August 2018 · Dr. Sven Herpig

Governmental Vulnerability Assessment and Management

Weighing Temporary Retention versus Immediate Disclosure of 0-Day Vulnerabilities

A proposal supported by the Transatlantic Cyber Forum

Stiftung Neue Verantwortung

Think Tank at the Intersection of Technology and Society

Introduction

Government-led acquisition, assessment and management of vulnerabilities to enable offensive cyber operations and a wide array of information gathering efforts is one of the most topical debates in cyber security today. Vulnerabilities are at the core of major discussions, namely hacking by law enforcement, private sector hacking back, military cyber operations and intelligence collection, because they greatly increase the operational value and efficiency of those activities. Mitigating and patching vulnerabilities is crucial, however, for protecting both public and private sector networks and critical infrastructures. In addition to government entities and critical infrastructure operators, the private sector and the general public have vested interests in this issue. Governments for their part are not only potential exploiters of vulnerability information, but also users of the technologies that may be impacted from a security standpoint by a failure to patch discovered vulnerabilities. In addition, they are regulators of entities required to protect critical infrastructure and sensitive personal information. Hardware vendors, software vendors and providers of online services have an incentive to identify and fix vulnerabilities in their hardware, software and online services, in order to provide a secure service and prevent reputational and financial harm. The private sector and the public would like to use secure devices and services to avoid falling victim to cyber espionage and crime, as well as to simply communicate freely and confidentially[1]. In general, the internet ecosystem benefits from patching vulnerabilities, and government policy should be to disclose them unless there is a specific, justifiable reason for retaining and using them in law enforcement, intelligence or military programs. Therefore, it is paramount to assess and manage the tradeoffs and various equities of civil liberties, commerce, public safety and IT security. The underlying mission statement has been eloquently summarized in the Vulnerabilities Equities Policy and Process for the United States Government (VEP): “The primary focus of this policy is to prioritize the public’s interest in cybersecurity and to protect core internet infrastructure, information systems, critical infrastructure systems, and the US economy through the disclosure of vulnerabilities discovered by the USG, absent a demonstrable, overriding interest in the use of the vulnerability for lawful intelligence, law enforcement, or national security purposes"[2].

1 Kate Musgrave, In Repressive Countries, Citizens Go ‘Dark’ to Share Independent News

2 White House, Vulnerabilities Equities Policy and Process for the United States Government (UNCLASSIFIED) https://www.whitehouse.gov/sites/whitehouse.gov/files/images/External%20-%20Unclassified%20VEP%20Charter%20FINAL.PDF

In November 2017, the United States Government published its VEP charter, which outlines the organizational structure, processes and respective indicators/equities which are to be applied to government-held vulnerabilities. Though details about this process still remain classified, the VEP publication set a new bar when it comes to transparency of governments around vulnerability handling and disclosure. The longstanding discussion about the VEP, even when the mechanism was still mainly discussed behind closed doors[3], as well as other information about international approaches[4], provided important source material for the development of this paper, and are relevant to governments around the world.

The main focus of the Transatlantic Cyber Forum's working group on encryption policy and government hacking is to urge the adoption of publicly disclosed policies for vulnerability handling and disclosure in the German and EU[5] debates, while continuing to identify and advocate for further improvements to the existing process in the United States. This paper highlights principles and criteria that could guide discussions on implementation across different countries. An overarching theme is that governments should develop processes that are weighted towards disclosure, with retention being authorized under specific circumstances and only for limited periods of time. The focus of these policies should be on "when" and "how" disclosure should occur rather than "whether" and "if".

3 Consolidated materials, e. g. those made available by EFF through FOIA https://epic.org/privacy/cybersecurity/vep/

4 Aspen Institute, Cyber Breakfast: The View from the White House with Rob Joyce

5 Mirja Gutheil et al., Legal Frameworks for Hacking by Law Enforcement: Identification, Evaluation and Comparison of Practices

Acknowledgement

This analysis has been supported by members of the Transatlantic Cyber Forum through online collaboration and joint workshops in Washington, D. C. and Berlin. The views and opinions expressed in this paper are those of the author and do not necessarily reflect the official policy or position of the working group members or that of their respective employer/s. Acknowledging essential contributions of:

1. Cathleen Berger, Mozilla
2. Ulf Buermeyer, Gesellschaft für Freiheitsrechte e.V. (GFF)
3. Betsy Cooper, Independent Consultant
4. Alan Duric, Wire
5. Marc Fliehe, Verband der TÜV e.V. (VdTÜV)
6. Sharon Bradford Franklin, New America's Open Technology Institute
7. Andrew Grotto, Stanford University
8. Trey Herr, Microsoft
9. Karsten-Kai König, CIPHRON
10. Andreas Kuehn, EastWest Institute
11. Susan Landau, Tufts University (affiliation for identification purposes only)
12. Daniel Moßbrucker, Reporters Without Borders Germany
13. Jan Neutze, Microsoft
14. Riana Pfefferkorn, Stanford Center for Internet and Society
15. Thomas Reinhold, cyber-peace.org / Institute for Peace Research and Security Policy Hamburg
16. Michelle Richardson, Center for Democracy and Technology
17. Volker Roth, Freie Universität Berlin
18. Julia Schütze, Stiftung Neue Verantwortung
19. Ari Schwartz, Cybersecurity Coalition
20. Megan Stifel, Public Knowledge
21. Eric Wenger, Cisco Systems, Inc.
22. Jessica Zucker, Microsoft
23. Christoph Zurheide, Deutsche Post DHL Group

Table of Contents

1. Vulnerabilities

Understanding "vulnerabilities" is a first and essential step towards assessing and managing them. According to ISO/IEC standard 30111:2013, vulnerabilities are "weaknesses of software, hardware, or online service that can be exploited[...] Regardless of cause, an exploitation of such vulnerabilities may result in real threats to mission-critical information systems"[6]. The United States National Telecommunications and Information Administration (NTIA) similarly states that "vulnerabilities are weaknesses of software, hardware, or online services that can be used to damage the confidentiality, integrity, or availability of those systems or the data they store. Finding these vulnerabilities and informing affected parties is essential to protect our economy and citizens"[7]. The White House National Security Council (NSC) defined vulnerability in a 2017 document as "a weakness in an information system or its components (e.g., system security procedures, hardware design, internal controls) that could be exploited or impact confidentiality, integrity, or availability of information"[8].

For the scope of this paper, a vulnerability is therefore defined as "a flaw in hardware or software which individually or chained with other vulnerabilities enables an exploitation by third parties to perform by the owner of the system unauthorized - and possibly covert - operations on one or many devices or online services." While there are several aspects that allow for further refinement of categories of vulnerabilities - such as the software or hardware they affect, their severity[9] and characteristics[10] - the most basic distinction is made between 0-days (or unknown vulnerabilities) and n-days (or known vulnerabilities).

6 ISO, ISO/IEC 30111:2013(en) Information technology — Security techniques — Vulnerability handling processes https://www.iso.org/obp/ui/#iso:std:iso-iec:30111:ed-1:v1:en

7 National Telecommunications and Information Administration, Vulnerability Disclosure Attitudes and Actions

8 White House, Vulnerabilities Equities Policy and Process for the United States Government (UNCLASSIFIED) https://www.whitehouse.gov/sites/whitehouse.gov/files/images/External%20-%20Unclassified%20VEP%20Charter%20FINAL.PDF

9 CVE Details

10 Such as whether the vulnerability can be exploited remotely over the internet or just locally when in possession of the vulnerable device. The severity and characteristics are important aspects of the vulnerability assessment and therefore further discussed in the respective sections of this paper.

