security configuration and network segregation.
- Limit lateral movement with endpoint firewalls, tools like Microsoft's free Local Administrator Password Solution and (again) security configuration and network segregation.
- Shield administrative users from lateral movement with tools and procedures based on "least privilege", "just-in-time" elevation of privileges and separate, offline administrative workstations.

PENETRATION TESTS IN THE CLOUD

When preforming penetration tests in a cloud environment, these three recommendations are still valid, and we give them all the time. However, if the critical information asset is located in a Software as a Service (SaaS) environment, like Salesforce, OneDrive or Box, then these requirements need to be communicated to and enforced by the cloud provider. In addition, users are more likely accessing these systems and services from an unmanaged, personal computer, meaning it is outside the perimeter of the organisation's security controls. In either case, the ownership of the systems that require change are outside of an organisation's control, and therefore the scenarios to be tested, remediation actions, and risk measurements must be taken back to the drawing board.

I have a sticker on my computer: "There is no cloud – it's just someone else's computer".

It points to the fact that the responsibility for security of cloud services is shared between the cloud vendors and the customer. The different cloud vendors have various matrices describing who, between the cloud vendor and customer, is responsible for what. What they fail to mention is that the customer is responsible for gathering enough information about the cloud vendor and performing adequate due diligence to ultimately trust that the cloud vendor is both competent enough to, and actually is enforcing the security controls that are listed as their responsibility.

Similarly, it is the customer's responsibility to make themselves aware of, understand and implement the built-in security controls provided by the cloud vendors. This is a topic that could fill its own article, but if we instead jump straight to the conclusion: the general rule is that larger global cloud vendors are most often better at managing security than the rest of us.

Figure based on Microsoft Azure's shared responsibility model

Figure based on Amazon Web Services' shared responsibility model

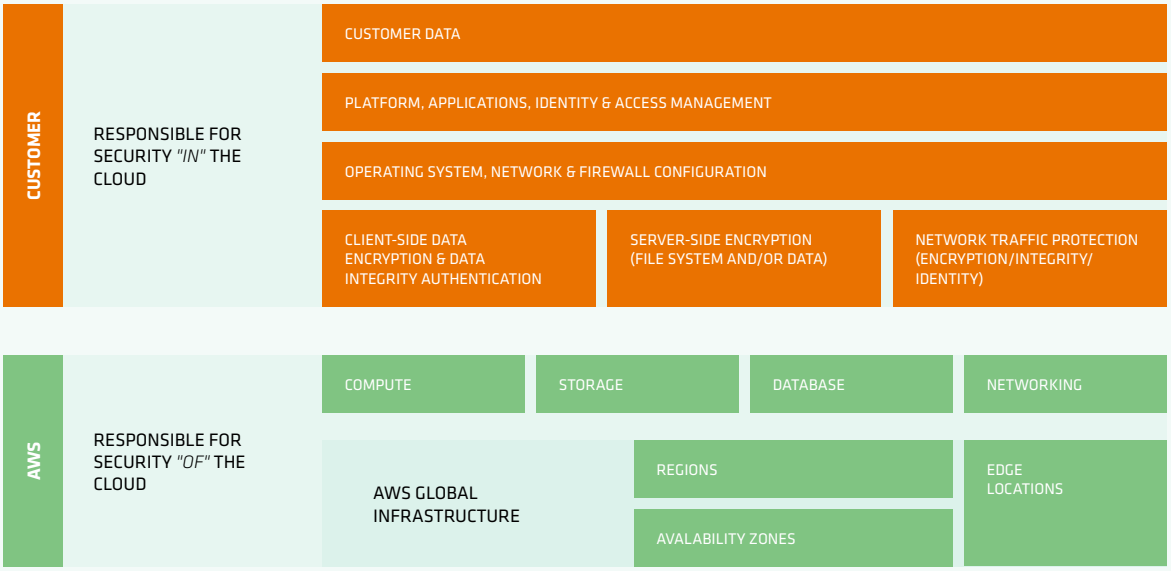

If we decide to trust the vendor, we can see from the figures that they only take complete responsibility for things like physical security, or what Amazon is calling "of" the cloud. The customer is still responsible for what is "in" the cloud, including deciding who should have access to what services.

