OECD*publishing*

ENCOURAGING VULNERABILITY TREATMENT

OVERVIEW FOR POLICY MAKERS

OECD DIGITAL ECONOMY PAPERS

February 2021 **No. 307**

OECD
BETTER POLICIES FOR BETTER LIVES

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 builds upon a separate in-depth background report providing an extensive discussion of digital security vulnerability management, handling and disclosure, and including a large number of references, detailed illustrations and explanations (OECD, 2021[2]). A full bibliography supporting this overview for policy makers is available in the background report.

This report has been 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.

Note to Delegations:

This document is also available on O.N.E. under the reference code:

DSTI/CDEP/SDE(2020)12/FINAL

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 have 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 often 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 access to a trusted co-ordinator**, who can help connect stakeholders when needed and 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. It requires an open mind, flexible solutions and case-by-case consideration, often on the basis of international standards and guidelines;
- **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 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 is also 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.

Code is the engine of digital transformation. An increasingly complex engine. Today, a typical smartphone application has tens of thousands of lines of code, while operating systems or systems embedded in cars have tens of millions. Software is also complex because it embeds numerous layers of code and components developed by third-parties.

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. All information systems also have vulnerabilities related to how software is implemented, configured, and updated. Criminals and other ill-intentioned actors actively seek to discover vulnerabilities in code and information systems. With this knowledge, they develop or use tools such as “malware” to exploit those vulnerabilities, and 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 other actors across the value chain, including users, in order to reduce digital security risk. Organisations should also monitor their information systems to ensure that patches are appropriately applied and to avoid misconfigurations. These are complex, burdensome and expensive endeavours. They are also never-ending tasks because threat actors evolve their techniques and continuously discover new vulnerabilities.

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. 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. The report aims to:

- Identify how public policies can promote the broad adoption of good practice for vulnerability treatment, which includes the discovery, handling, management and disclosure of vulnerabilities in products and information systems (definitions are provided in the Glossary);
- Raise policy makers' awareness about the importance of vulnerability treatment, and in particular co-ordinated vulnerability disclosure, in the era of digital transformation;
- Inform the development of international guidance by the OECD to encourage the adoption of good practice in this area.

The report explains stakeholders' roles, and describes existing good practice, challenges and obstacles to their adoption. It also provides directions on how public policy can help encourage the adoption of good practice and facilitate the international co-ordination of approaches. The structure of this report is as follows:

- Chapter 1 describes vulnerabilities and their mitigations;
- Chapter 2 explores vulnerability treatment
- Chapter 3 focuses on Coordinated Vulnerability Disclosure (CVD)
- Chapter 4 discusses possible public policy guidance

A Glossary provides definitions of key terms for the purpose of this report. Annexes 1 and 2 provide an overview, respectively, of good practice for CVD and of possible areas for future work.

This report is based on an in-depth background report on this subject containing more details and full references of sources (OECD, 2021[2]).

1. Understanding vulnerabilities

A digital security vulnerability (hereafter "vulnerability") is a weakness that, if exploited or triggered by a threat, can cause economic and social damages, by affecting availability, integrity, or confidentiality of a digital resource or asset. This report focuses on code and system vulnerabilities. It does not address other types of vulnerabilities such as the absence of backup procedures, or people clicking on malicious links in emails for lack of digital security awareness.

1.1. Code vulnerabilities

Code vulnerabilities affect the code embedded in a product's software and/or firmware components. Malicious actors can develop exploit code (or "exploit") to weaponise them, and use them to steal and extort money and data, disrupt processes, or spy on organisations and individuals. Exploits and exploitation services (i.e. attacks) can be purchased on black markets, including as turnkey attack solutions. As of mid-2020, the most well-known code vulnerability databases included between 129 000 and 150 000 unique code vulnerabilities grouped in over 800 categories. However, these databases do not cover all products available globally, and many discovered vulnerabilities are never disclosed. While only a fraction of vulnerabilities registered in vulnerability databases are present in real-world information systems, research by a security firm found that 100% of the applications scanned by its security tool contained at least one known vulnerability, with a median number of 15 vulnerabilities per application. Although harmful when used for offence, exploits can be useful when a defender uses them to test the effectiveness of its protections, for example through network penetration testing and "red-team" work[1], which some standards and regulations may require.

Code vulnerabilities are not created equal: discovered code vulnerabilities have different levels of risk and severity. The *risk* depends upon the use context of the product in which the vulnerability is located, which can vary considerably across users and use contexts for the same product. The *severity* of a vulnerability depends upon its exploitability and the amount of damages it can create regardless of the product's use context.

A code vulnerability can be latent, i.e. undiscovered. There is an unknown number of latent vulnerabilities in every piece of software. Code considered sufficiently secure when first released can become vulnerable subsequently, from a few minutes to many years later because attackers can discover vulnerabilities that went previously unnoticed. Some vulnerabilities were discovered in Microsoft's Windows XP more than ten years after its release, and years after the end of its commercial life. Attackers can also adopt new attack techniques that were unknown when the code was written. **Code is never perfect**, even when applying best security coding practice ("security by design").

Zero-day vulnerabilities ("zero-days") are code vulnerabilities for which no mitigation has yet been released or which are unknown to the code owner. Figure 1 provides an overview of zero-day related terminology. A zero-day is particularly difficult to mitigate. A "zero-day exploit" is an exploit based on a zero-day vulnerability. While very likely to succeed until a mitigation is available, zero-day attacks are rare in comparison with attacks exploiting known vulnerabilities.

Figure 1. Typical evolution of a code vulnerability

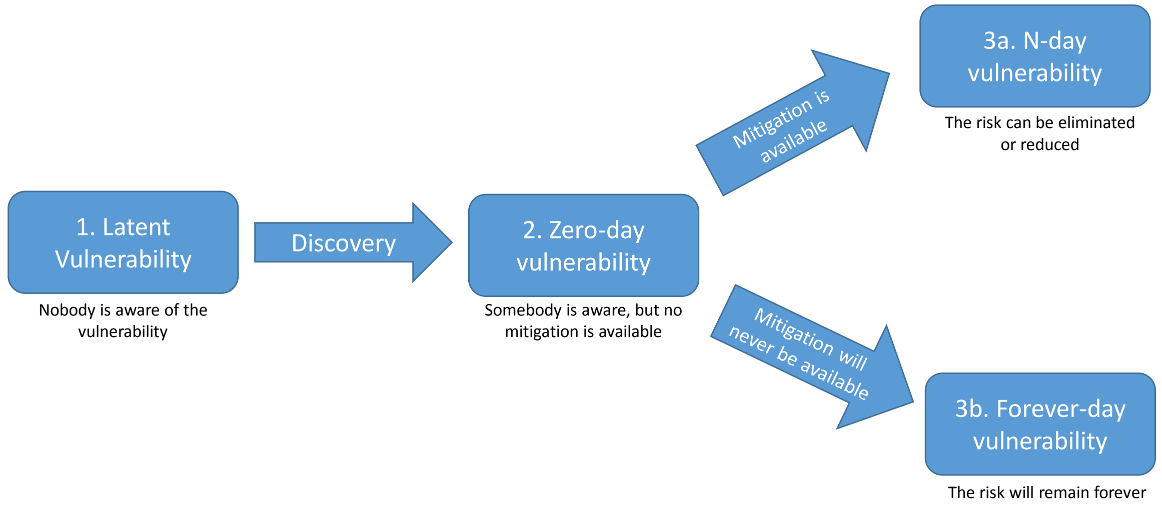
1. Latent Vulnerability
Nobody is aware of the vulnerability
Discovery
2. Zero-day vulnerability
Somebody is aware, but no mitigation is available
Mitigation is available
3a. N-day vulnerability
The risk can be eliminated or reduced
Mitigation will never be available
3b. Forever-day vulnerability
The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

Note: Some vulnerabilities remain latent (1) forever, while others are discovered and become a "zero-day" (2). When the code owner develops a mitigation, 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).
Source: OECD

1.2. System vulnerabilities

System vulnerabilities are weaknesses in the way products are implemented or configured in information systems. Failure to keep implemented products up-to-date with the latest update or patch is a major source of system vulnerabilities. Rather than spending time and resources to discover new code vulnerabilities, most attackers test their victims' systems against known vulnerabilities until they find one that has not been patched and thus can be exploited. According to a 2019 survey, 60% of respondents say one or more breaches they faced occurred because a patch that was available for a known vulnerability was not applied. There are many infamous examples of disasters that would have failed if security updates had been swiftly applied. They include the 2017 WannaCry and NotPetya attacks which resulted in multi-billions of dollars in global damages, the Equifax incident that affected 56% of all American adults, and 14 million British citizens, costing the company at least USD 1.4 billion, as well as the 2016 attack against the Ukrainian electricity operator which generated a blackout in Kiev.

