Organisation for Economic Co-operation and Development

DSTI/CDEP/SDE(2020)3/FINAL

Unclassified **English - Or. English**

3 February 2021

DIRECTORATE FOR SCIENCE, TECHNOLOGY AND INNOVATION
COMMITTEE ON DIGITAL ECONOMY POLICY

Working Party on Security in the Digital Economy

ENCOURAGING VULNERABILITY TREATMENT

Responsible management, handling and disclosure of vulnerabilities

JT03470859

This document, as well as any data and map included herein, are without prejudice to the status of or sovereignty over any territory, to the delimitation of international frontiers and boundaries and to the name of any territory, city or area.

Foreword

This report was prepared by the OECD Working Party on Security in the Digital Economy (SDE) following discussions held at the inaugural event of the OECD Global Forum on Digital Security for Prosperity (GFDSP) in 2018 (OECD, 2019[1]). It provides an in-depth discussion of vulnerability management, handling and disclosure. It served as a basis to develop a separate shorter "Overview for policy makers" on the same issue (OECD, 2021[2]).

This work was developed in parallel and should be read in conjunction with the OECD reports on "Understanding the digital security of products: an in-depth analysis" and "Enhancing the digital security of products: a policy discussion" (OECD, 2021[3]; OECD, 2021[4]). Both work streams on security of products and vulnerability treatment were meant to inform the review of the OECD *Recommendation on Digital Security Risk Management for Economic and Social Prosperity* (OECD, 2015[5]).

This report was approved and declassified by the OECD Committee on Digital Economy Policy on 30 November 2020. It was drafted by Laurent Bernat, with support from Ghislain de Salins, Matthew Nuding, and Marion Barberis of the OECD Secretariat. Delegates to the OECD SDE also provided valuable feedback and inputs on earlier drafts.

The Secretariat was supported by an international and informal advisory group comprising 94 experts from government, business, the technical community and civil society who sent written input, and met face-to-face in February and virtually in July 2020 under the auspices of the OECD GFDSP. The Secretariat wishes to thank all these experts for their valuable feedback on earlier drafts, and in particular: Alexander Botting, Christopher Boyer, Kwadwo Burgee, Kaja Ciglic, Amit Elazari, Nicolas Eyraud, Stefan Frei, Chris Gibson, Anastasiya Kazakova, Amélie Koran, Ariel Levite, Art Manion, Axel Petri, Lorenzo Pupillo, Nicolas Reichert, Sebastien Rummelhardt, Christine Runnegar, Fred Schneider, Rayna Stamboliyska, Jeroen Van der Ham, and Tarah Wheeler.

This document, as well as any data and any map included herein, are without prejudice to the status of or sovereignty over any territory, to the delimitation of international frontiers and boundaries and to the name of any territory, city or area.

Table of contents

Figures

Boxes

Executive Summary

Addressing vulnerabilities more effectively is key to a successful digital transformation

Digital security risk undermines trust in digital transformation and generates tremendous economic and social costs. Digital security risk is estimated to a yearly global cost ranging between USD 100 billion and 6 000 billion, and is increasingly threatening individuals' safety through vulnerable Internet of Things (IoT) devices.

Vulnerabilities are a major source of digital security risk. Vulnerabilities are weaknesses in products' code and information systems that can be exploited to damage economic and social activities, and harm individuals. Malicious actors exploit such vulnerabilities to steal money, personal data as well as trade and State secrets, disrupt business operations, and hold ransom firms, cities, and hospitals.

Code almost always contains vulnerabilities. It would be unrealistic to attempt to "free" all code of any vulnerability.

However, it is possible to treat vulnerabilities more effectively. Getting better at treating vulnerabilities is a major opportunity to reduce digital security risk and increase trust in the digital transformation era.

Vulnerability treatment deserves more policy attention

So far, vulnerabilities have not received enough policy attention. The acceleration of digital transformation, while bringing tremendous benefits, also relies dangerously on billions of potentially vulnerable IoT devices, and complex information systems cumulatively running hundreds of billions of lines of code. Although criminals and other attackers seize every opportunity to create harm, as they showed during the COVID-19 pandemic, there has been limited policy efforts to encourage stakeholders to treat vulnerabilities more effectively.

Vulnerability treatment includes discovery, handling, management and public disclosure. Vulnerabilities are first identified (discovery). Vulnerability owners then need to fix them by developing and distributing a patch or another mitigation (handling). System owners have to apply patches (management). Lastly, vulnerabilities often need to be disclosed publicly to enhance security knowledge and facilitate protection.

Treating vulnerabilities is a shared responsibility amongst vulnerability owners. In the era of digital transformation, it is grossly irresponsible to develop code and maintain systems while ignoring the consequences of the vulnerabilities that may emerge over time. Producers and system owners need to establish processes to treat vulnerabilities systematically and proactively in order to decrease risk for themselves and others.

Significant economic and social challenges prevent stakeholders from treating vulnerabilities effectively

Treating vulnerabilities is an economic as much as a technical issue. Many challenges to effective vulnerability treatment are economic in nature. They include a lack of co-operation amongst stakeholders, limited market incentives, legal barriers, and lack of resources and skills. This combination can be overwhelming for SMEs, public sector bodies, and organisations with low digital maturity, such as traditional manufacturers entering IoT markets.

Stakeholders often do not trust governments because in some cases law enforcement, intelligence and national security agencies look for vulnerabilities to exploit for their own purposes. Policies often allow them to discover vulnerabilities without reporting them to vulnerability owners, and to stockpile, weaponise and exploit them against public or private targets. These agencies can also buy vulnerabilities to carry out "offensive operations". In some cases, policies may permit governments to require developers to insert "backdoors" in their products, which are equivalent to intentional vulnerabilities, a practice unanimously condemned by other stakeholders, and by some governments. A government's ambiguity with respect to vulnerability exploitation can diminish the effectiveness of policies to promote vulnerability treatment by undermining other stakeholders' trust in government efforts to reduce risk.

A collective effort is needed to make vulnerability treatment more effective

Security researchers are a significant but underappreciated resource to help vulnerability owners assume their responsibility to find and disclose vulnerabilities before malicious actors. However, many vulnerability owners do not welcome vulnerability reports from security researchers. Vulnerability owners are not sufficiently aware of good practice to encourage security researchers to find vulnerabilities in their code or systems, such as vulnerability disclosure policies and bug bounty programmes.

In many countries, researchers face significant legal risk when reporting vulnerabilities to vulnerability owners. Vulnerability owners can threaten researchers with legal proceedings instead of welcoming their vulnerability reports. This legal risk, aggravated when stakeholders are located across borders, creates powerful disincentives and a chilling effect in the security community.

Co-ordinated Vulnerability Disclosure (CVD) is a key best practice to treat vulnerabilities effectively. In a CVD process, vulnerability owners and researchers work co-operatively to discover vulnerabilities, develop, disseminate and apply patches that fix them, and disclose vulnerability information broadly without giving attackers a chronological advantage. However, CVD may be complex, in particular when co-ordination involves numerous stakeholders, such as when the vulnerability is located in a component disseminated across many products. Furthermore, each discovery of a vulnerability is unique and CVD may not be appropriate nor possible in some cases.

Policy makers can play a decisive role

Public policies can encourage stakeholders to treat vulnerabilities more efficiently. For example, they can:

- **Change the culture and mind-set** by breaking the "vulnerability taboo", recognising that vulnerabilities are a "fact of digital life" that can be mitigated through the adoption of best practices;
- **Update imperfect cybercrime and intellectual property frameworks** to better protect security researchers, for example through "safe harbours";
- **Lead by example** by adopting vulnerability treatment within the public sector and leverage public procurement;

- **Include vulnerability treatment in regulation, standards and guidance**, including as an indicator of compliance;
- **Ensure stakeholders' access to a trusted co-ordinator**, who can help connect stakeholders, provide additional technical analysis and support;
- **Increase stakeholders' trust in the government**, for example by separating offensive functions from digital security agencies and CERTs, and establishing transparent processes regarding how the government processes vulnerability information;
- **Encourage international co-operation**, such as the establishment of a non-governmental international co-ordinator, the internationalisation of vulnerability databases, the development of common principles to establish safe harbours for researchers, and the development of international standards and best practices.

In taking action, policy makers need to keep in mind that:

- **There is no one-size-fits-all solution to vulnerability disclosure.** It is a "wicked problem", without a panacea, and requiring an open mind, flexible solutions and case-by-case consideration;
- **Governments should use mandatory regulation with caution**. For example, mandatory reporting of vulnerabilities to the government is particularly challenging and many experts suggest adopting a voluntary approach based on mutual trust.

Introduction

Code is at the core of digital transformation. Every digital device embeds code to perform its tasks. All computers and smartphones run code. Data can flow through the internet thanks to code in routers, gateways, modems, etc. Code also powers industrial and consumer Internet of Things (IoT) devices, ranging from electricity meters and medical equipment, to heating systems and children's toys. Code isalso called software, or firmware when embedded in hardware. All ICT revolutions over the last decades, from the invention of databases, to the internet, cloud computing, artificial intelligence and blockchain were either based on or turned into reality through code

If data is the oil of digital transformation, then code is its engine. An increasingly complex engine. Although code complexity is difficult to quantify, software size provides an approximation. Between 1992 and 2007, Microsoft Windows increased lines of code from 2 million to 40 million (from Windows 3.1 to Windows 7). Today's typical new car includes 100 million lines of code, a typical iPhone or Android application has tens of thousands of lines of code (Perlroth, 2017[6]; Wilson, 2013[7]) and the size of software source code is estimated to double every three and a half years (North, 2019[8]).

Information systems have also become extremely complex. Governments and large firms manage tens and hundreds of thousands of connected devices with each of them often running dozens of different applications. Tracking what devices and software are in operation, where they are, and what function they perform, is a colossal endeavour. Smaller organisations also struggle to manage their digital assets.

However, the code engine of digital transformation has issues. Code is never perfect. It almost always has vulnerabilities, namely weaknesses or bugs that can be exploited to damage economic and social activities. According to estimates, the average software development process usually results in 20 to 100 flaws every 2 000 lines of code, down to one flaw every 2 000 lines if security guidelines are followed (DHS and DoC, 2018[9]; Dean, 2018[10]), a number which is likely to increase in proportion with software complexity.

All information systems also have vulnerabilities related to how software is implemented, configured, and kept up-to-date. Threat actors, such as criminals and other ill-intentioned players, actively seek to discover those vulnerabilities and develop or use tools (such as "malware") to exploit them, in order to commit cybercrimes or perform other malicious activities. They steal money, personal data, trade and State secrets, interrupt business operations and supply chains, disrupt critical activities, and ransom firms, cities, and hospitals.

Developers should therefore look and test for vulnerabilities in their code, develop patches that fix them and distribute these patches to users to reduce digital security risk. Organisations should also monitor their information systems to ensure that patches are appropriately applied and to eliminate configuration weaknesses. However, addressing vulnerabilities is a complex, burdensome and expensive endeavour for both software makers and information system owners. It is also a never-ending task because threat actors never stop their efforts, continuously discovering new vulnerabilities and new attack techniques.

Nevertheless, producers of software and hardware, as well as system owners are not alone in this race against malicious actors. A broad international community of security researchers also hunts vulnerabilities and is eager to report and disclose them in order to contribute to digital security risk reduction.

However, vulnerability disclosure can become counterproductive if not managed appropriately. For example, when vulnerabilities are publicly disclosed, malicious actors can exploit them for attack purposes. If they are offered on the black or grey market rather than reported to the party best placed to mitigate them, threat actors can purchase and operationalise them for offensive purposes, increasing digital security risk for all legitimate stakeholders. Software and hardware producers can also fail to manage timely a vulnerability reported to them by a security researcher who may then consider disclosing it publicly as a means to put pressure on them to fix the vulnerability. For example, in 2017 a security researcher reported a serious vulnerability in the MySpace website through which an attacker could log in to any one of the 3.6 million active users' accounts in a few easy steps. After three months without any action from the company, the researcher publicly disclosed the vulnerability in a blog post and it was fixed within a few hours. The company never got back to the researcher (Spring, 2018[11]). When malicious actors are the only ones aware of a vulnerability or have a chronological advantage, digital security risk increases for all stakeholders, from the owners of the vulnerable systems, to their users who can face a disruption of service or some other harm, to third parties who can be attacked through compromised products.

When malicious actors discover a vulnerability first, digital security risk increases for all stakeholders, from the owners of the vulnerable systems, to their users who can face a disruption of service or other harm, to third parties who can be attacked through compromised products. At a macro level, the consequences of digital security incidents undermine trust and efforts to realise the benefits from digital transformation. When attacks target critical activities such as the delivery of energy, health care or emergency services, the society and economy as a whole can be disrupted. Furthermore, attacks targeting systems controlling physical devices can affect human safety.

This report aims to raise policy makers' awareness about the importance of responsible "vulnerability treatment", namely the discovery, management and handling as well as co-ordinated disclosure of digital security vulnerabilities in products and information systems (definitions are provided in the Glossary).

It analyses the roles of stakeholders, existing good practice as well as challenges and obstacles to their adoption, and how public policy can help address them. Recognising that vulnerability management, handling and disclosure very often take place across borders, this report also discusses possible avenues for international policy guidance including on the international co-ordination of approaches.

The primary target audience of the report is the community of digital security and digital economy policy makers. Digital security experts may benefit from this work as they interact with government policy makers as well as with business leaders and decision makers in their own organisation. However, this document is not a technical report or operational guide. For such technical guidance, experts should refer to the most recent version of internationally recognised standards, such as ISO/IEC 29147 on Vulnerability disclosure and ISO/IEC 30111 on Vulnerability handling processes. To keep the report focused on its primary target audience, some technical terms and definitions may not be fully aligned with such technical documents.

This report is organised in three chapters:

- The first chapter introduces key concepts such as vulnerabilities and their mitigations, zero-day and exploits, stakeholders, vulnerability disclosure, as well as vulnerability handling and management. This chapter also introduces the concept of vulnerability treatment to capture this area more holistically. It also discusses the key challenges related to vulnerability disclosure from a public policy perspective.
- The second chapter describes the Co-ordinated Vulnerability Disclosure (CVD) process that emerged among the technical community as a best practice to address these challenges. It discusses legal risk faced by researchers and introduces tools that contribute to mainstreaming CVD such as standards, vulnerability disclosure policies, bug bounty programmes and platforms, as well as co-ordinators. Lastly, this chapter provides a high-level overview of good practices for CVD based on an analysis of guidance documents listed in Annex 1.

- The third chapter introduces possible high-level public policy guidance, which could be advanced at the international level.

A glossary provides explanations of key terms used in this report. To keep the report focused on its primary target audience of public policy makers, some technical terms and definitions may not be fully aligned with technical documents. Annex 2 provides a list of possible areas for future work.

1. Key concepts and challenges

This chapter introduces the scope of this report through the description of the key concepts and challenges related to vulnerability treatment from a public policy perspective.

1.1. Key concepts

This report focuses on vulnerabilities in products' code and on the way products are implemented in information systems (hereafter "systems") by organisations. It does not address the management of vulnerabilities by consumers and individuals.

To ensure consistency with the reports on "Enhancing digital security of products", the term products refers to "smart" products, i.e. products that contain code and can interconnect. Code can be defined as the set of instructions forming a program executed by a processor. Products can be goods or services, tangible or intangible, hardware and/or software, rely on open source or proprietary code, and can be commercialised or available for free. A discussion of this definition is provided in (OECD, 2021[3]).

Unless specified otherwise, the term product in this report means products beyond the design stage, i.e. available to users or still in use although no longer available on the market. Broader issues related to the digital security of products throughout their lifecycle more generally are addressed in (OECD, 2021[3]) and (OECD, 2021[4]).

Products are used within information systems. Some products are core to information systems because most of the other products rely on them. They include operating systems as well as network devices and components (e.g. routers), etc. Other products include countless generic and specialised programmes and devices, ranging from desktop office software to sales, and HR applications, to industrial control systems monitoring hydroelectric valves and gates, to wirelessly connected medical devices, etc.

The following sections introduce digital security vulnerabilities and their mitigations (1.1.1), as well as zero-day vulnerabilities, the difference between vulnerability and vulnerability information, and the notion of exploit (1.1.2). The section also discusses the key stakeholders involved (1.1.3), the vulnerability disclosure lifecycle (1.1.4) as well as vulnerability handling and management (1.1.5). Lastly, this section introduces the concept of vulnerability treatment to overcome the lack of an expression covering this entire area (1.1.6).

1.1.1. Digital security vulnerabilities and mitigations

The term "vulnerable" derives from the Latin "vulnus", which means "wound". To be "vulner-able" means to be capable of being wounded or harmed. A digital security vulnerability is a weakness that, if exploited, triggered, or activated by a threat, has the potential to cause economic and social damages, by affecting availability, integrity, or confidentiality of a digital resource or asset.[1] Hereafter, the term "vulnerability" will refer to a digital security vulnerability.

This report focuses on two types of vulnerabilities affecting a product: vulnerabilities in the product's code (code vulnerabilities) and vulnerabilities related to a product's implementation within an information system (system vulnerabilities). Both are introduced below.

A mitigation measure, or simply “mitigation”, is the antidote to a vulnerability. Code and system mitgations are also introduced below.

This report does not address other types of digital security vulnerabilities that can affect information systems, typically the lack of security measures such as the absence of a backup procedure in an organisation, or poor digital security awareness that can lead a person to click on a malicious link in an email.

Hereafter, the standalone term “vulnerability” will refer to both code and system vulnerabilities.

Code vulnerabilities and their mitigation

Code vulnerabilities

For the purpose of this report, a code vulnerability is a vulnerability affecting the code embedded in a product. Such code can be found in the product’s software and hardware components (e.g. firmware).[2] This report does not specifically distinguish proprietary from Open Source Software (OSS). However, issues related to vulnerability disclosure in OSS raise special challenges that would deserve a separate analysis (cf. Annex 2).

There are hundreds of different types of code vulnerabilities (e.g. “Buffer Copy without Checking Size of Input ('Classic Buffer Overflow')”, also known as CWE-120). For example, as of January 2020, the United States MITRE Corporation’s Common Weakness Enumeration (CWE) identified 808 types in its effort to help developers and practitioners use a common language (n.d.[12]). The

Publicly known vulnerabilities are registered in vulnerability databases. The most well-known system for referring to a vulnerability is the Common Vulnerabilities and Exposures (CVE) list maintained by MITRE. The CVE list is used and incorporated in many other products and services, such as the US National Vulnerability Database (NVD) maintained by the US National Institute for Standards and Technology (NIST). As of January 2020, the CVE list contained over 129 000 entries which are all assigned a unique CVE number. The community-driven vulnerability database VulDB contained 146 981 entries covering 35 809 products (Vuldb.com, n.d.[13]).

While these figures give an idea of the number of vulnerabilities we are dealing with, they do not provide an accurate and comprehensive picture of the scope of the challenge. For example, many discovered vulnerabilities are never disclosed, databases do not cover all products in use globally, etc. Experts generally consider that there are undiscovered vulnerabilities in every piece of software (Schneier, 2018[14]; UK NCSC, 2016[15]). According to a security firm, 100% of the applications scanned by the company’s security products in 2018 and 2017 contained at least one vulnerability of a known type, with a median number of 15 vulnerabilities per application (Trustwave, 2019[16]). Naturally, these types of analysis cannot test applications against types of vulnerabilities that are yet to be discovered.

However, code vulnerabilities are not created equal. Discovered code vulnerabilities have different levels of risk and severity. The risk related to a code vulnerability depends upon the use context of the product in which it is located, which can vary considerably across users for the same product. The severity of a vulnerability depends upon its degree of exploitability and the amount of damages it can create regardless of the product’s use context. For example, a vulnerability that enables the unauthenticated remote execution of code typically has a critical level of severity. However, it does not necessarily mean that the related level of risk will be high in all use contexts. In some cases, the vulnerable product may benefit from other layers of security making it extremely unlikely that the vulnerability is reached and exploited by a threat actor. Conversely, the risk associated with a low severity vulnerability can be high in some contexts where the code component is particularly exposed to threats. Furthermore, some code vulnerabilities can be easily exploited while others can only be exploited in theory or in the lab, or the cost to develop an attack exploiting them would be prohibitive.

Severity is often rated against the Common Vulnerability Scoring System (CVSS), an open framework maintained by the international Forum of Incident Response and Security Teams (FIRST) for communicating the characteristics and severity of vulnerabilities. CVSS scores from zero (no severity) to 10 (critical severity) are calculated based on a formula that depends on metrics approximating the ease and impact of a potential exploit and other aspects such as the availability of mitigations (FIRST, n.d.[17]). As it does not take into account the use context, the CVSS score is not a risk metric (Spring et al., 2018[18]).

Many code vulnerabilities do not have a critical level of severity, and some do not necessarily require action. The European Union Agency for Cyber Security (ENISA) found that only 8.65% of a large dataset of vulnerabilities reported in 2018 and half of 2019 could be exploited in practice. 20% of the exploitable vulnerabilities in this research dataset had a critical CVSS severity score (492 out of 2 377) (ENISA, 2020[19]). Vulnerabilities registered in the US National Vulnerability Database (NVD) have an average CVSS score of 6.6 (see Figure 1).

Figure 1. Vulnerability distribution by CVSS Scores (March 2020)

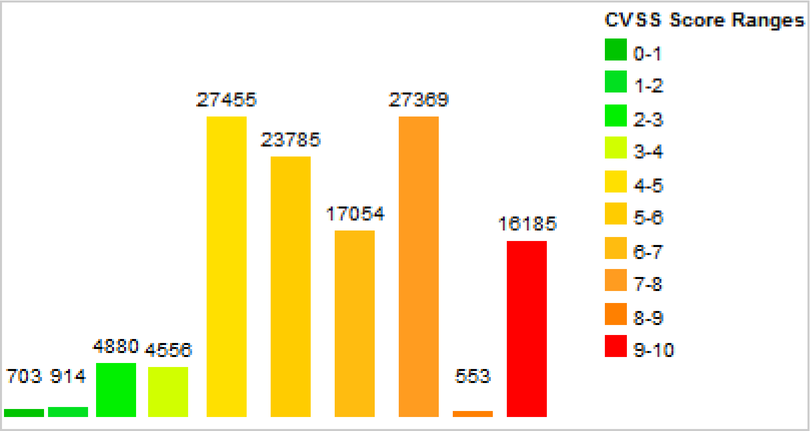

Note: This graph shows the distribution of vulnerabilities in the US National Vulnerability Database (NVD) according to their CVSS score, where 0 = no severity and 10 = critical severity (see https://nvd.nist.gov/vuln-metrics/cvss). The NVD only includes reported vulnerabilities rather than all existing vulnerabilities. The average CVSS in the NVD is about 6.6.
Source: cvedetails.com

In addition to these known vulnerabilities, there are many not yet discovered. They are latent and therefore cannot be exploited. Vulnerabilities can be present but unnoticed for years or decades: some were discovered in Microsoft's Windows XP more than ten years after its release, and years after the end of its commercial life. The ratio of the vulnerabilities that are not known to the vulnerabilities that are known is undetermined.

The prevalence of vulnerabilities among different types of products is difficult to measure. An automatic analysis of 1.4 million applications in 2019 found that 85% had at least one vulnerability of a known type and more than 13% at least one critical flaw (Veracode, 2019[20]). However, such figures are likely to reflect biases related to the sample of products tested, and the tools used to scan them. According to research carried out in 2017, the vast majority of known vulnerabilities are simply not present in real-world information systems. The population of distinct vulnerabilities that actually resides in enterprise environments represented only 23% of all entries in the CVE database in 2017. Nevertheless, this percentage still represented 22 625 out of 107 710 CVE entries (Tenable, 2018[21]).

Assuming it is possible, it would be highly time-consuming and resource-intensive to develop code without any weaknesses, in particular with modern software which is iteratively updated with new functionalities at a fast pace. It is possible to optimise software development in order to reduce the number of potential vulnerabilities from the outset, a good practice called "security by design".[3] However, following such good practice takes time and requires significant resources, which may be at odds with other objectives such as time-to-market or cost-reduction. Furthermore, attackers are innovative, use dynamic techniques and evolve rapidly. While code could be deemed sufficiently free of vulnerabilities when the product is released on the market, it may no longer be at a later stage, even a few days or hours after the product release, in light of new attack techniques and updated security practices.

Mitigations for code vulnerabilities

A code vulnerability can be addressed through a code mitigation called a patch that modifies the released code. Patches have to be implemented on each software instance through a security update, broader update, new release or upgrade (e.g. in mobile apps) comprising the patch and other code modifications for example improving functionalities.

However, it is not always possible to develop a patch. For example, some products have reached end of commercial life and are no longer supported; fixing certain code vulnerabilities can require redesigning the product which would be too costly, or would raise performance or compatibility issues for customers (Johnson and Millett, 2019[22]); and some products do not have update capabilities, such as certain low-cost IoT devices. In such cases, the mitigation can be a set of instructions and configuration requirements or documentation changes that help reduce the risk related to the vulnerability without necessarily eliminating it as a patch would most likely do. Sometimes, the vulnerability owner simply informs users that a newer version of the product is available and is being actively supported.[4]

Currently, the development of mitigations for code vulnerabilities is largely driven by market forces, with market failures at play.[5] In the United States, a recent report underlined research suggesting that "50% of vulnerabilities remain without a patch for more than 438 days after disclosure" (Cyberspace Solarium Commission, 2020[23]).

System vulnerabilities and their mitigations

For the purpose of this report, a system vulnerability is a weakness in the way a product is implemented or configured. System vulnerabilities include deficient vulnerability management and misconfiguration.

Deficient vulnerability management

Failure to keep implemented products up-to-date with the latest mitigation (a term discussed below) is a major source of system vulnerabilities. Rather than spending time and resources to discover new code vulnerabilities, most attackers exploit known vulnerabilities in products that system owners have not patched. If attackers know in advance which organisation or individual they want to compromise, they test their future victim's systems against known vulnerabilities until they find one that has not been patched and thus can be exploited (Verizon, 2019[24]). According to a 2019 survey of 2 900 IT professionals in nine countries,[6] 60% of respondents say one or more breaches they faced occurred because a patch was available for a known vulnerability but not applied (Ponemon Institute, 2019[25]).

While silent and automatic patching is a reasonable objective for consumer products, it is less so for more complex information systems in organisations. In contrast with home users, many organisations cannot simply apply all security updates as soon as they receive them, in particular in industrial environments. They often need to test them first to assess whether the patch itself is not going to disrupt operations or introduce new security, compatibility, performance or instability issues through domino effects in chained systems. To ensure that all the smart products in use are running with the most recent security update or mitigation, organisation implement a "vulnerability management" process, which includes "patch management".[7] Vulnerability management can be a heavy process requiring, for example, test environments and procedures. According to a survey, it takes organisations an average of 102 days to test and fully deploy patches, and more specifically 16 days to patch a critical vulnerability, and 151 days to patch a medium or low priority vulnerability (Ponemon Institute, 2018[26]; Ponemon Institute, 2019[25]). In many cases, organisations may have a risk-based rationale for long patching delays, or even no patching at all for example in cases where assembly lines or physical processes cannot be interrupted. In other cases, such delays just result from poor vulnerability management. For example, some security experts found publicly traded companies with up to three year-long average patch application cycles, which means that machines were, on average, exposed to all known vulnerabilities disclosed over the last three years!

There are many infamous cases of attacks that would have failed if the system owner had swiftly applied a security update. For example, hundreds of thousands of organisations lost assets during the 2017 WannaCry attack because they had not applied the patch released two months earlier by Microsoft. When NotPetya hit organisations one month after WannaCry, leveraging the same vulnerability, many organisations, including some global firms, had still not implemented the patch. The result was multi-billion dollars of global damages. Other examples include the 2017 Equifax incident that affected 56% of all American adults, 14 million British citizens and about 20 000 Canadians, costing the company at least USD 1.4 billion (Schwartz, 2019[27]). The attackers exploited a vulnerability in a product for which a patch had been available for two months prior to their intrusion in the system (Hay Newman, 2017[28]).

Vulnerability management is not yet the norm because it is both difficult and expensive. According to the same 2019 Ponemon survey, only 15% of the respondents considered that their organisation is at a middle stage of maturity with respect to vulnerability management, leaving 85% at an earlier stage. Keeping up with security updates was extremely challenging or challenging for 65% of the respondents whose company had a patch management process in place. Only 37% considered that the IT security function had adequate staffing to patch vulnerabilities in a timely manner. 45% and 40% pointed out, respectively, higher automation and increased IT security staff as key to improve patch management.

Misconfiguration

Many system vulnerabilities are related to weaknesses in products' configuration or settings. They may result from improper or outdated configuration caused by administrators lacking security awareness or knowledge of how to deploy or use products in a sufficiently secure manner in their digital environment. Administrators may also lack resources or time to configure their products appropriately, or they might deploy information systems without managing the security in a systematic and ongoing manner. For example, an administrator may set up a poorly secured server to test a product or respond to a one-time need, and then forget to discard it afterwards. By keeping the server live, the administrator enables threat actors to exploit it and potentially access the entire information system or use the server as a proxy to reach to other targets. Administrators need to implement a comprehensive technical risk management strategy (whether formal or informal depending on their organisation's size and complexity) to systematically review and update security settings throughout their organisation's systems lifecycle.

Many products are shipped with minimal security settings by default, making them "vulnerable by default". A typical example is a product shipped to all its users with the same weak default password and without a mechanism forcing users to change it at first installation. All users who do not know they are supposed to change this default password upon taking possession of the product would therefore be easy targets for any potential attacker. There are many possible reasons why numerous products have weak default security settings, from a lack of awareness to a misperception of risk and a misallocation of risk ownership. Many stakeholders and security experts now call for products to be "secure by default", i.e. provided with sufficiently high security settings when first installed or used, leaving it to users to take the responsibility for weakening security as appropriate.[8]

Mitigations for system vulnerabilities

To mitigate system vulnerabilities, system owners need to take appropriate action such as changing configuration settings or applying the existing patch that was previously set aside.

There is no way to eliminate all vulnerabilities

Addressing vulnerabilities is essential, but fixing all vulnerabilities would not be a realistic objective.

First, some vulnerabilities are latent. A latent vulnerability can mean that nobody has yet thought about a particular type of vulnerability, and therefore nobody can look for the presence of vulnerabilities pertaining to that type in any existing product or system. Alternatively, it can also mean that nobody has yet tested a specific product or service against a given known type of vulnerability and therefore the presence of this type of vulnerability in that product or service is a possibility.

Second, vulnerabilities have different levels of *i)* severity, according to factors such as the amount of damage they can create and the ease through which they can be exploited (exploitability), and *ii)* risk, depending upon the use context of the product in which they are located.

Third, in many circumstances, users may continue to use products that have reached their end of life without support from the vendor. In that case, vulnerabilities may still be discovered but never corrected, leaving the user without any mitigation instructions.

Fourth, while every vulnerability, even with a low level of exploitability, can be exploited in principle, most vulnerabilities will never be.

Lastly, it would be cost-prohibitive to remove every code vulnerability from a product. Some known code vulnerabilities might persist, as it can be technically complex or cost-prohibitive to develop a patch for them. Alternatively, users may not apply a patch because it would disrupt their operations, create compatibility issues and introduce additional risk. Some code vulnerabilities are not worth removing because the likelihood of their exploitation is too low, the cost of fixing them is prohibitive in light of their low severity, or the cost of fixing them in advance or after an exploitation is not particularly different. In some cases, another mitigation may exist, such as a configuration change in a firewall or some other workaround, that protects the system and the data in absence of a code mitigation or when a patch cannot be applied.

It is therefore important to approach vulnerability issues with the objective of making products and information systems "secure enough" rather than absolutely secure, in order to sufficiently *reduce,* rather than *eliminate,* security risk for users and third parties. Vulnerability owners have to perform a risk assessment to prioritise their efforts to address vulnerabilities they are aware of (cf. 1.2.5). In this respect, the above mentioned CVSS severity scoring standard, which does not take the use context of the vulnerable product into account, has been criticised as being often inappropriately used as a risk assessment and vulnerability prioritisation tool (Spring et al., 2018[18]).

1.1.2. Zero-day, Vulnerability Information, and Exploit

Zero-day

Depending on the context, a "zero-day vulnerability", "zero-day" or "0-day", can be a code vulnerability:

- For which no mitigation has yet been released, i.e. system owners have had zero-days to apply a patch; or
- Which is unknown to the code owner (a concept introduced in 1.1.3), i.e. the code owner has had zero-days to develop a mitigation.

The notion of zero-day is not relevant for a system vulnerability. Figure 2 provides an overview of zero-day related terminology. When nobody is aware of it, a vulnerability is latent (1). Some vulnerabilities remain latent forever, while others are discovered and become a "zero-day" (2). When the code owner develops a mitigation (e.g. patch, fix, instructions), the zero day becomes a "N-day" vulnerability, where N = the number of days since the mitigation has been available (3a). From then on, it is possible to reduce or eliminate the risk by using the mitigation. However, the code owner may also never fix the vulnerability, which then becomes a "forever-day" (3b).