How can a penetration tester use this division of responsibility to gain access to information assets protected by the most security-conscious, professional and financially strong cloud vendors in the world?

Simple. Walk in the front door disguised as the customer's own users.

Because more users are accessing cloud services from unmanaged devices a new attack vector and risk is introduced.

IDENTITY: YOUR PASSPORT TO INFORMATION

One of the main benefits from cloud services is that they are available from anywhere on the Internet. As long as you are authenticated to the service and have the necessary access rights, you have access to any information from any location.

Naturally, getting access to user credentials is key for a penetration tester: if we have the valid password for a user who is authorised to access sensitive information assets, we can just log in to the cloud service and download whatever the user has access to.

The good (or bad) news is that finding these passwords is what we have been doing all along. Only now, we do not have to worry as much about getting on the inside of the customer's network and the risk of detection that entails.

We can still use phishing attacks, and because most organisations using cloud services use some kind of password synchronisation between their on-premise network and the cloud, lateral movements from the cloud back to the on-premise network often work just as well as before. And even if the organisation isn't synchronising passwords this way, users tend to re-use passwords anyway. If we get on the inside of the network, we can still extract passwords and use them to log on to cloud services and vice-versa.

Because more users are accessing cloud services from unmanaged devices (that is, devices not managed by corporate IT, such as home PCs, or personal mobiles and tablets), a new attack vector and risk is introduced. Not only do unmanaged devices typically have weaker security controls than managed devices, but they are likely also being used by non-corporate users like family and friends who have no understanding or care for your IT usage policies. The unfortunate reality here is that this increases the risk of corporate passwords becoming compromised.

THE NEW RISK MANAGEMENT MATRIX

So we now have several new cloud scenarios to include in our risk management matrices: attacking vendor's networks, stealing passwords to cloud services, and access from unmanaged devices. So what can a penetration tester, as opposed to threat agents, do to help? What new recommendations are we giving?

AUTHENTICATION

The most straightforward recommendation is to never allow authentication to cloud services using only a (possibly stolen) password. In other words turn on multi-factor authentication. All the reputable cloud vendors have options for this, and it is generally included in the cheapest plans.

Enabling multi-factor authentication for a plethora of different cloud services can wear out the phone-PIN-entering thumb of the most understanding user. Various vendors offer authentication brokerage services to avoid that. You just log on to one service, and that service will automatically log you on to other services.

Authentication brokers often provide dashboards and advanced reporting of the usage of all the cloud services in the organisation. Some are even able to integrate with various services to provide detailed information of not only who logged into which service, but also list or control the information assets those users accessed or should be able to access.

Using advanced features like this can be very useful for security. Authentication brokers may be able to track sensitive documents between cloud vendors, and even delete documents that are shared where they shouldn't be. Collecting and retaining such logs are useful not only for audit purposes, but, and perhaps more importantly, when responding to a security incident.

Other advanced features are risk-based authentication, where the authentication broker requires more authentication factors, or steps, dependent on the risk of the activity being performed or user behaviour. So for example, an additional password challenge may be presented to a user when accessing the most valuable information assets, or if a user is accessing the cloud services from new or uncommon locations. A user whose location suddenly moves from Oslo to Shanghai in one hour should not be able to log in from a completely unknown device without additional authentication scrutiny, if at all.

> In the cloud, identity is key, and you'd better believe there's a line-up of people waiting to get in.

ADMINISTRATIVE ACCESS

Similar to on-premise and traditional solutions, administrators for cloud services will also have varying capabilities to create new users, adjust configuration settings, and access a slew of other functionalities. Therefore, the credentials you use to log in to the administrative console are the most attractive to threat agents and penetration testers alike.

Logging in to cloud services as an administrator generally provides some level of access to logs that can (and should) be used to monitor the usage of cloud services. The information contained in these logs and their usefulness, along with the configurability and options to export these logs will vary between cloud vendors, and also between subscription levels at any given cloud vendor. Nonetheless, such logs represent the fingerprints and breadcrumbs of user activity, and are essential to detect and respond to unwanted access. However, a critical and often unlooked step is to actually enable and configure these logging functions in the first place.

Also, be aware and mindful if you use directory synchronisation between your on-premise and cloud environment(s). Changes made by or to on-premise administrators can suddenly affect things like passwords and group memberships in the cloud.