Unknown vulnerabilities

0-days are vulnerabilities that exist unbeknownst to the respective "maintainer"[11] and the public[12]. Thus, no patch or mitigation method[13] has been provided by the maintainer. It makes this kind of vulnerability very potent and dangerous. More than one actor (e.g. intelligence agency) might know about the same vulnerability without being aware of any other party also knowing about it. This could allow for other intelligence agencies or criminals to exploit the 0-day. However, the vulnerability remains a 0-day vulnerability until it is known by the maintainer.

Known vulnerabilities

N-days, on the other hand, are vulnerabilities which are known to the maintainer. This does not necessarily mean that there is already a patch created by the maintainer, or that the maintainer intends to patch the vulnerability. This is typically seen in cases where the respective product reached end-of-life support (EOL), there is no central maintainer, the maintainer does not have the capacity to fix the vulnerability, the maintainer does not exist anymore[14], or the maintainer is simply not interested in fixing the vulnerability,

11 The "maintainer" terminology is borrowed from the open source environment and in this paper refers to the actor that can design a patch. Therefore applicable to either the manufacturer/ vendor of hard- or software, the provider of an online service or the persons/ community in charge of open source projects.
David "cdlu" Graham, OLS: Kernel documentation, and submitting kernel patches

12 (Prominent) static checkers can be used on vulnerable code to determine if it is an unknown vulnerability or not.
White House, Vulnerabilities Equities Policy and Process for the United States Government (UNCLASSIFIED) https://www.whitehouse.gov/sites/whitehouse.gov/files/images/External%20-%20Unclassified%20VEP%20Charter%20FINAL.PDF

13 Mitigation refers to "[...] detection and protection strategy used to safeguard networks, servers and applications by IT administrators in order to minimize the effect of malicious traffic and intrusion attempts while maintaining functionality for users" according to radware, Attack Mitigation. This is independent from the availability of a patch provided by the maintainer in whose hard-/ software or online service a vulnerability is found. There is not always a workaround that mitigates the effect of a vulnerability. Example for a mitigation method: if a vulnerability is known to exist in the macro handling of an word processing software (such as Microsoft's Word or LibreOffice's Writer), the administrators in a Windows operating software environment could disable the macro features through adjusting the respective group policies. As a result, no one in that network would be able to use macro functionality anymore but that would also mean that the vulnerability could not be exploited anymore by an attacker. In most cases, this might just be a short-term fix until the patch is available or security software is adjusted to the threat.

14 Dan Geer suggested at BlackHat 2014 that "[...] if Company X abandons a code base, then that code base must be open sourced". In that case, other actors could serve as maintainer.
Dan Geer, Cybersecurity as Realpolitik

meaning a patch may never be available to users[15]. Even if there is a patch available, it still has to be tested, rolled out to the affected devices and, outside of managed IT environments, installed by the actual user. The time between the disclosure of a vulnerability to the maintainer and the implementation of the patch is referred to as the "window of exposure"[16].

Implications

Patching is often inconsistent, especially with complex devices or older software. Such patching can be particularly complicated in the case of devices with multiple providers of the underlying software; for example, a 2016 security review by Google has shown that only half of the top-50 Android phone models have received the most recent security updates[17]. When adding the many EOL devices to this equation, the number of non-patched smartphone devices with n-day vulnerabilities becomes extremely worrisome. In other fields, such as IoT[18], ICS[19], and embedded or systems on a chip,[20] the situation might even be more dire, for example because devices lack update capabilities or channels.

This is good news for government agencies wishing to engage in hacking, as they do not necessarily need to acquire knowledge of 0-days - which is much harder and more costly; they can exploit widespread unfixed n-days[21]. For the public's overall cyber security it would also be better if governments would limit themselves to only exploiting n-days. A recent study even concluded that the vast majority of hacking attacks in 2015 exploited n-days which had patches readily available[22]. If security agencies can successfully leverage n-days against their targets, there is less downside to disclosing vulnerabilities, since retaining them may not add operational value. However, there is

15 This reasoning applies mainly to closed source solutions - open source solutions can also be patched by actors other than the maintainer due to the source code being open.

16 Bruce Schneier, Full Disclosure and the Window of Exposure
Trend Micro, Maintaining Vulnerable Servers - What's Your Window of Exposure?

17 Android Security 2016 Year In Review

18 Mario Ballano Barcena and Candid Wueest, Symantec Security Response - Insecurity in the Internet of Things
hp, HP Study Reveals 70 Percent of Internet of Things Devices Vulnerable to Attack https://www.symantec.com/content/dam/symantec/docs/white-papers/insecurity-in-the-internet-of-things-en.pdf

19 European Union Agency for Network and Information Security, Window of exposure... a real problem for SCADA systems?

20 Graz University of Technology, Meltdown and Spectre

21 CVE Details

22 Verizon, 2016 Data Breach Investigations Report http://www.verizonenterprise.com/resources/reports/rp_DBIR_2016_Report_en_xg.pdf

no empirical data available for this specific question about how many times governments, whether Russia, the U.K., China, Israel, the US, Iran or others, relied on 0-days in performing cyber operations, or how many could have been successful through other means. Absent this data, we should assume that disclosing 0-days - so they can be fixed by the maintainer - has a net potential for improving security and an unproven likelihood of increasing the ability of security agencies to access target devices by retaining and exploiting those vulnerabilities.

According to MITRE data, there have been 40 new vulnerabilities per day in 2017[23]. However, "researchers found that 77 percent of the vulnerabilities they analyzed did not have any exploits developed, and only 23 percent of published vulnerabilities had associated exploit code. Just 2 percent of published vulnerabilities have observed exploits in the wild"[24].

It is useful to put the exploitation of 0-days in perspective. David Hogue, a senior technical director for the NSA's Cybersecurity Threat Operations Center, said in April 2018 that "at NSA we have not responded to an intrusion response that's used a zero day vulnerability in over 24 months [...] The majority of incidents we see are a result of hardware and software updates that are not applying"[25]. Andreas Könen, Head of the IT and Cyber Security Department at Germany's Ministry of Interior, Building and Community, said in June 2018 that only about 5% of all vulnerabilities that are used in an exploit are 0-day vulnerabilities; the rest are n-day vulnerabilities[26].

2. Process

2.1 Principles

The vulnerability assessment and management process laid out in this paper is based on the following principles and further actions:

1. A government that wants to temporarily retain and use vulnerabilities in hardware, software and online services for law enforcement, intelligence or military operations must implement an open and transparent assess-

23 Fahmida Y. Rashid, Predict Which Security Flaws Will be Exploited, Patch Those Bugs

24 Fahmida Y. Rashid, Predict Which Security Flaws Will be Exploited, Patch Those Bugs

25 Chris Bing, Nation-state hackers attempted to use Equifax vulnerability against DoD, NSA official says

26 Public panel at the Cyber Security Policy in Germany conference on June 6, 2018 in Berlin. The conference was jointly organized by Germany's Federal Academy for Security Policy and the Stiftung Neue Verantwortung.

ment and management process, ideally in accordance with the provisions recommended in this paper.

