THE CYBER SECURITY PROJECT

Government's Role in Vulnerability Disclosure

Creating a Permanent and Accountable Vulnerability Equities Process

Ari Schwartz

Rob Knake

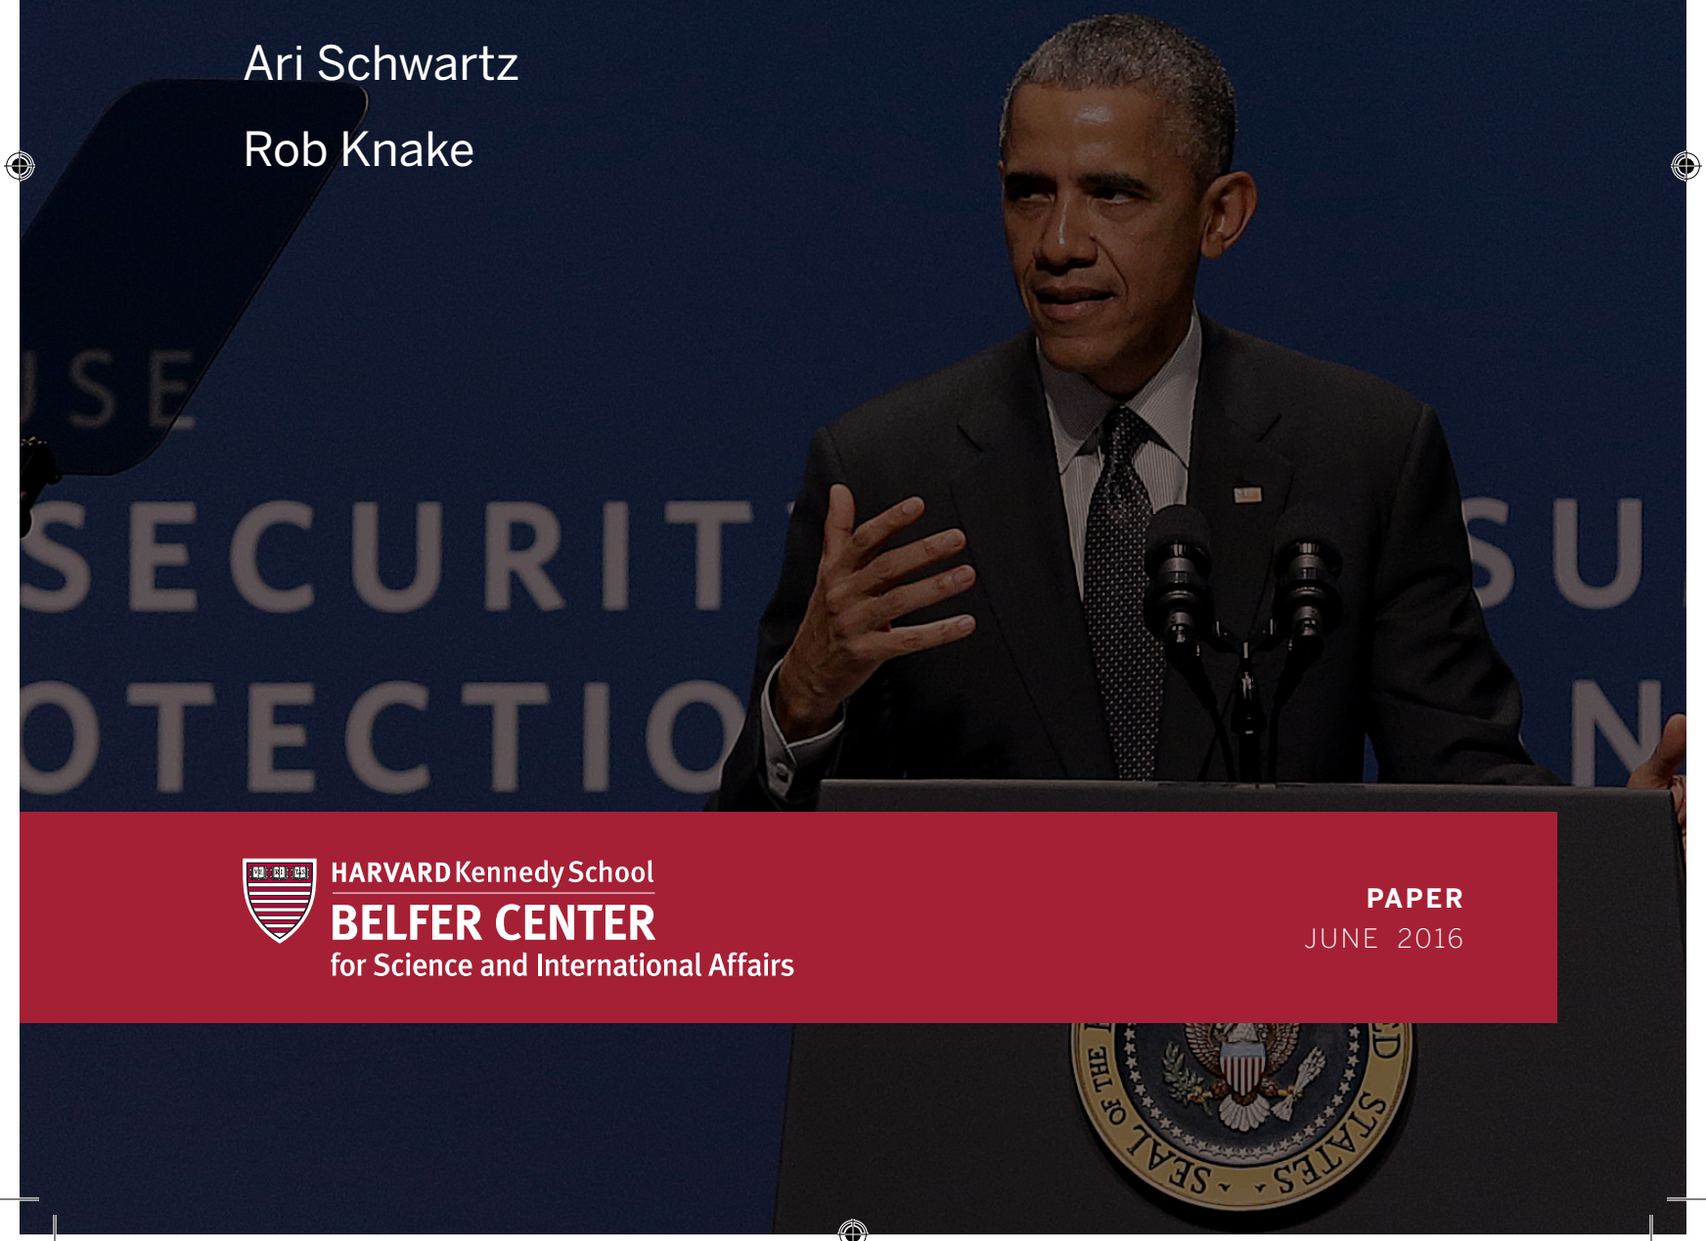

HARVARD Kennedy School
BELFER CENTER
for Science and International Affairs

PAPER
JUNE 2016

The Cyber Security Project
Belfer Center for Science and International Affairs
Harvard Kennedy School
79 JFK Street
Cambridge, MA 02138

www.belfercenter.org/cyber

Design & Layout by Andrew Facini

Cover photo: President Barack Obama speaks at the White House Summit on Cybersecurity and Consumer Protection in Stanford, Calif., Friday, Feb. 13, 2015. (AP Photo/Jeff Chiu)

THE CYBER SECURITY PROJECT

Government's Role in Vulnerability Disclosure

Creating a Permanent and Accountable Vulnerability Equities Process

Ari Schwartz

Rob Knake

HARVARD Kennedy School
BELFER CENTER
for Science and International Affairs

DISCUSSION PAPER 2016-03
JUNE 2016

About the Authors

Ari Schwartz is Managing Director for Cybersecurity Services at Venable LLP, Coordinator of the Coalition for Cybersecurity Policy and Law, and Former Special Assistant to the President and Senior Director for Cybersecurity Policy at the White House National Security Council.

Rob Knake is the Whitney Shepardson Senior Fellow at the Council on Foreign Relations and Former Director for Cybersecurity Policy at the White House National Security Council.

Acknowledgments

The authors would like to thank **Brian Zimmet** at Venable for his substantial input on this paper.

The views in this paper are the authors' and do not necessarily reflect the views of Venable, any of its clients, the members of the Coalition for Cybersecurity Policy and Law, or the Council on Foreign Relations or its members. This article was submitted to the National Security Council for review prior to publication.

Table of Contents

President Barack Obama speaks at the White House Summit on Cybersecurity and Consumer Protection in Stanford, Calif., Friday, Feb. 13, 2015.