Many system vulnerabilities are also related to improper or outdated product or system configuration or settings caused by administrators lacking security awareness, knowledge, resources or time to configure products appropriately, or systematically manage security. Many products are shipped with minimal security settings by default, making them "vulnerable by default". For example, a product may be shipped to its users with the same weak default password and without a mechanism incentivising users to change it at first installation. Many stakeholders and security experts are now calling for products to be "secure by default", i.e. provided with high security settings when first installed or used, leaving it to users to take the responsibility for weakening security as appropriate.[2]

1.3. Mitigations and patches

To reduce the risk, it is possible to apply a mitigation measure ("mitigation"). For *code vulnerabilities*, a mitigation called a "patch" modifies the code to fix the vulnerability. Patches need to be implemented on each software instance through a security update, broader update or new release (e.g. upgrades in mobile apps). However, it is not always possible to develop a patch, for example, when the product is no longer supported, does not have update capabilities, or would have to be redesigned to fix the vulnerability. In such cases, a set of instructions, configuration requirements or documentation can reduce the risk without necessarily eliminating it. In some smart products, such as certain low-cost IoT devices, the code cannot be updated. For *system vulnerabilities*, mitigations consist in actions that system owners can take, e.g. changing configuration settings or applying an existing patch previously set aside.

However, there is no way to eliminate all vulnerabilities. While addressing vulnerabilities is essential, fixing *all* vulnerabilities would not be a realistic objective, for many reasons including cost and technical feasibility. Furthermore, many code vulnerabilities will never be exploited, and some system owners may not apply a patch because it would disrupt operations, create compatibility issues or introduce additional risk. Moreover, vulnerabilities may be discovered in products that are still in use but no longer supported and will never be corrected. In absence of a code mitigation, or when a patch cannot be applied, a workaround may exist, such as a configuration change in a firewall.

Risk management, which enables discernment and flexibility, is the cornerstone of vulnerability treatment. The overarching objective of vulnerability treatment is to make products and information systems "secure enough" rather than absolutely secure, in order to sufficiently reduce, rather than entirely eliminate, security risk for users and third parties. Mitigation development is primarily a matter of prioritisation based on risk assessment.

2. Vulnerability treatment

Vulnerability treatment is a key building block of digital security. Digital security risk can be described as the likelihood of *threats* exploiting *vulnerabilities*, creating *incidents* that *impact* the economy and society, including human life, people's privacy and trust. Each of these risk factors, namely threats, vulnerabilities, incidents and impact, provides a lever for policy makers and stakeholders to enhance digital security. So far, policy makers have focused on the fight against threats (e.g. through cybercrime frameworks), incidents mitigation (e.g. through incident response capacity building), and impact reduction (e.g. through preparedness and resilience measures). Vulnerabilities however have received less policy attention. Nevertheless, it is generally easier to fix a vulnerability than to catch a criminal that could exploit it. In years to come, digital transformation will accelerate, making our economies and societies increasingly digitally dependent. Furthermore, disruptive technologies such as artificial intelligence, the Internet of Things and blockchain will spread and expand stakeholders' attack surfaces. As shown during the COVID-19 pandemic, attackers will continue to seize every opportunity to create harm (OECD, 2020[6]). Therefore, **it is time for policy makers to complete their digital security toolkit and encourage stakeholders to treat vulnerabilities**. To do so, it is necessary to understand the vulnerability treatment lifecycle as well as the core conditions and key challenges for effective vulnerability treatment.

2.1. The vulnerability treatment lifecycle

There is an "ideal" vulnerability treatment lifecycle ("vulnerability lifecycle") to reduce risk, which consists of the following stages:

- **Discovery**: to reduce a vulnerability's harmfulness, someone has to discover it in the first place. If a security researcher, namely someone who has the intention to reduce digital security risk, finds the vulnerability, they need to report the information to the "vulnerability owner", namely the organisation that owns the responsibility to mitigate the vulnerability because it is the best placed to fix it. However, if malicious actors ("threat sources") find it first, they will likely develop code to weaponise it ("exploit"), and use it against the product or system users, or sell it to other threat sources who will do the same.
- **Handling**: If the vulnerability is in a product's code (code vulnerability), the "code owner" needs to develop a mitigation and distribute it to all users ("vulnerability handling").
- **Management**: if the vulnerability is a misconfiguration or unapplied mitigation (often a patch) in an information system (system vulnerability), the "system owner" needs to manage it, i.e. apply the patch or change the system or product configuration as soon as possible. When this overall process is over, the risk is eliminated or sufficiently reduced for that particular vulnerability owner.
- **Disclosure**: In most cases, code vulnerability information needs to be disclosed publicly or at least to the security community or to targeted audiences in order to enable relevant stakeholders to discover related or similar vulnerabilities in other products, improve security tools (e.g. by feeding AI-based attack detection applications), and enhance the community's knowledge. This is often called "public disclosure".

The terms "handling" and "management" are not used interchangeably among experts. Taking a smart product's perspective[3], vulnerability handling takes place on the vulnerable product's supply-side, and vulnerability management on the vulnerable product's demand side. Unfortunately, the term disclosure can sometimes refer to the reporting of the vulnerability to the vulnerability owner in the discovery stage, which can create some confusion.

Box 1 provides an overview of the main stakeholders involved in the vulnerability lifecycle.

The vulnerability lifecycle is a race against the clock. 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. Failures in the treatment process benefit malicious actors who have more time to rediscover, weaponise, and exploit the vulnerability to create harm.

Box 1. The main stakeholders involved in vulnerability treatment

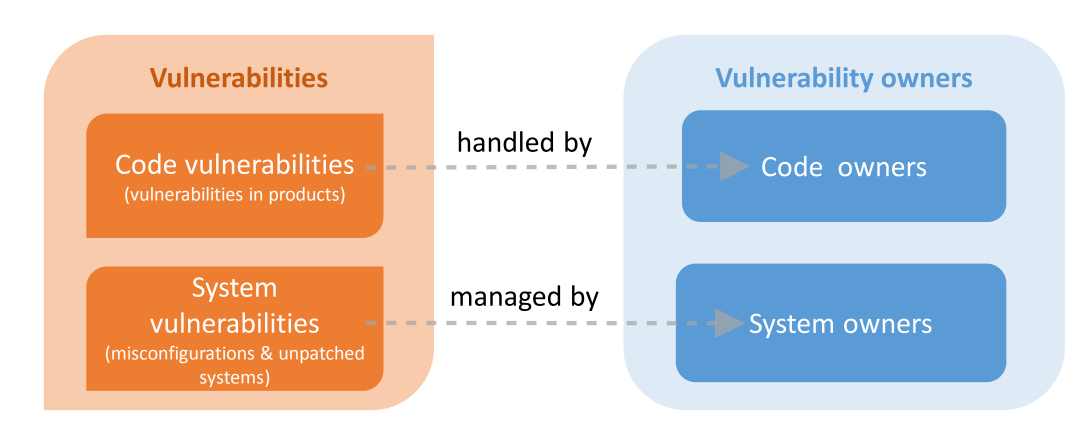

Vulnerability owners are stakeholders who own the responsibility to act upon a vulnerability they are aware of, in order to mitigate it. The term "owner" focuses on the responsibility to address vulnerabilities (as in "risk owner") and does not necessarily entail property right. They include code and system owners.

Code owners are supply-side actors, i.e. 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 handle code vulnerabilities ("vulnerability handling"). With digital transformation, numerous stakeholders not previously associated with the ICT industry are becoming code owners. They range from banks to grocery stores, local governments, newspapers, cars, toys and manufacturers of medical, cooking, and industrial devices, etc.

System owners are the demand-side actors, i.e. 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 likely to operate increasingly complex systems, without necessarily having sufficient awareness and understanding of digital security, and resources to manage digital security risk. In theory, individual users are also system owners. However, they are outside of the scope of this report because they generally do not receive vulnerability reports.

