mnemonic

Securing your business.

2020

SECURITY REPORT

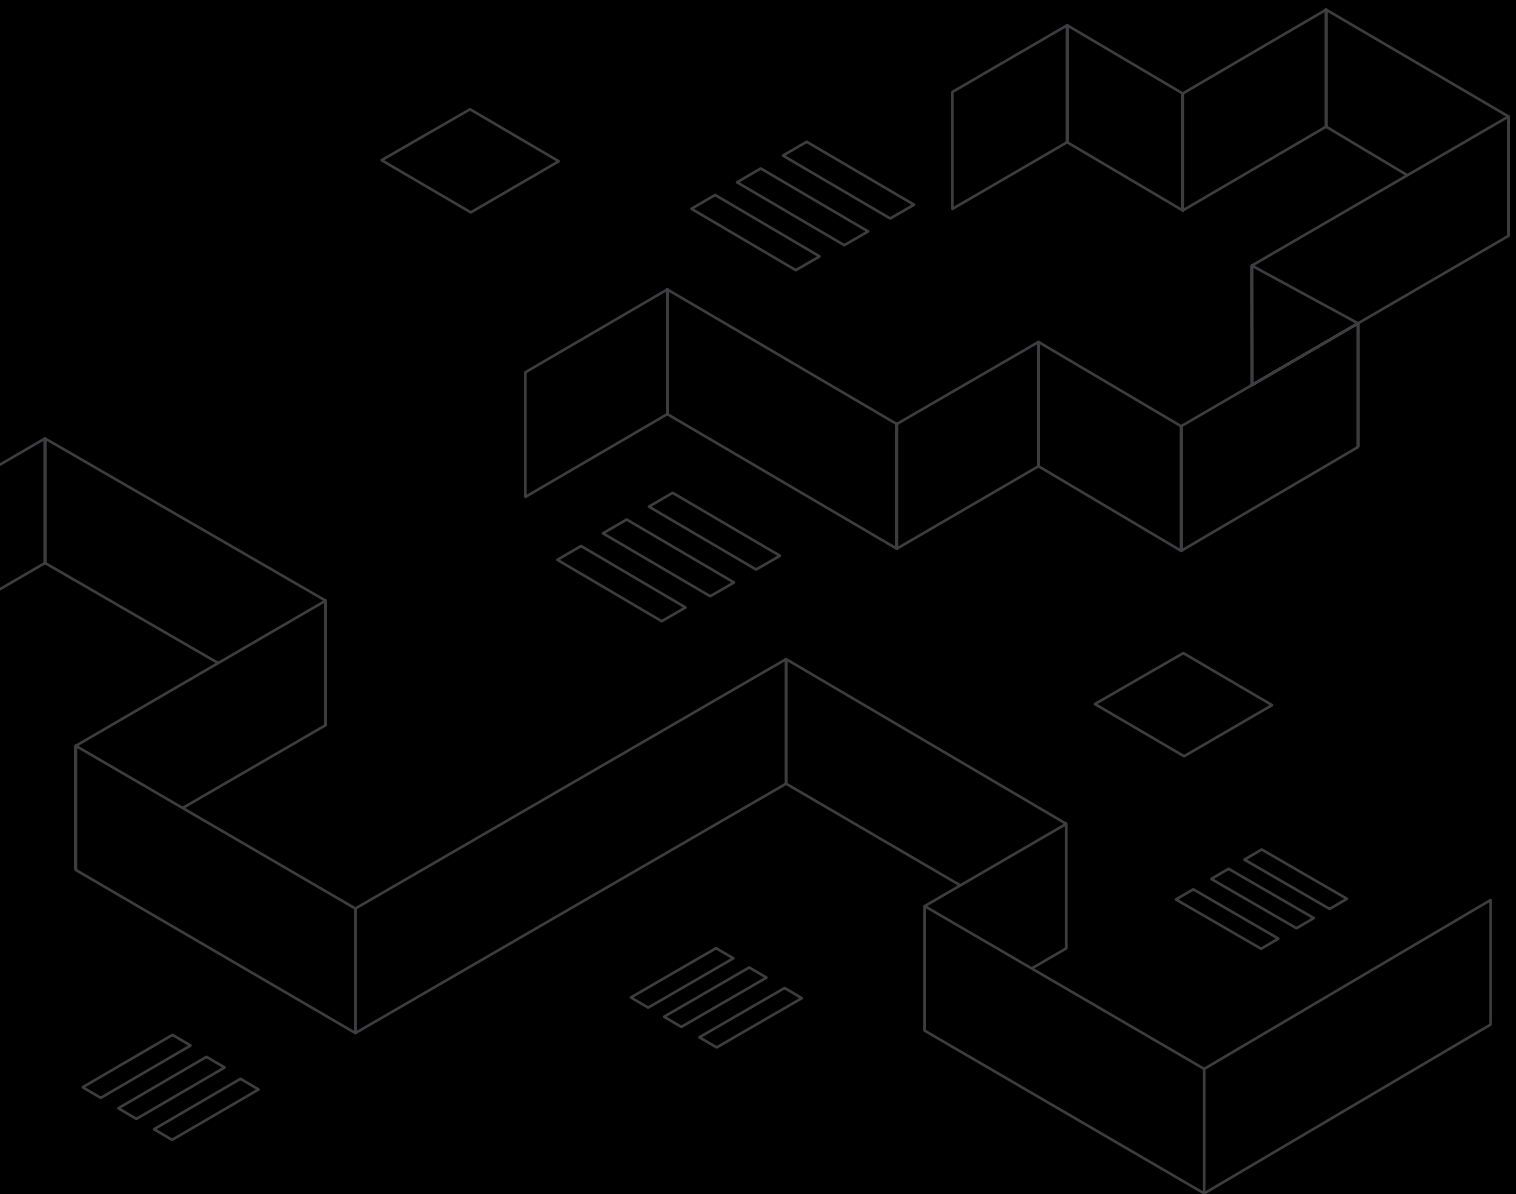

A Reflection on Two Decades in Cybersecurity

As mnemonic celebrates its 20th anniversary, I feel it's only appropriate to reflect on the years that have past, and what the future may hold.

When mnemonic was founded in 2000, the world was a very different place. The dot-com bubble was about to peak and subsequently burst in spectacular fashion, and the Internet itself a distant relative of how we recognise it today. This was a time before Facebook, before YouTube, before Wikipedia, and before mobile apps. Google was a start-up with less than 100 employees, Apple was recovering from near-bankruptcy, and the Internet was most commonly accessed on 56k modems.

During this time, cybersecurity was in its infancy. Few organisations had personnel dedicated to security – a task often undertaken by network teams or merely anyone who took the task, and a secure network consisted of a stateful firewall with anti-virus on the endpoints. At the time, no one would imagine that security would become a regular boardroom discussion, breaches a regular segment on the daily news cycle, or there being a global shortage of security professionals that is measured in the millions. The world was a different place, but looking back, it's hard not to see exactly where we were headed.

Over the years, we have seen waves of new technology adopted by our customers, new security solutions created by the market, and ever-rising demands from society for a technologically-driven future. Government policy appears to finally be catching up with technology, rather than falling behind, and users globally are becoming more aware of their digital rights and online presence in general.

We continue to adapt to these changes and evolution through developing technology, investing in research, and establishing partnerships throughout the security industry. One constant through these past two decades has been the need for people. Security has, and for the foreseeable future will continue to be a challenge created and solved by people.

At mnemonic we pride ourselves on the 230 high-skilled professionals we are lucky enough to count as part of our team. Our conscious effort to build a culture of continuous development, respect and autonomy was publicly recognised in 2019 as mnemonic was ranked as Norway's top workplace, and number 15 in all of Europe. This is an extraordinary honour and achievement that we will continue to harvest, invest in and improve on for the next two decades.

From our origins in Norway, to Sweden, the UK and now in 2019 the United States, it is this global team of professionals that has, and will continue to steer mnemonic to address the next two decades of security challenges.

Thank you for the past twenty years, and I hope you enjoy the eighth iteration of our Security Report.

TØNNES INGEBRIGTSEN
CEO, mnemonic

TABLE OF CONTENTS SECURITY REPORT 2020

ARTICLE

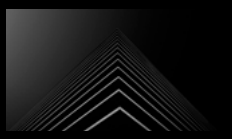

05

Security Predictions 2020

INTERVIEWS | STATISTICS

WORD ON THE STREET

17

Sensa

31

Storebrand

2019 IN NUMBERS

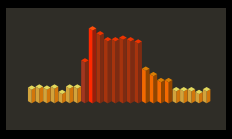

51

2019: A View From mnemonic's Security Operations Centre

ARTICLES

11

Strategic Software Security

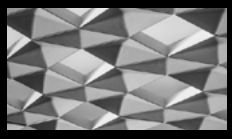

19

Security Risk Management in Supply Chains: How to Avoid Unacceptable Risks

25

Internet of Things and Its Firmware: A Tale of Memory Corruption Bugs

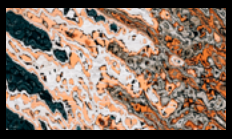

33

The Last Piece of the Puzzle: Incident Readiness

ARTICLES

39

The Value of Outsourcing Detection and Response: Making Informed Security Decisions

45

Integrating Security Controls Within a DevOps Pipeline

53

The NIS Directive: A Step in the Right Direction

61

The Missing Link in Email Security

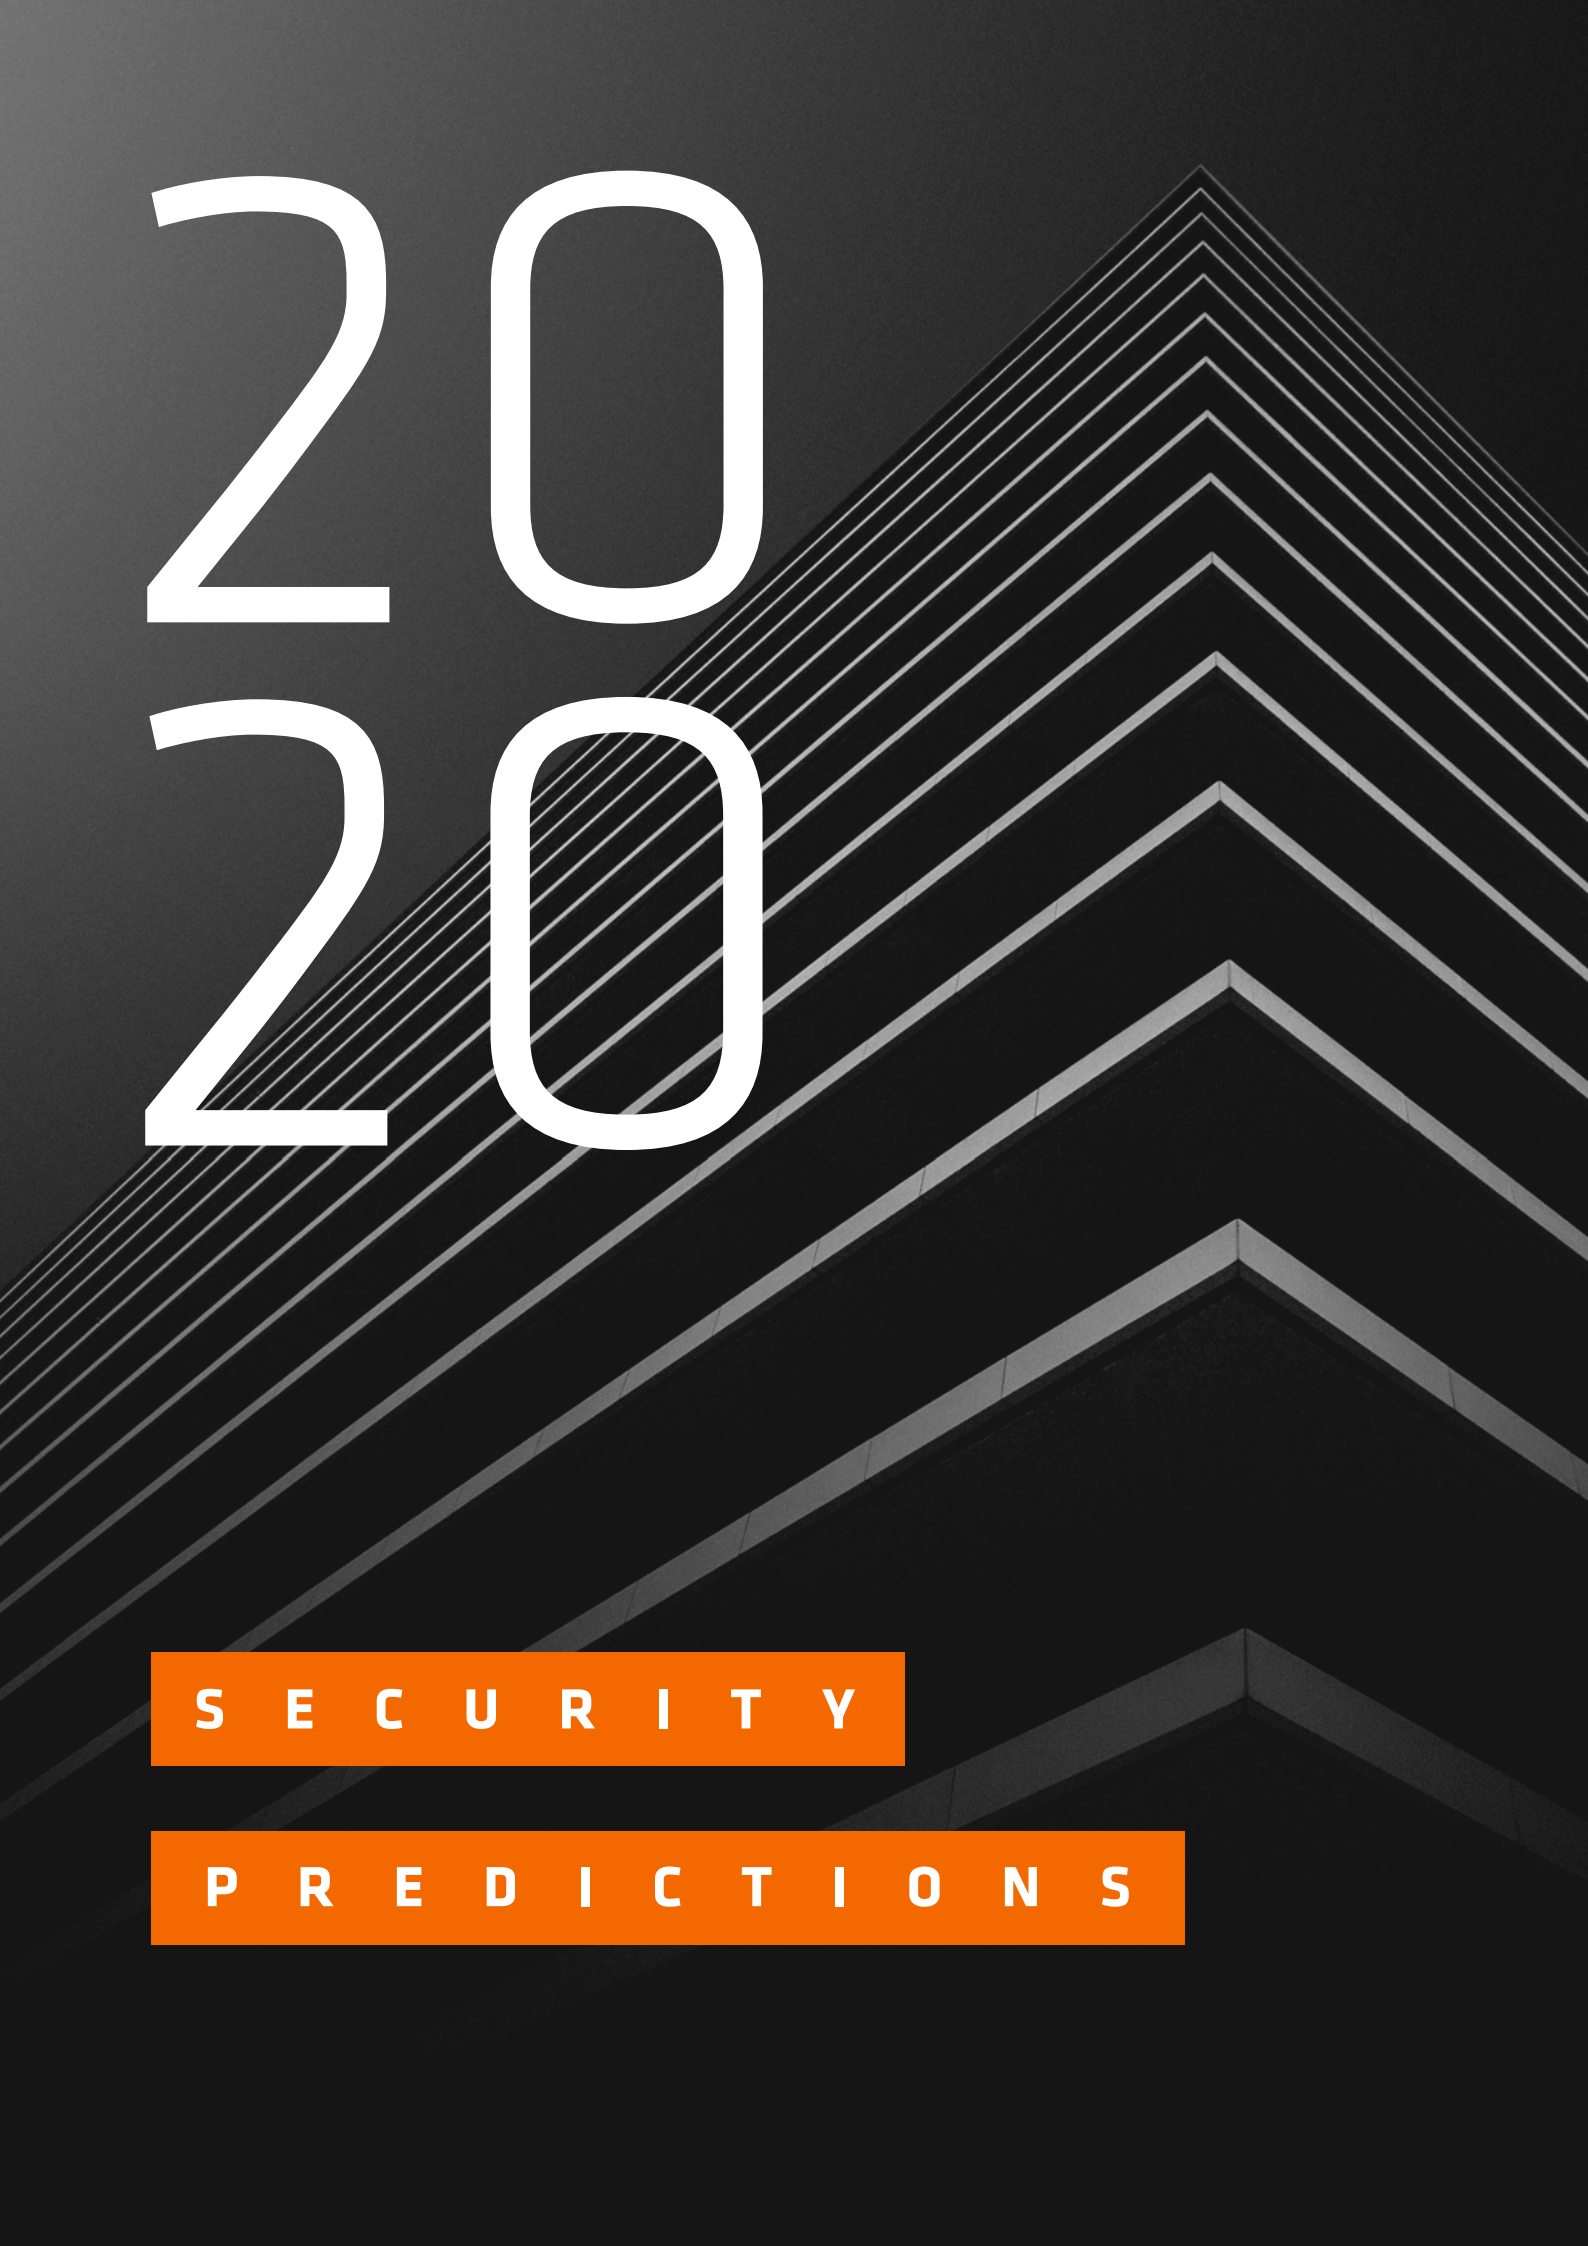
2020
SECURITY
PREDICTIONS

Jon Røgeberg
Head of Threat Intelligence,
mnemonic

Morten Weea
Security Consultant,
mnemonic

The good year of 2019 has delivered its concluding remarks, and the runic calendar shows that it is time to peer into the scrying bowl and interpret the omens for the year that is now upon us. We are ready to toss our bones and entrails for a glimpse of what the future will hold. As this year's haruspex, we have calibrated our foresight to 2020 and will be conveying the universal truths of the future on the next few pages.

However, to enter the right state of mind, we should take a look at what we predicted for 2019. Quite accurately, we foretold the coming of quantum computing. Google has our backs regarding this lofty prediction and declared "quantum supremacy" late in 2019 – thanks to their Sycamore processor, a toned-down 53-qubit processor version of the Bristlecone. Calculations with an estimated completion time of 10,000 years were completed in just 200 seconds, or just above 3 minutes.

The significance of the quantum breakthrough can be compared to the first flight of the Wright brothers, and, as history has taught us, practical use of this new technology will likely be years down the road. First, we need to design our equivalent to a stable plane before we can develop supersonic engines and stealth technology. However, demonstrating the capabilities of quantum computing will open up a completely new field of possibilities where both normal and mad scientists will be able to play around and invent currently unimaginable creations. These creations will be considered so natural that we won't be able to see how we could have lived without them in the past.

Despite Google claiming to have achieved quantum supremacy, RSA and ECC encryption algorithms will not immediately be at risk, and Public Key Infrastructure (PKI) can also be used in the immediate future. However, with quantum computing gaining more and more traction, it is only a matter of time before the current best-in-class encryption is broken. Therefore, we recommend that you exhibit great care when considering the purchase of IoT devices that do not include the possibility of upgrading to a new encryption standard, given their short security lifespan. ▸

Secondly, we also predicted the shift from destructive to less destructive approaches by cyber criminals. We were partially right about this, given that one of the biggest cybersecurity incidents of 2019 was the Norsk Hydro case. A destructive strain of ransomware hit Hydro, which ended up costing them a pretty penny. This leads us into our first prediction for 2020.

PROFESSIONALISATION OF CYBER CRIMINALS

After predicting that cyber criminals would shift their focus away from destructive methods to less destructive methods last year, we realised that these are not necessarily mutually exclusive strategies. Focusing more on increasing return on their investments, criminals will enter 2020 concentrating on what will give them more bang for their buck.

