Response to IT-ISAC EI-SIG Request for Information

Jack Cable

To whom it may concern:

I am writing in response to the IT-ISAC EI-SIG Request for Information regarding coordinated vulnerability programs for election systems. I am a security researcher with experience disclosing security vulnerabilities to hundreds of companies under vulnerability disclosure programs and bug bounties. As we have seen, the security of election systems is of increasing national concern, and it is crucial that we take steps to better security and transparency of election systems ahead of the 2020 elections. I firmly believe that instituting coordinated vulnerability disclosure policies is one of the most effective actions than can be taken to bolster election security, and am glad to see interest from election vendors in starting such policies. After engaging with election vendors and election officials at the July congressional roundtable, I am confident that such policies can be implemented and hope that this response can be of aid in structuring such policies. Such policies should apply to all layers of election technology, ranging from internet accessible systems like voter registration systems and election-night reporting software, to physical systems including voting machines and tabulation devices. In this response, I outline best practices for managing coordinated vulnerability disclosure programs, suggestions for testing sensitive and isolated election systems, and industry standards for public disclosure. I am always happy to discuss these topics, and my contact information is provided at the end.

I. Industry coordinated vulnerability disclosure best practices

Vulnerability disclosure policies are rapidly increasing in adoption across industry. Ranging from tech companies like Facebook and Google to financial services such as Chase and Capital One and government agencies including the U.S. Department of Defense, organizations across all sectors are incorporating public security scrutiny into their security model. These organizations are realizing that receiving public feedback on their security is not only beneficial, but necessary: no system is completely secure, and no organization can expect to find all of its vulnerabilities internally. Given that the presence of vulnerabilities are inevitable, the only way forward is through transparency in security and incorporating public feedback. Fortunately, the

vulnerability disclosure ecosystem has matured, and there exist common best practices that companies and researchers follow for smooth disclosure processes.

At its core, a vulnerability disclosure policy outlines systems in scope, contact methods, researcher expectations, and a safe harbor. As a reference, I have compiled an example vulnerability disclosure policy at https://github.com/cybertransparency/vdp-terms. As a baseline, a vulnerability disclosure policy should have an open scope, applying to all of an organization's assets, and should openly allow any security researcher to participate by default. I would particularly like to emphasize the presence of a safe harbor in a vulnerability disclosure policy. On both sides there should exist ground rules for what actions are and are not permitted. Researchers are aware that they must adhere to the scope and rules in place, and a safe harbor assures researchers that should they follow the scope and rules of the policy, their actions will be authorized and not penalized. This is important in encouraging productive research and gives researchers a safe method to disclose vulnerabilities. On the whole, a vulnerability disclosure policy need not be a complicated document; rather, it simply outlines how researchers should contact an organization upon discovering a vulnerability, and the steps they should take to ensure their testing is authorized.

II. Suggestions for testing sensitive, isolated election systems

Naturally, sensitive election systems like voting machines are likely to be disconnected from the internet and difficult to access for researchers. With that said, I would emphasize the importance of still establishing a coordinated disclosure policy for any vulnerability reports regardless of testing conditions. The role of a coordinated disclosure policy is to act as a contact channel for researchers who have discovered vulnerabilities in an organization's systems. The disclosure policy does not need to provide researchers a way to access these systems — should researchers notice a vulnerability by any means, provided they follow the rules laid out in the policy, they should be able to report it and be afforded safe harbor. This applies to security flaws in publicly accessible systems, such as election registration sites, and physical systems such as voting machines. Even without physical access to machines, researchers may discover, for instance, vulnerabilities in devices accidentally connected to the Internet or auxiliary Internet-accessible assets. Alternatively, a research may be able to legitimately buy a used voting machine

and test it, in which case they should have a way to report. Thus, it is crucial that a vulnerability be first instituted as a "see something, say something" for any and all flaws researchers come across.

Bug bounties provide the ability to provision access to trusted researchers for systems that would otherwise be inaccessible to researchers. For instance, all major bug bounty platforms have offerings that choose a select group of verified, trusted researchers and provision access to sensitive systems. To test a voting machine, for example, a machine vendor could provision VPN access into a voting machine's internal network and allow only the trusted researchers access. The Department of Defense follows this testing model: in addition to a public-facing Vulnerability Disclosure Policy, the DoD also operates a second, more sensitive bug bounty program that provisions access for trusted researchers to sensitive systems. I note that by no means is a private bug bounty program a substitute for a vulnerability disclosure policy. At minimum, every organization should maintain contact channels for any researcher, regardless of their background or identity, to report vulnerabilities. Bug bounty programs are an effective addition to incentive testing and provide additional access to trusted researchers.