Security researchers ("researchers") are individuals or organisations who identify a potential code or system vulnerability with the intention to reduce security risk. The media often calls them "finders", "discoverers", "ethical hackers", "white hats" or "friendly hackers". The term researcher is more neutral than the ambiguous term "hacker", which often carries negative connotations. Different goals drive different categories of researchers who operate under different constraints. Some research vulnerabilities as part of their professional activities, working for academia, commercial security companies, security teams, government agencies, or civil society. Others find vulnerabilities as a personal hobby in their spare time.

Other stakeholders include *co-ordinators*, who can assist vulnerability owners and researchers in the disclosure process, *market intermediaries*, such as bug bounty platforms and grey market vulnerability brokers, and *third party victims* of digital security incidents (e.g. patients whose surgery operations were delayed due to a ransomware attack on a hospital). In addition, *threat sources* are malicious actors who identify and exploit vulnerabilities for malicious or ill-intentioned purposes. They include governments (cf. Box 3), criminal groups and individuals.

The ideal vulnerability lifecycle is often difficult to reach. For such a scenario to unfold, everyone has to understand their role and take responsibility accordingly. Unfortunately, all incentives, capacities and perceptions are rarely aligned. Vulnerability owners can lack an effective vulnerability handling or management process, therefore they do not take appropriate action when a vulnerability is discovered or reported to them. In addition, stakeholders' perceptions are often divergent. At each step of the ideal lifecycle, stakeholders may take a different path and the whole process can slow down, stop, or take a different direction, increasing instead of decreasing risk for all.[4]

Sometimes, security researchers decide to disclose publicly a newly discovered vulnerability without co-ordinating with the vulnerability owner ("full disclosure"). They can have good reasons to do so. For example, it can incentivise an otherwise unmotivated code owner to begin or accelerate the development of a mitigation. Nevertheless, vulnerability information is sensitive because it can increase the risk of weaponisation. Therefore, full disclosure can in some cases increase risk. However, threat actors can also reverse-engineer patches to discover the underlying vulnerabilities they correct. Therefore most researchers and vulnerability owners prefer vulnerability information to be ultimately disclosed publicly because it is always possible that at least one threat actor has been exploiting the vulnerability without being detected. Keeping the information confidential would give them an advantage. Nevertheless, to the extent possible, information should be disseminated at a moment and in a manner that do not facilitate threat actors' efforts. These decisions are always a dilemma.

2.2. Key conditions for effective vulnerability treatment

For effective vulnerability treatment, vulnerability owners need to take responsibility for vulnerabilities in their products and systems, and co-ordinate with researchers the public disclosure of vulnerabilities.

Code and system owners need to take responsibility for addressing vulnerabilities in their products or systems.

Vulnerability owners need to recognise that their products and information systems are no exception and therefore have vulnerabilities. They have to establish a process to address them within their existing product or system's security frameworks (cf. left and right circles in Figure 2).

Figure 2. Overview of vulnerability treatment

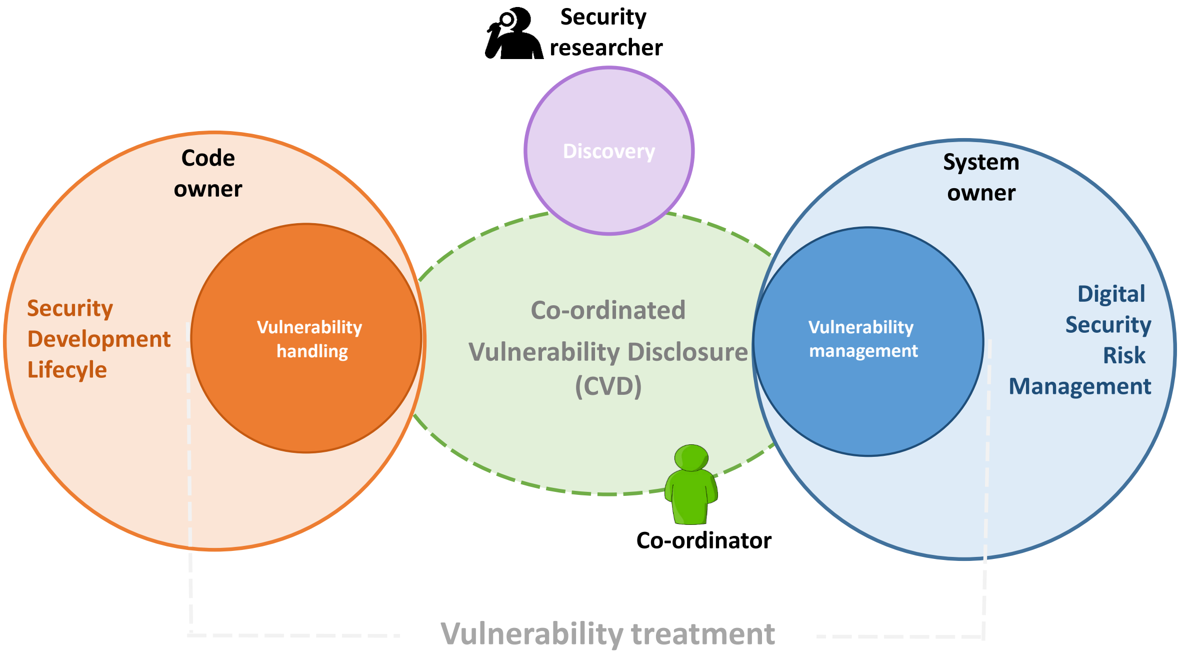

Note: Security researchers discover vulnerabilities and report them to the relevant code owner (on the vulnerable product's supply side) or system owner (on the vulnerable product's demand side). Vulnerability handling is part of a code owner's Security Development Lifecycle (SDL). It can involve a vast ecosystem of stakeholders, depending on the vulnerable product's value chain. Vulnerability management is part of a system owner's digital security risk management framework. Stakeholders co-ordinate their actions through CVD to ensure that risk is reduced for all. A co-ordinator can facilitate the CVD process.
Source: OECD

Code owners have to implement a "vulnerability handling" process to address code vulnerabilities. The reference international standard for vulnerability handling, ISO/IEC 30111[5], covers how code owners should process vulnerability information. 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.

System owners have to implement a "vulnerability management" process to address system vulnerabilities. Vulnerability management enables the organisation to know if there are vulnerabilities within their IT estate and take appropriate risk management decisions and actions. It includes vulnerability scanning, patch testing and deployment. 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.

Vulnerability handling and management are learning and improvement cycles that aim at progressively improving the security of products and systems, as well as the organisation's overall security practices, in order to develop products or manage systems with less vulnerabilities in the future.

These processes are costly and difficult to implement systematically and continuously, which may in part explain why many organisations do not have them yet fully in place. Misconceptions may also be at play, such as associating vulnerabilities with a failure that can lead to negative reactions from the leadership, shareholders or markets.

Vulnerability handling and management are business-led risk-based processes because they can affect the organisation's economic and social performance. As vulnerabilities are not created equal, vulnerability owners have to prioritise efforts and resources to develop, distribute, or apply a mitigation, based on the risk associated with the vulnerability.

For *code owners*, the risk depends upon the product's use context, which varies across users of the vulnerable product, and is not necessarily aligned with the vulnerability's severity. It also depends upon code owners' business risk, such as the consequences of publicly recognising the presence of vulnerabilities in their product. Therefore, vulnerability handling should be a business-led process, integrated in the product's business strategy, and consistent with users' needs, rather than only a technical matter. When a mitigation is available, *system owners* need to balance the risk of a successful attack exploiting the vulnerability with the risk of immediately applying the mitigation to their system. While silent and automatic patching is a reasonable objective for consumer products, it is less so for more complex information systems in organisations. Many organisations cannot simply apply all security updates as soon as they receive them, in particular in industrial or complex digital environments because it can impact business-critical systems. They need to test the mitigation first to assess whether it is not going to disrupt operations or introduce new security, compatibility, performance or instability issues. Vulnerability management requires business-technical co-ordination within the organisation. It can also be a heavy process.

According to a 2019 survey, it takes organisations an average of 102 days to test and fully deploy patches. While some organisations just have poor vulnerability management, others may have a risk-based rationale for long patching delays, or even no patching at all, for example where assembly lines or physical processes cannot be interrupted. Vulnerability management is not yet the norm in part because it is both difficult and expensive. 85% of the same survey's respondents considered that their organisation was below a middle stage of vulnerability management maturity. Keeping up with security updates was extremely challenging or challenging for 65% of those with a patch management process in place, and only 37% said they had adequate staffing to patch vulnerabilities in a timely manner. Because the window of exposure is open until the mitigation is applied, time is of the essence, but assessing risk is time consuming and therefore keeps the window open.