Criminal actors are already structuring their organisations like businesses and dividing responsibilities into appropriate departments. They have in-house developers maintaining their malware codebase and operational departments operating and optimising virtual "datacentres" of infected botnets. The natural next step would be to develop attack strategies yielding higher ROI to increase revenue.

Identifying the security measure threshold of their potential victims enables criminals to choose from their arsenal of attack vectors, giving them the possibility of attacking with a higher-yielding option.

So, what does this mean for us? In short, this means that the information superhighway we are all siphoning bits and bytes of will have cyber criminals present on all layers. Cyber highwaymen will be targeting the weakest victims, hitting passers-by with the shotgun approach. The old highwayman phrase "your money or your life" is now replaced by "your money or your data" – accompanied by low-effort, destructive ransomware.

During the end of 2019, we observed these criminals leveraging a new approach to ransomware. Instead of urging their victims to pay to get their data back, the datanappers now threaten to expose your data if you do not pay the ransom.

Consider the following scenario; you are a successful business, and you are taking your backup procedures seriously. You test everything regularly, and you have working backups with sufficient retention lengths. If things go sideways, you could always roll back to the day before and nothing would really be lost. You have eradicated the threat of ransomware – or so you think.

One day you get the otherwise dreaded message: "all your base are belong to us." Ransomware has hit you, and now

you need to pay. You initiate your incident response plan to determine the epicentre and fallout zone of the ransomware in parallel with your preparation of rollback to yesterday's backup.

Once you start reading the fine print of the ransom note, you see that this threat is a little different from the usual ransom notes. This time you do not have to pay to unlock your data. This time they have spiced things up, and you need to dig into your bitcoin wallet to prevent exposure of your data.

As you perform your investigation, you discover that the encrypted data are sensitive, and under no circumstances should become exposed to the public uncontrolled. Suddenly the consequence of being hit by this ransomware is grave and unmitigated. Paying the ransom becomes a more and more viable option.

CYBERSECURITY BEING A PART OF THE "BIG BOYS TABLE" AND CYBERSECURITY INSURANCE ON THE RISE

Acknowledging the grim scenario presented by the increasingly specialised and professionalised threat landscape leads to new challenges when it comes to tackling cybersecurity.

Firstly, we are going to see an increase in focus on cybersecurity. The CISO role is invited to "the big boys table" where it belongs. Investing money in cyber defence and countermeasures is as hard as budgeting for any other preventive and defensive arena. Why do we need extensive monitoring, and why should we establish and train an IRT or CERT if all we get is drive-by malvertising? What is *good enough*, and *what is good enough for us*?

Establishing an appropriate level of cybersecurity in your organisation fully depends on the transparency of the cybersecurity field. Businesses need to participate in relevant forums where like-minded professionals gather. There are two disciplines in particular that could assist in determining the sufficient level of cybersecurity: the cybersecurity consultancies and the insurance companies. Both have access to data from a wide variety of businesses across multiple industries. Assisted by tools like *Top 20 Critical Security Controls* from the Centre for Internet Security (CIS) and the Open Web Application Security Project (OWASP) Top 10, you can establish baselines for cybersecurity policies and procedures.

Cybersecurity insurance is on the rise, and companies like Norsk Hydro had such insurance prior to their incident. However, the catch with insurance companies is that they punish you hard for negligence. If you fail to try to prevent an incident, your insurance payout suffers massive reductions. This is similar to car insurance, where those who are involved in an accident will likely see an increase in their insurance premiums. Having the insurance discipline on the dictating side and the cybersecurity consultancies on the advising side means that expectations for your security levels are clearly defined and set.

However, black swans exist and even though they were unimaginable before we first observed them, the probability of them existing was never zero, which leads us to the next prediction.

CYBER WARFARE CONTINUES TO HEAT UP

As Russia annexed Crimea without any consequences, the leader of the free world nicknamed his opponents with names such as "Little Rocket Man". There is a full-blown trade war between the U.S. and China, civil unrest in former colonies like Hong Kong, general instability in the Middle East, and the Democratic People's Republic of Korea has stepped up its military game. The gloves have come off, and the world is a cyber spy playground resembling a game of Risk.

Nation states have proven capabilities in the cyber domain as well as a lack of respect for borders. Everyone is fair game, and friends do not necessarily keep away from friends. Edward Snowden exposed and disclosed comprehensive surveillance across friendly borders, reminding us that not even the "good guys" can be trusted. Companies operating in markets with threat actors representing foreign interests need to consider this.

When buying equipment from vendors, the vendor's origin should be considered prior to signing a long-term contract binding the customer to the vendor. Governments from both East and West have been caught red-handed with requests, demands, or actual backdoors into vendor products, forcing you to pick your own poison: Will you accept a possible backdoor from a government with whom your own government has a security cooperation, or not?

The discussion is raging globally, with many parts of the world upgrading their critical infrastructure to 5G. The global mobile scene is not the only arena for constant competition between vendors representing different interests. It is known that spies are used mainly in peace to secure a foothold with the possible future enemies. The ongoing infrastructure war could be considered a matter of positioning for the future.

What does this mean for us as consumers and users? There isn't much we can do on an individual level. If we want to participate in the society surrounding us, we need to obey the rules and live by the choices taken higher up the political ladder. ▸

Even so, we can participate in the choices we need to take and be aware of the pros and cons of all vendors. This will be more important than ever as more and more of our daily life is based around technology.

Another aspect of the increased tensions is that cyber superpowers are battling it out through proxies. Attacking smaller countries with less mature cyber defence capabilities is even more attractive if they are also already in various kinds of cooperation with the larger nation states. This makes them a privileged attack vector. Even if you didn't think you were of any particular interest to the cyber sluggers, you can still be abused as a stepping stone towards what they really want.

ARTIFICIAL INTELLIGENCE AND DEEPFAKE

Our next prediction is that artificial intelligence will be deployed by both cat and mouse in the cybersecurity game. With an enormous focus on machine learning and many of the big contenders on the tech scene investing serious money in machine learning startups, this technology will only grow in the coming decade.

Accordingly, the downside of Google Duplex is that this technology will not only be used for good. With more and more automation of homes and services comes an increased speed of AI development. It has already been confirmed that companies like Amazon, Google, and Apple have people listening in on the voice commands, with the explanation that they seek to be "improving voice recognition."

While we are not at the point where robots pass the Turing test left and right, we are at a point where you can customise and specialise an AI with a natural, synthesised voice to ask for opening hours and order tables at a foreign restaurant. Diversifying into specialised software to imitate CEOs and other high-ranking officers is not improbable.

To improve on these fraudulent phone calls, the bad guys could simply add a Skype video callwith deepfake technology and a "jittery connection" to cement the authenticity of the distress callfrom our bootleg executive.

In addition to the obvious combinations of AI and deepfake, AI can also be used to counter all improvements made by the security teams. Where security scientists are deploying AI to scan emails, systems, and networks, malware developers are using the same AI techniques to gather intelligence on their victims. AI can also be leveraged to trigger weaponisation of their malware at given points.

What makes AI so powerful in the everlasting game of cat and mouse is that AI learns from context and improves. Malware with AI can even determine where to propagate, and how to do it based on the information it gathered on its own. Where old-school malware was developed and deployed, new AI-supported malware now gathers information and honours the old military adage *improvise, adapt and overcome.*

The progress in AI development triggers some questions about ethics related to people unknowingly communicating with a computer, and Google vows to inform whenever their AI is being used for communication. Due to the nature of their work, adversaries do not share the same ethical concerns when it comes to leveraging AI to fool human beings, however. They very much count on the robot to be as convincing as necessary to make you do its bidding.

RETHINKING THE CLOUD APPROACH

Finally, having everything stored "in the cloud" the traditional way is becoming old-fashioned at the speed of light. We predict revision of this classical cloud approach, and therefore remind our returning readers of the sound and healthy approach to serverless security. Amazon and Microsoft are pushing more and more for serverless environments. This opens up for great opportunities, but also security-related challenges. The obvious benefit of going down this road is the fact that cloud providers are much better than the rest of us at keeping the infrastructure upgraded. Server software and operating systems are being delivered from cloud providers while also being kept updated and secure.

On the other hand, traditional approaches still prove to be effective. Serverless does not mean codeless, and all snippets of code that are not securely written can still be exploited. Cross-site scripting and injection vulnerabilities are still open to exploitation by an adversary. The increased complexity that comes with smaller, dedicated functions also increases the difficulty of monitoring. It will be increasingly difficult to monitor them all, and having dedicated functions for all operations means that you would need to step it up a notch when it comes to access management.

Just having functions running on dedicated systems means that you are sharing your perimeter with an unknown number of other customers. This makes it impossible to implement perimeter controls or perform attack surface vulnerability scanning the way we are accustomed to.

The combination of sharing the perimeter and increasing the transit points for your data and data calls means that you are also dramatically increasing your chances of having that data interrupted, manipulated, or leaked. ●

ARTICLE

Strategic Software Security

Tor E. Bjørstad, Ph.D.
Application Security Lead,
mnemonic

AFTER READING THIS ARTICLE, YOU WILL:

- Understand the need for a strategic approach to software security
- Have learned how to establish a software security initiative
- Have gained an overview of existing software security frameworks

"We believe that every industrial company will become a software company."
– Jeff Immelt

"Software is eating the world."
– Marc Andreessen

A convergence towards software

In 2020, the world runs on software, and software is running the world. No longer the sole domain of technology companies, software is now part of the core business of just about any modern company. When software malfunctions or crashes, trains stop, ports close, payment services cease to function, and industrial plants halt production. The Boeing 737 Max accidents in 2018–2019, caused by its MCAS software system, are just a single demonstration of how software defects may turn deadly.

The ability to rapidly produce and use high-quality software powers innovation, enables a shorter time to market, and saves effort otherwise spent on addressing problems and hunting bugs. However, despite the general reliance on software, many companies do not have a clear software strategy, and the capability to acquire, build, and operate software in an effective and efficient manner is often lacking.

Over the last 20 years, we have seen major changes in how software is made and used by leading organisations. A shift towards lightweight methods began as early as the 1990s, leading to the watershed publication of the *Agile Manifesto*[1] in 2001. Today, a vast majority of organisations follow (or claim to follow) agile software development methodologies, albeit with varying degrees of success. In 2009, the core concepts of DevOps began to be formalised, with organisations striving to break down barriers between software development ("Dev") and IT operations ("Ops"). As part of these developments, the distance between software and business is also decreasing. Today, it is common to have business stakeholders directly involved in software development activities, rather than being manifested through a static requirements specification or otherwise distantly engaged.

At the same time, the distance between software and technical infrastructure is also vanishing. Through trends such as cloud, automation, and infrastructure as code, IT operations are quickly becoming more like software development as well. In a cloud environment, an infrastructure that would have taken weeks to provision and configure manually can be defined in code and deployed reliably and repeatedly at the press of a button or even be triggered automatically. ▸

[1] *See Reference List at the end of the report*

Software and security

Though many organisations struggle with software, it is probably fair to say that even more organisations struggle with security. There are multiple reasons for this, not the least of which is that information security has traditionally been implemented within a mindset of audit, compliance, and control. This approach often leads to reactive and sometimes heavy external processes, which are not sufficiently integrated into the core business processes. Because of this, they may align poorly (or be perceived to align poorly) with the overarching business need for agility and rapid delivery.

Another challenge is the general lack of skilled information security professionals, which means that many organisations do not have access to the competence and manpower they need within this area. Security organisations are rarely large and well-staffed, and there is never enough time to keep up with the ever-flowing stream of issues and incidents.

To meet these challenges, security is unlikely to be successful as a reactive add-on component. On the contrary, security has to be an integral part of the enterprise's overall software strategy. A failure to deliver on software security will both serve to slow down the organisation's software development efforts, and also contribute to increased operational risk due to software insecurity and loss of control.

When decisions are made by small, autonomous development teams, security activities need to be part of the default workflows and processes. At the same time, there must be a clear connection between low-level security activities and overall security objectives and strategy. Establishing a software security initiative covering the whole organisation is a suitable way of coordinating, managing, and evolving software security activities and capabilities.

Road traffic safety has improved tremendously over the last 50 years due to improved safety measures. Software security needs to make similar improvements, and automotive metaphors abound. In the context of secure DevOps (SecDevOps), the conceptual model is often described as building a paved road with guardrails for the developers. In order to move really fast with software, the metaphorical equivalents of brakes and seatbelts are clearly needed.

Nonetheless, existing security practices cannot be neglected. Maintaining a continuously updated asset inventory is critical to maintain situational awareness as the rate of change within IT keeps growing. Periodic penetration testing remains as necessary as ever in order to ensure that security controls are working as intended. However, periodic tests are no longer sufficient as the only detective software security control when production deployment is something the organisation does on a daily basis. The need to demonstrate compliance does not disappear when the software development methodology changes.

Scaling the security organisation

In order to keep up with the increased pace, a software security initiative needs to look beyond the core security team. While everybody cannot be security experts, there are many people in a typical organisation who are curious about security and who may be interested in learning more. Being able to identify the right people and engage them with a positive and enabling message is key to extending the reach of security. A little bit of security evangelism can go a long way, both when aimed at technologists and when aimed at management and other parts of the business.

When development teams are asking themselves "what could possibly go wrong with this feature?" and are able to identify potential security risks as part of the regular development process, then something is going right. Injecting that little bit of security awareness into the process makes it much easier to identify, approve and expedite low-risk changes with more confidence than before. On the flip side, when a potential issue is identified by a team, it must be possible for developers to draw on additional security resources for guidance and quality assurance.

Automation can be another useful way to scale out security. Traditionally, a high false positive rate has been an Achilles heel of automated security testing, and it remains something any automation project must address. Poorly configured tools will do nothing but annoy, and any true positives are likely to be lost in the noise. Despite this, there are many types of security issues, human errors, and misconfigurations that can be identified accurately and automatically. Source code and configuration analysis, continuous vulnerability monitoring (CVM), and dependency analysis are three areas that are very suitable for automation. As long as it is possible to produce accurate and actionable results through automated testing, it can serve as a powerful complement to traditional test and QA activities.

Finally, a modern approach to software development also brings clear security benefits. When software is deployed rarely and manually, and changes are made in large increments, the risk of security defects increases. It also becomes harder to keep systems patched and up to date with security updates. By making automated deployments routine, it becomes easier to roll back changes or make emergency updates, the risk of botching a deployment goes down, and it becomes possible to prevent configuration drift. It also simplifies a lot of compliance issues.

Software security frameworks

While there is no "silver bullet" or one-size-fits-all solution to software security, there are many useful frameworks and resources publicly available. ▸

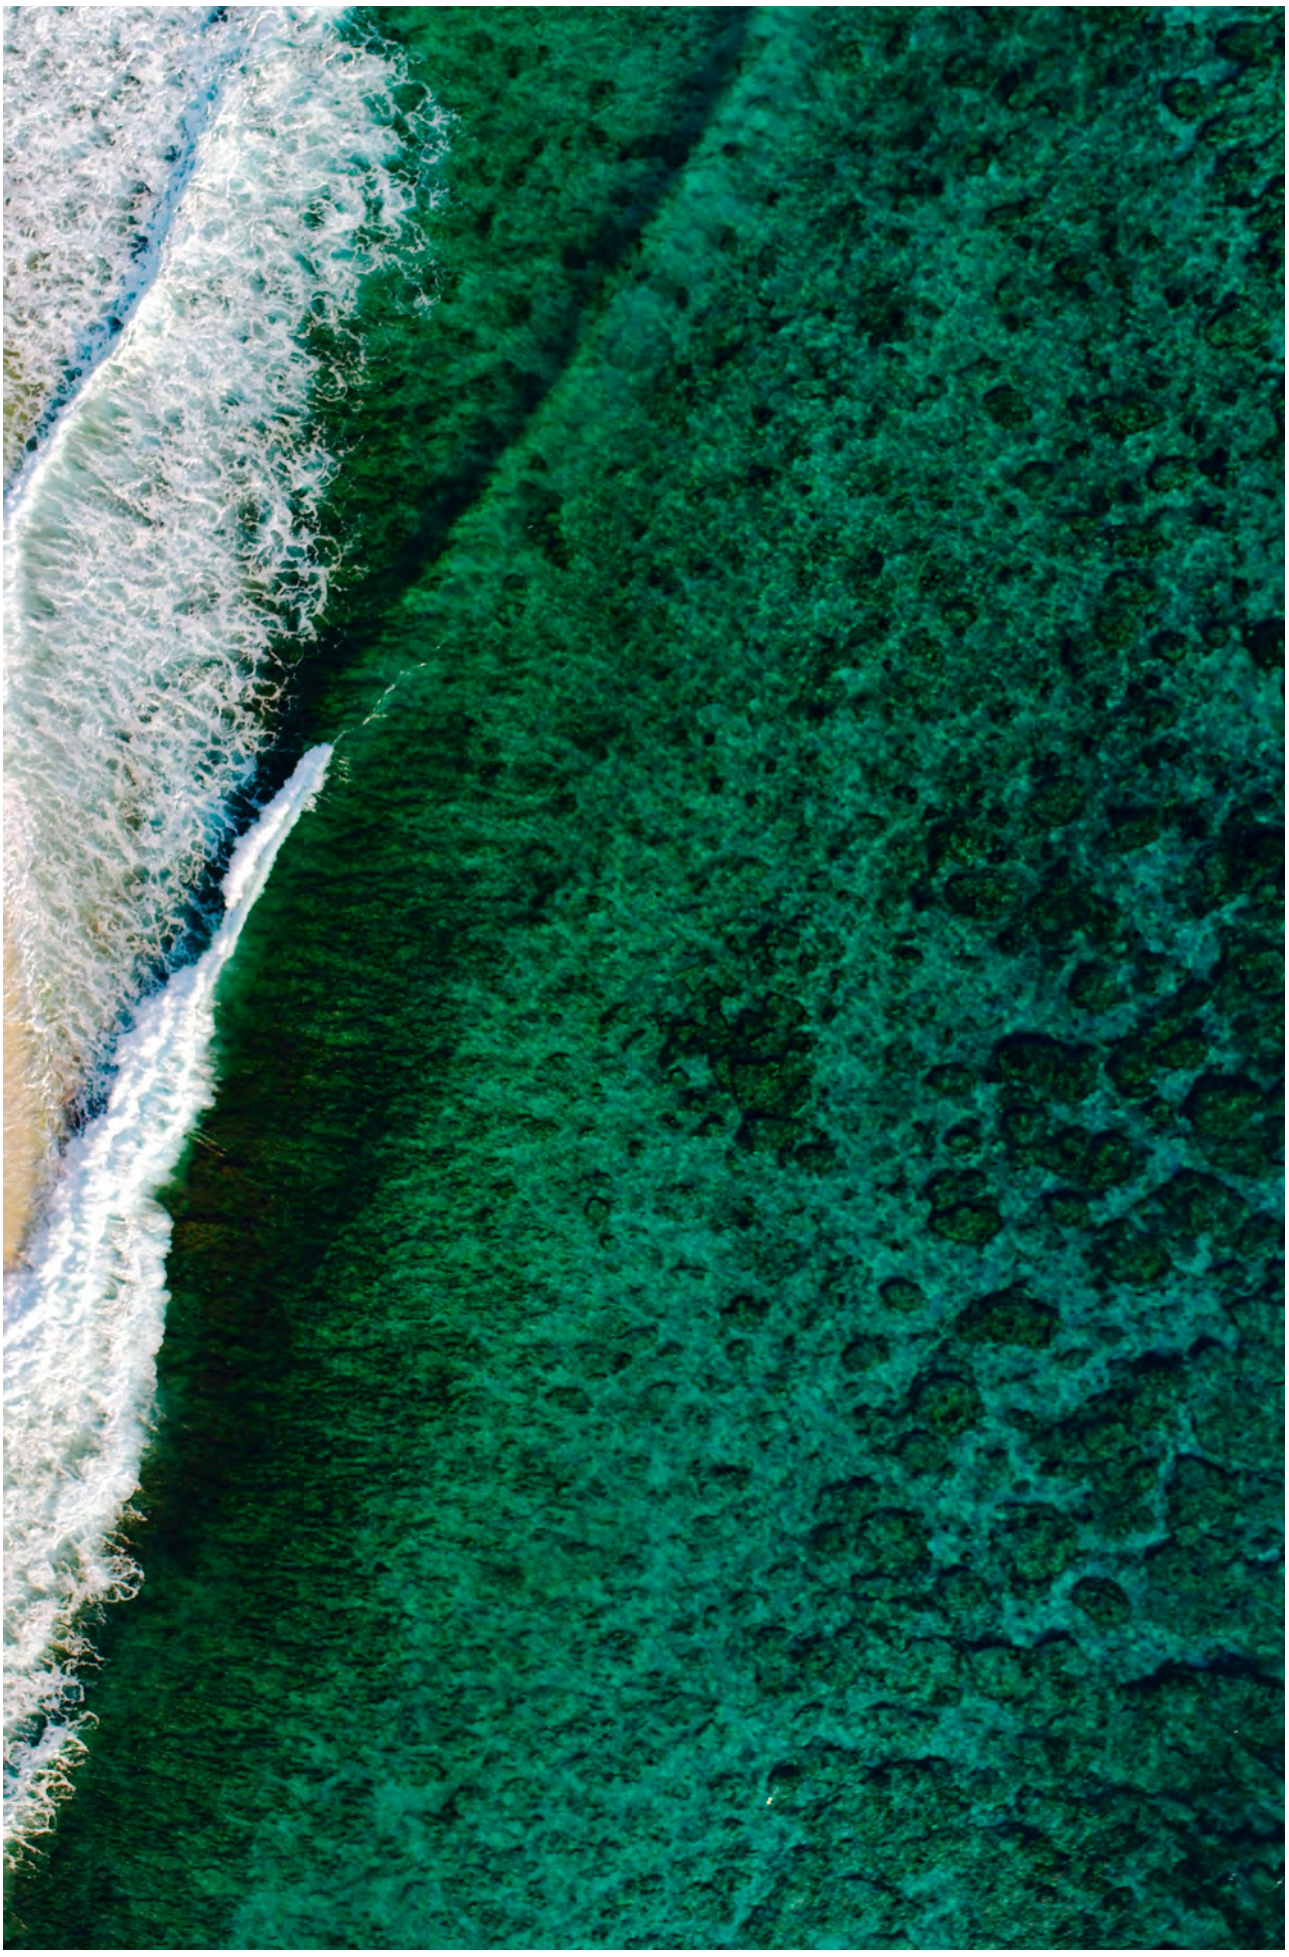

The frameworks can be used both to structure software security activities and as a source of possible activities and controls that have been tried and tested elsewhere.

