Hearing on "Complex Cybersecurity Vulnerabilities: Lessons Learned from Spectre and Meltdown"

Written Testimony of Art Manion
Vulnerability Analysis Technical Manager, CERT/CC
Carnegie Mellon University Software Engineering Institute

Before the U.S. Senate Committee on Commerce, Science, and Transportation

July 11, 2018

Introduction

Chairman Thune and Ranking Member Nelson, thank you for the opportunity to appear before the Senate Committee on Commerce, Science, and Transportation to discuss Complex Cybersecurity Vulnerabilities, specifically some of the challenges and lessons from the Meltdown and Spectre disclosures. I am currently the Vulnerability Analysis Technical Manager at the CERT Coordination Center (CERT/CC), part of Carnegie Mellon University's Software Engineering Institute (SEI).[1] The SEI is a Department of Defense Federally Funded Research and Development Center (FFRDC). The SEI conducts research and development in software engineering, systems engineering, cybersecurity, and many other areas of computing, working to transition new and emerging innovations into government and industry. The SEI holds a unique role as the only FFRDC sponsored by the DoD that is also authorized to work with organizations outside of the DoD. We work with partners throughout the U.S. government, the private sector, and academia.

Much of the vulnerability analysis work at the CERT/CC over the past 30 years has focused on Coordinated Vulnerability Disclosure (CVD). This is the practice of reporting newly discovered vulnerabilities to vendors (and/or third-party coordinators, like ourselves), working cooperatively to develop fixes and mitigations, and eventually publicly disclosing information for defensive purposes. The results of many coordinated disclosure cases we work on are published as Vulnerability Notes.[2] We work closely with the Department of Homeland Security (DHS) National Cybersecurity and Communications Integration Center (NCCIC) elements commonly known as US-CERT and ICS-CERT as well as other stakeholder communities including security researchers, vendors and other software development organizations, and more recently, policy makers and regulatory agencies.

In August 2017 my team published *The CERT Guide to Coordinated Vulnerability Disclosure*,[3] capturing decades of experience, observation, and advice. This testimony draws heavily on the *Guide* and the collective experience of my current team and past members.

[1] https://www.sei.cmu.edu/
[2] https://www.kb.cert.org/vuls
[3] https://resources.sei.cmu.edu/library/asset-view.cfm?assetid=503330

The CERT/CC is a founding member of the Forum of Incident Response and Security Teams (FIRST), and I co-chair two special interest groups within FIRST that deal with CVD. The Vulnerability Coordination SIG, collaborating with a National Telecommunications and Information Administration (NTIA) multistakeholder process,[4] published guidelines for multiparty CVD in June 2017.[5] My team provides advice to transportation (including the Department of Transportation) and medical device (including the Food and Drug Administration) sectors. We participate in other policy and community efforts to improve and advocate CVD processes, and we directly assist departments and agencies, helping to design and support the DoD Vulnerability Disclosure Program.[6] I also work in the International Standards Organization (ISO) where I am co-editor of ISO 29147 *Vulnerability disclosure*[7] and 30111 *Vulnerability handling processes*.[8]

Coordinated Vulnerability Disclosure (CVD)

We all depend on software and software-based systems. The devices we use to communicate and coordinate our lives, transport us from place to place, and keep us healthy include computers, network connections, and software. As a result, society has increased its dependence on software-based products and services that communicate both to each other and to the world at large.

One drawback: our modern and connected products and services have vulnerabilities—weaknesses that can compromise the security of the system in unexpected and undesirable ways. Vulnerabilities leave our devices and systems susceptible to attacks. Smart phones, ATMs, MRI machines, security cameras, cars, airplanes, and the like have become network-enabled software-dependent systems, making it nearly impossible to avoid participating in the world without the potential to be affected by cybersecurity vulnerabilities.

Essentially unavoidable, vulnerabilities have numerous origins. Implementation defects, unexpected interactions between systems, configuration or design decisions, and other factors all contribute to what is effectively an unlimited supply.[9] In order to maintain assurance in the systems and devices we use daily, we need clear public policy and socio-technical norms encouraging the discovery of vulnerabilities, notification of their existence, and cooperative defense in the form of repair or mitigation. Otherwise, adversaries can take advantage of vulnerabilities to achieve goals at odds with the creators and users of the systems we depend on.

