MORE SUNLIGHT, FEWER SHADOWS

Guidelines For Establishing & Strengthening Government Vulnerability Disclosure Policies

CENTER FOR CYBERSECURITY POLICY AND LAW

TABLE OF CONTENTS

GVD CHECKLIST

Implementing a **Government Vulnerability Disclosure (GVD)** policy for safely handling high-value cybersecurity vulnerabilities requires answering four key questions.

ARE THE CORRECT PEOPLE IN THE ROOM?	Deliberation of vulnerability disclosure decisions through a GVD program requires the participation of all government entities engaged in nationally critical functions that could be harmed or otherwise impacted by the improper handling or widespread exploitation of cybersecurity vulnerabilities. The involvement of technical experts is also required as a prerequisite to cross-government deliberation. These experts play a crucial role in assessing the risks posed by specific vulnerabilities.
ARE DECISIONS BEING MADE IN THE NATIONAL INTEREST?	The purpose of a GVD process is to weigh and reconcile distinct components of the national interest, including those that are often overlooked in national security policymaking. The GVD process should elevate the role of technical experts and be maximally inclusive of policymakers working on all relevant issue areas. The mission of the designated coordinating institution should be aligned with the goals of facilitating inter-agency cooperation and ensuring attention is paid to both offensive and defensive equities.
IS THE PROCESS TRANSPARENT & ACCOUNTABLE?	A degree of independence from existing centers of gravity in national security-related policymaking, assured through both design and oversight, is essential for a GVD process to function effectively. The gold standard for oversight of a GVD process is through the legislature, but general publication of GVD details including decision-making criteria is also desirable. An effective implementation will successfully reconcile the desire for maximal transparency with legitimate national security interests.
IS THE PROCESS SET UP CORRECTLY FOR THE GIVEN COUNTRY?	Each country's distinct institutional arrangements (e.g., ministerial purviews) will inherently lead to some heterogeneity in the selection of GVD participants, as well as the characteristics of the most appropriate coordinating body. Just as the GVD process itself should be iterative, with regular review of retention decisions to ensure that their justifications remain valid, so too should be the design of the GVD program itself.

CYBER THREAT ALLIANCE

CENTER FOR CYBERSECURITY POLICY AND LAW

The mission of the **Cyber Threat Alliance** (CTA) is to enhance the security of the digital ecosystem. We believe that achieving that goal requires timely, consensual sharing of relevant information about potential cyber threats among many different organizations.

One important type of information concerns vulnerabilities; flaws in software and hardware that can be exploited to enable third-party access to computer networks and systems. The proliferation of digital technology and existence of vulnerabilities together create significant cybersecurity risk. As such, information about vulnerabilities forms a critical input into managing that risk.

Many different actors look for and find vulnerabilities, including private researchers and governments. However, when governments act irresponsibly in handling the zero-day vulnerabilities they find or purchase, cyber risk — both in an acute sense and over the longer run — is heightened.

On the other hand, governments have legitimate interest in pursuing national security and law enforcement goals through the use of vulnerabilities because they can use those holes to identify and catch malicious actors and help to keep people safe. Therefore, the policy question becomes how governments should manage and share information about the vulnerabilities they find or purchase. Given their respective missions, CTA and the cybersecurity industry as a whole have a strong interest in governments effectively managing vulnerabilities that they acquire. This paper provides an important analysis of what makes for effective programs focused on that goal.

We hope that more governments will adopt this type of framework, and that governments with frameworks already in place will use this guidance to make those programs stronger. If they do, we will all be better off.

Michael Daniel
President & CEO

The media often sees vulnerabilities as a significant cybersecurity risk threatening the systems that underpin our modern society, while governments and law enforcement can be tempted to look for previously undisclosed vulnerabilities as a means to bolster national security and pursue bad actors. It is clear that there will continue to be vulnerabilities in hardware and software and finding the right way to deal with them needs to be a priority. It is the goal of this paper to help outline an effective and balanced approach for governments to make decisions on vulnerabilities.

As part of its mission, the **Center for Cybersecurity Policy and Law** develops, advances, and promotes cybersecurity best practices and educational opportunities among cybersecurity professionals. While we have tackled numerous issues during our existence, policy to properly deal with hardware and software vulnerabilities has remained a consistent theme. Our previous work includes publishing two white papers on vulnerability disclosure, **Improving Hardware Component Vulnerability Disclosure** and **Policy Priorities for Coordinated Vulnerability Disclosure and Handling**. The Cybersecurity Coalition, an initiative under the Center, has repeatedly engaged with government entities to share our expertise in crafting vulnerability policy. The Center also provided substantial feedback to drafts of the discussion paper that I wrote with my former colleague Rob Knake entitled **Government's Role in Vulnerability Disclosure: Creating a Permanent and Accountable Vulnerability Equities Process**, which was published by the Belfer Center at Harvard University.

I would like to thank the Cyber Threat Alliance for their lead on this paper, which helps to provide greater guidance in this policy area. We hope that governments everywhere will seriously consider the framework laid out here as a pathway to improving the cybersecurity ecosystem.

Ari Schwartz
Coordinator

INTRODUCTION

Software and hardware systems are rarely designed and manufactured without some mistakes being made. Some of these errors in computer code or hardware, which result from insecure design, intentional malicious actions, or human error, can introduce cybersecurity vulnerabilities into the technology. These can be exploited to enable access to the relevant software, system, or network, and used to cause damage, such as the installation and spread of ransomware. At the root of many major cyber incidents are one or more such vulnerabilities.

Cybersecurity vulnerabilities exist in all widely used software and hardware, including in the systems that we depend on for national critical infrastructure (CI), including healthcare, energy generation and distribution, telecommunications, food supply, and industrial control systems. Security researchers inside and outside the government use various technical methods to identify these vulnerabilities, which can be purchased from intermediaries for a price corresponding to both demand from those looking to use them and their availability (i.e. supply). If shared with the relevant vendors or operators, knowledge of these vulnerabilities can be used to create patches, remediations, or mitigations for the affected products — i.e. fixes. Conversely, in the hands of malicious actors, these vulnerabilities may be used to create exploits that enable clandestine access to the affected systems, causing myriad harms, for example violations of citizens' privacy, theft of intellectual property or other corporate assets, or disruption of CI. Unfixed vulnerabilities represent a serious and sustained security risk that cannot be easily gauged without detailed technical analysis encompassing context around use and interdependencies. Until a vulnerability is identified and fixed, it may be exploited by malicious actors to manipulate, deny, disrupt, degrade, or destroy vulnerable products or services; and allow for the compromise of connected networks and systems.