Microsoft's software security activities started in earnest with Bill Gates' famous trustworthy computing memo[2] in January 2002. Their *Secure Development Lifecycle* (SDL)[3], which was adopted internally in 2004 and released to the public in 2008, is one of the first large-scale secure software methodologies published. The SDL is still being used, revised, and refined by Microsoft, who provides ample public resources to support it. Adopting the entire SDL directly in one go may be too complex and invasive for most organisations, but there is nevertheless a lot to learn from its guidance, practices, tools, and processes.

Another good source of software security practices is the *BSIMM framework*[4]. This is not a prescriptive framework telling organisations how to organise their software security efforts. Instead, BSIMM takes the form of a repeated survey, describing which activities are commonly found in other development organisations of various sizes and across multiple sectors. Because of this, BSIMM has multiple uses. It can be applied as a yardstick for self-assessment ("how are we doing compared to other companies"), and as a list of possible activities ("other companies are doing X, we could try it and see if it works for us"). BSIMM is now in its 10th iteration, and its history provides a large amount of information about how security practices are evolving.

In order to ensure that software is built with the right security controls in place, OWASP's *Application Security Verification Standard* (ASVS)[5] is a good place to begin. For developers, it provides guidance on which types of security controls are commonly applicable, which can then be mapped to product features to identify potential gaps. During procurement processes, it provides a structured framework and a common language to specify the rigor and depth of testing that is required. This applies to more than just penetration testing engagements. For example, it is surprising how rarely software and SaaS vendors are required to document how rigorously their solutions have been tested for security issues. The ASVS is a good starting point for gathering such requirements.

Finally, the value of lightweight threat-modelling or risk-assessment activities for stimulating a conversation about security should not be underestimated. Mozilla's *Rapid Risk Assessment*[6] is one example of a lightweight activity that is easy to adopt, and it provides a structured way to identify and flag whether a specific feature requires additional review or a more in-depth risk assessment.

Preconditions for establishing a software security program

Software security is not easy, but it is necessary – at least for organisations that wish to use software to their advantage. Establishing a successful software security program requires time and investment, and a common understanding among stakeholders is that it is something worth doing on a strategic level. Somebody in the organisation must have a dedicated role and a clear mandate to bridge the gap between security and software development. This involves working closely with the software development teams and architects to build the processes, tools, guidelines, and knowledge that are needed, and serve as an evangelist and inspirator within the organisation.

As with the software development process itself, software security is not a one-off activity or a standalone project. On the contrary, a software security program should be iterative and have a short cycle length, just like the agile development processes it aims to align with, in order to utilise rapid feedback loops and find approaches that work well in practice.

To start out, establish business context and strategic direction, and get the necessary buy-in within the organisation. Follow this by creating and expanding security capabilities iteratively, in collaboration with the technical teams and other stakeholders. By evaluating the effectiveness of these measures continuously, the organisation will gain updated knowledge on where the pain points are, how they can be mitigated, and tools to identify measures that are not working as intended. Hopefully, performing these actions will help the organisation step up its effort to establish a successful software security program. ●

Despite the general reliance on software, many companies do not have a clear software strategy.

[2, 3, 4, 5, 6] *See Reference List at the end of the report*

SENSA

Iceland

Guðmundur Þór Jóhannsson
Senior Security Architect

Sensa is a professional managed services provider founded in Iceland in 2002. They specialise in and offer a wide range of digital solutions and technologies for networking, data centres, collaboration, security, hosting and more. Fully owned and financially backed by Iceland Telecom (XICE: SIMINN), Sensa has grown its revenues and maintained profitability during every year of operation.

What is your biggest cybersecurity concern?

Our biggest concern is that adversaries will abuse our systems, and thereby affect the relationships we have built and the trust we have earned with our clients. As a managed service provider, we host a large variety of solutions for our clients, many of which with custom or special configurations. Metaphorically speaking, we have so many doors and windows to secure, and are aware that it only takes one of them to be open for something bad to happen. Ransomware, business email compromise and general security monitoring are some of the main topics we've used a lot of our brains on this year.

In what areas of cybersecurity do you think we're falling behind?

In Iceland, as in many other countries, we see that "human" security resources are limited. This means that the client is heavily reliant on their relationships with vendors, as they simply do not have the time or resources to make sure everything is done properly and according to best practices. Relying on vendors is fine, as long as they themselves have invested in and obtained the adequate technical knowledge needed to advise their customers on security matters.

We are going to fall behind if we allow the commercial side of cybersecurity to control the roadmap as opposed the technical side. In addition, we feel that information security awareness is also a big concern in the Icelandic community, and is something that needs to be addressed through governmental involvement, for instance by issuing promotional materials to increase information security awareness.

What gives you hope for the future of cybersecurity?

Even though information security awareness is a big concern, the feedback we are getting from the management in companies is that security is being more highly prioritised than in the past. Recent examples of high-profile security breaches in Iceland, as well as in other countries, have given companies good reason to focus more on cybersecurity. We also see that user awareness is on the rise, though there is still a long way to go.

ARTICLE

Security Risk Management in Supply Chains

How to Avoid Unacceptable Risks

Torbjørn Øverland Amundsen
Senior Security Consultant,
mnemonic

AFTER READING THIS ARTICLE, YOU WILL:

- Understand the fundamental challenge of setting the right security requirements
- Have a generic framework and process for finding the right requirements
- Have learned how to better ensure that requirements are met

In a digital world where services and products are becoming increasingly more specialised, it is a fair assumption that organisations rely more on vendors in order to meet their business requirements. At the same time, it is also fair to assume that an increase in the number of any organisation's systems leads to increased complexity, something that is directly correlated with the number of risks the organisation must manage. Add to this the fact that it is inherently more difficult to manage risks that have its root cause outside of your organisation, and the conclusion should in most cases be that all organisations will have to manage more risks in the days to come. So how can we handle this complexity to ensure that the organisation does not end up with unacceptable risks as a consequence of trying to meet necessary business requirements?

Pyramids and computers – the same old challenge?

The rise of new information security requirements and the related threats has added new complexity to the procurement process. This has again caused more frustration and fear among many organisations. However, it is important to remember that although the specific challenges we are facing are of a newer date, both the fundamental challenge and its solution are very old, perhaps as old as our first civilised society. The challenge arose the moment humans decided to execute larger projects.

In principle, any procurement can be viewed as a change in a value chain. We either expand or shrink the value chain, or we are exchanging a part of it with a new part. Whether you need stones to build a pyramid, or you need a new application to manage CRM, the underlying challenges are in some fundamental ways similar. If we formulate them in risk terms, we can simplify and say that two main types of risks can affect a procurement or perhaps any generic change in the value chain:

1. The risk of setting wrong or no requirements
2. The risk of implementing correct requirements in the wrong way

We can use this way of understanding the challenges as an indicator for whether a suggested process is well defined or not: A well-defined process for handling a change in the value chain must be applicable across all types of knowledge domains. It should not matter if you are sending a rocket to Mars, building a hospital, or procuring an access management system. Consequently, the process that will be presented is one that is possible to apply to all knowledge domains, but to give concrete examples of challenges and solutions, examples from information security will be used.

A generic framework

To regard a procurement as a change in the value chain has certain logical consequences that are unavoidable but not necessarily intuitive. Please note that shrinking the value chain is omitted for practical reasons. We can present these consequences as a list:

1. Exchanging a part "A" in the value chain requires that we understand both part "A" and all the other parts connected to "A" in order to understand all ▸

the requirements for the new part "A*".

2. Expanding the value chain with a new part requires that we understand the requirements for the new part and all the other parts that the new part can be connected to.

3. If exchanging or expanding a part leads to any changes in the connected parts, the affected parts must also be treated as if they are to be exchanged or expanded.

Those of you familiar with algorithms will notice that the list is recursive. In theory, this means that it is capable of expanding itself until every part of the value chain is included. To analyse every part of a value chain is in many cases impossible, due to shortage of manpower and time. Unfortunately, this is not a purely theoretical problem. Certain changes will require a full analysis of the value chain.

Information security is a domain in which the overall attainment level often is defined by the lowest score, or as the saying goes: "A chain is only as strong as its weakest link." If we combine this insight with our knowledge of value chains, it becomes clear why information security related to supply chains can become quite challenging, insofar as we need to have complete control of a potentially long and complex chain.

A specific process for finding requirements

Let's look at an example: An organisation wants to replace a computer system. In order to discuss this challenge within the boundaries of this article, we need to make some simplifications:

1. We regard the system as only one part.
2. We assume that we have full control of all aspects of the old system, including what kind of data it stores.
3. We assume that we only have to analyse one new system, not a set of vendors.

The first step is to map all functionality in the new system. Note that we will also have to map the functionality we don't plan to use if it is integrated in the system and cannot be removed. We need to understand what all the functions do and how they do it. Further, we need to generate a list of the kinds of data types that can be processed/stored in the system. Since we assumed that we had full control of the old system, this implies we have a method of classifying data types. It also means that we have a set of requirements connected to these classifications. If the new system introduces new data types, these will also have to be classified.

The next step is to map the expected data interactions to and from the new system. Our assumption stated that we have full control of how data is communicated to and from the old system. If these data connections must remain, we must document them as dependencies. If the new system requires it, these interactions will also have to be mapped and considered dependencies. This also applies if the organisation wants to introduce new functionality that requires either new connections to old systems or connections to new systems. Please note that the term "data" is very broadly defined in this context. For example, a logical access connection between two systems, even if it is not used, must be analysed and documented.

Finally, when we believe we have complete control of what data will be processed/stored in the new system and how those data will flow to other systems, and we have classified those data, we can conclude that the new system is fully analysed with regard to its data assets. We will then have to repeat this process for the next parts, i.e. the adjacent parts with which the new system will exchange data. This is where recursiveness hits, and a system with many dependencies will require analysis of many parts.

Assuming that this work is successful, we will have obtained documentation that tells us exactly which assets are involved. Given that the assets have been classified, and that we for each classification type have a set of information security requirements, we will know which security controls we must implement in order to protect those assets.

Assumptions vs. Reality

What has just been described is in itself challenging work, even with the simplifications and assumptions that were made. Unfortunately, these assumptions are, in my experience, unrealistic. Especially the second assumption, "we assume that we have full control of all aspects of the old system, including what kind of data it stores," is problematic. In fact, it is a generic assumption that hides many specific assumptions that we should consider. The most important ones can be listed as follows:

1. There is a framework and a tool for handling master data and data flows.
2. There is a framework for classifying all types of assets in the organisation.
3. With each classification type, there is a set of associated information security requirements.
4. There is a framework for risk management and risk analysis, which is a key component in decision-making.
5. Existing systems are well documented.

Amongst these points, extra consideration should be given to number 3. If you have not established standards for security controls in your organisation, this means that you are doing the same types of evaluations repeatedly but not necessarily with the same results. Obviously, all of the listed conditions must be satisfied in a mature organisation.

The process described above can primarily be viewed as a way to avoid setting the wrong requirements. Setting requirements is exclusively the organisation's responsibility. ▸

But let's assume that the organisation sets the right requirements. What challenges await then?

How to ensure that requirements are met

Even if we assume that the organisation has identified the correct requirements, we do unfortunately not have any guarantee that it is possible to implement these requirements in the given context. As described earlier, the general risk in this phase is implementing the correct requirements in the wrong way. Unlike setting requirements, implementation is a task for which the organisation and the vendor must share responsibility.

We can break down this challenge into three different questions:

1. Which security controls can we allow the vendor to own?
2. How can we verify that the controls are possible to implement correctly before the vendor is chosen?
3. How can we verify that the controls are correctly implemented before the system goes live?

Before taking a closer look at all three, it is worth spending some time on a very specific observation. Sometimes, you will hear statements like "we need to have a certain degree of trust to the vendor." The idea of having some trust can often intuitively make sense, since it is unreasonable to expect that all vendors can or want to prove absolutely everything to us. Some people might therefore conclude that there will always be a degree of trust, and that this is how it must be. Unfortunately, this conclusion is incorrect, and the underlying mindset is also problematic.

First, trust is binary. Either I trust you, or I don't. If I say that I trust you given certain conditions, what I really say is this: I accept the risk involved, given this specific context. In other words, there is no degree of trust involved. As a principle, trust should not be used as an argument for anything. The Zero Trust model, rooted in the principle of "never trust, always verify," is becoming an established model in information security. The zero trust mindset must also be part of the procurement mindset. We do not need trust; we need risk analysis. This is also aligned with the mindset that we see emerging in new laws and regulations such as GDPR and the NIS Directive. In any case, let's go back to the three parts of the challenge.

Which security controls can we allow the vendor to own?

This question is more or less already answered. The organisation needs a risk analysis of the involved controls to evaluate whether the vendor can own them. In some cases the answer will be "no", but in most cases the answer will be "yes, given..." Usually, this will be handled by expanding the security control, or adding a new requirement. An example of a control that could be owned by both the organisation and the vendor is a security log.

How can we verify that the controls are possible to implement correctly before the vendor is chosen?

This question can be very difficult to answer. In the early phases of procurement, the verification method can be very theoretical. Often, the organisation will only have the means to perform a documentation verification. This is OK for some systems, because the system is based on established standards and the vendor can refer to earlier implementations, etc. However, for many procurement scenarios, this is irrelevant, and in the worst-case scenario the procurement is unique. Simplified, we can say that two types of requirements must be verified: *governance requirements* and *technical requirements.*

Governance systems exist formally in documents and informally in culture. Documents are easily available and easy to evaluate, but they can be misleading since they are not necessarily being followed. Culture provides good verification of the actual mindset, but culture is challenging to observe. Another challenge related to the evaluation of a vendor's governance system is that the evaluation is often limited and dependent on the person who evaluates it. Let's assume a maturity scale from 1 to 5. If the person who is doing the evaluation has never witnessed nor had much experience with any system higher than a 3, that person will often not be able to separate clearly between the upper levels 4 and 5.

Technical requirements can obviously vary from very simple to very complicated. The challenge is that even the simplest requirement can require quite specific domain knowledge. You can be an expert on firewalls, while at the same time know nothing about secure use of containers. It is nevertheless often assumed that one single person can have the full responsibility for information security and possess the knowledge needed to handle all the requirements. This is in many cases an unreasonable expectation. Information security knowledge is too complex in each domain for one person to be able to know everything that is relevant. This complexity will continue to increase in the foreseeable future.

With regard to complex technical requirements, we can say that only tested experiences can give any degree of verification. However, in this phase we can rarely do tests and experiments. This must not be interpreted to mean that technical requirements are not important in this phase. On the contrary, it is these types of requirements that ensure that the organisation can legitimately cancel contracts if the system cannot deliver as expected.

How can we verify that the controls are correctly implemented before the system goes live?

As this is the last gateway before the system is set in

production, it is obviously important. It is also important to note that it isn't necessarily the final gateway for the whole process. Some controls can only be confirmed when the system is in production. This task can also be very different from the previous ones in some important respects. First, this task isn't necessarily a part of the procurement process. Depending on the internal setup, the organisation can handle this as part of their normal change management process. Second, this task should be defined long before it takes place, by describing the verification process as a set of requirements. In an ideal world, this level of description would be used for all tasks, but in my opinion doing so is neither justifiable nor possible.

Rather, we must, as always, fall back on risk analysis, identifying all critical controls based on risk scenarios and prioritising them accordingly. As much as circumstances allow, we include the verification process described as a requirement. An example of this could be penetration testing. The organisation, vendor, or an agreed-upon third party performs a penetration test based on pre-agreed scope and technology. The description of the process must also include acceptance criteria and describe the consequence of deviations. The worst-case scenario would be cancellation of the contract, regardless of how much work and money the organisation has spent on the system. Another example could be random sampling of selected areas, where a representative from the organisation sits down with someone from the vendor and evaluates critical processes or verifies the content of critical logs in real time.

It is worth noting that the task of describing acceptance criteria and similar tasks are a specific part of a much larger challenge. How does the organisation ensure that all relevant requirements are included in the vendor contract? The scope of this article does not allow us to delve deeply into that challenge, but much is solved if the organisation establishes contract standards. In some cases, vendors can be forced to use those standards, but if that's not possible, the standards can be used as a checklist to do quality assurance of the vendor's proposal.

When all necessary requirements have been verified, the system can be set in production. We are nearing the end, but a very important task still remains. Everyone involved must ensure that the relevant information is transferred to the line organisation, in a format suitable for doing life cycle management. All requirements, including the rationale for setting them, must be documented in the relevant document management systems. Follow-up of the vendor should already have been decided and documented in the contract.

Some final advice

In the end, I would like to say something about dependencies, because they are often the true cause of the largest challenges to information security. Security always needs to consider the whole picture, meaning that security cannot finalise its work before almost everyone else is finished. At the same time, security is one of the control functions that have the mandate to stop a proposed solution. If an organisation follows the traditional waterfall models for their procurement process, it is not unlikely that issues concerning the allotted time for the procurement will arise at some point. The worst-case scenario is that security gets involved at the very end and ends up sending a project back to its beginning. To avoid this, it is important that the organisation takes an iterative approach in its processes, and that everyone involved knows exactly what is needed to pass a gateway. Security and the architecture function must work closely together from the moment it is established that the business side has a need that can only be solved by starting a procurement process. Often, this is referred to as a portfolio or a program management process.

In a way, this piece of advice can be generalised, and is as such perhaps the single best piece of advice to any organisation that wants to improve their procurement process: Try to move as many activities as possible in your process to an earlier point. Where it is possible, you will often get a quick win, and in those cases you can't, you should be able to identify the root cause and at least have the information necessary to address the problem. Often, the root cause lays outside the procurement process in some adjacent process, and this process cannot be finalised at an earlier point and therefore acts as a bottleneck.

The observant reader might notice that this piece of advice resembles the process described earlier for analysing the value chain. In fact, the underlying logic is exactly the same, but instead of mapping requirements, we will primarily be mapping time dependencies and/or constraints. Let us therefore conclude with the first piece of advice on processes: A well-defined process for handling a change in the value chain must be applicable across all types of knowledge domains.

Try to move as many activities as possible in your process to an earlier point.

ARTICLE

Internet of Things and Its Firmware

A Tale of Memory Corruption Bugs

Matteo Malvica
Senior Security Consultant,
mnemonic

AFTER READING THIS ARTICLE, YOU WILL:

- Be familiar with IoT best practices for security testing
- Understand key concepts related to memory corruption vulnerabilities
- Have learned how to improve IoT firmware safeguards

The concept of *Internet of Things* (IoT) came from the idea of connecting ordinary things such as lights, doors, and other household appliances to a computer network to make them somehow smarter than they had originally been designed. When we think of IoT today, our mind immediately goes to a smart home appliance, digital fitness tracker, or even a fully connected car. While a world like the one we are shaping today would have seemed merely a remote possibility a decade ago, International Data Corporation (IDC)predicts that we will have 41.6 billion connected devices generating 79.4 zettabytes (ZB) of data in 2025[1].

Threats, regulators, and the firmware

With the sudden urge among legacy vendors such as IBM, Apple, Intel, and Cisco to invest in the IoT market, and the relentless creation of devices among new startups, IoT has become a highly lucrative target for attackers. Only when the Mirai botnet[2] wreaked havoc by infecting about 600,000 appliances did it become clear that the effort invested in securing the IoT spectrum was insufficient.

Fortunately, in recent years a number of regulators have come to the fore. Indeed, American organisations like GSMA and the Federal Trade Commission (FTC), as well as European associations such as ENISA[3], BEUC[4] and the Norwegian Consumer Council[5], have established compliance guidelines to ensure the safety and security of devices. The adoption of the new regulation on certifying ICT products in the EU would also provide an increased level of regulatory pressure[6].

One of the most notable common denominators across IoT devices that need to be kept vetted, supervised, and secured by all these entities, is the firmware. This article will explain what that is and why it is remarkably important to the whole security ecosystem.

IBM has developed a well-known model for IoT layers[7], which can be used as a reference for better understanding the security issues associated with IoT. If we look at this layered model from a security researcher perspective, it really catches our attention that it points to plenty of entry points for different attack methods.

With the aim of helping both vendors and consumers understand IoT security issues, the Open Web Application Security Project (OWASP)[8] has assembled a Top 10 chart of the worst IoT security practices, which can come in handy during any IoT security assessment. By mapping the risks and vulnerabilities within the IBM model with the OWASP IoT Top 10 chart, we can gain a thorough understanding of the IoT security posture. The mapping would provide a small subset of possible vulnerabilities: ▸

[1,2,3,4,5,6,7,8] *See Reference List at the end of the report*

IoT threats and the IBM Internet of Things model

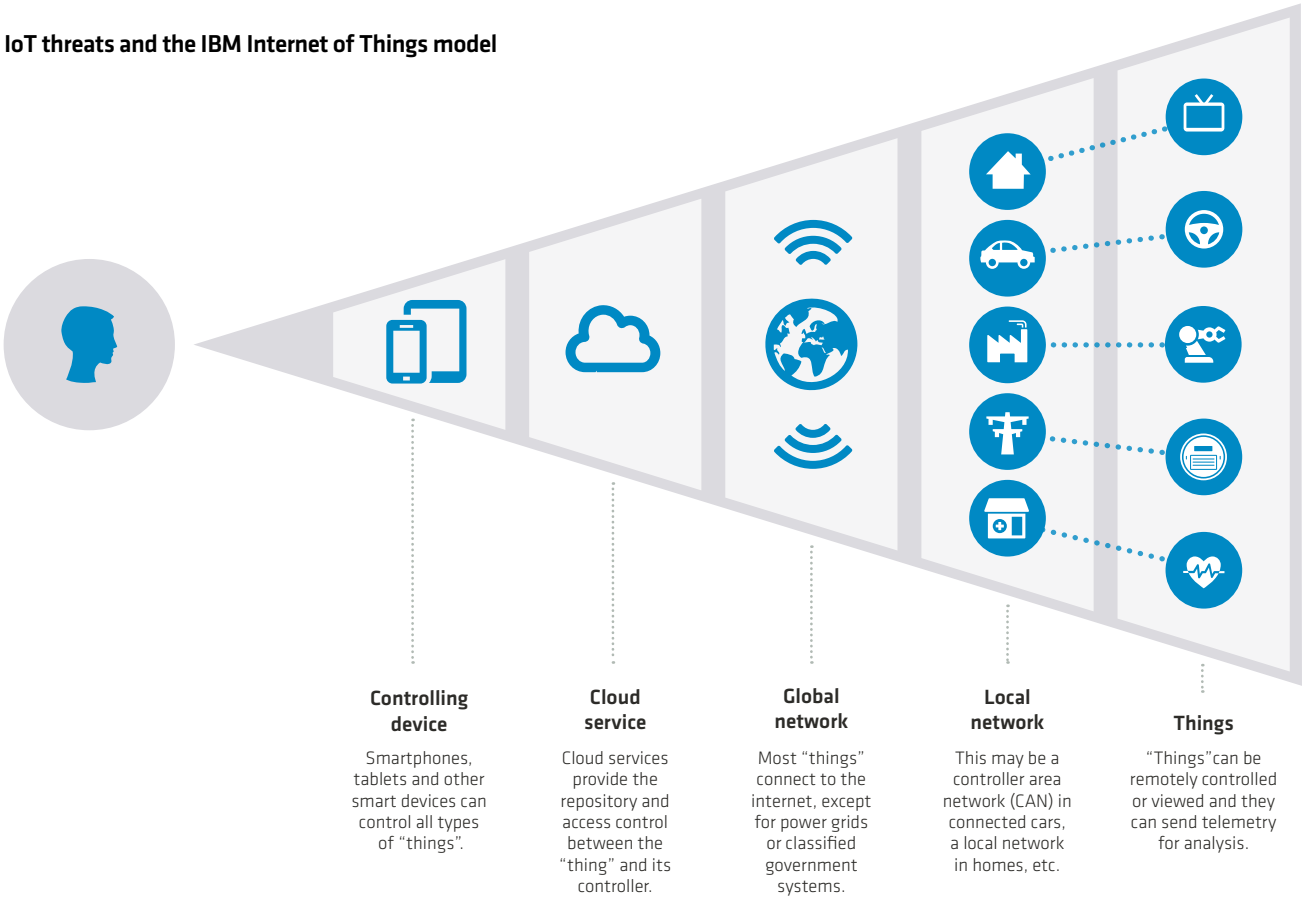

1. Controlling device
   - Insecure data transfer and storage
   - Use of insecure or outdated components
2. Cloud service
   - Insecure ecosystem interfaces
   - Insufficient privacy protection
3. Global network/local network
   - Insecure network services
   - Insecure default settings
4. Things
   - Weak, guessable, or hardcoded passwords
   - Lack of secure update mechanism

When performing IoT penetration testing, vulnerabilities are often found in each of the four layers. Covering all these security aspects would require several articles, so the present article will focus on the entity with the highest impact, which is probably also the least tested: the firmware.