AP Photo/Jeff Chiu

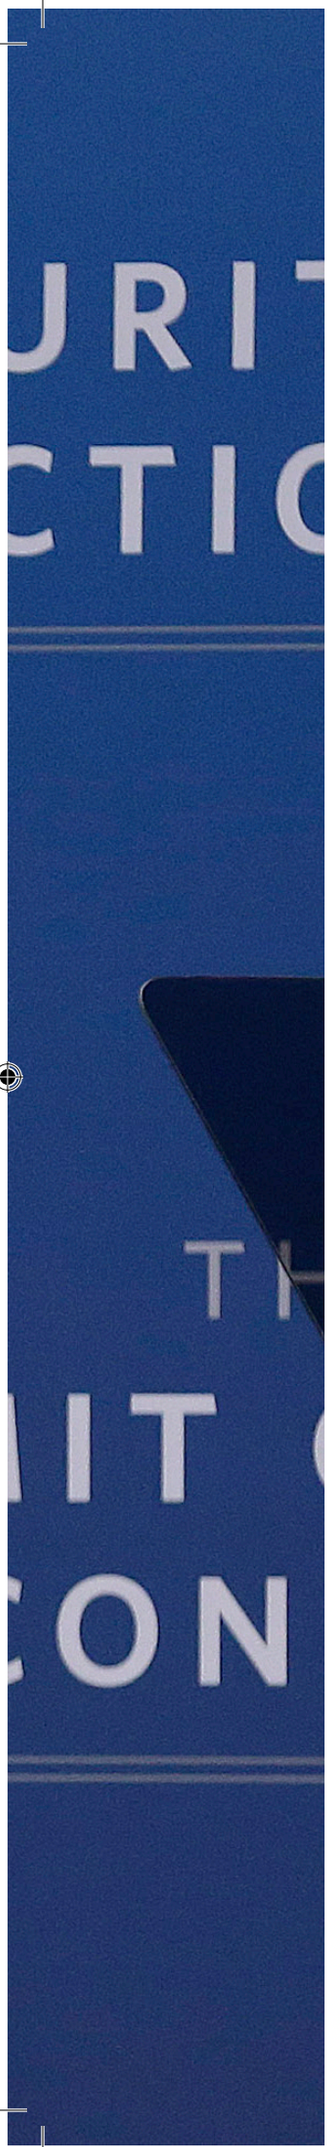

Introduction

When government agencies discover or purchase zero day vulnerabilities, they confront a dilemma: should the government disclose such vulnerabilities, and thus allow them to be fixed, or should the government retain them for national security purposes? This is a difficult question because the government is simultaneously charged with protecting the nation in cyberspace and with intelligence, law enforcement, and military missions that may require the use of such vulnerabilities. A decision by the government to retain a zero day vulnerability likely undercuts general cybersecurity, while disclosing information about a zero day vulnerability so vendors can patch it could undercut the ability of law enforcement to investigate crimes, intelligence agencies to gather intelligence, and the military to carry out offensive cyber operations.

The debate over this issue is complex. Some commentators take the position that the government should immediately release all zero day vulnerabilities, irrespective of their intelligence or national security value.[1] At the same time, there are circumstances where retention of a zero day vulnerability by the government for law enforcement or national security purposes is justified, as long as there are clear limits on and adequate oversight of the decision to retain and use such a vulnerability. For example, if a law enforcement agency has an ongoing investigation on a suspect and the only information is coming through communications legally intercepted through a previously unknown vulnerability, the balance may very well be for the agency to keep the vulnerability, at least until the end of the investigation.[2]

Only in recent years has the government created a Vulnerability Equities Process ("VEP"), and attempted to explain how the government determines whether to release or retain a zero day vulnerability. As explained by White House Cybersecurity Coordinator Michael Daniel,

1 *See, e.g.,* Bruce Schneier, "Disclosing vs. Hoarding Vulnerabilities", *Schneier on Security*, May 22, 2014, https://www.schneier.com/blog/archives/2014/05/disclosing_vs_h.html.

2 *See, e.g.,* Allan Friedman, Tyler Moore, and Ariel D. Procaccia, *Cyber-Sword v. Cyber-Shield: The Dynamics of US Cybersecurity Policy Priorities*, Center for Research on Computation & Society, Harvard University, http://web.mit.edu/ecir/pdf/Friedman%20cyberwar-governance.pdf.

the existing VEP uses a "deliberate process that is biased toward responsibly disclosing [a] vulnerability. . . ."[3] Daniel also explained, however, that there are "no hard and fast rules" governing the VEP, although he did outline a series of questions that he considers when presented with a zero day vulnerability disclosure issue.[4]

While the current VEP functions as intended, the guidelines articulated in the Daniel blog post may be undercut in a future administration unless formalized now. Some individual VEP decisions must remain classified, but the high-level criteria that informs disclosure or retention decisions should be subject to public debate and scrutiny. Furthermore, certain information about the implementation of the VEP, particularly the aggregate numbers of zero day vulnerabilities discovered, the aggregate numbers of such vulnerabilities disclosed (as opposed to retained for government use), and the length of time that vulnerabilities are kept before disclosure, do not compromise sources and methods of how these vulnerabilities may have been discovered. Public and official release of information about the process with clear oversight would increase public confidence in the program, and in the government's commitment to the core principles laid out by Administration to date, and could become a model for other nations around the world.

3 Michael Daniel, "Heartbleed: Understanding When We Disclose Cyber Vulnerabilities", *White House Blog*, April 28, 2014 ("Daniel Blog Post"), https://www.whitehouse.gov/blog/2014/04/28/heartbleed-understanding-when-we-disclose-cyber-vulnerabilities.

4 *Id.*

1. Overview: Genesis of the Vulnerability Equity Issue

Zero day vulnerabilities are software weaknesses that are unknown to the software's developer or users. When discovered by a third party, the developer literally has had "zero days" to develop a patch and users have had "zero days" to implement the patch or take other protective measures. These coding flaws can expose all users of the relevant software to security risks until those flaws are discovered and a patch can be developed and implemented. A vulnerability ceases to be a zero day and becomes a known vulnerability when the vulnerability becomes publicly known. Every piece of software contains vulnerabilities. The majority of vulnerabilities are benign and have no material impact on the functionality of the product. Some vulnerabilities, however, may be exploitable, enabling unintended (and potentially malicious) functionality.

