ELSEVIER

network SECURITY

ISSN 1353-4858 July 2019

www.networksecuritynewsletter.com

Featured in this issue:

The complexity of prioritising patching

Anyone working in or around vulnerability remediation knows that the apparently 'simple' task of applying a patch is anything but. The vulnerability lifecycle is filled with pitfalls.

The time and effort needed to remediate any single vulnerability across an entire enterprise are often underestimated. This creates an obvious and urgent demand for prioritisation, which requires we understand more about the world of vulnerabilities. Michael Roytman of Kenna Security and Jay Jacobs at the Cyentia Institute explore what the open vulnerability landscape looks like and investigate multiple factors contributing to the remediation efforts.

Full story on page 6…

Why IIoT should make businesses rethink security

Companies of all sizes are revolutionising the way modern businesses operate by taking advantage of embedded sensors and cloud computing. But securing technology and protecting networks has become increasingly difficult.

Businesses urgently need holistic solutions that create transparency and traceability at a technical and organisational level. With more-complex IT infrastructures spanning thousands of endpoints, automation has become essential to streamline the detection and patching of vulnerabilities in a timely manner by making sure that every application is fully patched against that specific bug quickly and effectively, argues Sean Herbert of Baramundi.

Full story on page 9…

The impact of GDPR one year on

Data Protection Authorities (DPAs) across Europe have worked diligently to enforce compliance and ensure that the core principles at the heart of the General Data Protection Regulation (GDPR) are met.

Meanwhile, organisations have worked to ensure compliance. Paul Breitbarth of Nymity looks at what has been learned in the past year. How have businesses responded? Has the GDPR impacted other national data protection regulations? And what impact will the UK's impending exit from the European Union (EU) have on regulatory compliance and data flows?

Full story on page 11…

War breaks out between US and Iran in cyberspace

The increasingly tense relationship between Iran and the US is spreading into the cyber realm, potentially confirming the prediction from many experts that all future wars will be preceded by cyber conflict.

According to the US Department of Homeland Security (DHS), Iranian state-backed hackers are targeting US companies and government agencies with malware designed to destroy data and take down systems.

"Iranian regime actors and proxies are increasingly using destructive 'wiper' attacks, looking to do much

Continued on page 2...

Contents

Editorial Office:
Elsevier Ltd
The Boulevard, Langford Lane, Kidlington,
Oxford, OX5 1GB, United Kingdom
Tel: +44 1865 843239
Web: www.networksecuritynewsletter.com

Publishing Director: Sarah Jenkins

Editor: Steve Mansfield-Devine
E-mail: infosec@webvivant.com

Senior Editor: Sarah Gordon

Columnists: Ian Goslin, Karen Renaud, Dave Spence, Colin Tankard

International Editoral Advisory Board:
Dario Forte, Edward Amoroso, AT&T Bell Laboratories; Fred Cohen, Fred Cohen & Associates; Jon David, The Fortress; Bill Hancock, Exodus Communications; Ken Lindup, Consultant at Cylink; Dennis Longley, Queensland University of Technology; Tim Myers, Novell; Tom Mulhall; Padget Petterson, Martin Marietta; Eugene Schultz, Hightower; Eugene Spafford, Purdue University; Winn Schwartau, Inter.Pact

Production Support Manager: Lin Lucas
E-mail: l.lucas@elsevier.com

Subscription Information
An annual subscription to Network Security includes 12 issues and online access for up to 5 users. Subscriptions run for 12 months, from the date payment is received.

More information: www.elsevier.com/journals/institutional/network-security/1353-4858

12987
Digitally Produced by
Mayfield Press (Oxford) Limited

...Continued from front page

more than just steal data and money," said Christopher Krebs, director of the DHS Cyber Security and Infrastructure Security Agency (CISA), in a statement. "These efforts are often enabled through common tactics like spear-phishing, password spraying and credential stuffing. What might start as an account compromise, where you think you might just lose data, can quickly become a situation where you've lost your whole network."

The use of wiper malware has been associated with Iran as far back as 2012 when oil firm Saudi Aramco suffered the destruction of data on hard drives in 30,000 PCs.

US Cyber Command has issued an alert that Iranian-based attackers are targeting a security bypass vulnerability in Microsoft Outlook (CVE-2017-11774, detailed here: http://bit.ly/2NOQa66). An attacker already in possession of a victim's Outlook credentials (usually obtained via phishing) can use the flaw to make changes that allow the downloading and execution of malware when Outlook is opened. A patch was issued in October 2017 but many systems remain vulnerable.

A number of security companies have warned of increased attacks emanating from Iran. CrowdStrike and FireEye both reported a campaign of spear-phishing emails aimed at government officials and people in key sectors such as finance, oil and gas. The new campaign started shortly after President Donald Trump imposed new sanctions on Iran's petrochemical sector.

The Associated Press reported that it had seen phishing emails shared with it by the two security firms. "One such email that was confirmed by FireEye appeared to come from the Executive Office of the President and seemed to be trying to recruit people for an economic adviser position," it said. "Another email was more generic and appeared to include details on updating Microsoft Outlook's global address book."

The Iranian group thought to be responsible for this campaign is well known to IT security companies, which have variously dubbed it APT33, Elfin, Magnallium and Refined Kitten. It is reputed to use destructive malware and to focus on targets in the petrochemical sector. Following a March 2019 report by Symantec, which detailed the group's infrastructure and tools, researchers at Recorded Future noted significant changes in APT33's attacks, which included the adoption of previously unseen remote access trojan (RAT) malware.

Following Iran's downing of a US drone, the Associated Press cited unnamed former US intelligence officials, who claimed that the US Cyber Command launched a retaliatory cyber attack against a group thought to have ties to the Iranian Revolutionary Guard Corps – this might also be APT33. This group, it's said, supported the limpet mine attacks against ships transiting the Strait of Hormuz. The attack coincided with Trump retreating from earlier threats to mount a physical attack on Iran. The cyber attacks were allegedly mounted against systems used to control missile launchers.

Mohammad Javad Azari Jahromi, Iran's minister for information and communications technology, took to Twitter to claim that the attacks were unsuccessful.

"If the reporting is accurate, this is a great example of when and how cyber operations should be deployed in response to kinetic operations," commented Dave Weinstein, CSO at Claroty. "It is both proportionate and limiting from a collateral damage perspective. Furthermore, it has deterrence value because it demonstrates not only to Iran but to other adversarial observers that the US is both capable and willing to project cyber force in a tailored fashion. It's also noteworthy that the US reportedly targeted what can be considered a strictly military target. As international norms of cyberspace evolve, it's important to demarcate military from civilian targets, particularly as it relates to dual-use infrastructure. Finally, this operation illustrates the advantages of cyberspace as an attractive alternative military domain to sea, air, or land – especially for conducting retaliatory strikes."

Meanwhile, another attack initially attributed to Iran has turned out to be the responsibility of a European. A piece of malware dubbed Silexbot has been destroying data and rendering systems inoperable (ie, 'bricking'). It seeks out any system running a Linux, Unix or similar

Threatwatch

Modular backdoor

Kaspersky has presented details of backdoor malware, dubbed Plurox, that can be modified and enhanced through the use of plugin modules. It's capable of spreading across a network to install malware on multiple PCs. The backdoor has two interfaces with its command and control (C&C) server: one is used to install crypto-mining and the other is multi-purpose – it might also be used to install mining code or it could download other forms of malware. Kaspersky noted a UPnP plugin that it speculates could be exploited to attack a network, and an SMB one that uses the EternalBlue exploit. The main indicators of compromise (IoCs) are the C&C server addresses, some of which are IP addresses plus a couple of domain names. There are full details here: http://bit.ly/2L9MXfy.

Torrent bot targets Korean TV

Criminals have adapted an existing piece of bot code to exploit film and TV fans in South Korea. According to ESET, the GoBotKR malware is derived from GoBot2, written in Go, and is being spread via pirated copies of movies and TV shows that fans download using torrent. Along with the expected MP4 file – which is often hidden in a sub-directory – the victim receives a .lnk file (a Windows shortcut) crafted to look like the video. Clicking on this actually opens a .pma file containing the malicious executable. The malware collects information about the victim's computer and contacts a command and control server for instructions. GoBotKR is also capable of seeding torrents, enabling it to spread to other users. There's more here: http://bit.ly/2NCP1hV.

First DoH malware spotted

Netlab says it has identified what it believes to be the first malware strain to exploit the DNS over HTTPS (DoH) protocol. Dubbed Godlua, and written in the Lua scripting language, it creates a backdoor on Linux servers. Two versions have been spotted in the wild, both of which use DoH requests to obtain the TXT record of a domain where the malicious actors are running a command and control (C&C) server. By using encrypted DoH communications, the malware is better able to hide its presence on infected systems, rendering useless defences that rely on the passive monitoring of DNS requests. There's full information here: http://bit.ly/2XBB0Wz.

Astaroth fileless malware

Microsoft has released details of a strain of fileless malware known as Astaroth. It's been around since 2017, mostly used to steal information from organisations in South America and Europe via spear-phishing. But Microsoft has seen a spike in the use of malware that exploits the Windows Management Instrumentation Command-line (WMIC) tool to 'live off the land', running in memory and not touching the filesystem as a way of avoiding detection. The company is keen to emphasise, however, that such malware is not undetectable, which is why it has provided extensive details about how Astaroth functions. The details are here: http://bit.ly/2JzHwDb.

SACK Panic

Security specialists at Netflix have identified a group of vulnerabilities affecting Linux and BSD machines that can be exploited to bring down machines. The flaws are related to the Maximum Segment Size (MSS) and TCP Selective Acknowledgement (SACK) functions in the kernels, with the most serious being the so-called SACK Panic (CVE-2019-11477) via which an attacker sending a maliciously crafted TCP packet could remotely cause a kernel panic and crash the computer. The most severe effects concern older Linux kernels (older than 4.15), but all are affected to some degree, as are some versions of FreeBSD. There are details here: http://bit.ly/2Sdoxm9.

operating system – primarily Internet of Things (IoT) devices that it co-opts into a botnet. It simply tries the username and password combination root:password. Once a device is infected, it reports back to a server in Iran. However, NewSky Security researcher Ankit Anubhav found that the malware is being managed and further developed by a 14-year-old teenager based in Europe who uses the handle 'Light Leafon'. There's more information here: https://zd.net/2YFriyV.

China attacks major tech firms

For years, major technology service providers have been suffering serious network intrusions by a Chinese hacker group dubbed Cloud Hopper, according to a recent report by Reuters.

The report names Ericsson, IBM, Hewlett Packard Enterprise (HPE), Fujitsu, Tata Consultancy Services, NTT Data, Dimension Data, Computer Sciences Corporation and DXC Technology (formerly part of HPE) as the affected firms. In many cases, the companies' networks were used to mount attacks against their customers. For example, according to Reuters: "Teams of hackers connected to the Chinese Ministry of State Security had penetrated HPE's cloud computing service and used it as a launchpad to attack customers, plundering reams of corporate and government secrets for years in what US prosecutors say was an effort to boost Chinese economic interests."

The existence of the Cloud Hopper group and its activities have been known about for some time. The group was named in indictments against two Chinese nationals handed down by US authorities in December 2018; Ericsson has been monitoring attacks since 2016; and HPE first discovered malicious activity in 2012 that forensic analysis showed had been going on for at least two years. However, this is the first time that the scale of the group's activities has been revealed.