2. The process needs to be enshrined in law, with an independent legislative review of its effectiveness and proportionality, or a sunset clause, after five (5) years.
3. The vulnerability assessment and management process laid out in this paper is applied to all 0-day vulnerabilities[27] acquired by all government entities – including hacking tools and services which leverage 0-days acquired by the government.

4. Vulnerabilities are never permanently retained. Their disclosure is merely delayed. Therefore they are temporarily retained by the government and must ultimately be disclosed through close coordination with the "maintainer".

5. There should be a strong presumption that disclosure of vulnerabilities is in the best interest of commerce, civil liberties, public safety, and IT security.

6. A government that wants to retain vulnerabilities must demonstrate a critical need that outweighs the security benefits of disclosure to the "maintainer," and a plan to minimize harm, including adequate protections to ensure a secure retention of vulnerabilities by properly protecting them from any unauthorized access during the period prior to disclosure. The protections should include a mechanism to accelerate disclosure in light of an event indicating that the existence of the vulnerability is known to others.

7. Governments that have implemented such a process should work towards promoting it as an international norm[28] through relevant international coordination mechanisms and cooperation networks.

27 Vulnerabilities that are procured with the purely defensive (read: IT-security) intention do not fall under this process but must only be used for the intended purpose.

28 Kate Charlet, Sasha Romanosky and Bert Thompson, It's Time for the International Community to Get Serious about Vulnerability Equities
Shaun Waterman, Responsible vulnerability disclosure is becoming an international norm
Centre for European Policy Studies (CEPS), Software Vulnerability Disclosure in Europe: Technology, Policies and Legal Challenges https://www.ceps.eu/system/files/CEPS%20TFRonSVD%20with%20cover_0.pdf

8. Research on collision rates[29] needs to be financially supported by the government in order to justify retaining vulnerabilities for their operational value, while at the same time not endangering civil liberties, commerce, public safety and IT security[30].
9. A better understanding of acquisition (e.g. transparency requirements for vulnerability markets and trading) and disclosure (e.g. implementation of coordinated vulnerability disclosure) needs to be developed.

2.2 Scope

Governments need to make difficult choices between vulnerability disclosure and retention; each carries risks. Therefore, governments need clear guidelines, processes and accountability, possibly including negative incentives, to enforce compliance. This paper thereby offers a process that facilitates the respective decision making by governments. This process consists of four stages and three cross-cutting safeguards. The stages are not necessarily linear, especially once the vulnerability assessment enters the assessment-management-disclosure cycle.

The initial stage is the acquisition of vulnerabilities, which involves a government entity having obtained access to a vulnerability which can be used for law enforcement, intelligence or military purposes. The second stage of this process is the institutional setup of this process, followed by the assessment of the vulnerability, which is the main focus of the process. The fourth stage is the management of the vulnerabilities: either disclosing, retaining and/ or mitigating them. In addition, technical and legal safeguards are part of the entire process.

When it comes to acquisition, the model presented does not go into depth regarding the method through which the vulnerability was obtained, such as international intelligence cooperation, security research, or the vulnerability markets. Some work has been done in this area already[31] but it requires more

29 Collision rates provide for an empirical mathematical model to calculate how long it takes two or more actors to find the same vulnerability independently from each other.

30 An economic take on the best time to disclose vulnerabilities has been conducted by Tristan Caulfield, Christos Ioannidis and David Pym, The U. S. Vulnerabilities Equities Process : An Economic Perspective

31 Luca Allodi, Economic Factors of Vulnerability Trade and Exploitation
Nicole Perlroth and David E. Sanger, Nations Buying as Hackers Sell Flaws in Computer Code
Lillian Ablon and Andy Bogart, Zero Days and Thousands of Nights
Joseph Cox and Lorenzo Franceschi-Bicchierai, How a Tiny Startup Became the Most Important Hacking Shop You've Never Heard Of

attention from a technical, ethical, political, economic and legislative perspective; this is a separate discussion from that of a vulnerability management process, however. The scope furthermore excludes details about the coordinated vulnerability disclosure (CVD). CVD "is the process of gathering information from vulnerability finders, coordinating the sharing of that infor-

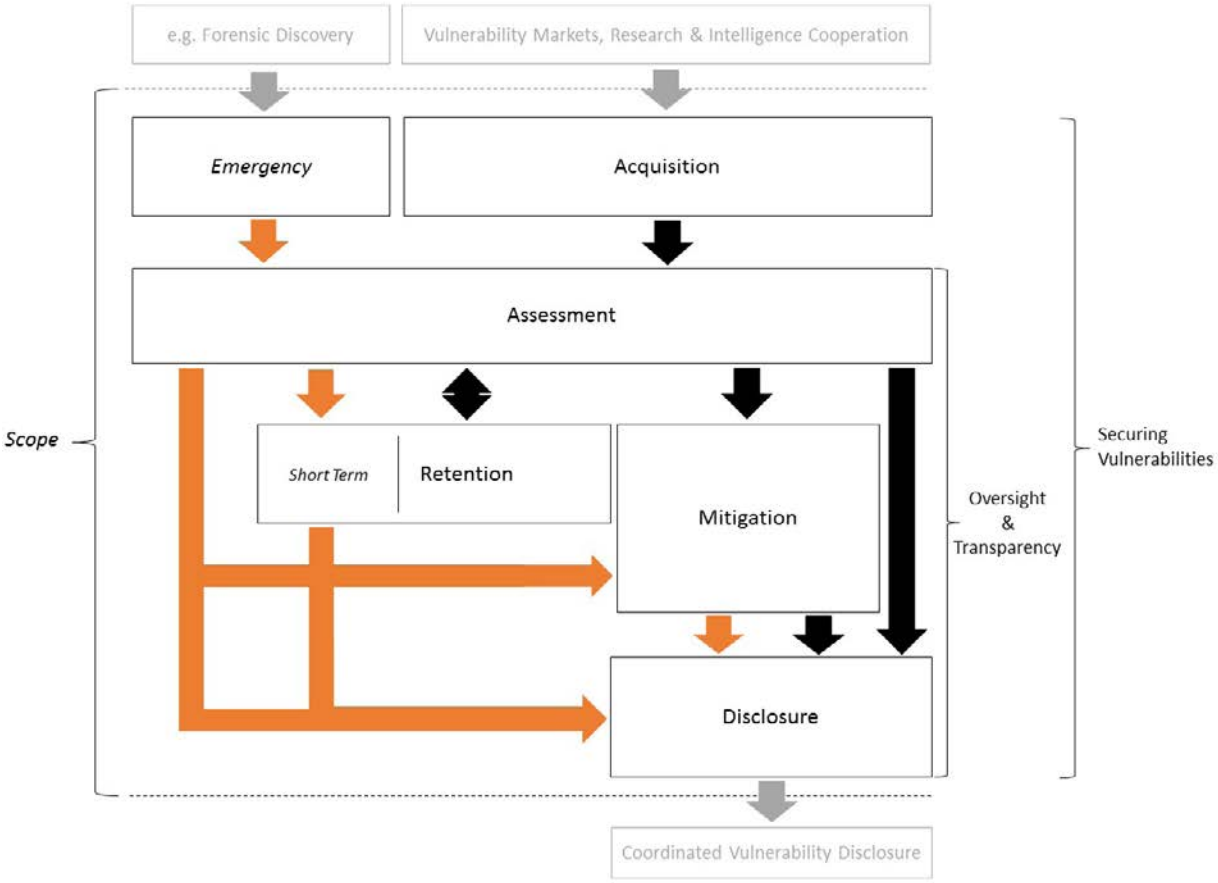

Figure 1. Visualization of the process

mation between relevant stakeholders, and disclosing the existence of software vulnerabilities and their mitigations to various stakeholders including

the public"[32]. A government should report vulnerabilities to and coordinate with maintainers in a way that is consistent with CVD. Governments do have to have a CVD policy and process in place so they can receive vulnerability information about government systems, but they also must report vulnerabilities to maintainers, thereby also functioning in the "finder" role for CVD[33].