There is a wide array of actors that search for and discover zero day vulnerabilities, including government agencies, software developers, security researchers, and a host of bad actors, from adversarial nation states to criminal gangs. Their motivations vary: some may engage in security research as a public good, disclosing vulnerabilities they discover to the vendor to improve the security of products in use throughout the world. There also exist a series of markets, with varying degrees of legality, within which zero day vulnerabilities can be sold and purchased. When federal agencies discover vulnerabilities as part of carrying out law enforcement and intelligence missions, the government must determine whether knowledge of these vulnerabilities should be restricted and used for these purposes or disclosed in the national interest of improving cybersecurity. The VEP is the process by which these decisions are made.

Origins of the VEP

The genesis and contours of the existing VEP are reflected in a series of documents obtained and made public in 2015 and 2016.[5] A basic overview of the origins of the VEP is set forth in a document entitled "Vulnerability Equities Process Highlights" ("Highlights Paper"). We can trace the origins of the VEP to a January 2008 directive, signed by President George W. Bush. This directive, known as National Security Policy Directive 54 (NSPD 54), established a US-government-wide effort called the Comprehensive National Cybersecurity Initiative ("CNCI"). One component of the CNCI required the Departments of State, Defense, Homeland Security, and Justice, as well as the Director of National Intelligence, to develop "a joint plan for the coordination and application of offensive capabilities to defend US information systems."[6]

This joint plan[7] noted, among other things, that the discovery of vulnerabilities "may present competing equities for [government] offensive and defensive mission interests" and recommended that "actions taken in response to knowledge of a specific vulnerability must be coordinated to ensure the needs of each of these 'equities' are addressed."[8] The plan recommended the development of a "Vulnerabilities Equities Process."

A working group, led by the Office of the Director of National Intelligence, developed the VEP during 2008 and 2009 in response to the Joint Plan's recommendation. It consisted of members from the National Security Council, Central Intelligence Agency, Defense Intelligence Agency, Justice Department, Federal Bureau of Investigation ("FBI"), Department of Defense, Department of State, Department of Energy, and Department of Homeland Security. This working group ultimately produced a document

5 As described later, the documents were made public by the Electronic Frontier Foundation (EFF) in 2015 and 2016 in response to a Freedom of Information Act request. *See* https://www.eff.org/cases/eff-v-nsa-odni-vulnerabilities-foia.

6 *See* National Security Policy Directive-54/ Homeland Security Policy Directive-23 at Paragraph (49), https://fas.org/irp/offdocs/nspd/nspd-54.pdf. The Highlights Paper, a single page document citing this provision of National Security Policy Directive-54/Homeland Security Policy Directive-23, can be referenced at http://www.wired.com/wp-content/uploads/2015/03/Vulnerability-Equities-Process-Highlights-7.8.10-DOC-65-redactions_Redacted1.pdf.

7 The Joint Plan is referenced in the main VEP document, the "Commercial and Government Information Technology and Industrial Control Product or System Vulnerability Equities Policy and Process" ("VEP Document"). The VEP Document can be found at https://www.eff.org/files/2015/09/04/document_71_-_vep_ocr.pdf.

8 VEP Document at 2 (quoting Joint Plan).

entitled "Commercial and Government Information Technology and Industrial Control Product or System Vulnerability Equities Policy and Process" ("VEP Document"), which lays out the process that the government apparently continues to follow today. The VEP Document is dated February 16, 2010.

VEP Document

The VEP Document demonstrates that the government established a process to determine whether vulnerabilities discovered or purchased by the government or its contractors should be retained for government use or revealed to the appropriate vendor for patching.[9]

Under the process as outlined by the VEP Document, the VEP is intended to "establish[] policy and responsibilities for disseminating information about vulnerabilities discovered by the United States Government (USG) or its contractors, or disclosed to the USG by the private sector or foreign allies in Government Off-The-Shelf (GOTS), Commercial Off-The-Shelf (COTS) or other commercial information technology or industrial control products or systems (to include both hardware and software)."[10] More specifically, the paper "defines a process to ensure that dissemination decisions regarding the existence of a vulnerability are made quickly, in full consultation with all concerned government organizations, and in the best interest of government missions of cybersecurity, information assurance, intelligence, counterintelligence, law enforcement, military operations, and critical infrastructure protection."[11] The policy "applies to all components, civilian and military personnel, and contractors of the United States Government . . . ."[12]

The VEP Document goes on to give directions for appropriate classified treatment for "vulnerabilities discovered by the USG or by non-USG entities under contracts with the USG, or disclosed to the USG by the private

9 The public version of the VEP Document, obtained and posted by EFF, is heavily redacted, but still has enough information to see the outlines of the VEP process.

10 VEP Document at 1.

11 *Id.*

12 *Id.*

sector prior to entry into this process" and then directs further that "USG entities shall introduce any such vulnerability discovered into the following Vulnerabilities Equities Process (VEP)."[13] The VEP Document clarifies further that a vulnerability should be put through the VEP if it is "both newly discovered and not publicly known."[14] The VEP Document exempts from the VEP any vulnerability discovered before its effective date (February 16, 2010).[15] It also exempts from the process any vulnerability discovered during the course of open and unclassified federally-sponsored research.[16]

The VEP Document establishes the Information Assurance Directorate of the NSA as the Executive Secretariat of the VEP.[17] It also establishes an interagency Equities Review Board ("ERB") for making decisions on whether to retain or disclose a vulnerability.[18] The composition of the ERB remains classified, however.