Most of the companies named in Reuters' report have declined to comment, simply said that they have defences in place or claimed that they have no evidence important data was compromised.

The Reuters report is here: https://reut.rs/30lyGzK.

Meanwhile, Israeli security firm Cybereason claims it has uncovered a China-based hacking campaign targeting Western telecoms firms dating back as far as 2012. According to the company, the activity – which it has dubbed Operation Soft Cell – attempted to compromise Active Directory installations to steal all the user credentials for an organisation, as well as "other personally identifiable information, billing data, call detail records, credentials, email servers, geolocation of users and more".

Cybereason said that the multi-wave attacks, "focused on obtaining data of specific, high-value targets and resulted in a complete takeover of the network".

Many of the attacks took place over months, with the threat actors returning with new tactics and tools whenever one form of attack failed. The attackers' tools, as well as their tactics, techniques and procedures (TTPs), are usually associated with Chinese hackers.

There's more information here: http://bit.ly/30qffpJ.

Report Analysis

Sikich: 2019 Manufacturing and Distribution Report

Cyber criminals take their opportunities where they find them. And they prefer rich targets – although the 'wealth' in this instance might be money or data. And that's why an increasing number of them are focusing on manufacturers.

According to the latest Sikich research, half of manufacturing and distribution firms have suffered at least one data breach in the past year – and of those that had been compromised, 11% described the incident as 'major'.

"Cyber criminals have moved on from focusing primarily on organisations rich in sensitive personal data, such as financial or healthcare institutions," said Brad Lutgen, the partner in charge of Sikich's cyber security practice. "Instead, they target any organisation with IT weaknesses and attempt to turn a profit through ransomware and other cyberextortion techniques. In response to this growing threat, manufacturing executives must make security a core corporate priority and push forward the implementation of preventative measures in their organisations."

Weirdly, however, manufacturing firms' confidence in their ability to withstand cyber attacks is high. More than half (54%) rate themselves 'very' or even 'extremely' confident in their ability to weather the effects of a data breach. This is in spite of the fact that many are clearly not doing enough in terms of cyber security defences – especially among firms with under $500m in revenues. Less than 40% of these smaller companies perform cyber audits (38%), penetration testing (33%), security assessments of vendors (32%) or phishing exercises on employees (31%).

In many ways, manufacturers have been surprisingly slow to adopt new information technologies. Often, their focus has been on operational technology rather than IT. That's changing rapidly now, not least with the rampant growth of the Industrial Internet of Things (IIoT). Alas, as in so many industries, the eagerness to reap the benefits of innovation is not matched by a willingness to ensure that this is done securely. Smaller manufacturers, in particular, seem to struggle to find the budget to properly protect their networks.

"Overall, our industrial sector is poorly prepared for today's cyber attacks," commented Saurabh Sharma, VP of business development at Virsec Systems. "While most competitive firms have invested in advanced manufacturing equipment, too few have extended this investment beyond basic security. Many also hang on to an outdated notion that their systems are effectively isolated or air-gapped. With today's exploding industrial IoT and connected systems, no businesses are immune from having data stolen or encrypted for ransom, or having sensitive industrial equipment disabled."

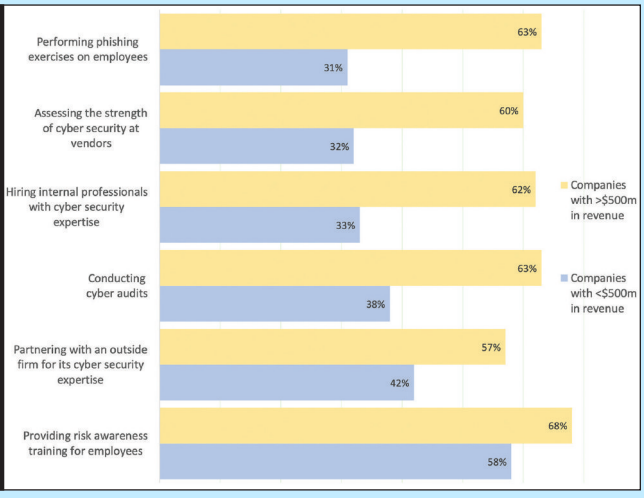

Actions being taken to enhance cyber security. Source: Sikich.

This is leaving organisations in this sector open to a wide variety of threats. For example, cyber criminals who once threw their ransomware at vulnerable individuals are now finding juicier targets, with deeper pockets, among manufacturing firms. Norsk Hydro was famously hit earlier this year and that incident has more recently been echoed with major disruption at aircraft parts manufacturer ASCO. These firms refused to pay the ransoms, but the damage – in lost productivity and remediation – can be measured in the tens of millions of dollars.

Manufacturers also represent enticing targets for those seeking to steal intellectual property. And these days, a manufacturer's supply chain is often long, complex and highly interconnected. Your own defences may be strong, but an attacker can go after one of your weaker suppliers.

Fixing the issues isn't easy. To some extent, firms can turn to external security consultants and services. But all firms need to have some security skills in-house, and this is proving challenging due to a lack of access to talent.

"The ability to attract professionals with required skills was a top issue in a variety of areas including implementing new technologies, cyber security, supply chain and fostering innovation," says Jerry Murphy, partner in charge, manufacturing and distribution services at Sikich. "Given the historically tight labour market prevailing in 2019, industrial companies will need to employ creative approaches to securing talent and make talent management central to their strategic decision-making."

Fewer than half (45%) of manufacturers have someone whose primary responsibility is managing cyber security. This drops to 15% among smaller companies.

All of this will only become more urgent as firms adopt more automation and more technology, from advanced manufacturing processes to big data. This is all built on information technology and will all be subject to potential vulnerabilities.

The report is available here: www.sikich.com/md-report/.

In brief

BA faces huge GDPR fine…

British Airways has been hit with the biggest fine ever imposed by the UK's Information Commissioner's Office. The £183m penalty is in response to an incident, first disclosed on 06 September 2018, in which the airline's website was compromised. Visitors were redirected to a fake site where, it's believed, the personal information of 500,000 people was harvested. By the time of its disclosure, the attack had probably been operating for around three months. This is the first major fine to be levied in the UK under the General Data Protection Regulation (GDPR). The amount corresponds to around 1.5% of BA's global turnover in 2017, so it is significantly less than the 4% maximum made possible under the GDPR. BA is likely to appeal the penalty.

…while King's College confesses

King's College London has admitted to a breach when it shared a list of student activists with the police and prevented the students from visiting the campus during a visit by the Queen. A letter sent by acting principal Evelyn Welch to all students and staff reported the findings of an independent enquiry, which concluded that the college had acted in a discriminatory manner and breached the GDPR. The college has sent a copy of the report to the Information Commissioner's Office. There letter is here: http://bit.ly/2XPwFy7.

Magecart blitz

More than 960 e-commerce sites have been infected with the Magecart skimming malware in just 24 hours, according to research by Sanguine Security. The firm, which specialises in malware scanning for the Magento platform, discovered what it believes to be "the largest automated campaign to date". Hackers may have exploited a SQL injection flaw in Magento that was revealed in March 2019 but which may not have been patched on many sites. The attack campaign seen by Sanguine probably used automated scripts to find vulnerable sites and inject malicious JavaScript into their pages. The malicious code has been decoded and published by Sanguine on GitHub: http://bit.ly/2XvAeWg.

Medical warnings

The US Food and Drug Administration (FDA) has issued a warning that certain insulin pumps sold by Medtronic can be hacked. The MiniMed 508 insulin pump and MiniMed Paradigm series, used by up to 4,000 people in the US, have flaws in their software that could allow an attacker on the same network to change the pump's settings, delivering too little or too much insulin. Medtronic has issued a recall for the products. There's more information here: http://bit.ly/2L7Pk2n.

Meanwhile, researchers have found vulnerabilities – one of them severe – in the Becton Dickson (BD) Alaris Gateway Workstation (AGW), used to control a variety of infusion and syringe pumps for delivering intravenous fluids and medications. Healthcare cyber security firm CyberMDX first found the problems in September 2018 and this has now led to an advisory being issued by the US Cyber security and Infrastructure Security Agency (CISA), part of the Department of Homeland Security. The most serious flaw, which could be exploited by an attacker on the same network, is in the firmware of the AGW computer that powers, monitors and controls the infusion pumps. The advisory (ICSMA-19-164-01) is here: http://bit.ly/2XysAzg.

Bangladesh bank heist

A Russian hacking group has just stolen at least $3m, and possibly more, from Bangladesh-based Dutch-Bangla Bank, according to Singaporean security firm Group-IB. The hacking group – called Silence by Group-IB and which may consist of just two people – has a track record of sophisticated attacks, particularly against banks. In this case, it seems to have gained access to Dutch-Bangla Bank's networks and installed malware on a number of PCs. This in turn gave access to the bank's card-processing system. From there, the hackers were able to set up a system that would allow ATMs to dispense cash without alerting the bank's main systems. Associates of the group – probably simple money mules – travelled to Bangladesh from Ukraine and visited the compromised ATMs to withdraw the money. It's probable the group will target other banks in a similar way. There's more information here: http://bit.ly/2S2v56E.

AMCA bankrupt

US billing company American Medical Collection Agency (AMCA) has filed for bankruptcy following the recent disclosure of major breaches. The incidents led to the leak of millions of records belonging to two of the firm's major customers, Quest Diagnostics (11.9 million records) and LabCorp (7.7 million). As a result, both firms stopped doing business with AMCA, as did two of its other chief clients – Conduent and CareCentrix. On top of the loss of business, AMCA faced an immediate bill of $3.8m purely in order to alert individuals whose information may have been compromised. The firm may also have filed for bankruptcy as a way of heading off a number of class action suits.

Xenotime targets US electric utilities

The Xenotime hacking group, which became infamous for attacks on oil and gas companies, and which nearly caused an explosion at a Saudi oil facility in 2017, has now turned its sights on US electricity firms, according to security firm Dragos. The company says the group is using similar tactics and tools, including the infamous Triton malware (also known as Trisis). Dragos has detected Xenotime probing the networks of US utilities since late 2018 and it may be preparing for a full-blown cyber attack. "Industrial control system (ICS) cyberthreats are proliferating," said Dragos. "More capable adversaries are investing heavily in the ability to disrupt critical infrastructure like oil and gas, electric power, water and more." There's more information here: http://bit.ly/32eLOs9.

New FIDO standards

The FIDO Alliance has announced two new standards and certification initiatives for identity verification and the Internet of Things (IoT). The aim is to strengthen identity verification assurance to support better account recovery and automate secure device onboarding to remove password use from IoT devices. The Alliance has formed two new working groups: the Identity Verification and Binding Working Group (IDWG) and the IoT Technical Working Group (IoT TWG) to establish guidelines and certification criteria in these areas. There's more information here: http://bit.ly/2JnJUON.

Small firms run old Windows

Two thirds of small and medium-size businesses (SMBs) are using versions of Microsoft Windows that are no longer supported or will lose support by the end of this year, according to research by Alert Logic. In fact, the majority of devices scanned during the six month-long research were running copies of Windows more than 10 years old. To make matters worse, of the many unpatched vulnerabilities found on endpoints, three-quarters of them were more than a year old, raising the likelihood that there are exploits for them in the wild. Nearly a third of SMBs rely on email servers running unsupported software – typically Exchange 2000. The figures come from analyses of more than 1.3 petabytes of data, 10.2 trillion log messages, 2.8 billion intrusion detection events and 8.2 million verified security incidents across Alert Logic's customer base. Its report is available here: http://bit.ly/2G0BVFh.