Thus, I will add these recommendations for installations involving the cloud:

- Protect on-premise administrative accounts and maintain control over how your most critical credentials are used – don't give attackers unnecessary opportunities to gain domain control

- Use separate cloud administrative accounts for high-risk services like authentication brokers

- Access administrative accounts from trusted systems and utilise additional security controls, such as multi-factor authentication, VPNs and/or access control lists (ACL)

THE FUNDAMENTALS REMAIN THE SAME

Risk management for the cloud is not fundamentally different from risk management in on-premise networks. You still have to identify what to protect, understand the risks to those assets, test how vulnerable you are and use appropriate protection measures.

However, a new challenge is introduced because you not only require permission from the cloud provider to preform penetration tests in the first place, but the cloud provider may even deny you such permission, forcing you to evaluate this risk in a new way. Because the cloud breaks the traditional security perimeter, it encourages an increased focus on detecting and preventing credential theft. Remember: in the cloud, identity is key, and you'd better believe there's a line-up of people waiting to get in.

Iceland

VALITOR

DR. REY LECLERC SVEINSSON
Chief Information Security Officer

Valitor is an international payment solutions company. It helps partners, merchants and consumers to make and receive payments. Leveraging thirty-four years of experience, Valitor provides issuing, acquiring and gateway services to partners and merchants across Europe. Valitor improves cash flows through next day settlements, increases financial control through real time reporting and analytics, as well as reduces risk of fraud. Valitor's headquarter is in Iceland with a strong presence in London, UK and Copenhagen, Denmark.

WHAT IS YOUR BIGGEST SECURITY CONCERN?

As smartphones are becoming the preferred source of authentication for many financial transactions, malware authors will increase their efforts to steal funds from consumers' Apple Pay, Google Wallet and other mobile payment systems.

Once attackers have learned to infiltrate consumers' mobile wallets, they may tap into your corporate networks through those smartphones. Emails, contacts, authentication measures and apps that access the corporate network from the phone can become a phenomenal source of intellectual property, insider information and other confidential business materials. This way, they can become easily obtainable and can give an attacker sizable gains.

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

With the dramatic increase in the number of breaches and the rapid spread of cybercrime, the pressures for corporate action and further regulation continue to mount. The catalyst for this change has been an environmental one: we used to process all the information on our own computers, in our own building and within our own controlled information ecosystem.

The arrival of the Internet and the changes in the IT environments impact two major realities: first, it is becoming ever more difficult to control the corporate IT environment; and second, the road to greater regulation is rapidly taking shape. While I can argue that regulatory mandates help toe the line for large and small corporations, it is becoming vastly overcomplicated, especially for small organizations that lack resources.

This year alone we are seeing several key regulations that need addressing: Payment Systems Directive (PSD2), General Data Protection Regulation (GDPR), Network and Information Security (NIS) Directive and the e-Privacy Regulation to name a few. Even SWIFT has launched the Customer Security Programme (CSP) that provides a cyber-security requirement framework.

We are drowning in regulatory compliance requirements. The fact is that small businesses are being choked by excessive compliance regulations and large, global firms are forced to increase resources to comply with regulations. Many of these complex regulations are redundant, with each placing a different spin on its meaning and wording.

WHAT GIVES YOU HOPE FOR THE FUTURE OF SECURITY?

There's a growing realization that cybersecurity requires budgetary commitment, sincere collaboration, and a solid strategy management. If enterprises can pull together, with the right expertise, we can build a bright future that's secure from cybercriminals. Companies are growing more aware of threats, and this is leading to a greater allocation of resources.

Of course there is also increased awareness. More businesses are starting to understand the value in educating their own workforces on security. Establishing programs to ensure that staff are aware of vulnerabilities and the potential for cyber-attacks is important. Companies can leverage much greater value from existing security systems and policies by teaching staff good habits, and it's also important that they understand the potential impact of a breach.