According to the above, each known vulnerability (i.e. N-day) in a product has been a zero-day prior to a mitigation becoming available. Surprisingly however, a known vulnerability can become zero-day again if it is in a product embedded as a component in another product which does not support an update mechanism for that component. For example, a video doorbell may embed a web server that allows remotely changing the doorbell's settings and reviewing pictures of visitors. If the doorbell designer did not include a mechanism for the embedded web server to receive security updates, the doorbell may include zero-days that are otherwise N-days for a standalone server.

Figure 2. Typical evolution of a code vulnerability

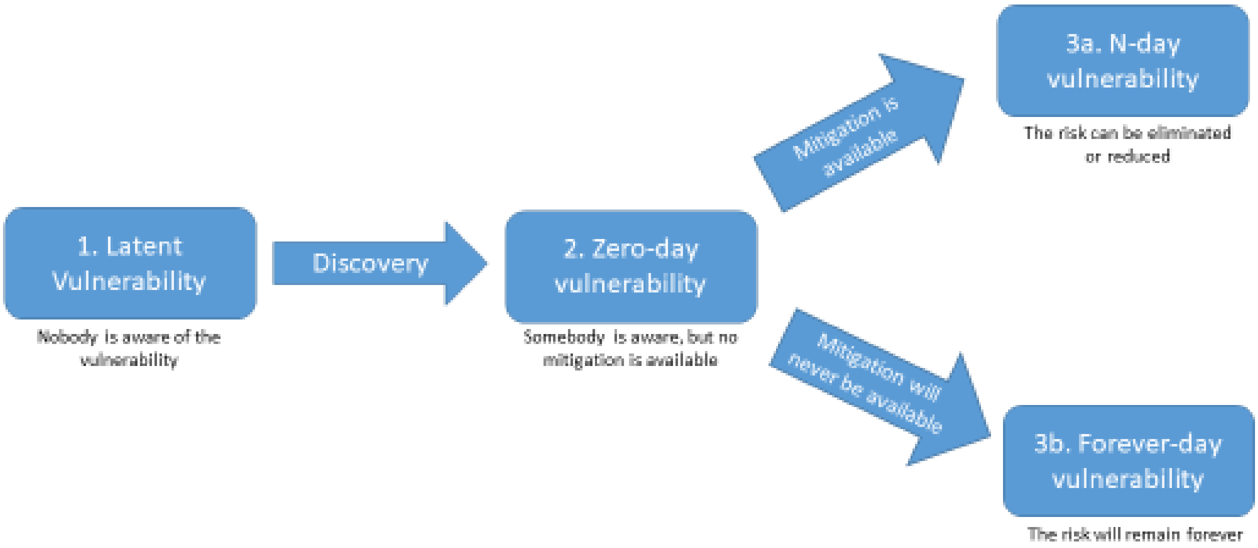
1. Latent Vulnerability
Nobody is aware of the vulnerability
Discovery
2. Zero-day vulnerability
Somebody is aware, but no mitigation is available
Mitigation is available
3a. N-day vulnerability
The risk can be eliminated or reduced
Mitigation will never be available
3b. Forever-day vulnerability
The risk will remain forever

Source: OECD

Vulnerability information

The term "vulnerability" can refer to the actual exploitable weakness (as introduced above) or to "vulnerability information", i.e. the set of information that describes a particular vulnerability and empowers the vulnerability owner to address it. Vulnerability information is most sensitive when known by a single actor and least sensitive when it is widely known in the security community and industry. In practice, code vulnerability information can often include functional *proof of concept* code, which can be used to create a programme to gain unauthorised access or otherwise interfere with the computer system (EFF, 2018[29]). With both the description of the code vulnerability and the proof of concept code, a code owner can more easily test and understand the vulnerability in order to more rapidly develop mitigations, and deliver them to end-users to decrease the likelihood of its exploitation.

Exploit

Exploit code, often called “exploit”, is code developed to weaponise a vulnerability. Developing an exploit requires skill and takes time and effort. Exploits can be harmful or beneficial depending on who uses them and for what purposes. A threat source can use vulnerability information to develop an exploit for offensive purposes. However, when an organisation patches a vulnerability, attempting exploitation is often an effective way to test its effectiveness. Therefore, defenders also need exploits to carry out network penetration testing and “red teaming”, a more advanced form of network penetration test, which may be required by some regulation and non-regulatory standards to identify an organisation's vulnerabilities for mitigation. Some widely used security tools such as the open source Metasploit Framework and proprietary Metasploit Project owned by Rapid7 aim at facilitating defence by automating the use of exploits. These tools facilitate the development and execution of exploit code against a remote target machine. The Metasploit Framework includes almost 2 000 exploits applicable to a large number of platforms. As noted above, many vulnerabilities cannot be weaponised.

An attack based on exploit code is called the *exploitation* of the vulnerability. Exploitation can also include commercialisation of the exploit code and of attack services on black markets. Depending on how the exploit has been developed, exploitation can be extremely easy, including through turnkey attack solutions. Weaponisation and exploitation are the two high-level stages of the *exploit lifecycle*, which often loops into a cycle of improvement.

A “zero-day exploit” is an exploit based on a zero-day vulnerability.

Zero-day exploits have a high likelihood of success until a mitigation is available because they are more difficult to detect and remediate. However, zero-day attacks are rather less common in comparison with attacks exploiting known vulnerabilities, although there is limited quantitative evidence to support this claim. This is likely because once an attack using a zero-day exploit is detected, security experts will share its characteristics to improve the effectiveness of attack detection tools. They will also reverse engineer the exploit code and inform the code owner who will develop a patch. The zero-day exploit will become a known exploit, decreasing its effectiveness and value over time as mitigations become available and are deployed. Therefore, attackers tend to use zero-day exploits for targets of highest value to them. In most cases, they will try to achieve their goals with exploits based on known vulnerabilities first, and use an expensive zero-day exploit only as a last resort and for targets that are worth it. Sophisticated threat sources commonly known as Advanced Persistent Threats (APTs), including State-sponsored groups, would follow the same approach and avoid spoiling zero-days when system vulnerabilities can achieve the same results.

1.1.3. Stakeholders

This section reviews the main categories of stakeholders who are directly concerned with vulnerability-related processes (cf. Table 1). Given the global nature of the digital environment, these stakeholders can be located anywhere in the world and their interactions can take place regardless of their geographical location.

Table 1. Categories of stakeholders

Vulnerability owners	Market intermediaries
○ Code owners	Third party victims
○ Systems owners	Defenders
Security researchers	Threat sources
Co-ordinators	Other stakeholders

Vulnerability owners

Many documents about vulnerability disclosure use the "vendor" concept to refer to stakeholders who can receive vulnerability reports related to a product and fix the related vulnerabilities. However, the term vendor is misleading for several reasons:

- The actual vendor (i.e. seller) who is in a commercial relationship with the products' users is often not the party receiving vulnerability reports and does not intervene in the product's code.
- The term vendor excludes products available for free, such as many open source products.
- The term vendor hides the complexity of most products that include many different layers of code and code components developed by different parties, which may or may not have established relationships. The provenance of some code components may even be difficult to track (e.g. code snippets). In other words, a product may embed many other products as components. A vulnerability may be located in a code layer that the final product "vendor" can often not address on its own.
- Some products are similar but sold under different brands by different vendors. It is not always clear which products are affected by a given known vulnerability.
- The vendor notion is useless for system vulnerabilities. The party who receives system vulnerability reports and can act upon them is the system owner, not the product vendor or the stakeholder responsible for the product code.

Therefore, this report rather considers the broad notion of "vulnerability owners", namely the stakeholders who own the responsibility to act upon a vulnerability they are aware of, in order to mitigate it. The term "owner" focuses on the responsibility to address vulnerabilities (as in "risk ownership") and does not necessarily entail property rights.

There are two kinds of vulnerability owners: code owners and system owners, introduced below. With digital transformation, all organisations tend to become both a code owner and a system owner, as "every company is [becoming] a software company", as noted by Microsoft's CEO S. Nadella (Microsoft, 2018[30]).

This report does not specifically address vulnerability owners according to their size (e.g. Small and Medium-sized Enterprises, local governments, etc), although it recognises that size can play an important role in how vulnerabilities are addressed. This is a potential area for future work (cf. Annex 2).

Code owners

Code owners are the individuals or organisations who developed the layer of code where the code vulnerability is located in a product or/and are best placed to fix it. They "own" the responsibility to address code vulnerabilities.

A code owner is not necessarily the vendor of the product, i.e. the organisation or individual providing the product to users.

Most products include layers of code developed by different stakeholders which are (or should be) responsible for addressing the code vulnerabilities in the layer they developed. The concept of code owner, borrowed from (FTC, 2018[31]) and further discussed in (OECD, 2021[3]), is useful to overcome the limitations of the term vendor. It also helps map responsibilities in products embedding multiple components developed by different stakeholders, which is common. In a smartphone, for example, the chip manufacturer owns the deep level code on hardware; the operating system designer owns the operating system code; the app designer owns the app code, etc. In a smartphone app, the app designer may have developed the main code, but external companies may have developed other functionalities such as analytics, social network sharing buttons or advertising integrations. Such code may contain vulnerabilities that can create risk to users of all applications embedding it.

With digital transformation, traditional ICT players are no longer the only code owners. Rather, with the proliferation of applications, and "smart" objects and IoT devices, code owners increasingly include stakeholders not previously associated with the ICT industry, ranging from banks, to grocery stores, local governments, newspapers and television channels, smart cooking device makers, car and tractor manufacturers, etc. As digital newcomers, many of these stakeholders have a lower level of digital maturity compared to ICT firms. Many of them, however, outsource code development to traditional ICT players, or reuse existing code previously developed by traditional ICT players.

In some cases, the identity of the code owner may be unclear. For example, users and researchers often only know the vendor and cannot interact with the many code owners involved in the various layers of the product, which are unknown to them. Furthermore, the code may have been voluntarily developed as a free software or open source project. In these cases, code ownership is voluntarily assumed and shared among a community of individuals and organisations, which may differ from time to time. The code owner of proprietary software may have gone out of business, or key contributors in an open source project may have stopped maintaining the code and keeping the project alive.

System owners

System owners are the organisations using products within their information system. They are responsible for these products' configuration and for applying security updates provided by code owners ("vulnerability management"). With digital transformation, almost every organisation is increasingly likely to rely on increasingly complex systems, without necessarily having sufficient awareness and understanding of digital security.

As with code owner, the term "system owner" focuses on the stakeholder's ownership of the responsibility to address vulnerabilities in the system rather than on property rights related to that system.

System owners can also be viewed as products' users. As such, they can suffer directly from the exploitation of a code vulnerability in a product they use. Their money or trade secrets can be stolen, their operations disrupted, and their reputation undermined. In contrast, code owners face indirect consequences from their products being used to attack victims, i.e. their reputation can be damaged and they can lose customer trust, depending on how they address the code vulnerability.

For the purpose of this report system owners do not include products' consumers as individuals in their personal capacity.

Security researchers

Security researchers ("researchers") are the individuals or organisations who identify a potential code or system vulnerability with the intention to reduce security risk.

Security researchers are often also called "finders" or "discoverers" in some security documents, as well as "ethical hackers", "white hats" and "friendly hackers" in internet slang and in the media. The term researcher is more neutral than the term "hacker", which has become ambiguous and often carries negative connotations.[9] The negative meaning of "hacker" is better conveyed by the term "cybercriminal", "criminal", or "malicious actor".

Different categories of researchers are driven by different goals and operate under different constraints. Many are security experts who research vulnerabilities as part of their professional or personal activities. They can work for academia, commercial security companies, product security teams, government agencies or civil society. They can also find vulnerabilities as a personal hobby in their spare time. Researchers may not have recognisable security credentials, as demonstrated by the anecdote of a 5-year-old child discovering a vulnerability in the Microsoft Xbox game console (Cluley, 2014[32]). There are many cases where end-users without significant security skills have discovered and reported vulnerabilities.

Code and system owners can take the role of security researchers with respect to products and systems they or other parties' own, such as with Google Project Zero (Google, n.d.[33]).

Co-ordinators

Co-ordinators are stakeholders who can assist code and system owners as well as researchers in the vulnerability disclosure process. They can help with expertise, language, time zone and cultural barriers between code owners and researchers. They can also act as a facilitator between different parties, including when multiple code owners or value chain stakeholders are involved (cf. 2.1.2).

Market intermediaries

Market intermediaries include bug bounty platforms (cf. 2.3.4) and grey market vulnerability brokers (cf. 1.2.6). Some market intermediaries can act as co-ordinators.

Third party victims

Third-party victims of digital security incidents are the individuals as well as public and private organisations anywhere in the world who can be harmed by the exploitation of vulnerabilities without being users of the vulnerable products themselves. For example, patients whose surgery operations were delayed due to ransomware attacks on hospitals, individuals caught in the 2015 and 2016 black outs in Ukraine, or whose privacy was violated following a data breach.

Defenders

Defenders are stakeholders who are responsible for defending against attacks or provide tools and services to do so. They include, for example, developers of security software that detect vulnerable systems, or detect and respond to, and firms selling network penetration testing and incident response services.

Threat sources

Threat sources are governments, groups or individuals who identify and exploit vulnerabilities for malicious or ill-intentioned purposes. Their motivations can vary, but typically include geopolitical goals for governments, profit making for criminals, ideology for hacktivists, violence for terrorists, personal aims for thrill seekers, and discontent for insider threats (Canadian Centre for Cybersecurity, n.d.[34]). Incidents can also result from unintentional threats such as a human error or a power cut.

Other stakeholders

Other actors include governments, including in their policy making and regulatory role, standards and certification bodies, insurance companies, civil society, the media, etc.

The multiple roles of government

Governments simultaneously play almost all the roles described above:

- *Users* of products;
- *Vulnerability owners* of their own products such as e-government services and web sites, and internal digital government platforms;

- *Security researchers* when an agency is tasked with discovering vulnerabilities;
- *Third-party victims* of incidents, for example when the public administration is disrupted by incidents taking place elsewhere;
- *Threat sources* if law enforcement and national security agencies develop or purchase offensive tools, and when they perform offensive "cyber operations".

Although it is beyond the scope of this report, it is important to understand the possible role of governments as threat sources because it may undermine the trust and confidence that stakeholders need to have in governments' efforts to reduce digital security risk. Box 1 provides further information about the role of governments as threat sources.

Box 1. The offensive role of governments with respect to vulnerabilities

Cybercriminals exploit vulnerabilities to attack their targets. Therefore, vulnerabilities are at the core of digital security risk. All stakeholders must focus their efforts on addressing them, as explained in this report. Among all stakeholders, however, some governments are in a special position. While they can adopt good vulnerability management practices for themselves and use public policy to encourage other stakeholders to do so, many governments are also looking for vulnerabilities to exploit as part of their law enforcement, intelligence and national security activities.

These governments may discover vulnerabilities, stockpile them without reporting them to vulnerability owners, weaponise and exploit them against public or private targets, civilian or military, domestic or foreign, targeted or in bulk. At times, they may even create vulnerabilities, or require or contract with others to do so on their behalf.

To carry out "offensive operations" (a soft term to refer to a digital security attack carried out by government), some governments can also buy vulnerabilities and exploits on the grey and black markets. Therefore, these governments contribute to price setting and legitimisation of businesses and other supply-side actors on these markets, who may in some cases also buy from or sell to cybercriminals. Some governments can also secretly require developers to insert "backdoors" in their products, which are similar to intentional vulnerabilities. The business, civil society and technical community have almost unanimously condemned such a practice, which increases risk to all stakeholders, as such backdoors are likely to be discovered and exploited by other offensive or criminal actors at some point. Some governments have also condemned backdoors.

Vulnerabilities and exploits that these governments stockpile can sometimes be detected when they are used in the course of offensive operations. Other offensive actors, such as criminals, foreign governments, or activists can steal them, and insiders can leak them to the public. Anyone who obtains these exploits can in turn use them against anyone in the world. They can also improve them, and/or commercialise them on the grey and black market. As a result, the global level of digital security risk increases. The digital security chronicle of the last decade includes several examples of such leaks, thefts, and repurposing of exploits initially kept or used by governments.

The targets and goals of offensive operations carried out by governments may be often legitimate. They include criminal investigations, the fight against terrorism, protection of sovereignty, counter espionage, etc. However, in certain cases, such operations can violate human rights and fundamental values, by targeting human rights advocates, whistle-blowers, journalists, lawyers, and simple citizens. They can also be used for economic espionage, theft of trade secrets, and other covert actions including preparation for armed conflicts. Most of the time, these operations are secret or top-secret. They sometimes become known to the public when they fail (e.g. when detected by targets), through errors, through whistle-blowers, or a long time after the fact.

In some cases, governments therefore appear to be part of both the solution and problem of digital security vulnerabilities, using the left hand to increase the risk while the right hand is struggling with other stakeholders to reduce it. This thorny issue helps understand the broader context of this report although it is partially outside of the report's scope. For example, it shades a light on the limits that some governments may have to drain the grey market and fight against the black market, such as when these governments' "national security" side is supporting a market that their "prosperity side" is trying to dry out.

Some government's ambiguity with respect to vulnerability exploitation can undermine other stakeholders' trust in governments' efforts to reduce risk. It can diminish the effectiveness of policies to promote co-ordinated vulnerability disclosure. Prior to sharing vulnerability information, or asking for assistance in a co-ordination process, stakeholders always assess whether they can trust the recipient side. Domestic and foreign governments are no exception, and most stakeholders are rather suspicious about some of them by default, for the reasons highlighted above. Therefore, from the non-governmental stakeholders' point of view, governments need to make a special effort to demonstrate that they can be trusted. For example, some countries have adopted a public governance model that strictly and transparently separates the government's defensive and offensive functions at the institutional level

1.1.4. Vulnerability disclosure

The meaning of "vulnerability disclosure" varies in the literature. Sometimes, the term covers the entire vulnerability lifecycle, from discovery to public disclosure, including some parts of vulnerability handling and management (1.1.5), which are inherently related to disclosure. "Vulnerability disclosure" can be used to refer to the provision of vulnerability information from one stakeholder to another, such as when a security researcher reports a vulnerability to a vulnerability owner or co-ordinator. It can also be used as an abbreviation for "public disclosure", i.e. the provision of vulnerability information to the public.

This imprecise use of the term is common in the technical community but can be confusing for the non-expert. It reveals that this area is still relatively nascent and rarely approached from a holistic perspective. Section 1.1.6 proposes the term "vulnerability treatment" to refer to the broad subject area.

In this document, the provision of vulnerability information to a vulnerability owner or co-ordinator is called "reporting".

This section introduces the types of vulnerability disclosure generally referred to in the literature, and describes the vulnerability disclosure lifecycle.

Types of vulnerability disclosure

Four types of vulnerability disclosure are often distinguished in the literature:

- *Non-disclosure*: vulnerability information is not disclosed to anyone.
- *Full disclosure*: vulnerability information is disclosed to the public unilaterally, i.e. without co-ordination.
- *Disclosure to third parties*: information is disclosed to other parties than those who can develop mitigations or assist in the development of a mitigation.
- *Limited disclosure*: the disclosure is limited in a manner that reduces risk to all parties, for instance to a vulnerability owner or co-ordinator, or to the public but with a low level of detail, and when mitigations are available.

In some cases, the disclosure may be unintended, such as when the vulnerability is stolen by a threat actor, which can happen when a stakeholder stockpiles vulnerabilities (cf. 1.2.8 and Box 1).

The technical community has long debated the merits of these approaches, struggling to define what a "responsible disclosure" should be (cf. 1.2.2). Overall, there is a consensus that, even when the vulnerability information is disclosed with the best intention to reduce risk, its disclosure can create damages and increase the risk depending on how it is carried out. Therefore limited disclosure, where vulnerabilities are disclosed in a way that reduces risk and minimises damages, appears as the best approach (ENISA, 2016[35]).

The vulnerability disclosure lifecycle

Vulnerability disclosure takes place within the broader vulnerability lifecycle representing events that can affect a vulnerability from its discovery to its disclosure to the public. Figure 3 provides an overview of such a lifecycle. Considering the multiple possible variations and scenarios, this figure is necessarily incomplete but sufficiently detailed for the purpose of this report.

The green steps from 1 to 7, including 2a and 2b, reflect the scenario which maximises the likelihood of the disclosure process reducing the risk level for all. Although they are represented as a linear sequence, they are in reality (or at least they should be) a learning and improvement cycle whereby each time a stakeholder goes through these steps, it learns and improves its processes and culture so as to better manage the next iteration of the cycle.

The grey, red and dotted arrows reflect a selection of alternative paths where detrimental outcomes are more likely. The figure includes the following steps:

1. Discovery

Both ill- and well-intentioned actors can first discover vulnerabilities.[10]

The discovery of a *code* vulnerability by a threat source triggers an exploit lifecycle whereby the threat source weaponises the vulnerability to develop exploit code (A. weaponisation). Then it can exploit this code in various ways (B. exploitation), such as carrying out attacks or selling, sharing, exchanging the exploit code with other threat sources who will use it themselves. Once malicious code is available, its creator can enhance it, and so can other threat sources who can access it. This creates an exploit improvement cycle. Several threat sources can independently discover the same vulnerability and trigger separate exploit lifecycles.

When a threat source discovers a *system* vulnerability, it can exploit that vulnerability directly (B), either by attacking the system or by selling the vulnerability information to another threat actor. The exploitation lifecycle leads to incidents that may be detected and mitigated if targets have an effective digital security risk management cycle in place, or otherwise lead to economic and social damage for them and, potentially, for third party victims.

Vulnerability owners, researchers and threat sources can all discover vulnerabilities. For example, researchers can reverse-engineer product code, exploit code detected in the course of attacks, or just stumble on a vulnerability by accident when encountering unusual behaviour. System owners can ask "offensive security" professionals (i.e. defenders) to carry out network penetration tests ("pen tests") to help test the effectiveness of the organisation's vulnerability management program within a defined scope. They can also practice "red teaming". When defenders discover a code vulnerability, they can trigger an exploit lifecycle that will feed the security community and reduce digital security risk (this is not represented in Figure 3 for the sake of simplicity). A threat source can rediscover a vulnerability already discovered by a researcher or vulnerability owner before it becomes public.

If the vulnerability owner discovered the vulnerability, step 2 is skipped.

If a researcher discovered the vulnerability, he/she needs to decide what to do among at least five options:

2a. Reporting to the vulnerability owner

The reporting kicks-off a process of limited disclosure whereby the researcher works with the vulnerability owner to minimise the risk related to the dissemination of vulnerability information.

The researcher may fail to report to the vulnerability owner, for example because it no longer exists, cannot be contacted, or does not respond. According to experts, lack of response from vulnerability owners is common and most vulnerability owners simply ignore vulnerability reports. In such cases, the researcher may end the process, report to a co-ordinator or other intermediary (2b), disclose the vulnerability information to the public (2c), or seek to monetise it on the grey market (2d).

2b. Reporting to a co-ordinator or other trusted intermediary

The researcher can also report to a third party mediator such as a co-ordinator or journalist[11] who can act as a proxy, facilitator and conflict resolver between the researcher and the vulnerability owner. For example, the researcher might not know how to report to the vulnerability owner who may lack a clear channel of communication; there may be language or cultural barriers; the researcher may not trust that the vulnerability owner will address the vulnerability; he/she may wish to remain anonymous and use the intermediary as a shield against legal action; etc. Bug bounty platforms may also act as intermediaries, although the vulnerability owner needs to have created a bug bounty programme on the platform in the first place (see 2.3.3).

2c. Full disclosure

The researcher makes the vulnerability information directly available to the public. The vulnerability owner can then start investigating the vulnerability information made public (step 3 below). Users and the digital security community are informed at the same time about the vulnerability, and defenders can develop exploits for testing purposes. All potential threat sources can simultaneously start a vulnerability exploit lifecycle.

2d. Monetisation through the grey or black market

The researcher communicates the vulnerability information to another entity than the vulnerability owner, generally to monetise it. This is likely to feed the vulnerability exploit lifecycle (see below).

2e. Non-disclosure

The researcher takes no action and the vulnerability remains unknown to the rest of the world, unless it is rediscovered. Research has shown that the probability of independent rediscovery within a year is from 15% to 20%, with important variations depending on the dataset. For example, 13.9% of vulnerabilities are rediscovered within 60 days, rising to 20% within 90 days, and above 21% within 120 days (Herr, Schneier and Morris, 2017[36]).

If the vulnerability owner receives vulnerability information, the disclosure lifecycle continues through what a code owner calls the "vulnerability handling process" and a system owner a "vulnerability management process" (further detailed in 1.1.5).

3. Investigation, also called verification, validation and triage.

The vulnerability owner analyses the reported information to confirm or refute the presence of the vulnerability in the product. If the vulnerability owner does not consider the information reported as a vulnerability, the process can stop there. The process can also stop if the product has reached end of support or end of life despite being still in use.[12]

The investigation phase should involve continued or at least regular communication and updates between the vulnerability owner and the researcher. When this does not happen, the disclosure process can derail and the researcher can decide to disclose the vulnerability to the public ("full disclosure"). He/she can also ask a co-ordinator to intervene (2b).

In case of a system vulnerability, the process moves to step 6. For a code vulnerability, it continues with step 4.

4. Development of the mitigation, also called resolution, or remediation.

The code owner decides whether it will develop a mitigation such as a patch modifying the code, and/or instructions and configuration or documentation changes that either remove or mitigate the vulnerability. This decision requires some risk-based prioritisation as not all vulnerabilities are equal and the code owner's resources as well as sometimes its digital security competences are limited.

The development of a mitigation can take place in several rounds, with the release of a temporary mitigation first and better solutions being provided later on. A third party co-ordinator and other trusted parties can be involved in the verification of the mitigation.

As explained above, not all code vulnerabilities can be fixed with a patch. When no patch is developed, the code owner can produce instructions or guidance for users to implement mitigation techniques and workarounds reducing likelihood of exploitation. In some cases, the code owner may regard the vulnerability as an accepted risk and not to remedy it, possibly in consultation with the researcher, ending the process (NCSC-NL, 2018[37]). If the communication with the researcher is not optimal, he/she may disagree with the code owner's conclusion that a patch cannot be developed or that the risk should be accepted, and fully disclose the vulnerability (2c), with or without agreement of the code owner, or seek to monetise it on the grey market (2d).

5. Release of the mitigation

The vulnerability owner releases the information about the code vulnerability and the mitigation to the public or to specific stakeholders. This enables threat actors to trigger an exploit lifecycle. Defenders can also start using this information to test their system and develop defensive exploits for network penetration testing.

6. Application of the mitigation

System owners apply mitigation measures to protect their systems and activities.

In case of a *code* vulnerability, they implement the newly released patch or follow the mitigation instructions from the code owner. If the patching process is automatic and transparent, as should be the case for consumer products, end users will not need to intervene in order to apply the mitigation. In professional environments, system owners are strongly encouraged to apply, where possible, a risk-based approach in deciding how quickly they should deploy mitigations when made available. The final decision should be business as opposed to only technically driven (FIRST, 2017[38]; UK NCSC, 2016[15]). When the code owner is the system owner, it takes care of steps 5 and 6 and applies the mitigation to its own product.

In case of a *system* vulnerability, the system owner changes the configuration settings or applies the patch or mitigation instructions that should have been applied earlier.

In both cases, the implementation of the mitigation either eliminates or reduces the risk related to the vulnerability to an acceptable level.

7. Post-release

Vulnerability owners take appropriate actions, taking into account users' feedback. For example, a mitigation may be incomplete or cause side effects that disrupt information systems. Users can report this information to the vulnerability owner and the latter can assist in resolving these issues, for example by providing additional information or modifying the mitigation.

There are many possible variations of Figure 3 lifecycle scenarios. For example, the researcher can publicly disclose some information about the vulnerability prior to reporting the vulnerability to the vulnerability owner, or in parallel with the reporting (limited or partial disclosure). This can happen, for example, when a known active exploitation is currently taking place. The researcher can also provide advance warning that he/she will disclose a vulnerability at a particular date (FIRST, 2017[38]).

Figure 3. Overview of the vulnerability disclosure lifecycle

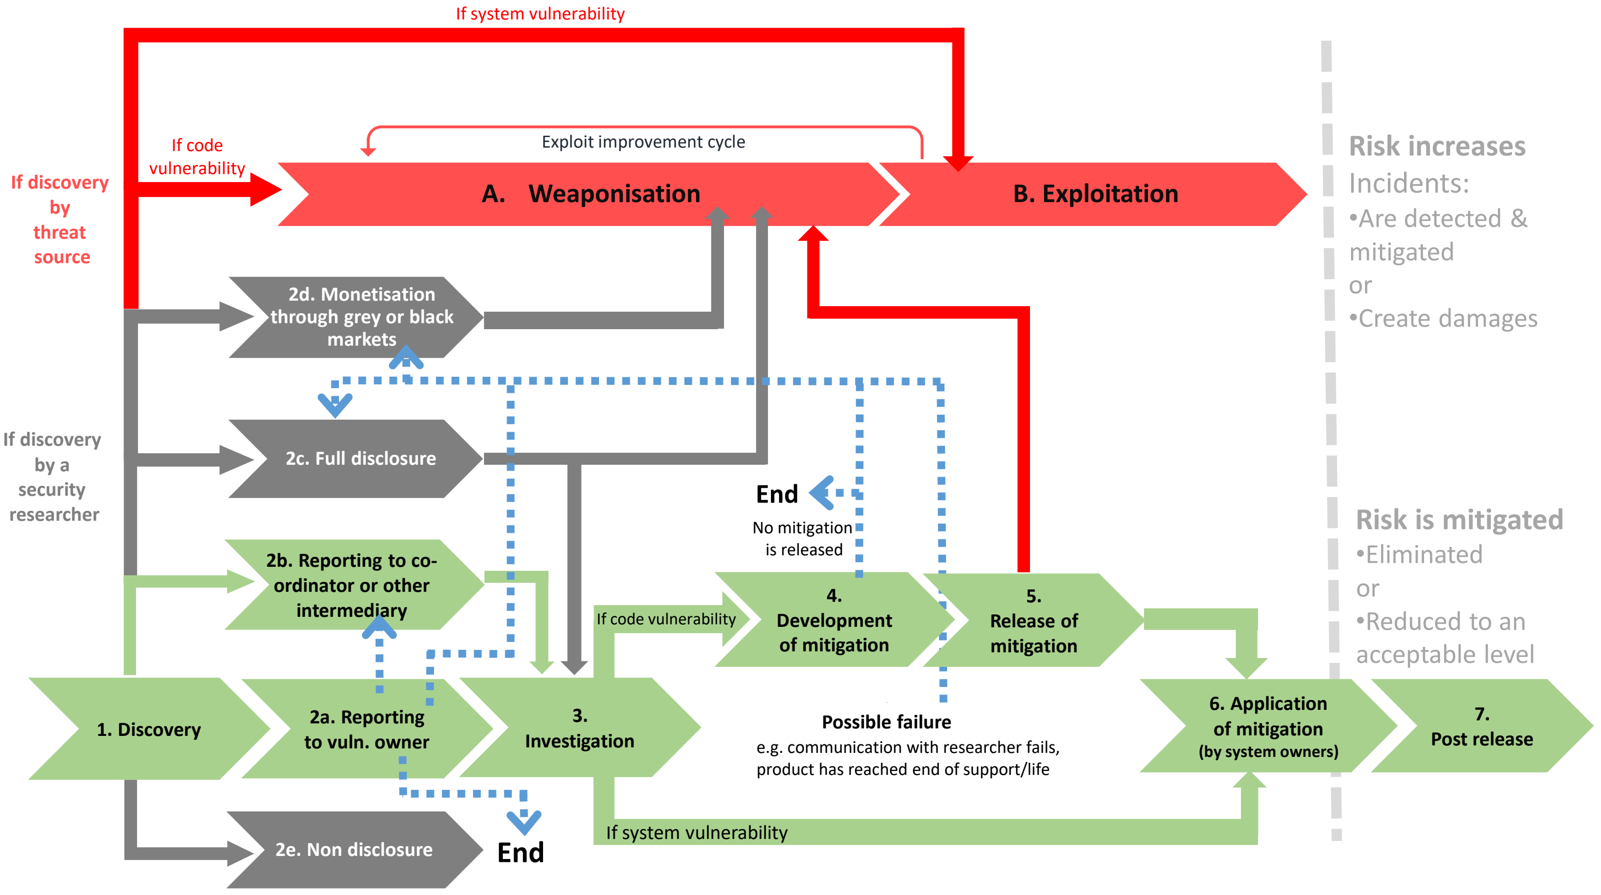
If system vulnerability
If code vulnerability
Exploit improvement cycle
If discovery by threat source
A. Weaponisation
B. Exploitation
Risk increases
Incidents:
•Are detected & mitigated
or
•Create damages
2d. Monetisation through grey or black markets
If discovery by a security researcher
2c. Full disclosure
End
No mitigation is released
2b. Reporting to co-ordinator or other intermediary
If code vulnerability
4. Development of mitigation
5. Release of mitigation
Risk is mitigated
•Eliminated
or
•Reduced to an acceptable level
1. Discovery
2a. Reporting to vuln. owner
3. Investigation
Possible failure
e.g. communication with researcher fails, product has reached end of support/life
6. Application of mitigation (by system owners)
7. Post release
2e. Non disclosure
End
If system vulnerability

Note: Events in green from 1 to 7 describe the vulnerability lifecycle. Blue dotted arrows = failure of the process. Events A and B represent the exploit lifecycle.
Source: OECD

1.1.5. Vulnerability handling and management

The previous section provides a description of the subject area from a vulnerability disclosure perspective, i.e. the flow of vulnerability information among stakeholders. In doing so, it touched on two steps that are essential to mitigate vulnerabilities rather than simply disclose them: vulnerability handling and vulnerability management.