Vulnerability owners and researchers need to co-ordinate vulnerability disclosure

Vulnerability disclosure is a "wicked problem", difficult to approach in a general or systematic manner. There is no simple and always valid solution. Each case is sufficiently different to prevent abstract and ideal solutions to work in most cases. Furthermore, stakeholders' perceptions often diverge, including on how best reducing risk. Therefore, after many years of discussions, the technical community came to the consensus that co-ordination among stakeholders is the best approach to disclose vulnerabilities in an effective and responsible manner. Co-ordinated Vulnerability Disclosure (CVD) is based on the assumption that the shared goal of reducing digital security risk can transcend differences of perceptions between vulnerability owners and researchers, and that co-ordination between them can iron out divergences and facilitate the adoption of the best solution in most cases. CVD is sometimes used as a term covering all stakeholders' co-ordination efforts throughout vulnerability treatment. Over time, CVD best practices and standards have emerged, such as ISO/IEC 29147 on vulnerability disclosure, which addresses how code owners should process vulnerability reports from and manage the relationship with security researchers.

Co-ordination between the vulnerability owner and the researcher is at the core of successful vulnerability disclosure. A dialogue based on good communication between them is essential to reduce misunderstandings and facilitate the process. For example, they have to agree on what information to disclose and when. Should they immediately share vulnerability information publicly, or wait for a patch or mitigation to be released? What should be the level of detail of the information to make public? This is often a dilemma because, depending on how and when vulnerability information is shared, it can give an advantage to attackers rather than defenders. Stakeholders need to agree on many other aspects such as the criticality of the vulnerability, and urgency of developing a mitigation and making vulnerability information public, etc.

CVD is further detailed in chapter 3.

2.3. Key challenges to vulnerability treatment

Key challenges to effective vulnerability treatment include the grey and black markets for vulnerabilities, as well as limited trust in the government.

Making defence more attractive than offence on the vulnerability market

Stakeholders can exchange vulnerabilities on white, grey or black markets. These markets are global as many transactions take place across borders.

- The *white (i.e. regulated) market* connects vulnerability owners and researchers, with or without a reward mechanism (e.g. bug bounties).
- On *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 tools purchased by police forces or intelligence agencies to access the content of mobile phones. Grey markets provide a means for buyers to externalise the vulnerability discovery phase to other actors and rapidly develop zero-day-based tools and exploits. Such grey markets may be legal or illegal, depending on the jurisdiction and context.
- On *black (i.e. illegal) markets*, buyers and sellers trade vulnerabilities for offensive use, generally on underground platforms in the dark web, through online chat rooms, or specialised marketplaces.

While the black market is illegitimate by definition, there is a debate about the grey market. 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 beyond the scope of this report. Many experts agree that when buyers are well-resourced intelligence, defence and law-enforcement agencies, the grey market distorts prices, diverts researchers from the white market and makes offense lawfully pay better than defence as these buyers can outbid code owners to acquire critical zero-day vulnerability information.

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).

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

Recent research based on data from the US National Vulnerability Database (NVD) 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.

The analysis, carried out with 2010 to 2020 NVD data, further 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. 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 in the NVD database, 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% of the US GDP). It would also represent much less than 0.5% of global cybercrime losses, assuming the total losses amount to USD 1000 billion (estimates of the global cost of cybercrime, which are always a matter of debate, range from USD 100 to 6 000 billion). Furthermore, the cost for the top 11 publicly tradedvendors, 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.

The idea of vendors buying all their products' vulnerabilities at an arbitrary price is rather unrealistic and unfeasible. However, as a research proposal, it is useful to put the grey and black market challenges into perspective. 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[7]).

Building trust in the government

Governments cumulate almost all possible roles in the vulnerability treatment landscape: product users, vulnerability owners, security researchers when an agency is tasked with discovering vulnerabilities, as well as victims of attacks. In some cases, governments can also be threat sources when law enforcement and national security agencies develop or purchase offensive tools, or perform offensive "cyber operations". The possible offensive role of governments, outlined in Box 3, can be considered at the root of a major trust challenge that can seriously undermine public policy efforts to encourage vulnerability treatment.

Governments can receive vulnerability reports from researchers or vulnerability owners reaching out to the government-operated CERT to receive co-ordination assistance. Furthermore, code owners and/or researchers may consider providing advance information to the government about code vulnerabilities that could affect critical activities, with impact on safety, national security, a large share of GDP, or large parts of the population[6]. In these cases, the government can help ensure that operators of such critical activities take measures prior to public disclosure.

Box 3. The offensive role of governments

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 them as part of their law enforcement, intelligence and national security activities. They may discover vulnerabilities, stockpile them without disclosing 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. Some governments can also buy vulnerabilities and exploits on the grey and black markets to carry out "offensive operations", a soft term to refer to a digital security attacks by government. In doing so, these governments contribute to price setting and legitimisation of these markets' supply side actors 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 equivalent to intentional vulnerabilities. The business, civil society and technical communities have almost unanimously condemned this practice, and some governments have condemned it as well. Backdoors are likely to be discovered and exploited by criminals and other offensive actors at some point. They increase risk to all stakeholders.

In some cases, governments therefore appear to be part of both the solution and problem of digital security vulnerabilities, using one hand to increase the risk while their other hand is struggling to reduce it jointly with other stakeholders.

Goals of governments' offensive operations may be often legitimate: criminal investigations, anti-terrorism, protection of sovereignty, counter espionage, etc. However, in certain cases, such operations can violate human rights, by targeting civil society advocates, whistle-blowers, journalists, lawyers, and simple citizens. They can also be used for economic espionage, theft of trade secrets, preparation for armed conflicts and other covert actions.

Stockpiled vulnerabilities and exploits can sometimes be detected when they are used in the course of offensive operations. Criminals, foreign governments, or activists can steal them, and insiders can leak them to the public. Anyone who obtains these exploits can turn them against anyone in the world, improve, and/or commercialise them on the grey and black markets. As a result, the global level of digital security risk increases. The digital security chronicle includes several examples of such leaks, thefts, and repurposing of exploits initially kept or used by governments.

Some government's ambiguity with respect to vulnerability exploitation can undermine other stakeholders' trust and 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, including domestic and foreign governments. For the reasons highlighted above, most stakeholders are rather suspicious about some of them by default. Therefore, from the non-governmental stakeholders' point of view, governments need to demonstrate that they can be trusted. How to do so is, however, not yet clear. A public governance model that strictly and transparently separates the government's defensive and offensive functions at the institutional level seems to be a step in the right direction.

A vulnerability owner operating or having users in several countries faces a dilemma when it wishes to provide such critical vulnerability information to government. If it provides the information to its own government and/or to trusted governments whose population could be affected by the vulnerability, other countries may accuse the vulnerability owner of putting their population in danger. If it reaches out to all governments with potentially affected populations, including those it does not trust, the latter could weaponise the vulnerability, or behave inappropriately, e.g. by immediately disclosing the information publicly. Currently, international CERT co-operation tends to resolve these issues even if their informal working methods can raise difficult challenges, for example, when CERTs are not independent from agencies with an offensive capability. Furthermore, in some countries, it can be dangerous for a security researcher to discover a vulnerability as the government can use legal threat, and/or coercion to discourage public disclosure or discredit the researcher. Lack of trust in the government can disincentivise security research, depriving the country from a useful means to increase security.

Most vulnerability owners and researchers will provide information in advance only if they trust that the receiving government agency will not weaponise it or communicate it to another entity that could use it for offensive purposes, and that it will handle it in a secure manner, avoiding leaks or communications to inappropriate third parties. To foster trust, government can adopt strict separation between the entity receiving information and other agencies, such as those with an offensive role, as well as clear and transparent processes about what the agency does with the information received. Furthermore, some experts have suggested that agencies receiving vulnerability information for co-ordination purposes should be independent, akin to data protection authorities, and that an international co-ordinator operating as a non-profit and non-governmental organisation could help address cross-border information exchanges.