Notifying the public that a problem exists without simultaneously providing defense leads to increased adversarial advantage. Because there is rarely one optimal formula for minimizing

---

[4] https://www.ntia.doc.gov/other-publication/2016/multistakeholder-process-cybersecurity-vulnerabilities
[5] https://www.first.org/global/sigs/vulnerability-coordination/multiparty/
[6] https://hackerone.com/deptofdefense
[7] ISO/IEC 29147 https://www.iso.org/standard/45170.html
[8] ISO/IEC 30111 https://www.iso.org/standard/53231.html
[9] Risk Based Security recorded "...over 20,000 vulnerabilities disclosed in 2017." The NIST National Vulnerability Database (NVD), based on the Common Vulnerabilities and Exposures (CVE) project, reports 14,650.

risk and harm—short of avoiding the introduction of vulnerabilities in the first place—the current best practice is a process called Coordinated Vulnerability Disclosure (CVD).

CVD is the process of gathering information from security researchers, coordinating the sharing of that information to vendors and other relevant parties, and disclosing the existence of software vulnerabilities along with updates or mitigations to various stakeholders—including the public. The CVD process concludes when updates and mitigations have been widely deployed.

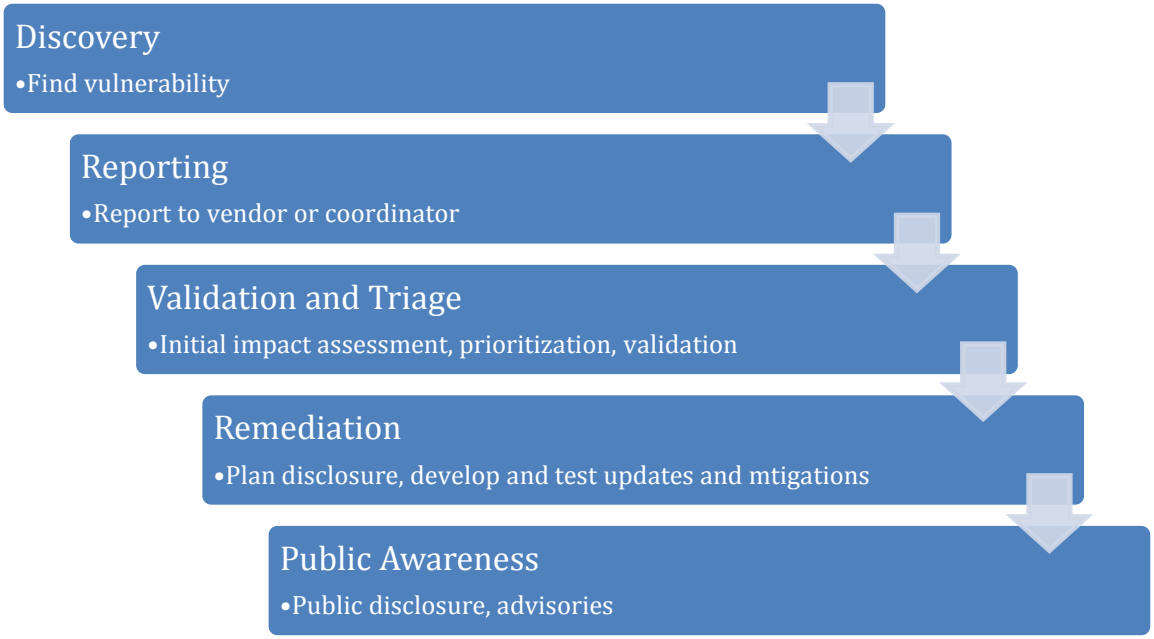

Figure 1: CVD phases

Figure 1 outlines a generally accepted set of basic CVD phases. As noted in the *Guide*, the more difficult questions are: "How much information should be released? To whom? And when?"

Bilateral CVD—between one researcher and one vendor—is largely a solved problem. That is to say, there exist established CVD processes that both parties can follow to a generally agreeable and optimal outcome.