FBI says don't trust the padlock

After years of telling people to look for the padlock symbol when visiting websites, the FBI has now issued advice to ignore it. The availability of cheap, easy-to-obtain SSL certificates has meant that cyber criminals, particularly those engaged in phishing, now add credibility to their fraudulent sites through the adoption of the HTTPS protocol. And so the presence of the padlock symbol is no longer a guarantee of legitimacy, the FBI says. Its warning, and advice of staying safe, is here: www.ic3.gov/media/2019/190610.aspx.

The complexity of prioritising patching

Michael Roytman, Kenna Security and Jay Jacobs, Cyentia Institute

Jay Jacobs

Michael Roytman

As American journalist and essayist HL Mencken once wrote: "For every complex problem there is a solution that is concise, clear, simple, and wrong." Anyone working in or around vulnerability remediation knows the apparently 'simple' task of applying a patch is anything but. The vulnerability lifecycle is filled with pitfalls and deceptively complex tasks.

The time and effort it takes to remediate any single vulnerability across an entire enterprise is often underestimated, compounded by the velocity and volume of newly discovered vulnerabilities. This creates an obvious and urgent demand for prioritisation, which requires we understand more about the world of vulnerabilities. Organisations everywhere are constantly trying to optimise the efficiency of their limited resources against the need for the broadest coverage in addressing the critical vulnerabilities in their remediation efforts.

Over the past few years, the Cyentia Institute and Kenna Security have partnered to study the vulnerability lifecycle. But calling it a lifecycle incorrectly implies that vulnerabilities are born, progress through a series of sequential phases and are eventually remediated. In reality, vulnerabilities exist in a world of non-exclusive states.

A vulnerability may or may not be made public and the vendor may or may not issue or patch. Maybe it's recorded on Mitre's Common Vulnerabilities and Exposures (CVE) list, maybe an exploit has been made public, maybe signatures are generated in vulnerability scanners and/or intrusion detection/prevention solutions (or not). And these states are not linear – it is, of course, possible for a vulnerability to be discovered and exploited in the wild before it's ever public (a so-called zero-day vulnerability). Some vulnerabilities are discovered and made public after a patch was released. Some vulnerabilities may have an associated CVE but are never exploited. These non-exclusive states, and especially our knowledge of these states, have an effect on our understanding and research into vulnerabilities.

Publishing the research

We have attempted to make sense of this complex landscape and have published our research across three publications known collectively as the 'Prioritisation to Prediction' reports.[1-3] The research is both enabled and bolstered by bringing together multiple data sets on vulnerabilities and exploitation centred around the de facto standardisation of the CVE ID – an ID issued by Mitre to disseminate information about vulnerabilities.

"Only a minority of published CVEs have published exploits and even fewer are being exploited in the wild"

When a CVE is initially published it will include a brief free-text summary and URLs referencing or discussing the vulnerability. The National Vulnerability Database (NVD) will add to the data for each CVE and include information from the Common Vulnerability Scoring System (CVSS), the Common Platform Enumeration (CPE) and the Common Weakness Enumerations (CWE). For evidence of real-world exploits and activity, we scrape Exploit DB and exploit frameworks (eg, Metasploit, D2 Elliot) for evidence of published and/or weaponised exploits. We then fold multiple commercial and public data sources recording exploitation in the wild (evidence of an exploit being executed against a target) into the data.

One last data source has helped us understand the vulnerability universe and that's the output from vulnerabilities scanners being used across hundreds of enterprises and being loaded into the Kenna Security platform. Not only does this provide insight into the existence and prevalence of vulnerabilities, it also provides vision into how vulnerabilities are being prioritised within and across organisations.

Exploit relationship

Our first volume focused on a simple question: What's the relationship between a published vulnerability and an associated exploit of that vulnerability? We found that only a minority of published CVEs have published exploits and that even fewer are being exploited in the wild. Specifically, we studied approximately 100,000 CVEs and found that only 23% had a published exploit and less than 3% were exploited in the wild.

We were also able to measure the relationship between an exploit being published and it being exploited in the wild: the probability of exploitation in the wild is seven times higher when an exploit is publicly released. This implies that any CVE that is either actively being exploited in the wild or has a published exploit should be a priority for remediation. This begs further questions: first, what metrics can help us understand the performance of various prioritisation strategies? With good metrics, we should be able to differentiate between good and bad remediation strategies. Second, can we build a predictive model that can help prioritise better than any existing strategy?

In an ideal world, we could perfectly predict which vulnerabilities will be exploited and focus all remediation efforts on those vulnerabilities. How accurately a prioritisation strategy is able to identify vulnerabilities that will have exploits in the wild (ie, how close it comes to this ideal world) may seem like a good measure of a strategy's effectiveness. However, because vulnerabilities with exploits in the wild are rare (only 3%), a strategy that predicted no vulnerabilities would ever be exploited would have an accuracy of 97%. Moreover, the cost of a false negative, failing to predict exploitability and therefore not remediating a vulnerability would likely be costly. False positives, while still costly, are less likely to have dire consequences.

Balanced strategies

Our goal then is a measure for strategies that balances coverage of exploited vulnerabilities and efficiency of a low false positive rate. Thankfully, the machine learning community has already defined good measures in precision and recall.

"The cost of a false negative, failing to predict exploitability and therefore not remediating a vulnerability, would likely be costly. False positives, while still costly, are less likely to have dire consequences"

Efficiency (precision) measures the proportion of prioritised vulnerabilities that are actively being exploited or have been weaponised (ie: an exploit is publicly released). It's calculated by dividing the true positives (vulnerabilities we correctly prioritise) by the sum of the true positives and the false positives (vulnerabilities we incorrectly prioritise).

Coverage (recall) measures the proportion of vulnerabilities being actively exploited/weaponised that we've prioritised. It is calculated by dividing the true positives by the sum of the true positives and false negatives (vulnerabilities that were not prioritised yet should be). In theory, we'd strive for 100% efficiency and 100% coverage, which requires us to find a strategy that only prioritises the vulnerabilities that are actively being exploited or are weaponised. In reality a trade-off will exist between the two, we will be able to raise the efficiency of our actions at the expense of a lower coverage, or we could get better coverage at the expense of lower efficiency. Risk-seeking or resource-constrained organisations may opt for the former, while risk-averse organisations may opt for the latter.

Remediate above CVSS Base Score	Remediated correctly (True Pos.)	Delayed incorrectly (False Neg.)	Remediated too soon (False Pos.)	Delayed correctly (True Neg.)	Efficiency (Precision)	Coverage (Recall)	Efficiency by Chance	Coverage by Chance
10	1,510	20,207	5,025	67,855	**23.1%**	**7%**	23%	7.1%
9	3,148	18,569	10,405	62,475	**23.2%**	**14.5%**	23%	14.7%
8	3,228	18,489	10,736	62,144	**23.1%**	**14.9%**	23%	15.1%
7	11,562	10,155	25,180	47,700	**31.5%**	**53.2%**	23%	39.8%
6	14,320	7,397	34,715	38,165	**29.2%**	**65.9%**	23%	53.2%
5	17,547	4,170	49,753	23,127	**26.1%**	**80.8%**	23%	73%

Table 1: Results for prioritisation strategies based on CVSS Base Scores. Source: Kenna/Cyentia.

"When accounting for the volume of each vulnerability, we found that just three vendors – Oracle, Microsoft and Adobe – accounted for almost seven out of every 10 of the open vulnerabilities"

How do existing strategies perform?

With two measures of prioritisation strategy established, we can begin to evaluate strategies and compare them. An obvious first strategy is the CVSS base score, which asks six multiple choice questions and assigns a score between 0 and 10, with 10 representing the most severe vulnerability. Many organisations use this as a starting point to prioritise vulnerabilities. For example, when a vulnerability is assigned a CVSS base score of 10 (as 6,535 were in our first study), it naturally sparks a sense of urgency and intuitively should be prioritised over other vulnerabilities.

Let's measure the performance of that strategy. If we prioritise the 6,535 CVEs with a CVSS base score of 10, we get 1,510 true positives and 5,025 false positives, while missing 20,207 false negatives. This calculates to an efficiency of 23.1% and coverage of 7%. But is this good? Let's compare that to a nonsensical strategy of a completely random approach. If we prioritise 6,535 CVEs at random, we would achieve on average, 23% efficiency and 7.1% coverage. In other words, prioritising CVSS 10 CVEs is no better than random chance. In fact, prioritising CVSS 9 and above and CVSS 8 and above are also no better than ran-

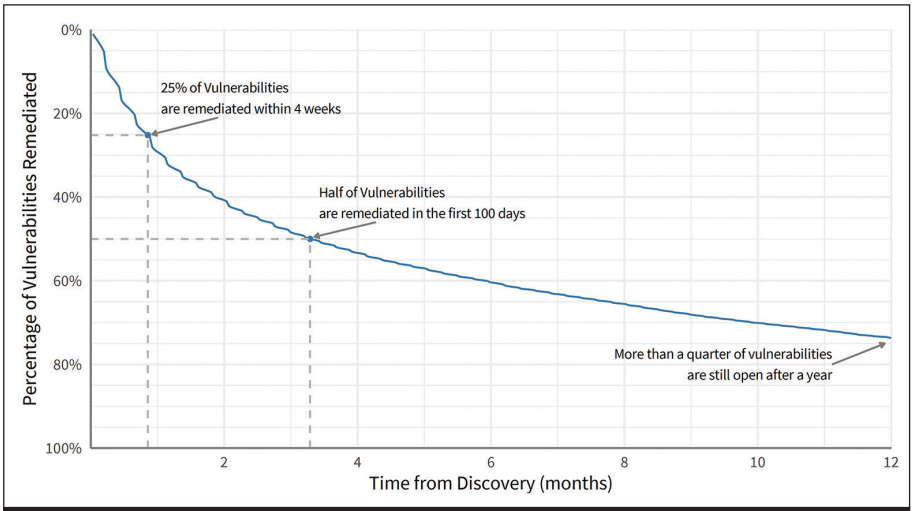

Figure 1: Overall vulnerability survival analysis across firms. Source: Kenna/Cyentia.

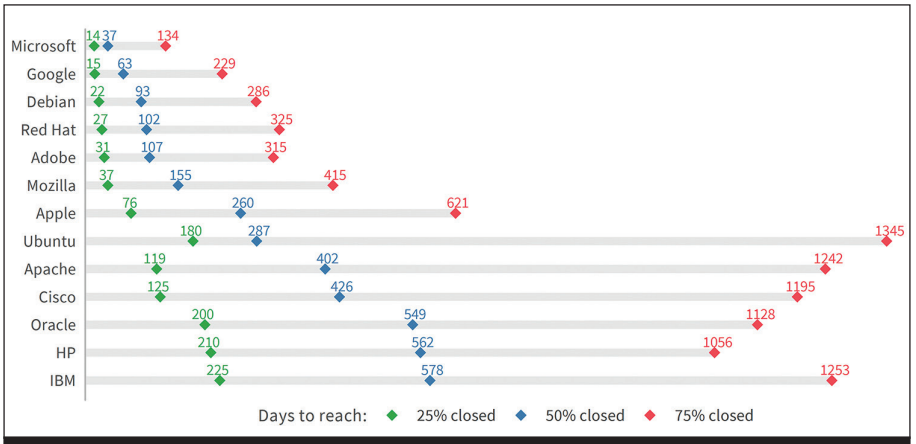

Figure 2: Remediation velocity for major product vendors. Source: Kenna/Cyentia.

dom chance. We do see an improvement over random chance at CVSS 7 and above, though, as shown in Table 1.