When governments discover zero-day vulnerabilities through their own research and internal efforts, they have to decide what to do with them. If the mandate of the discovering agency is limited to protection, the agency will have to immediately report the information to the vulnerability owner and trigger a CVD process in order to decrease risk for users, including the government itself and operators of critical activities. If its mandate also includes offense, or if the agency can legally share the vulnerability with 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 to enable offensive exploitation or stockpiling. The United States and the United Kingdom have published a “Vulnerabilities Equities Process” charter guiding how they assess risks and benefits to both the intelligence requirements and the digital security of the country. At the time of writing, the German government is working on developing a similar charter.

3. Co-ordinated Vulnerability Disclosure

CVD has become a key best practice to optimise the outcome of vulnerability treatment. This section explores the key conditions for CVD, the tools to facilitate CVD as well as the legal risk faced by researchers, which is a major obstacle to CVD.

3.1. Key conditions for CVD

CVD is a process through which vulnerability owners and researchers work co-operatively to reduce the risks associated with the disclosure of a vulnerability. The premise of CVD is that vulnerability owners and researchers need to co-operate to reduce users' window of exposure (see Glossary). The objective of CVD is to develop and distribute mitigations before vulnerability information becomes public. CVD is widely recognised as a good practice for researchers and vulnerability owners to address vulnerability disclosure responsibly. It is a co-ordination framework involving one or more security researcher(s) and vulnerability owner(s). CVD can, but does not necessarily, involve a third-party co-ordinator. International standards and guides provide guidance for CVD at the operational level. CVD is effective both for code and system vulnerabilities, and therefore an extremely large number of stakeholders can potentially use it to improve digital security.

The complementary, competing or conflicting incentives influencing the behaviour of organisations and individuals largely determine the success or failure of CVD. Externalities can also play a role when the costs incurred by the exploitation of vulnerabilities are not borne by the vulnerability owner (Cf. Figure 3). Several conditions, introduced below, are required for CVD.

Awareness and knowledge of CVD

To implement a CVD process, stakeholders need first to be aware that it is a recognised best practice to reduce digital security risk. They also need to understand how it works and what their role is, which may be difficult given the complexity of CVD. Many businesses lacking a strong digital and digital security culture, such as IoT devices producers and non-traditional ICT firms, are often unaware of CVD. In organisations, internal stakeholders lacking a sufficient understanding of CVD may discourage entering into a CVD process. For example, they may refuse to recognise that their products or system can have vulnerabilities, view CVD as a risk to their brand, interpret CVD as a broad encouragement for "hacking" the organisation, or fear that embracing CVD would attract criminals. When sufficiently trusted, the security and IT teams can educate departments with a role in decision-making or exposed to negative consequences of a vulnerability disclosure failure (leadership, marketing, public relations, legal, etc.).

Figure 3. Economic incentives, motivations and barriers in a co-ordinated vulnerability disclosure process

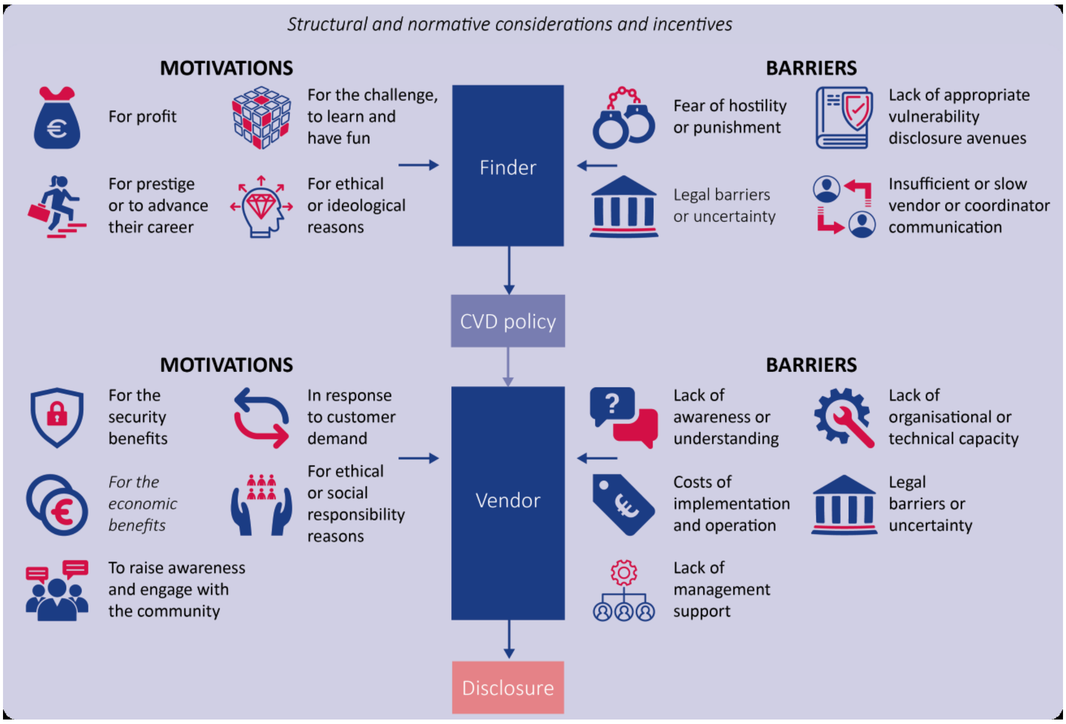
Structural and normative considerations and incentives
MOTIVATIONS
For profit
For the challenge, to learn and have fun
For prestige or to advance their career
For ethical or ideological reasons
Finder
BARRIERS
Fear of hostility or punishment
Lack of appropriate vulnerability disclosure avenues
Legal 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

Note: in ENISA's terminology, a security researcher is a finder and a vulnerability owner is a vendor.
Source: (ENISA, 2018[8])

Managing the sensitivity of vulnerability information

Newly discovered vulnerability information can be highly sensitive. It need to be reported only to the parties best placed to develop, test and deploy mitigations, and only to the extent necessary to enable them to do so. The possibility of leaks is a serious challenge. It increases if entities receive sensitive information they do not need to have, or when the number of entities informed is too high. As noted above, it may be necessary to involve governments early in the co-ordination process, for example, when the vulnerability can affect critical activities, which increases the difficulty of keeping the vulnerability secret. Exchanges of vulnerability information should use tools and techniques that reduce the risk of confidentiality or integrity breaches, such as end-to-end encryption.

Balancing patch availability and quality

As explained above, swift development and distribution of patches is key to minimise users' window of exposure to code vulnerabilities. However, developing reliable patches and related risks can take time, in particular in multi-party CVD. Code owners need to ensure that proposed patches are complete and effective based on sufficient tests in different usage scenarios and technical configurations prior to being distributed, and system owners have to test them prior to deployment. For example, some initial patches for the 2018 Spectre and Meltdown microprocessors vulnerabilities had negative effects on performances and compatibility with certain anti-virus software. While addressing software vulnerabilities can take a few weeks, hardware vulnerabilities may require 6 months or more. Rigid mandatory deadlines for patch availability and application are likely to impede code owners' ability to assess mitigations-related risks and create undesired side effects.

Establishing trust through good communications and clear expectations

Trust between vulnerability owners and researchers can be difficult to establish in a CVD process. The power dynamics is not in favour of security researchers. Vulnerability owners, including governments, can easily use legal threat and coercion against researchers. Many organisations can feel threatened by a vulnerability report, in particular when their communications and legal departments, as well as high-level decision makers have a low level of digital and digital security maturity. This is often the case in particular with IoT manufacturers who are new to the digital world. Researchers have less power but they can communicate anonymously with vulnerability owners and disclose the vulnerability unilaterally to the public, which can give them some leverage but it needs to be used carefully. Co-ordination is primarily a matter of communication and information exchange between parties, as well as management of expectations. Both parties can reach out to a co-ordinator to help mend relationships and resolve tensions. Well-recognised standards and good practices can be used as a basis for all stakeholders to develop a CVD culture.