Governments may also make use of these vulnerabilities for military, intelligence, or law enforcement operations, leveraging knowledge of security flaws acquired through internal research or acquisition via legal or illegal marketplaces. The risks posed by government mishandling or misuse of these vulnerabilities are significant and, as government hacking grows ever more common, each new, unfixed vulnerability represents a potential risk to a variety of national interests. These include CI protection, citizens' privacy and civil liberties, and trust in governments within and across countries. At the same time, governments often use these vulnerabilities to achieve important law enforcement, public safety, and national security goals. Thus, amidst governments' legitimate, competing policy goals, a policy tension exists in how governments handle vulnerabilities.

In theory, two 'corner solutions' exist to address this policy problem, wherein governments could adopt either a policy of universal disclosure, never retaining or using cybersecurity vulnerabilities, or a policy requiring near-universal, indefinite retention. However, governments' legitimate, competing policy goals render neither of these approaches practical or desirable in a world where digital technology is ubiquitous, cyber risk is increasing, and protection of democratic values is declining.[1,2]

While exploitation of certain vulnerabilities by governments may lead to second-order impacts that affect systems beyond the target network, disclosure of vulnerabilities that could not feasibly be fixed may also be counterproductive to the goal of minimizing cybersecurity risk.

It is crucial, therefore, that when government agencies come into possession of such vulnerabilities, that the government as a whole is able to weigh all relevant interests in determining whether to disclose those vulnerabilities to vendors immediately or to delay disclosure and use the retained vulnerabilities to advance specific operational goals. Doing so requires that governments have processes in place encompassing:

- **Formalization of the rules by which the decision-making process is to be operationalized**
- **Identification of national interests that could be impacted by retention or disclosure**
- **Consideration of how different stakeholders should be engaged or considered in the process**
- **Creation of policies to advance transparency, accountability, and public trust in the process**
- **Specification of how divergent interests or viewpoints are to be handled and weighed**

BACKGROUND

EXISTING GVD POLICIES

Throughout this paper, we use the term Government Vulnerability Disclosure (GVD) in reference to the kinds of internal policymaking structures that governments need to implement in order to adequately assess and weigh the potential costs and benefits of immediately disclosing knowledge of previously unidentified cybersecurity vulnerabilities, versus retaining that knowledge based upon carefully considered and time-limited justifications.

Although the precise institutional form and substructures of GVD programs will naturally vary across countries, based in part on existing institutional differences, common features and underlying principles should cut across jurisdictions. This includes ensuring that GVD programs can maintain a degree of independence from existing operational structures and have sufficient oversight to ensure that they function as intended. The recent publicization of GVD programs in the US ("Vulnerabilities Equities Process" – VEP) and UK ("Equities Process" – EP) has raised the prominence of this issue within academic and policy circles. However, to date, few other countries have made much progress towards adopting and publicizing similar policies.

The US VEP was first conceptualized under President George W. Bush in National Security Policy Directive 54 in January 2008 and was further developed by President Barack Obama in October 2012 through Presidential Policy Directive 20 and other National Security Council policy decisions. An unclassified summary of a VEP Charter was publicly released by the White House was released in November 2017, which is summarized on the following page.[3]

At the end of 2018, the UK government released details of its EP. This process is overseen by the UK's National Cybersecurity Center (NCSC) housed within the Government Communications Headquarters (GCHQ). As with the US VEP, it also relies on significant involvement from across the intelligence community, as well as input from other UK government agencies.[4] The Canadian Communication Security Establishment (CSE) also oversees a more limited "Equities Management Framework," which provides for joint participation of both CSE's defensive and intelligence arms; the Canadian Centre for Cyber Security and Signals Intelligence branches, respectively.[5]

GVD policies should not be confused with so-called Coordinated Vulnerability Disclosure (CVD) programs or Vulnerability Disclosure Policies (VDPs) that have been adopted by a range of organizations, including governments, to facilitate the reporting of vulnerabilities in those organizations' systems and networks by security research community. For example, while the US Cybersecurity and Infrastructure Security Agency's "Binding Operational Directive 20-01" requires all executive branch agencies to establish CVD / VDP mechanisms to enable the intake and processing of vulnerability information reported security researchers, these vulnerabilities are not processed through the US GVD policy structure (known as the VEP).[6]

While CVD or VDPs deliver widespread but incremental benefits to the security of covered systems and networks, GVD is valuable for managing high impact, low probability (long-tail) risks in cyberspace through its influence over the capabilities and decision-making of the most sophisticated and well-resourced offensive actors in cyberspace: governments and their proxies.

THE NEED FOR WIDER GVD ADOPTION

Other countries have so far resisted establishing or publicizing similar programs. For example, after the US publicized its VEP, promising discussions around Germany establishing a GVD program began taking place across government, academia, and civil society.[7] The Transatlantic Cyber Forum (TCF), among others, made detailed recommendations as to what factors needed to be considered and how the government should approach developing and implementing a GVD program.[8] However, this momentum has since seemingly been lost and the German government remains without a clear, publicized GVD process.

Similarly, discussions around the use of cyber vulnerabilities and consideration of GVD formalization have been ongoing in the Netherlands from as early as 2014. However, the Dutch government has consistently asserted its

US VULNERABILITIES EQUITIES PROCESS WORKFLOW

A US government agency that acquires a vulnerability, which it determines to be newly discovered and not publicly known, is required to provide timely notice of that vulnerability's existence to the VEP Executive Secretariat. This notice is accompanied by a technical assessment and preliminary recommendation on whether to retain or disclose the vulnerability.

Within one business day of receiving this notice, the VEP Executive Secretariat is required to notify the designated points of contact for the various agencies represented on the VEP Equities Review Board (ERB).

Any agency with an interest in the handling of that vulnerability is required to inform the VEP Executive Secretariat whether or not it opposes the initial recommendation to retain or disclose the vulnerability within 5 business days.

Any non-concurring agencies then hold discussions with subject matter experts from the agency that submitted the vulnerability, as well as the VEP Executive Secretariat, in order to try and identify a consensus approach.

If there are no objections to the preliminary recommendation or a consensus approach can be reached, the VEP ERB is then required to ratify the decision. If disagreement persists, the decision-making process can be elevated to the National Security Council level and no disclosure will take place until the disagreement can be resolved.