2.3 Acquisition

2.3.1 Sources and handling

A government can acquire vulnerabilities through its own research, forensics, from foreign (intelligence) partners, or from the private sector. How a vulnerability is acquired may influence evaluation parameters, and needs to be factored in at that stage. If a vulnerability is discovered through research, it might be a vulnerability which no one else knows about, meaning the likelihood of it being exploited against the country's IT-systems is lower. If the same vulnerability is acquired through forensics work, an exchange with an international partner or through sale of a private vendor, however, other actors necessarily have knowledge of this vulnerability, and might exploit it themselves. In those cases, the threat to equities such as public safety and IT security might be higher than in the aforementioned example. The same process extends beyond acquiring knowledge of a vulnerability and also applies to buying or renting hacking tools and services which exploit an unknown vulnerability.

Every 0-day vulnerability acquired by the government has to go through the government vulnerability disclosure process with one exception: vulnerabilities shared with the government by any entity for the purpose of CVD must not be retained, but instead must directly be reported to the affected maintainer as required by CVD; such government action is critical to not undermine this process and to enable the government to function in a trusted way in the CVD coordinator role, which ensures that vulnerabilities get reported

32 Allen D. Householder, Garret Wassermann, Art Manion and Chris King, The CERT® Guide to Coordinated Vulnerability Disclosure
More information about CVD:
FIRST, Multi-Party Coordination and Disclosure
NTIA, Vulnerability Disclosure Attitudes and Actions
Bundesamt für Sicherheit in der Informationstechnik, Handhabung von Schwachstellen
ISO, ISO/IEC Standard 29147:2014
ISO, ISO/IEC Standard 30111:2013

33 CEPS Task Force, Software Vulnerability Disclosure in Europe Technology, Policies and Legal Challenges https://www.ceps.eu/system/files/CEPS%20TFRonSVD%20with%20cover_0.pdf

to vendors or maintainers when direct reporting from finders to vendors or maintainers is disrupted.

In emergencies, regular retention or retention and mitigation (long-term mitigation with delayed disclosure) are not options. The vulnerability has to be either directly mitigated and disclosed, or retained for a short amount of time (short-term retention) and then mitigated and disclosed. An emergency could be the forensic recovery of an unknown vulnerability that is actively exploited against the state's IT systems. The reason for the short-term retention option is that there might still be conflicting equities at stake, such as an ongoing forensics/attribution operation, which could be exposed if the vulnerability is immediately disclosed.

2.3.2 Third party stipulation

One of the criticisms of the United States VEP is that it does not include a third party stipulation against nondisclosure agreements (NDAs)[34]. If a state acquires a vulnerability from a third party (other government, private sector, independent security researcher) it might be required to sign a NDA effectively prohibiting it from entering the vulnerability into the government vulnerability assessment and disclosure process, of which one outcome might be the vulnerability's disclosure. Thus, NDAs could be leveraged to effectively shut down most of the government vulnerability disclosure process; an exception would be vulnerabilities acquired through the government's own research and forensics. Therefore, vulnerabilities must not be subject to NDAs[35]. This includes government procurement of hacking tools and services which leverage those vulnerabilities. This might give governments which enter into NDAs an edge over those who do not, because companies would rather sell to the former, but monopsony conditions might however still allow governments that reject NDAs on principle to acquire vulnerabilities from third parties, and possibly at a higher price. In order to increase the pressure on vendors to abandon the requirement of NDAs, like-minded states could be convinced to reject NDAs as well.

34 Robert Knake, Grading the New Vulnerabilities Equities Policy: Pass

35 An alternative would be not to prohibit NDAs per se but require it to be the exception and last resort. For those cases, a paper trail has to be created explaining why signing the NDA was the last resort. To maintain this status, the agency in charge of the government vulnerability disclosure process must annually review this paper trail and evaluate whether it can still uphold this policy, add further restrictions to it or prohibit NDAs.

2.4 Structure

2.4.1 Equities

When creating the institutional structure for the government vulnerability assessment and management process, the roles and responsibilities assigned to each institution must align with their other roles within their respective political system. For this reason, institutional setup is one of the toughest challenges in designing such a process. To ease communications and organize the interagency deliberations, the process should have a central platform, a secretariat similar to the VEP Executive Secretariat[36]. Every involved agency with equity[37] should declare a point of contact (POC) towards the secretariat. The VEP charter[38], which features such a structure seems like a good example. Difficulties lie in the composition of the participating agencies and in the choice of the lead agency/secretariat. The VEP charter gave that function to the National Security Agency. As military foreign intelligence agency which conducts cyber operations, it might be not the best fit because it might be biased towards the retention of vulnerabilities[39]. Another option would be a major cyber security/information assurance agency. In Germany, for example, this might be the Federal Office for Information Security. Similar to the NSA, this kind of agency might be biased, in its case towards the disclosure of vulnerabilities. Additionally, trust is critical in the IT-security world, and having a cyber security authority managing vulnerabilities of which some might not be disclosed immediately will inevitably tarnish its reputation with (international) contacts (e. g. CERTs), the private sector, civil society and other government agencies. Other actors rely on the security guidelines and standards offered by that authority. It would diminish the trust in those guidelines, and ultimately the implementation rate

36 White House, Vulnerabilities Equities Policy and Process for the United States Government (UNCLASSIFIED) https://www.whitehouse.gov/sites/whitehouse.gov/files/images/External%20-%20Unclassified%20VEP%20Charter%20FINAL.PDF

37 White House, Vulnerabilities Equities Policy and Process for the United States Government (UNCLASSIFIED) https://www.whitehouse.gov/sites/whitehouse.gov/files/images/External%20-%20Unclassified%20VEP%20Charter%20FINAL.PDF
Involved agencies with an equity can for example be: commerce, justice, foreign affairs, law enforcement, intelligence, military, consumer protection, cyber security and specialized agencies (such as the Federal Communications Commission in the United States or the Federal Institute for Drugs and Medical Devices in Germany).

38 White House, Vulnerabilities Equities Policy and Process for the United States Government (UNCLASSIFIED) https://www.whitehouse.gov/sites/whitehouse.gov/files/images/External%20-%20Unclassified%20VEP%20Charter%20FINAL.PDF

39 Ari Schwartz and Rob Knake, Government's Role in Vulnerability Disclosure - Creating a Permanent and Accountable Vulnerability Equities Process

and therefore the state of cyber security[40]. Another disadvantage that arises from installing the secretariat in an agency such as the Federal Office for Information Security or the National Security Agency is that it might have difficulties facilitating interagency deliberations; the weighing of equities is a representation of that. Due to the interagency nature of the process and its importance as well as the diverse equities, this paper favors setting up the secretariat in the proximity of a high government position which has a coordinating function in the broader realm of digitalization and security, for example the White House Cybersecurity Coordinator in the United States, or the Head of the Federal Chancellery in Germany.