Where do IoT vulnerabilities come from?

Along with the previously mentioned four-layered interoperability, every major IoT device vendor initially decided to come up with its own protocols and standards, not to mention their very own operating systems, built specifically for the IoT landscape. Some examples include Homekit by Apple, IoTivity by Intel, Brillo by Google, and Jasper by Cisco. The need for customer-base lock-in has historically been the main reason for adopting a proprietary software strategy, which obviously penalises the end-user by inhibiting interoperability with products from other vendors. Today, things have become more standardised, especially with regard to network protocol and wireless technologies.

However, many other important IoT elements are still subject to arbitrary selection. Chief among them is the operating system and its compressed version, the firmware. There are at present more than a hundred different variants of embedded operating systems to choose from[9].

This highly heterogeneous and customised ecosystem has generated an exceptional amount of code, which in turn means that extensive effort is required to maintain and keep everything secure. Today, an unwritten but widely acknowledged rule states that the cost of securing software is somewhat proportional to the complexity of the code. As a consequence, with an increased codebase volume and intricacy, the chance of software errors increases. While some bugs might result in a crash or denial of service (DoS) in the affected software, others, more severe ones, can possibly lead to an unauthenticated remote code execution (RCE) through a well-crafted exploit.

What is the status on IoT firmware security?

"Firmware" can be defined as an operation-critical code running on its very own hardware, interacting with the low-level components and having the dreadful reputation of being infrequently updated, possibly because it is physically infeasible to do so. Also, the inflexible nature of this type of software helps

[9] *See Reference List at the end of the report*

explain the "firm" in firmware. The firmware is the heart of the IoT device and what we consider "firmware" can be a number of things, from the Industrial Control Systems (ICS) software, to an embedded Linux operating system.

Today, firmware lives in everything from smartphones to embedded devices that are so conventional and ordinary that you might not even think that they are computer-based. That raises an important question: how can we smoothly update the firmware once one or more bugs are fixed and the patches available in the latest and greatest release? For instance, the firmware in a smart light bulb may not need frequent updates, but the firmware in a smart thermostat may need to be updated more often to stay compatible with smartphone operating systems.

An extensive survey related to this question, focusing on IoT firmware security, has been conducted by the Cyber Independent Testing Lab (CITL)[10], an independent non-profit organisation. CITL researchers studied a fair amount of firmware images and checked them for any presence of standard security features. After evaluating 6000 firmware updates over the past 15 years, the survey showed no improvement in the security and hardening among the major IoT router vendors. CITL researcher Sarah Zatko commented:

"We found no consistency in a vendor or product line doing better or showing any improvement. There was no evidence that anybody is making a concerted effort to address the safety hygiene of their products."

What is a memory corruption vulnerability and why is it critical for IoT devices?

A memory corruption vulnerability is an unintended state of a program's memory which may arise when memory content is modified due to programming behaviour that sidesteps the original developer's intentions. Today's operating systems and browsers are written in programming languages like C and C++ that have specific memory features to enhance runtime performance and, if used inaccurately, may lead to unforeseen programming behaviour.

To help prevent further exploitation of vulnerabilities, many operating-system-level hardenings have been designed, built, and rolled out on the major mobile and desktop devices over the last decades. For example, safeguards such as non-executable stacks, Address Space Layout Randomisation (ASLR), and stack cookies help mitigate one of the oldest yet regrettably current memory corruption bug classes: buffer overflow.

CPUs and their unspoken language

A common way to exploit memory corruption bugs is to directly interact with the vulnerability affecting the running program. From an offensive perspective, the most efficient way to take control of the program is to operate at the CPU level, typically through a language called "assembly".

Assembly is a CPU-specific language, and each processor has its own unique variant. Commodity IoT is typically built on top of well-known CPU architectures like Intel, ARM, or MIPS. Over time, these processor models have been extensively researched and understood by the security community, something that has greatly benefitted strategies geared towards finding and fixing vulnerabilities, as opposed to more closed-source and proprietary hardware running on top of different technologies, such as ICS. We can now take a closer look at the anatomy of vulnerabilities and the necessary building blocks for building a strategy to discover and prevent them.

When memory corruption leads to exploits

Developers write code that is eventually shipped as part of the final firmware package. Although not every bug is a memory corruption bug, the resulting code will inevitably contain some flaws. This happens due to multiple factors such as human error, vulnerable third-party libraries, or unforeseen race conditions that cannot easily be predicted through standard code reviews or unit testing.

These bugs cause a corrupted state in memory. Often, they simply lead to a harmless program crash; at other times, security researchers and threat actors alike can take advantage of a program crash by hijacking the intended logical flow and steering it to execute arbitrary code. This is obviously a risky aftermath from a security and integrity standpoint. ▸

Security researchers and threat actors alike can take advantage of a program crash by hijacking the intended logical flow and steering it to execute arbitrary code.

[10] *See Reference List at the end of the report*

Anatomy of a running program in memory

Though we wish to avoid getting lost in too many details, it is necessary to give a brief overview of how a program is mapped into computer memory. Regardless of the operating system, each program features the following runtime memory layout:

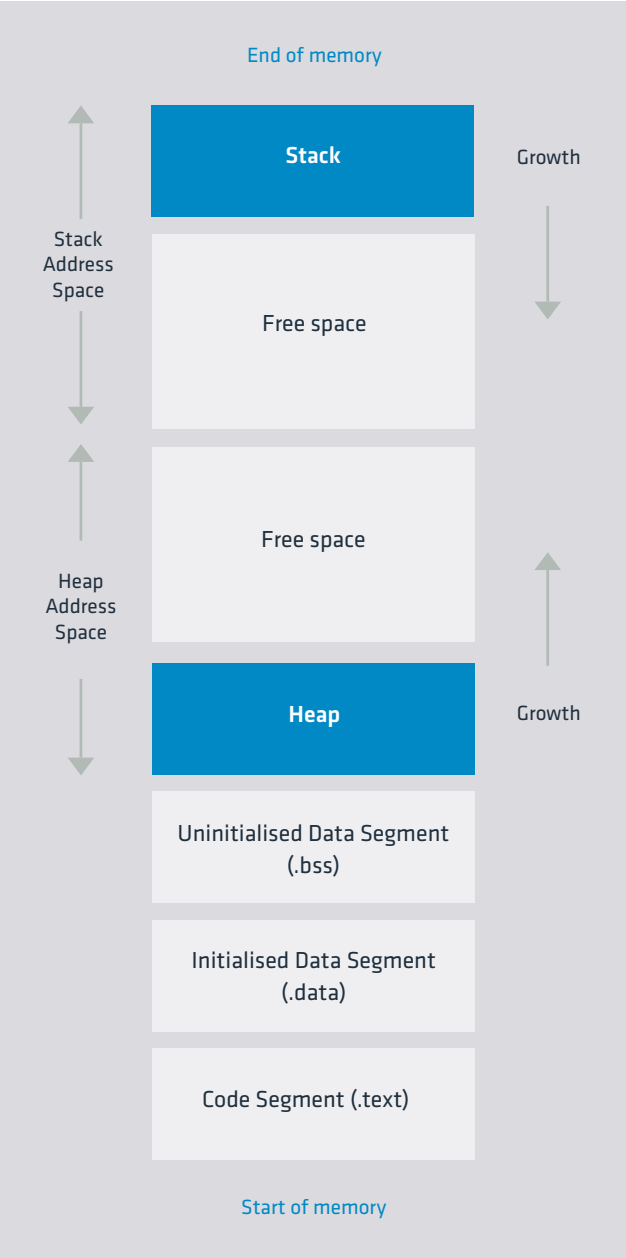

From this figure we can learn that both *stack* and *heap* are significant areas where user input will be stored. As a consequence, this is where most of the clashes and memory corruption are going to happen. These two memory areas are dynamic by design, and if a sensitive code region is overwritten, it could lead to one of the aforementioned scenarios: a crash, in the best case, or – worse – malicious code hijacking. As an example, a *program stack* might simultaneously contain both user data and program code. If the user input is not properly sanitised, it could potentially overflow into the code section and thus overwrite critical memory addresses responsible for the entire program behaviour.

Operating system mitigations to the rescue

As previously mentioned, history has taught us that because of multiple circumstances such as miscommunication, poorly documented code, software complexity, and development tools, vulnerabilities have regularly been incorporated as part of every development effort.Operating system vendors have thus developed numerous mitigations that can preemptively block malicious exploitation of preexisting flaws.

While these mitigations offer some kind of additional fortification, they only postpone the need for patching vulnerabilities, which should be the desired end goal. These memory safeguards can be summarised in the following table:

Mitigation	Purpose
Address Space Layout Randomisation (ASLR)	To make exploitation unpredictable, the memory address of a program is randomised at runtime.
Non-executable memory (NX)	Exploits are prevented from running by marking some memory areas, such as the stack and heap, as non-executable (read or write only).
Stack Cookies/Canaries	A randomly chosen value placed at the end of the stack that is checked before the function completes. If the value has been modified, it will force the program to crash and avoid malicious code execution.
Control Flow Guard (CFG)	A set of valid functions is pre-compiled and verified at runtime to prevent any malicious use.

The table above describes only the most important mitigations. However, many other defences have been developed to tackle corner-case attacks. If available, we recommend enabling them all at once, as these safeguards could dramatically impede an advanced and skilled threat actor. While it is common today to find these mitigations enabled on the majority of desktop and server vendors, the year 2020 is still looking obsolete from an IoT security perspective. Indeed, we can deem ourselves lucky if we, during our assessments, discover that two or more mitigations have been enabled.

Vulnerability discovery and automation

When a piece of software is open-sourced, meaning that its code is publicly available for vetting and scrutiny, it will greatly benefit from peer code reviews. This community-driven effort dramatically decreases the chance for vulnerabilities. Open-sourcing the codebase will not only be beneficial for the code reviewers, but will also greatly improve the quality and visibility of the *fuzzing process*. Fuzzing is an additional way to test a piece of software that involves sending unexpected or invalid data input to the target application. If the original program's code is unavailable to the fuzzer,

it will make it harder or even impossible to inspect and scrutinise every possible branch of the fuzzed software. A considerable amount of today's bugs are found through fuzzing avenues, and, insofar as we have the source code available, doing so will considerably boost the overall quality.

Nonetheless, most of the IoT firmware is shipped with proprietary and closed code, making it more difficult to find bugs through standard means. Alternative vulnerability researching approaches revolve around reverse engineering the firmware binary, where "reversing" means trying to decipher the original code by analysing machine-level assembly instructions. Unfortunately, due to the nature of a program compilation this process is a lossy one, making it unfeasible to fully restore the original developer's code.

As a matter of fact, when a software is compiled it loses most of the "human-related" data, such as variables and functions names, while preserving only the necessary parts that are needed by the CPU to understand the code and execute it efficiently. As a result, reversing can often be a time-consuming process, though it can sometimes be simplified and automated to grab the vulnerable low-hanging fruits, something that is usually good enough to raise the security bar.

A healing patch and update model

In addition to what we have learned so far, IoT vendors seldom alert their customers about new vulnerabilities and even more rarely automate the firmware update process. Up until recently, most devices sitting in a private network would most likely be left unpatched and exposed to security risks due to a lack of a self-update features. However, Over-the-Air (OTA) updates are slowly becoming a standard practice today, which means that IoT devices will soon be able to automatically fetch and install firmware updates without human intervention.

Even though it is crucial to keep an IoT device up-to-date, it is only one task among many on the broader security spectrum. For instance, when it comes to home routers Internet Service Providers (ISP) are typically the ones taking care of device security. Yet, this is one of the very few safety measures we have witnessed being put in place; security configuration and management duties are often left as a responsibility to the end user.

IoT security hardening best practices

- If available, verify that the device has an Over-the-Air (OTA) update mechanism enabled.
- Ask your ISP if they are taking care of your home routers management, updates and security configuration.
- Ensure that the cloud solution adopted by the device is widely recognised and audited.
- If the IoT device is handling sensitive data, consider contacting a security expert to perform a thorough analysis.

Awareness as the first step

Vendors, security researches, and CERTs report IoT vulnerabilities on a daily basis, and this trend does not appear to be abating anytime soon. We have thus learned that, because of their simplistic nature, embedded devices are more prone to attacks and exploitations due to the poverty of mitigations and code review currently being put in place. As educated consumers, we should be aware of the potential risk that unsecure IoT devices pose to our lives and personal data, especially when they are equipped with different kinds of environmental sensors such as microphones, cameras, or GPS. Consequently, awareness is the first step to start challenging the IoT security status quo.

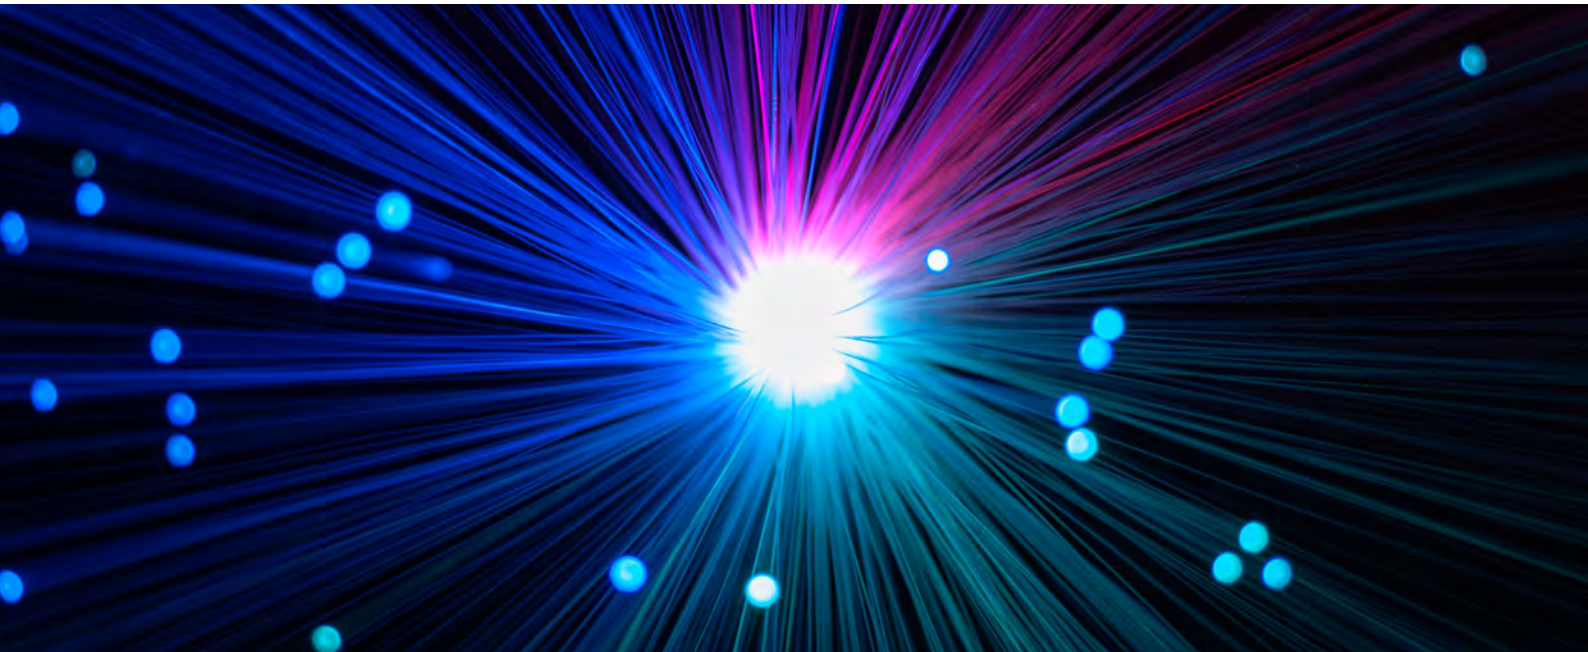

STOREBRAND


Norway

Bjørn R. Watne
SVP, Head of Group Security (CISO)

The Storebrand Group is a leading player in the Nordic market for long-term savings and insurance. The company manages more than NOK 750 billion, making Storebrand Norway's largest asset manager.

What is your biggest cybersecurity concern?

The threat landscape within cybersecurity is changing daily, and my concerns change accordingly. There are, however, some "universal truths" around that I never stop keeping an eye on. The days are gone when people punched in at eight AM in the same location five days a week, and when a company's data were confined to servers in the basement. Rapidly changing technology, increased outsourcing and offshoring, agile development and flexible work hours are the norm these days, and these developments are accompanied by a whole new set of security concerns. Thus, if I were to pinpoint my biggest concern these days I would say that it relates to information protection and access management – how we ensure that the right information is available at the right time only to the right people. For the Storebrand Group, our current value chains span ten countries over two continents, and the vendor landscape is always increasing. Making sure that security is tight from A to Z in this multi-layered web is something that requires constant attention.

In what areas of cybersecurity do you think we're falling behind?

In my opinion *falling behind* is not what we're doing – we are catching up! I've been working in cybersecurity for close to 20 years now, and we have never been as good as we are now. That being said, there is definitely still room for improvement. There are two things I would say are of equal importance to me to focus on going forward. First, there is *visibility*. To be able to detect and manage incidents we need transparency and visibility in our networks. Zero-day attacks and advanced persistent threats require that we constantly look for anomalies, and running around blindfolded would give us nothing.

The second is about *people* and *processes*. Yes, technology is important, and we need to have the right tools to do the job. Still, a hammer won't build a house – you need both the carpenter and the blueprints. When addressing cybersecurity, it is of utmost importance that we address both employees and ways of working in addition to technical tools.

What gives you hope for the future of cybersecurity?

The fact that cybersecurity is climbing the risk ladder and making its way into the boardrooms gives me hope. Management has to be on board if you want to get things moving, and cybersecurity is definitely not an "IT problem" managed by the guys over at the technology department. Cybersecurity needs to be aligned with the business, and that requires both strategy and budget, as well as a clear commitment from the top management. Looking back just five years the situation was completely different. Another thing that gives me hope is watching the younger crowd of newly graduated people entering the stage at conferences speaking about cybersecurity. Go back 15–20 years and security wasn't even part of the curriculum for an engineering degree in computer science. Fast forward to today, and you can even get your PhD specifically in cybersecurity. ●


COMFORT HOTEL

ARTICLE

The Last Piece of the Puzzle: Incident Readiness

Alexander Paulsen
Security Consultant,
mnemonic

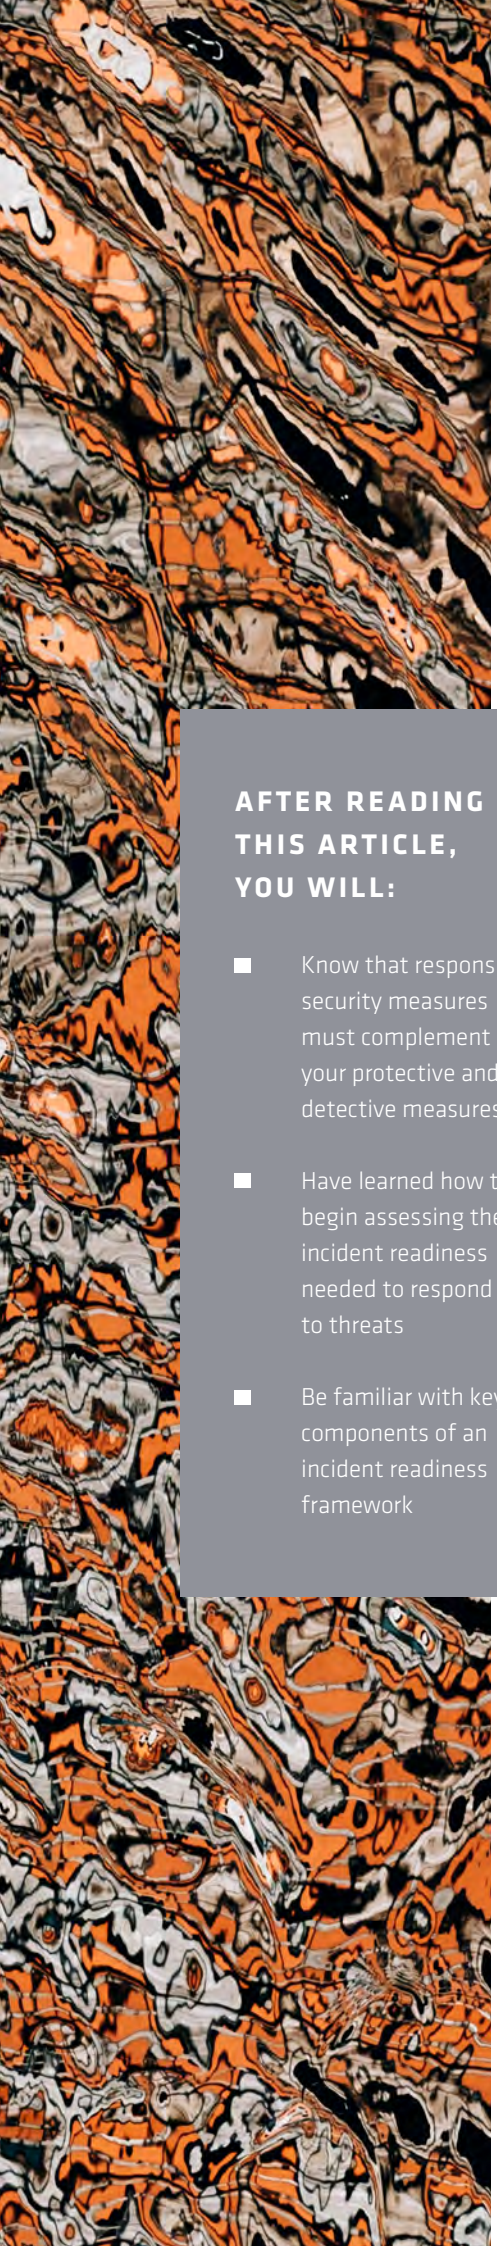

AFTER READING THIS ARTICLE, YOU WILL:

- Know that responsive security measures must complement your protective and detective measures
- Have learned how to begin assessing the incident readiness needed to respond to threats
- Be familiar with key components of an incident readiness framework

Presumably, your organisation has implemented numerous protective measures, including network segregation, firewalls and perhaps a proxy or two. Additionally, the most critical data are encrypted at rest and in transit, and the employees might have undergone some awareness training. However, you still don't feel completely secure and have therefore adopted detective measures, such as log collection from key systems and networks, and ongoing analysis of that data is, hopefully, being done. Moreover, you know that employees can be socially engineered or that other vulnerabilities can be exploited, and that is why capabilities to detect what cannot be stopped are essential. As a result of detecting suspicious or malicious activity, you must always follow through on responsive measures.