If the vulnerability is to be disclosed, the submitting agency is responsible for disseminating information about the vulnerability to the relevant vendor(s) within seven business days. The submitting agency is also expected to ensure that the vendor addresses the vulnerability expeditiously and to inform the VEP Executive Secretariat if the vendor is not sufficiently responsive. The US government may, at that point, undertake other steps to ensure adequate mitigation.

If a vulnerability is retained, its status is reevaluated, at a minimum, on an annual basis by the VEP ERB.

Equities Review Board (ERB)

The role of the ERB is to deliberate and determine how a vulnerability should be handled. The following agencies are represented on the ERB, although others may be included if appropriate:

Office of Management & Budget
Office of the Director of National Intelligence
Department of the Treasury
Department of State
Department of Justice
Department of Homeland Security
Department of Energy
Department of Defense
Department of Commerce
Central Intelligence Agency

Executive Secretariat

This body is responsible for facilitating the logistics of the review process and is typically staffed by the National Security Agency under the supervision of the Secretary of Defense.

confidence in ad hoc checks and constraints on its handling of vulnerabilities, including in a December 2019 response to proposed GVD legislation from Kees Verhoeven MP.[9] In this response, the Dutch government objected to the adoption of a more uniform approach; however, the rationales for these objections were not entirely clear.

Although run-of-the-mill cybercrime has also expanded dramatically over the last decade, in terms of both its scope and reach, we cannot afford to ignore the risks associated with nation-state actors' use of undiscovered cyber vulnerabilities for the advancement of their military, intelligence, or other security objectives. The number of governments with offensive cyber capabilities has been growing as the relevance of cyberspace as a domain of geopolitical conflict has become increasingly apparent. However, the legitimate interests of governments in conducting cyber operations using zero-day-based exploits, as well as exploits based on known vulnerabilities, must be balanced with the risks that those operations pose to the domestic and global public, as well as to the government's own broader interests. At present, however, governments' use of vulnerabilities for their offensive purposes too often goes unchecked. A continued lack of GVD policies worldwide, especially in an increasingly belligerent geopolitical environment, will have negative (and entirely predictable) consequences, with an increased risk of real-world harm arising from governments' mishandling or misuse of vulnerabilities.

Zero-day vulnerabilities are particularly problematic in this context. These vulnerabilities are unknown to those entities responsible for designing or maintaining the security of an affected network or system. As such, attacks based on zero-day vulnerabilities cannot be preempted and prevented. Given the current preeminence of European countries in global digital policymaking, we believe it is worth acknowledging the available public evidence of offensive capability cultivation by their governments leveraging these zero-day vulnerabilities. This information is presented to motivate more widespread adoption of GVD policies within Europe and, ultimately, broader adoption by other nations with robust institutions and a commitment to rule of law that either conduct or aspire to conduct offensive operations in cyberspace:

- The German Federal Office for Information Security (BSI) maintained a contract with VUPEN, a now-defunct but once well-established French company selling zero-day vulnerabilities and exploits, through September 2014.[10]
- After VUPEN shuttered in 2015, its founders went on to found Zerodium, which purchases and resells exploitable vulnerabilities, including to governments.[11] Original, exploitable vulnerabilities for mobile OS such as Android can fetch upwards of 2 million USD on this marketplace, while vulnerabilities in desktop OS's or software are currently sold for between 10,000 USD and 1 million USD.[12]
- Leaked documentation from 2015 listed a variety a globe-spanning range of government agencies as clients of the Italian cyber exploitation company, Hacking Team.[13] This list included agencies from several European states, namely Italy, Hungary, Luxembourg, Czechia, Spain, Cyprus, and Poland. Hacking Team specialized in custom, easy-to-use offensive technologies that could be used for surveillance activities, such as bypassing encryption, accessing message records, and remotely recording activity through webcams and microphones.[14]
- There has been outcry among civil society groups and human rights advocates over recent years amid concerns that the French government knowingly permitted the sale of systems for domestic cyber exploitation and surveillance produced by a French company, Amesys, to the Libyan and Egyptian governments.[15]

Zero-day-based exploits are different from many other tools used for cyber operations in that while the exploits are reusable, their value degrades rapidly after first use. This situation occurs because zero-days are known only to a limited number of parties, crucially not including the vendor or operator of the vulnerable technology, but once used have a tendency to become broadly known very quickly, enabling defenders to patch the vulnerability. Moreover, zero-days are assumed to have an unknown but inherently limited shelf-life due to the potential for rediscovery, either by the vendor / operator or other third parties. This dynamic has precipitated two similarly undesirable approaches among government actors who come to

acquire these vulnerabilities:

1. A "use-it-or-lose-it" mentality, wherein an attack is launched as soon as a zero-day can be exploited and operationalized in order to avoid its potential rediscovery and, consequently, loss.
2. Zero-day stockpiling for accomplishing future operational goals.

Critics of GVD may point out that because many organizations fail to patch known vulnerabilities for extended periods, governments could use already-known vulnerabilities or other techniques to gain access to target data or systems for their operational needs. However, because zero-days can enable intrusions against otherwise well-defended systems and networks, governments want (and need) to have zero-day vulnerabilities available. Thus, GVD policy adoption is vital for the protection of high-value targets, such as CI or military infrastructure.

In high-level terms, we agree with the assertion from Sven Herpig and the Transatlantic Cyber Forum that "government policy should be to disclose [zero-day vulnerabilities] unless there is a specific, justifiable reason for retaining and using them in law enforcement, intelligence or military programs."[7] The events of 2020 only highlighted our dependence on a broadly-defined understanding of CI, and we cannot afford to continue neglecting consideration of the tradeoffs involved in governments' acquisition, handling, and use of vulnerabilities.

In this context, GVD programs are essential. If governments are to continue leveraging software vulnerabilities to advance certain national security and law enforcement interests, robust GVD processes are a crucial means of mitigating associated risks of harm. The remainder of this paper focuses on a set of broadly applicable best practices in GVD program design encompassing three main aspects:

- **Stakeholders:** Who needs to be at the decision-making table, and why?
- **Decision-Making Structure:** How can GVD programs be designed to ensure transparent, accountable, and socially beneficial decision-making?
- **Decision-Making Criteria:** What risks do disclosure and retention, respectively, tend to pose and how can these be weighed against the operational value of vulnerabilities?