Volume 2 of our research looked into the vulnerabilities observed in the environments of hundreds of organisations that shared their vulnerability scanner data with Kenna Security. We found something that extended our previous findings; out of the 108,000 published CVEs at that time, we found that approximately 34% of the published CVEs were observed to be open in one or more organisations, and only about 5% of the CVEs were observed *and* known to be exploited in the wild.

We next asked how these vulnerabilities were distributed across software vendors. When accounting for the volume of each vulnerability, we found that just three vendors – Oracle, Microsoft and Adobe – accounted for almost seven out of every 10 of the open vulnerabilities. This isn't necessarily surprising given the market penetration of their products. What is interesting is the overall closure rate from those vendors. Only about 30% of open vulnerabilities in Oracle products have been closed, while 64% of Adobe products and an impressive 77.6% of Microsoft vulnerabilities have been closed. This makes more sense given the majority of Oracle vulnerabilities are on the Java platform. Ask anyone who has tried to update or patch a JVM and they'll be quick to explain why Oracle patches largely go unapplied.

We wrap up our research in Volume 2 by looking at remediation times (using survival analysis) for a small sample of 12 organisations. Volume 3 expands that sample to almost 300 organisations and creates an informative chart showing the overall remediation times for vulnerabilities:

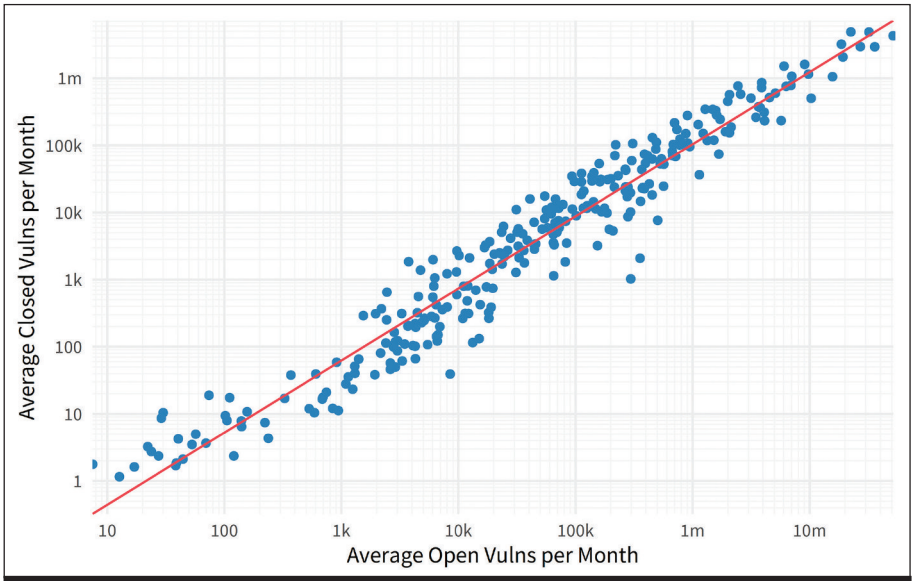

Figure 3: Ratio of open to closed vulnerabilities per month. Source: Kenna/Cyentia.

Notice the annotations in the chart identify how long it takes to close 25%, 50% and 75% of the open vulnerabilities? In the full research report, we visually compare those points across different variables in the vulnerabilities. For example, we break out remediation times by vendors shown in Figure 2.

Figure 2 could serve as an ambassador for the auto-update movement. The vendors at the top of the chart enforce or support auto-updates and the vendors towards the bottom support a more manual approach. Microsoft at the top of the chart takes two weeks (14 days) to close 25% of its open vulnerabilities, 50% is done within 37 days while 75% of its open vulnerabilities are closed within 134 days (about three and a half months). Compare that with Oracle products that take about five months to close 25% of the open vulnerabilities and over three years to reach a 75% closure rate. Clearly there is a huge difference between Microsoft's patching practices and Oracle's.

Struggling with patching

There is one last point we researched in Volume 3, looking at what we termed 'remediation capacity'. We looked at the average open vulnerability count per month for each organisation and compared that to the average number of closed vulnerabilities each month. The relationship between these two variables was truly surprising and we show this in Figure 3.

Notice that both of the scales are logarithmic, with each measure likely increasing in proportion with the exponential growth we see in other aspects of organisations (counts of employees, assets, breached records, etc). If we follow the regression line from the bottom left to upper right, we see how every tenfold increase in open vulnerabilities is met with a roughly tenfold increase in closed vulnerabilities.

That, in a nutshell, is why it feels like vulnerability management programs can never pull ahead in the race of remediation. A typical organisation, regardless of asset complexity, will have the capacity to remediate about one out of every 10

vulnerabilities in its environment within a given month. That seems to hold true for firms large, small and anywhere in between.

Where are we heading?

Vulnerability management is deceptively complex. The more we research the topic, the more questions we end up discovering. This article serves as a high-level summary of what we've uncovered so far. Overall, we have established that many of the common rule-based remediation strategies are no better than random chance.

"A typical organisation, regardless of asset complexity, will have the capacity to remediate about one out of every 10 vulnerabilities in its environment within a given month. That seems to hold true for firms large, small and anywhere in between"

We've begun to explore what the open vulnerability landscape looks like and investigated multiple factors contributing to the remediation efforts across hundreds of organisations. Our next report asks the questions that will reveal the causal factors behind our research to date. Our main observation is that, as always, if you ask the right questions, the data always reveal an interesting and complex story.

About the authors

Michael Roytman is chief data scientist at Kenna Security and a recognised expert in cyber security data science. At Kenna, he is responsible for building the company's core analytics functionality focusing on security metrics, risk measurement and vulnerability measurement. Named one of Forbes' 30 Under 30, Roytman's entrepreneurship skills include founding organisations such as Dharma, a cloud-based data management platform, and TruckSpotting, a mobile app for tracking food trucks. He also serves as an advisor to CryptoMove, as well as humanitarian organisations including Doctors without Borders, The World Health Organisation and the UN. In addition, he chairs the board of Dharma and is also a board member and programme director at the Society of Information Risk Analysts (SIRA). Roytman is a frequent speaker at security industry events, including BSides, Metricon, RSA, SIRACon, SOURCE and more. He holds a Master of Science in Operations Research degree from Georgia Institute of Technology.

Jay Jacobs is partner and co-founder at Cyentia Institute. He is a security data scientist with a deep-seated passion for using data to improve cyber security decisions, practice and products. Jacobs enjoys digging into data to find the insight and knowledge to tackle hard problems. Taking on numerous projects throughout his career, Jacobs is best known for his contributions to Verizon's annual 'Data Breach Investigations Report' series and his book Data-Driven Security: Analysis, Visualization and Dashboards. *He is a founding member of the Society of Information Risk Analysts and remains an active proponent of improving how we measure and manage risk.*

References

1. Prioritisation to Prediction, Volume 1: Analysing Vulnerability Remediation Strategies'. Kenna Security/Cyentia Institute. Accessed Jun 2019. www.kennasecurity.com/prioritization-to-prediction-report/images/Prioritization_to_Prediction.pdf.
2. 'Prioritisation to Prediction, Volume 2: Getting Real About Remediation'. Kenna Security/Cyentia Institute. Accessed Jun 2019. https://www.kennasecurity.com/prioritization-to-prediction-report-volume-two/images/Getting_Real_About_Remediation.pdf.
3. 'Prioritisation to Prediction, Volume 3: Winning the Remediation Race'. Kenna Security/Cyentia Institute. Accessed Jun 2019. https://www.kennasecurity.com/prioritization-to-prediction-report-volume-three/images/Prioritization_To_Prediction_Winning_the_Remediation_Race_digital.pdf.

Why IIoT should make businesses rethink security

Sean Herbert, Baramundi

Sean Herbert

Companies of all sizes are revolutionising the way modern businesses operate by taking advantage of embedded sensors and cloud computing, making it possible for machines, industrial plants, factories and even construction sites to be connected as part of Industry 4.0. But the practice of securing technology and protecting networks from cybercrime has become increasingly difficult, as demonstrated by the catastrophic threats posed by cyber attacks on nuclear power plants across the US and power grids in the UK energy sector in 2017.

The high level of connectivity between industrial components – referred to as Industrial Internet of Things (IIoT) – allows for employees at any level of the organisation to access and share vast amounts of data through operating sensors and applications in real time. Breaking open data silos is enabling businesses to connect seemingly isolated departments throughout the organisation and develop better products faster, as the production can quickly be adjusted to any change in customer demands or if equipment needs servicing or maintenance.

When the infrastructure of a production facility is Internet-connected, subcontractors and product developers as well as distributors and logistic bodies can take part in the manufacturing process and, thereby, optimise the entire supply chain. However, while IIoT enables businesses to minimise human errors and production costs, it also exposes them to greater security risks.

"Nine out of 10 organisations in the operational technology (OT) sector, including critical national infrastructure providers, experienced at least one damaging cyber attack over the past two years, half of which resulted in downtime to plants or operational equipment"

More connected devices result in increasingly complex IT infrastructures with multiple endpoints for cyber criminals to access. According to the Ponemon Institute, nine out of 10 organisations in the operational technology (OT) sector, including critical national infrastructure providers, experienced at least one damaging cyber attack over the past two years, half of which resulted in downtime to plants or operational equipment. Consequently, IT administrators now have to manage and secure employees' desktops, mobile devices and data stored in the cloud as well as all the connected devices and automated systems that are present in the manufacturing process.

Jeopardising production

In addition to the networking infrastructure of a production facility, other areas such as building technology – where intelligent smoke detectors, temperature sensors or building controls are present – can also pose a threat. Often, one contaminated device is enough to allow criminals to infiltrate the entire network and endanger the production process. This can also indirectly lead to partners working with the manufacturing company, such as suppliers or customers, being affected by the consequences of the cyber attack. The sheer number of devices affected by an attack could increase the potential damage enormously.

A recent example was when one of the world's largest producers of aluminium, Norwegian company Norsk Hydro, was hit by a ransomware attack. The LockerGoga virus contaminated Norsk Hydro's computers and spread throughout IT systems across most business areas. After encrypting company files, the ransomware demanded payment in crypto-currency in exchange for the decryption key. Norsk Hydro has yet to estimate the exact operational and financial impact of the attack, but as it caused production lines to shut down and forced the company to switch to manual operations, it is safe to say that the damage suffered was severe.

In addition to shutting down production lines, cyber criminals can also exploit the connected devices in industrial networks to carry out attacks on critical national infrastructure, such as the attacks on Ukraine's power grid. In 2015, hackers caused blackouts across the country by manually switching off power to electricity substations, whereas the blackout attack hitting Ukrainian capital Kiev in 2016 was fully automated. Ultimately, the advancement of this form of attack enables hackers to cause blackouts far more widespread in the future, which could have catastrophic consequences.

Separating IT and OT

The constant stream of cyber attacks and the increasing number of Internet-connected endpoints shows that IIoT is calling for a new approach in network security. Although many companies are still failing to determine whether IT teams or OT teams should be held accountable for securing the production network, they should start assessing what the most appropriate network architecture looks like.