Despite investing in a lot of security controls, the attacker got in. What do you do? Your organisation will realise that responding to security incidents requires routines and procedures that should have been established and thoroughly tested. Picking up the pieces of what organisations instinctively do while facing security incidents can be demanding, especially when you realise that there are numerous pitfalls. As an attempt to organise and solve this issue, this article will introduce an approach to implementing sufficient incident response capabilities. ▸

Cybersecurity framework

The three functions *protect, detect* and *respond* can be used to make a simplified security model based on the NIST Cybersecurity Framework:

- **Protect:** Supports the ability to limit or contain the impact of a potential cybersecurity event.
- **Detect:** Enables timely discovery of cybersecurity events.
- **Respond:** Supports the ability to contain the impact of potential security incidents.

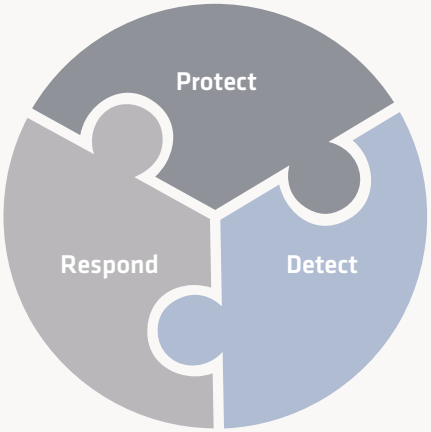

Planning for success

Author Alan Lakein says, famously, that "failing to plan is planning to fail." Another great quote from him is "planning is bringing the future into the present so that you can do something about it now." Obviously, nobody plans to fail, but if you believe that there is a chance that security incidents will happen, why wouldn't you plan for it? Planning for future incidents is a keystone to achieving security objectives when incidents occur.

So, what are the objectives of security? To *prevent* security incidents from happening? An unachievable objective, if you ask me. Take into account that security incidents can and will happen. Also, realise that we have not failed in our work when they do. However, we must be prepared to demonstrate that we have a structured approach to achieving one of our security objectives, namely to contain the impact of security incidents in order to minimise damage to operations, reputation, loss of market share, finances and more. Keep in mind that all eyes will be on us, and that expectations will be high.

When incidents happen, it is vital to have a plan in place for incident handling. Be aware that there is no such thing as the "perfect plan" or "one plan fits all," so what you need to do is to develop the right plan for your organisation. These are the main benefits of planning:

- **Increases efficiency:** Efficient handling of incidents is essential. The potential impacts increase dramatically as time goes by, and having a plan to lean on when things go wrong will certainly reduce the blast radius.

- **Facilitates coordination:** When major incidents happen, one or more teams will have to coordinate with each other to investigate and draw the big picture for the decision-makers.

- **Gives direction:** Without planning, no one will know what to do or when to do it. Planning helps us do the right thing at the right time.

Before starting the planning phase and defining the capabilities you need, you should know what you are protecting, and who you are protecting it from. Performing a business impact analysis and a threat actor assessment will help you answer these questions.

Business Impact Analysis

The business impact analysis, which identifies critical functions of the organisation and assesses the consequences of future events, might be the closest you get to predicting the future. Perhaps you have already done this analysis, as it is a recommended input to several security activities. The output is the potential impacts of security disruptions and recovery requirements. Incidents affecting key services will require significant attention and resources, and after performing this exercise, you will know which ones.

Threat actor assessment

You wouldn't have gone into the boxing ring without knowing something about your opponent. Is he or she a southpaw (left-handed boxer stance) or an orthodox (right-handed), does he or she have any signature moves, and what are his or her height and stamina like? Similarly, you should assess your cyber opponents to figure out who you are up against and get to know their skills, motivations and persistence. If you are up against Tyson, you better be holding your guard up.

Establishing a baseline

The outcomes from the business impact analysis and threat actor assessment will provide the intelligence necessary for establishing a baseline for the capabilities you need for incident handling. Organisations facing potentially disastrous events as a result of being targeted by an advanced and persistent threat actor will surely need greater capabilities for handling incidents than those who do not face the same level of risks.

Remember to keep your eyes on your security objective while having an ongoing security incident, namely to minimise the impact. Your next step is to establish and operationalise the capabilities defined in your baseline, and the rest of this article will present key principles which should be included in a framework for incident handling.

The incident framework

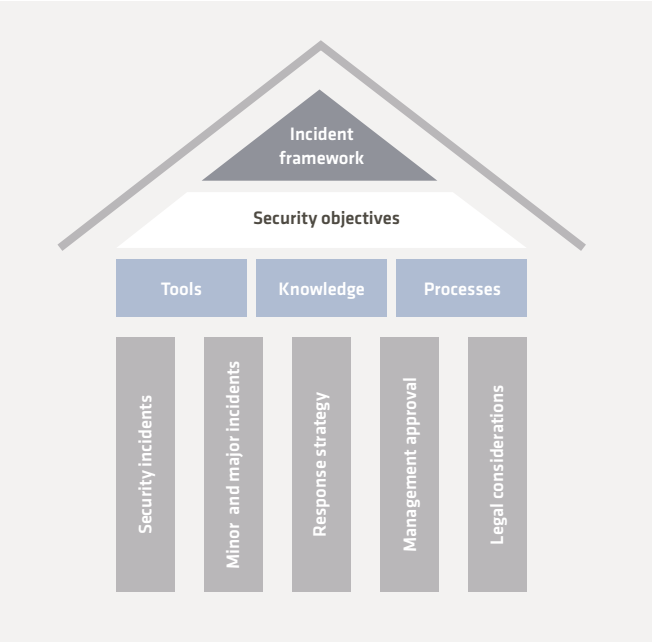

The framework pillars

The main pillars of your security incident handling framework should be established to provide a solid foundation, and in particular to make sure that the first critical hours and minutes of a detected security event or incident are handled appropriately. The five pillars presented below should be defined to ensure good conditions for future incident handling activities. ▸

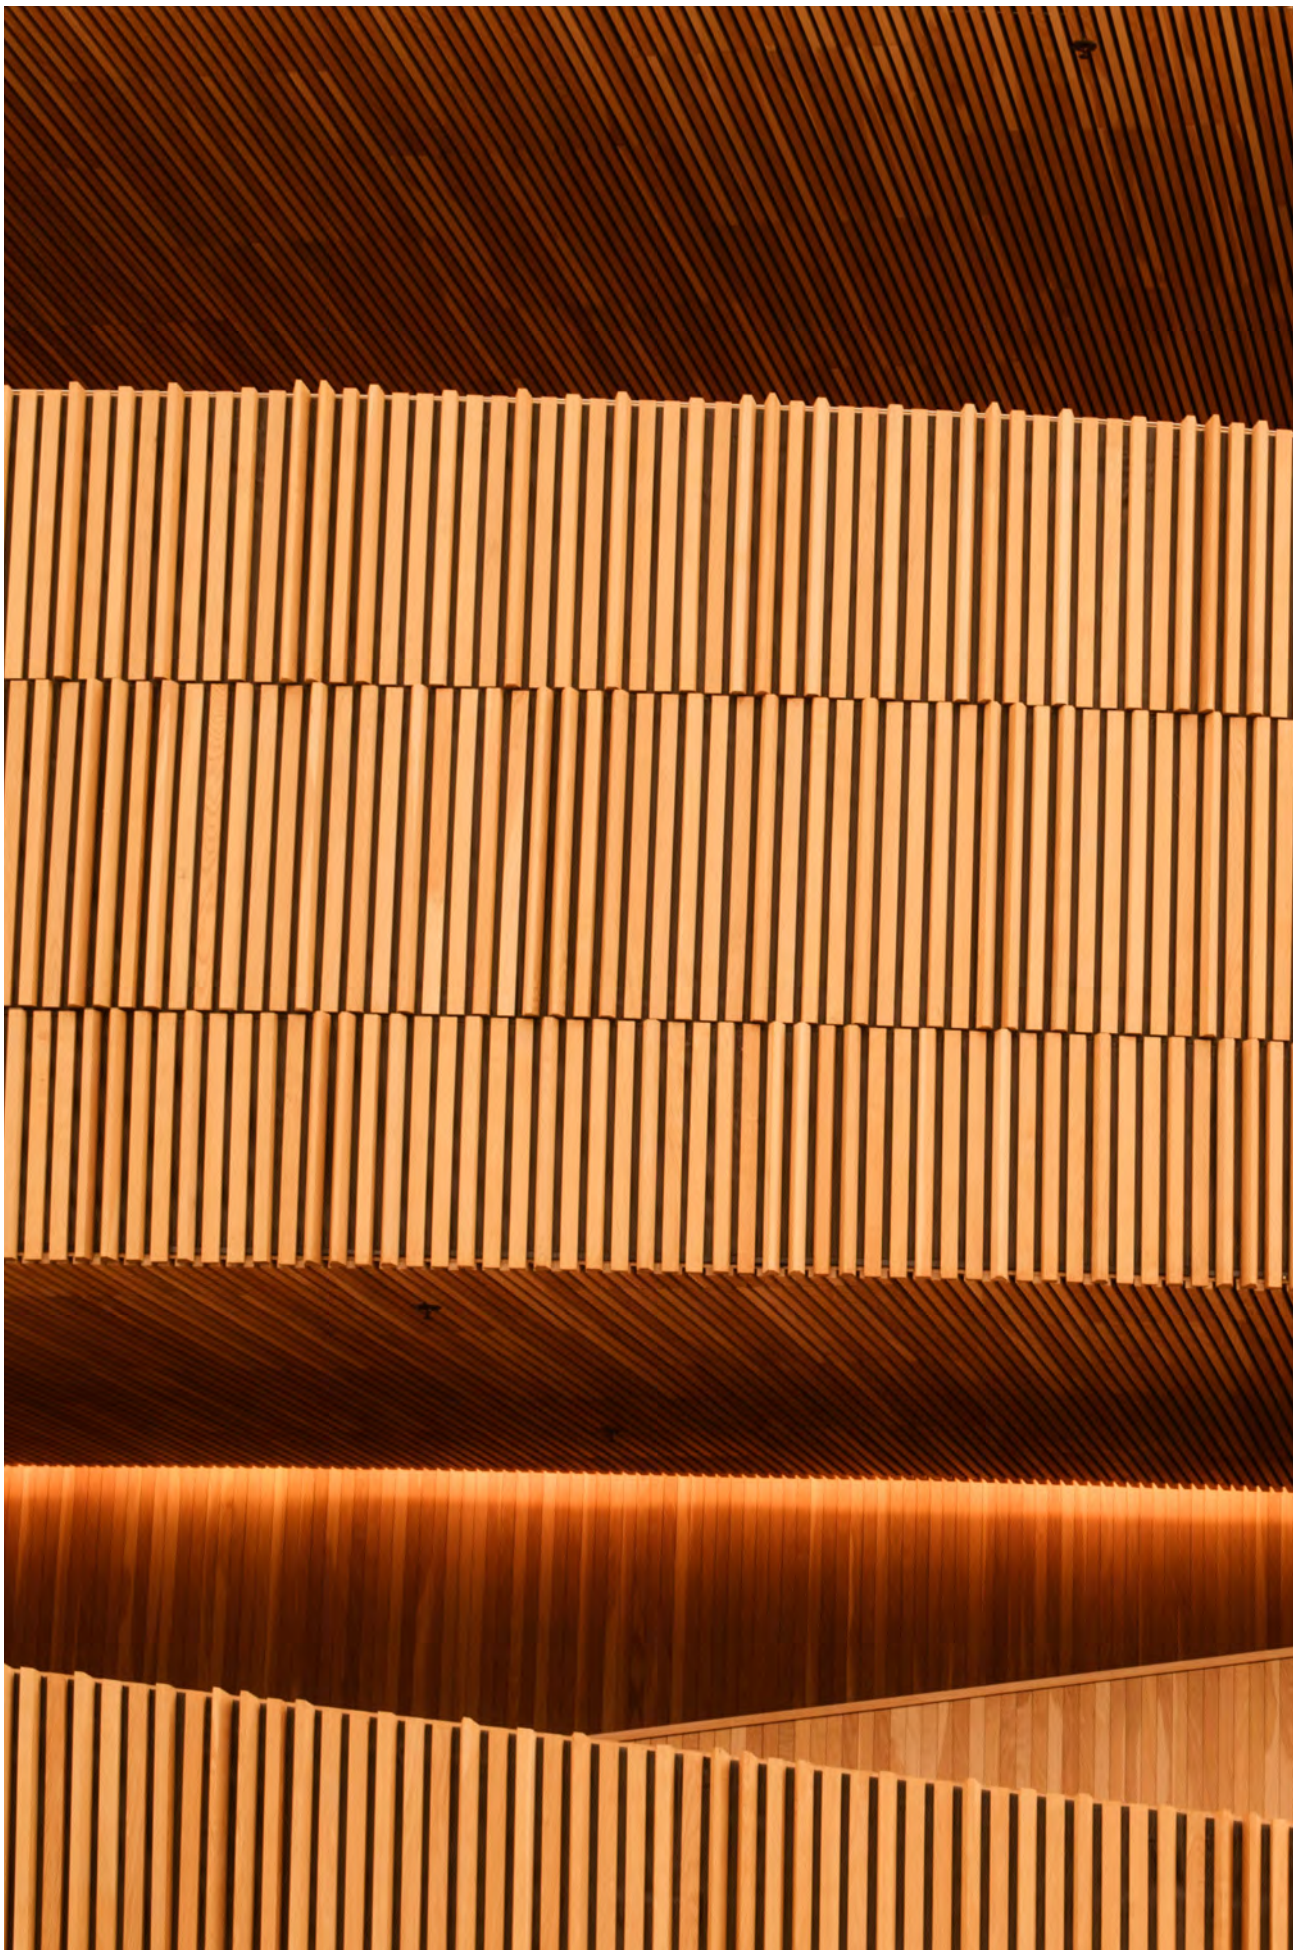

Defining "security incident"

A security incident could be any disruption to information, systems, applications and services. Regardless of what the incident looks like, you should know how it differs or relates to other established phrases such as "incidents," as defined the in the Information Technology Infrastructure Library (ITIL), or "personal data breaches," as defined by the General Data Protection Regulation (GDPR). The International Organization for Standardization (ISO) defines a security incident as one or more unwanted or unexpected information security event(s) that have a significant probability of compromising business operations and threatening the confidentiality, integrity and availability of information. The National Institute of Standards and Technology (NIST) describes security incidents as threats risking violation of security policies. I recommend finding the definition that best suits your organisation. Ultimately, you and your team should be able to categorise the incident at the time it is detected and recorded. Failure to do so can lead to incorrect responses.

Minor and major incidents

Just like a musical composition, a security incident should be categorised as minor or major. Minor security incidents are normally handled by a system and a network administrator or a security analyst, whilst major incidents are handled by a group of people: the incident response team (IRT). You should be able to define criteria for what security events fall into which category based on your understanding of the criticality of your information, systems and services as well as the types of security events you can identify and analyse. Underestimating the nature, scope and potential business impacts of a security incident can lead to improper handling, which in turn will significantly increase the impacts.

Response strategy

After a security incident is detected, there are two immediate response strategies: either to *watch and learn* or to *contain and clear*. The latter might seem like the most tempting option of the two, but in many cases it is not the best one. Imagine chasing a burglar out of your home in the middle of the night. You think you got rid of him, but the burglar knows something you don't: your terrace door can easily be opened from the outside with basic lock-picking skills. The following night he won't be doing the same mistakes that got him caught the night before. Consequently, you should know the movements of the threat actor and the scope and magnitude of the compromise before you make any hasty decisions. However, if you catch him right at the entrance, feel free to land some devastating punches.

Management approval

The senior management needs to be informed about your response strategies and their corresponding risks. At a glance, and in the middle of the heat, they will most likely demand immediate action to make the threat disappear. As explained in the previous section, what they need to understand is that attempting to remove a threat today can lead to a more threatening situation tomorrow, and maybe even a threat going undetected.

Legal considerations

The last pillar involves legal considerations such as evidence handling or notification requirements. For example, if you need to search for content on employees' personal spaces, such as their e-mailboxes, special requirements apply in certain jurisdictions. Furthermore, incidents should be notified to a competent authority if the organisation and the incident in question are subject to GDPR, the NIS Directive or a national security legislation.

The framework elements

When all of the important pillars are in place, you can get to the core of what this is all about, namely *how* you are planning to do the actual incident handling, *what you need* and *who you need* to carry out the incident handling itself. I have divided this into three vital elements.

Tools

Needless to say, tools are required to handle security incidents. Either you have the tools yourself, or you need to know how to get them in a timely manner. Tools include applications for threat hunting and forensics, threat intelligence platforms, log analysis systems, network detection systems, packet capturing services and more. The whole toolbox is probably not required, but you should at least know where to go and what to do in case you need tools you do not already have. What you should have, however, are workstations for your incident response team. Keep in mind that you might end up in a situation where normal workstations cannot be trusted. This also applies to communication channels such as e-mail.

All major security incidents should be documented thoroughly, and you will therefore need a tool for this as well. Remember that your SharePoint might be compromised, and you should know what other options you have.

Knowledge

Yet, what is the point of having all these tools if you cannot operate them? In order to benefit from the tools, you need skilled people. Unfortunately, skilled and specialised security personnel don't grow on trees. You might, however, have network and system administrators who know your systems and networks inside out. Maybe you have an employee responsible for backup and perhaps even a legal counsellor and a corporate spokesperson. All of these are useful resources that you may need to have on your incident response team. For specialised competence, many rely on external expertise who may also bring their own tools. You should get to know your options and who to call.

Processes and procedures

The armed soldiers are now in place, but now you need them to stay in line. The Roman Army was led with discipline and structure, which was crucial for its success. Most recognised standards, such as NIST or ISO, define the incident management process more or less like this:

Prepare ► Detect and report ► Contain, eradicate, recover ► Lessons learned

You should always carry out *all* of the abovementioned steps. Too often the first and last steps are skipped. Preparation constitutes all the things you have read in this article and more. One of them is frequent incident exercises, which are strongly recommended. Major security incidents are not likely to happen frequently, and that is exactly why training is essential to keep the knowledge alive. Lessons learned are meant to give feedback to your own framework for incident handling, but also to other parts of the organisation where security gaps have been detected as a result of an incident.

Finally – and maybe this is the most important part – how do you conduct the information gathering, analysis and decision-making? I suggest building your work around a cycle called the OODA-loop (Observe, Orient, Decide, Act), a concept used at the operational level of military campaigns.

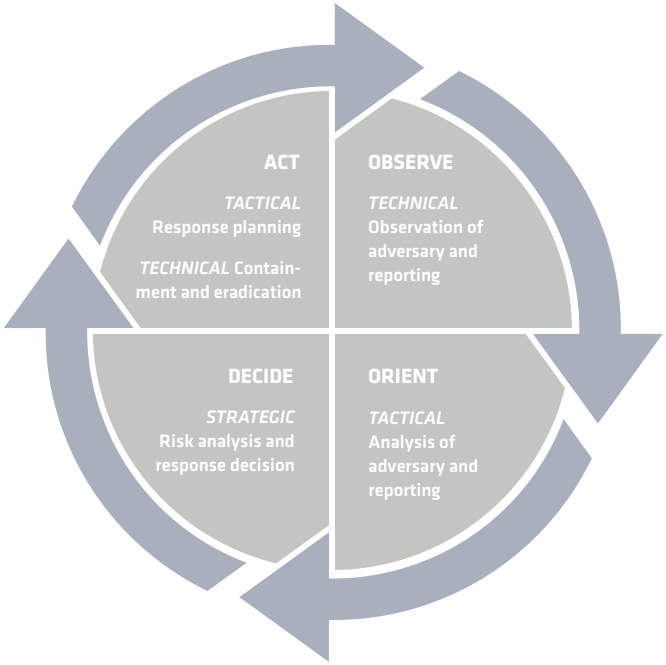

1. At the technical level, your technical analysts *observe* the operations carried out by the threat actor.

2. At the tactical level, your tactical analysts *orient* on what the threat actors want to achieve with their operations. For instance, dumping information from a database to a file can be a technique to stage the data for exfiltration.

3. The strategic level is where the major *decisions* are made. Based on the technical and tactical analysis you can assess the risk of the threat actor's operations. The decision concerns whether you *watch and learn* or *contain and clear*.

4. The decision is further communicated to the tactical level, who makes a plan to *act* on the decision. Finally, the plan goes to the technical level where it is carried out.