DESIGNING A GVD PROGRAM: STAKEHOLDERS

Government Ministries & Institutions

At the core of a successful GVD process are participation, collaboration, and compromise among major government agencies and other offices engaged in nationally critical functions that could be negatively impacted by the improper handling or misuse of cybersecurity vulnerabilities. Even if exploitation of a given vulnerability by the military, intelligence services, or law enforcement is seen as an operational necessity, risks associated with its use may still be significant enough to outweigh potential benefits. As such, various risks to government and society writ large must be afforded proper consideration and reconciled against immediate and tangible operational objectives that could be furthered through retention and use of acquired vulnerabilities.

Mutual respect across agencies and a widely shared understanding of the importance of weighing relevant interests are also of critical importance to establishing and maintaining a well-functioning GVD program. At present, in countries without a GVD program, decision-making around vulnerability disclosure versus retention is typically restricted to law enforcement agencies, intelligence services, the military chain of command and, ultimately, the head of government. Prior and subsequent to GVD adoption, a whole-of-government effort is needed to improve understanding of why the involvement of agencies with mandates covering the full range of equities explored in this paper is needed to protect the national interest. While still imperfect in this regard, the US VEP correctly includes on its Equities Review Board (ERB) the departments charged with overseeing protection of civil rights and liberties, foreign affairs, trade, and other foreign and domestic economic policy.[4]

At a minimum, the agency-level component of GVD deliberation should include permanent high-level representation from ministries or agencies tasked with overseeing:

- **Military operations**
- **Local and national law enforcement**
- **Intelligence and espionage**
- **International and domestic economic policy**
- **Civil rights and liberties protections (including data protection)**
- **National critical infrastructure**
- **Executive (e.g., Prime Ministerial) functions**

In addition, other government entities should have the opportunity to contribute to the process on an ad hoc basis where their function is relevant to decision-making for specific vulnerabilities. Clear and secure communication channels should be pre-established with this broader range of potential participants. Criteria for determining when these irregular participants should be engaged will naturally be specific to national contexts, but should still be standardized and applied consistently by each national government.

The coordinating body for GVD participation and communication should be housed in close proximity to the executive office (of the Prime Minister or President). This ensures that effective coordination and communication across government can be readily achieved. Adequate institutional checks are required to prevent abuse of the process (see **Oversight**) given this necessity of proximity to power.

Technical Experts

Assessment of the risks posed by vulnerabilities by technical experts from the relevant identifying agency is a critical prerequisite to the interagency phase of the decision-making process. Quantifying the risk posed by non-disclosure of vulnerabilities involves exploration, to the extent that such analysis is feasible, of the following factors:

- **The types and versions of software, hardware, or firmware containing the vulnerability**
- **The levels of access to affected systems that could be achieved via its exploitation**
- **How widespread the use of the affected products or services is (i) within government, (ii) in CI systems, (iii) by the general public, or (iv) within the systems of other countries**
- **Users' requirements and expectations regarding the security of the affected product(s)**
- **Whether a new or existing mitigation (such as configuration changes) could reduce the risk posed by the vulnerability**
- **How easily and quickly a fix could be developed for affected systems**
- **How widely and rapidly adoption of a fix, if released, would be by the users of affected products**
- **How likely it is that this vulnerability has already been or could be discovered by another (e.g., adversarial) actor, and whether they would use it**

In many instances, precise, quantitative evaluations of these technical factors may not be possible. However, even imperfect or qualitative assessments (e.g., use of a particular product is widespread / moderately widespread / not very widespread) will allow policymakers to more effectively determine the most appropriate course of action for a given vulnerability. Having the best information available is better than having no information at all.

In existing models, this assessment is generally conducted by technical experts within the agency that identified the vulnerability, regardless of whether it was identified in-house or through third-party acquisition. This agency then submits details of the vulnerability to an interagency committee alongside an initial recommendation regarding disclosure or retention. This is followed by further review by technical ("subject matter") experts from the other agencies included in the VEP. This

approach encourages the identifying agency to engage in thorough due diligence prior to making its recommendation so as to avoid the risk of delaying the decision-making process while extensive technical reassessment by other GVD participants are undertaken or policy-level disagreements are resolved.

Alternatively, and particularly in countries where technical cyber capability is more concentrated within a limited number of agencies, technical assessment and recommendation may be a function for an institutionalized technical assessment group or committee. In this arrangement, which mirrors that of the UK VEP process, it is especially crucial that technical representatives be encouraged to make holistic assessments of risk. This goal could be furthered through the hiring or designation of outside technical experts. The process of decision-making at the policymaking level may be used to initiate further technical review.

Civil Society & The Public

Maximizing what government can share with the general public regarding both its GVD process and the associated throughput is essential to preserving and maintaining public trust. Not only is trust an essential ingredient for continued institutional legitimacy at the national level, but the transparency conveyed by GVD adoption and enactment can also serve as a means of building mutual confidence between nations. Such confidence-building measures for cyber conflict are hard to construct because software, unlike nuclear warheads, is intangible, easily replicable, and highly portable. As such, GVD adoption may be one of few policy tools available to advance diplomatic efforts to establish rules and norms around governments' conduct in cyberspace.

It should be underscored that the equities involved in government handling of software vulnerabilities by definition extend beyond government agencies themselves, since exploitation of software vulnerabilities can elevate material cybersecurity risk to a wide range individuals and institutions beyond government. As such, GVD review processes should ensure meaningful transparency and accountability to public interest stakeholders. Although involving the general public directly in the routine functioning of a GVD program is largely infeasible, annual reports and legislative hearings on the process would serve as effective vehicles for focusing relevant agencies on a principal goal of GVD: preserving public trust in the cyber policy decisions made by the government. This holds true regardless of the specific institutional arrangements adopted and subsequent vulnerability handling decisions undertaken. As former White House Cybersecurity Coordinator, Michael Daniel, noted in his initial White House Blog post that opened up discussion of government vulnerability handling in the US, "Too little transparency and citizens can lose faith in their government and institutions, while exposing too much can make it impossible to collect the intelligence we need to protect the nation."[16]

Vendors & Manufacturers

Although not directly involved in the decision-making process, strong relationships at the technical level between invested agencies and manufacturers and vendors of digital products and services are desirable for enabling effective GVD deliberation. Given that these private-sector entities will have more direct knowledge of their own offerings and procedures for developing and distributing fixes, accurately gauging the risk and mitigatability of a vulnerability may require input from those companies. Moreover, vendors have a clear and immediate equity in the security of their products, hence the importance of ensuring that their interests can be appropriately represented and accounted for in the process (e.g., by representatives of agencies responsible for promoting economic growth or digital safety and security).