Overcoming co-ordination complexity

CVD can become particularly complex when many code owners are involved ("multi-party CVD"). Most modern software includes pre-existing third-party components, modules, and libraries from the open source and commercial software worlds. Complex products' value chains increase CVD complexity. For example, a vulnerability can affect a product included as a component in one or many (e.g. hundreds) other products, making it difficult to know which parties can be affected. A researcher can report a vulnerability to the entity owning the product's brand or vendor, which may not own the responsibility to address the vulnerable layer of code or component, developed by someone else. In some cases, such as the Spectre microchips vulnerability, CVD may entail a broad collaboration within the product ecosystem to validate the vulnerability, develop, test and finally deliver mitigations to end users. The code owner may be located several steps away from the consumer down the value chain, and there may not be a communication channel across these layers. This can make uncertain and complex the co-ordination of vulnerability handling, including the distribution of a mitigation. International standards provide guidance on how to manage multi-party CVD, for example suggesting that the affected or best-positioned code owner should lead the co-ordination effort, and that a third-party co-ordinator can assist in setting up a broad collaboration within the concerned ecosystem (see below).

3.2. Tools to facilitate CVD

Vulnerability Disclosure Policies (VDP)

A VDP is an essential tool to invite researchers to send reports, increase their confidence that reports will be handled seriously, and to reduce the risk of legal action. The VDP helps clarify researchers' expectations by setting clear rules of the game. 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. A basic VDP can be as short as a single paragraph indicating how to send a vulnerability report securely to the vulnerability owner. A more typical and more effective VDP explains the contact method for secure communication; preconditions for reporting parties; clear expectations for handling a report; and methods for rewarding a report. It also details the scope of the policy, i.e. the list of assets explicitly allowed and not allowed for testing; and conditions under which a researcher can disclose the details of a vulnerability to third parties. A VDP also 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. This is particularly useful to encourage reporting, including from researchers located in other jurisdictions. While

VDPs are a key 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.

If relevant, the policy can cover rules related to rewards, including their conditions and nature. Non-monetary rewards, credits, acknowledgements and marks of prestige can be particularly appreciated and act as an important incentive. For example, NCSC-NL's VDP has been particularly successful in the Dutch security researchers' community by rewarding researchers with a humorous T-shirt marked "I hacked the Dutch government and all I got was this lousy T-shirt", or a cup with a plate saying "I hacked the Dutch tax administration and never got a refund". Similarly, the Korean government credits security researchers in the public description of the vulnerability they report and includes them in a "hall of fame".

Co-ordinator

Stakeholders can turn to a co-ordinator to facilitate co-ordination, namely a trusted third party such as a Computer Emergency Response Teams (CERTs) or Computer Security Incident Response Teams (CSIRTs), or any other entity in a position to facilitate the process. A co-ordinator can assist in a variety of cases, from easing the researcher-vulnerability owner relationship, to orchestrating complex multi-party co-ordination, including vulnerability handling in the context of complex product value chains. 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.

Trust in the co-ordinator has multiple facets: technical competence, neutral judgement, strict respect of confidentiality, ability and capacity to make a balanced assessment of the reasonableness of the various parties' claims and demands, possibility to interact with trusted stakeholders across borders. The co-ordinator may receive confidential information from all parties and facilitate mutual understanding without sharing such information 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. 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 (cf. Box 3). Experts have highlighted cases of international co-ordination where governments have leaked sensitive 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.

Standards and good practice

Standards are a useful means to facilitate 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. ISO/IEC 29147 on Vulnerability disclosure provides requirements and recommendations to code owners (called vendors) on the disclosure of vulnerabilities in products. 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. There is currently no ISO/IEC standard on vulnerability management, but there are guides and domestic standards.

Annex 1 provides an overview of good practice based a selection of guidance documents that complement international standards.

Broader digital security standards and guides 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 also recently published two reports including voluntary guidance on vulnerability management for IoT devices. One provision of the ETSI Technical Specification on "Cyber Security for Consumer Internet of Things", which may become a higher-level European Standard, focuses on the need to implement a means to manage vulnerabilities reports.

Regulation can play an important role in promoting the adoption of good practice and standards. Current initiatives include draft regulation by the UK Government requiring that all new IoT products adhere to minimum digital security requirements based around the ETSI standard. Product manufacturers would have to provide a public point of contact for vulnerability reporting, and to continually monitor for, identify and rectify security vulnerabilities within their products and services. In 2019, the United States Department of Homeland Security (DHS) released a draft Binding Operational Directive (BOD) requiring federal agencies to develop and publish a vulnerability disclosure policy. The EU Cybersecurity Act, adopted in 2019, gives a mandate to ENISA to assist EU Members in establishing and implementing vulnerability disclosure policies on a voluntary basis. The Act also establishes a framework for EU cybersecurity certification schemes including rules on how to report and deal with vulnerabilities. A draft *Common Criteria based European candidate cybersecurity certification scheme* (EUCC scheme) released in July 2020 contains harmonised conditions for vulnerability handling and a fast track assessment procedure for patches.[7]

Bug bounty programmes and platforms

Bug bounty programmes (BBPs, or "bug bounties") 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. A bug bounty policy details the contractual rules governing the programme. Vulnerability owners can set precise expectations, as they would when purchasing security test services from a security firm. Bug bounty hunters are motivated by income as well as education, challenge and pleasure, according to surveys. BBPs bridge the gap between discoverers and vulnerability owners, structuring the CVD process. They can also contribute to drain the black market by providing an alternative option to researchers motivated by monetary gains, provided that they are established only under the authority of vulnerability owners, as opposed to grey market brokers.

Bug bounty platforms operate as a marketplace intermediary serving vulnerability owners and researchers. Vulnerability owners can use a web interface to rapidly design a VDP or a public or private bug bounty policy and publish it on the platform. Security researchers access a dedicated interface where they can search for ongoing programmes, submit their report, interact with the vulnerability owner, receive rewards as well as get visibility and reputation credits. Several platforms leverage their researchers community to offer additional services such as network penetration tests, red teaming, or standards compliance checks. Platforms are for-profit entities (e.g. HackerOne, Bugcrowd, SynHack, YesWeHack, Intigriti) with a few exceptions such as Open Bug Bounty and Trend Micro's Zero Day Initiative.

Many organisations underestimate the cost of launching a BBP. Maturity as well as human and financial resources are key conditions to create a BBP. The ability to handle or manage vulnerabilities in a rapid, consistent, effective and predictable manner are key preconditions for launching a BBP. **Vulnerability owners should approach BBPs as one tool among many others**, and use it in conjunction with others rather than as a turnkey security solution or digital security panacea. BBPs will not replace but are complementary to software code reviews, audits and network penetration tests, for instance. BBPs can motivate internal staff to think about security, and help keep a high-level of security awareness. They can act also as a recruitment tool, facilitating the identification of talents and helping researchers to select companies in line with their expectations. Nevertheless, a BBP 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 resources that could be allocated to preventative "security-by-design" measures, then BBPs might be counterproductive in the medium to long term.

3.3. Legal risk for researchers is a major impediment to CVD

Legal risk faced by researchers when they report a vulnerability to a vulnerability owner is one of the most significant obstacles to CVD. It is enabled by an overall legal environment that does not sufficiently protect security researchers and by the behaviour of many vulnerability owners threatening security researchers with legal proceedings when receiving reports. Overall, this situation creates a power imbalance and a chilling effect, limiting the adoption of CVD and undermining its potential benefits.

Areas of legal risk

Criminal law

According to the Cybercrime Convention, intentionally accessing a computer system without right is a criminal offense. Interpretations of this provision vary significantly across parties to the Convention. In Canada and Chile, cybercrime legislation requires a malicious intent for access to constitute a criminal offense. In the European Union, the Cybercrime Directive (2013/40/EU) sets minimum protections, leaving EU members to adopt stricter rules if they wish. Reflecting the Convention's provision, the United States' Computer Fraud and Abuse Act (CFAA) provides that it is illegal for an individual to intentionally access a computer without authorisation. However, system owners can authorise access, for example through a VDP, enabling safe harbours for security researchers. In 2017, the United States' Department of Justice published a framework to assist organisations instituting a formal vulnerability disclosure programme. Several State-level computer crime laws are also in force across the country.

Another risk of criminal legal proceedings can arise when a vulnerability owner interprets as an extortion attempt the behaviour of a researcher exploring the possibility of a reward for their work. Some vulnerability owners use this as a means to tell researchers that they would fall afoul of extortion law unless they provide them immediately with vulnerability information for free, and sign a non-disclosure agreement (also discussed below), which then muzzles them if they wish to disclose the vulnerability later in the public interest.