THE IMPORTANCE OF SECURITY RESEARCH

MARTIN EIAN, PH.D
Senior Security Analyst, Threat Intelligence
mnemonic

SIRI BROMANDER
Ph.D. candidate and Threat Intelligence Analyst
mnemonic

After reading this article, you will:

- See what benefits organisations can get from investing in security research
- Understand the importance of public-private collaboration in security research
- Gain a general overview of some of the relevant funding and collaboration opportunities made possible by the Research Council of Norway

Successfully detecting and defending against advanced threat actors requires cutting-edge technical, tactical, strategic and operational security. To stay ahead takes experience, competence and a lot of time spent on research. Through our experiences from security research, we have seen the benefits of combining real life experiences from both private and public sectors with the academic excellence of educational institutions. The Research Council of Norway is an essential part in facilitating these bridges between relevant players in the security community.

WHY IS SECURITY RESEARCH IMPORTANT?

Private organisations must have convincing and sound argumentation to justify why its employees should spend their time, and therefore company resources, on research. mnemonic, as a private company ourselves, use a substantial amount of resources supporting research. Some of our justification is as follows:

PARTNERS

Sharing knowledge in open research projects connects you to other knowledgeable individuals and organisations in a trustworthy manner. By attracting quality collaboration partners through research, we gain access to valuable experience and knowledge. From our existing partnerships, we receive knowledge from a diverse array of industries, academic institutions, governmental organisations and international law enforcement. This knowledge is absorbed not only in our research projects, but throughout the entire organisation.

RECRUITMENT

The best minds are no longer just looking for high salaries and great benefits, but for the best environments to challenge themselves, expand their ideas, and work with equally passionate and likeminded individuals. Knowing that there is a global cybersecurity skill shortage, we have long held a staffing strategy that basically acknowledges that we must build and develop our own cybersecurity experts. Our experience has shown that having a formal strategy towards research helps attract and retain the best talent.

EFFICIENCY AND IMPROVED ANALYSIS

New methods to perform threat analysis make us more efficient. Finding new ways to analyse the ever increasing amounts of data we are facing is not only making us more efficient, but also enables us to better protect our customers. An example of this is the new methods developed for analysing PassiveDNS data (that is, the historical relationship between IP addresses and their associated domain names). From research conducted together with our partner the Norwegian Computing Centre, we have identified new sinkholes and suggested machine learning models for identifying malicious domains from those that are benign.

Creating new tools based on research can reduce the need for an analyst to spend time on repeating manual tasks, or having to dig through several systems to find relevant information. Consequently, we see our analysts spending more time performing analytical tasks, using new systems to free up resources and gain valuable insights from large amounts of data.

STRENGTHENING YOUR SIDE

Our industry has a somewhat clear "good side/bad side" divide. Choosing to protect our society, democracy and citizens is an ethical stance. Our research, in combination with our practical experience from most sectors give us a unique position to culminate relevant knowledge strengthening our side of the divide.

WHAT MAKES SECURITY RESEARCH POSSIBLE?

First and foremost, security research requires cooperation. Without cooperation, research would be conducted in isolated bubbles - every person for themselves. The sharing of findings, experiences, methods, mistakes and knowledge in general is what allows us to make rapid progress.

Let's take Norway – our primary hub for research – as an example. The Research Council of Norway is a government agency that builds bridges between the public and private sector and academia by using cooperation as a criterion for funding. This incentivises different agents to work together, and returns useful results and benefits for all parties involved.

The Research Council of Norway has several ways of contributing to research. Research is costly and by contributing with funding and facilitating cross-sector cooperation, their initiatives enable large-scale research with complementary results. Here are a few examples of initiatives that we have experiences with:

SKATTEFUNN

The SkatteFUNN R&D tax incentive scheme is a government program designed to stimulate research and development in Norwegian trade and industry. Businesses and enterprises that are subject to taxation in Norway are eligible to apply for tax relief.

mnemonic began working with SkatteFUNN in 2009. The program has contributed to extending R&D in mnemonic with several positions and the launch of a professional R&D department. SkatteFUNN has been an entryway into the Research Council of Norway and has provided us with experience and motivation to prioritize research and innovation.

USER-DRIVEN RESEARCH BASED INNOVATION (BIA)

The projects must result in substantial value creation for the companies as well as for society-at-large, and must take an international perspective. The projects are organised in consortia whereby companies and R&D communities cooperate on achieving results.

In 2016, mnemonic launched the Semi-Automated Cyber Threat Intelligence (ACT) project. The project aims to create an open source platform for detection and defence against cyberthreats. The project benefits from cooperation with several project partners from both public and private sector as well as academia through BIA.