The *vulnerability handling process* takes place on the product side. It is the code owner's responsibility. It covers how code owners (often called vendors) should process vulnerability information from the investigation to the post-release phases, regardless of whether the information comes from an external source or the code owner's internal security team (phases 3, 4, 5, and 7 in Figure 3). ISO/IEC 30111[13] is the reference standard for vulnerability handling (ISO/IEC, 2019[39]). Vulnerability handling is a sub-element of a broader product security development lifecycle (SDL), which includes other elements such as secure design, secure coding practices, testing and validation, etc (Safecode, 2018[40]; Edison Group, 2013[41]).[14]

On the *user's* side, the *vulnerability management process* is the system owners' responsibility. It is the set of ongoing processes enabling all organisations to know if (N-days) vulnerabilities are present within their IT estate and take appropriate risk management decisions and actions (UK NCSC, 2016[15]). It includes vulnerability scanning, patch testing and deployment (step 6 in the above vulnerability lifecycle description), and overall vulnerability management strategy maximising the resources provided. Patch deployment is an important step during which the risk posed by a particular vulnerability can be concretely eliminated. Vulnerability management is one of the security controls of ISO/IEC 27002 (Code of practice for information security controls). It fits within an organisation's digital security risk management cycle as called for in the OECD Recommendation on digital security risk management for economic and social prosperity (OECD, 2015[5]). A vulnerability management process starts with a discovery phase often based on vulnerability scanning, but also including vulnerability reports received from security researchers, network penetration tests, or other forms of security audits.

Vulnerability handling and management are related: patches and mitigations are the outputs from the vulnerability handling process and are feeding the vulnerability management process.

1.1.6. Vulnerability treatment

The description of vulnerability disclosure, handling and management shows that these concepts are all interrelated. Each concept represents the perspective of a different stakeholder group with respect to broader issue of how vulnerabilities are (or should be) addressed, using a different verb to distinguish them: security researchers "disclose" vulnerabilities, code owners "handle" them, and system owners "manage" them, i.e. vulnerability owners deal with vulnerabilities in a manner that reduces risk, whether the vulnerability has been reported to them by a researcher or discovered through an internal process. Considering vulnerability disclosure without taking vulnerability handling and management into account is not sufficient because these processes are interdependent.

The lack of an expression covering this area holistically encourages policy makers and other stakeholders to approach the issue from a narrow angle, such as vulnerability disclosure, handling, or management. However, if policy makers' objective is to reduce digital security risk, considering one of these aspects in isolation from the others is insufficient.

Therefore, this report proposes to use the term "vulnerability treatment" to overcome the current lack of a holistic expression. Vulnerability treatment includes four interrelated and overlapping processes: discovery, handling, management and co-ordinated disclosure. Co-ordinated disclosure is further introduced below (1.2.2 and chapter 2.). When discovery is an internal process of a vulnerability owner, external researchers are not involved and there is no need for co-ordination with them. In this case, it would be useless to refer to vulnerability treatment because it would become synonymous with vulnerability handling or vulnerability management, respectively for code and system owners.

Figure 4 represents vulnerability treatment.

Figure 4. Vulnerability treatment

Vulnerability treatment = Discovery + Handling + Management + Co-ordination of disclosure (CVD)

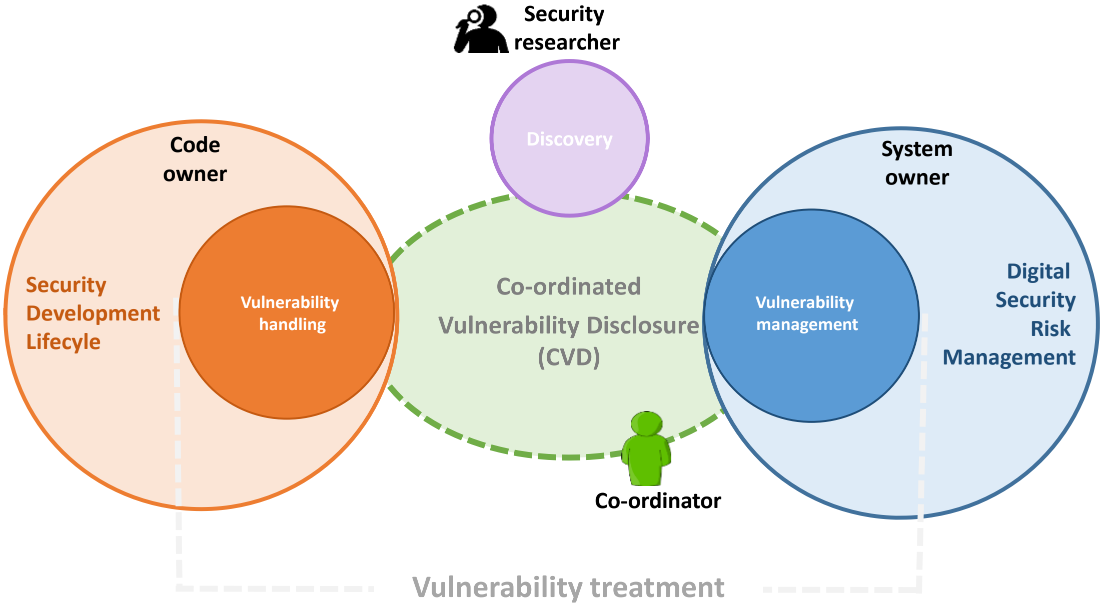

Source: OECD

1.2. Key challenges

Vulnerability treatment is a complex and dynamic challenge for which there is no objective "right" answers, only subjective "better" or "worse" solutions in a given context. It has the characteristics of a "wicked problem" (Rittel and Webber, 1973[42]), in that each case is essentially unique, there are no intrinsic criteria to indicate that a solution is sufficient, and pre-disposed ideology will influence stakeholders' judgment regarding a solution's fitness. It is also not possible to demonstrate that all possible solutions have been considered or even identified (CERT/CC, 2017[43]).

With this complexity in mind, this section introduces key challenges related to vulnerability disclosure from a public policy perspective.

1.2.1. Reducing the window of exposure

The ultimate objective of security professionals is to reduce users' *window of exposure* to vulnerabilities, which begins with the discovery of the vulnerability and ends with the application of the mitigation. Only the implementation of the mitigation decreases and potentially eliminates the risk, closing the window of exposure. Therefore, as soon as a vulnerability is discovered, the clock starts ticking. Time is paramount to optimise risk reduction because the level of risk does not stay the same within this window. The risk level usually increases over time as the probability of an exploit being developed grows, then decreases with the deployment of mitigations.

In the case of code vulnerabilities, two interrelated factors influence the level of risk.

First, the weaponisation of the code vulnerability by a threat source, where it is possible, and the availability of an exploit significantly increase the risk for all users. The exploit lifecycle can begin at any time after the discovery of the vulnerability. When a threat actor is the first discoverer of the vulnerability, the weaponisation is very likely to start immediately and rapidly lead to a zero-day exploit. Research by ENISA shows that "the exploit publication date of critical vulnerabilities is attracted near the vulnerability publication date, with the most exploits being published shortly before or after the vulnerability publication date" (2020[19]). According to research by the security firm Tenable based on a sample of the most prevalent vulnerabilities, an exploit was available on the same day that the vulnerability was disclosed in 34% of cases. The attacker had the first mover advantage in 76% of analysed vulnerabilities, meaning that they could attack without defenders being even aware of the risk. Attackers had a median 7-day window of opportunity to exploit a vulnerability before a defender is even aware they are vulnerable (2019[44]). Moreover, when the code owner or a researcher first discovers the vulnerability, a threat actor can rediscover and weaponise it before a mitigation is available, or use publicly available vulnerability information after its disclosure to weaponise it. Threat actors can also steal vulnerabilities hoarded by another actor.

Second, the public disclosure of code vulnerability information increases the risk of weaponisation because threat sources can use the vulnerability information to develop exploits, typically faster than code owners release a mitigation. However, public disclosure can incentivise an otherwise unmotivated code owner, or the security community at large, to begin or accelerate the development of a mitigation. If a mitigation is already available, public disclosure of vulnerability information is nevertheless essential for the security community to increase its knowledge (e.g. discovering related or similar vulnerabilities in other products, or improving attack detection tools) and for system owners to deploy mitigations. In fact, mitigation measures (e.g. security updates) can be, and usually are, reverse-engineered by threat actors to discover the underlying vulnerability they aim to correct. Eventually, the mitigated vulnerability will be known and potentially weaponised by threat actors.

In the case of system vulnerabilities, system owners are exposed as soon as the vulnerability information is available to anyone. The risk level increases with the number of individuals accessing the information. Other factors can further increase it. For example, when a well-known system owner consistently demonstrates poor vulnerability management practices, it can become a recurrent target for attackers, such as the case with the Sony Playstation Network in 2011, which was successfully attacked on more than 10 occasions over a relatively short period of time (Lee, 2011[45]).

This suggests that two factors should be tightly managed for optimal vulnerability disclosure:

- *Time:*
  - Vulnerability owners should create conditions to facilitate swift reporting by researchers to disclose vulnerabilities and expedite the vulnerability handling or management process (investigation, mitigation development and release);
  - Vulnerability owners and security researchers should co-operate to ensure rapid remediation;
  - System owners should apply mitigations as soon as possible;
  - Public disclosure of code vulnerability information should generally take place at the same time or after the mitigation's release.
- *Content:* sensitive code vulnerability information should not be provided to the public or its publication should be delayed. Vulnerability information is sensitive when it could facilitate weaponisation or when it relates to vulnerabilities that cannot be rapidly mitigated and are typically found in critical systems. Nevertheless, there are exceptions where earlier public disclosure may be warranted. They are often based on the recognition that attackers may already be aware of the vulnerability or can re-discover it and that leaving system owners without information would give attackers an advantage.

As shown in the URGENT/11 case explained in Box 3, the window of exposure varies depending on the parties concerned and can sometimes be very long.

Code owners who receive many valid vulnerability reports or have limited resources have to prioritise the development of mitigations according to each vulnerability's criticality and risk level. However, as noted above, a vulnerability's risk level depends on its use context by each system owner, therefore code and system owners may come to a different conclusion with respect to a given vulnerability's level of risk.

1.2.2. Encouraging co-ordinated vulnerability disclosure

For many years, the technical community has attempted to determine what a *responsible disclosure* mechanism should look like, namely how researchers should manage the dilemma of reporting the vulnerability to the vulnerability owner and making it public in the general interest. In this context, the term "responsible" refers to the need for vulnerability owners and researchers to "act responsibly, based on their role, ability to act and the context, and to take into account the potential impact of their decisions on others" (OECD, 2015[5]).

For example, there has been an ongoing discussion on the potential merits of full disclosure. Opponents of full disclosure claim that it creates the conditions for malicious actors to exploit the code vulnerability until a patch is available, thereby increasing the risk level for all users. Proponents argue that full disclosure aims to raise public awareness of the vulnerability so users and defenders can take self-protective action. It also creates an incentive for code owners to develop patches and increases the opportunity for public scrutiny (ENISA, 2016[35]). There are cases where full disclosure before the availability of a mitigation is more likely to take place, for better or worse, such as when the researcher is unable to locate a vulnerability owner's contact, when the vulnerability owner does not respond to the researcher, or when they no longer exist. In other cases, the researcher and the vulnerability owner may disagree that it is a vulnerability, or on its severity, and the researcher may want to put pressure on the vulnerability owner to take the vulnerability seriously and rapidly publish or implement a mitigation. The code owner may decide not to fix the vulnerability, for example because it no longer supports the product, or does not have the required skills and/or financial resources, or there may be compatibility issues affecting the fix. The researcher may be concerned by the legal implications of dealing with the vulnerability owner or believe that the latter is insensitive to users' security concerns (FIRST, 2017[38]). International standards such as ISO/IEC 30111 and 29147 recognise the merits of full disclosure in some cases. Ultimately, whether full disclosure is a responsible action depends upon the context, which is why the Cybersecurity Coalition, for example, does not recommend that policymakers seek new legal prohibitions on public disclosure of unpatched vulnerabilities (2019[46]).

Debates around "*responsible"* disclosure showed that definitions of who should be responsible for what in the public disclosure process were generally predicated on diverging value judgements and opinions, rather than objective and neutral reasoning (CERT/CC, 2017[43]). Furthermore, the term implied that if one type of disclosure was responsible, then all the others had to be irresponsible (451 Research and HackerOne, 2017[47]). The reality is more complex, and what seems unreasonable in some cases may be appropriate in others, and vice versa. Therefore, security researchers and civil society tend to discourage the use of the expression "responsible disclosure" to avoid unnecessarily aggravating the legal pressure and uncertainty on security researchers (cf. 2.2).

These debates resulted in the emergence of co-ordination as a shared principle to reduce risk related to vulnerability disclosure. For example, one-to-one co-ordination between the vulnerability owner and researcher, or through a trusted third-party co-ordinator may help avoid harmful full disclosure in many of the above examples. Co-ordination among vulnerability owners is also essential in complex multi-party situations such as those involving vulnerabilities in products' components. The technical community now recognises *co-ordinated vulnerability disclosure* (CVD) as the best approach to address vulnerabilities without increasing risk, whenever it is possible. CVD is further explained in Chapter 2.

1.2.3. Encouraging responsible handling and management of vulnerabilities

In order for co-ordinated vulnerability disclosure to be effective, vulnerability owners need first to take responsibility for addressing the vulnerabilities that can be found in their products or systems. Therefore, vulnerability handling and management are essential building blocks for vulnerability treatment, as explained above (1.1.5). There may be basic obstacles to taking such responsibility, such as the association of vulnerabilities to a failure of the organisation that can lead to negative reactions from the leadership, shareholders, markets, etc.

Figure 5 provides another representation of vulnerability treatment, showing where CVD stands between the code owners' vulnerability handling and system owners' vulnerability management processes. It also shows handling and management as part of broader security development lifecycle and digital security risk management frameworks. The top orange process is one possible representation of a product's security development lifecycle (Edison Group, 2013[41]). One of the steps addresses vulnerability handling. The blue box at the bottom represents an organisation's digital security risk management cycle. Vulnerability management is one of the many security measures or controls the organisation can take to reduce risk. The green box in the middle is the CVD process where a researcher reports to a vulnerability owner. This connects to the vulnerability handling process above or the vulnerability management process below. For CVD to work, the top and bottom building blocks have to be in place.

It is interesting to note that if the researcher discloses the vulnerability to the public rather than to the vulnerability owner, the same process will apply but in a crisis mode, during which the co-ordination of the communication may for instance include temporary workarounds.

Vulnerability handling and management are also learning and improvement cycles. They aim at improving the security of products and systems, as well as the organisation's overall security practices on a vulnerability-by-vulnerability basis, in order to develop products or manage systems with less vulnerabilities in the future.

Figure 5 shows that the effectiveness of a CVD process depends upon the capacity of a vulnerability owner to handle or manage vulnerabilities, as part of its broader security development lifecycle or digital security risk management framework. From a public policy perspective, this suggests that CVD is an important tool to reduce risk only for code owners that already have implemented a security development lifecycle and for system owners that already have a vulnerability management process. For example, an organisation which patches its systems occasionally and irregularly instead of systematically and cyclically should first improve its vulnerability management process prior to diverting resources and attention in encouraging security researchers, such as through a bug bounty programme (2.3.3). Nevertheless, it should always be able to receive and address a vulnerability report spontaneously sent by security researchers. The most basic method of receiving security reports is to have an email address at security@company.com, which should be monitored for reports from external sources.

To use a metaphor, vulnerability treatment can be viewed as a digital security engine receiving fuel (i.e. vulnerabilities) pumped by researchers and flowing through a CVD pipeline. The digital security engine processes the fuel through its vulnerability handling and management process, up to the secure development lifecycle or risk management framework, in order to improve the organisation's products or systems' digital security. If the organisation does not have such a digital security engine in the first place, it can be counterproductive to fuel it with vulnerabilities. It is more effective in that case to set up or improve the engine rather than diverting resources in a sophisticated CVD pipeline. However, if the engine is working well, then fuelling it through such a pipeline can significantly improve its performance.

In other words, policy makers should promote CVD as an effective additional mechanism for stakeholders whose vulnerability management or handling is already in place. They should not encourage those who have yet to manage vulnerabilities to rush into CVD, or to replace their processes with CVD.

Figure 5. CVD as part of the broader product and system security

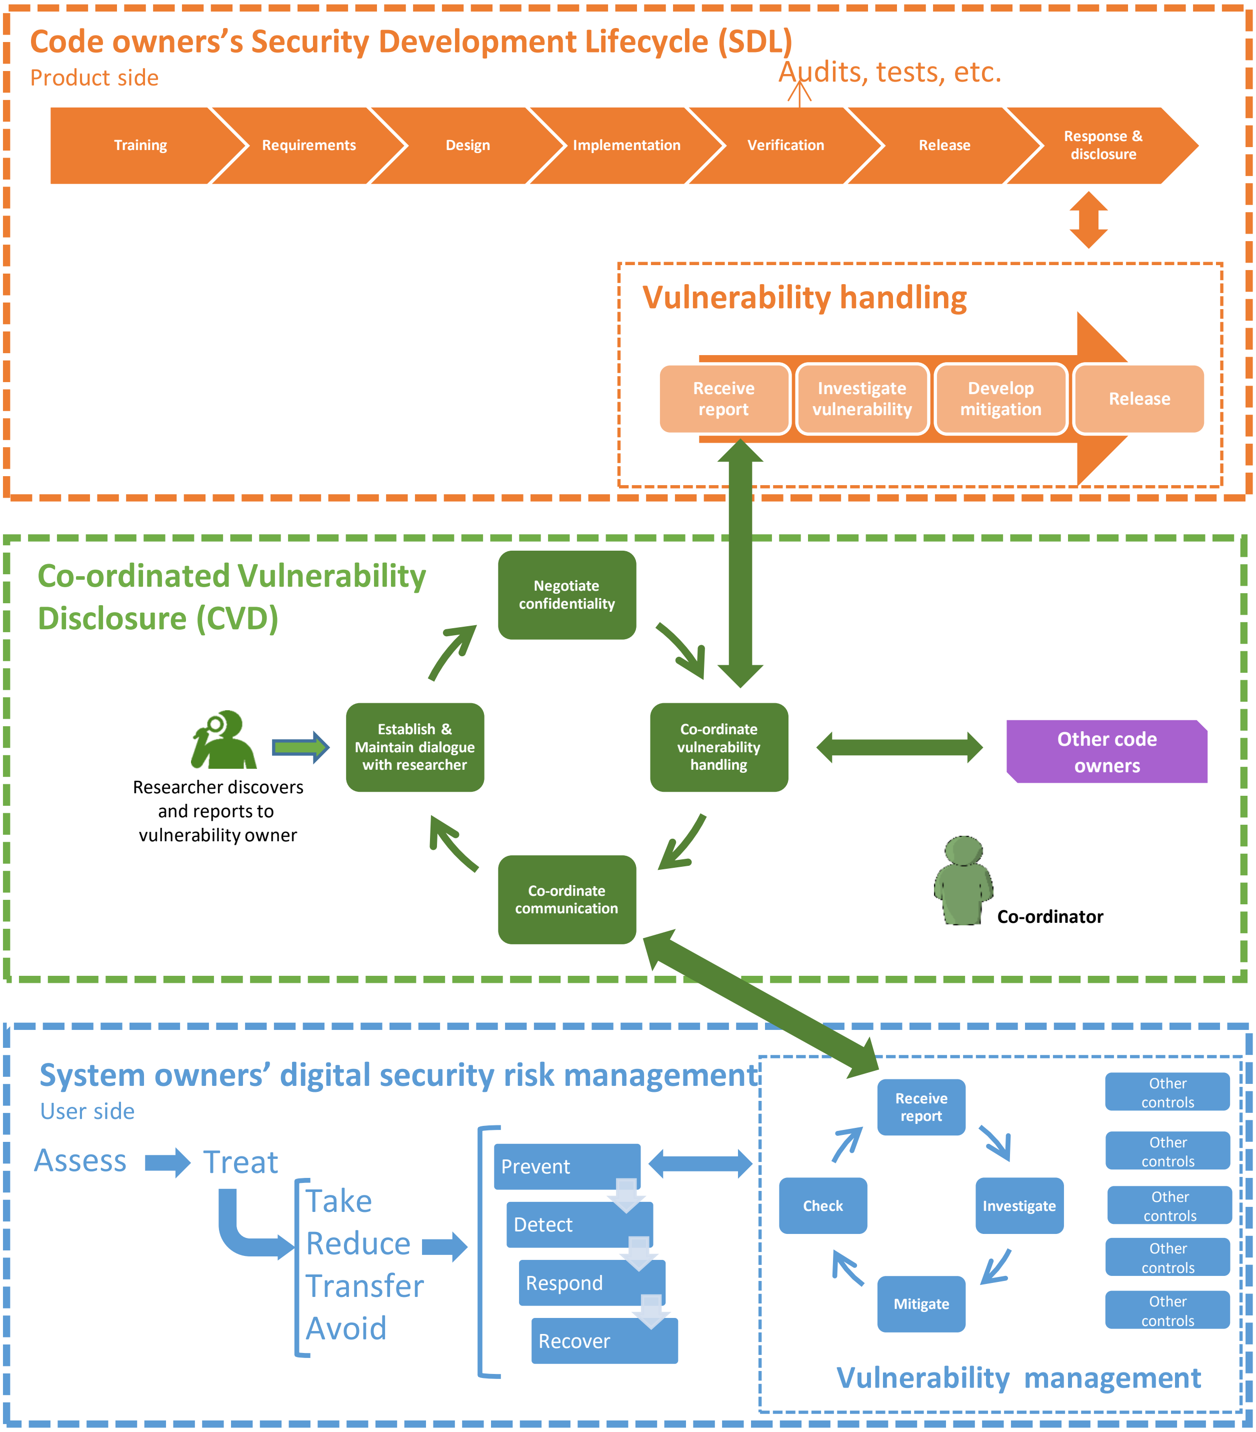

Source: OECD

1.2.4. Agreeing on what information to disclose publicly and when

Assuming that the researcher engages in a CVD process, the code owner and the researcher must co-operate to decide *i)* whether to broadly share vulnerability information as soon as possible, or wait and publicly disclose the information only when a patch is released; and *ii)* what information to disclose.

Sharing vulnerability information can empower security firms (defenders, cf. 1.1.3) to update their detection systems, and system owners to look for evidence of an exploit and take mitigation measures while waiting for a potential fix to reduce the risk. If the information is not shared, defenders and system owners cannot do anything to protect themselves in case an exploit is already being used. This can be particularly important when mitigation development is expected to take time, such as when many parties are involved in the disclosure process (cf. 2.1.2 below). However, sharing information can also empower threat actors to develop exploits. Most security researchers and vulnerability owners prefer to assume that at least one threat actor may have already discovered a newly discovered vulnerability and has been exploiting it in the wild for a certain time without being detected. In that case, waiting to share vulnerability information would publicly protract threat actors' advantage instead of abbreviating it. Even without a patch or fix, system owners can still apply temporary workarounds until patch deployment, if they consider that the risk is sufficiently high.

1.2.5. Managing the risk

Vulnerability owners often have to prioritise vulnerabilities upon which spending their efforts and resources. Prioritisation should be based on the risk associated with the vulnerability. Vulnerability handling and management are risk-based decisions making processes.

This is difficult to do for the code owner because it depends upon the product's use context, which varies considerably across users of the vulnerable product. Code owners also have to assess their own business risk with respect to the development of a patch, namely the potential positive and negative effects of publicly recognising the presence of vulnerabilities in their product, or the consequences of allocating resources to address the vulnerability. This suggests that vulnerability handling should be a business-led process, well integrated in the product business strategy, rather than only a technical matter. Public disclosure decisions are difficult to make because the severity of a vulnerability does not necessarily equal its risk.

System owners must also make risk-based decisions when code vulnerability information or patches are made public. They need to balance the risk of a successful attack exploiting the vulnerability with the risk of applying the mitigation to their system, as explained above (cf. System vulnerabilities in 1.1.1). Time is of the essence in this process. The decisions to apply a patch may be easy to make when it fixes the most severe vulnerabilities. For example, a patch fixing the so-called EternalBlue vulnerability that affected almost all Microsoft operating systems when it was released and "allowed remote attackers to execute arbitrary code via crafted network packets" (CVE-2017-0144) is to be applied immediately to avoid multi-million disasters such as Wannacry and NotPetya. Organisations may for example adopt a policy requiring any vulnerability scoring the highest levels of criticality to be immediately mitigated upon patch or workaround availability. However, most other vulnerabilities require testing to assess the risk in many cases, a time-consuming process, which increases the window of exposure. Ultimately, vulnerability management is a business as much as technical risk management process since it can affect the organisation's economic and social performance.

Code and vulnerability owners may have different interests and perceive risk differently. In the case of a code vulnerability, a product user (i.e. system owner) faces a direct security risk which is reduced when mitigations are implemented to protect its information systems. In most cases, a code owner does not directly face a digital security risk but rather a business risk related to the possible commercial repercussions of its vulnerability handling decisions on its product reputation and position on the market.

If the code owner has some critical customers, it could be in its business interest to undertake proactive measures to ensure that they are protected. The government also has a different perspective. For example, it may aim to ensure that the operators of critical activities have all deployed the mitigation on all systems integrating the vulnerable product.

Lastly, security researchers also need to make risk-based decisions when they investigate a product or a system to find vulnerabilities, and when they disclosure their findings. In particular, they need to understand the legal risk they are facing and adjust their testing techniques accordingly (cf. 2.2).

1.2.6. Co-ordinating stakeholders

Co-ordination between the vulnerability owner and security researcher is at the core of successful vulnerability treatment. A dialogue between them is essential to reduce misunderstandings and facilitate the process.

In many cases however, the party receiving the vulnerability report from a researcher is not the only one with a role to address the vulnerability. Co-ordination can sometimes involve many other stakeholders taking part in the product value chain. For example, a security researcher can report a vulnerability to the entity owning the product's brand or to the party that commercialises the product (i.e. the vendor). However, these entities may not really own the responsibility to address the vulnerable layer of code. The code owner may be located several steps away from the consumer down the value chain, making the distribution of a mitigation an uncertain and complex endeavour.

Therefore, co-ordination is not limited to the relationship between the researcher and vulnerability owner. It can extend to co-ordination among supply-side actors. The disclosure of the "Spectre" vulnerability in 2018, which affected microprocessors, highlighted many limitations to the ideal process described in Figure 3. In that case, processor manufacturers whose products were found vulnerable knew which corporate clients (e.g. computer manufacturers) bought processor chips or boards from them, but had no information on who end users were. Therefore, they could not alert them about remediation. Furthermore, some mitigations involved changes to products created down-stream from the vulnerability owners, such as patching operating systems (OS) or compilers to avoid certain instruction sequences. In such complex cases, co-ordination is significantly more challenging than described above. [15]

In many cases, a third-party co-ordinator can facilitate the process. Co-ordinators are described in 2.3.5.

1.2.7. Making defence more attractive than offence

Vulnerability markets are global as many transactions take place across borders. Three types of markets for vulnerabilities can be distinguished (Fidler, n.d.[48]; ENISA, 2018[49]):

- White - i.e. regulated – markets, which connect vulnerability researchers and vulnerability owners. They include**:**
  - *Co-ordinated disclosure markets:* the vulnerability is disclosed to the public after a co-operative process between the researcher and vulnerability owner, with or without the participation of a co-ordinator;
  - *Vulnerability reward markets:* the researcher reports the vulnerability to the vulnerability owner in exchange of a reward, either directly to the vulnerability owner through a bug bounty programme, or to a trusted-third party through a bug bounty platform (see 2.3.4);
  - *Captive markets:* where a vulnerability is not disclosed to the public after its communication to the vulnerability owner. Another case is when the vulnerability remains within the researcher's host organisation such as a government defence or intelligence agency.
- Grey – i.e. partially regulated – markets: vulnerability brokers connect sellers with buyers who are not the vulnerability owners and whose objective is not to fix the vulnerability. They include

government intelligence and defence agencies as well as companies developing and selling tools based on the exploitation of vulnerabilities, such as devices purchased by police forces or intelligence agencies to access the content of mobile phones. These markets provide a means for buyers to bypass the vulnerability discovery phase and rapidly develop zero-day-based tools and exploits. Some brokers launch bug bounties with pre-defined payouts for vulnerabilities in specific products to attract researchers. When they take place across borders, which is often the case, such transactions can violate domestic legislations implementing the Wassenaar Arrangement.

- Black – i.e. illegal – markets: buyers and sellers trade vulnerabilities, generally on underground platforms in the dark web, through online chat rooms, or specialised marketplaces.

With the rise of bug bounty programmes and platforms (further described in 2.3.3 and 2.3.4), the white market has considerably increased over the last few years, apart from captive markets about which there is little information.

The grey market has been pointed out as illegitimate, despite being legal in some countries, because it contributes to increasing the overall level of digital security risk globally, and to the surveillance of populations including human rights' activists, issues which are beyond the scope of this report (Fidler[48]; Lee, 2019[50]).

Nevertheless, the grey market also has a negative influence on vulnerability disclosure because it can divert researchers away from co-ordinated disclosure as buyers can outbid vulnerability owners to acquire critical zero-day vulnerability information. This is particularly the case when buyers are well-resourced intelligence, defence and law-enforcement agencies. Several experts consider that the development of the grey market makes offense lawfully pay better than defence (OECD, 2019[1]; Manion, 2014[51]).

This influence is difficult to measure as transactions on the grey market are generally kept confidential. However, advertised payout prices provide some indications on the importance of the problem. For example, the US-based vulnerability reseller Zerodium indicates that payouts for one zero-day vulnerability can reach USD 2.5 M for mobile and up to USD 1 M for desktops and servers' operating systems, depending on the popularity of the affected products, the criticality of the vulnerability and the quality of the exploit (Zerodium, n.d.[52]). The company also targets products such as WhatsApp, iMessage and SMS/MMS applications with rewards up to USD 1 M. It claims it can pay even higher rewards for exceptional exploits and research (Franceschi-Bicchierai, 2019[53]). In 2018, the UAE-based company Crowdfense advertised a USD 10 M fund to buy zero-day vulnerabilities that increased to 15 M in 2019, with payouts ranging from USD 100k to 3 M (Crowdfense, n.d.[54]). As mobile and operating systems' designers have significantly stepped up their products' security, both companies are now also targeting Internet routers with payouts up to USD 100k for remote execution exploits (Cox, 2019[55]).

In an attempt to measure the size of the market, researchers have concluded that while prices of high-end vulnerabilities may look high at first sight, the entire market is very small in comparison with the cost of these vulnerabilities being exploited in the wild. They suggest that if the industry would internalise the cost of a programme to buy all vulnerabilities available, the total cost would be less than the commonly accepted rate of "pilferage" in other industries (Box 2).

Some code owners have started to increase their bug bounty rewards to compete with these companies on the vulnerability market, an effort that not all vulnerability owners are able to pursue. For example, in December 2019, Apple and Google increased payouts up to USD 1 M for extremely critical vulnerabilities, plus a 50% bonus if the vulnerability affects a product available in public beta version (Apple, 2019[56]; Lin, 2019[57]). The same month, Mozilla doubled all its payouts, and tripled payouts for remote code execution vulnerabilities found on a list of its critical web sites (Bennetts, 2019[58]).

Box 2. How much would it cost to buy most vulnerabilities?

According to recent research (Frei and Rochfort, 2021[59]), the analysis of CVE data from the US National Vulnerability Database (NVD) provides useful economic insights on the vulnerability grey and black markets challenges.

First, the data shows that over the last ten years, a few vendors have owned the responsibility for the majority of vulnerabilities disclosed per year: only 50 vendors accounted for about 50%, and 500 vendors for at least 72% of the vulnerabilities disclosed each year.

Second, it shows that the scale of the challenge might be lower than it seems, if approached from an economic angle. The analysis, carried out with 2010 to 2020 NVD data, explores how much it would cost to buy these top vendors' vulnerabilities at a price depending on their CVSS score, such as USD 250 K, 150 K, and 50 K respectively for critical, high and medium severity vulnerabilities. The idea of vendors buying all their products' vulnerabilities at an arbitrary price is rather unrealistic and unfeasible. However, as a research hypothesis, it can be useful to put the grey and black market challenges into perspective. The findings are as follows:

Buying all vulnerabilities from the top-50 vendors, accounting for 57% of all vulnerabilities, would cost approximately USD 1.165 billion in 2020.