Under the process, the agency that comes into possession of a vulnerability that is newly discovered and not publicly known is required to notify the Executive Secretary, which then disseminates the vulnerability to all relevant agency Points of Contact ("POCs"), who "are responsible for ensuring that applicable cybersecurity, cyber defense, information assurance, intelligence, counterintelligence, law enforcement, or other offensive cyber operations equities of their organization are appropriately represented in the process."[19] Each such agency then is responsible for designating one or more Subject Matter Experts ("SMEs") to participate in a discussion convened by the Executive Secretary to arrive at a consensus on whether the vulnerability should be retained by the government or disclosed for patching.[20] Ultimately, the ERB is charged with making the decision on whether to disclose or retain a discovered vulnerability, and the ERB acts

13 *Id.*at 2.

14 *Id.* at 5.

15 *Id.*at 4.

16 *Id.*

17 *Id.*at 5.

18 *Id.* at 7.

19 *Id.* at 5-6.

20 *Id.*at 7.

by majority vote.[21] An affected agency is entitled to appeal the ERB's decision, although the appeals process itself remains redacted.[22]

The VEP Document provides guidance for implementation of ERB decisions. The guidance on implementation of a decision not to disclose a vulnerability continues to be classified.[23] For implementation of decisions to disseminate information on vulnerabilities, the guidance requires the ERB to establish "guidelines for disseminating that information, including mitigation strategies, to the cyber security centers that are primarily responsible for defending or coordinating the defense of networks and systems, as well as offensive entities."[24] The VEP Document also provides for an annual oversight mechanism involving the production of an annual report by the Executive Secretariat, although the identity of the overseer of the process remains classified.[25]

Reinvigoration of the Vulnerability Equities Process

Although the VEP was established in 2010, its existence was revealed publicly only after questions were raised about the government's use of zero day vulnerabilities for intelligence and offensive purposes, including the government's practice of purchasing zero day vulnerabilities, following the leaks of classified information by Edward Snowden in 2013. [26]

The report issued by the President's Review Group on Intelligence and Communication Technologies ("President's Review Group") in December 2013 did not mention an existing VEP, but did contain recommendations about how the government should manage vulnerability equities. Specifically, Recommendation 30 of the President's Review Group report stated:

21 *Id.*

22 *Id.*

23 *Id.* at 7-8.

24 *Id.*at 8.

25 *Id.* at 8-9.

26 Brian Fung, "The NSA hacks other countries by buying millions of dollars' worth of computer vulnerabilities", *Washington Post*, August 31, 2013, https://www.washingtonpost.com/news/the-switch/wp/2013/08/31/the-nsa-hacks-other-countries-by-buying-millions-of-dollars-worth-of-computer-vulnerabilities/.

> *We recommend that the National Security Council staff should manage an interagency process to review on a regular basis the activities of the government regarding attacks that exploit a previously unknown vulnerability in a computer application or system. These are often called 'Zero Day' attacks because developers have had zero days to address and patch the vulnerability. US policy should generally move to ensure that Zero Days are quickly blocked, so that the underlying vulnerabilities are patched on government and other networks. In rare instances, US policy may briefly authorize using a Zero Day for high priority intelligence collection, following senior, interagency review involving all appropriate departments.*[27]

The Obama Administration released information about the VEP only after Bloomberg News alleged in April 2014 that the NSA had known about the then-recently revealed Heartbleed vulnerability, and exploited Heartbleed for its own purposes instead of disclosing the vulnerability to be patched.[28] In response to that allegation, which the NSA vigorously denied, the White House acknowledged that the government sometimes relies on zero day vulnerabilities for intelligence and other, related purposes, rather than disclosing such vulnerabilities and allowing them to be patched.[29] However, the White House also asserted that it had reviewed the recommendations of the President's Review Group and had determined, as of January 2014, that the government's policy should be that there is a "bias" toward disclosure to vendors for patching rather than retention by the government.[30] The exception to this government "bias" toward disclosure is if there is a "clear national security or law enforcement need. . . ."[31] The Administration

---

27 *Liberty and Security in a Changing World: Report and Recommendations of the President's Review Group on Intelligence and Communications Technologies*, December 12, 2013, at https://www.whitehouse.gov/sites/default/files/docs/2013-12-12_rg_final_report.pdf.

28 Michael Riley, "NSA Said to Have Used Heartbleed Bug, Exposing Consumers", *Bloomberg News*, April 12, 2014, http://www.bloomberg.com/news/articles/2014-04-11/nsa-said-to-have-used-heartbleed-bug-exposing-consumers.

29 David E. Sanger, "Obama Lets N.S.A. Exploit Some Internet Flaws, Officials Say", *New York Times*, April 12, 2014, http://www.nytimes.com/2014/04/13/us/politics/obama-lets-nsa-exploit-some-internet-flaws-officials-say.html?_r=0.