The collaboration requirements in BIA give other organisations the opportunity to participate without having to run their own project. The feedback from project participants has been positive, and the experiences from our partners are immensely important for the results. Through this incentive, a considerable part of the security community in Norway has been able to review and evaluate the results of ACT.

THE INDUSTRIAL PH.D. PROGRAM

Under the Industrial Ph.D. scheme companies may apply for support for a three-year period for an employee seeking to pursue an ordinary doctoral degree. The doctoral candidate must be employed by the company and the doctoral research project must be of clear relevance to the company's activities.

mnemonic participates in a research project complementary to ACT called Threat Ontologies for Cyber Security Analytics (TOCSA). The criteria for funding is the same as other similar industry PhD programs; initiation and funding by a private company, participation by an academic institution, and the Ph.D. candidate dividing his or her time between the private company and the academic institution. The program provides all involved with new competences, and creates a bridge between private sector and academia that works well in practice.

IKTPLUSS

The primary objective of the IKTPLUSS initiative is to enhance quality, promote boldness in thinking and increase the relevance of Norwegian ICT research by linking R&D investments to national frameworks and needs for ICT research and innovation.

The IKTPLUSS initiative is the Research Council of Norway's large-scale initiative on information technology and digital innovation. mnemonic participates in two IKTPLUSS projects: Ars Forensica, run by the Center for Cyber and Information Security (CCIS) and the National Police Directorate, and Oslo Analytics, run by the University of Oslo.

> There are significant, real-world effects from these efforts that are being leveraged by not only mnemonic, but the security community around the globe.

TANGIBLE EFFECTS OF SECURITY RESEARCH

Our results from the funding, initiatives and collaborations mentioned above include the actual development of threat intelligence and security analytics platforms, and models and methods for use in these platforms. In other words, there are significant, real-world effects from these efforts that are being leveraged by not only mnemonic, but the security community around the globe.

While earlier we listed some of the less tangible, but equally important justifications for our investments into research, we thought it also important to share a few examples of the tangible, real-world impact this research is having.

THE ACT PLATFORM

Through the ACT project, we have successfully built, from the ground up, a now functional threat intelligence platform. The platform is in an early beta form to be used for testing and the basis for further development. An image containing the platform and the possibility to bootstrap some initial data is freely available for testing and has been distributed to project partners and other interested parties. The core platform's source code is also published on mnemonic's GitHub.

PASSIVEDNS

mnemonic has an open PassiveDNS project (https://passivedns.mnemonic.no) that collects data on the historical relationship between IP addresses and domain names. Using this data, which sees over 125 million operations daily, the Oslo Analytics project has used machine learning techniques to identify unknown sinkholes and malicious domains.

STANDARDIZATION OF CYBERSECURITY ONTOLOGIES

At the 29th Annual FIRST Conference mnemonic and Trend Micro organised a Birds of a Feather (BoF) session on "Ontologies". Ontologies is a formal way of representing knowledge, and includes descriptions of types, properties and relationships within a given domain. At the session mnemonic contributed with findings from TOCSA and ACT. The next step is establishing a FIRST Special Interest Group (SIG) for the standardization of cybersecurity ontologies. Afterwards, once the standard is mature, an IETF (Internet Engineering Task Force) standardisation initiative may follow.

THE ARGUS PLATFORM

Since its inception in 2002, Argus – our purpose-built security analytics platform – has had an incredible journey. With the help of SkatteFUNN, continuous development and years of application against real-world threats, the solution has grown to analyse over 6.5 billion events every day, with this number doubling year-over-year. Oslo Analytics, ACT and TOCSA have led to the development of models and new analytics capabilities for incident response. Among other things, the research has identified links between threat actors in incidents, and enabled low-level indicators to be analysed in a high-level manner – an effective threat detection technique that simply would have not been possible without applying the learnings of these research efforts.

FORUMS FOR SHARED KNOWLEDGE

The research projects mentioned have opened doors for our researchers and led to further cooperation in the security field. Because of our researchers' findings, we have for instance joined Europol EC3 (European Cybercrime Centre) as part of the Advisory Group on Internet Security. Such cooperation and affiliations have furthered our participation in leading security forums, such as Interpol, CERT-EU, FIRST and international universities, all of whom have also invited our project group to present and discuss our research.