Buying all vulnerabilities from the top-500 vendors, accounting for 81% of all vulnerabilities, would cost USD 1.732 billion. This represents 0.003% of the cumulated GDP of OECD Members (or 0.011% of the cumulated GDP of EU Members, or 0.008% the US GDP). USD 1.732 billion would also represent less than 0.5% of global cybercrime losses, assuming the total losses amount to USD 1 000 billion (estimates of the global cost of cybercrime, which are always a matter of debate, range from USD 100 to 6 000 billion).

The cost for the top 11 publicly listed vendors, in number of known vulnerabilities, to purchase all vulnerabilities in their products at USD 250 K, 150 K and 50 K per unit, respectively for critical, high and medium severity vulnerabilities, would account on average for less than 0.5% of the vendor's yearly revenues. In the United States' retail sector, the accepted rate of "pilferage" or "inventory shrinkage" (considered a cost of doing business) is between 1.5% and 2.0% of annual sales.

According to the authors of this research, these findings suggest that the majority of vendors with the highest numbers of vulnerabilities, which are highly profitable organisations, are dumping the cost of the security defects in their products on society while pocketing the profits (liability dumping). The authors stress that there is considerable room for these vendors to take the responsibility for digital and invest more into the security of their products without a risk to their business.

Source: (Frei and Rochfort, 2021[59]).

1.2.8. Trust in the government

Vulnerabilities reported to the government

Some governments can receive vulnerability reports from security researchers, as in the case of many national CERTs operated by governments that have a vulnerability co-ordination function. In these cases, security researchers need to have a high level of certainty that the government agency will not use the vulnerability information to develop exploits or communicate it to another government entity that could weaponise it. Strict and transparent separation between the government entity receiving reports and agencies with an offensive role is an important condition for trust in this area. To ensure stakeholders' trust, some experts have suggested that government agencies receiving vulnerability information for co-ordination could be established as independent bodies akin to data protection authorities.

In some cases, code vulnerabilities can carry a high-level of risk for the economy and society, for example when their exploitation could affect safety, national security, a large share of GDP, or a very large number of people, etc.[16] Vulnerability owners and/or researchers should consider providing advance information about such vulnerabilities prior to public disclosure. Informing the government can help ensure that affected vulnerability owners take appropriate measures prior to public disclosure, thereby reducing the likelihood of exploitation.

Nevertheless, vulnerability owners and researchers are likely to provide advance information only if they trust that the government will not misuse it, for example by exploiting it for offensive purposes, and that they will deal with it in a secure manner, avoiding leaks or communications to inappropriate third parties. According to experts, transparency can help establish trust. For example, a party who communicates such vulnerability information to the government can also inform the public that it is doing so without providing details about the vulnerability itself, and the government could acknowledge receipt and commit to certain communications milestones.

Nevertheless, in some countries, it can be dangerous for a security researcher to discover a sensitive and potentially embarrassing vulnerability as the government can use legal threat to discourage public disclosure or discredit the researcher (cf. the case of J. Sorianello in Argentina in Box 4).

Advance communication of vulnerability information by a vulnerability owner to government can also become a delicate challenge if the vulnerability owner operates or has users in several countries. Should the vulnerability owner provide the vulnerability information to its own government, to the governments that it trusts and whose population could be affected by the vulnerability, or to all governments with potentially affected populations, including those it does not trust?

In the first and second cases, other countries may accuse the vulnerability owner of putting their population in danger. In the third case, some governments in the list could weaponise the vulnerability, or behave inappropriately, e.g. by immediately disclosing the information publicly. Currently, these type of issues tend to be resolved through international CERT co-operation, even if their informal working methods can raise difficult challenges, for example when CERTs are not independent from agencies with an offensive capacity.

Governments are not necessarily trusted by all stakeholders across borders. For example, they may not understand or respect the need to keep information confidential, or who can pass confidential information to untrusted participants, inside or outside the government. This suggests that governments need to have a clear and transparent process about where and how they receive vulnerability information, and what they do with that information once received. To address this issue, some experts are calling for the establishment of an international co-ordinator operating as a non-profit and without links to a government.

Vulnerabilities discovered by the Government

Governments can also discover zero-day vulnerabilities through their own research and other efforts. They will have to decide what to do with such vulnerabilities. If the mandate of the agency that discovered the vulnerability is limited to protection, the agency will have to immediately report it to the vulnerability owner and trigger a CVD process in order to decrease the risk faced by users, including the government itself, operators of critical activities, and other stakeholders.

However, governments have a special role with respect to vulnerabilities, as explained in Box 1. If the government agency’s mandate also includes offense, or if the agency can legally share the vulnerability with other agencies in charge of offensive operations, intelligence or law enforcement, it will have to decide whether to trigger a CVD process, or delay the report to the vulnerability owner in order to enable the exploitation of the vulnerability for offensive activities. The agency could also stockpile it for the same purposes, a behaviour that has raised significant concerns in the digital security community considering the risk that stockpiled vulnerabilities be stolen or leaked and turned against legitimate stakeholders.

Charlet, Romanosky and Thomson argued that this decision making process is extremely important. They recommend that every nation should openly acknowledge that decisions regarding retaining or releasing zero-day vulnerabilities are not taken lightly, and that such decisions weigh both the national security gains of keeping the vulnerability secret and the digital security benefit of reporting it to the vulnerability owner and subsequently disclosing it the public (2017[60]). They called for governments to adopt a process to ensure that decisions about all zero-day vulnerabilities involve participants who represent commercial, critical infrastructure, and public digital security interests, and are informed by a range of viewpoints. The United States and the United Kingdom have both published a charter describing their “Vulnerabilities Equities Process” (VEP) (US White House, 2017[61]; UK GCHQ, 2018[62]) and the German government is working on developing and publishing one (Herpig and Schwartz, 2019[63]). The term “equity” refers to these governments’ recognition of the need to assess fairly risks and benefits to both the intelligence requirements and the digital security of the country. The Centre for European Policy Studies (2018[64]) also called for the generalisation of such policies which, as an element of national security governance, are beyond the scope of this report.

1.2.9. Other challenges

Vulnerabilities affecting critical activities

Some vulnerabilities may affect critical activities, such as when the product is typically used in sensitive industrial, medical or defence environments (e.g. industrial control systems) or when its customer base is so widespread that incidents leveraging the vulnerability would create systemic damages for one, several or all nations (e.g. microprocessors, operating systems, etc.).

In such cases, it may become necessary to inform governments about the vulnerability. However, vulnerability owners or co-ordinators who become aware of such vulnerabilities need to decide which governments to inform. Vulnerability owners such as large multinational corporations may be reluctant to share the information only with some governments, as their customer base is global. However, they may also not trust some other governments who they believe could weaponise the vulnerability and contribute to harm some of their customers and damage their product’s reputation. In some cases, this could lead to these vulnerabilities not being shared with any government.

2. Co-ordinated Vulnerability Disclosure

For at least two decades, the digital security community has explored how to optimise the vulnerability lifecycle in the best interest of all stakeholders. Over time, the need for increased co-operation between stakeholders appeared more and more essential. CVD is emerging as a best practice to incentivise the digital security community to work together despite potential differences of views between stakeholders. This chapter is a deep-dive into CVD. After introducing the concept (2.1), it discusses legal risk for researchers, which is a key obstacle to more widespread adoption (2.2), as well as tools to facilitate CVD, including standards, vulnerability disclosure policies, bug bounty programmes and platforms, and co-ordinators (2.3). This chapter also includes good practice (2.4) gathered from various guidance documents listed in Annex 1.

2.1. What is CVD?

2.1.1. Overview

CVD is a process through which vulnerability owners and researchers work co-operatively in finding solutions that reduce the risk associated with a vulnerability. The primary objective of CVD is to ensure that the vulnerability information is publicly disclosed only after mitigations are available to end users in order to reduce their window of exposure and the related risk (NCSC-NL, 2018[37]). CVD is widely recognised as a good practice to ensure that researchers and vulnerability owners act in a responsible manner for vulnerability disclosure. It is an overarching co-ordination framework involving several stakeholders (one or more security researcher(s) and vulnerability owner(s)) rather than an operational process specific to a single stakeholder. A CVD process can, but does not necessarily, involve a third-party co-ordinator (2.3.5).

The premise of CVD is that all vulnerability owners and researchers co-operate to:

- Rapidly develop and distribute mitigations;
- Decrease the risk that vulnerability information becomes public before mitigations are available.

Delaying the publication of vulnerability information enables the vulnerability owner to evaluate the issue and handle or manage respectively the code or system vulnerability. In the absence of a mitigation from the code vulnerability owner, CVD can also provide users with sufficient information to evaluate risk from vulnerabilities in their systems and help to reduce them.

One condition for co-operation to take place is that vulnerability owners commit not to pursue legal action against security researchers who participate in the process in good faith and respect the pre-defined rules (2.2).

In contrast with full disclosure, i.e. the release of the vulnerability information to the public without co-ordination (cf. 1.1.4), CVD is not a single event but rather a process, i.e. a series of events involving relationships and information exchanges between stakeholders, as well as decisions and actions (CERT/CC, 2017[43]). As explained above, CVD encompasses a process related to the relationships between the researcher and the vulnerability owner (addressed in ISO/IEC 29147), and a process related to the code owner's internal vulnerability handling process (addressed in ISO/IEC 30111) or the system owner vulnerability management. In Figure 3, the CVD process covers all the steps in green, from 1 to 7. Figure 4 and Figure 5 position CVD in relation to vulnerability handling and vulnerability management.

Standards and guides provide detailed descriptions and guidance, taking into account different cases and introducing additional stakeholders and steps.

CVD is effective both for code and system vulnerabilities, and therefore number of potential stakeholders that can potentially use it to improve digital security is extremely large. First, the app economy has led to a significant increase in the size and heterogeneity of the community of coders. Second, because all organisations maintaining an information system can potentially benefit from CVD, virtually all businesses, government agencies and non-profit, regardless of their size and mission can use it. For example, in the Netherlands, a vast number of organisations ranging from banks, to ISPs, DIY stores, and supermarkets have established CVD programmes as a result of NCSC-NL's promotion of the good practice.

2.1.2. Multi-party CVD

Most modern software includes pre-existing third-party components, modules, and libraries from the open source and commercial software worlds. The complexity of the CVD process increases when many code owners are involved in a product's value chain, for example when the vulnerability affects a product included as a component in one or many other products. For example, it can be particularly challenging to understand which parties can be affected by a code vulnerability in a product typically used as a component in many other products (cf. 1.2.6).

In certain technical environments, such as hardware, or with certain types of vulnerabilities, such as in widely-used protocols, CVD may entail a broader collaboration within the ecosystem to validate the vulnerability, develop and test mitigations and finally deliver and make them available to end users (Center for Cybersecurity Policy and Law, 2019[65]).

In addition to the Spectre vulnerability discussed above, Box 3 shows another case where multi-party co-ordination would have been necessary. This case suggests that:

- Code owners should notify downstream vendors of the vulnerability, although this can be an issue for open source software;
- All parties should organise public disclosure of vulnerability information within a reasonable timeframe so that most vendors can prepare the mitigation, taking the time of deployment into account;
- The time required to work on a mitigation and pass the validation tests prior to a release, and for users to plan a deployment, are very different between the software, hardware and IoT industries.

FIRST vulnerability co-ordination Special Interest Group has been working on guidelines and practices for multi-party vulnerability co-ordination and disclosure (2017[38]). They distinguish upstream (i.e. component) from downstream vendors. For example, suboptimal communication between upstream and downstream vendors can block the vulnerability disclosure process. Considering the importance of third-party components in digital products, there is a need to encourage the development of a co-ordination framework derived from ISO/IEC 29147 which would address multi-party vulnerabilities, and consider the involvement of national or sectoral co-ordinators to address cases where critical infrastructure security is at stake.

According to international standards, the affected code owner, often called "vendor", should create processes to support vulnerability reporting as well as manage and lead the co-ordination effort. Should multiple vendors be concerned by the vulnerability, researchers are encouraged to report vulnerability information to the potentially affected vendor who is best positioned to lead the co-ordination efforts to validate the vulnerability, develop and deliver mitigations to users (UK DCMS, 2018[66]; ISO/IEC, 2018[67]; CSDE, 2019[68]).

Box 3. URGENT/11 – A case showing that multi-party co-ordination is key when critical activities are involved

In July 2019, a security firm announced the discovery of URGENT/11, a set of 11 vulnerabilities, including 6 with critical severity, in the IP stack of VxWorks, Wind River Systems' Real-Time Operating System (RTOS). VxWorks is an operating system used by over 2 billion devices deployed in critical industrial, medical, automotive, and enterprise environments. Impacted products include industrial controllers, SCADA devices, patient monitors and other healthcare devices such as infusion pumps, as well as firewalls, VOIP phones, printers, etc.

Wind River Systems co-ordinated the vulnerability disclosure with the security researchers to ensure that a security fix was available in time. However, as of March 2020, many downstream vendors were still working on implementing the updated version of VxWorks in their own products and releasing updates to their own users.

However, the story does not stop there. Upon further investigation, it appeared that the vulnerable code in VxWorks RTOS had initially been developed by Interpeak, a Swedish company acquired by Wind River Systems in 2006, and had been licensed by Interpeak to many customers including numerous other RTOS developers. After its acquisition by Wind River Systems, Interpeak was dissolved and support to these licensees was terminated. As Interpeak's code remained in its customers' products, the list of products potentially affected by URGENT/11 is much larger than products embedding VxWorks. They include platforms such as ENEA's OSE, INTEGRITY by Green Hills, ITRON, Microsoft's ThreadX, Mentor's Nucleus RTOS, and zebOS. Mitigations provided by VxWorks are unlikely to work for them.

In a legislative proposal, the US Food and Drug Administration (FDA) has advocated that manufacturers adopt a "software bill of materials" (SBOM) outlining which stacks, libraries, and open source components are in devices to allow for tracking vulnerabilities such as URGENT/11 across all sorts of devices. The development of SBOMs has also been analysed by the US NTIA (Newman, 2019[69]; NTIA, 2019[70]).

This illustrates how complex value chains can significantly extend the window of exposure.

In such case, a multi-party co-ordination would have been helpful to arrange for a more optimal disclosure timeframe from two perspectives:

- Notify downstream vendors to let them prepare their security fix, considering that this will be possible for commercial off-the-shelf (COTS) but not for open source software as these downstream vendors are not identified by the upstream vendor in most cases;
- Notify critical users in advance to let them investigate and define treatment plans.

When there is no clear vendor positioned to lead the co-ordination, a co-ordinator can assist in setting up a broad collaboration within the concerned ecosystem. Information Sharing and Analysis Centres (ISACs) can provide the venue for stakeholders to reach out to their peers, including competitors (CSDE, 2019[68]).

The resources and time required to lead co-ordination with dozens or even hundreds of other vendors may be very high. Small code owners, including from the open source world, should carefully assess whether they have the capacity to lead such a process prior to accepting the challenge. For example, OpenSSL only had two people to write, maintain, test, and review 500 000 lines of business critical code when the infamous Heartbleed vulnerability was discovered (Walsh, 2014[71]). Some experts have suggested the establishment of a well-resourced, international and not-for-profit vulnerability co-ordinator to address this issue as well as some of the trust-related co-ordination challenges highlighted above (cf. 1.2.6).

In general, the co-ordination process first focuses on the development of a mitigation, and only stakeholders who can facilitate the technical development of the mitigation need to be involved. Then once the tested mitigation is ready, more stakeholders can be join the process to facilitate the mitigation's distribution and deployment.

Figure 6. Stakeholders' roles and communication paths in multi-party CVD (FIRST, 2017)

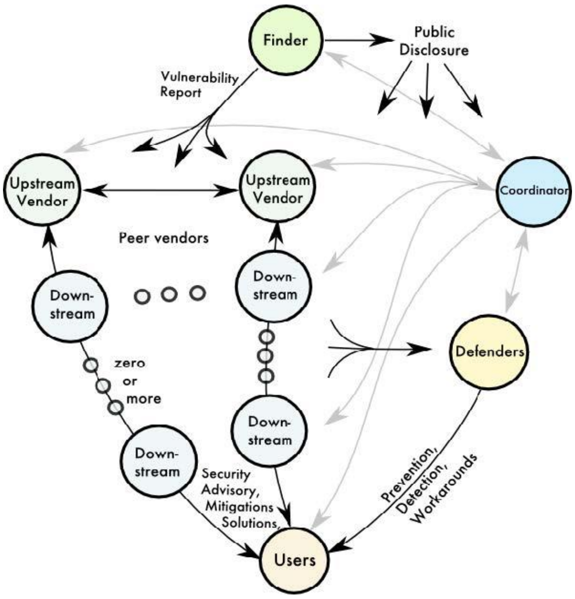

Note: Each guide and standard defines stakeholders slightly differently. In this diagram developed by FIRST, the researcher is called finder. Upstream vendors provide a product to a downstream vendor. Zero or more downstream vendors receive a product from an upstream vendor for use in the downstream vendors' product. Defenders are third parties who are responsible for defending against attacks, such as a system administrator, vendor, provider of defensive technologies or services. They may detect vulnerable systems, detect and respond to attacks, etc.
Source: (FIRST, 2017[38])

Private sector partnerships initiatives can also facilitate co-ordination in the context of multi-party vulnerability disclosure. For example, the Industry Consortium for Advancement of Security on the Internet (ICASI) in the United States brings together vendors in a trusted forum, such as Cisco, Intel, Juniper, Microsoft, Amazon and Oracle, where they developed a Unified Security Incident Response Plan (USIRP) to harmonise their internal security incident response procedure. ICASI members can trigger a USIRP event such as the receipt of a vulnerability report, share information about it and work together on a co-ordinated response. ICASI also developed the Common Vulnerability Reporting framework in 2012, a standard that enables different stakeholders across different organisations to speed up critical vulnerability-related information exchange and digestion by sharing it in a single format. In 2016, this framework was transferred to OASIS, a non-profit open standards organisation (OASIS, n.d.[72]).

While FIRST and other groups have clarified multi-party CVD, further work is needed, including regarding how and when to decide which stakeholders to involve in the process, how far down the supply chain the co-ordination process should extend to, as well as if and which government agencies should be involved (ENISA, 2018[49]).

2.1.3. Conditions for CVD

This section provides an overview of some interrelated conditions for CVD. The success or failure of CVD is largely determined by the complementary, competing or conflicting incentives that influence the behaviour of organisations and individuals, as pointed out in a report on the "economics of vulnerability disclosure" by ENISA (2018[49]). According to the authors, the net effect of incentives and barriers often leads to pervasive sub-optimal outcomes in this area. For example, different categories of researchers are driven by different goals and operate under different constraints. However, they note that it is possible to change these outcomes through policy mechanisms that can influence the behaviour of vulnerability disclosure stakeholders, such as legislation and regulation. Figure 7 represents the factors that influence stakeholders' decisions in the vulnerability disclosure lifecycle, i.e. incentives and barriers.

The ENISA report also underlines that externalities can be at play in some cases where the costs incurred by the exploitation of vulnerabilities are not borne by the vulnerability owner. This may explain why some code owners, for example, are reluctant to patch vulnerabilities reported to them and even to enter into CVD. Other possible causes for slow or absence of mitigation development by code owners include liability dumping or shifting between different stakeholders across the supply chains.

Figure 7. Economic incentives, motivations and barriers in a co-ordinated vulnerability disclosure process (ENISA, 2018)

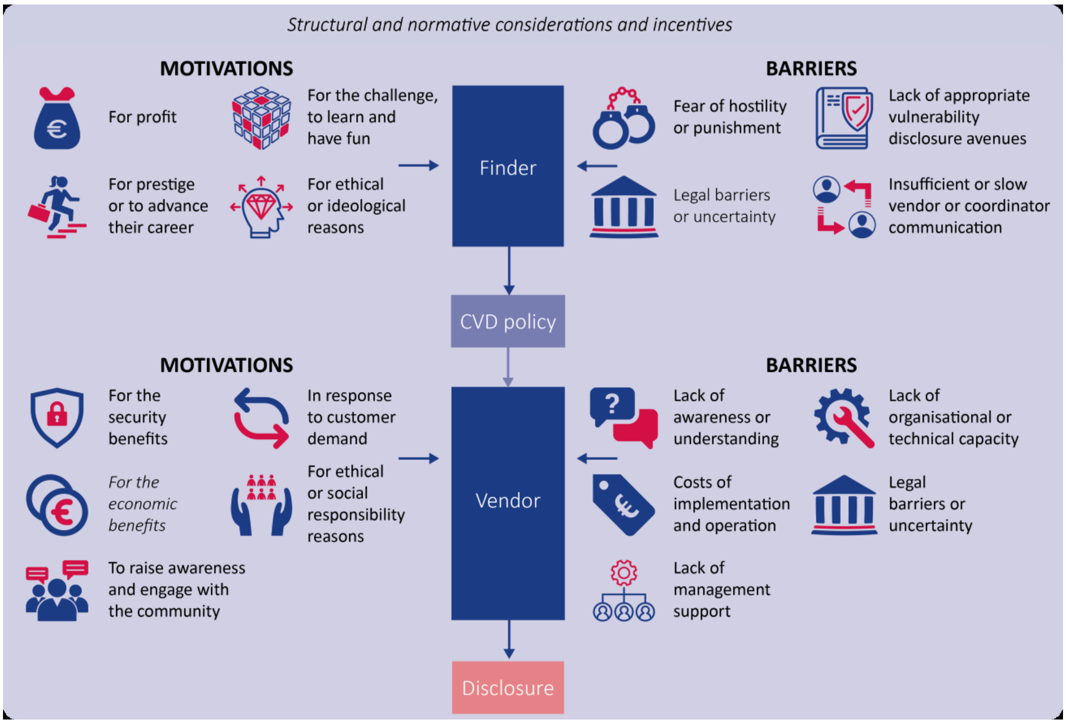
Structural and normative considerations and incentives
MOTIVATIONS
For profit
For the challenge, to learn and have fun
For prestige or to advance their career
For ethical or ideological reasons
Finder
BARRIERS
Fear of hostility or punishment
Lack of appropriate vulnerability disclosure avenues
Legal barriers or uncertainty
Insufficient or slow vendor or coordinator communication
CVD policy
MOTIVATIONS
For the security benefits
In response to customer demand
For the economic benefits
For ethical or social responsibility reasons
To raise awareness and engage with the community
Vendor
BARRIERS
Lack of awareness or understanding
Lack of organisational or technical capacity
Costs of implementation and operation
Legal barriers or uncertainty
Lack of management support
Disclosure

Source: (ENISA, 2018[49])

Awareness and knowledge of CVD

To implement a CVD process, vulnerability owners, security researchers and other stakeholders need first to be aware that it is available as an option that can contribute towards reduced digital security risk in general. They also need to understand how it works and what their role is within this process, which may be difficult given the complexity of CVD.

As indicated above, security researchers are anything but a homogeneous group. They include a wide variety of sociological profiles (e.g. professionals, talented amateurs, etc.) and people operating in different contexts (e.g. professional, hobby). Some may lack awareness and knowledge of what a good CVD process looks like, what their role is to make it successful, or how to manage complex situations, etc.

Many businesses operating in various industries are increasingly engaged in the digital transformation, for example by developing or embedding IoT devices in their products. Yet, they have a low digital and digital security maturity and are therefore unlikely to be sufficiently aware of responsible management and CVD to embrace good practice.

More generally, some internal stakeholders lacking a good understanding of this issue in organisations may discourage entering into a CVD process or vulnerability reporting to the organisation for a variety of reasons. For example:

- Marketing departments and business leaders may refuse to recognise that their products and/or system can have vulnerabilities, although it is the case for all products and systems.
- They may think that embracing CVD or adopting a vulnerability disclosure policy could damage their brand, although an objective of CVD is precisely to increase brand reputation by showing that the organisation is responsible with respect to addressing vulnerabilities.
- Some may interpret CVD as a broad encouragement for “hacking” the organisation, potentially leading to an ongoing digital security crisis, whereas having a vulnerability disclosure policy aims precisely at leveraging people of good will to reduce the likelihood of such crises happening.
- Others may fear that embracing CVD would attract criminals looking for opportunities for ransom and blackmail, although it would actually attract individuals motivated to protect the organisation from such threats rather than criminals.
- Some may also fear that some researchers may try to abuse the CVD process, for example by not adopting the expected behaviour at some point of the process, leading to possible damages for the organisation. However, according to experts, the vast majority of security researchers who enter into a CVD process are well intentioned and genuinely trying to reduce risk.

This means that basic knowledge about digital security vulnerabilities and CVD should be shared among the core security team as well as all departments who may have a role in decision-making or face consequences of a vulnerability disclosure failure (leadership, IT, marketing, public relations, legal, etc.).

Managing the sensitivity of vulnerability information

Vulnerability information can be highly sensitive. Where it is the case, international standards and industry best practices recommend only reporting it to the parties absolutely required to develop, test and deploy mitigation or remedial measures, and only to the extent necessary to enable them to do so. They also recommend that information concerning the vulnerability be kept in confidence (Center for Cybersecurity Policy and Law, 2019[65]; ENISA, 2016[35]). Communicating information concerning the vulnerability to other entities could increase the risk that information will leak, allowing threat sources to exploit the vulnerability. In some cases, such as when critical activities can be affected, it may be necessary to involve governments early in the co-ordination process, increasing the difficulty of keeping the vulnerability secret from potential threat sources as the number of people with access to this information increases (Johnson and Millett, 2019[22]).

Exchanges of vulnerability information should also use tools and techniques that reduce risks of confidentiality or integrity breaches, such as end-to-end encryption.

Balancing swift availability vs. mitigation quality

A key priority of CVD is to minimise the window of exposure in order to reduce the likelihood of the vulnerability being exploited. Therefore, co-ordination to ensure swift development and distribution of mitigations is key.

However, code owners need to ensure the completeness and effectiveness of a proposed mitigation, and assess the risk associated with its distribution. Mitigations introduce changes in products and therefore can create risks for users, in particular if they are not sufficiently tested in different usage scenarios and technical configurations prior to being distributed. For example, some patches initially provided by code owners to mitigate the 2018 Spectre and Meltdown vulnerabilities affecting microprocessors had negative effects on performances and compatibility with certain anti-virus software.

Assessing mitigation-related risks can take time, in particular in multi-party CVD. Requiring adherence to rigid deadlines and immature pre-disclosure to parties that do not take part in mitigation development would impede code owners' ability to assess mitigations-related risks. The Dutch National Cyber Security Centre (NCSC-NL) uses a period of 60 days between the reporting and the public disclosure, while highlighting that there may be circumstances in which this period is extended or shortened to consider the context. Hardware vulnerabilities, for example, may typically require 6 months (NCSC-NL, 2018[37]), or more in some cases. The difficulty and lengthy timeline of mitigating hardware vulnerabilities compared to the lifecycle of the physical product often means that hardware vulnerabilities are mitigated in practice by purchasing the next product available instead of repairing or patching the products currently owned.

Establishing trust

For co-operation to take place between parties, each stakeholder needs to trust that the relationship will yield positive outcomes.

In a CVD process, security researchers are facing different types of legal threats when they report vulnerabilities (discussed in 2.2). According to a 2015 survey by NTIA, 60% of researchers in the United States cited the threat of legal action as a reason for not working with a vulnerability owner to disclose a vulnerability (2016[73]). As pointed out by NTIA, fear of legal action is not a barrier *per se*, but may cause researchers to deviate from their default choices on disclosure. This suggests that increasing legal certainty may improve adoption of best practice.

In theory, the balance of power between parties can drive both sides to adopt a reasonable behaviour, and progressively trust each other. However, as noted by civil society, security researchers are usually at the weaker end of the spectrum when it comes to power dynamics in a relationship involving vulnerability owners, other private stakeholders and governments. If security researchers are recognised as making an important contribution to reducing digital security risk and elevating trust, governments should address the legal obstacles that can discourage them.

Most organisations have a low maturity with respect to vulnerability disclosure, and while some digital security teams may be more knowledgeable, it may often not be the case for IT, communications, legal departments, and high-level decision makers in general. This issue can be exacerbated in the context of IoT products, as IoT manufacturers often do not have a sufficiently mature digital and digital security culture.

In the absence of a basic digital security culture, vulnerability owners can feel threatened when receiving a report from a security researcher. Furthermore, a researcher may use language, intentionally or not, that could be interpreted as threatening the vulnerability owner to fully disclose the vulnerability, or by informing the media. He/she could also provide limited information about the vulnerability and ask for a reward as a condition for further co-operation. Some researchers also need to understand that they need to gain trust from vulnerability owners.

Well-recognised standards and good practices are essential to facilitate conflict resolution and decision-making. They can be used as a basis for all stakeholders to develop a CVD culture in the organisations and in the security researchers' community. Co-ordinators acting as trusted third parties are also key to smooth relationships and to resolve tensions. A security researcher lacking trust in a vulnerability owner can also launch a CVD process anonymously, for example by using a platform such as Zerodisclo.com, which was designed for this purpose by the bug bounty platform YesWeHack, using a mix of encryption and blockchain timestamp and signature.

Optimising communications and managing expectations

Co-ordination is primarily a matter of communication and information exchange between parties, as well as management of expectations.

A basic condition for CVD is for vulnerability owners to maintain an open and sufficiently secure communication channel to receive security reports and interact with researchers. This channel should also be easy to find, and its security measures sufficiently documented for researchers to use it with confidence that the vulnerability information will be protected from confidentiality breach. As mentioned above, the minimum expected requirement is that there is an easily available email address to report vulnerabilities and that it is monitored (security@example-company.com). Many researchers will begin their attempts to report by emailing that address. If there are no resources to monitor that email continuously, setting an auto-response can help to direct security researchers to the correct person, organization, or entity that can continue the process of CVD.

It is also essential to manage communications in a manner that sustains trust and drives the process to a positive outcome. In the case of CVD, the human factor is capital. The security researcher can be an individual, often isolated or in a small team, whereas the vulnerability owner is often an organisation with complex decision-making processes, limited agility, and controlled external communications. The cultural and motivational gaps between the researcher and the vulnerability owner are often wide, and can be exacerbated when the parties are located in different parts of the world and speak different languages. Reducing the probability of miscommunication is essential to ensure a successful CVD process, and requires an effort on both sides.

Setting clear expectations from the outset and meeting them is key. According to the NTIA survey, 84% of researchers involved in a CVD were available to answer questions about their report. 67% expected regular updates on the investigation and progress in mitigating their vulnerabilities. While 95% expected to be notified when the issue is resolved, only 58% were. 54% experienced frustrations during the process. As shown in Figure 8, researchers view communication with vulnerability owners not just as important to eliminate bugs more efficiently, but also as a recompense for the time that researchers put into vulnerability discovery. In return for their report, only 15% of them expected a compensation, but 70% expected regular communication, 57% expected being involved in testing mitigations, and 53% a simple acknowledgement.

Frustrated expectations, mostly around communication, are a major source of alternative behaviours to CVD such as full disclosure. Despite planning initially to disclose a vulnerability in a co-ordinated manner, nearly half of all researchers at some point considered disclosing publicly due to frustrated expectations, a behaviour also noticed by NCSC-NL. Furthermore, 32% shared publicly due to unmet timelines (NTIA, 2016[73]). Frustrations can lead to provocative reactions, such as the full disclosure by a Russian researcher of zero-day vulnerability information related to the widely used virtualisation software VirtualBox after the vendor (Oracle) took 15 months to fix a previous similar issue (Cimpanu, 2018[74]).

Figure 8. Researchers' expectations (NTIA, 2016)

What did researchers expected in return when reporting a vulnerability?

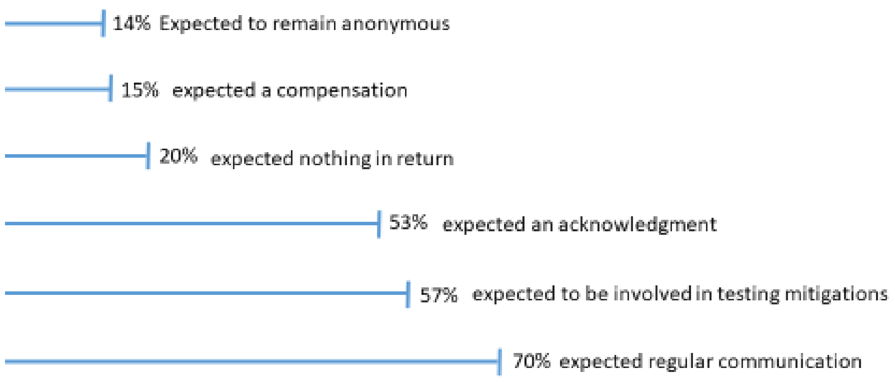

Source: (NTIA, 2016[73])

Different types of researchers are driven by different expectations. For example, in many cases, academic researchers are seeking visibility and aim to write and publish papers at academic conferences. Their draft paper is first submitted to a program committee prior to the researchers being invited to speak at the conference and the paper being published, a process typically taking 3 to 6 months. If the vulnerability owner needs more time to develop the mitigation or cannot commit six months in advance to a time when the public disclosure at the conference will take place, the researcher is put in a situation where he/she must put career advancement on hold. Researchers focused on protecting civil society organisations, for example as part of the CiviCERT[17] ecosystem, are likely to have different expectations, which are yet to be studied, and so do researchers in commercial security firms.