In contrast, knowing that governments have their own procedures in place to appropriately handle new and potentially high-impact vulnerabilities may actually incentivize vendors to shore up their own processes for intake and patching of disclosed vulnerabilities and to induce the development of more robust disclosure pathways between government and the private sector.

DESIGNING A GVD PROGRAM: DECISION-MAKING STRUCTURE

Oversight

A degree of independence from existing centers of gravity in national security-related policymaking, assured through effective oversight mechanisms, is essential for a GVD process to function effectively. What this looks like in practice will vary from country to country as a function of institutional structure and political feasibility. Establishing effective oversight and requiring regular communication of the processes and outcomes of the GVD program represents the best possible assurance against the risk of policy failure.

The gold standard for oversight of a GVD process is through the legislature, with additional accountability coming through ombudspersons or inspectors general. In some countries, legislators have sought to establish such direct oversight of government vulnerability decision-making. For example, in 2017, the US Congress briefly sought to pass the PATCH Act, which would have established GVD-related oversight procedures and obligations in law.[17] Similarly, in the Netherlands, the parliamentarian, Kees Verhoeven of the liberal D66 party has been attempting to advance legislation that would assert parliamentary control over the government's handling and use of zero-day vulnerabilities.[8]

Information disclosed through intra-governmental oversight channels may be treated as classified to enable more expansive reporting, including around ongoing operations. The Transatlantic Cyber Forum report recommends scrutiny of GVD procedural outputs along lines that overlap with the following suggested criteria:[7]

- **The number of retained vulnerabilities**
- **The average retention period**
- **The operational value created through retention and use of those vulnerabilities (for concluded operations)**
- **The nature / effectiveness of mitigations**
- **The technical assessments of all retained vulnerabilities, even if those flaws were ultimately disclosed to the vendor / manufacturer**
- **For undisclosed vulnerabilities found to have been subsequently exploited by another actor, the average length of time between the government's first use and discovery of other actor's use of those vulnerabilities**

The aforementioned entities may also benefit from an understanding of the actions undertaken by the private sector following disclosure of vulnerabilities by the government. This is necessary to ensure that the GVD processes, once in place, are actually helping to advance the goal of a stronger cybersecurity ecosystem. Specifically, such oversight would entail understanding the efficacy and speed of patch rollout, extent of patch adoption, and impact on relevant companies' product development and security assurance processes.

Although intra-executive oversight — and even parliamentary scrutiny — can be an imperfect means of ensuring adherence to laws and best practices, with sufficient attention to institutional design in the GVD process, such an approach would still represent a significant improvement on the status quo for most democratic countries. In particular, this could be achieved through intentional levelling of competing equities at the technical and political levels of the process, as well as incorporation of an institutional bias towards disclosure (and mitigation). Naturally, there will be instances where the GVD is actioned for which absolute transparency (e.g., around intelligence sources and methods) is not viable. However, the successful management of such concerns in the context of oversight and transparency has been a long-standing practice across governments in a range of contexts, suggesting that GVD oversight mechanisms could similarly be crafted in such a way as to be consistent with established, legitimate, and appropriately deployed government classification systems.

In addition, public-facing transparency regarding the acquisition and use of vulnerabilities is also important, including for countries with weaker

democratic institutions or less robust rule of law. The value of such transparency goes beyond the largely intangible benefits to public trust and global confidence-building referenced in the **Stakeholders** section. Publishing unclassified, redacted versions of oversight-related materials — specifically excluding "[s]ubstantive descriptions of all currently-retained vulnerabilities," as suggested in the TCF report — would also help to support research into the comparative efficacy of different approaches to GVD policy and vulnerability handling more widely.

Reconciling Competing Interests

Assessment of prospective impacts from retention and use is highly context-specific with respect to the prevalence of affected software or hardware across government and non-government systems. This is why GVD frameworks are important in practical terms, and why those governments that have adopted GVD policies have done so with a procedural preference towards disclosure when the risks arising from use are high.

Adopting an institutional framework that elevates the role of technical experts and is maximally inclusive of policymakers working on diverse national interests is at the heart of an effective GVD program. In the US, this has meant a system where the identifying agency's recommendation is scrutinized by agency-level technical experts, with subsequent deliberation by a board of high-level agency representatives. This process, as well as internal reporting and communication of vulnerability information, is facilitated by an Executive Secretariat, which in turn receives its own oversight through the National Security Council process and, ultimately, the president.

It is not enough to have all relevant stakeholders nominally included in the process; formal points of contact (POCs) for each involved agency at the technical and policymaking level need to be identified and regular channels of communication for these POCs established through a central coordinating body. Technical representatives should meet regularly (at a minimum, monthly) to facilitate continued assessment and exchange of information, which will enable the development of trust and build "muscle memory" for ongoing coordination. Parallel meetings should occur at the policymaking level even when deliberation of vulnerabilities is not required, for example to facilitate ongoing policy development. Regular business should also be complemented by the possibility for ad hoc deliberations around vulnerability disclosure conducted on short notice. Such exceptional meetings would allow vulnerabilities with time-limited operational value or that otherwise merit such expeditious assessment to be properly processed under a consistent GVD approach.

Depending on the institutional landscape of the implementing country, the designation of a coordinating agency, ministry, or executive office may pose a number of challenges; however, the optimum designation would house a GVD program "in the proximity of a high government position [with] a coordinating function in the broader realm of digitalization and security."[7] The mission of that coordinating institution should be aligned with the goal of facilitating inter-agency cooperation and ensure an equivalent interest in both offensive and defensive equities. This is critical for fostering public, industry, and international trust in the GVD. Specifically, given that an undue preference for retention is more likely to come from within agencies that have an offensive or espionage function, it is desirable where possible to avoid housing the coordinating institution under these umbrellas. However, it is also important to ensure that technical cybersecurity expertise is available to or within the designated coordinating body. Ergo, it is generally desirable to place that coordinating function directly within the remit of the head of government to ensure that the coordinator is motivated by the broadest possible conception of the national interest.

Voting

There is significant disagreement among GVD practitioners and experts regarding the desirability of a voting threshold for retention or disclosure versus a consensus-oriented approach. This debate is further complicated by the idea that voting systems could, by design or by accident, consistently bias decision making towards a particular outcome in ways that impede the advancement of a broadly defined national interest and thereby prevent the GVD process overall from functioning as intended.