30 *Id.*

31 *Id.*

described these decisions as the implementation of a "reinvigorated" process for balancing the equities surrounding zero day vulnerabilities.[32]

Daniel Blog Post

Responding to the Heartbleed allegation against the NSA, White House Cybersecurity Coordinator Michael Daniel authored a White House blogpost in late April 2014 that further outlined the Obama Administration's policy regarding zero day vulnerabilities. Echoing the White House's statement in early April 2014, Daniel stated that "[t]his spring, we re-invigorated our efforts to implement existing policy with respect to disclosing vulnerabilities—so that everyone can have confidence in the integrity of the process that we use to make these decisions."[33] Addressing allegations that the government hoards zero day vulnerabilities, Daniel asserted that "[b]uilding up a huge stockpile of undisclosed vulnerabilities while leaving the Internet vulnerable and the American people unprotected would not be in our national security interest."[34] Nonetheless, while asserting that disclosure of zero day vulnerabilities is in the national interest "in the majority of cases," he categorically rejected the suggestion that the government should "completely forego this tool as a way to conduct intelligence collection, and better protect our country in the long-run."[35]

Daniel explained that the government has a "disciplined, rigorous, and high-level decision-making process for vulnerability disclosure" that is inter-agency in nature, and that explores all the pros and cons of disclosure. He emphasized that "there are no hard and fast rules" governing the process, but outlined the following factors, lightly edited for clarity, that the ERB considers before deciding whether to disclose a zero day vulnerability:[36]

32 *Id.*

33 Daniel Blog Post.

34 *Id.*

35 *Id.*

36 *Id.*

- The extent of the vulnerable system's use in the Internet infrastructure;
- The risks posed and the harm that could be done if the vulnerability is left unpatched;
- Whether the Administration would know if another government or organization was exploiting the vulnerability;
- Whether the vulnerability is needed to obtain intelligence (i.e., how badly does the US government need the information, and are there alternative means of obtaining it);
- How likely it is that others will discover the vulnerability;
- Whether the government can use the vulnerability for a short period of time before disclosing it; and
- Whether the vulnerability can be patched or otherwise mitigated.

Daniel concluded by asserting that the government "weigh[s] these considerations through a deliberate process that is biased toward responsibly disclosing the vulnerability . . . ."[37] He also emphasized the need for sufficient transparency in the process to "instill some confidence that your government is acting responsibly in the handling of this important issue."[38]

Subsequent Disclosures

The disclosures by the White House in response to the Heartbleed allegations made clear that the government has a VEP, but provided no additional information on the VEP itself, or when it was developed. In response to these disclosures, the Electronic Frontier Foundation (EFF) filed a Freedom of Information Act request for documents related to the VEP described in the Daniel Blog Post, and then a lawsuit to compel the disclosure of those documents. The result of that suit was the release of a series of documents, the most significant of which are the Highlights Paper and the VEP Document.

37 *Id.*

38 *Id.*

In the period between the filing of the EFF suit and the release of the Highlights Paper and the VEP Document, there were several additional statements by government officials that shed further light on the VEP and its implementation. In a speech delivered in November 2014, Admiral Michael Rogers, the head of the NSA, discussed the government's vulnerability disclosure policies, and stated that "by orders of magnitude, the greatest number of vulnerabilities we find, we share."[39]

While the Obama Administration deserves credit for re-invigorating the process and for demonstrating a clear bias toward disclosure, the fact that the process fell into disuse from when it went into effect in 2010 until the Intelligence Review Group made its recommendations in 2014 is troubling. In an interview given to WIRED Magazine, Daniel asserted that the "default-disclosure policy was established in 2010" but that "it 'had not been implemented to the full degree that it should have been,' hence the government's use of the term 'reinvigorated' to describe this new phase."[40] In particular, the "relevant agencies . . . 'had not been doing sufficient interagency communications and ensuring that everybody had the right level of visibility across the entire government' about vulnerabilities that were discovered."[41] Daniel asserted that "although 'they probably were disclosing the vulnerability' by default, they 'may not have been communicating that to all the relevant agencies as regular as they should have been.'"[42] Agencies "might have been communicating 'at the subject-matter expert level,' but the communication may not have been happening as consistently, in as coordinated a fashion or within the timelines that the policy dictated."[43]

Thus, from the public statements made after the issuance of the Daniel Blog Post, it can be concluded that the VEP has at least once in its short history fallen out of use. While the process appears now to be functioning well, should the issue once again fade from attention, the policy in its current form gives few guarantees that adherence to it will not lapse. After several years of experience with a functioning process, the Obama Administration

39 Kim Zetter, "U.S. Gov Insists it Doesn't Stockpile Zero-Day Exploits to Hack Enemies", *WIRED*, November 17, 2014, http://www.wired.com/2014/11/michael-daniel-no-zero-day-stockpile/.

40 *Id.*

41 *Id.*

42 *Id.*

43 *Id.*

should take steps in its final months to ensure that the policy and process are carried over into the next Administration.

2. Recommended Improvements to the Vulnerability Equities Process

The core concepts around vulnerability equity that the government has articulated to date, particularly those set forth in the Daniel Blog Post, are unquestionably a helpful starting point for cementing the VEP into the future. In specifically-defined circumstances, government use of such vulnerabilities may be justified. In the recent debate between the FBI and Apple over how the FBI could access information on a deceased terrorist's iPhone, many in the technical community have signaled a strong preference for so-called "lawful hacking" over government mandated back doors.[44] Such lawful hacking will require the use of unknown vulnerabilities some of the time. For these reasons, the disclosure issue is not as one-sided as many interests in the debate make it out to be. The potential options for approaching and resolving this issue are not binary, but fall along a spectrum. The question for the government is where along the spectrum its approach will fall.