2.2. Legal risk for researchers

Legal risk faced by researchers when they report a vulnerability to a vulnerability owner is one of the most significant obstacles to CVD, as illustrated by the above-mentioned NTIA survey. It is enabled by an overall legal environment that does not sufficiently protect security researchers and by the behaviour of many vulnerability owners who threaten security researchers with legal proceedings when receiving reports. Overall, this situation creates a power imbalance that has been described as creating a chilling effect, limiting the adoption of CVD and undermining its potential benefits. The first section below provides an overview of the various facets of legal uncertainty, recognising that further work is needed to better understand legal obstacles to digital security research in general, and CVD in particular, across countries. The second section introduces existing public policy and private sector initiatives to address these obstacles.

While this section primarily focuses on legal risk faced by researchers, code owners may also face legal risk when they receive vulnerability reports. For example, acknowledging receipt of a vulnerability report without taking action to remediate it can expose the company to liability risk. Some experts report that there are companies that prefer not acknowledging receipt of some reports to mitigate this legal risk.

2.2.1. Areas of legal risk

Security researchers can face different types of legal risk when reporting vulnerabilities to vulnerability owners. Legal risk generally stems from the areas addressed in this section: criminal law, intellectual property law, contract law and, possibly, export controls regulations, with details varying significantly across jurisdictions. A security researcher may face a combination of these legal risks, aggravated by the complexity stemming from cross-border aspects and conflicts of jurisdiction when different legal regimes from several countries are at play, for example in the case of vulnerabilities in cloud services.

Criminal law

One could argue that the behaviour of a researcher who looks for and/or discovers a vulnerability is similar to that of a threat source, except for the intent, which is honest and benevolent for the former, and malicious for the latter.

According to the Cybercrime Convention, accessing the whole or any part of a computer system without right is a criminal offense when committed intentionally. The Convention requires signatory countries to adopt legislative measures establishing this criminal offense under their domestic law, while noting that they may require that the offence be committed with the intent of obtaining computer data or other dishonest intent (Council of Europe, 2001[75]).[18]

Interpretations of this provision vary significantly across countries that are parties to the Convention. In some countries such as Canada and Chile, cybercrime legislation requires a malicious intent for the access to constitute a criminal offense (EFF, 2018[29]). In the European Union, the Cybercrime Directive

(2013/40/EU) sets minimum protections, leaving EU members to adopt stricter rules if they wish. The requirement of intent is however not an effective protection for researchers since the intent is not necessarily apparent at all stages of the vulnerability discovery process.

In the United States, the Computer Fraud and Abuse Act (CFAA) provides that it is illegal for an individual to intentionally access a computer without authorisation or exceed authorised access and obtain information from any protected computer.[19] That language has been described as vague, having the potential to capture a large amount of research behaviour (Etcovitch and van der Merwe, 2018[76]), leading to inconsistent interpretations (CDT, 2018[77]) and failing to legitimise security research, often creating barriers for researchers (Elazari, 2018[78]). However, system owners can authorise access, enabling vulnerability owners to create safe harbours for researchers. VDPs provide a means for vulnerability owners to authorise testing by third parties, in which case security researchers are required to abide by the limitations they may contain. Furthermore, in 2017, the United States Department of Justice (DoJ) published a framework to assist organisations interested in instituting a formal vulnerability disclosure programme. The framework outlines a process for designing a vulnerability disclosure programme that will "clearly describe authorised vulnerability disclosure and discovery conduct, thereby substantially reducing the likelihood that such described activities will result in a civil or criminal violation of law under the CFAA" (US Department of Justice, 2017[79]). VDPs may take the form of contracts between parties in some cases. Through this mechanism, the CFAA provides a means for vulnerability owners to create a safe harbour for security researchers through VDPs, although limited to CFAA-related legal risk. The CFAA is a Federal law. Several State-level computer crime laws are also in force across the United States.

Another rarely mentioned risk of criminal legal proceedings is when a vulnerability owner interprets the behaviour of a researcher exploring the possibility of a reward for their work as an extortion attempt. There is a very thin line between negotiation and extortion in such contexts, and researchers must be very cautious not to cross it. However, some vulnerability owners tell researchers that they would fall afoul of extortion law unless they report the vulnerability for free, immediately, and sign a non-disclosure agreement (also discussed below), which then muzzles them if they wish to disclose the vulnerability more broadly in the public interest.

Intellectual property law

Vulnerability owners can claim that the researcher breached at least three areas of intellectual property law: copyright, trade secrets and patents.

Copyright law

Copyright law can be breached when the information disclosed contains portions of copyrighted software code. Such copyright protection could restrict sharing vulnerability information with the original vendor, making CVD difficult to implement in many cases.

The United States' Digital Millennium Copyright Act (DMCA) includes anti-circumvention requirements originally designed to protect media publishers against unauthorised copyright violations which have been interpreted to encompass a wide range of software protection mechanisms typically encountered when performing security audits (Gamero-Garrido et al., 2017[80]). The DMCA has been updated regularly to include exemptions for security testing under certain circumstances (Adams, 2018[81]). However, these exemptions are still pointed out as insufficient (Elazari, 2018[78]; Etcovitch and van der Merwe, 2018[76]; CDT, 2018[77]). As for the CFAA, these exemptions can be established by the vulnerability owner, enabling the possibility of safe harbours. European law does not provide exemptions (CEPS, 2018[64]).

CFAA and DMCA are Federal laws, and several State-level intellectual property laws are also in force across the United States.

Trade agreements can have unexpected effects in developing countries which often lack domestic expertise, resources and multi-stakeholder dialogue. For example, free trade agreements led to the adoption by some developing countries of inflexible anti-circumvention measures inspired by the first version of the US DMCA. These countries then rarely updated their frameworks to reflect subsequent improvements in the US law (Rimmer, 2017[82]; Lopez Romero, 2006[83]; Lerman, 2015[84]).

Trade secret law

Law on the protection of trade secrets can be breached when the vulnerability owner can prove that the researcher's prior knowledge led him to his discovery, such as when he was a former employee or consultant (CEPS, 2018[64]), or when a non-disclosure agreement was breached (EFF, n.d.[85]).

Patents law

At least one researcher in the US was threatened based on a patent infringement because he had created a homebrew device to demonstrate a vulnerability that arguably worked as an existing patented device (EFF, n.d.[85]).

Data protection law

Researchers who discover a vulnerability in an online system can access personal data, which could be interpreted as a breach of data protection law in some jurisdictions, unless the applicable law includes an exemption for digital security research, which is not the case (at least explicitly) in the EU GDPR (CEPS, 2018[64]).

Contract law

Bug bounty policies, and in some cases VDPs, constitute the terms of a contract between the vulnerability owner and the researcher. Breaching the terms of the contract entails legal liability and risks for researchers.

A review of bug bounty policies showed that many are confusing and difficult to analyse for security researchers who typically lack legal expertise. Furthermore, they often include language that shifts the legal risk to researchers (Elazari, 2018[78]):

- Some programmes do not refer to the legal terms as binding legal contracts and present them separately from the technical guidelines, which can lead to confusion;
- If the programme takes place through a bug bounty platform, both the vulnerability owner and the platform's policies apply, creating potential conflicts leading to confusion;
- Many policies require security researchers to comply with all applicable laws instead of granting researchers' authorisation to test their systems under laws such as DMCA and CFAA, thereby creating a safe harbour. Other policies do not mention compliance with laws, creating uncertainty.
- In some cases, the bug bounty policy indicates that researchers must not breach the terms of the product's End User Licence Agreement (EULA), while at the same time the EULA prohibits the use of security techniques (e.g. reverse-engineering) and even the mere attempt to gain unauthorised access. There are even cases where the bug bounty policy explicitly notes that it does not give any permission to penetrate the vulnerability owner's systems.

The reasons for such shortcomings are unclear. Some may be organisational, such as when teams responsible for the bug bounty programme and the firm's lawyers do not sufficiently communicate, or share the same understanding of the programme's purpose. Lawyers will typically look for ways to minimise the legal risk for the company, while security experts can neglect potential legal implications of their disclosure programmes as well as impact on the firm's reputation in case of legal tensions with a researcher. These elements suggest that BBP are often not sufficiently integrated into products' business strategy.

Non-disclosure agreements, already mentioned above, are another area of concern for researchers as they can be construed as prohibiting any future disclosure of a vulnerability, thus preventing academic publication and presenting the researcher's work at a conference. These agreements can discourage researchers from reaching out to vulnerability owners, and lead to full disclosures.

Export controls

Export controls legislation and regulation can also create legal uncertainty for security researchers as they may apply to tools, techniques and even knowledge that are typically used to discover vulnerabilities.

The Wassenaar arrangement is an overarching international framework for export controls. It gathers 41 countries that meet regularly and agree to control certain technologies by imposing an export licence requirement at the national level to transfer these technologies abroad. Part of this arrangement relates to surveillance software (e.g. "technology for the development of intrusion software"). It could be interpreted until 2018 as covering digital security technologies used for reverse engineering, as well as vulnerability information. In 2018, a modification of the Arrangement promoted by the United States explicitly excluded vulnerability disclosure and incident response from the technologies concerned (Wassenaar Arrangement Secretariat, 2018[86]). While many security experts welcomed the modification as an improvement, some viewed this exemption as "a line in the sand", noting that depending on the interpretation and circumstances, vulnerabilities and exploits may be exempted or may satisfy the definition given for intrusion software (Ruohonen and Kimppa, 2019[87]).

Some experts have also expressed concerns that vulnerability information exchanged across borders and involving individuals or organisations in countries targeted by extra-territorial sanctions could create legal uncertainty for security researchers.

2.2.2. Addressing legal risk

The threat of legal proceedings by vulnerability owners against security researchers is not rare and is well known in the security community. Box 4 provides some examples among the numerous cases regularly reported. They show that legal threats can come from vulnerability owners as well as from governments, including for political reasons in some countries.

According to some authors, security researchers are rarely prosecuted after reporting vulnerabilities, and even more rarely are they successfully convicted of criminal charges, nor do they often lose civil suits based on these statutes, at least in the United States. However, legal threat without actual prosecution is sufficient to undermine their ability to publish research and to create a chilling effect acting as a powerful disincentive for CVD (Etcovitch and van der Merwe, 2017[88]). An empirical study showed that most product manufacturers are reluctant to surrender legal recourse and either are unwilling to engage on questions of permission or will impose significant restrictions on researchers who do so. There is also significant difference in the responsiveness afforded to academic versus independent security researchers. Furthermore, the study confirms earlier findings by NTIA that legal concerns are significant for many vulnerability researchers, with almost a quarter of researchers surveyed reporting experiences of legal threats or action in the course of their research (Gamero-Garrido et al., 2017[80]).

Fear of prosecution can have different effects depending on whether, for example, the security researcher is a hobbyist, a professional penetration tester or an academic researcher. In some cases, the latter may get support from his/her university's legal department, altering the power balance with the vulnerability owner. However, universities lacking a well-resourced legal department may simply discourage vulnerability research to avoid this type of legal pressure.

This section examines initiatives carried out by governments and private sector to reduce the legal risk faced by researchers, and introduces some key considerations for further reducing legal risk.

Box 4. Examples of researchers threatened with legal proceedings

These examples illustrate cases of researchers threatened with legal proceedings.

In 2011, a teenager showed the Finnish online game platform Habbo (273 million users in 150 countries) how he had been able to log into its helpdesk system, whereupon the company brought criminal charges. Two years later, the courts ruled that there was no case to answer (van't Hof, 2015[89]).

In 2012, three security researchers at Radboud University (Netherlands) discovered weaknesses in a chip widely used in immobilisers for various car brands. They informed the chip manufacturer and wrote a scientific article that was accepted for publication at a digital security symposium. However, in June 2013 an English court, acting at the request of Volkswagen, ruled that the article had to be withdrawn. In August 2015, Volkswagen ultimately agreed to the release of the publication after accepting the authors' proposal to remove one sentence from the manuscript (Radboud Universiteit, 2015[90]).

In 2015, security researcher Joaquin Sorianello reported a vulnerability to Magic Software Argentina (MSA), the producer of an e-voting application in Argentina that would be used for elections the following week. Three days before the elections, the police raided his apartment, and seized his electronic equipment based on the criminal charges presented by MSA. The case was dismissed one year later on the ground that he had not accessed MSA systems unlawfully or caused any harm.

In 2016, researchers at a US security company received a cease-and-desist letter three days after reporting a serious vulnerability to the global consulting and auditing company PwC (whittaker, 2016[91]). Another researcher had his home raided and was arrested by the FBI after he reported that a dental software company left unencrypted sensitive health information of 22 000 patients at risk of access by others (Doe, 2016[92]).

In 2017, out of curiosity, a Danish citizen discovered a vulnerability in the Frederiksberg Municipality web site that enabled the harvesting of personal information of any citizen by entering their birth date in a form. He automated the process to demonstrate the flaw and reported the vulnerability to the municipality. The service provider discreetly fixed the vulnerability and reported the researcher to the police (Andersen, 2017[93]).

In 2017, Javier Smaldone reported how the Argentine Federal Police suffered a leak of information from their email accounts that included the dissemination and publication of a huge amount of information, including personal data of law enforcement personnel and their families, data about witnesses in judicial investigations, data of complaints and judicial eavesdropping. In 2019, the attack known as “The cap leaks” was repeated and federal forces raided Smaldone’s home, and seized his computers and phones. No criminal charge was ever brought against him (AccessNow, 2020[94]). He was still trying to recover his computer equipment as of 2020 according to civil society sources.

In 2018, the FBI investigated a student from the University of Michigan who had been reported by the mobile voting company Voatz for illegally attempting to hack its application. The company claimed that the system tested by the researcher was outside the scope of its HackerOne bug bounty programme. However, this exclusion clause was inserted by the company on the bug bounty programme’s page after the student had found the vulnerability as part of a college class assignment.

In 2020, MIT researchers uncovered other vulnerabilities in Voatz’s system that could allow hackers to "alter, stop, or expose how an individual user has voted”. The application had already been used in several local and State elections in the United States. The researchers reported their findings to the

Department of Homeland Security's Cybersecurity and Infrastructure Security Agency (CISA) in order to avoid the experience of previous researchers interacting with the firm. The company disputed the severity of the vulnerabilities, making quite aggressive public statements against the researchers. Ultimately, an independent audit requested by Voatz confirmed the MIT's findings. HackerOne terminated its partnership with Voatz in March 2020 (Grauer, 2020[95]; Freed, 2020[96]; Cimpanu, 2020[97]).

Initiatives by governments

Some countries are taking action to create a favourable legal environment for security research and CVD. According to the CEPS report on Software Vulnerability disclosure in Europe (2018[64]), ten EU countries had plans to develop a policy to support vulnerability disclosure but had not yet reached a consensus at the political or legislative level, and seventeen countries had no such plans. The Netherlands, France and, to a lesser extent, Lithuania were the only EU members providing some protection to researchers. The CEPS report calls the EU to amend the Cybercrime Directive to allow for the smooth and rapid development of CVD and to clarify the legal responsibility of researchers.

In the Netherlands, the Public Prosecution Service developed guidelines on how to decide whether to initiate a criminal investigation and/or to prosecute, which take into account compliance with an existing CVD policy or, in absence of a policy, the general concept of CVD. However, the Public Prosecution Service and the police would investigate a case further if there are indications that the researcher has consciously or unconsciously gone too far in their actions and/or failed to comply with the CVD policy (NCSC-NL, 2018[98]). The prosecutor's guidelines encompass most cases, and court decisions provide further clarifications on specific aspects, progressively reducing residual uncertainty.

In Belgium, the Centre for Cyber Security Belgium (CCB) has also recently published guidelines to encourage the adoption of CVD policy or bug bounty for private and public entities (CCB-BE, 2020[99]). This documentation, developed in collaboration with the Public Prosecution Service and the "ethical hackers" community, provides certainty for researchers when a vulnerability owner has adopted a CVD policy and also attributes a role to the CCB (CERT.be team) as a CVD coordinator by default, even when there are no CVD policy in place.

In France, the law protects a researcher when he/she reports a vulnerability to the government national cybersecurity agency (ANSSI), which can inform the vulnerability owner without disclosing the researcher's identity (République Française, 2016[100]; ANSSI, n.d.[101]). However, this mechanism does not protect the researcher from prosecution if he/she has gone too far in his/her actions such as the exploitation of the vulnerability or breach of intellectual property laws, or if he/she decides to disclose vulnerability information later on, such as for academic purposes. It also requires the researcher to trust ANSSI in the first place. In practice, many researchers who report a vulnerability to ANSSI do so after failing to initiate a CVD process with the vulnerability owner.

Lithuania adopted a vulnerability disclosure framework limited to providers of public communications networks. It includes a disclosure deadline, scheduled resolution and an acknowledgment report.

The Latvian experience illustrates the difficulty that policy makers can face in trying to develop a framework to protect researchers. In 2017, the government tried to develop amendments to its cybercrime and IT law in order to protect security researchers. However, the process failed as the State police insisted on the creation of a register of researchers, eliminating the possibility of anonymity for researchers. In addition, some stakeholders erroneously feared that the amendments would have enabled anonymous actors to attack governmental systems without the possibility of suing them (CEPS, 2018[64]).

In the United States, the possibility of safe harbours already exists in the DMCA and CFAA. However, much effort is still needed to overcome the existing chilling effect (CDT, 2018[77]). While some authors call for revising these acts, others emphasise the need to encourage vulnerability owners to use clearer vulnerability disclosure policies that take advantage of the possibility to create safe harbours. All stakeholders can play a role. The government provides guidance on what is likely to constitute a good practice according to the law (i.e. 2017 DoJ Framework). Lead vendors and platforms can work with other stakeholders to improve and promote their good practice. As suggested by Amit Elazari, stakeholders can co-operate to standardise bug bounty policy legal terms, for example using a model akin to Creative Common. The researchers' community can also organise itself to negotiate improved bug bounty policies or create a reputation system for legal terms which are often imperfect (2018[78]).

In general, an effective and ongoing dialogue between the government and the digital security community seems to be a useful prerequisite to foster the adoption of a favourable legal environment, as illustrated in the United States (e.g. improvement of the export controls mechanism) and in the Netherlands (e.g. NCSC-NL CVD guidelines) or Belgium (CCB-BE Guide). Furthermore, when a government decides to promote CVD, it should ensure that agencies in charge of enforcing legislation and regulation that could create legal risk for researchers (e.g. cybercrime laws) understand CVD. It should also ensure that they are trained to differentiate well-intentioned security researchers from malicious actors, whether they are located in their jurisdiction or abroad.

The inclusion of CVD in the US NIST Cybersecurity Framework and EU cybersecurity certification schemes through the EU Cybersecurity Act will certainly help mainstream CVD and facilitate its recognition as a best practice. ENISA's mandate to help EU members develop CVD policies on a voluntary basis is also an important step forward.

Private sector initiatives

Businesses can also take action to create a clearer and more secure legal environment for security researchers. For example, they can adopt a safe harbour policy whereby they commit not to prosecute people who investigate vulnerabilities in good faith. For example, Microsoft's Bounty Legal Safe Harbour states that the company "will not pursue civil or criminal action, or send notice to law enforcement for accidental or good faith violations of Microsoft Bug Bounty Terms and Conditions" and encourages researchers to contact the company before engaging in conduct that may be inconsistent with or unaddressed by this policy. If the can vulnerability affect a third party, Microsoft commits to limit the amount of information about the researcher that it will share with the third party (n.d.[102]).

Businesses can also systematically review their terms of service and licenses to ensure consistency with their vulnerability disclosure policy and bug bounty programmes.

Reducing legal risk

Unjustified legal threats by vulnerability owners against security researchers can damage the public interest and create a chilling effect. They are enabled by a power imbalance between vulnerability owners and security researchers in most cases. To facilitate security research, public policies can aim to discourage unjustified legal threats, and update legal frameworks to reduce the possibility of their occurrence. While many experts agree on the need to address the chilling effect, there is limited agreement on how to proceed. This section introduces considerations for future work in this area.

Recognising security researchers' responsibility

According to the OECD Digital Security Recommendation, "all stakeholders should act responsibly and be accountable, based on their roles […] for the management of digital security risk and for taking into account the potential impact of their decisions on others". Therefore, the recognition that security researchers can play an important role to reduce digital security risk implies that, as other stakeholders, they should take responsibility for their action when engaging in product and system vulnerability testing and when disclosing their findings. Furthermore, the reduction of digital security risk cannot be raised as a justification for breaching the law, especially if others are harmed, and vulnerability testing without authorisation can harm others. Therefore researchers assume the risk of lawsuits or prosecution when they conduct testing on products or systems owned by another party.

This is unlikely to change. It is difficult to develop legislation to explicitly protect researchers who follow good practice against prosecution because it is not possible to comprehensively and objectively describe such good practice in detail. Each case of vulnerability disclosure is different and existing good practice, although valid at a high-level, suffers many exceptions in practice and is likely to evolve over time. Therefore, legal frameworks need to account for a degree of interpretation, which inevitably leaves some uncertainty for researchers.

Understanding the notion of safe harbour

According to Amit Elazari (2018[78]), security researchers “want to play by the rules, but the rules often don’t let them. Therefore the rules should change”. Together with other experts, she highlights the need to create safe harbours for security researchers. This may require different efforts in different countries with different legal regimes, and in different areas of law.

The term “safe harbour” refers to a specific legal mechanism within a given legislation (e.g. against cybercrime) that provides a way for researchers to be protected under certain conditions defined by the vulnerability owner, generally in a VDP.

However, while safe harbours increase legal uncertainty, they do not create a blanket exemption for researchers, and do not mean that researchers will not be sued, even if they respect the VDP.

For example, in the United States, the government recognises that good faith security researchers should be able to engage in a CVD process because vulnerability information sharing can reduce digital security risk. The CFAA and DMCA safe harbour provisions are based on the possibility for the vulnerability owners and researchers involved in CVD to develop sufficient mutual trust to allow for testing while staying within the boundaries of the law. CFAA safe harbours protect researchers who test systems covered by a VDP that provides authorisation for testing, and only within the limits detailed in that VDP. It is the responsibility of the researcher to ensure that the system they test is covered by the VDP and to respect its conditions. It means that if there is no VDP, there is no safe harbour. It also means that if there is a VDP with ambiguous content, or in contradiction with terms of service (for example), the safe harbour may not be so safe.

Furthermore, a safe harbour is never an absolute protection against legal risk. For example, an organisation's VDP can promise that the organisation will not sue the researcher, but this may not shield the researcher from third party lawsuits or prosecution, such as if there is a violation of export laws or restrictive disclosure laws in other countries. It is crucial for researchers to understand the limited protective effect of safe harbours for them not to overstep the permissions they are granted. It is also crucial that policymakers understand that safe harbours and VDPs will not protect good faith researchers in all circumstances, and that perhaps additional policies would need to be developed to complement them.

Understanding the risk to tailor policy action

Levels of risk for researchers vary depending on the legal area concerned. For example, testing online systems exposes to a breach of cybercrime and, potentially, data protection legislations, while reverse-engineering software products exposes to intellectual property lawsuits. While the domestic legal framework may create more or less opportunities to mitigate these legal risks, there may be additional complexity when researchers and vulnerability owners are located in different countries or regions. To formulate a set of high-level policy objectives aiming at reducing risk, it might be useful to understand the gaps and according to different basic risk scenarios, such as:

- The researcher engages in a CVD process with a vulnerability owner who supports CVD and has a public VDP matching the researcher's intentions. This is the most favourable case for researchers because the achievement of mutual trust is expected from the outset on each side. Yet it is imperfect, as explained above: the VDP can be ambiguous or conflict with terms of service, researchers may violate third parties rights, or other countries laws, etc.
- The researcher engages in a CVD process with a vulnerability owner who agrees to follow-up but does not have a public VDP. This is a more uncertain scenario for the researcher because in addition to the uncertainty from the previous scenario, the vulnerability owner could change its mind in the course of the process and does not publicly commit to exclude legal proceedings. The Dutch legal approach seems to address this case elegantly by recommending that the public prosecutor consider good CVD practice when deciding on whether to prosecute.
- The researcher engages in a full disclosure (i.e. no CVD). This is the most complex and uncertain scenario for security researchers, but it cannot be excluded because of the wicked nature of vulnerability disclosure making full disclosure sometimes the best solution. Risk mitigation would consist in exploring avenues for co-ordination, consulting a co-ordinator, etc.

With such a risk assessment, it would be possible to carry out a gap analysis and tailor policy action according to the domestic context. For example, in some countries, it may be necessary to amend legal frameworks, for example if the cybercrime legislation does not provide for an exception for research. In other countries, it may be sufficient to follow the Dutch example and interpret existing legislation in a manner that takes into account the positive role of security researchers and existing good practice for CVD.

2.3. Facilitating CVD

2.3.1. Standards and certification schemes

Standards are a useful means of facilitating co-ordination by providing parties with a shared understanding of processes and procedures. International standards are particularly relevant with respect to CVD since the co-ordination often takes place across borders and local standards, social norms and laws may create confusion and uncertainty among stakeholders.

The International Organization for Standardization (ISO) and International Electrotechnical Commission (IEC) have developed two joint standards which together partially address CVD. ISO/IEC 29147:2018 on Vulnerability disclosure provides requirements and recommendations to vulnerability owners (called vendors) on the disclosure of vulnerabilities in products (ISO/IEC, 2018[67]). Its first version was released in 2014. This standard approaches vulnerability disclosure primarily from the perspective of code vulnerabilities. ISO/IEC 30111:2019 on Vulnerability handling processes provides requirements and recommendations for how to process and remediate reported potential vulnerabilities in a product or service (ISO/IEC, 2019[39]). Its first version was released in 2011. There is currently no ISO/IEC standard on vulnerability management, but there are guides and domestic standards.[20]

The CERT Guide to Co-ordinated Vulnerability Disclosure provides guidance on how to implement CVD (CERT/CC, 2017[43]).

At a more technical level, the OASIS Common Vulnerability Reporting Framework (CVRF) provides a language supporting the creation, update, and interoperable exchange of security advisories as structured machine-readable content (OASIS, 2017[103]).

Governments can play an important role in promoting the adoption of good practice and standards. In November 2019, the United States Department of Homeland Security (DHS) published a draft Binding Operational Directive (BOD) applicable by US Federal Government, executive branch, departments and agencies, to develop and publish a vulnerability disclosure policy (DHS, 2019[104]). The draft Directive includes useful guidance on how a government agency should handle vulnerabilities.

The 2019 EU Cybersecurity Act recognises that “co-ordinated vulnerability disclosure policies could play an important role in Member States’ efforts to enhance cybersecurity” and gives a mandate to ENISA to assist EU Members’ institutions, bodies, offices and agencies in establishing and implementing vulnerability disclosure policies on a voluntary basis. The Act also establishes a framework for EU cybersecurity certification schemes, which shall include “rules concerning how previously undetected cybersecurity vulnerabilities in ICT products, ICT services and ICT processes are to be reported and dealt with”. Manufacturers or providers of certified products shall make public their contact information and accepted methods for receiving vulnerability information from end users and security researchers, information which ENISA will make available on a web site (European Union, 2019[105]).[21] In July 2020, ENISA launched a public consultation on the *Common Criteria based European candidate cybersecurity certification scheme* (EUCC scheme), which contains harmonised conditions for vulnerability handling and a fast track assessment procedure for patches (ENISA, 2020[106]). Some voices have highlighted that it might be more effective to adopt general rules for vulnerability handling that would apply to various schemes rather than insert vulnerability handling rules in each scheme. The articulation between standards, certification and digital security of products in general is further discussed in (OECD, 2021[3]) and (OECD, 2021[4]).

Some broader digital security standards can also promote and facilitate the adoption of CVD. For example, the NIST Cybersecurity Framework version 1.1 (2018) includes a subcategory related to the vulnerability disclosure lifecycle (NIST, 2018[107]). NIST also recently published two Interagency Reports (IR) that include voluntary guidance on vulnerability management in the context of IoT devices (NIST, 2020[108]; NIST, 2020[109]). One of the thirteen provisions of the ETSI Technical Specification “Cyber Security for Consumer Internet of Things” focuses on the need to implement a means to manage reports of vulnerabilities (ETSI, 2019[110]), initially based on the UK Code of Practice for Consumer IoT (UK DCMS, 2018[66]). At the time of writing ETSI is discussing the evolution of this “Technical Standard” into a higher-level “European Standard”. Furthermore, the UK Government is pursuing regulation requiring that all new products adhere to minimum digital security requirements based around aspects of the Code of Practice and the ETSI standard. The draft regulation includes a provision stating that “all companies that provide internet-connected devices and services shall provide a public point of contact as part of a vulnerability disclosure policy in order that security researchers and others are able to report issues." Device manufacturers, IoT device providers and mobile application developers would be required to continually monitor for, identify and rectify security vulnerabilities within their own products and services as part of the product security lifecycle.

2.3.2. Vulnerability Disclosure Policies (VDP)

A VDP is an essential tool for vulnerability owners to invite researchers to send reports, increase their confidence that reports will be welcome and handled seriously, and that the reporters will not be subject to legal action if they stay within the policy’s boundaries. The VDP helps clarify researchers’ expectations by setting clear rules of the game (DHS, 2019[104]). If they are sufficiently readable and visible, VDPs can reduce the likelihood of vulnerability reports exploring the possibility of a reward to be interpreted as extortion attempts. However, VDPs are still emerging and far from widespread: a 2018 study of 331 consumer IoT products in the UK showed that 90% of the manufacturers lack a vulnerability disclosure policy (IoT Security Foundation, 2018[111]).

While a basic VDP can be as short as a single paragraph indicating how to securely send a vulnerability report to the vulnerability owner, a typical and more effective VDP explains:

- The contact method for secure communication;
- Preconditions for reporting parties;

- Clear expectations for handling a report;
- Methods for rewarding a report (NCSC-NL, 2018[37]).

It is important that vulnerability owners carefully consider the content and tone of communications, as well as set and meet clear expectations regarding communications' frequency, for example acknowledging receipt within 3 days, providing an assessment of the vulnerability within 7 days, communicating about the resolution within 90 days and including regular updates in the meantime.[22]

The VDP has a legal dimension. For example, the VDP needs to provide a high level of certainty to the researcher that he/she will not face legal proceedings if he/she respects the terms of the VDP, a matter further discussed in 2.2. This is particularly useful to encourage reporting, including from researchers located in other jurisdictions.

Disclose.io provides standardised best practice language for VDPs that set a safe harbour to enable good-faith security research, with accessible and understandable language. The terms are available under a Creative Common licence, with international versions and versions tailored for Canada and the United States. They suggest including sections on:

- Scope, i.e. the list of assets for which the organisation is explicitly allowing and encouraging security research, and, optionally, a non-exhaustive list of systems and security testing activities that the organisation does not authorise testing against,
- Rewards, if relevant, i.e. information on whether or not the program offers payment for valid, unique issues, as well as the type and parameters of that compensation;
- Official communication channels, i.e. a full list of the communication methods available to receive and communicate about vulnerability submissions.
- Disclosure policy, i.e. a clear policy outlining the conditions under which a researcher can disclose the details of a reported issue to the public or third parties.

Some best practices provide guidance on practical aspects related to VDPs (EdOverflow, n.d.[112]). For example, an IETF informational specification is currently being developed to standardise the format of a VDP accessible on a web site as a /security.txt file (Foudil, Shafranovich and Nightwatch Cybersecurity, 2020[113]; EdOverflow and Shafranovich, n.d.[114]).

VDPs can include non-monetary rewards and credits. For example, NCSC-NL's VDP indicates that the agency "provides a reward by way of thanks for the assistance. Depending on the severity of the security problem and the quality of the report, the reward can vary from a T-shirt or a gift voucher to a maximum of EUR 300. It must relate to a serious security problem of which the NCSC is not yet aware" (NCSC-NL, n.d.[115]). The Dutch government T-shirt became quite popular among security researchers, with some displaying it as a trophy on their social media account, together with the letter from the NCSC (Figure 9). Acknowledgement letters and other sign of recognition are also important non-monetary rewards that can fulfil many researchers' expectations. For many researchers, prestige may be an important incentive. For example, the Korean government credits security researchers in the public description of the vulnerability they report (known as CVE entry) and includes them in a "hall of fame". If they are sufficiently visible and clear with respect to rewards, VDPs can reduce the likelihood of the vulnerability owner interpreting a report as an extortion schemes.