Multiparty or multivendor CVD, as illustrated by Meltdown and Spectre, remains unsolved. The CERT/CC focuses our efforts on multiparty CVD, both directly handling cases and researching ways to make lasting improvements. As described in the *Guide*, CVD is a wicked problem,[10] and multiparty CVD even more so.

[10] H. W. Rittel and M. M. Webber, "Dilemmas in a General Theory of Planning," Policy Sciences, vol. 4, no. 1973, pp. 155-169, June 1973.

Meltdown and Spectre

Meltdown and Spectre are the widely used names for the first three instances of a class of vulnerabilities that arise from speculative execution[11] and shared caches,[12] features designed into modern CPU hardware for improved performance. These side-channel vulnerabilities allow attackers to infer the contents of memory without having direct access to the memory. Meltdown and Spectre were initially reported in June 2017 and three variants were publicly disclosed on January 3, 2018. Since then, three additional variants have been published, and further public disclosures are expected. It is interesting to note that these security issues were previously discussed in 1995.[13]

Meltdown and Spectre allow attackers to read memory that they shouldn't have access to, memory that could contain users' passwords, trade secrets, encryption keys, or the contents of private documents. Users of shared cloud infrastructure are particularly at risk. Attacks can also be performed against web browsers that visit malicious sites.

Such access to another's data is remarkable. Modern CPU hardware and operating system software separate running programs from each other and from the operating system, for stability and security reasons. This separation was meant to ensure that one user of the computer cannot read memory in use by another user or the privileged operating system, which could contain sensitive or secret information.

Because Meltdown and Spectre are intrinsic to CPU hardware, they are different from much more common software vulnerabilities. Consequently, while some of the Meltdown and Spectre variants can be mitigated with operating system and CPU microcode updates, newly designed CPU hardware will be required to fully resolve the majority of the vulnerabilities.[14]

Challenges and Lessons Learned

Overall, the vendor-led CVD process followed for Meltdown and Spectre was reasonably successful. Major vendors (including competitors Intel, AMD, and Arm) cooperated on security, and major software and service providers applied updates that protected many users en masse. The vendors involved followed current CVD practices, arguably tuned too far in favor of attempting to prevent premature public disclosure. This tuning introduced challenges, particularly for those tasked with defending critical infrastructure and public safety. Due to a number of factors, including the vendor-led CVD process and the novelty and complexity of the technology involved, Meltdown and Spectre garnered public attention that arguably exceeded the "actual" risk of the vulnerabilities.

[11] occurs when the CPU, which would otherwise be sitting idle, instead makes an informed guess as to what instructions a running program will take next

[12] areas of memory with quick access

[13] https://pdfs.semanticscholar.org/2209/42809262c17b6631c0f6536c91aaf7756857.pdf

[14] https://energycommerce.house.gov/wp-content/uploads/2018/02/Intel-Corp-response-HEC-FINAL.pdf

The following are a set of challenges and lessons brought to light by the Meltdown and Spectre disclosures.

CVD should follow the supply chain

At its most effective, CVD follows the supply chain affected by the vulnerability. Many products today are not developed by a single vendor. Instead, they are assembled from components sourced from other vendors. For example, software libraries are often licensed for inclusion into other products. When a vulnerability is discovered in a library component, it is very likely that not only does the originating vendor of the library component need to take action, but all the downstream vendors whose products use it need to take action as well. Complex supply chains can increase confusion regarding who is responsible for coordinating, communicating, and ultimately fixing vulnerabilities, leading to delays and systems exposed to unnecessary risk. Because of the underlying nature of the vulnerabilities, Meltdown and Spectre exacerbated these concerns.