One line of thinking suggests that a "supermajority"

voting threshold could be used to prevent retention of a vulnerability given a sufficient minority of participating entities preferring disclosure. The TCF report, for example, recommends that only 15% support for disclosure among all designated points should be required. However, while a formal voting framework of this kind may have the conceptual advantage of bolstering external trust in the GVD process and ensuring that divergent interests within government are considered, both practical and political impediments lead us to reject such an approach.

Few governments have regular, institutionalized voting process within their executive branches. As a result, using a voting process for this topic would be anomalous compared to how other policy issues are typically resolved. Further, the effectiveness and reliability of threshold voting models would be exceptionally inconsistent given institutional heterogeneity across countries and over time. Variation in the number and nature of voting participants would in turn impact the overall balance of voting entities leaning more or less towards prioritizing offensive or non-offensive considerations. Such concerns hold true even if similar sets of equities are represented in the process overall. Certainly, a one-size-fits-all approach across countries' highly diversified institutional arrangements is not advisable, and the outcomes of tailored approaches to GVD voting should be based on experience with an existing process, rather than assumed ahead of initial implementation.

The opposing view focuses on default positions and consensus. Under this approach, governments would establish a default policy for disclosure, requiring those agencies that wish to retain the vulnerability to make an affirmative case. The decision-making process should use whatever high-level decision-making model the government typically employs. This approach ensures that retention is only pursued when its benefits obviously outweigh its costs. For intractable divisions among GVD participants about which course of action to pursue, escalation to a political office where the public interest is most acutely experienced (e.g., the office of a president or prime minister) should ensure that expectations of transparency and the norm of disclosure-by-default are upheld despite objections from pro-retention entities within government. Most importantly, in a worst-case scenario where a retention or disclosure decision leads to significant, negative consequences, such a focused process also lends itself to ex-post accountability in a much clearer manner than would a system of threshold-based voting, in which responsibility could be divided so widely that none is every really felt at all.

Iteration & Improvement

In circumstances other than a decision to disclose and subsequently patch a vulnerability, that retained (or retained and mitigated) vulnerability should be reassessed under a regularized schedule appropriate to the vulnerability's sensitivity. This is an essential characteristic of an effective GVD program — since the value of a retained vulnerability may change with evolving context or further investigation of risks and potential benefits. This ongoing review should be facilitated by continued technical reassessment and interagency dialogue. In the specific instance where a retained vulnerability has been operationalized and used, the GVD infrastructure should be activated to support risk mitigation during the operation and patching (or at a minimum reassessment) upon its conclusion. As each country's GVD process matures and such structures become more well-established worldwide, the appropriate cadence for such reviews will become clearer. While, for example, a monthly review process may be desirable, it may also be logistically infeasible, at least initially. Frequent but informal reviews that focus more narrowly on potential changes with respect to key decision-making criteria, in combination with less frequent, comprehensive reviews (e.g., every six months), would offer a realistic approach to procedural accountability following initial implementation of a GVD program.

DESIGNING A GVD PROGRAM: DECISION-MAKING CRITERIA

When deciding whether to disclose or retain the acquired vulnerability, the costs and benefits of each action should be analyzed from diverse perspectives that go beyond the interests of a single government stakeholder or group of stakeholders. While differences in context and vulnerabilities' technical profiles necessitate a case-by-case approach to decision-making criteria, there are a number of perspectives that should serve as baseline criteria in most instances. Those perspectives can be grouped into three broad categories:

1. **Risks around retention**
2. **Risks around disclosure**
3. **Utility & necessity of retention**

The criteria listed below, which are based in part on those publicized by the US, are not exhaustive. It should also be noted that both technical interpretation and holistic consideration of context, including possible interactions across agency interests, should shape the detailed implementation of decision-making criteria. Nonetheless, this baseline will help decision makers clarify factors indispensable for assessing the risk that each option entails and achieve comprehensive cost-benefit evaluation among conflicting national interests.

Retention Risks

It is crucial to assess the risks and potential costs to various government interests that may arise from retention of an acquired vulnerability. The following criteria set a baseline for such analysis:

1. **Does the affected product or service have a high prevalence?** (As measured by market penetration; number of users potentially affected; their geographic distribution; and the range of affected product versions.)
2. **Are users heavily reliant on the product's security?** (For example, to protect highly sensitive personal, government, or corporate data.)
3. **Could exploitation of the identified vulnerability lead to direct harm?** (Financial, reputational, physical, etc.)
4. **Do adversaries possess or are they likely to become aware of the vulnerability?**
5. **Would adversaries be likely to exploit the vulnerability if they become aware of it?** (Attention should also be paid to adversaries' means and objectives in potential exploitation.)
6. **Would adversaries have a strong incentive to exploit or publicize the vulnerability?**

The above criteria are useful for assessing the scale and likelihood of risks associated with exploitation of a retained vulnerability by both the government in question and other actors upon its release into the digital ecosystem. Even if fully accurate quantitative data is not available for these questions, qualitative assessments would still form useful inputs to the decision-making process.

7. **Is there any potential mitigation mechanism that would adequately protect affected products and services against exploitation where such protection would be in the national interest?** (This is especially crucial where affected systems are part of the country's CI or government systems, as well as those of partners and allies.)
8. **Do communication channels exist that would enable the government to (i) inform relevant stakeholders of exploitation risk and (ii) enable swift implementation of targeted mitigation or patching?**

It is important that decision makers are able to consider whether mitigation measures exist that would be feasibly implementable and effective enough to protect the country's own critical systems,

as well as those of its international partners and allies. If such risks cannot be adequately mitigated, retention should not be pursued.