While VDPs are a very useful tool, they do not always match researchers' expectations or intentions. In practice, researchers can always research and disclose vulnerabilities in the manner they want (full disclosure, anonymous disclosure, etc.), and face the positive and negative consequences of doing so (cf. 2.2).

Figure 9. The Dutch example of a non-monetary reward

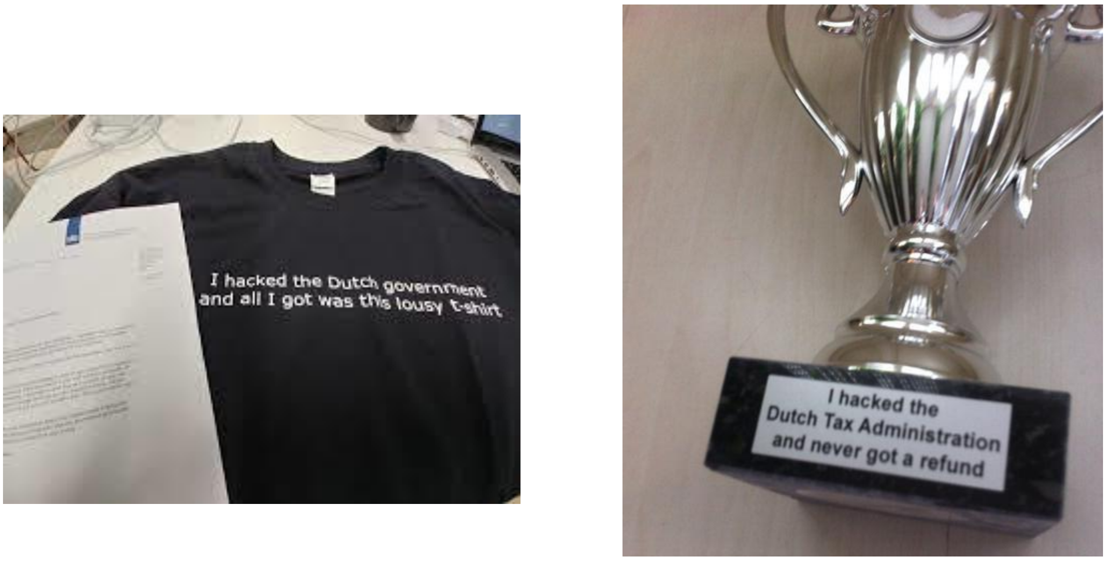
I hacked the Dutch government and all I got was this lousy t-shirt
I hacked the Dutch Tax Administration and never got a refund

Note: The T-shirt text is "I hacked the Dutch government and all I got was this lousy t-shirt". The cup's text is "I hacked the Dutch Tax Administration and never got a refund".
Sources: left: Amal Thamban, www.linkedin.com/pulse/i-hacked-dutch-xss-amal-thamban/, right: Jeroen van der Ham.

2.3.3. Bug bounty programmes (BBP)

Bug bounty programmes, also called "bug bounties" or vulnerability rewards programmes, are crowdsourcing initiatives that reward individuals for discovering and reporting vulnerabilities to vulnerability owners. They can be viewed as an open contract to research vulnerabilities that vulnerability owners put on the market for any interested individuals to enter into. They represent a shift from a passive to a proactive approach, whereby vulnerability owners publicly call security researchers to find vulnerabilities in their products or systems instead of only welcoming possible reporting through a VDP. Unlike VDPs, BBPs allow vulnerability owners to set precise expectations, as they would when purchasing security test services from a security firm.

Vulnerability owners can launch public BBPs (i.e. open-to-all) or private BBPs (i.e. by invitation only), before or after product release. BBPs can take a white-box approach, where bounty hunters have access to documents and software code to ease the identification of vulnerabilities, as opposed to a black-box-approach where the bounty hunter is in the same position as an attacker and does not have access to the vulnerability owner's internal information.

Some large organisations are often contacted by researchers who have found a vulnerability and offer to report it in exchange for a reward. Unless this takes place as part of and is in line with an existing bug bounty programme established by the organisation, it is likely to be interpreted as a form of extortion and can lead to legal proceedings.

The first BBP was launched by Netscape in 1995. In 2002, security firm iDefense (now part of Accenture Security) launched a BBP, followed by the Mozilla Foundation in 2004. Trend Micro created the Zero Day Initiative in 2005. Google, Barracuda Networks, and the Deutsche Post launched their programmes in 2010, Facebook in 2011, and Microsoft in 2013 (Friis-Jensen, 2014[116]). As of 2014, many other organisations joined these early adopters.

Apart from open source products, BBPs have to be organised under the authority of the entity responsible for the tested product or system, otherwise they provide a means for vulnerability brokers to harvest vulnerabilities, feeding the grey market.

There is no comprehensive inventory of public vulnerability disclosure policies and BBPs in the world. However, the vulnerability disclosure search engine Firebounty.com found 2 527 English language VDPs and BBPs on the Internet as of March 2020, including 67% VDPs accessible through a security.txt file, 25% public BBPs, and 7% standalone VDPs. These numbers include public BBPs intermediated through bug bounty platforms. Disclose.io maintains a community-powered index of known public bug bounty and vulnerability disclosure programmes. As of April 2020, 44% of the 886 entries were BBPs and 66% were VDPs. This index, however, only includes some programmes carried by some platforms (Disclose.io, 2020[117]).

Several governments are also launching BBPs. For example:

- The US Department of Defense (DoD) launched a two month-long pilot in 2016 targeting vulnerabilities in five public-facing DoD web sites called “Hack the Pentagon”. A total of 1 410 researchers sent 1 189 reports (i.e. one every 30 minute in average), the first one 13 minutes after the programme kick-off, and more than 200 within 6 hours. Another programme included “Hack the Army” later in 2016 (HackerOne, 2017[118]). 138 report were deemed valid and unique vulnerabilities, with researchers receiving a total of USD 75 K. Building on this success, the department invested USD 34 M in bug bounty programmes in 2018 (US Department of Defense, 2016[119]; Boyd, 2018[120]).
- The Singapore Government Technology Agency (GovTech) and Cyber Security Agency (CSA) carried out three BBPs in 2018 and 2019, covering internet-facing government systems. In the second BBP, the agencies received 4 high and 27 medium and low severity vulnerabilities. 290 researchers, including 70 Singaporeans, participated and earned a total of USD 26 K. A single 24-year-old Singaporean found nine vulnerabilities and was awarded USD 8.5 K. The third BBP was expanded to include mobile applications (GovTech Singapore, 2019[121]).
- The Swiss government offered a total of USD 150 K for vulnerability reports in its internet-based e-voting system in 2019, with rewards ranging from USD 100 for examples of best practices not being followed, up to USD 50 K for undetectable vote manipulation (Porter, 2019[122]). It also carried out a bug bounty programme for its “SwissCovid Proximity Tracing System” in 2020 (Switzerland National Cyber Security Centre (NCSC), 2020[123]).
- The French government launched a BBP focusing on its Tchap instant messaging application dedicated to civil servants. Bounties ranged from EUR 50 to 1500 (DINSIC, 2019[124]).
- The Korean Ministry of Science and ICT (MSIT) and the Korean Information Security Agency (KISA) have established a BBP called “Hack the challenge” focusing on the web sites of volunteer private companies as well as KISA. They also collect code vulnerabilities through a BBP and send the vulnerability information to the code owner with a request to develop a mitigation.

Some BBPs launched by companies and governments support vulnerability discovery in open source products. For example, Google has been distributing bounties ranging from USD 500 to 20k since 2013, totalising hundreds of thousands of dollars (Google, 2013[125]; Mohit Kumar, 2019[126]). Microsoft, Facebook, the Ford Foundation and Github sponsored the Internet Bug Bounty which rewarded over USD 733k to 202 researchers for uncovering 827 vulnerabilities in internet-related open source products, including Heartbleed (USD 15k) and Shellshock (USD 20k) (Internetbugbounty.org, n.d.[127]). In 2018, the European Union also set up a BBP targeting 15 open source software with bounties from USD 30k to 100k (Reda, 2019[128]; Mayersen, 2018[129]).

As an indication that BBP are in the air, a dark web marketplace offering illegal products launched a BBP in 2017 to identify and mitigate security issues that might allow other hackers or law enforcement to identify and de-anonymise the site's owners and users. The owners offered between 0.05 to 10 bitcoins (over EUR 90k as of Q1 2020) for valid vulnerability reports (Cimpanu, 2017[130]).

The two main US bug bounty platforms have released surveys which provide indications of bug bounty hunters' profile (HackerOne, 2020[131]; Bugcrowd, 2019[132]). In general they:

- Are male: 91% (Bugcrowd);
- Are young: 71% are 18-29 year old (Bugcrowd), 42% are 18-24 and 41% are 25-34 year old (HackerOne);
- Have completed some form of higher education (80%, including 18% holding a masters, according to Bugcrowd);
- Do not hunt bugs full time (77%, Bugcrowd). Only a small percentage practice bug hunting full time, spending more than 40 hours (16% for HackerOne, 6.8% for Bugcrowd) or between 31 and 40 hours (8% for HackerOne, 4.7% for Bugcrowd).

The average yearly payouts of the top 50 hackers on the Bugcrowd platform in 2019 was USD 145 000. The average submission payout per vulnerability across the platform was USD 783, representing a 73% increase year on year.

BBPs' rewards vary considerably across organisations, both in terms of total amount paid and maximum reward for a single report. Further work would be necessary to enable meaningful comparisons, taking into account companies' size, revenues, etc. and to understand the pricing models for bounties. In this respect, corporate BBP organisers can theoretically value vulnerability information by taking into account the cost of mitigating the vulnerability and the potential revenues losses for not doing it. However, governments must almost always evaluate the worth of this information negatively by assessing the value of the risk avoided by fixing this vulnerability.[23]

Income is far from the sole motivation for bug bounty hunters, as shown in Figure 10 and Figure 11. Education, being challenged and taking pleasure are equally important.

Figure 10. Researchers' motivations : why do you hack ? (single response)

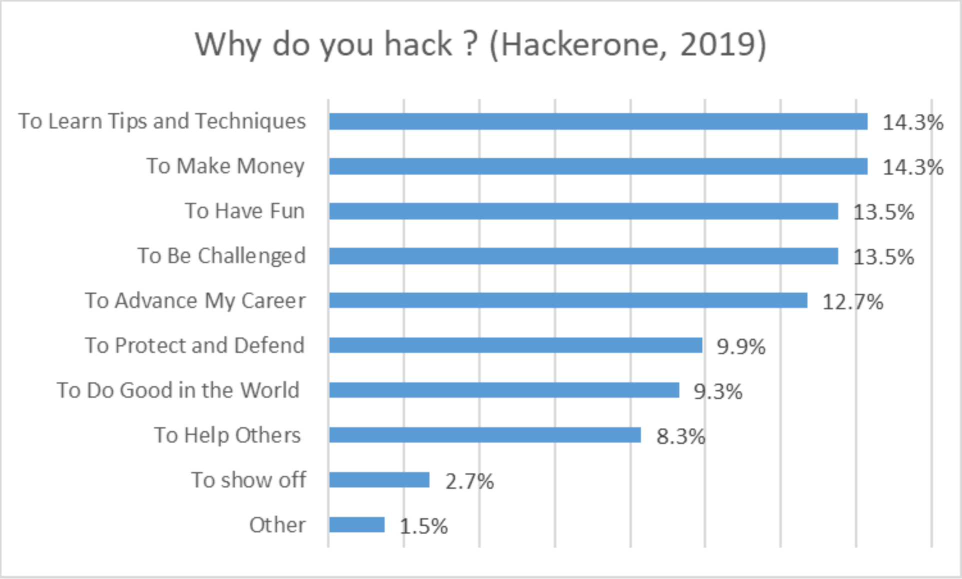

Source: (HackerOne, 2019[133]).

Figure 11. Researchers' motivations: why do you hunt bugs (multiple responses)

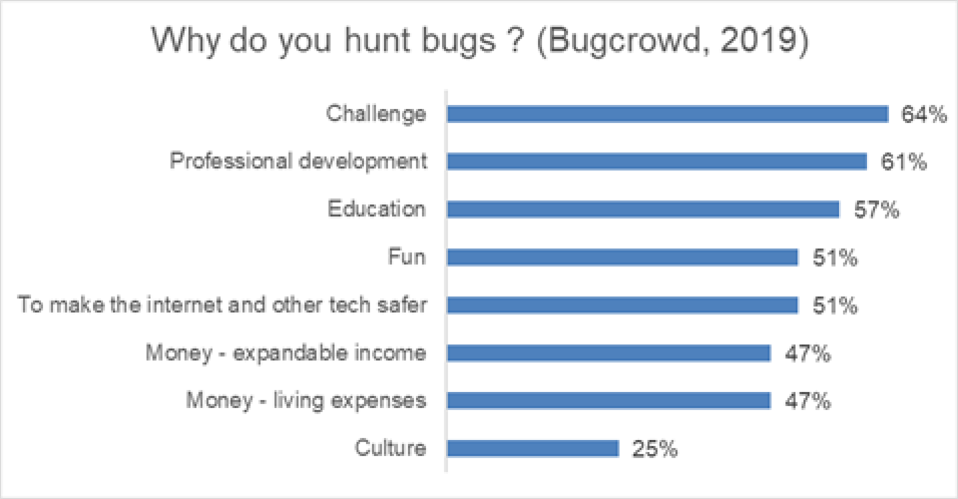

Note: respondents could select several responses.
Source: (Bugcrowd, 2019[132])

Bug bounty policy

To carry out a BBP, the vulnerability owner publicly provides a bug bounty policy with all the details about the programme. The policy sets the contractual rules governing the programme. It details the vulnerability owner's expectations regarding the behaviour of researchers engaged in the programme and the vulnerability owner's commitments to researchers. Such rules typically include the:

- Programme's technical scope, including accepted and non-eligible vulnerabilities, list of targeted products or servers, commitment to maintain system integrity, to minimise risk and harm to users, etc,
- Eligibility guidelines, e.g. originality, novelty, ineligibility of social engineering and automatic tools, denial of service attacks, phishing or malware attacks on employees, physical attacks on people, building and devices, lateral movement after a compromise of a system, etc.
- Reward guidelines, including nature and value of rewards according to the types of discovered vulnerabilities,
- Reporting guidelines, e.g. template report, proof of concept requirements,
- Conditions for and modalities of public disclosure,
- Conditions of communication between the researcher and the vulnerability owner,
- Legal terms, including protection of the researchers against legal proceedings (see 2.2) (Elazari, 2018[78]).

A bug bounty policy is similar to a VDP, although for BBPs the reward section is mandatory and the overall nature of the relationship is explicitly of a contractual nature. Bug bounty policies are also differentiators on the vulnerability market where all vulnerability owners compete for researchers' time and skills. In this context, rewards are one among several factors researchers take into account to decide which product they will focus on. Other factors may include, for example, the knowledge and experience they can get from the relationship, and the responsiveness of the vulnerability owner. According to Laszka et al. (2018[134]), effective bug bounty policies ensure the alignment of vulnerability owners' and researchers' interests when they explain *i)* how duplicates are addressed, *ii)* the conditions under which the vulnerability owner may or may not bring a lawsuit against the researcher, and *iii)* the conditions for public disclosure. Their analysis of 111 policies found that only 51 of them contained at least one of these conditions, and that only 10 covered the three, which seems to indicate a major issue of misalignment of incentives in this area. They also assessed these policies' readability and concluded that improvement could be undertaken to make these documents more approachable.

Value and effectiveness

BBPs bridge the gap between discoverers and vulnerability owners, structuring the CVD process. They offer benefits for both. Vulnerability owners can increase the likelihood of finding vulnerabilities. They are also less likely to experience unexpected and often costly full disclosures. Researchers have a higher likelihood of reward for their efforts. BBPs can also contribute to drain the black market by providing an alternative option to researchers motivated by monetary gains (Finifter, Akhawe and Wagner, 2013[135]), provided that they are established only under the authority of vulnerability owners, as opposed to grey market brokers.

The economic viability of BBPs has been a matter of debate since 1995. Microsoft, for example, argued against their effectiveness in terms of return on investment, prior to launching its own programme in 2013. The same year, an empirical study of Google and Mozilla BBPs suggested that they are more cost-effective in finding vulnerabilities than hiring full-time security researchers (Finifter, Akhawe and Wagner, 2013[135]).

The budget allocated to pay researchers is only one aspect of the vulnerability owner's cost equation. It also includes, for example, the management of the relationships with researchers and the legal advice for designing the programme and resolving disputes that could potentially arise (see 2.2). Vulnerability owners should also take into account operational costs, such as the need to manage the signal-to-noise ratio in the reporting, in particular when the total number of submissions is very high (e.g. 12 000 submissions to Facebook's programme in 2017). In average, invalid reports account for 35% to 55% of submissions across different platforms. For example, a BBP that was run by Uber through a bug bounty platform received 2030 reports, with a 1:6 signal to noise ratio (451 Research and HackerOne, 2017[47]).

Invalid reports can have different causes. Since researchers are interested in receiving as many bounties as possible while minimising efforts, they may for example use outputs from automated vulnerability scanners without spending enough time to analyse them, sending false positives to vulnerability owners. They may also not pay sufficient attention to bug bounty policies and submit out-of-scope reports. (Zhao, Laszka and Grossklags, 2017[136]; Laszka, Zhao and Grossklags, 2016[137]).

Another challenge is the high probability of duplicates, i.e. multiple reports by different researchers concerning the same vulnerability. For example, the number of duplicates has been higher than the number of valid reports for both Google's BBP and on the Bugcrowd platform (Laszka et al., 2018[134]).

Vulnerability owners need sufficient in-house security skills and resources to manage their BBP. As observed by security expert Katie Moussouris, a bug bounty can lead to an unpleasant experience if it is approached by the vulnerability owner "like going to an all-you-can-eat buffet without a working digestive system"! (Nichols, 2019[138])

Debates on the effectiveness of BBPs as a means to increase security are ongoing (Muncaster, 2019[139]; Trail of Bits Blog, 2019[140]). The recent steep increase in BBPs has led some observers to underline that they might be "a little overhyped" and should not be viewed as a digital security panacea (Heckman, 2019[141]; Nichols, 2019[138]). Some experts have highlighted that BBPs can also create challenges. For example, some companies might require that security researchers do not disclose vulnerabilities to other parties as a condition to receiving the bounty. This could limit the dissemination of the vulnerability information, potentially leaving stakeholders vulnerable, in particular when the vulnerability has a systemic nature (e.g. can affect multiple systems, products, implementations, sectors, etc.). It could also lead to a bidding escalation. Companies operating BBPs may also undermine their reputation if they refuse to pay when the vulnerabilities are not new to them or concern another partner in their supply chain.

Maturity and resource readiness are key conditions for organisations to create a BBP. Prior to launching a BBP, and in addition to having sufficient financial and human resources, vulnerability owners need to be already able to handle or manage vulnerabilities in a rapid, consistent, effective and predictable manner. They should operate a systematic, documented and tested process based on well-recognised international standards when they engage in crowdsourcing, which requires time and experience to reach a sufficient level of maturity.

Nevertheless, there seems to be a broad recognition that crowdsourcing vulnerability discovery can be useful as part of a sufficiently prepared overall product secure development lifecycle, in addition to other good practices such as internal code audits and penetration tests. Some experts highlighted that a BBP attracts new researchers with fresh ideas, thinking outside of the box, when the results of repetitive penetration tests can tend to decrease because they apply the same methodologies or involve the same experts. BBP can also attract security experts working for a product's client willing to better apprehend and assess a product's security without involving their company. BBPs can motivate internal staff to think about security, and keep a high-level of awareness that an error could cost the company time and money. Last, but not least, they can act as a recruitment channel, facilitating the identification of talent by vulnerability owners and helping researchers to select the companies they fine most in line with their expectations.

However, vulnerability owners should not approach BBPs as a silver bullet (ENISA, 2018, p. 63[49]) but rather as one tool among many others that they should consider, and use it in conjunction with others rather than as a turnkey security solution. BBPs will not replace but are complementary to software code reviews, for instance. Code reviews can also reveal design issues, or bad practices, that would lead to future vulnerabilities if not detected and corrected. BBPs is a reactive measure and, alone, is unlikely to improve the underlying design security limitations in a product or product line. If the budget allocated to a BBP reduces the resources that could be allocated to preventative security measures (e.g. security by design approach), then BBPs might be counterproductive in the medium to long term.

2.3.4. Bug bounty platforms

The OECD defines platforms as "online entities that serve at least two different sets of users simultaneously, bringing them together and enabling interactions between them that can benefit the users as well as the platform itself" (OECD, 2019[142]). The first bug bounty platform operating as a marketplace intermediary serving vulnerability owners and researchers appeared in 2005 with the launch of the Zero Day Initiative (ZDI) by the antivirus company Trend Micro. In 2010, a former Baidu employee launched Wooyun in China, which gathered about 20 000 researchers in 2016 when Wooyun's founder was arrested and the platform shut down (Chin, 2016[143]; Cao, 2016[144]). In the United States, HackerOne and Bugcrowd were created in 2012, and Synack in 2013. In the European Union, YesWeHack was launched in 2013 and Intigriti in 2017. Open Bug Bounty was created in 2014.

Bug bounty platforms offer services for vulnerability owners and security researchers. Vulnerability owners can use a web interface to rapidly design a VDP or bug bounty policy and publish it on the platform. Bug bounty programmes can be public or private, i.e. by invitation only. Researchers are invited to private bug bounty programmes according to their skills, which are known from their reporting track record on the platform. Researchers send reports through the platform's web site. For simple VDP services, the platform's team does not process the reports and the vulnerability owner can simply use the web interface to manage communications with the researchers. For bug bounty services, the platform's team triages reports to determine their validity and severity. The vulnerability owner then processes them and pays the reward to researchers who met the bug bounty criteria. Security researchers access a dedicated interface where they can search for ongoing programmes, submit their report, receive rewards as well as get visibility and reputation credits through a "hall of fame" or a credit system. The platform also serves as a communication channel between the participants.

Several platforms leverage their access to a large community of profiled security researchers to offer additional services such as penetration tests, or attack surface management where researchers are incentivised to find forgotten or missed assets. Some platforms use AI-based software and offer services such as vulnerability assessment, red teaming and standards compliance checks.

Most platforms are for-profit entities. Exceptions include Open Bug Bounty, a not-for-profit platform focusing on web site vulnerabilities and managed by a group of independent security researchers since 2014. As of April 2020, the platform handled 538 300 co-ordinated disclosures from 15 141 researchers, helped fix 291 130 vulnerabilities through 747 bug bounties covering 1 467 web sites. Another exception is Trend Micro's Zero Day Initiative (ZDI). The security software company created the ZDI in 2005 to feed its security products with the most up-to-date filters. The ZDI receives code vulnerability reports from researchers, contacts the code owner and, simultaneously distributes filters to Trend Micro's customers. Then it works collaboratively with the code owner to notify the vulnerability through a joint advisory.

2.3.5. Co-ordinators

To facilitate co-ordination (1.2.6), stakeholders can turn to a co-ordinator, namely a trusted third party such as a Computer Emergency Response Teams (CERTs) or Computer Security Incident Response Teams (CSIRTs). A co-ordinator can assist in a variety of cases, from easing the researcher-vulnerability owner relationship, to orchestrating complex multi-party co-ordination, as in the case of Spectre. Co-ordinators can also facilitate relationships between stakeholders across borders. Over time, CERTs and CSIRTs have established trusted relationships, including through FIRST, the international Forum of Incident Response Teams.

Some CERT and CSIRTs offer vulnerability co-ordination services as part of their core mission, while others do it on a case-by-case basis. CERT/CC was the first entity offering co-ordination services. It was created as the CERT division of the Software Engineering Institute (SEI), a non-profit, public–private partnership that conducts research for the United States government at the United States' Carnegie Mellon University. Other co-ordinators include government bodies such as the Dutch NCSC-NL, the French ANSSI, the Latvian CERT.LV, and the United States' CISA at the Department of Homeland Security (DHS). In France, the law enables ANSSI to receive a researchers' report, submit it to the vulnerability owner without revealing the identity of the researcher, and put pressure on the vulnerability owner if necessary.

Trust in the co-ordinator has multiple facets: technical competence, neutral judgement, strict respect of confidentiality, ability and capacity to understand and make a balanced assessment of the reasonableness of the various parties' claims and demands, such as severity of a vulnerability or timelines for public disclosure, possibility to interact with trusted stakeholders across borders. Standing at the centre of the relationship, the co-ordinator may receive confidential information from all parties and facilitate mutual understanding without such information being shared between the parties. Parties need to have a high degree of confidence that the co-ordinator will preserve the confidentiality of vulnerability and other sensitive information. In particular, when the co-ordinator is a government agency, parties need to trust that confidential information will not reach other parts of the government who could weaponise it for offensive use.

Trust can become very challenging when vulnerability information needs to cross borders, in particular when governments are involved. Experts have highlighted cases of international co-ordination where governments have leaked information to the press, or to third parties with poor security practices. Some experts suggest that an international, not-for-profit and well-resourced vulnerability co-ordinator should be established to address such issues.

2.3.6. Other initiatives

Different from BBP, some audits of open source software have been undertaken such as the Open Crypto Audit project, which assessed the security level of Truecrypt.[24]

Security researchers, in particular businesses carrying out vulnerability research, can also adopt a policy whereby they commit to following a set of good practice. For example, the security company Kaspersky published a set of "Ethical Principles in Responsible Vulnerability Disclosure", which include: 1. Build trust, 2. Inform the affected party first, 3. Co-ordinate efforts, 4. Maintain confidentiality where appropriate, 5. Incentivise desired behavior (2020[145]).

2.4. Good practice for CVD

This section brings together existing good practices for CVD from various documents listed in Annex 1. It aims to inform public policy makers. Technical experts are invited to consult appropriate and up-to-date technical standards.

2.4.1. Common understanding

All stakeholders should share the following basic common understanding:

- All products that contain code also contain vulnerabilities; all systems have a high likelihood of containing vulnerabilities related to misconfiguration or unpatched software, including firmware.
- These vulnerabilities represent a danger because threats actors can exploit them and create damages for all stakeholders, and, in some cases, for the economy and society as a whole.
- Not all vulnerabilities can be eliminated; however, it is possible to mitigate many of them to reduce digital security risk and the potential for harm, especially those that pose the greatest risk.
- Co-ordinated Vulnerability Disclosure (CVD) is a process whereby stakeholders who own the responsibility to eliminate vulnerabilities in products or systems (vulnerability owners) and security researchers who have found a vulnerability in these products or systems combine efforts and work collaboratively towards the common goal of increasing security of (or reducing digital security risk to) all stakeholders.
- There is no one-size-fits-all in vulnerability reporting and disclosure. Stakeholders should agree to follow good practice and accepted policies while recognising that they reflect intended paths for general cases and may not be optimal in all circumstances. Therefore, they should work co-operatively both to address each situation according to good practice, and to determine the best approach when good practice is not the best solution to reduce risk in specific cases.
- Effective vulnerability disclosure is as much a matter of trust between humans as a technical challenge.

2.4.2. Taking responsibility

Vulnerability owners and security researchers should work together to ensure the swift treatment of vulnerabilities and sharing of information with other stakeholders for the common objective of reducing digital security risk.

Vulnerability owners

Vulnerability owners should:

- Be responsible for the security of the system they operate or product they developed, and address related vulnerabilities according to the risk they raise to themselves, users, third parties and the economy and society as a whole;
- Be prepared to receive and address unsolicited vulnerability reports as part of their normal duty of care and responsibility;
- Adopt a public Vulnerability Disclosure Policy (VDP).

Code owners should handle any known vulnerabilities as part of the basic support of their products, and secure product development lifecycle.

System owners should maintain systematic vulnerability management cycles to ensure that configuration errors are corrected swiftly and that mitigations as well as security updates are applied as quickly as possible after their release, while managing the business and technical risk inherent to vulnerability management.

Vulnerability owners that have the capacity to process more than occasional reports and understand the related potential costs and benefits of doing so should **adopt CVD as a standard component of their digital security framework,** i.e. Security Development Lifecycle for code owners and digital security risk management policy for system owners. In this case, their VDP should express the organisation's willingness to receive vulnerability reports, commitment to co-ordinated vulnerability disclosure, and related conditions.

Researchers

Researchers should be responsible for their own actions, including for the way in which they discover a vulnerability and disclose it. They should not do more than what is necessary to demonstrate a vulnerability. They should:

- Report the vulnerability to the vulnerability owner first and as soon as possible after its discovery.
- Provide clear documentation and artefacts to the vulnerability owner to support verification processes.
- Contact a co-ordinator if the vulnerability owner cannot be reached or the process is not satisfactory.
- Respect the conditions set by a vulnerability owner in its vulnerability disclosure policy (VDP), and, in absence of a VDP, follow good practice for co-ordinated vulnerability disclosure.
- Use sufficiently secure means of communication to communicate about discovered vulnerabilities.
- Not require a reward as a condition to report a vulnerability. The initiative for granting a reward should lay with the vulnerability owner.

2.4.3. Creating sustainable trust

All CVD Stakeholders should build and maintain trust, including by:

- **Presuming benevolence**, good intent and good will from other CVD stakeholders.
- **Clearly communicating intentions**, and making a good faith effort to understand respective expectations and perspectives.
- **Maintaining continual or frequent communication** characterised by quality, mutual respect, patience and transparency, and using sufficiently secure communication channels and handling of sensitive information.
- **Being transparent** about expected processes and milestones (timelines), including the remediation and disclosure process.
- **Reducing uncertainty, surprise** and potential for dissatisfaction in other CVD stakeholders.

- **Negotiating expectations and timelines** if standard processes are not appropriate.
- **Avoiding legal or other coercive pressure or threat**, actual or perceived, as well as escalation, including legal action, to any extent possible, to prevent chilling effect on desired security research.
- **Leveraging a co-ordinator** as appropriate in case of dissatisfaction of one or more parties.

2.4.4. Adopting a vulnerability disclosure policy

At the minimum, a VDP should contain a point of contact for researchers to report vulnerabilities securely. The most basic method of receiving security reports is to have and monitor an email address at security@company.com.

A VDP should:

- Be public and use plain, easily understood terms, without jargon or ambiguous language.
- Capture the organisation's intent accurately and unambiguously.
- Describe authorised and unauthorised conduct.
- Explain the consequences of complying and not complying with the policy, including legal protections offered to compliant researchers.
- Encourage participants to contact the organisation for clarification before engaging in conduct that may be inconsistent with or unaddressed by the policy.
- Explain how reported vulnerabilities will be processed.
- Not evolve frequently, and track, explain and document changes made.

A VDP should also:

- **Create a safe harbour** for compliant security researchers.
- **Define the scope** of the vulnerability disclosure programme, which should be proportionate to the vulnerability owner's capacity to effectively process reports.
- **Indicate clear modalities** for secure and possibly anonymous communication.
- **Clarify expectations** with respect to acknowledgments and rewards (as appropriate), timelines, response times and follow-up communications during the process, confidentiality, and for code vulnerabilities: expectations related to the development of a remediation, and public disclosure.
- **Highlight the possibility to contact a co-ordinator** to facilitate the process.

2.4.5. Establishing appropriate internal processes for CVD

Vulnerability owners should approach CVD as a complement rather than as a substitute to or replacement for other security measures such as internally driven security testing.

When engaging in CVD, vulnerability owners should:

- **Co-ordinate internally** with the business, legal and communications teams and integrate CVD as part of business decision making processes rather than keeping it as an isolated technical process.
- **Allocate sufficient internal resources and define appropriate governance** to handle vulnerability analysis and communications tasks.
- **Consider third parties' interests**, for example by excluding from the scope of the policy any components or data implicating third-party interests or seek the authorisation of the third parties before including them in the policy.
- **Establish a strong foundation of tested processes and relationships,** following existing international standards, guidance documents and best practice, to ensure predictable response and relationships with researchers and third parties, to operate a clear, publicly known, regularly

monitored, and adequately secure intake mechanism, and appropriate communication channels with researchers.

- **Using automated tools and technical standards**, where appropriate, for example to facilitate exchange of information with third parties.
- **Anticipate challenges**, for example, by deciding, in advance of launching the CVD programme, how it will handle accidental, good faith violations of the VDP, as well as intentional, malicious violations.
- **Progressively scale their vulnerability disclosure programme** according to their learning curve, maturity and internal capacity to process reports.
- **Establish a cycle of improvement**, by capturing lessons learned from vulnerability reports to enable improvement of their overall security practices, including the CVD process itself.

2.4.6. Leveraging a co-ordinator

Stakeholders who face difficulties in establishing or carrying out a co-ordinated vulnerability disclosure process should seek assistance from a trusted third-party co-ordinator. The co-ordinator can help connect stakeholders, provide additional technical analysis and other support, particularly when there is disagreement among the parties. It can also help address cross-border challenges. Some co-ordinators can also help sharing knowledge about the vulnerability with the technical security community.