2.4.2 Secretariat structure

In order to run a vulnerability assessment and management process, agencies send their POCs to partake in the process on a regular basis. In addition to government POCs and a small management team, adding external experts, including technical ones, on an ad-hoc basis (possibly from a predefined pool due to the security clearances needed) should be considered. This could be leveraged to lend particularly needed expertise during the deliberation[41], contributing views on benefits and harms of disclosure and retention[42] as well as lead to increased public acceptance of and trust in the process. The secretariat would also be tasked with reporting oversight and transparency efforts; managing and supporting the needed research on the topic of vulnerability assessment in the public and private sector; and taking stock of hardware, software and online services being used in government networks, critical infrastructures and other sensitive areas to better inform the assessment process.

2.4.3 Workflow

The procuring agency submits information about the vulnerability following the parameters laid out below to the secretariat immediately after acquisition. The entire secretariat, including POCs and possible additional experts,

40 A similar discussion was sparked by the controversial role the US National Institute of Standards and Technology has played promoting an encryption algorithm that the NSA could break, see Nicole Perlroth, Government Announces Steps to Restore Confidence on Encryption Standards

41 Apart from their technical expertise, this would be particularly helpful to assess parameters such as the likelihood of patch provision and adoption for the specific vulnerability.

42 Similar to the FISA court amicus.
US Congress, UNITING AND STRENGTHENING AMERICA BY FULFILLING RIGHTS AND ENSURING EFFECTIVE DISCIPLINE OVER MONITORING ACT OF 2015

convenes every month[43], or earlier if immediate need arises, and decides on the disclosure of vulnerabilities submitted, as well as those up for re-evaluation. The results of the deliberations can be "disclosure", "retention", or "mitigation and retention" (if mitigation is possible), with only the POCs being eligible to submit a vote. To account for a likely bias towards retention and to negate the power of numbers (e.g. the presence of more security and intelligence agencies with equity than other agencies), a POC championing disclosure (likely to be the representatives of commerce, foreign policy or national cyber security agency) has to create a robust minority by convincing a small number of additional POCs (~15%) to vote for disclosure. If there is no robust minority for disclosure and mitigation is an option, the vote

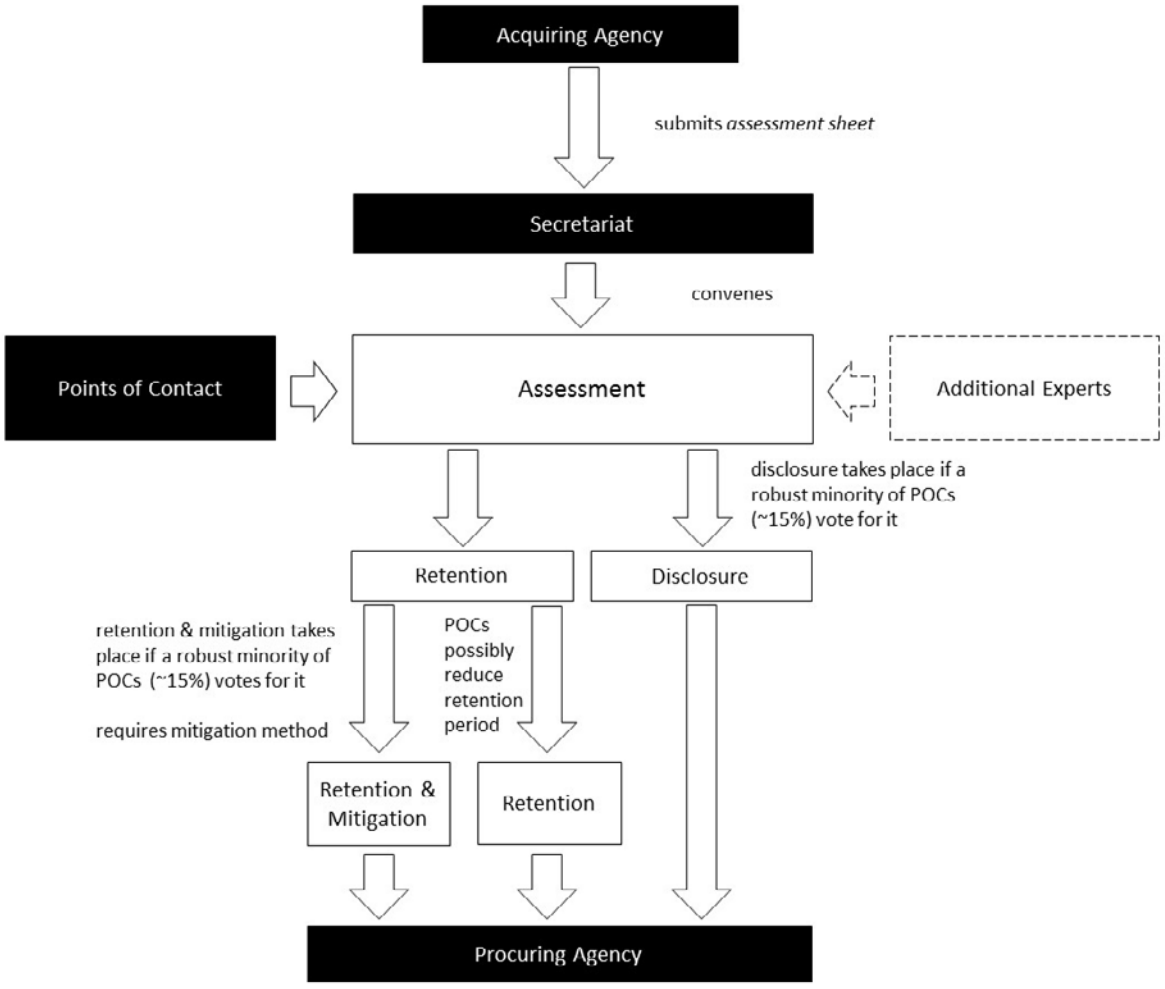

Figure 2. Visualization of the institutional setup and workflow

43 Germany's foreign intelligence oversight commission ("G-10 Kommission") for example convenes every month.
Bundesministerium der Justiz und für Verbraucherschutz, Gesetz zur Beschränkung des Brief-, Post- und Fernmeldegeheimnisses (Artikel 10-Gesetz - G 10) § 15 G 10-Kommission

will be repeated with the same mechanism to decide between retention and retention and mitigation. In this case the principle of a robust minority applies to retention and mitigation. If a vulnerability is being retained, the POCs have the option to shorten the re-evaluation period (for which the standard is 1 year) if the assessment has indicated that this might be necessary. The decision is submitted to the procuring agency, which is tasked to coordinate the further handling of the vulnerability in accordance with the decision of the secretariat.