At this point, the government must formalize the process and publicly adopt the principles outlined in the Daniel Blog Post. The government's process should continue to be biased towards disclosure, and retention of vulnerabilities for government use should be permissible in defined circumstances. By affirming existing policy in higher-level, unclassified governing principles, the government would add clarity to the process and help set a model for the world. If all of the countries with capabilities to collect vulnerabilities had a policy of leaning toward disclosure, it would be valuable to the protection of critical infrastructure and consumers alike as well as US corporate interests.

44 Selina Wong, "Security Experts: FBI Doesn't Need Apple's Help to Hack into iPhone", *Bloomberg*, March 4, 2016, http://www.expressnews.com/business/national/article/Security-experts-FBI-doesn-t-need-Apple-s-6871493.php.

To be clear, more formal affirmation does not equate to publicizing individual disclosure decisions or the deliberations of the government agencies charged with providing input into such decisions. In many cases, it likely would not serve the interests of national security to make such information public. However, the principles guiding these decisions, as well as a high-level map of the process that will be used to make such decisions, can and should be public. Some critics believe that information not contained in the blog post suggests that there are large loopholes in the VEP and others suggest that the lack of detailed disclosure and redactions in the version of the document released to EFF under FOIA demonstrate that there are secrets that are purposely being withheld.[45] The fact that the process is only public through the Daniel Blog Post and a heavily redacted document has simply led to an unnecessary lack of public confidence in the policy and adds to a lack of trust in essential efforts of government to work cooperatively with businesses to secure networks.

In particular, the Obama Administration should take the following steps with respect to the existing VEP:

- **Issue an executive order to formalize and require government-wide compliance with the VEP.** While the existing VEP is a good starting point, the document is an agreement among the participating agencies and does not carry the weight of an executive order signed by the President. Thus, there are few consequences for agencies that choose not to participate in the process. An executive order would make the policy binding on agencies; the public nature of such a document would be a significant demonstration that the government is acting in good faith.
- **Make public the high-level criteria that will be used to determine whether to disclose to a vendor a zero day vulnerability in their product, or to retain the vulnerability for government use.** The criteria set forth in the Daniel Blog Post are a good start for the governing principles, but, as Daniel emphasized, "there are no hard

45 *See* Mathew J. Schwartz, "White House Details Zero-Day Bug Policy", *Information Week's Dark Reading*, April 15, 2014, http://www.darkreading.com/analytics/white-house-details-zero-day-bug-policy/d/d-id/1204483 and Ashley Carman, "Documents on NSA Zero-Day Policy Provide Little Insight", *SC Magazine*, March 30, 2015, http://www.scmagazine.com/electronic-frontier-foundation-obtains-zero-day-documents/article/406230/.

and fast rules" for making disclosure decisions. However, we think it is possible to formalize guidelines for disclosure decisions while preserving flexibility in the decision-making process. Furthermore, Daniel stressed that the principles set forth in the blog post are ones that he would look at if making a disclosure decision, thus implying that they are not necessarily applicable to other agencies in the government, or even to other officials in the White House itself. The criteria used by the government to make zero day vulnerability disclosure decisions have very important implications for national security, and thus should be made definitive and formalized as part of the executive order.

- **Define clearly the process to be followed in making a disclosure decision with respect to a zero day vulnerability.** The new executive order should define the agencies that are to be involved in the process, the manner in which the issue is to be raised and debated, the manner in which recommendations to the ultimate decision-maker are developed, and—perhaps most importantly—who will be making the ultimate decision with respect to disclosure or retention. It may be that the nature of the process, and even the identity of the decision-maker, will vary depending on the nature of the vulnerability at issue. That is, in some instances, the process may be more intensive, and the decision-maker might be at a higher-level of the government, while in other instances, the process might be somewhat less intensive, and the decision might be made at a lower level. Indeed, there is precedent for such processes with respect to offensive cyber activities, and it may be worthwhile to implement a similar set of processes to govern vulnerability disclosures.

- **Ensure that any decision to retain a zero day vulnerability for government use is subject to periodic review.** The executive order should require vulnerabilities be disclosed to the responsible party once (1) the government has achieved its desired national security objectives or (2) the balance of equities dictate that the vulnerability should be disclosed. While government use of zero day vulnerabilities may be justified in limited circumstances, it is imperative that such vulnerabilities not be retained any longer than necessary. The

circumstance to be avoided is one where the government is allowed to retain zero day vulnerabilities indefinitely, without periodic re-review of the justification for retaining those vulnerabilities, or of the equities of retaining, as opposed to disclosing, such vulnerabilities. The VEP should have a review process built into it in order to avoid this outcome.

- **Prohibit agencies from entering into non-disclosure agreements with vulnerability researchers and resellers.** When the government purchases a zero day vulnerability or a tool to exploit such a vulnerability, the seller should be legally obligated to foreswear reselling the vulnerability or tool to a third party. The government must have exclusive rights to the vulnerability or tool. If it does not obtain these rights, including the right to disclose the vulnerability, it runs the risk that it could be sold or shared with other actors working against the national security interest of the United States. Thus, the executive order should require agencies to obtain the ability to disclose any vulnerability they purchase.