Intellectual property law

Researchers can breach copyright law when the information disclosed contains portions of copyrighted code. Copyright protection could restrict sharing vulnerability information with the original code owner, making CVD difficult to implement. The United States' Digital Millennium Copyright Act (DMCA) includes anti-circumvention requirements originally designed to protect media publishers against unauthorised copyright violations, which were interpreted to encompass a wide range of software protection mechanisms typically encountered when performing security audits. Regular updates to the DMCA have included exemptions for security testing under certain circumstances. The vulnerability owner can establish such exemptions, enabling the possibility of safe harbours. Several experts however point out the current exemptions are insufficient. Nevertheless, there are no such exemptions in European law. The DMCA is a federal law. Several State-level laws are also in force in the United States.

Some other countries adopted legislation similar to the initial version of the DMCA as a result of trade agreements. In countries lacking domestic expertise, resources and multi-stakeholder dialogue, these laws have been rarely updated to reflect subsequent improvements in the US law, and typically lack exemptions for research and security purposes.

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. Lastly, 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.

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. For example, there is an exemption for academic research in the Japanese privacy law but not in the EU General Data Protection Regulation (GDPR).

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. Typical issues include VDPs contradicting End User Licence Agreements (EULA) or terms of use, conflicts between a bug bounty platform terms of service and its customers' bug bounty policies, etc. Confusing contractual requirements may result from internal miscommunications on the part of the vulnerability owner between the legal and security teams, suggesting that vulnerability treatment is often not sufficiently integrated in products and organisations' 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.

Cross-border legal risks

Security researchers often have to share vulnerability information across borders as vulnerability owners or other stakeholders may be located in other jurisdictions, for example in the case of vulnerabilities in cloud services. The increased complexity stemming from cross-border aspects and conflicts of jurisdiction can aggravate legal risk.

Export controls legislation and regulation is often cited as a legal risk for researchers as it may apply to tools, techniques and even knowledge typically used to discover vulnerabilities. Part of Wassenaar arrangement, the overarching international framework for export controls, could be interpreted as covering digital security technologies used for reverse engineering, as well as vulnerability information. In 2018, a modification of the Arrangement explicitly excluded vulnerability disclosure and incident response from the technologies concerned. However, some experts have criticised this improvement as not providing enough protection.

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.

Further work might be needed to better understand the scope and scale of cross-border legal risk faced by security researchers, as well as how it could be mitigated.

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 and governments, including for political reasons in some countries. Legal threat, even without prosecution, is sufficient to undermine researchers' ability to publish research, and to create a chilling effect acting as a powerful disincentive for CVD.

Box 4. Examples of researchers threatened legal proceedings

These examples illustrate cases of researchers threatened with legal proceedings.

In 2011, the Finnish online game platform Habbo (273 million users in 150 countries) brought criminal charges against a teenager who reported how he could log into the site's helpdesk system. Two years later, the courts ruled that there was no case to answer.

In 2013, academic researchers informed a chip manufacturer about weaknesses in a chip widely used in immobilisers for various brands of cars. The same year, a British court, acting at the request of Volkswagen, ruled that the scientific article detailing the vulnerability had to be withdrawn. Two years later, Volkswagen ultimately agreed to the release of the publication.

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

In 2016, researchers received a cease-and-desist letter three days after reporting a serious vulnerability to the global consulting and auditing company PwC. 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.

In 2017, a Danish citizen discovered a vulnerability in a municipality web site that enabled the harvesting of personal information of any citizen by entering their birth date in a form. He reported the vulnerability to the municipality. The service provider discreetly fixed the vulnerability and reported the researcher to the police.

In 2018, the FBI investigated a university student who had been reported by the mobile voting company Voatz for illegally attempting to hack its application. In 2020, MIT researchers uncovered vulnerabilities in Voatz's e-voting 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 Cybersecurity and Infrastructure Security Agency (CISA). The company disputed the severity of the vulnerabilities, making aggressive public statements against the researchers. Ultimately, an independent audit requested by Voatz confirmed the MIT's findings.

Few countries have adopted measures creating a favourable legal environment for security research and CVD. According to research by the Centre for European Policy Studies (CEPS), the Netherlands, France and, to a lesser extent, Lithuania were the only EU members providing some protection to researchers. The Netherlands appears as the leading country, with guidelines by the Public Prosecution Service on how to decide whether to initiate a criminal investigation and/or to prosecute taking into account the general concept of CVD or compliance with an existing CVD policy. In the United States, efforts are needed to overcome the existing chilling effect, despite the CFAA and DMCA enabling VDPs to create safe harbours.

Latvia tried to update its legal framework to increase researchers' protection, but failed because some parts of the government did not understand CVD or were not willing to support it.

While many experts agree on the need to address the chilling effect, there is limited agreement on how to proceed. The following elements need to be taken into account when developing new policies

First, as other stakeholders, researchers should take responsibility for their action when engaging in products and systems vulnerability testing and when disclosing their findings. Digital security risk reduction cannot be 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.

Second, safe harbours can protect researchers under certain conditions defined by the vulnerability owner, generally in a VDP. However, a safe harbour is never an absolute protection against legal risk. Safe harbours do not create a blanket exemption for researchers, and do not mean that they will not be sued, even if they respect the VDP. For example, a VDP can indicate that the organisation will not sue the researcher, but 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. Researchers need to understand the limited protective effect of safe harbours for not to overstep the permissions they are granted. Additional policies and further coherence of approaches across countries may be needed to complement safe harbours.

Third, to develop measures that better protect researchers in their jurisdiction, policy makers need to analyse the gaps in their legal frameworks that create legal uncertainty. However, levels of risk for researchers vary according to many factors such as the legal area concerned (e.g. testing online systems exposes to cybercrime lawsuit, reverse-engineering products exposes to intellectual property lawsuits), the possible cross-border nature of the case, the presence of a VDP, the intervention of a co-ordinator, etc. Policy makers can use different basic scenarios with variations around these factors to better understand the risk and how to fill the gaps. For example, in some countries, it may be necessary to amend legal frameworks 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.

4. Public Policy Guidance

This section provides possible guidance to help policy makers facilitate vulnerability treatment. As stakeholders involved in CVD are often located in different countries, it is expected that such guidance would help reduce potential fragmentation of approaches across jurisdictions and contribute to reducing digital security risk globally.

Public policy development and implementation in this area should leverage all stakeholders' communities. In addition to government agencies, communities include:

- Supply-side actors: businesses developing goods and services that contain code;
- Demand-side actors, including public and private organisations managing information systems (based on goods and services that contain code);
- Digital security researchers working in a personal or professional capacity, in an academic, business or not-for-profit context, with a professional digital security background or not ;
- Bug bounty platforms;
- Civil society;
- Other stakeholders such as lawyers and insurers.

A good trusted relationship between the government and sufficiently organised technical, business and civil society communities can greatly facilitate policymaking, including the assessment of current frameworks and identification of the best path to policy improvement.

The suggestions below are interrelated. They are not provided by order of importance or priority.

4.1. Sharing a common understanding

Changing the culture and raising awareness

Breaking the "vulnerability taboo"

Organisations that view digital security as an ideal state where there are no vulnerabilities are unlikely to welcome researchers who report vulnerabilities in their products or systems. This may be the case in particular for organisations with little experience of the digital environment and digital security, currently accelerating the digital transformation of their business, such as IoT manufacturers. To successfully embrace digital transformation, the business leadership of these organisations needs to adopt a culture recognising that:

- *All products that contain code also contain vulnerabilities, and all information systems have a high likelihood of containing vulnerabilities* related to misconfiguration or unpatched software, including firmware. Breaking this "vulnerability taboo" is an essential first step to promote vulnerability treatment.

- *Digital security is a continuous effort to manage risk* rather than a state that is reached once and for all. The digital environment is fundamentally dynamic. It is also open by default and closed by exception. Security must be agile and based on systematic processes.
- *Organisations' business leaders need to own the risk and cannot simply delegate it to technical security experts* because digital security is an economic and social risk management issue rather than only a technical problem. The consequences of security incidents affect revenues, competitiveness, business operations, reputation, innovation, safety, privacy and trust, in addition to breaching the availability, integrity, and confidentiality of data, systems and networks.

In the digital era, leaders and decision makers need to change the way they think they can create trust with their customers and partners. Organisations should abandon the idea of a perfectly secure digital environment, recognise that their products and information systems can be vulnerable, and demonstrate that they are responsible for monitoring vulnerabilities and swiftly addressing them.