2.5 Assessment

The basis for this assessment is that the internet ecosystem benefits from patching vulnerabilities, and government policy should be to disclose them unless there is a specific, justifiable reason for retaining and using them in law enforcement, intelligence or military programs. The actor (e.g. an intelligence agency) submitting a vulnerability should not submit the vulnerability itself but an assessment sheet; the assessment sheet includes responses to the parameters mentioned below[44] and is easier to understand than the vulnerability itself. The parameters help to assess the technical description of the vulnerability to determine whether it should be disclosed or not. Here are the factors to consider when determining whether there is an overriding government interest in keeping the vulnerability for its own use:

1. Level of dissemination
2. Likelihood of patch provision
3. Likelihood of patch adoption
4. Mitigation
5. Technological sovereignty
6. Usage of the affected product
7. Likelihood of detection of exploitation
8. Severity of the vulnerability
9. Collision rates
10. Operational value

44 Takes into consideration the VEP's "Defensive Equity Considerations"
White House, Fact Sheet: Vulnerabilities Equities Process https://www.whitehouse.gov/sites/whitehouse.gov/files/images/External%20-%20White%20House%20Fact%20Sheet%20on%20VEP%20-%20FINAL%2011152017.PDF

2.5.1 Level of dissemination

The more that actors (e.g. foreign intelligence agencies or cyber-crime groups) are aware of the vulnerability, the more likely it is that the vulnerability will be used against the state itself. One indicator therefore is to determine whether only the acquiring agency knows about the vulnerability or not. This is interrelated with the mode of how the vulnerability was acquired. If it was acquired through a third party or through forensics of past adversarial cyber operations, at least that actor necessarily knows about it. The higher the level of dissemination of a vulnerability, the more likely it should be disclosed to the maintainer, as the operational value decreases and the risk of another actor using it against the country's IT-systems increases.

2.5.2 Likelihood of patch provision

Another issue for the disclosure process to consider is to whom a vulnerability can be disclosed, and whether the product or software in question is being actively maintained. Orphaned codebases or hardware that no longer have a maintainer present challenges to timely patch development. In other cases, a product may have reached EOL and thus no longer be supported in the same way as a current offering. Products that have reached EOL, or where the maintainer no longer exists, are less likely to be patched; in this case, disclosure might not yield good results. When it comes to EOL however, it depends on the maintainer (and the severity of the vulnerability). For example, Microsoft provided patches even for products past their EOL when the EternalBlue/DoublePulsar vulnerabilities were exploited by WannaCry in 2017[45]. Stakeholders involved in the vulnerability assessment process should consider whether there is a responsible party in a position to, and willing to, identify, develop, and issue a patch. The less likely the patch provision, the more reasonable it would be for the government to retain a vulnerability. However, the government could still issue a public warning to increase risk awareness of companies and users, enabling them to mitigate the risk, for example by stopping use of the device or service.

2.5.3 Likelihood of patch adoption

Apart from the likelihood of patch provision, stakeholders involved in the vulnerability assessment process might also consider the complexity of patch

45 Steve Ragan, Microsoft patches Windows XP and Server 2003 due to WannaCrypt attacks

adoption[46]. This would be challenging to assess in great detail, but at a high level this distinction can help identify how slow a patch might propagate: for example, distinguishing between a patch that requires recall of a hardware product versus an over the air software update. It is useful to know how likely it is that an existing patch will be adopted and in what timeframe[47]. Similar to the likelihood of patch provision, a low likelihood of patch adoption does not automatically mean that a vulnerability should be retained. In limited very high-risk scenarios, a government could issue a public warning to increase likelihood of patch adoption or take other precautions, such as recommending against the use of software, hardware or the online service in question.

2.5.4 Mitigation

The possibility of mitigation of a vulnerability is another useful parameter. If mitigation is possible, then precautions can be taken for the government and other actors to secure them against an exploitation of that vulnerability without necessarily disclosing the vulnerability to the maintainer. In Germany, for example, these "other actors" could be a predefined set of institutions of particular interest to the state including defense contractors and critical infrastructures. These are called "Institutionen im besonderen staatlichen Interesse" (INSI)[48]. This does however not mean that just because mitigation is possible, a vulnerability should be retained, because all other actors are still left vulnerable. The option of mitigation solely decreases the risk of being successfully targeted with this vulnerability for the mitigation-adopting actors. In rare cases it might make sense to decide to retain a vulnerability for a short period of time while implementing mitigation techniques for certain actors. If mitigation is not possible, the only two options left are retention and disclosure.

46 Patch adoption is a complex issue, for example because there is more than just maintainer and end user. Original Equipment Manufacturers (OEM) for example might provide a patch but it has to go through intermediaries before it reaches the end user. Additionally, large IT-infrastructures are complex and heterogeneous. Oftentimes, patches have to be thoroughly tested before the can be rolled out for the entire infrastructure, to make sure that they do not cause unexpected problems.

47 If the patch provides useful information about a vulnerability to an adversary but the patch is not adopted and no mitigation method exists, the consequences might be more severe than if the vulnerability is retained.

48 Bundesamt für Sicherheit in der Informationstechnik, Allianz für Cybersicherheit Jahresbericht 2012/2013

2.5.5 Technological sovereignty

If the vulnerability that is being assessed affects products (hardware, software or online services) which are in the hands of a domestically maintained company, economic aspects should be factored in the vulnerability assessment. For the government not to undermine the security of domestically maintained hardware, software and online services, special additional emphasis should be applied to the disclosure of vulnerabilities in domestically-maintained products. The users of these products may be disproportionately domestic, including the government, and thus at risk for however long the vulnerability is retained. Second, the reputation of the affected brand may be negatively impacted by the government's decision to withhold information about a vulnerability, thereby delaying the ability to mitigate or patch the vulnerability. Governments should recognize that this effect on the trustworthiness of domestic brands can negatively impact the economic security of the nation, which is part of national security.

2.5.6 Usage of the affected product

A key aspect for assessment is to know where the affected hardware, software or online service is in use, for example in the country's own military hardware, in allies' military hardware, government networks, critical infrastructures, companies of particular interest to the state, safety industry, private sector at-large or consumer products. It is important to note that general purpose and consumer products are used by government and business personnel at home and abroad (especially by military and diplomatic staff), as well as for privileged communication (e.g. journalists and lawyers), which could therefore possibly change the threat level. This parameter needs to factor in how widely those systems are used as well.

2.5.7 Likelihood of detection of exploitation

The lower the likelihood that security and defense authorities are able to detect an ongoing exploitation of the vulnerability, the greater importance of mitigating and disclosing the vulnerability. If the vulnerability can be exploited by third parties undetected by the national security apparatus, it would be detrimental for national security to retain it. In that case, the decision regarding the handling of the vulnerability should favor disclosure.

This assumption must not always be symmetrical; a vulnerability for which exploitation might be difficult to detect by an adversarial actor might be

easy to detect by the own security agencies or be too difficult to exploit for adversaries, the Nobody But Us (NOBUS) concept[49].

2.5.8 Severity of the vulnerability

Several factors determine the severity of the vulnerability. One aspect is whether it can be exploited remotely, requiring physical proximity to the device (e. g. WLAN/Bluetooth), or if it requires actual direct contact with the device. A vulnerability that can be exploited remotely is in that respect more severe than one that requires one to be in the possession of it. Additionally, a vulnerability that can be exploited without user interaction is more severe than one that requires user interaction, for example by opening attachments or clicking links in emails. Another point is whether it needs only that one vulnerability to successfully exploit hardware, software or an online service, or whether a chain of vulnerabilities is required. Just because a vulnerability in itself does not lead to an exploitation or leads to a less critical exploitation as when chained with other vulnerabilities does not necessarily make it less problematic to retain it: it has to be put into perspective with the current vulnerability landscape. Additionally, the form of damage that can be done by exploiting this vulnerability is crucial: can it lead to disruption of the service or full backdoor access to all data on the system, including the ability to monitor traffic and manipulate sensors? This aspect needs to be considered, because a disruption of online email service represents a different threat than full backdoor access to said email service, allowing the attacker to read all the current and past exchanges. The Common Vulnerability Scoring System (CVSS)[50] can serve as an example for assessing the severity of a vulnerability.