A SPECIAL THANKS

With this article, we hope to have demonstrated just some of the many benefits brought about from engaging in security research. Whether through collaborations or by initiating one's own project, security research, collaboration and partnerships are an important contributor to secure our digital society.

We are thankful to all our partners, and want to give a special thanks to the Research Council of Norway for making it possible for researchers from Norway to contribute alongside other international cybersecurity researchers. We hope this article will encourage others to do the same.

Recommended reading

"Semantic cyberthreat modelling", October 2016 (ACT, TOCSA, Oslo Analytics)
http://stids.c4i.gmu.edu/papers/STIDSPapers/STIDS2016_A2_BromanderJosangEian.pdf

"Cyber Threat Intelligence Model: An Evaluation of Taxonomies, Sharing Standards, and Ontologies within Cyber Threat Intelligence", September 2017 (TOCSA, Oslo Analytics)
http://ieeexplore.ieee.org/document/8240774/

"Ethical considerations in sharing cyber threat intelligence" November 2017 (TOCSA)
www.mnemonic.no/ethical-considerations-sharing-cyber-threat-intelligence

"Automatic Detection of Malware-Generated Domains with Recurrent Neural Models" 2017 (Oslo Analytics)
https://pdfs.semanticscholar.org/187b/3af9006ee4cc039b2b97fe03099a1c4133b2.pdf

"Neural Reputation Models learned from Passive DNS Data" 2017 (Oslo Analytics)
http://publications.nr.no/1515491568/neuralreputation-plison.pdf

You can also find these research articles at www.mnemonic.no/references-2018

FINN MYRSTAD
Director of Digital Policy
Norwegian Consumer Council

GUEST ARTICLE

WATCHOUT!

CONSUMER CHALLENGES IN THE INTERNET OF THINGS

In late 2017, the Norwegian Consumer Council and mnemonic cooperated on analysing the terms of service, privacy and security in smartwatches designed for children.

The project, later named WatchOut!, uncovered serious security and privacy flaws. Significant security weaknesses were discovered in three of the four watches tested, leaving sensitive personal data such as children's locations, pictures, and communication vulnerable to various attacks. More bluntly, attackers could easily seize control of the watches and use them to track and eavesdrop on children.

THE CHALLENGE WITH IOT

A vast range of products are being fitted with sensors and internet connections, in what is commonly referred to as the Internet of Things (IoT). It is easy to be blinded by all the possibilities this can give to increased well-being and efficiency. For example, the underlying premise of a smartwatch for children may seem practical for a parent – track and communicate with your child without having to buy them a mobile phone. However, the Internet of Things also introduces and amplifies a number of challenges that need to be addressed. Most of these issues relate to security, privacy and ownership.

Through testing fitness wearables and connected toys, the Norwegian Consumer Council have previously identified serious consumer challenges in connected devices. Nevertheless, the findings in the smartwatches for children were the most alarming.

CRITICAL SECURITY FLAWS

Smartwatches for children, also known as GPS watches, are wearable mobile phones that allow parents to use an app on their smartphones to keep in touch with and track the location of their children. Since the main purpose of these devices is to give parents peace of mind, it is crucial that they maintain adequate security and privacy standards. This turned out to not be the case for several of the devices.

In two of the devices (the Gator and SeTracker family of watches), mnemonic identified flaws that allowed a potential attacker to take control of the apps, thus gaining access to children's real-time and historical location and personal details, as well as enabling them to contact the children directly, all without the parents' knowledge.

Several of the devices transmitted personal data to servers located in North America and East Asia, in some cases without any encryption in place. The SeTracker family of watches also functioned as a listening device, allowing the parent or a stranger with some technical knowledge to audio monitor the surroundings of the child.

Additionally, the abundance of smartwatches for children available internationally, with cheap Chinese products being imported and rebranded by a vast number of local retailers, makes it difficult to obtain a clear picture of who is responsible for the various products. For example, several different smartwatches use the same app and hardware, but are sold worldwide under many different names and brands.