9. **Are affected products or services integral to ensuring the operation or resilience of critical national functions?** (For example, CI or other "essential" services.)
10. **Would a worst-case exploitation scenario risk negative impacts on the national interest of debilitating scale or severity?** (This includes second-order and third-order effects, as well as "known unknown" impacts.)
11. **Could similar impacts occur affecting supranational institutions, international allies or partners, international organizations, or, based on geopolitical context, adversaries / competitors?**

These three criteria direct decision makers' attention to the identification of a wide range of national interests possibly affected by retention and exploitation. This is a primary motivation for involving various government stakeholders in the disclosure decision-making process. Although specific interests of relevance can only be determined on a case-by-case basis, four perspectives carry across a broad range of circumstances:

1. Any threat posed to the security and integrity of CI from the potential retention and exploitation of a vulnerability should be thoroughly evaluated. In our present era of rapid technological advancement, an unprecedented number and variety of systems are interconnected, expanding the cyber attack surface. Such infrastructure includes energy infrastructure, healthcare facilities, transportation systems, and telecommunications. Once a vulnerability is acquired, it is imperative to precisely identify, by involving a wide range of subject-matter actors, which infrastructure and sectors are likely to be affected — both directly and indirectly. A failure to identify vulnerable CI at an early stage can allow unexpected and undesirable outcomes to arise from retention, even when that decision is made through a GVD-like process.
2. Diplomatic relations and potential impacts should also be taken into account. If a retained vulnerability is exploited and that exploitation causes harm to the national interests of other uninformed, unprepared countries, it would not only disrupt the immediate relationship with those countries but could also impair diplomatic trust and cooperation over the longer-term.
3. A retention decision could harm trust in the government from private sector actors if they eventually learn that the government made such a decision without due consideration of private sector interests (e.g., those of technology vendors). If the government chooses not to inform a company and it therefore fails to prevent or mitigate the damage caused by the exploitation, the company in question might suffer reputational damage as well as economic harm from lost business. If the company is headquartered in another country, then such harms could also damage bilateral diplomatic relations between the two nations.
4. As data protection and privacy laws continue to evolve worldwide, a decision to retain and use a vulnerability carries an ever-changing risk that the government operationalizing the vulnerability may commit legal infractions under domestic or international laws. Cyber exploitation often entails or leads to data breaches; thus, governments should be extremely cautious in any decision to retain and use a vulnerability in a manner that could fall afoul of these laws.

Disclosure Risks

While one of the primary objectives of GVD implementation is to open up the decision calculus around vulnerability handling to support disclosure when doing so is in the national interest, the risk and potential costs of such disclosure should also be considered to ensure a decision that maximizes and balances components of the national interest. The analysis should consider, at a minimum, the subsequent criteria:

1. **Would the relevant vendor or another entity be unwilling, unable, or otherwise disincentivized to rapidly develop and release an effective patch or mitigation?**
2. **Would adversaries be likely to reverse-engineer a patch to discover the vulnerability and attack unpatched systems?**

Even if the government discloses the zero-day information, the relevant vendor(s) or manufacturer(s) may not promptly patch the vulnerability. Alternatively, poor patching practices or other impediments to updating the affected software may mean that users do not install a released patch in a timely manner. The second criterion is of particular importance if the released patch or mitigation is unlikely to be applied to vulnerable systems. In these circumstances, disclosure can increase the chance of cyberattacks against still-vulnerable systems.

3. **Could disclosure reveal intelligence sources or methods, or otherwise disrupt other ongoing intelligence, military, or law enforcement operations?**
4. **Could disclosure negatively impact other operational capabilities?**
5. **Could similar impacts occur affecting supranational institutions, international allies or partners, international organizations, or, based on geopolitical context, adversaries / competitors?**

Despite the inherent challenges of doing so, attention should also be paid to the possibility that disclosure may harm national and / or international security in the ways described above. Especially when a zero-day is newly identified through in-house R&D by military, intelligence, or law enforcement services, a rush to disclose could ultimately provide adversaries with clues about the disclosing country's cyber capabilities and strategic approach. Disclosure could undermine ongoing operations among security allies if such considerations are not adequately identified and addressed. In order to avoid disruption of diplomatic relations, communication with partners in advance of disclosure is essential.

Utility & Necessity of Retention

Beyond the identification and minimization of risks to national interests, GVD processes should only allow for retention when an acquired vulnerability will concretely advance specific and immediate operational objectives. Stockpiling of vulnerabilities for future use is neither responsible nor even particularly strategic, given their limited shelf-life. Sample criteria for making this assessment are as follows:

1. **Can the vulnerability be exploited for specific purposes such as intelligence collection, cyber offensive operations, and law enforcement evidence collection?**
2. **Is there an urgent need to conduct the operation that would exploit the vulnerability?**
3. **Are there no alternative, less disruptive means of achieving the same operational benefits?**
4. **How significant is the expected short-term and / or long-term value of the expected benefit of conducting the operation?**

In many cases, military, intelligence and law enforcement communities are the principal stakeholders that will be making the case for a vulnerability's operational utility and the necessity of its exploitation. Zero-days can be leveraged by law enforcement agencies into investigative tools that facilitate access to information that would be otherwise accessible and may be crucial to law enforcement goals. Military and intelligence communities can also benefit from the exploitation of zero-days especially if their doing so could reduce or eliminate risk to human life or other undesirable characteristics of conflict. To legitimately claim that these operational advantages outweigh other national interests that would indicate a preference for

disclosure, the benefit and necessity of exploitation must be clear and immediate. When such benefits are concretely anticipated, the best course of action may be to retain the vulnerability until its exploitation and use in the relevant offensive operation has been completed. Subsequent disclosure would still require careful handling, again through the GVD framework, to avoid the risks explored throughout this section.

Exceptions

For some governments, there may be limited circumstances in which consideration of certain cybersecurity flaws under a GVD framework may not be desirable. These instances may introduce unnecessary complexity into the GVD workflow or make proper risk assessment either impossible or an inefficient use of limited resources.

The following situations, wherein GVD consideration and subsequent disclosure would not contribute to enhancing the security of the flawed product, may be considered:

1. When an affected product can no longer be patched by the relevant vendor or manufacturer, for instance if that company is defunct.
2. When an acquired vulnerability is determined to have been introduced "by design," suggesting that the product may still remain vulnerable to other as-yet undiscovered flaws introduced with similar malintent even if the specific vulnerability were itself to be patched.
3. Misconfiguration or configuration of a device that sacrifices security in lieu of availability, ease of use, or operational functionality.
4. Misuse of available device features that enables non-standard operation.
5. Misuse of engineering or configuration tools, techniques, and scripts for altering the functionality of the device.
6. Instances where a device or system has no inherent security features by-design.

Determining how to handle these kinds of particularly challenging situations will be important for countries with lower technical capacity, given the complexity and resources likely required to fully understand and manage them. While technical assessments, as would be required for a GVD, may still be desirable in such instances, actual decision-making could be driven through different policy frameworks or on an ad hoc basis. Although each government is encouraged to consider what specific circumstances should characterize any exceptions to its GVD policy, these exceptions should be determined ahead of time and limited in scope.