Even though all parameters need to be considered and weighed together anyway, there is a strong link between the usage of an affected product and the severity of a vulnerability in that product. A vulnerability that allows remote shutdown of Windows machines could be a strong candidate for disclo-

49 "The premise of the NOBUS approach is simple: when there is tension between offense and defense, the United States aspires to secure communications against all forms of signals intelligence collection—except those forms of interception that are so complex, hard, or inaccessible that only the United States uses them", Ben Buchanan, Nobody But Us Michael Hayden in 2013: "If there's a vulnerability here that weakens encryption but you still need four acres of Cray computers in the basement in order to work it you kind of think 'NOBUS' and that's a vulnerability we are not ethically or legally compelled to try to patch — it's one that ethically and legally we could try to exploit in order to keep Americans safe from others", Andrea Peterson, Why everyone is left less secure when the NSA doesn't help fix security flaws cited in Tristan Caulfield, Christos Ioannidis and David Pym, The US Vulnerabilities Equities Process: An Economic Perspective

50 FIRST, Common Vulnerability Scoring System SIG

sure, because the software is widely used, even though the damage might not be so severe. A vulnerability in a pacemaker that has been sold hundreds of times could also be a strong indicator for disclosure if exploiting that vulnerability would allow to shut it down remotely. A vulnerability enabling full backdoor access to a product that is only used by foreign (adversarial) military on the other hand would most likely warrant retention no matter what harm could be done with it.

This parameter gains additional importance when looking at it from the perspective of what would happen if security mechanisms to contain a retained vulnerability fail and it was stolen and/or leaked by a third party, as happened with the disclosure of the US National Security Agency (NSA)'s Eternal Blue vulnerability by the purported group *The Shadow Brokers*[51].

2.5.9 Collision rates

When vulnerabilities are discovered by two or more parties, such as the Heartbleed bug in OpenSSL or more recently Spectre and Meltdown vulnerabilities, it is called a collision. Calculating the rate of collision in different types of vulnerabilities in different kinds of software can inform (but not outright determine) the likelihood that someone other than the government in question might in future, or may have already, discovered the same vulnerability. The higher the collision rates, the more likely that another actor finds out about the vulnerability and can subsequently use it, making vulnerability retention more risky. To the extent possible, governments should work to calculate or reference collision rates for similar vulnerabilities to those under consideration for disclosure. Initial research in this area is limited and suggests different conclusions[52]. Without further research, broad empirical data will not be available for specific vulnerabilities. Further independent research on collision rates should be incentivized by the government. Until then, collision rates should be used as a criterium contributing to the decision to disclose, but not a core component.

51 Check Point IPS Research Team, BROKERS IN THE SHADOWS: Analyzing vulnerabilities and attacks spawned by the leaked NSA hacking tools

52 Trey Herr, Bruce Schneier and Christopher Morris, Taking Stock: Estimating Vulnerability Rediscovery
Lillian Ablon and Andy Bogart, Zero Days Thousands of Nights
Katie Moussouris and Michael Siegel, The Wolves of Vuln Street:The 1st System Dynamics Model of the 0day Market presented at RSA Conference 2015

2.5.10 Operational value

Another crucial parameter is concerned with the operational use for security and intelligence agencies of the vulnerability. This parameter takes into account aspects such as the operational area where it would be used (intelligence collection, military cyber operations or law enforcement investigations), whether vulnerabilities which fall in a similar category and provide the same value are already retained or in use, as well as the operational effectiveness of this vulnerability. Identifying its effectiveness must include an assessment of whether there are other less intrusive ways to collect a target's information, how important this technique would be to an operation, and how critical that operation is to national security. Again, the internet ecosystem benefits from patching vulnerabilities, and government policy should be to disclose them unless there is a specific, justifiable reason for retaining and using them.

2.6 Management

2.6.1 Handling and retention

Vulnerabilities can only be safely handled by government actors if those actors have the capacity to secure them appropriately (see "Securing Vulnerabilities"). Any agency that cannot secure vulnerabilities must not be allowed to retain them. Any agency from which retained vulnerabilities were stolen must undergo a security audit before it can handle vulnerabilities again.

2.6.2 Re-evaluation

Due to possible changes in the parameters and therefore outcome of the assessment, retained vulnerabilities must be subjected to a regular re-evaluation as well as a re-evaluation whenever they are operationalized. Re-evaluation must take place after 1 year[53], unless the POCs specified a shorter time span during their assessment. As noted above, use of a vulnerability is an example of a triggering event that should be presumed to result in discovery by an adversary and result in re-evaluation.

53 White House, Vulnerabilities Equities Policy and Process for the United States Government (UNCLASSIFIED) https://www.whitehouse.gov/sites/whitehouse.gov/files/images/External%20-%20Unclassified%20VEP%20Charter%20FINAL.PDF

2.6.3 Short-term retention

All vulnerabilities can only be retained temporarily and must be disclosed at some point: it is just a question of when. Short-term retention is a special kind of retention designed to open a very short window for the government to retain a vulnerability, while at the same time putting mitigation mechanisms into effect. Short-term retention can only be implemented if the following conditions are met:

- the vulnerability is obtained through forensics
- a short-term mitigation method exists
- the mitigation method is being rolled-out as soon as possible
- the assessment concludes in favor for retention
- retention is bound to a specific current purpose (i.e. for a specific forensics/ attribution operation[54]).

Short-term retention ends when the specific purpose is fulfilled, e. g. the operation is concluded. A vulnerability that is retained short-term must be re-evaluated during every subsequent assessment.

2.6.4 Mitigation

Mitigation refers to "[...] detection and protection strategy used to safeguard networks, servers and applications by IT administrators in order to minimize the effect of malicious traffic and intrusion attempts while maintaining functionality for users"[55]. This is independent from the availability of a patch provided by the maintainer who has a vulnerability in their hardware, software or online service. There is not always a workaround that mitigates the effect of a vulnerability. Within this process, mitigation methods should be developed and propagated in cooperation and/or through the national cybersecurity authority.

Mitigation during delayed disclosure
Delayed disclosure might sometimes be necessary when there are ongoing operations that would otherwise be exposed by the disclosure. Therefore, one application of mitigation is during short-term retention of a vulnerability with a slightly prolonged window of exposure, which allows those operations

54 For example, in late 2017/ early 2018 German security agencies did not disclose any information about an ongoing cyber operation against the Federal Foreign Office in order to observe the attackers for their methods, techniques, motivation and attribution. taz, Cyberangriff auf Ministerien Hacker noch im Bundesnetz

55 radware, Attack Mitigation

(e.g. forensics/ attribution) to be concluded. If the decision is made to delay the disclosure in favor of short-term retention and operational value, starting mitigation efforts could be part of this decision. This increases the risk that the retained vulnerability will become public as part of the mitigation efforts, but simultaneously increases the level of protection for the own systems.