III. Establishing trust with researchers

Trust is mutual when operating a coordinated vulnerability disclosure program. To first address the question of whether a CVD program can aid malicious actors, I stress that this is not the case and any restrictions on participation in a CVD program will only prevent good-faith researchers from productively disclosing vulnerabilities. A coordinated vulnerability disclosure policy must be open to all. It is understood that malicious actors are already targeting election systems, and a legitimate concern that they may try to conceal their actions by any means possible. However, a vulnerability disclosure policy does not aid attackers for two reasons. First, it is assumed all good-faith hackers will submit a vulnerability report upon discovering a vulnerability. If potentially malicious activity is noticed on a system, it can be deconflicted with vulnerability reports to determine if such activity was indeed malicious. Second, by defining rules of engagement, there is a clear line between testing authorized under a vulnerability disclosure policy and malicious testing. For instance, if an actor begins to exfiltrate or modify mass amounts of information, it is clear that this testing is malicious and not authorized under the

vulnerability disclosure policy. Indeed, as these systems are already targeted by malicious actors, it is unlikely that malicious activity would increase under a vulnerability disclosure policy. Instead, given vulnerability reports and testing from good-faith researchers, an organization can hope to better its security and avoid exploitation by adversaries.

At the same time, an organization must achieve trust with good-faith researchers in order to gain their assistance in testing systems. Due to the current legal environment around security research, researchers are already in a precarious situation when it comes to testing and disclosing vulnerabilities in systems. To the furthest extent possible, authorized testing under coordinated vulnerability disclosure policies should be encouraged. A coordinated vulnerability program must establish trust with good-faith researchers, primarily in outlining expectations through a safe harbor. Researchers understand that following rules of a vulnerability disclosure program is crucial, and extending safe harbor to abiding researchers has the effect of encouraging productive testing. Even small gestures such as timely and open communication play a large role in encouraging researcher participation and establishing trusted relationships.

IV. Public disclosure

A central component of coordinated vulnerability disclosure programs is expectations surrounding disclosure. Whereas without CVD, researchers may publish findings without sufficient coordination with the affected vendor, CVD affords a streamlined internal and external communication process. It is standard for organizations to set expectations for disclosure in the program policies, for instance to “keep vulnerability details confidential until they are fixed” or to “only disclose vulnerability details with the permission of the organization”. These rules are part of the program terms, and researchers understand that safe harbor provisions only apply if they abide by these ground rules. As a result, under CVD, both the vendor and researcher may agree on a disclosure timeline, with corresponding technical details. This provides more accurate and representative disclosure which signals to other researchers to engage in the organization’s vulnerability disclosure program.

V. Remediation

Of course, coordinated vulnerability disclosure is only effective when submitted vulnerabilities can be remediated in a timely manner. Election vendors should allocate staff for

processing and responding to vulnerability reports, and establish processes for patching and tracking vulnerability reports. It is common for vulnerabilities to resurface after being patched, or for there to exist bypasses to a patch. One standard way to prevent this is to ask the original researcher to re-test the vulnerability after a patch has been issued. Furthermore, vendors should investigate the root cause of all vulnerabilities, both to detect occurrences of the same vulnerability in other areas, and to integrate into the development process to prevent such vulnerabilities from occurring in the future. Additionally, the certification process must be improved to allow for quick patching of voting systems. As is, the certification process is timely and may interfere with the patching process. It should not be the case that a voting machine is decertified because a critical patch is issued. Regulators should establish exceptions to certification processes for vulnerability patches, either allowing a fast recertification process or for vendors to self-certify machines after issuing a patch.

VI. Next steps

I encourage all vendors to establish coordinated vulnerability disclosure programs, and hope that this response can aid in such processes. Should you have any questions or wish to discuss further, I can be reached at cablej@stanford.edu. On behalf of the security researcher community, I hope that through coordinated vulnerability disclosure we can collectively work to secure American elections before 2020.