Final stage

Although the project for incident planning is finalised, the work does not end there. What you don't want to happen is that your great work on planning becomes a paper-only exercise. The policies, requirements, activities and more must be followed and materialised. *Now* is when the real capability-building begins.

John F. Kennedy once said that "there are risks and costs to action. But they are far less than the long-range risk of comfortable inaction." Inability to respond to an incident can make all your security investments go to waste. Protection mechanisms have to succeed every time, while attackers only have to succeed once. In other words, your protective measures must at all times be ahead of the attackers' tools and techniques. If they ever get so lucky as to get in, you will have to detect the threat, then respond to it – every time. Typically, the organisation invests in protective and detective security measures, only leaving the budget for responsive measures to a minimum. That is why incident readiness really is the last piece of the puzzle. ●

Inability to respond to an incident can make all your security investments go to waste.

ARTICLE

The Value of Outsourcing Detection and Response

Making Informed Security Decisions

Tommy Steensnæs
Senior Security Consultant,
mnemonic

Outsourcing IT services is common practice for organisations everywhere. The majority of organisations already outsource significant parts of their IT services to third parties[1,2], and the remainder are likely to have at least considered it. Cost is still one of the key reasons for outsourcing[3], but other value-related reasons are becoming more important.

According to the Global Sourcing Association (GSA), outsourcing will continue to grow in the coming years[4]. However, GSA also claims that outsourcing has an image problem[5]. Personally, I have heard multiple stories that corroborate that claim, including stories about outsourcing leading to critical vulnerabilities and about costs escalating out of control.

From my discussions with customers, it seems that many are not completely satisfied with their outsourcing partners. The general feeling is that they don't get the flexibility and proactivity they expect or are promised, and there are far too many hidden costs. This article will not dive into a general discussion of the problems with outsourcing, but rather focus on how some security challenges may be successfully solved by outsourcing.

Making outsourcing decisions is hard. The general advice related to outsourcing is to keep core business in-house and outsource the rest. But what does that mean in practice? How does this advice relate to outsourcing security? Is security considered a core business?

Every organisation is unique but looking holistically at security challenges I think most would agree that there is a lot of ground to cover. The field of information security is vast, sometimes even intimidating. As a CISO or someone else who is responsible for the information security in an organisation, having the capacity to cover everything is extremely challenging and probably even unrealistic. I don't have all the answers, but I will address some of the common challenges and provide some advice with a main focus on detection and response.

AFTER READING THIS ARTICLE, YOU WILL:

- Understand some of the key challenges for CISOs and other security professionals
- Have gained insight into the elements of detection and response that will benefit from outsourcing
- Have knowledge that is helpful for making informed decisions about what to outsource and for choosing a security partner that fits your organisation

Breaking Down The Problem

Business priorities and security objectives

Whether they work for a small local company or for a multinational enterprise, most security professionals I have talked to are in a constant state of having too many tasks and not enough time to do them. Another way to put it is to say that organisations are under-resourced. Ultimately, that means that some tasks will not be completed, at least not to a satisfactory level. How do we ensure that the right tasks are prioritised? The key is to look for the driver of prioritisation, as it is easy to become trapped in security concerns and personal convictions and lose sight of the organisation's common goals.

Consequently, one fundamental task for anyone responsible for an organisation's security is to understand the business objectives, key ▸

[1, 2, 3, 4, 5] *See Reference List at the end of the report*

processes, values, and drivers. Once understood, these should be translated into key business drivers for security. The key business drivers for security could be seen as the high-level security objectives that will help enable the business to reach its goals. Generally, a thorough understanding of the organisation's core business is necessary to make informed security decisions.

When security professionals understand the business priorities and how they translate into security objectives, the next step is to identify the main risks and evaluate mitigating controls. Cybersecurity risk is an integral part of the total business risk, and handling cybersecurity risk has a direct impact on the success of business initiatives.

The process of identifying risks and prioritising controls should result in a comprehensive security strategy, and the organisation can apply frameworks such as the NIST Cybersecurity Framework, the ISO 27000-series and guides such as Center for Internet Security's CIS Controls to help structure and implement the security strategy. The ultimate goal of the security strategy is to minimise the consequences of security incidents to the business. Although the standards operate with different categories, these can be divided into three high-level categories: preventing incidents from happening, detecting and responding to incidents, and recovering after an incident has occurred.

Challenges when running an efficient security operation

Currently, there is a significant gap between supply and demand for security professionals. The actual number of unfilled positions varies between sources, such as (ISC)[2] [6] and CSIS[7], yet all agree that the number of unfilled positions will continue to grow and in Europe alone will be in the hundreds of thousands within 2–3 years. Knowing that the demand for cybersecurity skills will only increase, it seems essential to retain the talent you already have. A study from ISACA shows that 64% of the respondents have challenges related to retaining their cybersecurity specialists[8].

When looking for threats in an organisation's infrastructure, data overload and alert fatigue are amongst the biggest challenges. Cisco's annual Cybersecurity Report 2018 points out that only a small percentage, around 20%, of alerts are legitimate[9]. Of these, less than half are actually mitigated. One reason for this is that security analysts are spending too much time on triage, and not enough time on actual response. The goal should be to move the time spent on triage towards time spent on response and mitigation. Ideally, hunting for threats that are not detected by other means should also be prioritised.

With challenges related to both recruiting and retaining talent, security resources must be given tasks that are motivating and aligned with their competence. These challenges are highly interconnected. If you are struggling with recruiting, it is likely that, regardless of your security budget, there are tasks in the organisation that are not being done simply because there are not enough people. The high demand makes it easy for security professionals to change jobs and chase more challenging opportunities and higher salaries. These people will not accept that their job is to look at alerts that are mostly irrelevant, knowing that their chance of discovering anything important or interesting is minimal.

Outsourcing detection and response

The detection of and response to security incidents in an IT infrastructure, or OT infrastructure for that matter, is a cornerstone of most security strategies. This involve a multi-step process consisting of various tasks. There are also related functions that are important inputs, outputs, and control functions for detection and response. We'll start by addressing the different phases of this process and discuss if and why they could be outsourced. Later, we'll explore guidelines for how to make outsourcing decisions.

In short, detection and response could be divided into two distinct phases, each with three parts, as shown in the figure below.

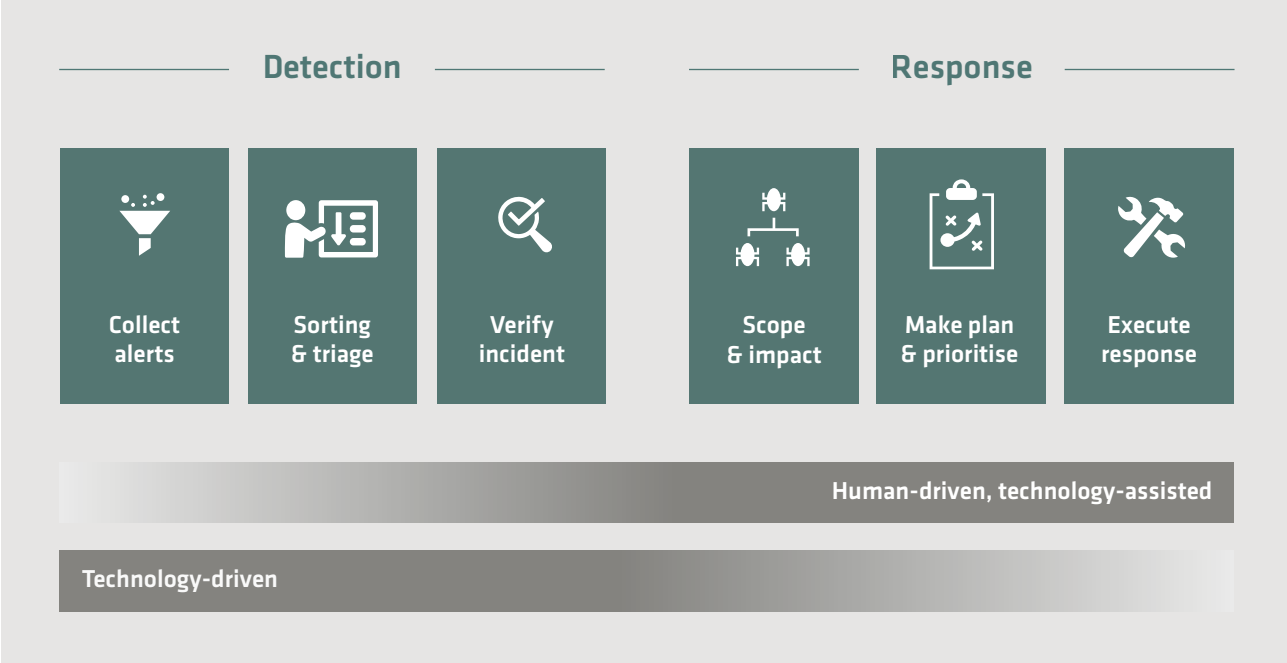

Detection

Detecting threats is a cumbersome task that requires technology and expertise, and the task involves reviewing an insurmountable number of events looking for the few that matter. This is where the challenge of alert fatigue comes in.

To enable detection, we need to collect data. Most of the collected data will not be relevant from a security perspective and should be filtered out from further analysis. The sheer amount of data makes it evident that this filtering needs to be automated. This requires the right technology, perfectly tuned for this task. There is a significant risk involved in filtering: if it is too strict, important events will be missed; if it is too open, however, the analysts will be overwhelmed and miss the important events. To achieve proper filtering, constant tuning that requires both knowledge and the right tools is needed.

When irrelevant events are filtered out, the analysts need to look at the remainder, verify incidents and decide on what to respond to. Even with proper filtering, this is very demanding. The analysts will need an excellent understanding of different threats and be able to make decisions quickly. In order to make decisions efficiently, the analysts need as much context as possible, including context for the threat, the business itself, and the overall threat landscape. This should all be supported by the tools the analysts use.

The entire process should be executed continuously 24/7. Such a capability is at best costly and, given the challenges discussed above, maybe even impossible to build and maintain in-house. It requires custom tools that need to be continuously tuned to be effective.

Additionally, the analysts will have a relatively stressful and monotone job and be prone to become bored and make mistakes, and may eventually quit their jobs. A professional security partner will have the economy of scale, both with regard to building and maintaining tools, and with regard to being able to provide analysts with more varied tasks. Detection of potential threats is therefore a task that could be outsourced as a whole.

Response

A verified security incident needs to be responded to, which triggers the response phase. The goal of the response phase is to mitigate the incident and move to recovery as quickly as possible. However, for mitigation to be effective, we need a more thorough analysis. The difference between analysis in the detection phase and the response phase is that the first looks to verify that *something has happened*, whilst the latter is intended to answer *what exactly has happened*. Knowing what happened makes us able to choose the correct mitigating actions. ▸

[6,7,8,9] *See Reference List at the end of the report*

As the business impact can be too high in some cases, the response phase can never be fully outsourced. For example, can an outsourced partner be trusted to make the sole decision to take a critical business process offline because it is involved in a severe security incident? However, if the entire response phase is internalised, one easily ends up with too many detected incidents with too little context to be dealt with efficiently.

Incidents will have varying degrees of severity and business impact, and it makes sense to spend most resources on incidents with the highest severity. For most security professionals, this is motivating work, and being able to spend time on incidents that matter will give internal resources a more meaningful workday. To achieve this, some of the burden of response could be put on an outsourcing partner. They can perform most of the analysis and scoping, leaving only the last part – the part that requires internal knowledge – to be covered by the recipient. Accordingly, enough context would be provided and alert fatigue could be eliminated, while leaving ownership of important decisions with the organisation.

In many cases, low-severity events can be dispatched directly from a security partner to the internal operations team. For pre-described incidents, the security partner could even perform the mitigation.

Business risk and threat picture

Detection and response are actions taken in response to risks that cannot be realistically mitigated by prevention. Risk and threat are highly interconnected as some understanding of threats is required to fully understand business risks. The unmitigated risks are important input to the detection and response processes. However, detailed knowledge about the threat landscape is even more important. Anyone delivering detection and response services, internally or outsourced, must have a clear picture of the risks and threats the organisation is facing.

Risk should always be owned internally, but an outsourcing partner can benefit from economy of scale and most of the time provide significantly better threat intelligence. Building the procedures, acquiring data sources, collecting and analysing data, and applying new intelligence require specialised skills and are time-consuming tasks.

In particular, if detection and response are outsourced, it makes sense to outsource parts of the threat intelligence as well. In that case, an ongoing dialogue between the outsourcing partner and the risk owner is necessary. From an internal standpoint, it is also important to ensure that the understanding of the risks is unified across the organisation.

Threat hunting

Threat hunting is related to detection in that the goal is to find threats. The difference is that it looks backwards to discover indications of a threat that was initially missed by detection. This is a discipline that requires knowledge about threats, infrastructure, and the organisation. To hunt effectively, you need an extensive skillset and a proper methodology. Whether to keep this capability internal or outsourcing it depends on a few factors.

Again, it becomes a question about resources. If you have the capacity to dedicate people to this task, it could be kept internal. For many organisations, though, the tasks will require too many resources to perform successfully, and it will be beneficial to outsource this capability.

A specialised security partner will have the knowledge and skills to build scenarios for hunting that can be hard to achieve internally. They will have to rely on the organisation to provide business knowledge when analysing findings.

Alternatively, it is possible to choose a hybrid model where internal capabilities hunt for specific threats, while the security partner hunts for others. In this case, it will be critical to establish a well-functioning partnership between the capabilities in order to extract the full value.

Cost considerations

While cost shouldn't be the only driver for outsourcing, there is a significant potential for cost savings in solving these challenges by outsourcing.

Below I have listed areas with high potential for cost savings:

- **Recruitment:** New hires are expensive, both when it comes to searching for candidates and training new people. Likewise, the knowledge that is lost when an employee leaves and the disruption to other team members while a replacement is found and trained, involve significant costs. Outsourcing some of the functions that are most prone to turnover will save significantly on recruitment.

- **Head count:** Reducing the number of employees will in most cases provide significant cost-saving benefits. This is particularly true when it comes to running a 24/7 operation.

- **Research, development, and maintenance of detection and response tools:** In order to efficiently detect and respond to incidents you will need continuous development in both technology and processes to support the 24/7 staff. This cost is frequently overlooked, but research and development, and maintenance of detection and response tools, are necessary to be successful. This cost spans IT operations and security operations.

■ **Training:** Maintaining and improving the skills of your security resources is a critical success factor in running a security organisation. Reducing the number of people, and the tasks and systems they are responsible for, reduces the cost of training. Yet it is still important to train the people you keep in skills needed for their responsibilities. In addition to improving the quality of the work they do, training might help retain them as well.

Making the right outsourcing decisions

Now that we've covered some of the primary challenges and their suitability for outsourcing, how do we arrive at the final decision of what is right to outsource for any specific organisation?

Finding the right balance between in-house and outsourced capabilities comes down to the business drivers for security. As discussed earlier, the security strategy should point to necessary capabilities. Any capabilities that require intimate business knowledge or policy decisions, or which involves actions that can have a direct impact on business processes, should always be internalised. Business knowledge can be shared, and a suitable outsourcing partner will strive to understand the business, but never gain the same understanding as the business owners. How these concerns translate to specific capabilities will vary between organisations.

There is no exact answer as to what should be outsourced and what has to be kept internal. However, I have arrived at a methodology to help you arrive at a conclusion that is suitable for your organisation. There are a few steps to this methodology that will be described in the next section.

Assign outsourcing tasks and choose a security partner

Self-assessment

Before deciding on what to outsource and what to keep in-house, it is important to know where you are. Based on the security strategy, identify which capabilities and skills are required to implement your strategy. Note that this may change over time.

Define ambition

When you know where you are, the next step is to decide where you want to be. Again, this should be based on your security strategy. There may be capabilities that are entirely lacking, or some of your capabilities may need training to reach the right maturity. It may even be the case that you have some excess capabilities that can be decommissioned.

Identify GAP

When you have identified both your current situation and where you want to be, it should be reasonably straightforward to establish the current gap. This gap is what you need to fill, either with internal resources or by outsourcing.

Evaluate capabilities for outsourcing

Based on the gap you have identified, it should be possible to specify which tasks should be outsourced. With the security strategy and the business drivers for security in mind, start by working out which tasks and capabilities should not be outsourced. By that I mean that you should focus on key competencies that can be achieved internally, and which will be central in supporting the business. Anything not in this category will be candidates for outsourcing. In this process, make sure you are realistic when evaluating your internal capabilities and the capacity to improve.

Choosing a security partner

Depending on what you decide to outsource, you may look for different capabilities and traits in your security partner. However, if you are outsourcing a significant part of your security process, you may want to consider looking for a partner with a broader set of capabilities. Your requirements may change over time, and your security partner must be able to accommodate these changes.

One of the goals of outsourcing is to lighten the load on your internal organisation. An important factor to consider in this context is what kind of output you will get from this service. The output should be precise and complete enough to actually solve your challenges. A security partner should be able to augment your internal security organisation, not just provide it with additional work.

Another thing to consider is technology and competency. It is important to evaluate the competency of the analysts and other personnel, as well as the technology used in the service delivery. All of it should fit your strategy and help close the gap for you to reach your ambition. Research, development and continuous training are indicators that the security partner will also stay relevant in the future.

No matter which partner you choose, make sure that you have a clear picture of what you expect from the service. Without clearly defined requirements, or at least a well-aligned ambition, there is a high risk that the service will not provide the expected value. In my opinion, it is a worthwhile exercise to also define what you are *not* getting. Doing so helps ensure that you don't miss important aspects and that you enter into a partnership with clear expectations on both sides. When you choose a security partner, be sure to establish clear lines of responsibility.

Finally, always plan for the future. Keep in mind that if you have an agile security strategy, your requirements may change. While the future is unknown, one absolute certainty is that the threat and security landscape will change and evolve over time. Whatever the future brings, adaptability will be crucial, no matter whether you rely on internal or external capabilities. ●

ARTICLE

Integrating Security Controls Within a DevOps Pipeline

Morten Marstrander
Senior Security Consultant,
mnemonic

Espen Hovind
Senior Security Consultant,
mnemonic

AFTER READING THIS ARTICLE, YOU WILL:

- Understand the role of automated and integrated security controls in a DevOps pipeline
- Have received an overview of the various categories of security controls in a pipeline
- Be able to choose security controls that will provide quick wins

Since 2009, DevOps has been adopted as the standard software development methodology by a growing number of companies. Pioneered by early tech unicorns such as Etsy and Twitter, DevOps and agile principles have contributed to meeting the ever-present business requirement *short time to market*. Arguably, the reason for DevOps's popularity is that it enables enterprises to produce software at a faster rate than traditional software development methods.

However, conventional security controls can drastically impede the high delivery rate that DevOps offers. The need for security is still of vital importance given rapid changes in ephemeral and complex production environments. To meet this challenge, the organisation needs a new approach to security from both a cultural and a technical perspective.

Adapting an organisational culture for Secure DevOps is a prerequisite for starting the technical process of implementing security controls. In short, collaboration and communication are essential. While security teams are concerned with risk, they should also meet the developers halfway to a solution without compromising security. They should accept that implementing security into a DevOps workflow introduces a new approach to mitigating risk and a different set of security controls. On the other hand, the developers should embrace security controls as enablers for delivering timely changes to their applications. Ideally, one or more of the developers should transition into the role of a *security champion* and establish a productive relationship with the security teams.

From a technical perspective, the well-known and established principle of a *layered security approach* still applies. The implementation of a *layered security approach* in Secure DevOps is based on the concept of automating security controls in a pipeline containing sequential stages. The pipeline and controls are defined entirely as code, adhering to the principle of *Everything as Code*. The pipeline is already a part of the DevOps toolchain, which aligns with the practice of using a unified set of automated tools and processes across the DevOps teams.

Due to the complexity of modern production environments and the availability of numerous security controls, it is crucial to know where to start. To determine this, one should have a clear understanding of the available options within the Secure DevOps pipeline. We suggest an approach in which you first select and implement controls that produce quick wins. Moving forward, you choose additional controls until you reach an acceptable level of risk tolerance.

The next pages showcase our recommended approach and some food for thought when implementing Secure DevOps in your environment. ▸

DEVOPS FLOW AT A GLANCE

When deploying changes to production, the code has to traverse stages in the DevOps pipeline. A stage includes different types of security controls that examine the code, dependencies, and other parts of the deployment. If the test results are within a defined risk threshold, the process proceeds to the next stage. The changes are deployed to production if the code is verified through all the stages.

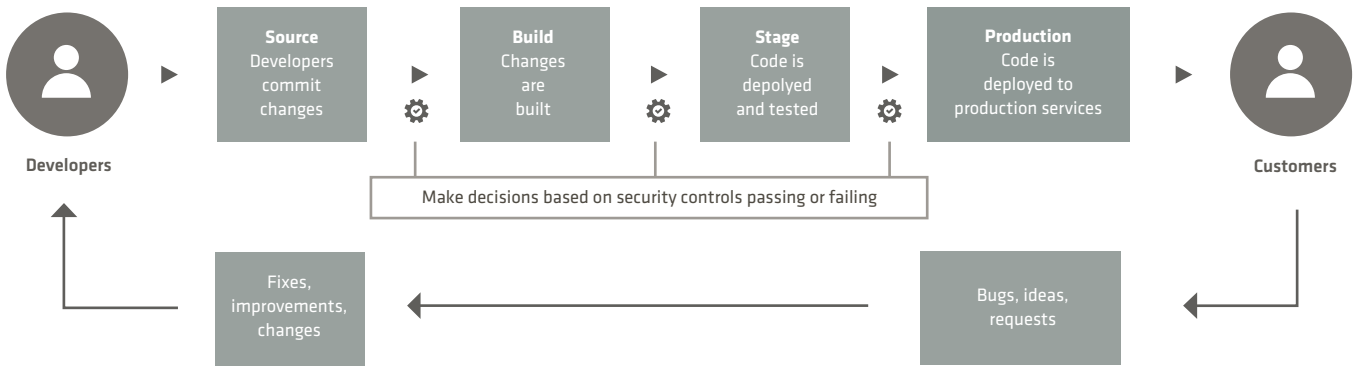

The Secure DevOps Pipeline

The figure below describes the different stages in a Secure DevOps pipeline, its various security controls, and concrete examples of implementation.

SOURCE
Before and when developers commit source code

BUILD
Automated build and continuous integration

The Paved Road

Help the developers to make the right choice.

The "Sec" in SecDevOps can advise and collaborate with the team on how to configure built-in framework security features, define compliance requirements, and configure risk thresholds. Consider using frameworks like the CIS controls. Apply secure coding standard principles.

Static Application Security Testing (SAST)

Detect bugs or security issues in source code.