IT managers should ideally design their network architecture in such a way that not all connected devices are accessible in the event of a break-in. In principle, IT should be kept separate from OT. The production facilities are combined in the latter. In such a concept, the processes can be controlled by servers located in the Demilitarized Zone (DMZ) devices that have access to both the IT and OT networks.

There are four layers in the OT network that map the production environment. The lower areas, from Level 0 to Level 2, refer to the operative production area and contain, for example, the production cells or plants there. Sensors and actuators are located at the lowest level, the so-called field level. Level 1 is the control level. Level 2 addresses process control. Production flow and comprehensive monitoring take place at Level 3. Above the DMZ is the IT network for the conventional office functions as well as the enterprise and management level

Even though IT and OT should be kept separated in the network architecture, it is necessary to bring together the knowledge from IT security and the specific requirements from OT (such as safety) to prevent security threats. Furthermore, it is important to implement IT security standards for all organisational levels as well as sharpen awareness of security among production employees.

Automation is a must

In the past, cyber criminals usually targeted the firewall as a first point of entry to the network. As defences at the network's outer borders have become more efficient, attackers are now choosing more subtle methods. For example, employees are lured to prepared pages that distribute malicious code. Here, manipulated files such as DOCs or PDFs are used. These exploit vulnerabilities in the programs used to edit or display them and implant malicious code that can also affect the production area. As a rule, the firewall can no longer detect this process because the connection is established from inside the corporate network. Therefore, vulnerability management should not only focus on the operating system, but also on the applications.

Due to the heterogeneous and complex structure in a modern production environment, it is difficult to identify risks and harmful vulnerabilities across thousands of devices in real time. Consequently, protection measures in Industry 4.0 should include an automated approach to vulnerability manage-

ment for intelligent production plants with different operating systems and different interfaces. A centralised approach that highlights vulnerabilities in the current state of software on devices or the network would enable administrators to identify and respond promptly to risks posed by outdated software versions, thus preventing potential damage.

In this context, a traditional approach to vulnerability management is no longer practicable, especially with the need to address new regulation such as the General Data Protection Regulation (GDPR). Last year, more than 16,000 new vulnerabilities were registered in the National Vulnerability Database, which amounts to an astonishing 300 a week. The administrator would have to manually search and evaluate databases and blogs for relevant information about vulnerabilities.

Mind the gap

The sheer number of connected devices and vulnerabilities suggests that it is just not possible for IT administrators to maintain a consistent and reliable overview of security operations in a production facility without an automation strategy. If administrators monitor endpoints manually, human error may occur unnoticed, exposing the network to incumbent threats. This is especially true within industrial environments, where OT managers are more reluctant to use software patches as changes to machinery can potentially involve re-certification and production delays.

Accordingly, both IT and OT managers should rely on a vulnerability management solution allowing them to map out the network setup and structure by carrying out a complete inventory of all network devices, configurations, installed software and the drivers for endpoint subsystems in a few seconds. This is the first step towards increasing security, because it is only through knowledge of a company's resources and how they interact that IT managers will be able to protect them effectively.

Once a certain vulnerability has been detected by the manufacturer, it normally does not take long for a patch to be released. However, the security threat will persist until that patch has been installed by the end user. The speedy installation of the patch will lessen the chances for intruders to penetrate the infrastructure. The longer the gap is left uncovered, the more vulnerable the IT environment is, which is the reason why an automated approach is crucial to ensure the highest security standards.

To minimise the risks posed by cyber attacks in Industry 4.0, businesses urgently need holistic solutions that create transparency and traceability on a technical and organisational level, which is necessary to ensure efficient corporate IT. With more complex IT infrastructures spanning thousands of endpoints, automation has become essential to streamline detection and patching of vulnerabilities in a timely manner by making sure that every application is fully patched against that specific bug quickly and effectively.

About the author

Sean Herbert is UK country manager at Baramundi and is experienced in developing strong relationships with clients – from small businesses to global enterprises – and helping them to address the challenges that IT departments face today in keeping their infrastructure up to date, safe and under control. He specialises in efficient methods for endpoint management and security and has an in-depth knowledge of endpoint management.

The impact of GDPR one year on

Paul Breitbarth

Paul Breitbarth, Nymity

The one-year anniversary since the European Union's General Data Protection Regulation (GDPR) came into effect has recently passed (25 May 2019). During the past year or so, Data Protection Authorities (DPAs) across different countries have worked diligently to enforce compliance and ensure that the core principles at the heart of the GDPR are met – namely responsible and transparent handling and protection of individuals' personal data.

But what has been learned in the past year? How have businesses responded? Has the GDPR impacted other national data protection regulations? And what impact will the UK's impending exit from the European Union (EU) have on regulatory compliance and data flows?

What have we learned?

There has been a clear shift in mindset from the DPAs as time has progressed. In the initial months after the GDPR came into operation, most DPAs began exploratory investigations. This mainly saw them offering guidance and advice to companies in breach of the regulations and the wider business community. This approach saw DPAs allow some leeway and, crucially, the chance for organisations to quickly address gaps in their policies.

However, this phase is now over. DPAs are ramping up enforcement and we have seen numerous examples of contraventions being sanctioned in the past year, ranging from high-profile fines levied against Internet giants to other examples involving smaller, less well-known organisations.[1]

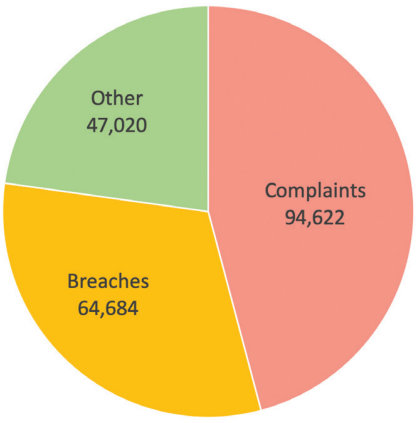

The nature of cases reported to national Data Protection Authorities (DPAs) in the first nine months of the General Data Protection Regulation (GDPR). Source: European Data Protection Board.

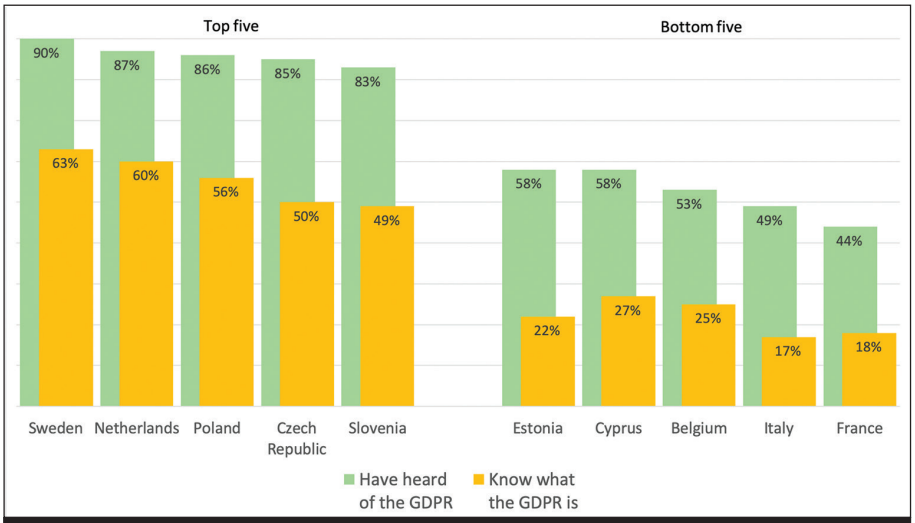

The top five and bottom five EU countries in terms of awareness of the GDPR. Source: European Commission.

The European Data Protection Board, which is made up of the European Data Protection Supervisor and representatives from national DPAs, released its first overview of the implementation and enforcement of the GDPR in February 2019.[2] The findings revealed that there were 206,326 cases reported from the DPAs in the 31 countries in the European Economic Area during the first nine months since the GDPR came into effect. Close to half of these cases (96,622) were related to complaints, while over a quarter (64,684) were related to specific data breaches.

The European Commission (EC) has also recently released the results of its Eurobarometer on data protection, which includes the views of over 27,000 people across the EU ahead of the one-year anniversary of the GDPR. The report explores awareness, compliance and enforcement of the rules, and reveals that over two thirds (67%) of Europeans have heard of the GDPR.[3] And by June, this figure had increased to 73%.[4] It also shows that more than half (57%) of Europeans know there is a dedicated public authority in their country that is responsible for protecting their data and personal rights surrounding it. However, only one in five people know which public authority is specifically responsible.

According to the EC, the most frequent type of complaints reported in the past year have been in response to telemarketing, promotional emails and video surveillance or CCTV. The Commission detailed the full results of its Eurobarometer at a special anniversary event on 13 June.

How have businesses responded?

Compliance with the GDPR demands ongoing attention. This in itself continues to bring myriad challenges that are of critical importance, including the issue of resourcing, both financial and staffing. The demand for seasoned privacy professionals continues to increase and identifying these individuals to deliver the work required is difficult.

Related to this, securing buy-in from the board for ongoing compliance has been another major challenge for companies. When conversations around the impact of GDPR began, the threat of significant fines and sanctions was commonplace and understandably captured the attention of boards. However, as the weeks and months have passed, keeping data protection front of mind continues to be an ongoing challenge for businesses.

Nevertheless, there has been progress, according to the latest IAPP-EY Annual

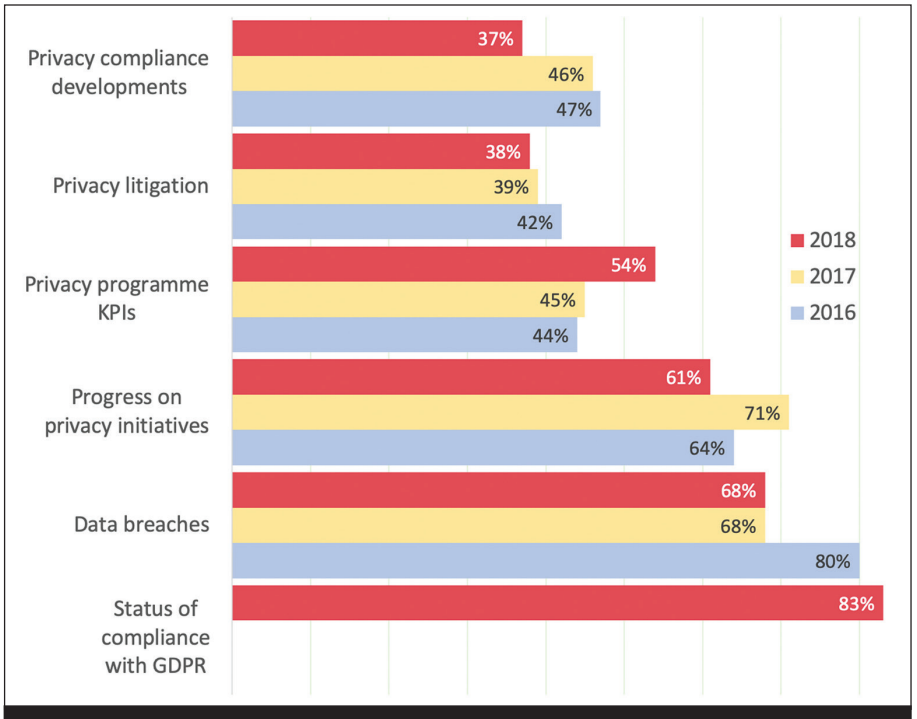

Privacy-related issues reported to the board. Source: IAPP.