- **Transfer the Executive Secretary function from NSA to the Department of Homeland Security.** Under the current publicly available version of the VEP, the Information Assurance Directorate (IAD) within the NSA serves as the Executive Secretary for the process. While NSA responsibility for managing this process has been raised in the past as a concern that might tilt it toward retention, IAD has historically been a strong voice for cybersecurity over the agency's intelligence collection mission. The decision by Admiral Rogers to merge IAD and the Signals Intelligence Directorate into a single organization, however, throws into question whether NSA can serve as the neutral manager of the process.[46] Even if the NSA can internally find a means to manage this process in an even-handed manner, there is still an appearance of conflict that raises unnecessary questions about the impartiality of the VEP. The executive order should transfer this responsibility to the Department of Homeland Security, which has developed a strong capability in vulnerability research and software assurance.

46 Ellen Nakashima, "National Security Agency Plans Major Reorganization", *Washington Post*, February 2, 2016, https://www.washingtonpost.com/world/national-security/national-security-agency-plans-major-reorganization/2016/02/02/2a66555e-c960-11e5-a7b2-5a2f824b02c9_story.html.

- **Direct the Executive Secretary to issue a public report on an annual basis on the status of the program.** The executive order should require the report to include the aggregate numbers with respect to zero day vulnerabilities discovered, aggregate numbers of disclosed vulnerabilities,[47] the average length of time that vulnerabilities were kept before disclosure, and aggregate numbers of vulnerabilities retained for government use. While publicizing individual decisions on vulnerability disclosure would have adverse consequences for national security, the disclosure of aggregate numbers for a given period (for example, a year) should not raise the "sources or methods" issues that release of an individual vulnerability could have. On the other hand, it would address charges that the government hoards large numbers of zero day vulnerabilities. Indeed, this kind of transparency will increase confidence in government decision-making.

Beyond the executive order, Congress should take further action to strengthen oversight of the government's use of vulnerabilities and ensure that the process results in a more secure cyber ecosystem. Specifically, Congress should:

- **Expand Congressional oversight of the government's use of vulnerabilities:** The relevant Congressional committees in both houses should take steps to oversee the implementation of the VEP, and to ensure that it is operating as intended. Committees with oversight of law enforcement, intelligence, and offensive operations all should increase their focus on this topic.

- **Mandate oversight by independent bodies within the Executive Branch:** Congressional oversight will require that Congress is adequately informed on agency adherence to the policy. Congress should direct the Inspectors General and the Privacy and Civil Liberties Oversight Board to audit the implementation of the VEP, review agency adherence to the VEP, and ensure that VEP decisions are being made in accordance with the applicable formalized process and criteria.

47 There is precedence for release of aggregate number of disclosed vulnerabilities. The British intelligence agency Government Communications Headquarters (GCHQ) announced in April that it disclosed 20 previously unknown vulnerabilities during the first quarter of 2016. *See Motherboard (http://motherboard.vice.com/read/gchq-vulnerabilities-mozilla-apple).*

- **Expand funding for both offensive and defensive vulnerability discovery and research.** In order for Federal agencies to end the retention of a vulnerability that is being used for law enforcement or intelligence purposes, more often than not they will need a new vulnerability to replace it. When a decision to retain a vulnerability is made, work should immediately begin on finding a new vulnerability that could replace it. Thus, a cycle of vulnerability discovery, exploitation, and disclosure could be initiated that would promote a more secure ecosystem through the discovery and disclosure of more vulnerabilities while allowing national security use of vulnerabilities for short durations. Such a cycle, however, would require significant funding so that agencies are discovering significantly more vulnerabilities than necessary. Congress should also increase funding for programs at the Commerce Department and the Department of Homeland Security to strengthen vulnerability disclosure and vulnerability mitigation in the private sector, particularly with respect to open source software.

3. Conclusion

The reinvigorated VEP process as constituted has served its function well during the current Administration. It should be strengthened by formalizing the process and by putting the full weight of the President behind it. Stronger oversight is necessary to ensure that the equities in favor of disclosure are being taken into account by the process and by each individual agency participant. Finally, more funding is needed so that agencies do not have to retain vulnerabilities for an extended period merely because they do not have the funding necessary to obtain replacements.

In an ideal world, the Federal government would have no need or use for zero day vulnerabilities. Yet in the digital age, completely foregoing the use of zero day vulnerabilities would amount to ending signals intelligence. While some may advocate for such a step, continued threats to the national security of the United States from foreign nations and terrorist groups demand that we maintain the ability to collect intelligence. At the same time, our dependence on digital infrastructure requires that the government take every step possible to secure this infrastructure against threats to it. The VEP is a policy that seeks to balance these two national security interests. By taking the steps recommended in this paper, the Federal government will have a better chance at serving national security, commercial, and personal computing security interests.

The Cyber Security Project
Belfer Center for Science and International Affairs
Harvard Kennedy School
79 JFK Street
Cambridge, MA 02138

www.belfercenter.org/cyber

Printed in the United States of America