CONCLUSION

Governments are not singular entities. They encompass and have a duty to protect diverse and often conflicting interests across a wide range of policy domains. Just as they are able to authorize the extraction and use of natural resources while simultaneously protecting our natural environment, they can also make use of cyber vulnerabilities to advance legitimate law enforcement, military, or intelligence interests while minimizing disruption to our global digital ecosystem. However, when governments make use of cybersecurity vulnerabilities as a means of advancing narrowly defined national security interests without constraints that consider of the broader picture, they can create serious risks to a wide variety of social, economic, and public sector interests.

Extreme solutions, either a mandate for or a moratorium against disclosure of cybersecurity vulnerabilities, are both impractical and undesirable. Even beyond the considerations around actually operationalizing vulnerabilities, neither disclosure nor retention is risk-free across every possible circumstance. For these reasons, widespread adoption of inclusive, accountable, transparent, flexible, and balanced decision-making processes for government handling of vulnerabilities is more than desirable; it is a necessity. This paper extends existing best practices in this space and makes a credible argument for swift and widespread adoption of GVD frameworks tailored to the specific institutional contexts to which they would be applied.

We hope that policymakers will take on board the core messages of this position paper. In addition, we encourage cyber policy experts to freely build upon and extend these best practices to suit specific national contexts.

REFERENCES

1 European Systemic Risk Board. (February 2020). "Systemic Cyber Risk." Available from: https://www.esrb.europa.eu/pub/pdf/reports/esrb.report200219_systemiccyberrisk~101a09685e.en.pdf

2 Sarah Repucci & Amy Slipowitz. (October 2020). "Democracy Under Lockdown." Freedom House. Available from: https://freedomhouse.org/sites/default/files/2020-10/COVID-19_Special_Report_Final_.pdf

3 The White House. (November 2017). "Vulnerabilities Equities Policy and Process for the United States Government." Archived at: https://web.archive.org/web/20201207235503/https://www.whitehouse.gov/sites/whitehouse.gov/files/images/External%20-%20Unclassified%20VEP%20Charter%20FINAL.PDF

4 GCHQ. (November 2018). "The Equities Process." Available from: https://www.gchq.gov.uk/information/equities-process/

5 Canadian Security Establishment. (March 2019). "CSE's Equities Management Framework." Available from: https://cse-cst.gc.ca/en/media/media-2019-03-08/

6 Cybersecurity and Infrastructure Security Agency. (September 2020). "Binding Operational Directive 20-01: Develop and Publish a Vulnerability Disclosure Policy." US Department of Homeland Security. Available from: https://cyber.dhs.gov/bod/20-01/

7 Sven Herpig & Ari Schwartz. (January 2019). "The Future of Vulnerabilities Equities Processes Around the World." Lawfare. Available from: https://www.lawfareblog.com/future-vulnerabilities-equities-processes-around-world

8 Sven Herpig. (August 2018). "Governmental Vulnerability Assessment and Management." Stiftung Neue Verantwortung / Transatlantic Cyber Forum. Available from: https://www.stiftung-nv.de/sites/default/files/vulnerability_management.pdf

9 Tweede Kamer der Staten-Generaal. (December 2019). "Kamerdossier 35257, Nr. 4: Advies Van De Afdeling Advisering Van De Raad Van State En Reactie Van De Initiatiefnemer" [Translated from Dutch with Google Translate]. Available from: https://zoek.officielebekendmakingen.nl/kst-35257-4.html

10 Der Spiegel. (November 2014). "BND Will Informationen Über Software-Sicherheitslücken Einkaufen" [Translated from German with Google Translate]. Available from: https://www.spiegel.de/spiegel/vorab/bnd-will-informationen-ueber-software-sicherheitsluecken-einkaufen-a-1001771.html

11 Dennis Fisher. (July 2015). "VUPEN Founder Launches New Zero-Day Acquisition Firm Zerodium." ThreatPost. Available from: https://threatpost.com/vupen-launches-new-zero-day-acquisition-firm-zerodium/113933/

12 Zerodium. (Accessed December 2020). "Our Exploit Acquisition Program." Available from: https://zerodium.com/program.html

13 Ryan Gallagher. (July 2015). "Hacking Team Emails Expose Proposed Death Squad Deal, Secret U.K. Sales Push and Much More." The Intercept. Available from: https://theintercept.com/2015/07/08/hacking-team-emails-exposed-death-squad-uk-spying/

14 Joaquín Gil. (July 2015). "Hacking Team: 'Ofrecemos Tecnología Ofensiva Para la Policía'" [Translated from Spanish with Google Translate]. El País. Available from: https://elpais.com/internacional/2015/07/08/actualidad/1436343710_282978.html

15 Damien Leloup. (July 2017). "Après la Libye de Kadhafi, Amesys a Vendu son Système de Surveillance à l'Egypte de Sissi" [Translated from French with Google Translate]. Le Monde. Available from: https://www.lemonde.fr/pixels/article/2017/07/05/apres-la-libye-de-kadhafi-amesys-a-vendu-des-outils-de-surveillance-de-masse-a-l-egypte-de-sissi_5156085_4408996.html

16 Michael Daniel. (April 2014). "Heartbleed: Understanding When We Disclose Cyber Vulnerabilities." The White House. Available from: https://obamawhitehouse.archives.gov/blog/2014/04/28/heartbleed-understanding-when-we-disclose-cyber-vulnerabilities

17 US Congress. (May 2017). "H.R. 2481: A Bill To Establish The Vulnerability Equities Review Board, and For Other Purposes." US Government Publishing Office. Available from: https://www.congress.gov/115/bills/hr2481/BILLS-115hr2481ih.pdf

This paper was conceived and drafted by Josh Kenway, Maho Sugihara, Asaf Zilberfarb, and Pablo Tortolero as a research project for the Ford Dorsey Master's Program in International Policy at Stanford University's Freeman Spogli Institute for International Studies, with those authors' preliminary research undertaken with mentoring support by the Mozilla Corporation.

Subsequent refinement was undertaken by Josh Kenway and Michael Daniel at the Cyber Threat Alliance, and Ari Schwartz and John Banghart from the Center for Cybersecurity Policy and Law.

The authors are grateful to the many additional experts and practitioners who provided guidance and feedback to inform this project's highly collaborative development process.

https://cyberthreatalliance.org
PR@cyberthreatalliance.org
@CyberAlliance
cyber-threat-alliance

CENTER FOR CYBERSECURITY
POLICY AND LAW

https://centerforcybersecuritypolicy.org
jlamel@insight-dc.com
@CyberSecCenter