Privacy Governance Report.[5] Some 89% of EU respondents to the survey stated that they have appointed a data protection officer in response to the GDPR, while awareness on the issues of data protection has risen. Progress on compliance (83%), data breaches (68%) and on privacy initiatives (61%) feature highest on the agenda among boardrooms, while investment in training is on the rise. Nearly eight in 10 respondents noted training investments as their top GDPR compliance priority for the coming year.

Following the GDPR's lead

The GDPR has created a surge in privacy regulations. The most common aspect of the legislation being replicated globally is the guidance around data subject rights, accountability requirements and data breaches, which have all generated widespread public interest and awareness of how personal data is handled by organisations.

While it is true that not every law across other global legislations is fully comparable to the GDPR, the majority do all share the same goal – enabling individuals to have more control and ownership of their personal data. For example, new legislation coming into effect in both South and North America next year has been at least in part influenced by the GDPR.

In Brazil, the country's first General Data Protection Law, the LGPD, will enter into force on 15 August 2020. Just as with the GDPR, the LGPD is an omnibus law covering numerous principles of data protection, including data transfer, data breaches and data security. It currently has around 133 amendments in process.

In California, the California Consumer Privacy Act (CCPA) will enter into application on 1 January 2020.[6] This legislation is not identical to the GDPR but has been inspired by it, particularly around data subject rights. However, it applies only in the State of California and is still subject to clarification from lawmakers on various elements. Nevertheless, it has set the agenda for a period of significant change to the privacy compliance landscape in the US, with 17 other states also proposing similar bills to the CCPA.

In Europe, countries not in the EU, including Norway, Switzerland, Liechtenstein and Iceland, have all aligned their respective regulations almost identically with the GDPR to facilitate access to the internal market. Meanwhile, several countries in Africa and South East Asia have similarly strengthened data protection legislation in order to continue doing business with Europe.

South Korea is updating its regulations with the goal of achieving adequacy in the coming year, with many of its data privacy laws potentially being combined into one omnibus law almost identical to the GDPR. The Indian Parliament is also currently debating data protection legislation reflecting multiple aspects of the GDPR.

The impact of Brexit

The uncertainty that continues to rumble on around the UK's exit from the European Union is a serious concern in terms of the implications for the GDPR. If the Government negotiates a deal, then the free flow of data from which the UK has benefited as a member of the EU under the GDPR will continue during the transitionary period before a final adequacy agreement is agreed. Data laws would likely also remain much as they are currently.

However, a 'no deal' scenario would undoubtedly have more severe and far-reaching implications. While, on the basis of the Information Commissioner's guidance, data transfers between the UK and EU would be unaffected for the time being, there would be an urgency to implement contractual clauses to legitimise transfers from the EU to the UK. This is because the UK would in effect be viewed as a 'third country'.

Regardless of how things develop between now and 31 October 2019, organisations and businesses alike must ensure they keep up to date with guidance offered and enforcement decisions made by their country's DPA. The reputational risks and financial penalties that come with non-compliance with the GDPR simply cannot be ignored.

About the author

Paul Breitbarth is director of strategic research and regulator outreach at Nymity (www.nymity.com). Based in The Hague, Netherlands, he leads relations with regulators and key customers across Europe. He is also senior visiting fellow at Maastricht University's European Centre for Privacy and Cybersecurity and, before joining Nymity, served as senior international officer at the Dutch Data Protection Authority. While there, Breitbarth was an active member of various Article 29 Working Party subgroups, co-authoring a large number of opinions, including on the data protection reform, surveillance and the Privacy Shield.

References

1. 'The CNIL's restricted committee imposes a financial penalty of 50 Million euros against GOOGLE LLC'. CNIL, 21 Jan 2019. Accessed Jun 2019. www.cnil.fr/en/cnils-restricted-committee-imposes-financial-penalty-50-million-euros-against-google-llc.
2. 'First overview on the implementation of the GDPR and the roles and means of the national supervisory authorities'. European Data Protection Board, Feb 2019. Accessed Jun 2019. https://edpb.europa.eu/sites/edpb/files/files/file1/19_2019_edpb_written_report_to_lib.
3. 'Data Protection Regulation: one year on'. European Commission, 22 May 2019. Accessed Jun 2019. http://europa.eu/rapid/press-release_IP-19-2610_en.htm.
4. 'Data Protection Regulation one year on: 73% of Europeans have heard of at least one of their rights'. European Commission, 13 Jun 2019. Accessed Jun 2019. http://europa.eu/rapid/press-release_IP-19-2956_en.htm.
5. 'IAPP-EY Annual Governance Report 2018'. International Association of Privacy Professionals. Accessed Jun 2019. https://iapp.org/resources/article/iapp-ey-annual-governance-report-2018/.
6. Californians for Consumer Privacy, home page. Accessed Jun 2019. www.caprivacy.org/.

Visual hacking – why it matters and how to prevent it

Peter Barker

Peter Barker, 3M

To be truly comprehensive, security management strategies need to include prevention of visual security breaches – the ability to physically view and even photograph sensitive or confidential information. Visual privacy is either specified or implicit within a variety of regulations and industry-specific guidelines, and one area that is gaining considerable attention is preventing unauthorised viewing of sensitive or confidential information on digital screens.

Abuse of that information to perpetuate a fraud, carry out corporate espionage or create a data breach is often referred to as 'visual hacking' or 'shoulder-surfing'. Compared to sophisticated cyber attacks, these breaches of security do not require specialist computer skills – anyone can carry out a visual hack and likewise, anyone with an unprotected digital screen is a potential victim.

Visual hacking experiment

The exact scale of the risk is hard to estimate, but several studies and anecdotal evidence demonstrate just how easy – and fast – it is to achieve a visual hack. For instance, back in 2016, the Ponemon Institute carried out the Global Visual Hacking Experiment, involving businesses in eight countries: China, France, Germany, Japan, India, South Korea, the UK and the US.[1] The study involved 157 trials in offices with between 25 and 100 employees and in all cases, the participating companies were given two days' notice that the trials were to take place. On the day itself, a white hat hacker posed as a temporary office worker, complete with a valid and visible security badge. The trials took place in full view of other workers.

The results were a stark illustration of just how rapid and easy visual hacking can be, with attempts successful at an average of 91% worldwide, 49% taking less than 15 minutes and 66 less than 30 minutes. An average of 3.9 pieces of sensitive data were obtained on each occasion and the white-hat hacker was only confronted an average of 32% during the trials.

"In France, visual hackers acquired an average of 5.3 pieces of sensitive data per experiment and employees only questioned or reported the hacker on 20% of occasions"

Information obtained was varied and included personal identification information, customer and employee details, general business correspondence, access and log-in credentials, confidential or classified documents, attorney-client privileged documents, plus financial, accounting and budgeting information. While content was obtained in several ways – including viewing paper documents and even putting them in briefcases – 52% of sensitive data was obtained by viewing people's screens.

The results in the UK and mainland Europe were on a par with global averages (and in some cases, marginally better), but the data suggests there is room for improvement. For instance, in France, visual hackers acquired an average of 5.3 pieces of sensitive data per experiment and employees only questioned or reported the hacker on 20% of occasions. Germany's scorecard was better than most countries, with only two pieces of data obtained per experiment, with the visual hacker confronted in 59% of the experiments. However, overall, Germany's visual hacker still achieved a success rate of more than 88%. In the UK, 44% of sensitive data came from viewing people's screens,

Visual hacking, or 'shoulder surfing', requires no specialist skills.

with an average of 3.1 pieces of sensitive data obtained per attempt and the visual hacker only confronted in 39% of the experiments. The overall success rate was high, at 87%.

Interestingly, while all the main functions within an organisation were affected, customer service and sales management were the two most successfully 'hacked'. This might indicate that these are departments less focused on the privacy of the information they are handling, compared to, say, the legal or finance teams.

Taking the risk outside

The Ponemon Global Hacking Experiment took place within office walls, but of course, these days workers are increasingly mobile. Estimates vary according to different sources, but the respected industry analyst firm Strategy Analytics predicts that the global mobile workforce will be 1.87 billion in 2022, accounting for 42.5% of the global workforce.[2]

"While we may already think to protect these machines with security software, biometric access and other techniques, do we also think about visibility of content when the screen is 'live'?"

Another Ponemon Institute study called Open Spaces focused on visual hacking risks when employees worked in public spaces.[3] Some 87% of people questioned had caught someone looking at data on their laptops in public, and 76% of them admitted to inadvertently seeing something important on someone else's screen. Only half of them said they had taken any steps to protect on-screen information.

Many of us can probably identify with these findings: after all, most of us will have caught a glimpse of someone's screen while seated behind them at a conference, or beside them on a train, or even walking past the back of their desk. Plus, we often have multiple digital devices at our disposal, each of which presents yet another extension to the security threat landscape, contributing to a growing volume of endpoints in a company's network, whether they are owned by the organisation or 'bring your own device' (BYOD). While we may already think to protect these machines with security software, biometric access and other techniques, do we also think about visibility of content when the screen is 'live'?

Furthermore, this information is easy to visually record, thanks to the sophisticated, high-quality camera apps embed-

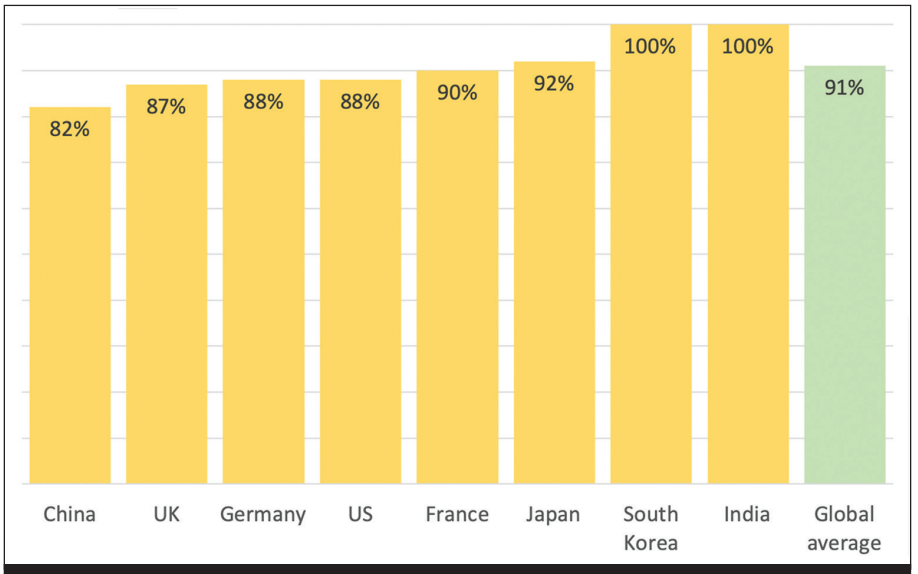

Visual hacking success rate by country. Source: Ponemon Institute/3M.