Soon after the findings were published, all of the companies reported that the flaws had been fixed. Based on the severity of the issues, and due to the flaws being potentially very difficult to repair, the Norwegian Consumer Council commissioned a new technical test from mnemonic. The results showed that not only were the problems still present, additional issues had appeared. Voice messages sent between thousands of parents and children were openly available online in a service that consumers had been promised was secure.

Attackers could easily seize control of the watches and use them to track and eavesdrop on children.

THE DATA PROTECTION AUTHORITY TAKES ACTION

Ahead of publication of the first report, the Norwegian Consumer Council alerted the Norwegian Data Protection Authority, which in turn notified the importers and manufacturers in question to allow them to rectify the issues.

Based on the tests and the response from the companies, the Data Protection Authority ordered companies to discontinue all processing of personal information about its customers by mid-January 2018. As far as we can understand, a possible result of this is that thousands of smartwatches for children will stop functioning, since almost all functionality relies on the processing of personal data.

The findings in the GPS watches are not unique. We saw similar issues with the "smart" toys Cayla and i-Que, where cheap components and a lack of security measures left the toys vulnerable to attacks. A plethora of cheap internet-connected electronics are available online. Sometimes the products are imported and marked up in price before being shipped into the European market. Meanwhile, importers and vendors do not seem to know how to ensure that the products they sell are secure.

UPDATING REGULATION TO INCLUDE SECURITY

There are potentially huge benefits for consumers in the Internet of Things, but this will only be achieved if services and products can be designed with trust, privacy and security built in, so that consumers feel that they are fair and safe to use.

If consumers are to embrace these devices, there must be a basic trust among those who buy and sell the products. Therefore, it should be in the interests of everyone to impose strict requirements on security, basic privacy and consumer protection.

Importers and stores should obviously know what they import and offer before they start selling the products. To ensure this, voluntary measures are not sufficient from a consumer advocacy point of view.

Consumers are not aware of what to look for, and are therefore unable to make informed choices regarding cybersecurity in connected products. This is, however, not a responsibility that should be placed on the consumers.

In the short term, vendors must take responsibility for the products they sell, implement control mechanisms and measures that include cybersecurity, and remove products that are not secure from their shelfs. Consumers who have bought products that are not secure, should get a refund based on the lack of security.

In the long term, the current product security legislation and standards meant to cover the safety of individual devices must also protect consumers in the Internet of Things environment. Additional provisions and standards will need to be adopted to ensure the safety of the system as a whole.

THE NEXT STEP

Consumers' need for security, proper support and privacy should be front and centre of product development - not bolted on as an afterthought. This is a lesson some of the importers of smartwatches for children, and unfortunately the owners, are learning the hard way. In an effort to address some of the security flaws, owners of one particular smartwatch are encouraged to follow a 33-step process to install encryption on their watch.

The cooperation between the Norwegian Consumer Council and mnemonic has had considerable effects on the products that were analysed. The terms of service are changed for the better, (some of the) security findings are fixed, and routines seem to have improved for at least some of the companies involved.

More importantly, the work on WatchOut! is part of an ongoing process aiming to improve the security of consumers with IoT products in Norway, Europe and globally. The project has gotten the attention of decision makers and regulators both nationally and internationally, and the political processes following in the wake of this work has just begun.

For more information on the WatchOut! project, including the full report with findings, visit
https://www.forbrukerradet.no/side/significant-security-flaws-in-smartwatches-for-children

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

CORPORATE HEADQUARTERS

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

STAVANGER

mnemonic AS
Solaveien 88
4316 Sandnes
Norway
+47 2320 4700
contact@mnemonic.no

STOCKHOLM

mnemonic AB
Borgarfjordsgatan 6c
SE-164 55 Kista
Sweden
+46 08 444 8990
contact@mnemonic.se

Lead editor: Rikke Klüver Voll, mnemonic AS

Publication design: Bjørnar Løvtangen, Make Noise AS

Photo credits: Charlotte Sverdrup Photography (page 4, 6/7, 12/13, 15, 24, 28/29, 38/39, 48, 60/61, 64/65), Norwegian Consumer Council and Unsplash.com.

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

© 2018 mnemonic AS. All rights reserved. mnemonic and Argus are registered trademarks of mnemonic AS. All other brands, products, or service names are or may be trademarks or service marks of their respective owners.