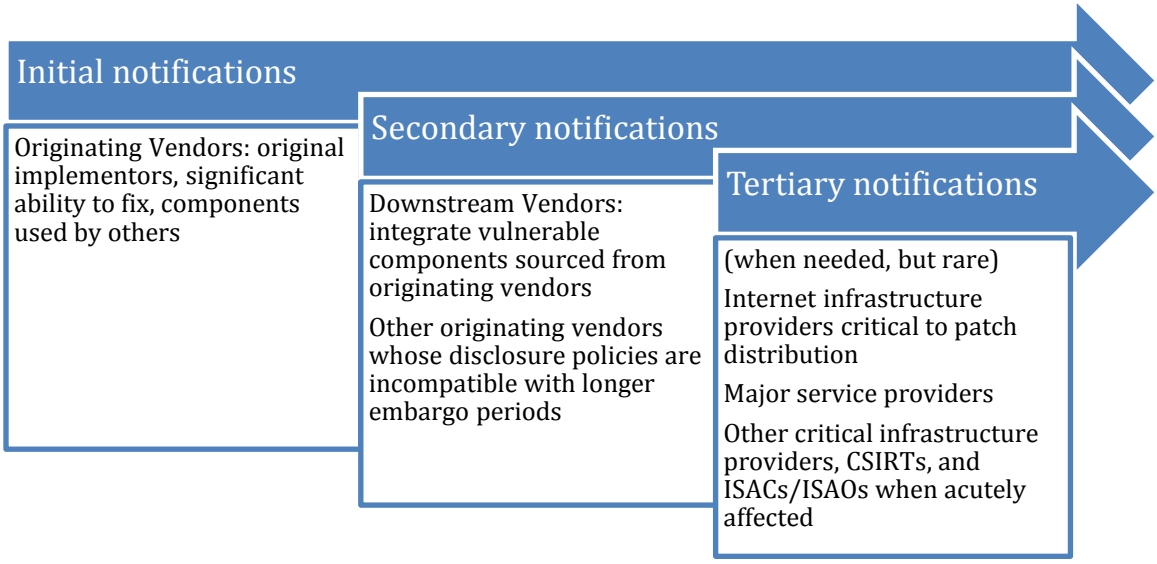

Figure 2: Notional multiparty CVD process

Initial notifications

When considering private notification and embargo, there is general agreement that those who have the ability to make changes that remove or substantially mitigate vulnerabilities need to be informed. This usually means vendors who produce original implementations. For Meltdown and Spectre, these vendors were CPU manufacturers: Intel, AMD, and Arm.

Secondary notifications

Operating system and virtualization software is tightly bound to CPU hardware, and mitigations for Meltdown and Spectre generally require both CPU microcode and software updates. Thus,

another set of vendors with the ability to fix included Microsoft, Google, and Apple. For vendors who advertise short embargo periods, it is possible to delay notification until shortly before public disclosure.

Tertiary notifications

Depending on the nature, scope, and potential impact of the vulnerability, it may make sense to notify major service providers, in this case, the cloud computing elements of Amazon, Microsoft, and Google. Consideration should also be given to notifying critical infrastructure protection stakeholders, at least to reduce the harm associated with surprise disclosure.

Conceptually, it can be useful to think of the supply chain as horizontal or vertical. A horizontal supply chain implies that many vendors need to independently make changes to their products in order to fix a vulnerability. A vertical supply chain implies that one vendor might originate the fix, but many other vendors may need to update their products after the original fix is available. In these terms, Meltdown and Spectre exhibit aspects of both horizontal and vertical supply chains, making the coordination process even more complex.

Fairness in CVD is desirable but difficult to achieve

From a coordinator's perspective, it can be difficult to be fair when coordinating a multiparty CVD case, because it's almost inevitable to either miss some downstream vendor or wind up with one or more vendors ready to release while everyone is waiting for the other vendors to catch up. The CERT/CC's practice is to notify a wider selection of vendors and other stakeholders than those included in the Meltdown and Spectre CVD process, acknowledging that this increases the risk of premature public disclosure.