Policy makers' role

Policy makers, jointly with other stakeholders, have a key role to play. They can help change mindsets about vulnerabilities, and encourage vulnerability owners to take responsibility. 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. Furthermore, security researchers would benefit from forming a more organised community, including developing a public discourse that can reach and influence policy makers. Their voices need to be heard, including at the international level.

Such an evolution of mindset will take time, but the consequences of slow action or inaction will become increasingly severe as digital transformation unfolds at a fast pace. 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. Initiatives such as the Cybersecurity Tech Accord and the Paris Call for Trust and Security in Cyberspace are steps in the right direction. Bug bounty programmes and public recognition of vulnerabilities by prestigious but rather traditionally conservative public institutions contributes greatly to changing mindsets in the society.

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, the recent US NIST report for IoT manufacturers describes recommended digital security activities they should consider before their IoT devices are sold to customers.

Public policy needs to reward organisations that adopt a good vulnerability treatment practice, including systematic vulnerability handling or management, transparent vulnerability disclosure policies encouraging vulnerability discovery and reporting, providing a safe harbour for researchers, and engaging in effective CVD processes resulting in swift resolution and disclosure. Awareness-raising campaigns, labels and public procurement can be effective tools to achieve these goals.

Clarifying roles and responsibilities to treat vulnerabilities

Policy makers need to take a broader approach to how all stakeholders "treat" vulnerabilities rather than focusing primarily on disclosure. 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 researchers, ignoring that, from a public policy perspective, the other lifecycle stages are equally important, and even prerequisite for effective disclosure.

All stakeholders involved in vulnerability treatment should take responsibility and be accountable, based on their role, ability to act and the context, for addressing vulnerabilities and sharing information in a timely manner, and for taking into account the potential impact of their decisions on others. In practice, it means that:

- On the supply-side, code owners should take (i.e. own) responsibility for vulnerabilities in products they put on the market. They should adopt a security development lifecycle whereby vulnerability handling is an integral part of basic product maintenance and support;
- On the demand-side, system owners should take responsibility for the vulnerabilities in the information systems they manage. They should 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.
- Security researchers should take responsibility for their actions.

Annex 1 provides good practice for CVD.

All stakeholders should aim at disclosing vulnerabilities in a co-ordinated manner. All vulnerability owners should have a VDP and a vulnerability disclosure process in place to receive and address vulnerabilities spontaneously reported to them (cf. best practice above). 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. Security researchers should do their utmost to co-ordinate with vulnerability owners the disclosure of vulnerabilities they find.

Governments need to ensure that agencies and processes can be trusted by stakeholders. 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 it or use it for offensive purposes. In some countries, this may require adjusting the public governance for digital security, in particular when mandatory reporting of vulnerabilities to the government is considered.

4.2. Mainstreaming good practice

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 requiring US Government federal bodies to develop and publish a vulnerability disclosure policy is an example of such an approach.
- *Adopting vulnerability handling and management within the government.*
- *Leverage public procurement to encourage CVD as well as vulnerability handling and management*. 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.

Including vulnerability treatment in regulation and guidance

Governments can promote vulnerability treatment by including vulnerability management, handling and CVD in regulation, standards and guidance, or using it as indicators of compliance. This may include, for example, product regulation; regulation related to critical activities, such as the EU NIS Directive (currently being reviewed); certification schemes, such as those established by the EU Cybersecurity Act; government-supported standards (e.g. NIST Cybersecurity Framework 1.1, ETSI Technical Specification "Cyber Security for Consumer Internet of Things"); IoT regulation (e.g. draft UK regulation); and privacy regulation (e.g. EU General Data Protection Regulation and the US Health Insurance Portability and Accountability Act). 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.

Voluntary rather than mandatory reporting of vulnerabilities related to critical activities is the generally preferred approach among experts. They 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. Since governments can play an ambiguous role in this area (cf. Box 3), they must build trust and demonstrate that vulnerability information they receive is handled appropriately. Mandatory reporting to a government entity would require the receiving entity to operate in a transparent manner, with the sole objective of remediating vulnerabilities, and independently from other government entities. 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 are communicated to the authorities has raised this type of concerns.

Providing tools, encouraging standards development and adoption

All stakeholder groups can facilitate CVD adoption with template Vulnerability Disclosure Policies (VDPs), quick start guides, and other best practice documents. These can be aimed at specific communities to address their needs and concerns. These documents 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 and the NCSC-NL Guidelines provide examples of such initiatives.

4.3. Fostering trust and removing obstacles

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 the CVD process. 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 stakeholders. For example, they can turn to a foreign, regional, international 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 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 the OECD 2019 Recommendation on digital security of critical activities).

Co-ordinators should be trustworthy, for example by (in no particular order):

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

Protecting researchers

Policy makers need to change 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. 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 vulnerability owners and researchers. A key aspect of the power imbalance is the legal pressure that the former can place on the latter.

Governments can 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 in the process (e.g. revision of a national digital security strategy, or development of an implementation plan).

Governments could consider developing and agreeing at the international level 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.

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. 3.3).

Addressing the grey market

Governments can 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 Box 3. As a first step, more research would be needed to better understand the size of the grey and black market, as well as the incentives and disincentives that could be leveraged to change actors' behaviours. Analytical 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. One option to drain this market would be, for example, to ensure that bug bounties could only be 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. Governments could also co-ordinate the regulation of this market at the international level. Another more thought provoking idea would be to establish an international fund to buy vulnerabilities, at least for open source software. More generally, an international dialogue covering prosperity and sovereignty issues seems essential to address the issue, as these markets are borderless by nature.

4.4. Encouraging international co-operation and standards

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

- Governments can 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 can 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).
- Governments could work together to develop a common ground for addressing the grey market.
- The Budapest Convention could be updated to add safe harbour provisions for researchers.

References

Full bibliography of all references supporting this work steam is available in (OECD, 2021[2]).

ENISA (2018), *Economics of Vulnerability Disclosure*, https://www.enisa.europa.eu/publications/economics-of-vulnerability-disclosure (accessed on 31 January 2020). [8]

Frei, S. and O. Rochfort (2021), *The Case for a Bug Bounty Program of Last Resort*, https://techzoom.net/bug-bounty-reloaded/. [7]

OECD (2021), *Encouraging vulnerability treatment: background report - Responsible management, handling and disclosure of vulnerabilities*, https://one.oecd.org/document/DSTI/CDEP/SDE(2020)3/FINAL/en/pdf. [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 (2020), *Seven lessons learned about digital security during the COVID-19 crisis*, https://www.oecd.org/coronavirus/policy-responses/seven-lessons-learned-about-digital-security-during-the-covid-19-crisis-e55a6b9a/. [6]

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). [9]

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 (2019), *Summary Report of the Inaugural Event Global Forum on Digital Security for Prosperity*, https://doi.org/10.1787/20716826. [10]

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]

Annex 1. Good practice for CVD

This annex brings together existing good practices for CVD from a set of guidance documents listed hereafter. It aims to inform public policy makers. Technical experts are invited to consult appropriate and up-to-date technical standards.

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 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. 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 to 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 swiftly corrected 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 that 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, report and disclose a vulnerability. 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.

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.

4. Adopting a vulnerability disclosure policy

At 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.

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 communication 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.

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 for example 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.

The above overview of good practice is based on an analysis of the following documents:

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

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 to the post-release (cf. ISO/IEC 30111).
- **Vulnerability management**: ongoing processes enabling all organisations to know if vulnerabilities are present within their IT estate 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, vulnerability 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.

Notes

[1] A red team is a contracted or in-house team emulating an advanced threat actor to identify gaps in the organisation's defensive strategy (as opposed to the blue team of defenders).

[2] Cf. (OECD, 2021[4])

[3] As defined in (OECD, 2021[3]) and (OECD, 2021[4]).

[4] There are numerous possibilities of failure. For example, the researcher may not find the vulnerability owner; the vulnerability owner may not be responsive or understand the report or have a policy; the researcher may contact the product maker while the vulnerability is located in a component developed by another party, requiring additional communications flows; the researcher may feel threatened by the vulnerability owner or/and vice versa; the communication between them may be difficult and lead to confusion; the vulnerability owner and researchers may disagree; the researcher may decide to monetise the vulnerability on the grey or black market (see below); etc.

[5] ISO/IEC 30111 on vulnerability handling 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.

[6] Cf. for example the 2019 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[9]).

[7] The articulation between standards, certification and digital security of products in general is further discussed in (OECD, 2021[3]) and (OECD, 2021[4]).