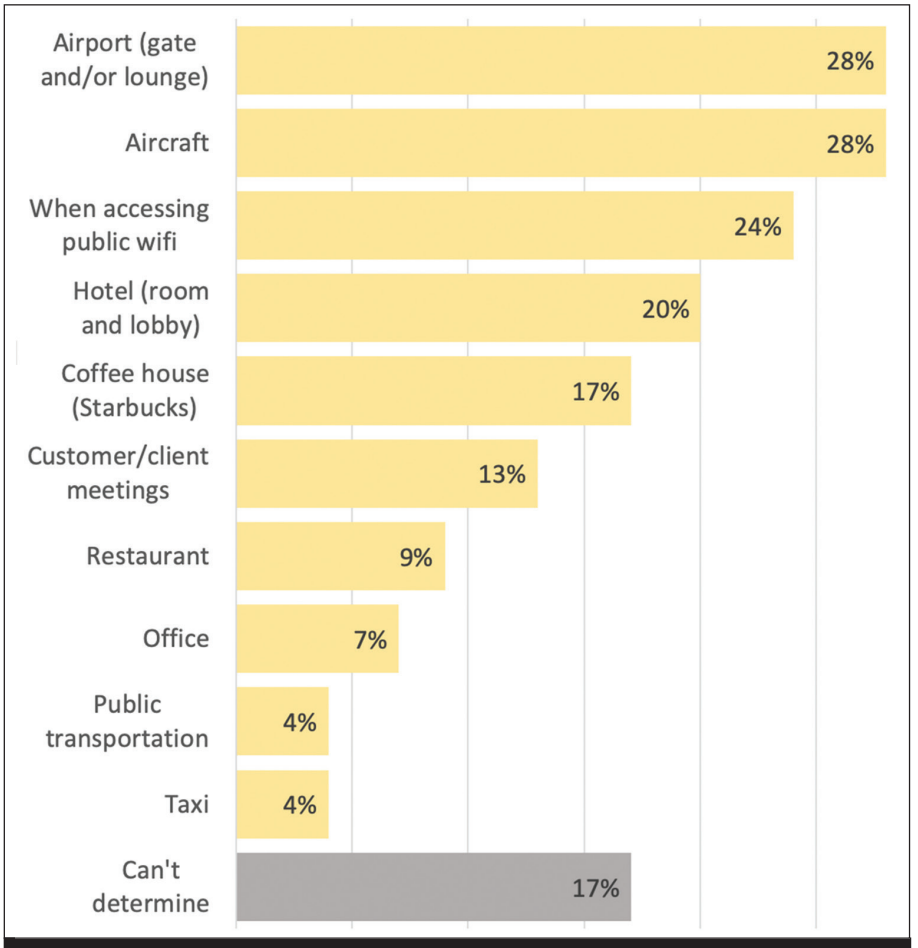

Places where organisations believe their employees are most vulnerable to visual hacking. Source: Ponemon Institute/3M.

ded within most modern smartphones. It takes just seconds to snap someone else's screen and often without them even being aware. Those images can then be instantly forwarded and shared around the world.

Mandatory privacy

Visual hacking is a very real risk and while there are many organisations or individuals yet to take preventative action, many already have taken steps to improve visual privacy. This might be as part of ISO27001 processes, but visual privacy has also got the attention of a variety of official bodies and government departments.

For instance, the UK Government's Security Policy Framework says that government departments and agencies must adopt 'clear computer screen' policies in areas where sensitive assets are handled.[4] The Department of Work and Pensions and the Foreign and Commonwealth Office have both specifically referenced the need to protect screens and they mention privacy filters in this context. Within the education sector, the Joint Council for Qualifications (JCQ) regulations state that each workstation within examination conditions must be isolated by a minimum space of 1.25 metres, unless monitors are positioned back-to-back, separated by dividers or protected by privacy filters.

"It is important to remember that the GDPR is a principle-based regulation. This means regulators don't provide organisations with a set of definitive actions to follow. Instead, organisations should think about GDPR requirements as a sort of 'desired state'"

Financial services are one commercial market sector leading the charge, which is understandable given the sensitive information they handle, plus the ability of the Information Commissioner's Office (ICO) to levy heavy fines for data breaches. Visual privacy is an implicit part of the Financial Conduct Authority's industry guidelines. Within the legal sector, the Bar Council and Law Society have both recognised the growing visual privacy issue, with the Bar Council issuing best practice guidelines including: "Where possible, computers should not be placed so that their screens can be overlooked, especially in public places" and "You should use appropriate security technologies suitable for the particular device or application".

Then, of course, there is the General Data Protection Regulation (GDPR) and this has almost certainly been one of the main catalysts for the increasing focus on visual privacy over the past 18 months. While so much of the focus on GDPR is around the content held on systems, networks and digital devices, it is important to remember that the GDPR is a principle-based regulation. This means regulators don't provide organisations with a set of definitive actions to follow. Instead, organisations should think about GDPR requirements as a sort of 'desired state' for their data-handling practices. In practice, this means it does not matter whether an unauthorised data disclosure happens because a hacker launches a sophisticated cyber attack on a company's website, or because a stranger takes a picture of highly sensitive data displayed on an employee's laptop screen.

Stopping visual hacks

All these activities, together with an overall awareness of how illegally obtained information can have catastrophic consequences, has put visual privacy much higher on the security agenda. The good news is that compared to other aspects of security management, it is relatively simple, fast and cost-effective to reduce visual hacking risks. Here are some of the actions that we have seen organisations around the world adopt.

The first step is to make sure that employees and anyone else responsible for handling valuable data – for instance, contractors or suppliers – are aware of not just the risks, but also their roles and responsibilities around visual privacy. Lobby senior-management support for any initiatives and consider appointing individuals within the organisation as champions who encourage and train teams around visual privacy. For example, they can help make clear that it is perfectly acceptable to politely question anyone in the building not displaying security clearance or who is unaccompanied.

Reducing the risk of visual hacking does not have to be very high-tech. Encourage staff to clear their desks at the end of the day and lock away any documents deemed sensitive or confidential. Make sure that the mailroom, photocopier and printer trays are checked to ensure that important documents are not left in full view. Multi-functional printers with a 'pull printing' feature mean that a paper document is only released into the hands of an authorised person at the point of collection.

Routine shredding of documents and avoidance of unnecessary printing or copying should already be standard office procedures. When working away from the office, avoid carrying printed documents unless strictly necessary. Ensure that briefcases or luggage can be securely fastened or even locked, though of course, it is important that they can be accessible on-demand to airport security officials.

"For instance, when working in public places such as cafes, airport or hotel lounges, staff displaying confidential information on their screens should always try to sit with their backs to a wall or similar barrier"

They may be 'old-school', but automatic screensavers or re-login requirements that are activated after a couple of minutes' activity are effective ways to reduce the amount of time a screen might be exposed to prying eyes. Another, very simple step is to make sure that a screen is angled so that it cannot be easily seen. For instance, when working in public places such as cafes, airport or hotel lounges, staff displaying

confidential information on their screens should always try to sit with their backs to a wall or similar barrier.

Also consider applying film-based privacy filters, which prevent on-screen data from being viewable except straight-on or at very close range. Someone taking a sideways glance or standing several feet behind will merely see a blank image. These filters also help to prevent scuffs and other damage to displays and compared to other security products, are comparatively easy, rapid and cost-effective to implement. Recent developments in film technology mean that there does not have to be any compromise to visual clarity – quite the opposite: the latest generation of privacy filters also reduce unwanted screen glare. They can also be easily slipped on and off in a matter of moments.

These filters are already in widespread use within financial institutions, legal firms, government departments and other parts of the public sector, particularly the health service, education and police. They are also increasingly being adopted by a variety of organisations – large and small – in other industries and while this is often to meet standards or compliance requirements, equally overall awareness of the need for better visual privacy is a driving force.

Of course, mitigating the risks of visual hacks is just one of many different elements to consider as part of security and risk management. On its own, implementing better visual privacy is not going to crack cybercrime. However, it addresses one of the most potentially vulnerable areas of information security, while also being one of the easiest to improve. Ensuring that visual privacy is built into security and privacy policies is a smart decision that any organisation should take to protect itself, its employees, its partners and its customers.

About the author

Peter Barker is EMEA market development manager, display materials and systems division, at 3M, the science-based technology company, and a provider of privacy filters.

References

1. 'New study exposes visual hacking is a global problem'. 3M/Ponemon Institute. Accessed Jun 2019. https://multimedia.3m.com/mws/media/1254330O/global-visual-hacking-experiment-whitepaper.pdf.
2. Luk, Gina. 'Global Mobile Workforce Forecast Update 2016-2022'. Strategy Analytics, 28 Oct 2016. Accessed Jun 2019. www.strategyanalytics.com/access-services/enterprise/mobile-workforce/market-data/report-detail/global-mobile-workforce-forecast-update-2016-2022.

Using artificial intelligence in the fight against spam


Jan Oetjen

Jan Oetjen, GMX

More than 40 years ago, on 3 May 1978, a computer vendor in the US sent the first spam email in history. It was sent by marketing manager Gary Thuerk to a list of 320 people who were active at that time on Arpanet, a predecessor of today's Internet, and was to invite them to the launch of a new computer in Los Angeles and San Mateo.

Even then, many recipients weren't exactly happy about the unwanted advertising, and one even said that it had crashed his computer.[1] Nevertheless, the email campaign was a complete success. The newly launched computer became a bestseller.

"The vast majority of spam emails have far less chance of making it into an email user's inbox because spam filters are constantly evolving"

While the first 'spammer' earned about $14m, spammers today can only dream of achieving such a high return for a single spam mail. They receive only one response for every 12.5 million emails that they send, but can still earn around $3.5m over the course of a year and, during that same period, businesses will suffer a $20.5 billion loss in productivity as a result.[2,3] Email providers fight this unwanted flood of messages with highly specialised staff and the most up-to-date spam filter systems available. And now there is a new buzz word: artificial intelligence. AI may currently be touted as the greatest thing since sliced bread – and it is certainly a major player in the fight against spam – but is it all it's cracked up to be? And can it ever take over from humans in combatting spam completely?

Professional spam attacks

Since its first occurrence, spam has changed a lot. For decades spam could be easily recognised by its poor design, clumsy sales pitch and numerous spelling mistakes. But today, spam mails are professionally designed and cover a wide range of topics.

Spam senders are increasingly picking up on trends such as the emergence of crypto currencies and messages that are intended to intimidate, frighten or appeal to the recipient's greed, desperation or just curiosity. These include

```
YEH@LLL-COMP
YONKE@USC-ISIB
YOUNGBERG@SRI-KA
ZEGERS@SRI-KL
ZOLOTOW@SRI-KL
ZOSEL@LLL-COMP
DIGITAL WILL BE GIVING A PRODUCT PRESENTATION OF THE NEWEST MEMBERS OF THE
DECSYSTEM-20 FAMILY; THE DECSYSTEM-2020, 2020T, 2060, AND 2060T.  THE
DECSYSTEM-20 FAMILY OF COMPUTERS HAS EVOLVED FROM THE TENEX OPERATING SYSTEM
AND THE DECSYSTEM-10 <PDP-10> COMPUTER ARCHITECTURE.  BOTH THE DECSYSTEM-2060T
AND 2020T OFFER FULL ARPANET SUPPORT UNDER THE TOPS-20 OPERATING SYSTEM.
THE DECSYSTEM-2060 IS AN UPWARD EXTENSION OF THE CURRENT DECSYSTEM 2040
AND 2050 FAMILY. THE DECSYSTEM-2020 IS A NEW LOW END MEMBER OF THE
DECSYSTEM-20 FAMILY AND FULLY SOFTWARE COMPATIBLE WITH ALL OF THE OTHER
DECSYSTEM-20 MODELS.

WE INVITE YOU TO COME SEE THE 2020 AND HEAR ABOUT THE DECSYSTEM-20 FAMILY
AT THE TWO PRODUCT PRESENTATIONS WE WILL BE GIVING IN CALIFORNIA THIS
MONTH.  THE LOCATIONS WILL BE:

               TUESDAY, MAY 9, 1978 - 2 PM
                    HYATT HOUSE (NEAR THE L.A. AIRPORT)
                    LOS ANGELES, CA

               THURSDAY, MAY 11, 1978 - 2 PM
                    DUNFEY'S ROYAL COACH
                    SAN MATEO, CA
                    (4 MILES SOUTH OF S.F. AIRPORT AT BAYSHORE, RT 101 AND RT 92)

A 2020 WILL BE THERE FOR YOU TO VIEW. ALSO TERMINALS ON-LINE TO OTHER
DECSYSTEM-20 SYSTEMS THROUGH THE ARPANET. IF YOU ARE UNABLE TO ATTEND,
PLEASE FEEL FREE TO CONTACT THE NEAREST DEC OFFICE
FOR MORE INFORMATION ABOUT THE EXCITING DECSYSTEM-20 FAMILY.
```

The message, sent in 1978, that is regarded as the first 'spam' email.

threatening emails from senders pretending to be lawyers or debt collection agencies, fake order confirmations from online shops or notifications from social networks that a message has been received.