The available tools range from special purpose code analysers like Bandit (Python) and gosec (Go) to advanced modelling and query frameworks like Checkmarx and Semmle. SAST is considered a powerful but complex security control.

Rapid Risk Assessment

Perform risk analysis when changing high-risk code.

Utilise open source frameworks like Mozilla Risk Assessment and Microsoft Threat Modelling Tool to assess risk as part of the DevOps workflow.

Dependency Analysis /Software Components

Establish a secure supply chain.

Manage known vulnerabilities introduced by dependencies in application code or container images. Consider tools with extensive framework support and high-quality vulnerability databases.

Code Reviews

Peer review of code before merging into release branch.

Educate the developers on the benefits of introducing merge requests as part of the workflow from a collaboration and security perspective. Consider tools like Atlassian Crucible and GitLab.

Security Unit Tests (high-risk code)

Scanners are usually not able to detect flaws in business logic.

Complement automated scanners and penetration testing by writing security-focused unit tests. Use language libraries like unittest for Python or JUnit for Java.

Pursuing the quick wins

The Secure DevOps Pipeline includes a myriad of tools and security controls which may seem overwhelming to plan and implement in a holistic manner. As a starting point, we recommend implementing a subset of these to achieve quick wins. The following tools and controls in the pipeline provide a significantly improved security posture, ease of implementation, and non-intrusiveness in a production environment:

The Paved Road:
Harden your environment by enabling security functionality.
Examples:

- Define HTTP security headers (excluding Content Security Policy).
- Enforce input validation in your development framework.
- Utilise Pod Security Policies for Kubernetes to prevent privileged containers.

It's all about YAML, and Everything as Code

The security control below is implemented with the *as Code* principle. The output from the dependency scan can be verified against conditions on whether to fail or pass the deployment process. The results can be displayed in monitoring systems or stored as audit logs.

```
- job: Application_Scanning
  dependsOn:  Build_Application_Image
  steps:
    # Scan application for known vulnerabilities in libraries
    - bash: |
        echo "Run Snyk Scan" ;  mkdir $(reports)

        # Install snyk.
        npm install snyk ; export SNYK_TOKEN=$(SNYK_TOKEN)

        # Scan application dependencies.
        snyk test --json > $(reports)/$(snyk_scan_report)

        # If HIGH vulnerabilities are found. Exit the step with a non-zero error code.
        jq '{vulnerabilities} | .[] | .[] | {severity}' $(reports)/$(snyk_scan_report) | egrep '(High)'
        if [ $(echo $?) -eq 0 ]; then exit 127; fi

    # Publish an artifact with the scan results.
    - publish: $(reports)/$(snyk_scan_report)
      displayName: 'Publish image scanning tests results'
      artifact: $(snyk_scan_report)
```

STAGE
Continuous delivery

PRODUCTION
Code is deployed to production

Dynamic Application Security Testing (DAST)

Simulate bad traffic destined for running web applications.

Asynchronously scan your application in a stage environment using automated vulnerability scanners like Burp Suite. Differentiate on passive or active scans and consider the effects of the availability and integrity of your data.

Runtime Protection

Ensure visibility, compliance, forensics, and monitoring of running code.

Deploy Sysdig/Falco, or Palo Alto Prisma Cloud to protect the application runtime, enhance visibility, and create audit trails in container and serverless environments.

Penetration Testing

Automated scanning cannot replace penetration tests with human interaction.

Perform periodical security testing to discover flaws in architecture, configuration, and business logic. Access to the source code produces a better result.

Secrets Management

Securely store and access secrets.

Azure Key Vault and AWS KMS provide storage services for secrets. The primary benefit is the native integration with other cloud services from the same vendor. Alternatively, deploy dedicated platforms like HashiCorp Vault or CyberArk. These solutions may be better suited for multi-cloud environments.

Automated Security Attacks

Automate security testing using customisable frameworks.

Create attack scenarios by using open source frameworks like Gauntlt, utilising external tools such as sslyze, nmap, and sqlmap.

Continuous Security Monitoring

Increase your ability to detect and alert on security incidents.

Use centralised logging of all data sources, adding the ability to correlate events and graph metrics. Combined with a runtime protection tool, this will further your forensic capability.

Dependency Analysis:
Scan your container images and analyse your applications dependencies for known vulnerabilities.
Examples:

- Use Snyk or Aqua Trivy to assess dependencies.

Continuous Security Monitoring:
Monitor your production environment for unauthorised events.
Examples:

- Use Sysdig Secure or Palo Alto Prisma Cloud for low-level visibility and protection of container environments.
- Utilise Palo Alto Prisma Cloud for monitoring and protecting serverless applications.
- Deploy tools like Splunk for log centralisation and SIEM functionality.

The road ahead

The implementation of security controls into a DevOps workflow is crucial for maintaining the security posture of your applications while still being able to deploy new functionality rapidly. When doing so there are some considerations to take into account.

Firstly, fine-tuning the security controls and defining the thresholds may be easy for some types of controls due to the lower complexity and intrusiveness of the control itself. As an example, consider extensive security monitoring which yields increased visibility into the environment while not affecting the deployment. By contrast, advanced static application security testing (SAST) is inherently complex on large code bases and prone to produce false positives. In addition, the scan results may be difficult for the developers to interpret. Simpler SAST tools may produce useful results at a lower cost but probably not the same level of security. Subsequently, there is often a correlation between the complexity of the implementation and the achieved level of security.

It is also important to consider performance, which is highly dependent on architectural and implementation details. A general recommendation is to utilise pre-built containers that execute the actual tests. The ephemeral nature of containers may, however, be an issue for security controls that require initialisation of their environment. As an example, deploying dependency analysis tools with extensive vulnerability databases requires significant start-up time and should be deployed as a long-running process. There is also a relationship between the rigorousness of the test and the elapsed time. For example, dynamic application security testing (DAST) should often be implemented as an asynchronous process, considering the time required to finish the test.

Nevertheless, Secure DevOps should be viewed as an enabler of frequent and secure deployments. By utilising the power of automation, APIs, *Everything as Code*, and portable data definitions like YAML and JSON, Secure DevOps can provide efficient security testing and create valuable audit log trails. This, in turn, can help organisations satisfy their compliance requirements.

The implementation process should be iterative, and gradually enforce security. Selecting and implementing technical security controls requires knowledge about available tools and how they compare and complement each other. The goal is to reach a balance of performance, usability, and security, and experience is an important factor in achieving it.

Moving forwards, one should utilise the strengths of a DevOps culture, communication and collaboration. This leads us to one of the fundamental principles: that the integration of security into DevOps should be a joint effort between “Dev,” “Ops,” and Security. ●

2019: A VIEW FROM MNEMONIC'S SECURITY OPERATIONS CENTRE

All statistics are collected from the analysis of nearly 4 trillion security events and over 25 000 real customer cases detected in our Security Operations Centre.

WHEN ARE SECURITY INCIDENTS HAPPENING?

It is no surprise that security incidents continue to occur 24-hours a day. The highest volume of security incidents occurs during regular 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 often cause security incidents.

The peak of security incidents between 08 – 09 is likely attributed to users logging onto their computers when first arriving at the office and quickly working through their collection of email from the previous evening. Our observations have repeatedly shown that users are more prone to inadvertently clicking malicious links, opening hostile attachments or visiting suspicious websites when they are tired, hungry, or likely to be paying less attention to individual emails, such as clearing their inbox first thing in the morning.

82%

targeted attacks occurred during business hours of 07 – 16.

39%

targeted attacks occurred between 07 – 10.

TIME OF DAY

THE TALE OF TARGETED ATTACKS

The vast majority of targeted attacks were detected between the business hours of 07 and 16. Due to the nature of targeted attacks, it can be expected that attackers will target their victims when these users are most likely to be online.

We have repeatedly observed that users are most likely to click a malicious email when returning to their computer after some extended break – whether it is first thing in the morning, or returning to their desk after lunch. In 2019 this behaviour spiked between 07 – 10.

INCIDENTS WITH UNAUTHORISED ACCESS TO RESOURCES

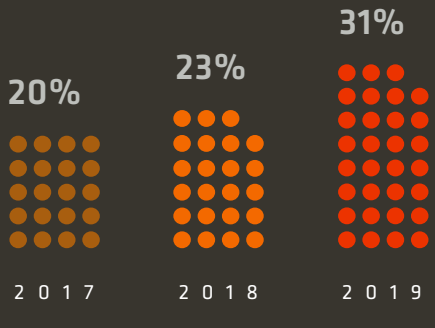

WHAT DAYS ARE SECURITY INCIDENTS HAPPENING?

Security incidents occur every day of the week, though as expected, there is a significant increase on weekdays. The slight decrease in incidents occurring on Mondays, Thursdays and Fridays can most probably be attributed to users working less frequently on these days than when compared to Tuesdays and Wednesdays – namely due to public holidays and users taking vacation days that fall on either side of the weekend.

TIME OF WEEK

6%
SUNDAY

17%
MONDAY

7%
SATURDAY

16%
FRIDAY

19%
TUESDAY

17%
THURSDAY

18%
WEDNESDAY

PASSWORDS ARE A HOT COMMODITY

There is a rising trend in security incidents where a user or attacker has gained unauthorised access to some resource. This behaviour was observed in 31% of incidents in 2019, up from 23% in 2018 and 20% in 2017. This increase is partially connected to the continued adoption of cloud services, and attackers targeting cloud environments where they will see more success using stolen credentials from users and admins than attempting to bypass the security controls of the cloud providers themselves. Attackers continue to see success with reconnaissance scanning and wide-spread exploitation attempts on vulnerable services exposed to the Internet, both those hosted on-premise and in the cloud.

Despite regulations like GDPR that promote security awareness in the application development lifecycle, we are observing an increase of incidents where usernames and passwords are being transmitted in clear text. We have observed a tripling of such cases since 2017. This can partially be attributed to the increase in IoT devices, which all too often consider security as an afterthought in favour of producing convenient and low-cost devices, and the increase in the use of cloud services.

ARTICLE

The NIS Directive

A Step in the Right Direction

Julian Fongen
Jurist,
mnemonic

AFTER READING THIS ARTICLE, YOU WILL:

- Have learned about the actors subject to the NIS Directive
- Have an overview of the security requirements imposed on operators of essential services and digital service providers
- Understand some of the differences and commonalities in the implementation of the Directive in EU countries

Protecting critical infrastructure is an important duty and a strategic task for any sovereign country. Disruption of critical infrastructure could cause inconveniences and financial losses for society, and destruction or incapacitation of infrastructures could eliminate the country's capability to protect itself from external threats, resulting in social unrest, significant economic harm, and even loss of life[1].

Directive (EU) 2016/1148 of the European Parliament and of the Council of 6 July 2016 concerning measures for maintaining a high common level of security of network and information systems across the Union (NIS Directive) is the first EU-wide legislation on cybersecurity, and entered into force 9 May 2018. Prior to the implementation of the Directive, the existing security tools and procedures were not sufficiently developed or common across the EU, something that established a rationale for constructing a comprehensive regulation at the Union level.

PURPOSE OF THE DIRECTIVE

The NIS Directive has three main objectives: improving national cybersecurity capabilities, building cooperation at the EU level, and promoting a culture of risk management and incident reporting. The purpose of the Directive is to achieve a high common level of security of network and information systems in the Union. However, the Directive allows Member States to voluntarily adopt further obligations that would help achieve a higher level of security.

The Directive claims to promote a culture of risk management, in which "risk" is defined as "any reasonably identifiable circumstance or event having a potential adverse effect on the security of network and information systems."

According to the Directive, "security of network and information systems" is:

"The ability of network and information systems to resist, at a given level of confidence, any action that compromises the availability, authenticity, integrity or confidentiality of stored or transmitted or processed data or the related services offered by, or accessible via, those network and information systems."

ACTORS SUBJECT TO THE DIRECTIVE

The principle of *lex specialis* is applicable to the Directive, meaning that sector-specific regulations imposing security requirements that are at least equivalent to those of the Directive will take precedence over the Directive. ▸

[1] *See Reference List at the end of the report*

Operators of essential services

An operator of essential services (OES) is a public or private entity which is *essential* for the maintenance of critical societal and/or economic activities, and is *dependent* on network and information systems. A potential incident affecting this entity has to result in *significant disruptive effects* on service delivery.

Firstly, whether the service is *essential* has to be assessed individually, but typical examples are provided in Annex II. However, the Member States are permitted to go beyond the scope of Annex II and include additional sectors and sub-sectors. Secondly, that the entity is *dependent* on network and information systems is a criterion many EU countries consider obvious, while others perform assessments of whether there is a potential dependency. Typically, essential sectors rely on such systems. Thirdly, the potential incident affecting the entity must result in *significant disruptive effects*. The assessment of whether the potential incident results in such effects is based on:

- the number of users relying on the service
- the dependency between the potential essential service and other essential services
- the impact that incidents could have
- the market share of the entity
- the geographic spread of an incident
- the importance of the entity for maintaining a sufficient level of service

Types of OES in Annex II:

- Energy (sub-sectors: electricity, oil and gas)
- Transport (sub-sectors: air transport, rail transport, water transport and road transport)
- Banking
- Financial market infrastructures
- Health sector
- Drinking water supply and distribution
- Digital infrastructure

A report from the EU Commission[2] illustrates that methodological approaches vary significantly among Member States, including which authorities that shall identify, assessments of the dependence of network and information systems, the definition of OES, and the application of thresholds.

The consistency differences are a result of the different implementation of the Directive and the minimum harmonisation approach, but it does not entail that Member States have implemented the Directive incorrectly.

The degree to which the identification process is centralised varies between EU countries, but the most common practice is to delegate some of the process to sectoral authorities, and give a single authority the responsibility for providing guidance to sectoral authorities. Such practice seems logical as usually sectoral authorities understand their sub-sectors better than the main authority. Furthermore, the identification process may be either a top-down approach, in which public authorities perform the identification process, or a bottom-up identification, in which operators themselves determine whether they are considered OES or not. Although the top-down approach seems to be most common, the authorities are dependent on some self-assessment exercises from the potential OES. Finally, as part of the identification process, the EU countries have to assess the OES' dependence on network and information systems, and what is involved in this evaluation varies between countries conducting detailed assessments, and those referring to the potential OES to self-assess their dependence.

Definition of OES and application of thresholds

The number of identified services varies between Member States both in terms of the total amount of services, but also on the amount of individual entities in each sector, which corresponds with the degree of granularity across the Union. To ensure convergent implementation, definitions of sectors should be applied similarly, as significant variations between EU countries can lead to an uneven playing field between OES across the Union.

The figure on the next page describes the inconsistencies on defining OES in the EU, and shows that some States (Estonia, for instance) have chosen a broad and general definition, which opens up the possibility for basically identifying any operator in the electricity subsector as an OES, while Bulgaria, on the other hand, identifies OES based on a very detailed list of services, also adding a sector outside of Annex II to its list. The Commission uses the expression "consistency gaps," which could be misleading, as a "gap" constitutes a break in continuity[3], which is not evident in this case as it is up to each Member State to determine continuity. "Consistency differences" is a more precise term. ▸

On average, Member States have identified 35 services per country, and the number of identified services ranges from 12 to 87[4].

[2, 3, 4] *See Reference List at the end of the report*

CONSISTENCY DIFFERENCES IN THE DEFINITION OF OES

ESTONIA (LEAST GRANULAR APPROACH)	PORTUGAL	DENMARK	BULGARIA (MOST GRANULAR APPROACH)
Electricity supply	Distribution system operators	Electricity distribution	Distribution of electricity
*	*	*	Ensuring the functioning and maintenance of a distribution system for electrical energy
*	Transmission system operator	Electricity transmission	Transmission of electricity
*	*	*	Operation, maintenance, and development of an electricity transmission system
*	*	Electricity production	Electricity production
*	*	*	Electricity market

*Consistency differences.

Furthermore, the thresholds for identifying OES also vary greatly between Member States, both qualitatively and quantitatively. Thresholds are applied differently across the Union, and can be based on a single quantitative factor, e.g. the number of systems supporting the service; a larger set of quantitative factors, e.g. the number of systems plus the market share; or a combination of quantitative and qualitative factors. The figure below describes some of these differences.

THRESHOLD DIFFERENCES

Country	Internet Exchange Points (IXP)	DNS Providers	Top-level-domain registries
SECTOR-SPECIFIC THRESHOLDS			
Austria	Connected autonomous systems > 100	DNS resolvers: 88 000 users; Author. DNS: 50 000 domains	50 000 domains
Malta	25% of market share	DNS resolvers: 78 000 requests/day; Author. DNS: 7 800 domains	750 000 requests/day
United Kingdom	Market share > 50%, or interconnectivity to global internet routes ≥ 50%	DNS resolvers: 2 000 000 clients/day; Author. DNS: 250 000 domains	TLD registries ≥ 2 billion queries/day
CROSS-SECTORAL THRESHOLDS			
Cyprus	50 000 users, or 5% of subscribers of the market	50 000 users, or 5% of subscribers of the market	50 000 users, or 5% of subscribers of the market
Lithuania	Inhabitants > 145 000	Inhabitants > 145 000	Inhabitants > 145 000

EU countries have their unique challenges and characteristics, therefore the criteria for identifying essential services should reflect country-specific factors, and the countries should consequently have the ability to apply thresholds differently.

Digital service providers

Digital service providers (DSPs) are any legal persons providing a digital service. A "legal person" is typically an entity, such as a corporation, with a set of rights and responsibilities. A "digital service" is any service normally provided for remuneration, at a distance, which uses electronic means and at the individual request of a recipient of services, which is of a type referred to in Annex III. Member States do not identify DSPs, which means that the DSP has to self-assess whether it has to comply with the Directive or not.

The country in which the DSP has its main establishment and head office is the country in which it has to comply with the national legislation. In cases where a DSP is not established in the Union but provides services there, the DSP shall designate a representative in the Member State where the services are offered. The implementation of the Directive demonstrates that all EU countries have classified DSPs as the three sectors provided in Annex III, except Finland, which is the only country categorising the three DSPs as OES.

Types of DSPs are detailed in Annex III:

- **Online marketplace**
  The final destination for the conclusion of online sales or service contracts between consumers and traders, but does not cover online services acting as an intermediary to a third-party in which the contract would be concluded, or price-comparing services.

- **Online search engine**
  A platform where the user can search all websites on the basis of a query on a subject, but does not cover the search functions limited to a specific website, or price-comparing services.

- **Cloud computing services**
  Services that "allow access to a scalable and elastic pool of shareable computing resources," including storage, applications, networks, servers, or other infrastructure and services.

SECURITY REQUIREMENTS IN THE NIS DIRECTIVE

Common security requirements for OES and DSPs are one of the central measures for responding effectively to the challenges of securing network and information systems, and the responsibilities for ensuring such security lie, to a great extent, on these actors. Both OES and DSPs shall identify and take appropriate and proportionate technical and organisational measures to manage the risks posed to the security of network and information systems.

In the process of implementing appropriate and proportionate technical and organisational measures, they shall have in regard the state of the art and ensure a level of security appropriate to the risk. Furthermore, appropriate measures shall be implemented to prevent and minimise the impact of incidents. Security measures appropriate to the risks faced involve risk management measures to identify risks; to prevent, detect and handle incidents; and to mitigate the impact of incidents. Security requirements apply to OES and DSPs regardless of whether the network and information systems are managed internally or have been outsourced.

Security requirements imposed on OES

OES are only subject to specific security requirements for the essential part of their service, which excludes the non-essential operations. The main objective of these security requirements is to ensure continuity of deliverance of essential services[5]. The national competent authority (NCA) or a qualified auditor performs an information security audit to make sure the OES complies with the Directive. The Commission encourages Member States to follow the NIS Cooperation Group Reference Document to align the national provisions to the greatest extent possible.

Some technical and organisational measures that the NIS Cooperation Group Reference Document recommends OES to implement:

- Evaluate its risks and possible threats in light of a regularly updated risk analysis, information system security policy (ISSP), information security management system (ISMS), and information security audit
- Security awareness training, and a framework for asset management
- Establish an IT security architecture
- Protective security measures
- Identity and access management
- Security incident detection system

Security requirements imposed on DSPs

Safeguarding a level of security appropriate to the risk requires the DSP to consider some important elements, and the implemented measures mitigating incidents affecting their service delivery shall ensure the continuity of those services. However, the security requirements do not apply to micro- and small enterprises[6].

Member States are strongly discouraged from imposing further requirements on DSPs, except when this is required to safeguard essential State functions. Furthermore, the DSPs should remain free to take measures they consider appropriate to manage the risks as long as those measures ensure an appropriate level of security[7]. ▸

[5, 6, 7] *See Reference List at the end of the report*

The competent authority has no general obligation to supervise DSPs, and the NCA only takes ex-post supervisory measures if DSPs show non-compliance with the Directive.

DSPs shall take into account the following elements (the Implementing Regulation clarifies NISD Art. 16):

"Security of systems and facilities" means:
- systematic management of network and information systems
- physical and environmental security
- security of supplies
- restrictive administrative security of network and information systems

"Incident handling" means:
- timely and adequate detection of anomalous events
- incident reporting
- vulnerability identification
- adequate response
- providing documentation and lessons learned

"Business continuity management" means:
- maintaining or restoring the service delivery after a disruptive incident by creating contingency plans based on business impact analysis
- regularly assessing disaster recovery capabilities

"Monitoring, auditing and testing" means:
- analysis based on a sequence of observations on whether network and information systems function as intended
- verifying that the DSP complies with a set of guidelines
- establishing processes to expose security faults

"International standards" means:
- standards that are applicable for security requirements, such as ISO27000-series

The differentiation approach

The Directive differentiates between OES and DSPs in that stricter requirements can be imposed on OES and lighter and more harmonised requirements can be imposed on DSPs. The lighter approach towards DSPs is justified by its less essential service delivery. Moreover, DSPs have more freedom to conduct business, which is a crucial factor for their success. ENISA has also concluded that the EU aims to react efficiently to cybersecurity incidents without overburdening the DSPs by having a light-touch approach.

If there is a need for DSPs to increase their security level, such as in situations where public administrations in Member States use digital services provided by DSPs, the Directive recommends stipulating these obligations in a contract.

THE IMPLEMENTATION OF SECURITY REQUIREMENTS

Member States shall encourage OES and DSPs to use relevant European or internationally accepted standards and specifications in order to promote "convergent implementation," and not impose an obligation to use a specific type of technology. We experience that most organisations implementing the internationally recognised standards, such as ISO27001, already have many of the required mechanisms and systems in place to also comply with the NIS Directive. Furthermore, ENISA provides detailed advice and guidelines related to the technical field of security requirements for OES and DSPs, and also assists the NIS Cooperation Group with implementing necessary policies to satisfy the legal requirements[8].

What is "appropriate" and "proportionate"?