3. Possible Public Policy Guidance

This chapter provides a list of possible objectives for public policy makers willing to encourage vulnerability treatment, based on good practice such as those introduced in 2.4. As stakeholders involved in CVD are often located in different countries, such guidance may help reduce potential fragmentation of approaches across jurisdictions and contribute to globally reducing digital security risk.

Public policy development and implementation in this area should leverage all stakeholders' communities. In addition to government agencies in charge of digital security policy making and potentially other public bodies, communities include businesses and researchers:

- developing products,
- operating information systems,
- offering digital security solutions and services,
- working in academia,
- researching vulnerabilities on their spare time, with a professional digital security background or not;
- bug bounty operators and platforms;
- lawyers.

Civil society can also play an important role. For example, it can help bridge some parts of the security researchers' community with the government, as illustrated in Box 5. Good trusted relationships between the government and a sufficiently organised technical, business and civil society community can greatly facilitate policy making, including the assessment of existing frameworks and identification of the easiest and most effective path to policy improvement.

Box 5. A case of co-regulation with civil society: Karisma Foundation (Colombia)

For the past 4 years, the Karisma Foundation, a Colombian civil society organisation has been analysing government web sites to evaluate the information they provide to citizens, their level of digital security, and how they protect privacy. The purpose of these analyses is to contribute to the improvement of websites to benefit both the citizens and the entities responsible for these sites. As part of this project, Karisma has conducted audits of 4 government websites and found vulnerabilities in each of them. In 2017, Karisma submitted the report to the Ministry of ICT and the Unit for Integral Attention and Reparation to Victims (UARIV). The Governments reacted positively and the report was used as a basis for an implementation plan to improve the site's digital security. The exercise was publicised during the National Digital Security Forum organised by the Government in 2017 as an example of collaboration and co-responsibility between civil society and the Government, which was deployed as a good practice under international standards (Karisma Foundation, 2017[146]). The development of the new digital security strategy provided the momentum for the government to review Karisma's analysis of the state of the art in the matter in Colombia (Labarthe, 2019[147]) and to consider including specific actions to enable responsible disclosure in the country and ensure response and due diligence in responding to vulnerabilities.

Source: Karisma Foundation; Civil Society Internet Society Advisory Council (CSISAC)

The suggestions below are interrelated. They are not provided by order of importance or priority. Table 2 provides an overview of possible policy guidance discussed in this chapter.

Table 2. Overview of possible policy guidance

<table>
<tr><td><b>Sharing a common understanding</b>
• Changing the culture and raising awareness
• Clarifying roles and responsibilities</td><td><b>Mainstreaming good practice</b>
• Leading by example
• Including vulnerability treatment in regulation and guidance
• Providing tools, encouraging standards development & adoption</td></tr>
<tr><td colspan="2"><b>Fostering trust and removing obstacles</b>
• Ensuring access to a vulnerability co-ordinator
• Protecting researchers
• Addressing the grey market
• Encouraging international co-operation</td></tr>
</table>

3.1. Sharing a common understanding

3.1.1. Changing the culture and raising awareness

The "vulnerability taboo" and other digital security basics

Many organisations view digital security as an ideal state where there are no vulnerabilities. They make marketing claims that their products or systems are "safe and secure"; therefore, customers and users should trust them. With this mindset, these organisations are unlikely to welcome individuals who find vulnerabilities in their products or systems, share information with them or disclose it to the public. Rather, they view vulnerability disclosure as a threat to their marketing claims, reputation, and customers and partners' trust. In the worst cases, which are not rare, they believe in their own marketing claims to the point of not handling or managing vulnerabilities. It is not surprising that these organisations try to protect themselves by threatening researchers who attempt to interact with them. Such a mindset is particularly common in organisations with low digital maturity, which are currently accelerating the digital transformation of their business. New IoT manufacturers who approach digital transformation with concepts borrowed from the physical world often fall in this category. However, this is not a sustainable approach in an increasingly digitalised economy.

To successfully embrace digital transformation, the business leadership of these organisations needs to adopt a culture that recognises the three fundamental aspects of digital security:

- *All products that contain code also contain vulnerabilities*, including in firmware, and all information systems have a high likelihood of containing vulnerabilities related to misconfiguration or unpatched software. It is illusory to believe that digital products or systems can be perfectly secure. Breaking this "vulnerability taboo" is an essential first step prior to adopting strategies that can mitigate this challenge.
- *Digital security is a continuous effort to manage risk* rather than a state that is reached once and for all. This is a consequence of the fundamentally dynamic nature of the digital environment. Digital technologies, networks, data, systems, vulnerabilities, threats, incidents, etc. are all continuously changing. Nothing is static. The digital environment is also open by default and closed by exception. It is considerably easier to attack an information system or a smart product than to protect it because information systems are only as strong as their weakest component. Therefore, security must be agile and based on systematic processes.

- *Digital security is an economic and social risk management challenge related to a technical issue* rather than only a technical problem. The consequences of security incidents are economic and social. They affect revenues, competitiveness, business operations, reputation, innovation, and trust, by breaching the availability, integrity, and confidentiality of data, systems and networks. With the IoT, they can also affect safety. Therefore, organisations' business leaders need to own the risk. As risk owners, they need to work with technical security experts to address the risk rather than simply delegating responsibility to them.

Leaders and decision makers' approach to the way they can build trust with their customers and partners needs to change in the digital era. Rather than clinging to the utopian vision of a perfectly secure digital environment, they should recognise that their products and information systems can be vulnerable, and demonstrate that they take responsibility for swiftly addressing vulnerabilities when they become aware of them.

However, this cultural challenge is not limited to vulnerability owners. If the demand side continues to believe in unrealistic "safe and secure" claims, which supplier will dare break the taboo, challenge the common wisdom, and recognise that "safe and secure" is a promise nobody can keep?

Policy makers' role to help change the culture

Policy makers, together with other stakeholders, have a key role to play to tell the uncomfortable truth and explain the direction to take. They can also help change how security researchers are perceived, and raise awareness about their contribution to our collective security and privacy.

Governments need to reach out to the security researchers' community to understand and take into account their point of view. However, in most countries, there are few opinion leaders among security researchers who can speak to governments, businesses and the media, and contribute to educating the society about digital security. Security researchers would benefit from forming a more organised community, including by developing a public discourse that can reach and influence policy makers. The voice of security researchers need to be heard, including at the international level.

Like all cultural changes, such an evolution of mindset will take time. However, digital transformation unfolds at a fast pace and the consequences of slow action or inaction will become increasingly severe.

All stakeholders need to play their part in breaking the vulnerability taboo, from firms with high brand reputation who already demonstrate digital security excellence, to opinion leaders in the security researchers' community, to civil society and consumer organisations. Initiatives such as the Cybersecurity Tech Accord (cf. Box 6) and the Paris Call for Trust and Security in Cyberspace are steps in the right direction (French Ministry for Europe and Foreign Affairs, 2018[148]). The organisation of BBPs and public recognition of vulnerabilities by prestigious but rather traditionally conservative public institutions contributes greatly to changing mindsets in the society. For example, following one of its BBPs, the US Department of Defense publicly disclosed details about four critical vulnerabilities on its infrastructure in September 2020 (Ilascu, 2020[149]).

Policy makers can promote the fundamental concepts of digital security to new entrants and industry players who are likely to lack appropriate digital security culture. For example, NIST recently published a report targeting IoT manufacturers which describes recommended digital security activities that they should consider performing before their IoT devices are sold to customers (Fagan et al., 2020[150]).

Policy makers need to promote the idea that the digital security of products should be an iterative process that can benefit from multi-stakeholder co-operation, such as co-ordinated vulnerability disclosures and the appropriate handling or management of vulnerabilities. Public policy should reward organisations that adopt a transparent vulnerability disclosure policy, encourage vulnerability discovery and reporting, provide a safe harbour for researchers, and engage in an effective CVD process resulting in swift resolution and disclosure. Awareness-raising campaigns, labels and public procurement can be effective tools to achieve these goals (see (OECD, 2021[3]) and (OECD, 2021[4])).

Box 6. Cybersecurity Tech Accord's commitment to vulnerability handling

All stakeholders can contribute to mainstreaming good practice for vulnerability disclosure. For example, the Cybersecurity Tech Accord is a group of 144 ICT companies (as of July 2020) who collaborate on initiatives that improve the security, stability and resilience of cyberspace. The Accord's first principle commits its signatories to design, develop, and deliver products and services that prioritise security, privacy, integrity and reliability. In upholding this principle, the group has promoted, from the outset, the adoption of vulnerability disclosure policies by companies throughout the technology industry. In 2019, they committed to having every signatory work towards putting their own vulnerability disclosure policy in place. This commitment by all members of the largest coalition of global technology firms dedicated to improving the cybersecurity ecosystem, is a significant step forward.

As of July 2020, over 80 signatories had a vulnerability disclosure policy in place.

Source: Cybersecurity Tech Accord, https://cybertechaccord.org.

3.1.2. Clarifying roles and responsibilities to treat vulnerabilities

The security community has long debated "responsible disclosure", probably because disclosure is a particularly sensitive stage in the vulnerability lifecycle, where stakeholders need to make important decisions. However, this framing of the issue places the attention on security researchers rather than on the responsibility of vulnerability owners. It also ignores the complex role of governments, underlined in **Error! Reference source not found.**.

"Responsible disclosure" debates assume that vulnerability owners have an effective vulnerability management or handling process in place whereas it is often not the case. Therefore, policy makers need to take a broader approach to how all stakeholders "treat" vulnerabilities rather than focusing primarily on disclosure. It is proposed to call "vulnerability treatment" the overarching area covering vulnerability discovery, handling, management and disclosure (1.1.6).

Building upon the Responsibility principle of the 2015 Recommendation on Digital Security Risk Management, all stakeholders involved in vulnerability treatment should take responsibility and be accountable, based on their role, ability to act and the context, for treating vulnerabilities and sharing information in a timely manner, and for taking into account the potential impact of their decisions on others.

Government agencies which may receive vulnerability information, for example in a co-ordinator capacity or as a regulatory body, should provide assurance that vulnerability information will only be shared with or accessed by the vulnerability owner (i.e. who can fix the vulnerability) and not with any other party, including those who could stockpile or use it for offensive purposes. In some countries, this may require adjusting the public governance for digital security, in particular when considering mandatory reporting of vulnerabilities to the government.

In line with the above good practice (2.4) organisations need to take (i.e. own) responsibility for security vulnerabilities in products they put on the market and information systems they manage, and security researchers should take responsibility for their actions.

In practice, it means that code owners need to adopt a security development lifecycle whereby vulnerability handling is an integral part of basic product maintenance and support; and that system owners adopt a digital security risk management framework including a vulnerability management process covering all digital assets to detect and address misconfigured and unpatched software and devices. Vulnerability management cycles should be sufficiently rapid to minimise the window of exposure. There might be a need to promote collectively recognised acceptable timelines.

In addition, all stakeholders should aim at disclosing vulnerabilities in a co-ordinated manner. All vulnerability owners should have a public VDP and maintain a vulnerability disclosure process to receive and address vulnerabilities spontaneously reported to them (cf. 2.4). According to their capacity and appetite, they may encourage security researchers to search vulnerabilities in their products and systems through a more detailed VDP and, potentially, a bug bounty programme. Vulnerability owners should also establish a continuous cycle of improvement by feeding their vulnerability handling or management respectively into their security development lifecycle and security risk management framework.

3.2. Mainstreaming good practice

3.2.1. Leading by example

Governments can play a key role in encouraging the adoption of CVD and promoting a cultural shift with respect to vulnerability treatment.

Governments can lead by example, for example by:

- *Adopting CVD within the government.* This would however require that agencies have appropriate capacity, funding, and resources necessary to receive and analyse disclosures, mitigate vulnerabilities, and manage communications with stakeholders. The draft Binding Operational Directive applicable by US Government federal, executive branch, departments and agencies, to develop and publish a vulnerability disclosure policy is an example of such an approach (DHS, 2019[104]).
- *Adopting vulnerability handling and management within the government.*
- *Using public procurement to encourage vulnerability treatment*. Governments can for example include them as conditions for public procurement.

These initiatives would need to be based on a government-approved set of good practice standards or guidance (see below).

3.2.2. Including vulnerability treatment in regulation and guidance

Governments can also promote vulnerability treatment by including vulnerability management, handling and CVD in regulation, standards and guidance, or using them as indicators of compliance with other regulations. This may include, for example:

- Product regulation;
- Regulation related to critical activities, such as the EU NIS Directive. This would follow the OECD Recommendation on Digital Security of Critical Activities, which recommends that governments "build capacity to support digital security risk management and resilience of critical activities by […] adopting and encouraging the adoption of responsible and co-ordinated vulnerability disclosure and management processes, as well as encouraging and protecting security researchers" (OECD, 2019[151]). At the time of writing, a public consultation for the review the NIS Directive is ongoing and includes questions about vulnerability disclosure.
- Certification schemes, such as those established by the EU Cybersecurity Act (cf. 2.3.1);
- Government-supported standards, e.g. as it is the case in the NIST Cybersecurity Framework 1.1, and ETSI Technical Specification "Cyber Security for Consumer Internet of Things".

- IoT regulation. For example, the UK is preparing draft regulation on IoT requiring that all companies providing internet-connected devices and services provide a public point of contact as part of a vulnerability disclosure policy. Device manufacturers, IoT device providers and mobile application developers would also be required to continually monitor for, identify and rectify security vulnerabilities within their own products and services as part of the product security lifecycle.
- Privacy regulation, such as the EU General Data Protection Regulation (GDPR) and the US Health Insurance Portability and Accountability Act (HIPAA), for which good vulnerability management and CVD practices could be indicators of compliance.

When considering regulation, it is important to ensure that new measures are both aligned with existing international standards and good practice, although sufficiently high-level and flexible to accommodate their possible future evolution.

Many experts agree that it may be useful to report to the government vulnerabilities in products and systems supporting critical activities. However, there is a general preference for voluntary rather than mandatory reporting in this case. Experts agree that the most effective approach is to build a trusted relationship between the government and all stakeholders, based on the transparency of how the government uses vulnerability information received from other parties.

As governments can play an ambiguous role in this area (cf. **Error! Reference source not found.**), they must build trust and demonstrate that vulnerability information they receive is handled appropriately. For example, experts who recommend mandatory reporting of vulnerabilities to a government entity strongly emphasise the need for the receiving entity to operate in a transparent manner, with the sole objective of remediating vulnerabilities, and independently from other government entities.

More generally, regulation that would place stakeholders in the position of being required by law to do something they believe could increase digital security risk is likely to discourage vulnerability research, and increase mistrust. A recent draft regulation by the Chinese government limiting public disclosure of vulnerabilities before they have been communicated to the authorities has raised this type of concerns (Yin, 2019[152]; Creemers and Webster, 2019[153]; Udemans, 2019[154]).

3.2.3. Providing tools

Policy makers, industry organisations, and other stakeholder groups can facilitate CVD adoption with template Vulnerability Disclosure Policies (VDPs), quick start guides, and other best practice documents. These can target specific communities to address their needs and concerns. Such documents can make it easier and cheaper for organisations to take the first critical steps towards a CVD-ready posture. The “early stage” US NTIA CVD template aimed at safety-critical industries (NTIA, 2016[155]) and the NCSC-NL Guidelines (2018[98]) provide examples of such initiatives by governments.

3.3. Fostering trust and removing obstacles

3.3.1. Ensuring access to a vulnerability co-ordinator function

Vulnerability owners and researchers should have the possibility to turn to a co-ordinator, i.e. a trusted third party who can facilitate vulnerability treatment when several stakeholders need to co-ordinate their action. This is typically the case in CVD, as well as in complex vulnerability handling and management scenarios.

Co-ordinators need to have enough resources to accomplish their task, which may be demanding in some cases. It is not necessary for every country to have at least one domestic co-ordinator. In many cases, the nationality of the co-ordinator does not matter, as long as it is trusted by the stakeholders (see below). For example, they can turn to a foreign, regional, or industry-led co-ordinator. To address issues of resources and trust, some experts have suggested that stakeholders explore the feasibility of establishing an international co-ordination function.

Co-ordinators may be public or private sector, general or sectoral, and domestic, regional or international bodies. CERTs often provide co-ordination services. Some can assist in co-ordination upon request without explicitly calling themselves co-ordinators. Several government CERTs act as a last resort co-ordinator, in particular when the vulnerability could affect critical activities as defined in (OECD, 2019[151]).

Nevertheless, a co-ordinator has to be trusted by all parties. To be trusted, it needs to (in no particular order):

- Have a high level of technical competence and expertise, in order to swiftly understand technical aspects at stake and the perspectives of the participants;
- Have well organised, predictable, and reliable processes, in particular with respect to the security, clarity and regularity of its communications with stakeholders;
- Provide assurance that vulnerability information will only be shared with or accessed by the vulnerability owner (i.e. who can fix the vulnerability) and not with any other party, including those who could use or stockpile it for offensive purposes;
- Have established and trusted relationships with other co-ordinators in third countries to overcome possible cross-border challenges, including practical (e.g. language), legal, political, or cultural;
- Respect the researcher's willingness to remain anonymous;
- Provide legal protection to the researcher (incl. by respecting its anonymity).

3.3.2. Protecting researchers

One role of the government is to build norms and institutions that promote co-ordinated disclosure by reducing the cost of entry and balancing the power dynamic between researchers and vulnerability owners. A key aspect of the power imbalance is the legal pressure that the latter can place on the former. This can be adjusted by changing the legal environment to better protect responsible security researchers and reduce the risk of lawsuits and criminal prosecution wherever it is an obstacle to CVD.

Legal regimes vary across countries and legal uncertainty generally affects several legal areas. Therefore, it is not possible to describe one-size-fits-all measures that would resolve this issue. For example, the approach taken by the Dutch government or the US safe harbour mechanisms may work in some countries but not in others (CEPS, 2018[64]). As shown in the case of Latvia, an agency willing to promote legal changes may face unexpected negative reactions from other parts of the government who may not really understand CVD and its public benefit, or they may pursue different agendas.

Therefore, governments need to take stock of legal risk for researchers in their jurisdiction, develop a plan to reduce it, and ensure that any new legislative or regulatory frameworks do not create new obstacles. This could take place as part of a broader strategic review where all intra- and extra-governmental stakeholders are involved, such as the revision of a national digital security strategy, or development of an implementation plan.

At the international level, governments could consider developing and agreeing upon a set of high-level criteria that would trigger adverse actions against a researcher or vulnerability owner that proved to be outside a generally held international legal understanding for appropriate behaviour.

Governments need to keep in mind that free trade agreements can also have undesired effects when exporting static legal frameworks to developing countries lacking the capacity to review and update them regularly (cf. 2.2.1).

3.3.3. Addressing the grey market

Governments should take action to address the grey market for code vulnerabilities, in order to prevent it from distorting prices, providing incentives for some researchers to keep vulnerabilities secret, and preventing vulnerability owners from developing mitigations and protecting users. However, more work might be needed to identify viable avenues to do so, in particular to overcome issues underlined in **Error! Reference source not found.**

As a first step, one option to drain this market is to ensure that bug bounties are only organised under the authority of the vulnerability owner, or a vendor whose product relies on the code covered by the bug bounty programme, or, in the case of open source products, by actors committing to addressing the vulnerability.

Another, more thought provoking idea, would be for governments to establish an international fund to systematically buy vulnerabilities, at least for open source software (cf. Box 2).

Further work would be needed to better understand the issue and identify other options to address it (cf. Annex 2). For example, such work could explore whether and under which conditions regulating this market could help address its negative effects on CVD. For example, governments could agree on a set of criteria for legitimate supply and demand-side actors.

3.3.4. Encouraging international co-operation and standards development

Public policy should encourage co-operation across borders to remove obstacles to and facilitate vulnerability treatment. In particular:

- Governments should work together to facilitate the exchange of vulnerability information across borders between security researchers and vulnerability owners. Governments should not create obstacles to such information exchanges.
- Public policy should encourage stakeholders' cross-border co-operation for the co-ordinated disclosure of highly sensitive vulnerability information that could affect critical activities. Governments and other stakeholders should work together to explore possible means to improve such cross-border co-ordination, for example by strengthening international collaboration between CERTs.
- Public policy should encourage all stakeholders to participate in the improvement of existing international standards (e.g. integration of security controls concerning vulnerability management in ISO/IEC 27000 and references to other ISO/IEC standards to better integrate CVD in vulnerability owners' information systems) and development of new ones (e.g. on multi-party vulnerabilities).

Glossary

This glossary provides simple explanations of terms for the purpose of this report. For technical and operational definitions, please refer to relevant standards documents.

Bug: error, flaw or fault in a computer program or system that causes it to produce an incorrect or unexpected result, or to behave in unintended ways. A bug is a vulnerability when it can be exploited by a threat.

Bug bounty programme (or “bug bounty”): crowdsourcing initiative that reward individuals for discovering and reporting vulnerabilities to the vulnerability owner.

Bug bounty platform: marketplace intermediary facilitating the relationship between vulnerability owners who launch bug bounty programmes and security researchers.

Code vulnerability: vulnerability affecting the code embedded in a product.

Code owner: individuals or organisations who developed the layer of code where a code vulnerability is located in a product or/and are best placed to fix it. They “own” the responsibility to address code vulnerabilities. They are often called “vendors” in the literature.

Co-ordinated vulnerability disclosure (CVD): process through which vulnerability owners and researchers work co-operatively in finding solutions that reduce the risks associated with a vulnerability.

Co-ordinator: stakeholders who can assist code and system owners as well as researchers in the vulnerability disclosure process.

Disclosure: publication or broad circulation of vulnerability information (different from reporting).

Exploit (or exploit code): code developed to weaponise a vulnerability.

Exploitation: use of an exploit against an information system.

Mitigation: solution to address a vulnerability, such as a patch, or a process to follow.

Patch: piece of code that modifies software to mitigate a vulnerability.

Red team or red teaming: a more advanced form of network penetration testing where a contracted or in-house red team (as opposed to the defending blue team) emulates an advanced threat actor using physical, digital, and human vectors to identify gaps in the organisation’s defensive strategy.

Remediation: cf. mitigation.

Reporting: communication of vulnerability information to the vulnerability owner (different from disclosure) or a co-ordinator.

Researcher (or security researcher): individuals or organisations who identify potential code or system vulnerabilities with the intention to reduce related security risk. They are sometimes also called “ethical hackers”, “white hat”, “finders”, and “discoverers”.

Security update: mechanism to distribute patches to users of vulnerable software.

System owners: organisations using products within their information system. They “own” the responsibility to configure these products and apply security updates provided by code owners.

System vulnerability: weakness in the way a product is implemented or configured (i.e. deficient vulnerability management and misconfiguration).

Vulnerability (or digital security vulnerability): weakness, bug or flaw that, if exploited, triggered, or activated by a threat, has the potential to cause economic and social damages, by affecting availability, integrity, or confidentiality of a digital resource or asset.

Vulnerability disclosure policy: information publicly provided by a vulnerability owner to explain how to securely report a vulnerability, what to expect upon reporting (incl. reward where appropriate), as well as legal conditions, implications, and protections.

Vulnerability handling: process followed by code owners to address code vulnerability information from its reception or discovery to the post-release (cf. ISO/IEC 30111).

Vulnerability management: process followed by an organisation to know if vulnerabilities are present within their digital environment and take appropriate risk management decisions and actions.

Vulnerability owner: stakeholders who own the responsibility to act upon a vulnerability they are aware of, in order to mitigate it.

Vulnerability treatment: policy area covering vulnerability discovery, handling, management and co-ordinated vulnerability disclosure.

Weaponisation: development of an exploit by using a vulnerability (the vulnerability is weaponised).

Window of exposure: period during which a stakeholder is exposed to digital security risk related to a vulnerability, beginning with the discovery of the vulnerability and ending with the application of the mitigation to the product or system.

Zero-day (or zero-day vulnerability): code vulnerability for which no mitigation has yet been released, or which is unknown to the code owner.

Zero-day exploit: exploit based on a zero-day.

Annex 1. List of good CVD practice documents

The following documents were used to develop the Good Practice section of the report (2.4).

1. Christey Steve, Wysopal Chris (2002), Responsible Vulnerability Disclosure Process.
2. US-CERT (2012), Common Industrial Control System Vulnerability Disclosure Framework.
3. ENISA (2016), Good Practice Guide on Vulnerability Disclosure. From challenges to recommendations.
4. Rabobank, CIO Platform Nederland (2016), Manifesto on Co-ordinated Responsibility Disclosure.
5. NTIA Safety Working Group (2016), “Early Stage” Co-ordinated Vulnerability Disclosure Template. Version 1.1.
6. Householder D., Wassermann G., Manion A., King C., (2017), The CERT® Guide to Co-ordinated Vulnerability Disclosure. CMU/SEI-2017-SR-022.
7. FIRST (2017), Guidelines and Practices for Multi-Party Vulnerability Co-ordination and Disclosure, v 1.0.
8. US Department of Justice (2017), A Framework for a Vulnerability Disclosure Program for Online Systems, version 1.0.
9. Dutch National Cyber Security Centre (2018), Co-ordinated Vulnerability Disclosure: The Guideline.
10. Cybersecurity Coalition (2019), Policy Priorities for Co-ordinated Vulnerability Disclosure and Handling.
11. Center for Cybersecurity Policy and Law (2019), Improving Hardware Component Vulnerability Disclosure.
12. Business Software Alliance (BSA) (2019), Guiding Principles for Co-ordinated Vulnerability Disclosure.

Annex 2. Possible areas for future work

- Good practice on vulnerability management and handling
- Vulnerability disclosure in open source products
- Vulnerability management: why are so many systems never patched and what can we do about it?
- Improving multi-party CVD (e.g. co-ordination, international co-operation, open source products, etc.)
- Artificial Intelligence and vulnerabilities
- System vulnerabilities and IoT devices
- Challenges faced by small entities to address vulnerabilities (code and system owners) such as Small and Medium-sized Enterprises (SMEs) and local governments
- How to drain the grey market? How does this market work? What is its size and pricing mechanisms? Should governments regulate it and if so, how?
- Relationships between vulnerability disclosure and product liability insurance
- Addressing “failure” cases in the vulnerability disclosure lifecycle, e.g. when the product has reached end of life/support, when code owner will not fix the vulnerability or no longer exist, etc.
- Developing vulnerability-related metrics, e.g. how many actors use CVD, and manage/handle vulnerabilities
- Legal obstacles to digital security research in general, and vulnerability disclosure in particular, across countries
- Addressing obstacles to advanced sharing of vulnerability information with governments.

References

451 Research and HackerOne (2017), *Bug Bounties and the Path to Secure Software*, https://www.hackerone.com/sites/default/files/2017-06/451-pathfinder-report.pdf (accessed on 15 April 2020). [47]

AccessNow (2020), *Human rights organizations reject arbitrary measures against digital security researcher Javier Smaldone*, https://www.accessnow.org/human-rights-organizations-reject-arbitrary-measures-against-digital-security-researcher-javier-smaldone/. [94]

Adams, S. (2018), *Getting Better All the Time: Security Research and the DMCA*, https://cdt.org/insights/getting-better-all-the-time-security-research-and-the-dmca/. [81]

Andersen, T. (2017), *Politianmeldt af KMD for hacking: »Jeg er totalt uskyldig«*, https://www.version2.dk/artikel/interview-hacker-tiltalt-jeg-totalt-uskyldig-1077581 (accessed on 21 April 2020). [93]

Andreessen, M. (2011), *Marc Andreessen on Why Software Is Eating the World - WSJ*, https://www.wsj.com/articles/SB10001424053111903480904576512250915629460 (accessed on 11 May 2020). [162]

ANSSI (n.d.), *Alertes aux vulnérabilités et failles de sécurité*, https://www.ssi.gouv.fr/en-cas-dincident/vous-souhaitez-declarer-une-faille-de-securite-ou-une-vulnerabilite/ (accessed on 5 May 2020). [101]

Apple (2019), *Apple Security Bounty - Payouts - Apple Developer*, https://developer.apple.com/security-bounty/payouts/ (accessed on 31 January 2020). [56]

Bennetts, S. (2019), *Updates to the Mozilla Web Security Bounty Program - Mozilla Security Blog*, https://blog.mozilla.org/security/2019/11/19/updates-to-the-mozilla-web-security-bounty-program/ (accessed on 5 May 2020). [58]

Boyd, A. (2018), *DOD Invests $34 Million in Hack the Pentagon Expansion - Nextgov*, Nextgov, https://www.nextgov.com/cybersecurity/2018/10/dod-invests-34-million-hack-pentagon-expansion/152267/ (accessed on 9 April 2020). [120]

Bugcrowd (2019), *Inside the Mind of the Hacker 2019*, https://www.bugcrowd.com/resources/guides/inside-the-mind-of-a-hacker-2019/ (accessed on 15 April 2020). [132]

Canadian Centre for Cybersecurity (n.d.), *Cyber Threat and Cyber Threat Actors - Canadian Centre for Cyber Security*, https://cyber.gc.ca/en/guidance/cyber-threat-and-cyber-threat-actors (accessed on 5 May 2020). [34]

Cao, S. (2016), *Helpful hacker forums close after arrest for revealing vulnerabilities*, http://www.globaltimes.cn/content/1001271.shtml (accessed on 5 May 2020). [144]

Carnegie Mellon University (2016), *Vulnerability Management Version 1.1*, https://www.us-cert.gov/sites/default/files/c3vp/crr_resources_guides/CRR_Resource_Guide-VM.pdf (accessed on 8 April 2020). [159]

CCB-BE (2020), *Politique de divulgation coordonnée de vulnérabilités et programme de récompense pour la découverte de vulnérabilités*, https://ccb.belgium.be/fr/politique-de-divulgation-coordonn%C3%A9e-de-vuln%C3%A9rabilit%C3%A9s-et-programme-de-r%C3%A9compense-pour-la-d%C3%A9couverte. [99]

CDT (2018), *Taking the Pulse of Hacking. A risk Basis for Security Research.*, https://cdt.org/wp-content/uploads/2018/04/2018-03-27-Risk-Basis-for-Security-Research-FNL.pdf (accessed on 8 October 2019). [77]

Center for Cybersecurity Policy and Law (2019), *Improving Hardware Component Vulnerability Disclosure The Center for Cybersecurity Policy and Law*, https://static1.squarespace.com/static/5acbb666f407b432519ab15e/t/5cc86f37c830251f28d258fc/1556639544235/The+Center+for+Cybersecurity+Policy+and+Law_Improving+Hardware+Component+Vulnerability+Disclosure_April+2019.pdf (accessed on 3 February 2020). [65]

CEPS (2018), *Software Vulnerability Disclosure in Europe Software Vulnerability Disclosure in Europe*. [64]

CERT/CC (2017), *The CERT ® Guide to Coordinated Vulnerability Disclosure*, http://www.sei.cmu.edu (accessed on 4 February 2020). [43]

Charlet, K., S. Romanosky and B. Thompson (2017), "It's Time for the International Community to Get Serious about Vulnerability Equities - Lawfare", https://www.lawfareblog.com/its-time-international-community-get-serious-about-vulnerability-equities. [60]

Chin, J. (2016), *China's 'White-Hat' Hackers Fear Dark Times After Community Founder Is Detained - China Real Time Report - WSJ*, https://blogs.wsj.com/chinarealtime/2016/08/01/chinas-white-hat-hackers-fear-dark-times-after-community-founder-is-detained/ (accessed on 5 May 2020). [143]

Cimpanu, C. (2020), *MIT researchers disclose vulnerabilities in Voatz mobile voting election app*, https://www.zdnet.com/article/mit-researchers-disclose-vulnerabilities-in-voatz-mobile-voting-election-app/ (accessed on 5 May 2020). [97]

Cimpanu, C. (2018), "VirtualBox zero-day published by disgruntled researcher", *ZDNet*, https://www.zdnet.com/article/virtualbox-zero-day-published-by-disgruntled-researcher/ (accessed on 16 April 2020). [74]

Cimpanu, C. (2017), *Dark Web Marketplace Launches Bug Bounty Program with $10,000 Rewards*, https://www.bleepingcomputer.com/news/security/dark-web-marketplace-launches-bug-bounty-program-with-10-000-rewards/ (accessed on 5 May 2020). [130]

Cluley, G. (2014), *How a five-year-old hacked his dad's Xbox One, only to be rewarded by Microsoft [VIDEO] – HOTforSecurity*, https://hotforsecurity.bitdefender.com/blog/how-a-five-year-old-hacked-his-dads-xbox-one-only-to-be-rewarded-by-microsoft-video-8316.html (accessed on 5 May 2020). [32]

Council of Europe (2001), *Convention on Cybercrime*, https://www.coe.int/en/web/conventions/full-list/-/conventions/rms/0900001680081561 (accessed on 5 May 2020). [75]

Cox, J. (2019), *As Phones Get Harder to Hack, Zero Day Vendors Hunt for Router Exploits - VICE*, https://www.vice.com/en_us/article/evek9z/phones-harder-to-hack-crowdfense-zerodium-buy-router-zero-days-exploits (accessed on 31 January 2020). [55]