Broader CVD cases require shorter embargo periods

The CVD process for Meltdown and Spectre was complicated by the nature of the supply chain and the premature public disclosure which caught many by surprise. Our experience shows that problems can arise when the multiple parties involved in CVD function at different operational tempos. In both the vertical and horizontal supply chain cases discussed above, synchronized timing of disclosure to the public can be difficult to coordinate. The originating vendor(s) will usually want to release a patch announcement to the public as soon as it is ready. This can, however, put users of downstream products at increased risk. As a result, coordinators sometimes find it necessary to make the difficult choice to withhold notification from a vendor in a complicated multiparty disclosure case if that vendor's disclosure policy is incompatible with the embargo or otherwise cannot be trusted to cooperate with the coordination effort. This may have been a factor for Meltdown and Spectre and was illustrated by the CVD process for the KRACK Wi-Fi vulnerabilities.[15]

Vendors are not the only stakeholders with a role to play prior to public disclosure

In situations where a vulnerability has the potential for major impact to critical infrastructure, it may be necessary to coordinate not only with vendors to fix the vulnerable products, but also

[15] https://www.krackattacks.com/#openbsd

with major deployers—those responsible for applying updates and other mitigations that affect large populations of users. One important concern in these cases is to ensure that internet and other critical infrastructure remains available so that deployers and other network defenders can acquire and deploy the necessary information and patches. Another important concern is that critical infrastructure protection stakeholders are prepared to provide accurate and actionable information before public disclosure.

Luckily this scenario is rare, but vulnerabilities like Meltdown and Spectre, or those that affect basic internet services such as the domain name system (DNS), can affect a large number of vendors. In these cases, the involvement of a coordinator such as the CERT/CC can often help contact and disseminate information to vendors, service providers, and other key stakeholders. Note that the CERT/CC was not engaged in the coordination of Meltdown and Spectre prior to their public disclosure.

Rushed solutions can increase risk

The Meltdown and Spectre disclosures generated a lot of public attention. They were the results of cutting-edge research from multiple sources; a lengthy embargo period (roughly 6 months) and closely held CVD process among major vendors; and in the end, the public disclosure happened one week earlier than planned. Many organizations were surprised by the public disclosure and spent considerable effort trying to understand the nature of the vulnerabilities and their impact.

Due to the fundamental technical nature of the vulnerabilities, the complexity of CPU and operating system interaction, and in some cases the lack of lead time, many of the updates and mitigations caused significant negative side effects. A partial list follows.

- Intel microcode updates caused instability.[16]
- Initial Meltdown updates for Microsoft Windows 7 and Server 2008 mistakenly allowed any user to read kernel memory and gain complete control of a computer.[17]
- Microsoft Windows updates caused some AMD systems not to boot.[18]
- Architectural changes caused some antivirus software running on Microsoft Windows to not work. The changes also had serious implications for receiving future security updates.[19]
- Lenovo systems running SUSE can become inoperable.[20]
- Pulse VPN client on Microsoft Windows would not connect.[21]

[16] https://newsroom.intel.com/news/intel-security-issue-update-addressing-reboot-issues/
[17] https://www.kb.cert.org/vuls/id/277400
[18] https://support.microsoft.com/en-us/help/4056892/windows-10-update-kb4056892
[19] https://support.microsoft.com/en-us/help/4072699/windows-security-updates-and-antivirus-software
[20] https://support.lenovo.com/us/en/solutions/len-18282
[21] https://kb.pulsesecure.net/articles/Pulse_Secure_Article/KB43600

- Dell systems experienced unpredictable behavior.[22]

Independent of the unintentional side effects, the updates decrease performance, because the CPU and operating system have to spend more time clearing out the remnants of speculative execution left in the cache. While an individual user may not notice, busy server systems are significantly impacted by the performance decrease.[23] This may require the purchase of additional server capacity to maintain performance equivalent to pre-update levels.

Given the side effects and performance penalties, users should carefully consider the need to install Meltdown and Spectre updates. The vulnerabilities pose the greatest risk to systems that allow multiple users to run code, for example, cloud-based shared or multi-tenant hosting providers. Individual users may not substantially improve their security by installing updates. Systems that require high availability and reliability, such as industrial control and other safety critical systems, should not install updates or make other changes without significant testing.