Our experience supports the promotion of a culture of risk management, as the threshold for implementing appropriate and proportionate technical and organisational measures requires that an organisation identifies its risks, and applies measures that are appropriate to the risks posed.

Management of risks depends on analysing and evaluating the risks and systemising these evaluations into an information security management system (ISMS). Such a system enables the organisation to analyse, assess and handle weak points, values, effects and threats, take control over residual risk, and continually optimise the overall risk exposure[9].

The mitigation of risks should further take into account that an appropriate level of security should be maintained, as well as the state of the art, which can be described as: "subject's best performance available on the market to achieve an object. The subject is the IT security measure; the object is the statutory IT security objective[10]."

As mentioned above, the risk assessments varies between OES and DSPs, as the DSPs can self-assess their security postures while the NCA audits OES regularly. Although DSPs can choose whether to take measures or not themselves, they are legally bound to comply with the Implementing Regulation, which specifies the elements that DSPs have to consider.

Transposition in the EU and Norway

The NIS Directive is transposed in all 28 EU countries. In an internal survey conducted among our senior security risk experts, 65% thought the Directive attracted little or no attention at all, while 35% thought this level of attention was moderate. Accordingly, this supports the general opinion that the NIS Directive has generated less attention than the GDPR did in 2018.

[8.9.10] *See Reference List at the end of the report*

One of the reasons could be that the nature and purpose of the regulations varies, as the GDPR is protecting a fundamental freedom, namely protecting personal data for all EU inhabitants, while the NIS Directive is securing network and information systems in certain critical sectors.

Furthermore, the mechanisms for compliance and penalties for non-compliance also differ between the two regulations. No country has imposed a penalty for non-compliance with the NIS Directive yet, but statements from governments across the Union indicate that the penalty for non-compliance with the NIS Directive will be lower than with the GDPR. Additionally, the requirements provided by the Directive are already well known by some critical sectors, as regulations in sectors such as energy, financial, and health have had similar requirements several years before the Directive was implemented.

Understanding the requirements of the Directive requires an analysis of the national transpositions. The key finding is that the Directive has been implemented in four distinct patterns:

- Security requirements in national legislations are written in a language similar or identical to the language of the Directive (Ireland, Luxembourg, Malta, the Netherlands, Portugal, Spain, Sweden, and the UK).
- The risk management obligations and the preventative obligations are merged into one obligation (Austria).
- Security requirements are sector-specific (Denmark, Finland, Germany, and Hungary).
- Security requirements have a higher degree of detail (Estonia, Greece, France, Latvia, Poland, and Slovenia).

We observe that regulations and directives quickly become outdated, as technology improves much faster than the time needed to adopt legislations. Thus, the requirements should be dynamical and not too detailed in order to be effective. Furthermore, as some sectors are more critical than others, the flexibility for countries to impose sector-specific requirements as supplements to common requirements is pivotal for achieving a high level of security. Additionally, the degree of risk can also differ between countries, which is a reason for having the possibility of imposing stricter nationwide security requirements.

In Norway, it is somewhat relevant to compare the Directive with the Norwegian Security Act (Sikkerhetsloven), even though the natures and purposes of the regulations are different. The Justice Department has concluded that the Directive is EEA relevant, and that no existing cross-sectorial or sectorial-specific laws provide levels of security requirements equivalent to those of the NIS Directive.

Whether the NIS Directive will succeed in providing a sufficient framework for achieving a high common level of security in the EU depends on the implementation effort from the EU countries.

GUEST ARTICLE

The Missing Link in Email Security

Erlend Andreas Gjære
CEO and Co-founder,
Secure Practice

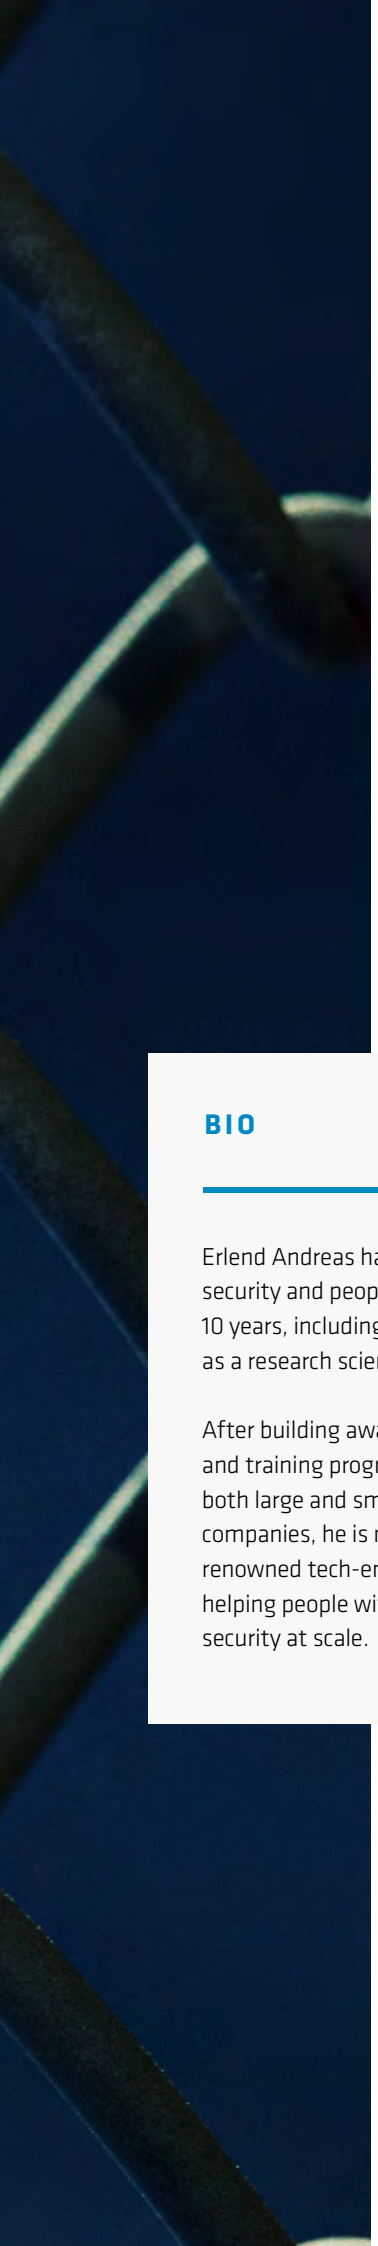

BIO

Erlend Andreas has studied security and people for nearly 10 years, including six years as a research scientist.

After building awareness and training programs at both large and small companies, he is now a renowned tech-entrepreneur helping people with digital security at scale.

Globally, more than 290 billion emails are sent every day, and the number is steadily growing[1]. A conservative estimate suggests that 1 per cent of these emails, which would amount to a daily average of 2.9 billion emails, represent a serious threat. If we expect current email security technology to block around 99 per cent, there will still be more than 10 billion dangerous emails delivered to people's inboxes on an annual basis.

It is commonly agreed that a large majority of security breaches relate to end user behaviour, and that indeed a large majority of these involve malicious emails. Leaving users on their own with suspicious emails is obviously risky, as it may only take one victim who clicks the wrong link or attachment to inflict potentially devastating consequences on an entire organisation.

People are often called out as the *weakest link* in cybersecurity, possibly in frustration over not finding reliable mitigations to human risk. Telling users to never click on untrusted links is unfortunately not very helpful, since URLs are inherently complex and difficult to parse for non-technical people. We simply cannot expect *everyone* to *always* understand and remember *everything* about internet security. While technology has also come short of eradicating the problem with malicious emails, users are frequently blamed on a somewhat unfair basis for the success of cybercriminals. Human risk is nevertheless closely related to technical risk, and we must seek to bridge the gap between these areas.

Consequently, there is a *missing link* between security and people, and solving this challenge requires a deeper understanding of the various reasons why users are prone to error.

Human errors in seven flavours

First of all, keep in mind that users are real human beings, and not only sources of failure. We all make mistakes, forget things, and deal with uncertainty in unpredictable ways. Occasionally, we even violate rules, and yet we may still create great value for our employer.

James Reason, Professor Emeritus of Psychology at the University of Manchester, has studied how situations can go wrong when reliability depends on people[2]. Although his work is primarily based on safety-critical contexts, such as process plants, aviation, and healthcare, his principles are also highly relevant for cybersecurity. In all these contexts, it is useful to distinguish between intentional and unintentional errors, and further distinguish the different errors into categories, according to studies by Reason and others. ▸

[1, 2] *See Reference List at the end of the report*

Unintentional errors can be divided into four categories, and none of the following types of errors is a result of malicious intent:

- **Slip:** A frequent action, which requires little conscious attention, goes wrong.
- **Lapse:** A particular action was omitted because it was forgotten.
- **Rule-based mistake:** A routine is followed, but an ineffective rule is applied, or a good rule is applied wrong.
- **Knowledge-based mistake:** A routine is not available, and the application of knowledge and experience is not sufficient to carry out the action safely.

Slips and *lapses* are simple actions that went wrong, and not according to plan. An example would be sending a sensitive document to the wrong recipient because of misspelling their email address, or simply due to choosing a poor autocomplete suggestion. Appropriate system boundaries and safety mechanisms are key to mitigating the risk of such events, including both failure detection as well as plain old checklists.

Mistakes are made based on a conscious but flawed plan; that is, by performing the correct actions according to the wrong plan. For example, employees unsuccessfully try to encrypt a highly sensitive email without succeeding, and the message is inadvertently sent in an insecure format. While encrypted email has yet to solve major usability challenges, human-centred design is key to make almost any task more resilient to mistakes.

Contrary to unintentional errors, intentional errors are actions characterised by non-compliance with intent, often called violations. Such errors can be further divided into three categories:

- **Routine:** A rule is so poorly implemented that its omission has become the norm.
- **Situational:** A shortcut is sometimes taken to get the job done.
- **Exceptional:** A calculated risk is taken due to special circumstances, to solve an otherwise impossible task.

Routine violations are more specifically related to how policies are defined, communicated and supported by technical systems. Perhaps people are not aware of certain rules, or at least not why a rule exists in the first place. If the rule is too vague, such as requiring users to never click on untrusted links, people could get used to breaking it. Routine violations come with an additional risk that people get comfortable with breaking other rules as well. A rule that is well justified may, on the other hand, become socially unacceptable to violate, and the norms embedded in company culture play a crucial role in securing behaviour.

Situational violations are often the result of limited time and resources, stress, or a lack of proper tools to get the job done. Take an example of someone receiving a suspicious email, and for whom policy states that phishing emails should be reported to IT. The reporting procedure requires users to forward the email as an Outlook attachment to a specific email address to avoid losing original email metadata. Very few users are able to remember and perform the necessary steps correctly, and reporting is accordingly omitted. People's perception of risk could also be inaccurate, and security training may be combined with technology support for significantly better results. The problem here may also be a cultural one: Security might be considered an issue for "security professionals" by those who are not, and accordingly receive lower priority than other goals they have during the day.

Exceptional violations occur when people take calculated risks and step outside the defined rules due to special circumstances. For example, somebody receives an email claiming the recipient has somehow failed to fulfil his or her duties, including a link or attachment with the so-called "evidence". While this scenario is quite commonly abused by cybercriminals and not exactly exceptional, it may still be subjectively perceived as unpleasant enough to warrant an exception from common advice of not clicking. The result may be users taking the risk of checking out the contents on their own, and even trying to cover up the mess when realising the embarrassing fact that they have been tricked and possibly infected. Some kind of trusted "phishing hotline" could instead provide an invaluable opportunity for IT and security to build trust with fellow employees in a positive and supportive manner. A measure like this provides people with a viable alternative to just clicking on links or attachments if they believe they need to figure out whether an email can be trusted or not. In essence, preparedness, trust and support are key to handling exceptional situations safely.

Four steps to increase security reporting

Unless any absence of reported incidents is a sign of perfect compliance, your organisation is at risk in ways that you are not aware of. Believe it or not, most of your colleagues can be valuable contributors to your company's security efforts. Although very few are security professionals, many can still spot a scam when they see one. Such colleagues are very useful resources for strengthening the company's resilience on behalf of those who are unable to spot a scam.

When people are not reporting unwanted events, critical data are lost. If you want people to report more than they currently do, you are not alone, however. To a large extent, the rate of security-related reporting naturally depends on employees ▸

being willing and able to report. Luckily, some insightful research has also been done in this area.

Sidney Dekker, Professor at Lund University in Sweden, is known for his work in human factors and safety research. Several findings from these areas may also be applied to cybersecurity, including how we can encourage employees to report incidents[3]. Based on organisations with a mature culture for reporting, we can take away four key approaches to increase reporting, which will in turn increase your company's resilience.

Mitigate negative impact

Reporting should not cause trouble for the person who reports. If there is too much work associated with reporting and following up on the report afterwards, some may rather keep quiet about incidents. Managers may receive a report and silently agree with their subordinate that it was simply a matter of "human error," and leave it at that. When this is the case, there is no opportunity for the organisation to learn from it, and people may get the idea that it is okay to cover up incidents. The situation becomes even worse if people believe that reporting involves a risk of blame, stigma, trouble, or even punishment. If an employee clicks on a malicious link, do you want to blame this person and give the person a reprimand, or do you want to fix the actual problem as soon as possible? Over time, you cannot have both.

So, are punitive responses never appropriate? Are end users never at fault for security failures? "Blame-free" does not mean that nobody can be held accountable, although there are indeed better alternatives to punishment. Dekker emphasises that accountability means getting people actively involved in creating a better system, which also requires that the organisation is open to learning. Therefore, you could begin by asking why the user was clicking on that link in the first place. Learning from this event requires listening and empathy. Maybe the actual problem was not the employee after all, but a lack of training, or lack of an appropriate support channel for determining whether the email could be trusted or not? The situation should then be improved for all employees going forwards.

Highlight positive impact

People desire a great workplace and will usually appreciate an opportunity to exercise influence in this sphere. If employees experience that reporting contributes to a safer environment, it will soon become a valuable cultural aspect of the organisation. For cybersecurity, it means that everybody knows that reporting suspicious activity will help the company protect itself against cybercriminals. If anyone flags something as suspicious or reports an incident, their efforts should always be welcomed. It also means that if something is reported, it will be taken care of and not just be put in a bin, and that should even include false positives.

Collected data should in turn be used to show employees what greater good they are contributing to. People will respond positively to seeing that their efforts yield useful results, such as actively removing or blocking detected threats. Moreover, their efforts' visibility allows the reporters to become active participants in the company's improvement process. Positive user involvement creates credible "wins" for a part of the company often associated with paranoia.

Minimise fear

Defining precisely what an incident is in advance may sometimes be difficult. This can be reason enough for people not to report what happened, because they do not want to cause any trouble or extra work. Dekker clearly states that reporting must be voluntary. If reporting is mandatory, it would mean that the company claims the right to define what is worthy of reporting. Explicit rules would however become either too specific, or too general to make them easily applicable for employees.

Although anonymity can be required for reporting certain irregularities, the opposite is often required for following up on concrete security events. To ensure that information does not get lost due to fear of repercussions, we can learn from Norwegian Air Law which states that reports cannot be used as evidence in criminal proceedings against the persons providing the evidence. If an incident is discovered in retrospect, without anyone having reported on it, the impression may be that someone wanted to cover up the issue. By contrast, being

Most people do not report incidents very often, so the user experience with doing so must be excellent.

[3] *See Reference List at the end of the report*

transparent and reporting a problem as soon as possible will effectively transfer responsibility of the situation from the employee to the organisation. This aspect is well worth emphasising to your colleagues.

Maximise accessibility

Regardless of incident type, it must be straightforward for employees to report. Whether it is an email address, a person that everybody knows, or a dedicated software tool, any channel should be readily available to everyone when it is needed. Relying on a single individual may be risky, however, insofar as this person may be busy, absent, or otherwise unable to respond to such reports. Most people do not report incidents very often, so the user experience with doing so must be excellent.

To sustain contributions over time, it is also necessary to ensure that people can get appropriate and timely feedback when they have requested help with an email or reported something as suspicious. The reporting mechanism should further make efficient use of the data we have available so that redundant reporting and communication is reduced to a minimum. This will make how the process works predictable to users, and turn reporting into a habit. Being able to repeat the process with ease will make it a natural part of people's workflow.

Resolving the missing link

An important finding related to human error research is that no major accident has ever been caused by a single error alone. James Reason's *swiss cheese* metaphor highlights this: Some holes are due to active failures, while other holes are latent conditions. Cybersecurity strategies must accordingly take all of these into account, and this is why we should apply a barrier-based approach to security that includes technology, processes, and people. Indeed, people are not only holes in the cheese!

The National Cyber Security Centre (NCSC) in the UK has published a guide for businesses for defending against email-based threats. They recommend that "(...) your people layer should put much more emphasis on reporting suspected phish as soon as possible, so your experts can investigate it[4]." We witness first-hand, every day, that people are able to detect suspicious activity when technology on its own has already failed.

Implementing an effective process for helping people handle suspicious emails in a safe way will make it clear how users could in fact be a great asset in defending your company against cyber threats. Integrating the reporting mechanism with your Security Operations Centre (SOC) to provide users with feedback in a timely manner, maybe even 24/7, will further increase the return on the investment.

By enabling people to take appropriate action and get help when needed, human suspicions can be leveraged to facilitate prevention, early detection, and effective response. Achieve this at scale, and the missing link of email security can finally be resolved.

[4] *See Reference List at the end of the report*

REFERENCE LIST

Strategic Software Security

1. https://agilemanifesto.org/
2. https://www.wired.com/2002/01/bill-gates-trustworthy-computing/
3. https://www.microsoft.com/en-us/securityengineering/sdl
4. https://www.bsimm.com/
5. https://github.com/OWASP/ASVS
6. https://infosec.mozilla.org/guidelines/risk/rapid_risk_assessment.html

Internet of Things and its Firmware: A Tale of Memory Corruption Bugs

1. https://www.idc.com/getdoc.jsp?containerId=prUS45213219
2. https://www.csoonline.com/article/3258748/the-mirai-botnet-explained-how-teen-scammers-and-cctv-cameras-almost-brought-down-the-internet.html
3. https://www.enisa.europa.eu/topics/iot-and-smart-infrastructures/iot
4. https://www.beuc.eu/general/internet-things
5. https://www.forbrukerradet.no/side/todays-legislative-tools-are-not-fit-for-connected-objects/
6. https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:32019R0881&from=EN
7. https://developer.ibm.com/articles/iot-lp201-iot-architectures/
8. https://owasp.org/www-project-internet-of-things/
9. https://en.wikipedia.org/wiki/List_of_operating_systems#Embedded
10. https://securityledger.com/2019/08/huge-survey-of-firmware-finds-no-security-gains-in-15-years/

The Value of Outsourcing Detection and Response: Making Informed Security Decisions

1. https://www.statista.com/statistics/662991/worldwide-cio-survey-outsourced-it-functions/
2. https://www2.deloitte.com/content/dam/Deloitte/us/Documents/process-and-operations/us-cons-global-outsourcing-survey.pdf
3. https://www.gsa-uk.com/uploads/attachments/cjx3a5y2l030knjh7xw7u3bcs-value-beyond-cost-noa-research.pdf
4. https://www.gsa-uk.com/uploads/attachments/cjx39ouso02xgnjh72mait6yh-outsourcing-in-2020-research-report.pdf
5. https://www.gsa-uk.com/uploads/attachments/cjx3a5m3o0308njh7q80zg1tm-the-public-perception-of-outsourcing.pdf
6. https://www.isc2.org/-/media/ISC2/Research/2019-Cybersecurity-Workforce-Study/ISC2-Cybersecurity-Workforce-Study-2019.ashx?la=en&hash=D087F6468B4991E0BEFFC017BC1ADF59CD5A2EF7X
7. https://www.csis.org/analysis/cybersecurity-workforce-gap
8. https://cybersecurity.isaca.org/state-of-cybersecurity
9. https://www.cisco.com/c/dam/m/digital/elq-cmcglobal/witb/acr2018/acr2018final.pdf

The NIS Directive: A Step in the Right Direction

1. Direktoratet for samfunnssikkerhet og beredskap (DSB). (2019). Analyser av krisescenarioer 2019.
Haber, E., & Zarsky, T. (2017). Cybersecurity for Infrastructure: A Critical Analysis. Florida State University Law Review, pp. 516-576.
2. REPORT FROM THE COMMISSION TO THE EUROPEAN PARLIAMENT AND THE COUNCIL assessing the consistency of the approaches taken by Member States in the identification of operators of essential services in accordance with Article 23(1) of NIS Directive.
3. Ibid.
4. https://www.merriam-webster.com/dictionary/gap
5. Markopoulou, D., Papakonstantinou, V., & de Hert, P. (2019). The new EU cybersecurity framework: The NIS Directive, ENISA's role and the General Data Protection Regulation. Computer Law & Security Review: The International Journal of Technology Law and Practice.
6. See definition of “micro-“ and “small enterprises” in EU COMMISSION RECOMMENDATION of 6 May 2003 concerning the definition of micro-, small and medium-sized enterprises.
7. EU COMMISSION IMPLEMENTING REGULATION (EU) 2018/151 of 30 January 2018 laying down rules for application of NIS Directive as regards further specification of the elements to be taken into account by digital service providers, Recital 1.
8. REGULATION (EU) 2019/881 OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL of 17 April 2019 on ENISA (the European Union Agency for Cybersecurity) and on information and communications technology cybersecurity certification and repealing Cybersecurity Act, Art. 5 (2).
9. IT Security Association Germany (TeleTrusT) and ENISA. (2019). Guideline "State of the Art": Technical and organisational measures. p. 63.
10. Ibid. page 11.

The Missing Link in Email Security

1. https://www.lifewire.com/how-many-emails-are-sent-every-day-1171210
2. Reason, J. (1990). Human Error. Leiden: Cambridge University Press.
3. Dekker, S. (2016). Just Culture Balancing: Safety and Accountability. Taylor & Francis Ltd.
4. https://www.ncsc.gov.uk/guidance/phishing

You can also find the references at
www.mnemonic.no/references-2020

CONTACT

CORPORATE HEADQUARTERS

mnemonic AS
Henrik Ibsens gate 100
0255 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

LONDON

mnemonic Cybersecurity
Level 39
One Canada Square,
Canary Wharf
London E14 5AB
United Kingdom

contact@mnemonic.co.uk

PALO ALTO

mnemonic
470 Ramona Street
Palo Alto, CA, 94301
USA

contact@mnemonic.no

CREDITS

Lead Editor
Alexandra Stenersen, mnemonic AS

Publication Design
Inventas AS, Oslo

Photo Credits
Charlotte Sverdrup Photography (page 60 and 65), Adobe stock 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.

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