Creemers, R. and G. Webster (2019), *Translation: China's 'Cybersecurity Threat Information Publication Management Measures (Draft for Comment)'*, https://www.newamerica.org/cybersecurity-initiative/digichina/blog/translation-chinas-cybersecurity-threat-information-publication-management-measures-draft-comment/. [153]

Crowdfense (n.d.), *Bug Bounty program*, https://www.crowdfense.com/bug-bounty-program.html (accessed on 31 January 2020). [54]

CSDE (2019), *Cyber Crisis: Foundations of Multi-Stakeholder Coordination*, https://www.ustelecom.org/wp-content/uploads/2019/09/CSDE-Report-Cyber-Crisis-Foundations-of-Multi-Stakeholder-Coordination.pdf. [68]

Cybersecurity Coalition (2019), *Policy Priorities for Coordinated Vulnerability Disclosure and Handling*, https://www.cybersecuritycoalition.org/policy-priorities. [46]

Cyberspace Solarium Commission (2020), *Cyberspace Solarium Commission Report*, https://www.solarium.gov/. [23]

Dean, B. (2018), *Strict Products Liability and the Internet of Things*, Center for Democracy and Technology. [10]

DHS (2019), *Binding Operational Directive 20-01 - Develop and Publish a Vulnerability Disclosure Policy*, https://cyber.dhs.gov/bod/20-01/ (accessed on 8 April 2020). [104]

DHS and DoC (2018), *Report on "Enhancing the Resilience of the Internet and Communications Ecosystem Against Botnets"*, https://www.commerce.gov/page/report-president-enhancing-resilience-against-botnets. [9]

DINSIC (2019), *TChap bug bounty program - Yes We Hack*, https://yeswehack.com/programs/tchap (accessed on 5 May 2020). [124]

Disclose.io (2020), *disclose/program-list.csv at master · disclose/disclose · GitHub*, https://github.com/disclose/disclose/blob/master/program-list/program-list.csv (accessed on 9 April 2020). [117]

Doe, D. (2016), *FBI raids dental software researcher who discovered private patient data on public server | The Daily Dot*, https://www.dailydot.com/debug/justin-shafer-fbi-raid/ (accessed on 5 May 2020). [92]

Edison Group (2013), *Microsoft Security Development Lifecycle Adoption: Why and How*, https://download.microsoft.com/download/F/C/6/FC624A57-8E38-40D7-88B9-C3C256E9940A/Microsoft-Security-Development-Lifecycle-Adoption-Whitepaper.pdf?WT.z_evt=WhitePaperClick. [41]

EdOverflow (n.d.), *How do I write a good security policy?*, https://bugbountyguide.com/programs/writing-security-policies.html (accessed on 9 April 2020). [112]

EdOverflow and Y. Shafranovich (n.d.), *security.txt | A proposed standard which allows websites to define security policies*, https://securitytxt.org/ (accessed on 9 April 2020). [114]

EFF (2018), *Protecting Security Researchers' Rights in the Americas*, https://www.eff.org/wp/protecting-security-researchers-rights-americas#executive_summary (accessed on 8 October 2019). [29]

EFF (n.d.), *Coders' Rights Project Vulnerability Reporting FAQ | Electronic Frontier Foundation*, https://www.eff.org/issues/coders/vulnerability-reporting-faq (accessed on 10 February 2020). [85]

Elazari, A. (2018), "Private Ordering Shaping Cybersecurity Policy: The Case of Bug Bounties", https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3161758 (accessed on 7 February 2020). [78]

ENISA (2020), *Cybersecurity Certification: EUCC Candidate Scheme*, https://www.enisa.europa.eu/publications/cybersecurity-certification-eucc-candidate-scheme. [106]

ENISA (2020), *State of Vulnerabilities 2018/2019 - Analysis of Events in the life of Vulnerabilities*, https://www.enisa.europa.eu/publications/technical-reports-on-cybersecurity-situation-the-state-of-cyber-security-vulnerabilities/ (accessed on 8 April 2020). [19]

ENISA (2018), *Economics of Vulnerability Disclosure*, https://www.enisa.europa.eu/publications/economics-of-vulnerability-disclosure (accessed on 31 January 2020). [49]

ENISA (2016), *Good Practice Guide on Vulnerability Disclosure. From challenges to recommendations*, https://www.enisa.europa.eu/publications/vulnerability-disclosure (accessed on 25 January 2020). [35]

Etcovitch, D. and T. van der Merwe (2018), *COMING IN FROM THE COLD A SAFE HARBOR FROM THE CFAA AND THE DMCA §1201 FOR SECURITY RESEARCHERS*. [76]

Etcovitch, D. and T. van der Merwe (2017), "Coming in from the Cold: A Safe Harbor from the CFAA and the DMCA 1201 for Security Researchers", *SSRN Electronic Journal*, http://dx.doi.org/10.2139/ssrn.3055814. [88]

ETSI (2019), *TS 103 645 - V1.1.1 - CYBER; Cyber Security for Consumer Internet of Things*, https://portal.etsi.org/TB/ETSIDeliverableStatus.aspx (accessed on 22 April 2020). [110]

European Union (2019), *Regulation (EU) 2019/881 of the European Parliament and of the Council of 17 April 2019 on ENISA (the European Union Agency for Cybersecurity) and on information and communications technology cybersecurity certification and repealing Regulation (EU) No 526/2013 (Cybersecurity Act)*, https://eur-lex.europa.eu/eli/reg/2019/881/oj (accessed on 22 April 2020). [105]

Fagan, M. et al. (2020), *Foundational cybersecurity activities for IoT device manufacturers*, National Institute of Standards and Technology, Gaithersburg, MD, http://dx.doi.org/10.6028/nist.ir.8259. [150]

Fidler, M. (n.d.), *Regulating the Zero-Day Vulnerability Trade: a Preliminary Analysis*, https://fsi-live.s3.us-west-1.amazonaws.com/s3fs-public/fidler-second-review-changes-made.pdf (accessed on 31 January 2020). [48]

Finifter, M., D. Akhawe and D. Wagner (2013), "An Empirical Study of Vulnerability Rewards Programs An Empirical Study of Vulnerability Rewards Programs", https://www.usenix.org/system/files/conference/usenixsecurity13/sec13-paper_finifter.pdf (accessed on 6 February 2020). [135]

FIRST (2017), *Guidelines and Practices for Multi-Party Vulnerability Coordination and Disclosure*, https://www.first.org/global/sigs/vulnerability-coordination/multiparty/FIRST-Multiparty-Vulnerability-Coordination-latest.pdf (accessed on 8 October 2019). [38]

FIRST (n.d.), *Common Vulnerability Scoring System SIG*, https://www.first.org/cvss/ (accessed on 5 May 2020). [17]

Foudil, E., Y. Shafranovich and Nightwatch Cybersecurity (2020), *A File Format to Aid in Security Vulnerability Disclosure*, https://datatracker.ietf.org/doc/draft-foudil-securitytxt/ (accessed on 9 April 2020). [113]

Franceschi-Bicchierai, L. (2019), *You Can Now Get $1 Million for Hacking WhatsApp and iMessage*, https://www.vice.com/en_us/article/qvqq97/whatsapp-imessage-exploits-zero-days-1-million (accessed on 5 May 2020). [53]

Freed, B. (2020), *Audit finds severe vulnerabilities in Voatz mobile voting app*, https://statescoop.com/audit-finds-severe-vulnerabilities-voatz-mobile-voting-app/ (accessed on 5 May 2020). [96]

Frei, S. and F. Artes (2013), “International Vulnerability Purchase Program (IVPP). Why buying all vulnerabilities above black market prices is economically sound”, https://www.researchgate.net/publication/259442053_International_Vulnerability_Purchase_Program_IVPP (accessed on 24 June 2020). [163]

Frei, S. and O. Rochfort (2021), *The Case for a Bug Bounty Program of Last Resort*, https://techzoom.net/bug-bounty-reloaded/. [59]

French Ministry for Europe and Foreign Affairs (2018), *Paris call for trust and security in cyberspace*, https://pariscall.international/en/call. [148]

Friis-Jensen, E. (2014), *The History of Bug Bounty Programs - Cobalt.io*, https://blog.cobalt.io/the-history-of-bug-bounty-programs-50def4dcaab3 (accessed on 5 May 2020). [116]

FTC (2018), *Mobile Security Updates: Understanding the Issues*, https://www.ftc.gov/system/files/documents/reports/mobile-security-updates-understanding-issues/mobile_security_updates_understanding_the_issues_publication_final.pdf. [31]

Gamero-Garrido, A. et al. (2017), “Quantifying the Pressure of Legal Risks on Third-party Vulnerability Research”, http://dx.doi.org/10.1145/3133956.3134047. [80]

Google (2013), *Patch Rewards*, https://www.google.com/about/appsecurity/patch-rewards/ (accessed on 5 May 2020). [125]

Google (n.d.), *Project Zero*, https://googleprojectzero.blogspot.com/ (accessed on 5 May 2020). [33]

GovTech Singapore (2019), *31 vulnerabilities remediated in second Government Bug Bounty Programme*, https://www.tech.gov.sg/media/media-releases/31-vulnerabilities-remediated-in-second-government-bug-bounty-programme (accessed on 9 April 2020). [121]

Grauer, Y. (2020), *Voatz Bug Bounty Kicked Off of HackerOne Platform*, https://cointelegraph.com/news/voatz-bug-bounty-kicked-off-of-hackerone-platform (accessed on 5 May 2020). [95]

HackerOne (2020), *The 2020 Hacker Report*, https://www.hackerone.com/resources/reporting/the-2020-hacker-report (accessed on 15 April 2020). [131]

HackerOne (2019), *The Hacker-Powered Security Report 2019*, https://www.hackerone.com/resources/reporting/the-hacker-powered-security-report-2019 (accessed on 14 January 2020). [133]

HackerOne (2017), *Hack The Army Results Are In*, https://www.hackerone.com/blog/Hack-The-Army-Results-Are-In (accessed on 5 May 2020). [118]

Hay Newman, L. (2017), *Equifax officially has no excuse*, https://www.wired.com/story/equifax-breach-no-excuse/ (accessed on 5 May 2020). [28]

Heckman, J. (2019), "FBI senior IT official: Bug bounties still useful, but 'a little over-hyped'", *Federal News Network*, https://federalnewsnetwork.com/cybersecurity/2019/07/fbi-senior-it-official-bug-bounties-still-useful-but-a-little-over-hyped/ (accessed on 15 April 2020). [141]

Herpig, S. and A. Schwartz (2019), *The Future of Vulnerabilities Equities Processes Around the World - Lawfare*, https://www.lawfareblog.com/future-vulnerabilities-equities-processes-around-world (accessed on 2 February 2020). [63]

Herr, T., B. Schneier and C. Morris (2017), *Taking Stock: Estimating Vulnerability Rediscovery | Belfer Center for Science and International Affairs*, https://www.belfercenter.org/publication/taking-stock-estimating-vulnerability-rediscovery (accessed on 8 October 2019). [36]

Ilascu, I. (2020), "U.S. Department of Defense discloses critical and high severity bugs", https://www.bleepingcomputer.com/news/security/us-department-of-defense-discloses-critical-and-high-severity-bugs/. [149]

Internetbugbounty.org (n.d.), *The Internet Bug Bounty*, https://internetbugbounty.org/ (accessed on 5 May 2020). [127]

IoT Security Foundation (2018), *Crazy! Less than 10% of consumer IoT companies follow Vulnerability Disclosure guidelines*, https://www.iotsecurityfoundation.org/less-than-10-of-consumer-iot-companies-follow-vulnerability-disclosure-guidelines/ (accessed on 22 October 2020). [111]

ISACA (2017), *Vulnerability assessment*, https://www.scribd.com/document/416535066/ISACA-WP-Vulnerability-Assessment-1117. [161]

ISO/IEC (2019), *ISO/IEC 30111:2019 - Information technology — Security techniques — Vulnerability handling processes*, https://www.iso.org/standard/69725.html (accessed on 8 April 2020). [39]

ISO/IEC (2018), *ISO/IEC 29147:2018 - Information technology — Security techniques — Vulnerability disclosure*, https://www.iso.org/standard/72311.html (accessed on 8 April 2020). [67]

Johnson, A. and L. Millett (eds.) (2019), *Beyond Spectre: Confronting New Technical and Policy Challenges*, National Academies Press, Washington, D.C., http://dx.doi.org/10.17226/25418. [22]

Karisma Foundation (2017), *La corresponsabilidad en acción : Fundación Karisma*, https://web.karisma.org.co/la-corresponsabilidad-en-accion/ (accessed on 6 May 2020). [146]

Kaspersky (2020), *Ethical principles of vulnerability disclosure*, https://www.kaspersky.com/blog/vulnerability-disclosure-ethics/35581/. [145]

Labarthe, S. (2019), *Estudio Sobre Rutas de Divulgacion en Seguridad Digital*, Karisma Foundation, https://web.karisma.org.co/aportes-para-un-entorno-seguro-y-confiable/ (accessed on 6 May 2020). [147]

Laszka, A., M. Zhao and J. Grossklags (2016), *Banishing Misaligned Incentives for Validating Reports in Bug-Bounty Platforms*, http://aronlaszka.com/papers/laszka2016banishing.pdf (accessed on 6 February 2020). [137]

Laszka, A. et al. (2018), *The Rules of Engagement for Bug Bounty Programs*, http://aronlaszka.com/papers/laszka2018rules.pdf (accessed on 6 February 2020). [134]

Lee, A. (2011), *Why Does Sony Keep Getting Hacked?*, https://www.huffpost.com/entry/sony-hack-problems_n_873443?guccounter=1 (accessed on 5 May 2020). [45]

Lee, D. (2019), *WhatsApp discovers 'targeted' surveillance attack - BBC News*, https://www.bbc.com/news/technology-48262681 (accessed on 5 May 2020). [50]

Lerman, C. (2015), *Impact of Free Trade Agreements on InternetPolicy, a Latin America Case Study*, https://repository.upenn.edu/internetpolicyobservatory/12/. [84]

Lin, J. (2019), *Google Online Security Blog: Expanding the Android Security Rewards Program*, https://security.googleblog.com/2019/11/expanding-android-security-rewards.html (accessed on 11 February 2020). [57]

Lopez Romero, T. (2006), *Internet Service Providers' liability for online copyright infringement: the US approach*, https://www.redalyc.org/pdf/825/82511207.pdf. [83]

Manion, A. (2014), *A Survey of Vulnerability Markets*, https://www.first.org/resources/papers/conference2014/first_2014_-_manion-_art_-_certcc.pdf (accessed on 31 January 2020). [51]

Mayersen, I. (2018), *EU to fund bug bounties for open source projects including PuTTY, Notepad++, KeePass, Filezilla and VLC - TechSpot*, https://www.techspot.com/amp/news/78051-eu-fund-bug-bounties-open-source-projects-including.html (accessed on 5 May 2020). [129]

Microsoft (2018), *Microsoft CEO Satya Nadella on fuelling 'tech intensity' in the UK*, https://news.microsoft.com/en-gb/2018/11/07/microsoft-ceo-satya-nadella-on-fuelling-tech-intensity-in-the-uk/ (accessed on 11 May 2020). [30]

Microsoft (n.d.), *Microsoft Bounty Legal Safe Harbor*, https://www.microsoft.com/en-us/msrc/bounty-safe-harbor (accessed on 29 June 2020). [102]

MITRE (n.d.), *CWE - CWE List Version 4.0*, https://cwe.mitre.org/data/index.html (accessed on 5 May 2020). [12]

Mohit Kumar (2019), *Google Offers Financial Support to Open Source Projects for Cybersecurity*, https://thehackernews.com/2019/12/google-open-source-projects.html (accessed on 5 May 2020). [126]

Muncaster, P. (2019), *Bug Bounties Aren't Silver Bullet for Better Security: Report*, Infosecurity Magazine, https://www.infosecurity-magazine.com/news/bug-bounties-silver-bullet-for/ (accessed on 15 April 2020). [139]

NCSC-NL (2018), *Coordinated Vulnerability Disclosure: The Guideline*, https://www.enisa.europa.eu/news/member-states/WEB_115207_BrochureNCSC_EN_A4.pdf (accessed on 10 February 2020). [98]

NCSC-NL (2018), *Coordinated Vulnerability Disclosure: The Guideline*, https://www.enisa.europa.eu/news/member-states/WEB_115207_BrochureNCSC_EN_A4.pdf (accessed on 7 April 2020). [37]

NCSC-NL (n.d.), *Reporting a vulnerability (CVD)*, https://english.ncsc.nl/contact/reporting-a-vulnerability-cvd (accessed on 5 May 2020). [115]

Newman, L. (2019), "Decades-Old Code Is Putting Millions of Critical Devices at Risk", *Wired*, https://www.wired.com/story/urgent-11-ipnet-vulnerable-devices/. [69]

Nichols, S. (2019), *Before you high-five yourselves for setting up that bug bounty, you've got the staff in place to actually deal with security, right?*, The Register, https://www.theregister.co.uk/2019/11/06/disclosure_bug_bounties/ (accessed on 15 April 2020). [138]

NIST (2020), *Foundational Cybersecurity Activities for IoT Device Manufacturers (IR 8259)*, https://nvlpubs.nist.gov/nistpubs/ir/2020/NIST.IR.8259A.pdf. [108]

NIST (2020), *IoT Device Cybersecurity Capability Core Baseline (IR 8259A)*, https://nvlpubs.nist.gov/nistpubs/ir/2020/NIST.IR.8259A.pdf. [109]

NIST (2018), *Cybersecurity Framework version 1.1*, https://www.nist.gov/cyberframework/framework. [107]

NIST (2013), *Guide to Enterprise Patch Management Technologies. NIST Special Publication 800-40 Revision 3*, http://dx.doi.org/10.6028/NIST.SP.800-40r3. [164]

North, O. (2019), *Understanding the economics of bug bounty programs*. [8]

NTIA (2019), *Roles and Benefits for SBOM Across the Supply Chain*, https://www.ntia.gov/files/ntia/publications/ntia_sbom_use_cases_roles_benefits-nov2019.pdf. [70]

NTIA (2018), "Multistakeholder Process on Promoting Software Component Transparency", *Federal Register / Vol. 83, No. 110 / Thursday, June 7, 2018 / Notices*, Vol. 83/110, p. 26434, https://www.ntia.doc.gov/files/ntia/publications/fr-notice-07192018-meeting-software-component-transparency.pdf (accessed on 31 January 2020). [156]

NTIA (2016), *"Early Stage" Coordinated Vulnerability Disclosure TemplateVersion v1.1*, https://www.ntia.doc.gov/files/ntia/publications/ntia_vuln_disclosure_early_stage_template.pdf . [155]

NTIA (2016), *Vulnerability Disclosure Attitudes and Actions*, http://www.iso.org/iso/catalogue_detail.htm?csnumber=45170 (accessed on 8 October 2019). [73]

OASIS (2017), *Common Vulnerability Reporting Framework*, https://oasis-open.github.io/csaf-documentation/index.html (accessed on 8 April 2020). [103]

OASIS (n.d.), *OASIS Common Security Advisory Framework (CSAF) TC | OASIS*, https://www.oasis-open.org/committees/tc_home.php?wg_abbrev=csaf (accessed on 5 May 2020). [72]

OECD (2021), "Encouraging vulnerability treatment: overview for policy makers", *OECD Digital Economy Papers*, OECD Publishing, Paris, https://doi.org/10.1787/20716826. [2]

OECD (2021), "Enhancing the digital security of products: a policy discussion", *OECD Digital Economy Papers*, OECD Publishing, Paris, https://doi.org/10.1787/20716826. [4]

OECD (2021), "Understanding the digital security of products: an in-depth analysis", *OECD Digital Economy Papers*, OECD Publishing, Paris, https://doi.org/10.1787/20716826. [3]

OECD (2019), *An Introduction to Online Platforms and Their Role in the Digital Transformation*, OECD Publishing, Paris, https://dx.doi.org/10.1787/53e5f593-en. [142]

OECD (2019), *Recommendation of the Council on Digital Security of Critical Activities*, https://legalinstruments.oecd.org/en/instruments/OECD-LEGAL-0456 (accessed on 10 April 2020). [151]

OECD (2019), "Roles and responsibilities of actors for digital security", *OECD Digital Economy Papers*, No. 286, OECD Publishing, Paris, https://dx.doi.org/10.1787/3206c421-en. [1]

OECD (2015), *Digital Security Risk Management for Economic and Social Prosperity: OECD Recommendation and Companion Document*, OECD Publishing, Paris, https://dx.doi.org/10.1787/9789264245471-en. [5]

Perlroth, N. (2017), "Why Car Companies Are Hiring Computer Security Experts", *The New York Times*, https://www.nytimes.com/2017/06/07/technology/why-car-companies-are-hiring-computer-security-experts.html (accessed on 6 October 2019). [6]

Ponemon Institute (2019), *Costs and consequences of gaps in vulnerability response*, https://www.servicenow.com/content/dam/servicenow-assets/public/en-us/doc-type/resource-center/analyst-report/ponemon-state-of-vulnerability-response.pdf (accessed on 30 March 2020). [25]

Ponemon Institute (2018), *The 2018 State of Endpoint Security Risk*, https://cdn2.hubspot.net/hubfs/468115/whitepapers/state-of-endpoint-security-2018.pdf (accessed on 6 October 2019). [26]

Porter, J. (2019), *Swiss e-voting trial offers $150,000 in bug bounties to hackers*, https://www.theverge.com/2019/2/12/18221570/swiss-e-electronic-voting-public-intrusion-test-hacking-white-hack-bug-bounties (accessed on 5 May 2020). [122]

Radboud Universiteit (2015), *'Banned' article about faulty immobiliser chip published after two years*, https://www.ru.nl/@1007355/megamos_2015/ (accessed on 21 April 2020). [90]

Reda, J. (2019), *Julia Reda – In January, the EU starts running Bug Bounties on Free and Open Source Software*, https://juliareda.eu/2018/12/eu-fossa-bug-bounties/ (accessed on 5 May 2020). [128]

République Française (2016), *LOI n° 2016-1321 du 7 octobre 2016 pour une République numérique - Article 47*, https://www.legifrance.gouv.fr/affichTexteArticle.do?idArticle=JORFARTI000033203174&cidTexte=JORFTEXT000033202746&categorieLien=id (accessed on 5 May 2020). [100]

Rimmer, M. (2017), "Back to the Future: The Digital Millennium Copyright Act and the Trans-Pacific Partnership", *Laws*, Vol. 6/3, p. 11, http://dx.doi.org/10.3390/laws6030011. [82]

Rittel, H. and M. Webber (1973), "Dilemmas in a general theory of planning", *Policy Sciences*, Vol. 4/2, pp. 155-169, http://dx.doi.org/10.1007/bf01405730. [42]

Ruohonen, J. and K. Kimppa (2019), "Updating the Wassenaar debate once again: Surveillance, intrusion software, and ambiguity", *Journal of Information Technology & Politics*, Vol. 16/2, pp. 169-186, http://dx.doi.org/10.1080/19331681.2019.1616646. [87]

Safecode (2018), *Fundamental Practices for Secure Software Development Essential Elements of a Secure Development Lifecycle Program Third Edition Fundamental Practices for Secure Software Development*, https://safecode.org/wp-content/uploads/2018/03/SAFECode_Fundamental_Practices_for_Secure_Software_Development_March_2018.pdf (accessed on 27 April 2020). [40]

Schneier, B. (2018), *Patching Is Failing as a Security Paradigm - VICE*, https://www.vice.com/en_us/article/439wbw/patching-is-failing-as-a-security-paradigm (accessed on 14 January 2020). [14]

Schwartz, M. (2019), *Equifax's Data Breach Costs Hit $1.4 Billion - BankInfoSecurity*, https://www.bankinfosecurity.com/equifaxs-data-breach-costs-hit-14-billion-a-12473 (accessed on 5 May 2020). [27]

Spring, J. et al. (2018), *Towards improving CVSS*, https://resources.sei.cmu.edu/asset_files/WhitePaper/2018_019_001_538372.pdf. [18]

Spring, T. (2018), *The Vulnerability Disclosure Process: Still Broken*, https://threatpost.com/the-vulnerability-disclosure-process-still-broken/137180/ (accessed on 14 January 2020). [11]

Switzerland National Cyber Security Centre (NCSC) (2020), *Public security test «SwissCovid Proximity Tracing System»*, https://www.ncsc.admin.ch/melani/en/home/public-security-test/infos.html. [123]

Tenable (2019), *Quantifying the attacker's first-mover advantage*, https://static.tenable.com/marketing/research-reports/Research-Report-Quantifying_the_Attackers_First-Mover_Advantage.pdf. [44]

Tenable (2018), *Vulnerability intelligence report*, https://static.tenable.com/translations/en/Vulnerability_Intelligence_Report-ENG.pdf. [21]

Thompson, C. (2019), *Penetration Testing Versus Red Teaming: Clearing the Confusion*, https://securityintelligence.com/posts/penetration-testing-versus-red-teaming-clearing-the-confusion/ (accessed on 5 May 2020). [160]

Trail of Bits Blog (2019), *On Bounties and Boffins*, https://blog.trailofbits.com/2019/01/14/on-bounties-and-boffins/ (accessed on 15 April 2020). [140]

Trustwave (2019), *Trustwave Global Security Report 2019*, Trustwave, http://trustwave.azureedge.net/media/16096/2019-trustwave-global-security-report.pdf (accessed on 19 January 2020). [16]

Udemans, C. (2019), *China working on rules to regulate vulnerability disclosures*, https://technode.com/2019/11/22/china-vulnerability-disclosures-risks/. [154]

UK DCMS (2018), *Code of Practice for Consumer IoT Security*, https://www.gov.uk/government/publications/secure-by-design (accessed on 22 April 2020). [66]

UK GCHQ (2018), *The Equities Process*, https://www.gchq.gov.uk/information/equities-process (accessed on 2 February 2020). [62]

UK NCSC (2016), *Vulnerability management*, https://www.ncsc.gov.uk/guidance/vulnerability-management (accessed on 25 January 2020). [15]

US Department of Defense (2016), ""Hack the Pentagon" Fact Sheet - June 17, 2016", https://dod.defense.gov/Portals/1/Documents/Fact_Sheet_Hack_the_Pentagon.pdf (accessed on 9 April 2020). [119]

US Department of Justice (2017), *A Framework for a Vulnerability Disclosure Program for Online Systems*, https://www.justice.gov/criminal-ccips/page/file/983996/download (accessed on 10 February 2020). [79]

US White House (2017), *Vulnerabilities Equities Policy and Process for the United States Government*, https://www.whitehouse.gov/sites/whitehouse.gov/files/images/External%20-%20Unclassified%20VEP%20Charter%20FINAL.PDF (accessed on 2 February 2020). [61]

van't Hof, C. (2015), *Helpful Hackers. How the Dutch do Responsible Disclosure*, https://cvth.nl/hhe.htm (accessed on 3 April 2020). [89]

Veracode (2019), *The State of Software Security Today*, https://www.veracode.com/sites/default/files/pdf/resources/reports/state-of-software-security-volume-9-veracode-report.pdf (accessed on 24 December 2019). [20]

Verint CIS (2019), *The Top 20 Vulnerabilities to Patch before 2020*, https://cis.verint.com/2019/12/19/the-top-20-vulnerabilities-to-patch-before-2020/ (accessed on 31 March 2020). [157]

Verizon (2019), *2019 Data Breach Investigations Report*, https://enterprise.verizon.com/resources/reports/2019-data-breach-investigations-report.pdf (accessed on 30 March 2020). [24]

Vuldb.com (n.d.), *https://vuldb.com*, https://vuldb.com. [13]

Walsh, J. (2014), *Free Can Make You Bleed*, http://blog.ssh.com/free-can-make-you-bleed. [71]

Wassenaar Arrangement Secretariat (2018), *Wassenaar Arrangement on Export Controls for Conventional Arms and Dual-Use Goods and Technologies PUBLIC DOCUMENTS Volume II List of Dual-Use Goods and Technologies and Munitions List Compiled by the Wassenaar Arrangement Secretariat*, https://www.wassenaar.org/app/uploads/2019/consolidated/WA-DOC-18-PUB-001-Public-Docs-Vol-II-2018-List-of-DU-Goods-and-Technologies-and-Munitions-List-Dec-18.pdf (accessed on 20 April 2020). [86]

whittaker, Z. (2016), *PwC sends ’cease and desist’ letters to researchers who found critical flaw*, https://www.zdnet.com/article/pwc-sends-security-researchers-cease-and-desist-letter-instead-of-fixing-security-flaw/ (accessed on 5 May 2020). [91]

Wilson, M. (2013), *Infographic: How Many Lines Of Code Is Your Favorite App?*, https://www.fastcompany.com/3021256/infographic-how-many-lines-of-code-is-your-favorite-app (accessed on 5 May 2020). [7]

Yin, D. (2019), *Making Reference to ‘State Secrets,’ China Moves to Restrict Vulnerability Disclosures*, https://www.caixinglobal.com/2019-11-21/making-reference-to-state-secrets-china-moves-to-restrict-vulnerability-disclosures-101485876.html. [152]

Zataz (n.d.), *ZATAZ*, https://www.zataz.com/ (accessed on 3 April 2020). [158]

Zerodium (n.d.), *How to Sell Your 0day Exploit to ZERODIUM*, https://zerodium.com/program.html (accessed on 5 May 2020). [52]

Zhao, Laszka and Grossklags (2017), “Devising Effective Policies for Bug-Bounty Platforms and Security Vulnerability Discovery”, *Journal of Information Policy*, Vol. 7, p. 372, http://dx.doi.org/10.5325/jinfopoli.7.2017.0372. [136]

Notes

[1] A threat can exploit a vulnerability without necessarily leading to economic and social damages. In some cases, there may be no or only technical impact.

[2] Borrowing partially from (FIRST, 2017[38]) and from the CVE definition of vulnerabilities. https://cve.mitre.org/about/terminology.html#vulnerability.

[3] See (OECD, 2021[3]) and (OECD, 2021[4]).

[4] For an in-depth discussion of issues related to "end of life", see (OECD, 2021[3]) and (OECD, 2021[4]).

[5] See (OECD, 2021[3]) and (OECD, 2021[4]).

[6] Australia, France, Germany, Japan, the Netherlands, New Zealand, Singapore, UK, and US.

[7] For details on patch management, see for example (NIST, 2013[164])

[8] Benefits from, obstacles to and ways to encourage "security by default" are further discussed in (OECD, 2021[3]) and (OECD, 2021[4]).

[9] The term "hacker" was popularised by Steven Levy in his 1984 book "Hackers: Heroes of the Computer Revolution". It did not have negative connotations but rather referred to a skilled computer expert who uses their technical knowledge to overcome problems.

[10] The discovery of vulnerabilities by governments in addressed in **Error! Reference source not found.** and 1.2.8.

[11] See for example Zataz in France (Zataz, n.d.[158]) and in (van't Hof, 2015[89])

[12] The end of life gap is discussed in (OECD, 2021[3]) and (OECD, 2021[4]).

[13] In ISO/IEC 29147, *vulnerability disclosure* covers what code owners (called vendors in the standard) should do when they receive information about a possible vulnerability in their products, with a focus on the relationship with security researchers. It is complementary to *vulnerability handling,* addressed in ISO/IEC 30111, which covers how code owners should process vulnerability information from the investigation to the post-release phases, regardless of whether the information comes from an external source or the vendor's internal security team.

[14] Standards and guidance to integrate digital security in design & development are discussed in (OECD, 2021[3]) and (OECD, 2021[4]).

[15] Co-ordination for vulnerability handling is further discussed in (OECD, 2021[3]) and (OECD, 2021[4]).

[16] Cf. for example the OECD Recommendation on digital security of critical activities for a definition of activities that could justify advanced communication of a vulnerability to government (OECD, 2019[151]).

[17] Cf. www.civicert.org.

[18] "Article 2 – Illegal access. Each Party shall adopt such legislative and other measures as may be necessary to establish as criminal offences under its domestic law, when committed intentionally, the access to the whole or any part of a computer system without right. A Party may require that the offence be committed by infringing security measures, with the intent of obtaining computer data or other dishonest intent, or in relation to a computer system that is connected to another computer system."

[19] 18 U.S.C. § 1030(a)(2)(C).

[20] For a guide on vulnerability management, see (Carnegie Mellon University, 2016[159]).

[21] Cf. para 30 and articles 6(1)(b), 54(1)(m), 50, and 55(1)(c).

[22] J. van der Ham, presentation at OECD SDE meeting, November 2019.

[23] The US Government is required to categorize, identify, implement, and assess all operational information and information systems per NIST SP 800-37 based on FIPS 199 & NIST SP 800-53 rev.4 security controls. Part of the Risk Management Framework process (RMF) is the Plan of Action and Milestone (POA&M) where each weakness (vulnerability) must accounted for cost, resource, and mitigation timeline based on its criticality rating. When it comes to VDP, the system owner will have to take into account their FISMA system, review their System Security Plan (SSP) and conduct a risk assessment.

[24] https://opencryptoaudit.org