Surprise leads to misplaced effort and opportunity cost

As with most situations in which multiple parties are engaged in a potentially stressful and contentious negotiation, surprise in CVD tends to increase the risk of a negative outcome. For technically complex vulnerabilities like Meltdown and Spectre, there is a need for stakeholders to understand the problem before it is possible to make good decisions about the appropriate response. Because so many vendors, deployers, and other stakeholders were caught off guard with the public disclosure of the Meltdown and Spectre vulnerabilities, much attention was diverted from potentially more pressing and immediate cybersecurity issues.

CVD Improvements

As previously stated, the Meltdown and Spectre CVD process was reasonably successful. Without any changes to existing practices, the process could have been tuned to include more vendors and to notify more stakeholder organizations before public disclosure. This recommendation comes with the understanding that the chance of premature disclosure increases with the number of people and organizations brought into the circle. I am aware of zero premature disclosures or other leaks caused by the NCCIC, US-CERT, or ICS-CERT. I am aware of a few leaks caused by organizations privately notified by the CERT/CC, but over 30 years and tens of thousands of CVD cases, I am comfortable with the balance we have chosen. Public information about Meltdown and Spectre started appearing in November 2017. In my experience, a case of this magnitude was unlikely to survive a lengthy embargo period.

Meltdown and Spectre set an inflection point in the history of CVD and internet security. The researchers, and more importantly the coordinating vendors, could have recognized the need to at least reduce surprise by informing the U.S. government (and possibly other governments)

[22] https://www.dell.com/support/article/us/en/04/sln308588/microprocessor-side-channel-vulnerabilities-cve-2017-5715-cve-2017-5753-cve-2017-5754-impact-on-dell-emc-products-dell-enterprise-servers-storage-and-networking-?lang=en

[23] https://cloudblogs.microsoft.com/microsoftsecure/2018/01/09/understanding-the-performance-impact-of-spectre-and-meltdown-mitigations-on-windows-systems/

sooner. Such a decision is already accounted for in existing CVD guidance; implementing it is a matter of tuning known parameters.

Aside from the increased risk of premature disclosure, vendors have cited the need to act fairly as another reason not to notify governments in advance. Governments can be both customers and regulators, wielding purchasing and legal power. Which government(s) should a CVD process include before public disclosure?

There are options. Vendors could choose to inform governments based on confidence that the government will maintain the embargo and only use the information for defensive purposes. Microsoft, for example, offers a Government Security Program[24] to qualified governments that includes advanced notice of vulnerability disclosures. It is not clear, however, how conflicting sharing agreements are resolved. Also, the program notifies governments five days before public disclosure, Meltdown and Spectre leaked six days early.

Despite, or because of, our long history of handling multiparty CVD cases, we do not believe that a single global coordinator designed to handle nearly every multiparty case will scale. The CERT/CC is part of a loosely-affiliated multinational network of coordinators, with whom we share common CVD practices. This network sometimes shares vulnerability information in order to reach a wider global selection of vendors. These coordinators are related to their respective national governments: Japan, Finland, and the Netherlands.[25] One solution to scalable, multiparty CVD may be a more formal network of coordinators. Another option could be a more formal collection of national government computer security incident response teams (CSIRTs)[26] that agree to follow suitable embargo and information sharing restrictions. However, as mentioned above, government involvement in CVD may be a concern for vendors and inhibit their willingness to participate. Other ideas include non-governmental organizations (NGOs) or commercial businesses that are sufficiently independent of any one government.

Conclusion

CVD is a process of coordinating human behaviors. Success at multiparty Coordinated Vulnerability Disclosure has more to do with understanding human communication and organization phenomena than with the technical details of the vulnerability. The hard parts are nearly always about coordinating the behavior of individuals and organizations with diverse values, motives, constraints, beliefs, feelings, and available energy and time. Technical vulnerability details may dictate the “what” of the response, but to a large degree, human organizational and social behaviors decide the “how.” Optimal CVD operation requires carefully balancing “How much information should be released? To whom? And when?”

Thank you again for the opportunity to appear today before the committee.

[24] https://enterprise.microsoft.com/en-us/trends/government-security-program-available-to-qualified-governments/
[25] JPCERT/CC, NCSC-FI, and NCSC-NL
[26] https://www.sei.cmu.edu/education-outreach/computer-security-incident-response-teams/index.cfm