The evolution of spam filters

Today, however, the vast majority of spam emails have far less chance of making it into an email user's inbox because spam filters are constantly evolving. In their simplest form, they work as follows – simple rules filter out messages with suspect words such as 'online pharmacy', 'Viagra' or 'Lottery Win' that come from unknown or blacklisted IP addresses.

But spammers can quickly update their messages to work around these barriers. By just adjusting the spelling of a word, they can outwit these simple filter rules. Depending on the font used, the difference between a lowercase 'L', an uppercase 'I' and the numeral '1' can hardly be recognised. From the word 'Viagra' you only have to make 'V1agra' and the word is no longer recognised by the algorithm. To make the spam filters recognise this unwanted message correctly, a new rule must be added to the filter system – and this has to be done for each new filter evasion that the spammer comes up with. This is complex and complicated. Nowadays, the analysis of individual words alone is no longer sufficient for reliable spam detection.

"The technology is very advanced and improving all the time, but the best possible spam filter at the moment relies on human beings and machines working together, not in isolation"

And this is where machine learning (ML), a branch of AI, comes into play – it allows computers to process data and learn for themselves without being manually programmed. An ML-based spam filter can learn in several ways, but you have to train it. This can be done, for example, by using a large amount of data from already recognised spam mails. These are examined by ML for patterns that occur repeatedly and are highly likely to be an indicator of spam. The ML algorithm then automatically creates a new rule for the spam filter.

A second way to train spam filters with the help of ML is user feedback. If many users mark emails containing the word 'V1agra' as unwanted, the filter learns that the changed spelling is a new criterion for spam and automatically creates a new rule for it.

So, can artificial intelligence replace humans when it comes to fighting spam? The technology is very advanced and improving all the time, but the best possible spam filter at the moment relies on human beings and machines working together, not in isolation. Why is that?

Human intelligence

An experienced email security expert or 'spam cop' can assess the individual potential of spam emails much more comprehensively than a machine to determine whether there is a genuine danger by identifying the possible 'value chain' – that is, how spam ultimately gets converted into cash. The spammer has one aim and one aim only – to get paid. The spam cop is able to ask: 'What happens if a link in a phishing email is clicked?', 'How will the online fraudster get his money in real life?', 'What banking method will they use?'. There is a lot of experience and some very specific expertise involved in thinking this through and these qualities are currently possessed only by humans.

The experts programme the algorithms that automatically analyse questions such as 'Where/who does the email come from?', 'Has the recipient bought from this retailer before?', 'Does the recipient usually receive a lot of mails from this IP address?', 'Have there been any malicious communications from this domain reported by other users before?' and so on.

Of course humans also have the ability to dig a bit deeper and try to identify patterns or see if they can cross-reference with other activities. So, for example, 'Was there a major data breach recently where private data could have been hacked – maybe from a well-known company with millions of subscribers?'.

Another example is phishing trends. Trends may come and go and it's up to humans, not machines, to identify them. Last year, for example saw, among other scams: malicious emails sent from compromised Mailchimp accounts; phish-

ing scams related to the EU's General Data Protection Regulation (GDPR), where criminals sent emails to presumed Airbnb hosts and told them that they could not accept any more guests or send messages until they had accepted/clicked on the new EU privacy policy; and a rash of sextortion emails where recipients were threatened with their contacts being sent compromising videos unless a ransom was paid. All this information is put together by developers to create algorithms that are combined with the strengths of ML to build a highly effective protection against spam.

Challenging 'graymail'

In addition to the anti-spam specialists, there is a second human factor in the evaluation of spam – the user. From the user point of view, spam can be classified into three categories. First there is black spam. This is spam which is either not accepted by the provider's email servers (because it is delivered by servers on blacklists) or can be detected as unwanted spam by spam filters – eg, illegal advertising. Second, there is red spam, which contains malicious links (eg, phishing) or even malware. For both categories, the recognition rate is very good across all major email providers, so that users hardly ever see these emails.

Then there is a third category – 'greymail'. Users currently have an edge over machines when assessing this third category. Called 'grey' because it is neither on the blacklist of blocked senders or on the user's whitelist of approved senders, this is email that your spam filter isn't quite sure what to do with until it has learned a bit more about it, because some users mark it as spam and others don't.

A good example is emails from retailers. The recipient technically opted in to receive those emails by 'engaging' with the company when making a purchase, but after that isn't really interested in the subsequent marketing emails. These emails are always moved to the 'Junk' folder and maybe the recipient also selects the 'block the sender' option. Over time, the spam filter will learn what the recipient considers to be greymail based on these actions as well as by the actions of all other recipients of emails sent from that particular domain name. AI may in the future be able to adjust and improve its reaction to this sort of spam proactively, based on such continuous feedback.

Man and machine

AI accelerates spam detection and at the same time increases the hit rate because it evaluates huge amounts of data almost in real time. As mentioned before, it is based on machine learning that relies on algorithms to learn from experience.

"Deep learning, a sub-discipline of machine learning, uses artificial neural networks built like the human brain. They can be trained in such a way that they independently recognise patterns in the input data and learn from mistakes"

There is further potential on offer from deep learning, a sub-discipline of machine learning that uses artificial neural networks built like the human brain. They can be trained in such a way that they independently recognise patterns in the input data and learn from mistakes. However, there are limits and these are where human expertise and strategic and creative thinking are indispensable. In addition, hackers are becoming increasingly sophisticated in overcoming defence systems, which makes it more difficult to defend against attacks, especially since the attackers also use AI.

Although AI is sometimes seen as a threat to human autonomy, humans and machines should be viewed in the context of enhancing each other's strengths: 'humans *plus* machines', not 'humans *versus* machines'. This hybrid intelligence based on human values is the best way to increase AI adoption and to boost productivity.

About the author

Jan Oetjen, CEO of GMX (www.gmx.com), joined the United Internet Group in October 2008 and is responsible for the Mail and Portal business of the company. Before that he was managing director at the online travel agency Travelocity Group, heading its operations in Germany and France. He has a diploma in business administration and economics. Since 2014 he has been a board member at United Internet and is responsible for consumer applications.

References

1. 'Reaction to the DEC Spam of 1978'. Brad Templeton, blog. Accessed Jun 2019. www.templetons.com/brad/spamreact.html#reaction.
2. Hartley, Adam. 'Spam gets 1 response per 12,500,000 emails'. TechRadar, 10 Nov 2008. Accessed Jun 2019. www.techradar.com/news/Internet/computing/spam-gets-1-response-per-12-500-000-emails-483381.
3. Bauer, Emily. '15 outrageous email spam statistics that still ring true in 2018'. Propeller, 1 Feb 2018. Accessed Jun 2019. www.propeller-crm.com/blog/email-spam-statistics.

The Firewall

Why you need Cyber Essentials

Colin Tankard, Digital Pathways

The UK Government's Cyber Essentials programme was developed in collaboration with industry and is intended to help businesses mitigate common online threats. Operated by the National Cyber Security Centre (NCSC), it was launched in 2014 and has become a key element of excellence for cyber security.

Applicable to all sizes of organisations, it offers help to those seeking to implement a robust data security strategy to protect both themselves and their clients and it does this by encouraging organisations to adopt good practice in information security: it includes a simple set of security controls, protecting information from external and internal threats.

The controls suggested by the scheme are designed to prevent basic cyber attacks and come in two formats: Cyber Essentials – a self-assessment application that addresses basic threats and helps to prevent the most common attacks; and Cyber Essentials Plus (CE+) – the same as for Cyber Essentials but rather than being self-assessed, it requires verification of cyber security carried out independently by a certification body.

Cyber Essentials offers a sound foundation of basic hygiene elements that all types of businesses can implement and potentially build upon. The Government believes that implementing these measures can significantly reduce vulnerability. However, it isn't a silver bullet to remove all cyber security risk; for example, it is not designed to address more advanced, targeted attacks and hence organisations will need to implement additional measures as part of their security strategy.

The Assurance Framework, leading to the awarding of Cyber Essentials Plus certification, has been designed to be light of touch and achievable at low cost. It is important to recognise that certification only provides a snapshot of cyber security practices at the time of assessment.

It is always advisable to have an internal and external network scan before a certification test is booked, as the scan will highlight any areas of weakness, giving time to fix issues, avoiding having a failure on certification day, or few 'last minute' fixes while the assessor is on site!

The CE+ process falls into two sections, external and internal. For external systems, the assessor carries out the following: a review of customer questionnaire information on ports; full-service scan, plus TCP and UDP service scans; an external vulnerability scan; and web application testing for common, known vulnerabilities, if in scope.

Internal tests cover greater ground, including: an internal vulnerability scan; a facility walkthrough; manual system checks covering unnecessary user accounts, weak passwords, user access control (privileges check), unnecessary software, an auto run feature check, security firewall and malware protection checks, and a review of password, Internet security, starter and leaver policies. Internal checks also include email system checks to test possible weaknesses and mobile device checks.

During the test, evidence is required such as audit logs from firewalls and servers.

For businesses that are willing to adopt these measures, the benefits can be substantial, including the ability to tender for contracts that require Cyber Essentials Certified supplier status.

Becoming accredited also helps to meet the needs of the General Data Protection Regulation (GDPR) as it covers the requirement to understand where personally identifiable information (PII) is held and can therefore provide evidence for GDPR statements/policies, showing that as an organisation you have considered such issues and had controls verified by an independent assessor.

EVENTS CALENDAR

3–8 August 2019
Black Hat USA
Las Vegas, US
www.blackhat.com

6–7 August 2019
Cyber Security in Government
Canberra, Australia
http://bit.ly/2WmOmRs

8–11 August 2019
Def Con
Las Vegas, US
www.defcon.org

14–16 August 2019
USENIX Security Symposium
Santa Clara, CA, US
www.usenix.org/conference/usenixsecurity19

19–22 August 2019
International Workshop on Securing IoT Networks (SITN)
London, UK
http://collaboratecom.org/workshop-sitn/

26–29 August 2019
International Conference on Availability, Reliability and Security (ARES)
Canterbury, UK
www.ares-conference.eu

26–30 August 2019
HITB Security Conference
Singapore
https://conference.hitb.org

4–8 September 2019
DerbyCon
Louisville, KY, US
www.derbycon.com

4–7 September 2019
r2con
Barcelona, Spain
www.radare.org/con/2019