Long-term mitigation with delayed disclosure
Another application of mitigation could be a long-term mitigation and retention of a vulnerability. If the outcome of the assessment is in favor of retention and a mitigation mechanism exists, retention and mitigation can be triggered simultaneously without immediately being followed-up by disclosure. The vulnerability is temporarily retained but a mitigation method is propagated (and implemented) to a clearly defined limited number of actors deemed critical for national security (see "INSI"). That way it is possible to balance the operational value of the vulnerability with a lower risk of the vulnerability becoming public as compared to disclosure. However, mitigation always carries the risk of the vulnerability eventually becoming disclosed.

In addition, both short and long-term retention with mitigation must consider that the government's ability to identify such mitigations are far lower than the maintainer's. Workarounds such as closing off access to a network or instructing users not to use an affected feature can be short-term fixes but do nothing to repair the underlying code base.

2.6.5 Disclosure

Once a vulnerability is marked for disclosure, it is to be handled by the authority responsible for coordinating vulnerability disclosures. In most instances this will be the national cybersecurity authority. The disclosure process should be managed in a way that is consistent with CVD practices.

2.7 Safeguards

2.7.1 Securing vulnerabilities

Securing vulnerabilities is a vital cross-cutting theme that has to accompany the entire process, from the acquisition of a vulnerability until its disclosure. Failure to secure retained vulnerabilities can yield catastrophic results, as shown by the EternalBlue vulnerability, which was taken from the NSA

and used in the WannaCry and NotPetya malware[56]. Current approaches may not be enough - while it is impossible to say from information available in the public domain what percentage of vulnerabilities have been revealed from government stockpiles like EternalBlue, even a very small number provokes serious concern. The difficulty of perfectly protecting these vulnerabilities at all times highlights the dangers of maintaining large stockpiles. The goal is to prevent any unauthorized party from obtaining information about the vulnerability. One aspect which has not been successful on its own so far is to classify vulnerabilities and related documents and secure them according to the respective classification guidelines. If even powerful intelligence agencies such as the NSA cannot adequately protect retained vulnerabilities, governments have to develop new (out-of-the-box) ways of doing so by supporting relevant research and looking at lessons learned from other sectors.

2.7.2 Legislative oversight

Since retaining vulnerabilities in hardware, software and online services has potentially large ramifications for commerce, civil liberties, IT security and public safety, legislative oversight is paramount[57]. The secretariat should be required to submit a classified annual report to the respective legislative bodies (such as in the US, the House of Representatives Homeland Security Committee, the Senate Homeland Security and Governmental Affairs Committee, the House Permanent Select Committee on Intelligence, the Senate Select Committee on Intelligence[58] or in the German case the Defense Committee, the Interior Committee, the Digitization Committee or the Parliamen-

56 Lily Hay Newman, Why Governments Won't Let Go of Secret Software Bugs

57 Even though the Vulnerabilities Equities Process does not include legislative oversight at the time of writing, corresponding steps have been undertaken
US House of Representatives Permanent Select Committee on Intelligence, AMENDMENT IN THE NATURE OF A SUBSTITUTE TO H.R. 6237 OFFERED BY MR. NUNES OF CALIFORNIA
US Congress, A BILL To authorize appropriations for fiscal years 2018 and 2019 for intelligence and intelligence-related activities of the United States Government, the Community Management Account, and the Central Intelligence Agency Retirement and Disability System, and for other purpose.

58 Chase Gunter, House passes vulnerability disclosure oversight bill https://fcw.com/articles/2018/01/11/house-passes-vep-bill.aspx

tary Control Panel). The report needs to include:

- the number of retained vulnerabilities
- the average retainment period
- the operational value of concluded operations
- the effectiveness of mitigations
- the assessment sheets of all retained vulnerabilities that have been disclosed already.

2.7.3 Transparency

Transparency reporting should be done annually and in a manner that enables outside experts to assess whether on balance, the vulnerability assessment and management process is increasing the overall security of the internet ecosystem by prioritizing disclosure and permitting well-justified instances of vulnerability retention. Substantive descriptions of all currently-retained vulnerabilities are to be excluded from the transparency report. The report is to be compiled by the secretariat, must be published online, and should include:

- the number of vulnerabilities procured and directly disclosed
- the number of vulnerabilities procured, mitigated and retained
- the number of vulnerabilities procured and retained
- the number of vulnerabilities procured and operationalized
- the average retainment period of retained vulnerabilities
- the number of vulnerabilities which were retained and during that time successfully exploited by a third party.

3. Conclusion

With states competing for all kinds of offensive operations in cyberspace such as espionage, criminal investigations, warfare or hack backs, government vulnerability assessment and management becomes a core challenge. States must face the question of how to deal with it and prudently balance the respective equities such as civil liberties, commerce, public safety and IT security. The United States has undertaken several important steps in the right direction by developing, implementing and, to a certain degree, showcasing their vulnerabilities equities process. It is now incumbent upon other states[59] to follow suit and develop similar mechanisms or to adopt and improve on existing ones, and to be transparent about what they are doing in this area. However, the US process can also still benefit from further improvements such as a mandatory reporting to Congress and the public.

This paper built on the existing US mechanism for government vulnerability assessment and management, and benefited from valuable input from discourse and criticism that has been voiced regarding the process over the years, as well as opinions and assessments from the Transatlantic Cyber Forum's cybersecurity expert network. The suggested mechanism and its elements are meant to further enrich the debate about government vulnerability management and offer potential solutions to existing challenges.

59 Centre for European Policy Studies (CEPS), Software Vulnerability Disclosure in Europe: Technology, Policies and Legal Challenges https://www.ceps.eu/system/files/CEPS%20TFRonSVD%20with%20cover_0.pdf

About the Stiftung Neue Verantwortung

The Stiftung Neue Verantwortung (SNV) is an independent think tank that develops concrete ideas as to how German politics can shape technological change in society, the economy and the state. In order to guarantee the independence of its work, the organisation adopted a concept of mixed funding sources that include foundations, public funds and businesses.

Issues of digital infrastructure, the changing pattern of employment, IT security or internet surveillance now affect key areas of economic and social policy, domestic security or the protection of the fundamental rights of individuals. The experts of the SNV formulate analyses, develop policy proposals and organise conferences that address these issues and further subject areas.

About the Transatlantic Cyber Forum (TCF)

The Transatlantic Cyber Forum (TCF) has been established by the Berlin based think tank Stiftung Neue Verantwortung (SNV).

The Transatlantic Cyber Forum is a network of cyber security experts and practitioners from civil society, academia and private sector. It was made possible with the financial support from the Robert Bosch Stiftung and the William and Flora Hewlett Foundation.

About the Author

Sven Herpig is the project director of the Transatlantic Cyber Forum (TCF), bringing together American, German and other EU-experts to collaborate on cyber security policies.

Contact the Author

Dr. Sven Herpig
Project Director Transatlantic Cyber Forum
sherpig@stiftung-nv.de
Twitter: @z_edian
+49 (0)30 81 45 03 78 91

Imprint

Stiftung Neue Verantwortung e. V.

Beisheim Center
Berliner Freiheit 2
10785 Berlin

T: +49 (0) 30 81 45 03 78 80
F: +49 (0) 30 81 45 03 78 97

www.stiftung-nv.de
info@stiftung-nv.de

Design:
Make Studio
www.make-studio.net

Layout:
Johanna Famulok

Free Download:

www.stiftung-nv.de