OPEN TECHNOLOGY INSTITUTE

Cybersecurity Research Should Not Be A Crime

Why We Need Clear, Permanent CFAA and DMCA Exemptions

Nat Meysenburg

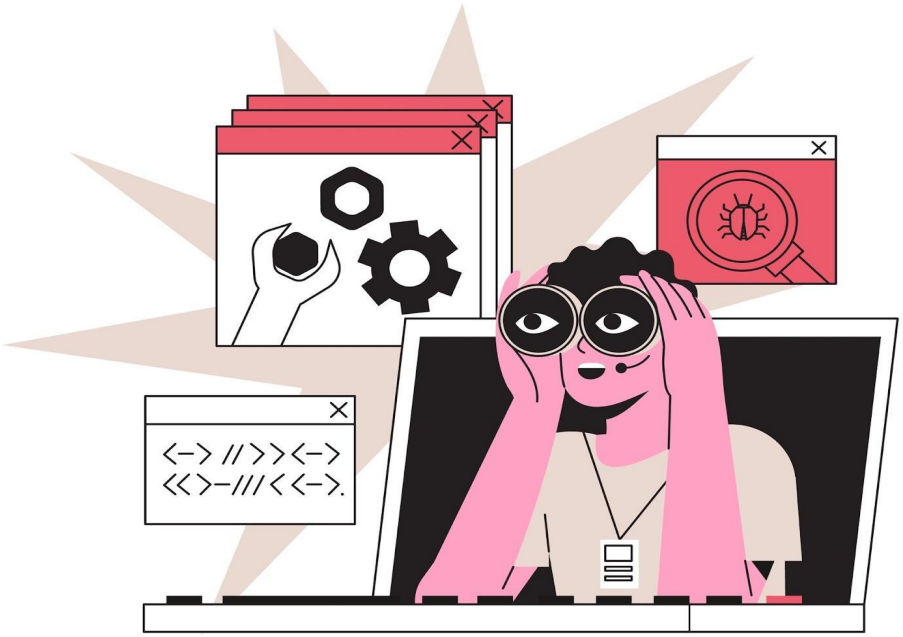

Introduction

In early 2021, the fate of the Bidens' Peloton became a brief sub-plot to the Presidential transition. The question of whether the soon-to-be First Family's exercise bike could introduce a cybersecurity risk onto the very secure White House grounds was mostly seen as a "presidential bubble" story in which security experts speculated about ways that this incredibly popular internet-connected stationary bike could be vulnerable to hackers, and how it might be hardened. Among the concerns were whether the bike's cameras and microphones could be used for spying, if it could leak other sensitive information, or if it could be used as a stepping stone for gaining access to other White House systems. As it turns out, those concerns were not misplaced. On the day President Biden took his oath of office, a team of British security researchers privately reached out to Peloton and disclosed a number of serious information leaks they had found on Peloton's servers.

The President's unique digital security needs are evaluated by top-notch cybersecurity professionals, who can take the necessary steps to mitigate any security and privacy risks. These experts are tasked with figuring out if the First Family's digital devices are safe for use in secure locations, and even modifying them to make them more secure. They also may decide that certain devices are just too insecure to introduce into such a high-risk environment. But lurking

beneath the question of whether a Peloton is secure enough for the White House is the question of whether it is secure enough for everyone else's house. Most American consumers don't have access to a team of security experts who could answer that question for them, and absent that team

of experts, whether or not a product is digitally vulnerable is a question left to product manufacturers.

A piece of tech is considered "vulnerable" when a part of the code that runs it can be exploited by an attacker. A software vulnerability is like a hole in an otherwise sturdy fence, it may be hard to find, but it has the potential to let out what should be inside (or let in what should be outside). There can also be vulnerabilities in a product's hardware, which similarly provide entrance points for technological attacks.

Uncovering a "zero-day"—the term for an undiscovered vulnerability—requires an understanding of common coding mistakes, previous vulnerabilities, and some knowledge of how to look for all of that in code. Discovering a new vulnerability is valuable for the malicious actor who finds it, but not everyone looking for vulnerabilities has terrible intentions. Often labeled as "white hat hackers," there is a community of security experts working to find those holes and responsibly disclose that information to manufacturers so vulnerabilities can be patched before they are maliciously exploited. This allows vendors to make their products more secure without spending the time and money necessary to conduct their own testing. Encouragingly, there has also been a growing trend among companies to implement processes for responsible vulnerability disclosure, and in some cases even reward researchers for the vulnerabilities they find.

In the case of Peloton, vulnerabilities in the app's code left the data of millions of exercise class participants exposed. While the vulnerability was present, it could have allowed nefarious actors to gather those users' data and use it for their own ends. Researchers found this leak, and disclosed that information to Peloton through the company's disclosure program. When those researchers released their report about the Peloton vulnerability in May of 2021, the company had already fixed the vulnerabilities in their system. Were it not for the responsible disclosures made by the researchers at Pen Test Partners, the private data of Peloton riders might still be broadly available.

Editorial disclosure: This brief discusses policies by Google and Microsoft (including LinkedIn, whose co-founder is on New America's Board), all of which are funders of work at New America but did not contribute funds directly to the research or writing of this piece. View our full list of donors at ***www.newamerica.org/our-funding****.*

Criminalization of Security Research

Despite some recent changes in norms around vulnerability disclosure, researchers have not always found a receptive audience with the companies whose products they test. Good-faith researchers work in fear of their efforts being met with legal threats, lawsuits, and even criminal penalties.

This fear is largely rooted in how two pieces of federal law—the Computer Fraud and Abuse Act of 1986 (CFAA), and the Digital Millennium Copyright Act (DMCA)—have been used to go after researchers. Both laws were written to address new forms of crime enabled by new uses of technology, but they were written for a fundamentally different digital world. Each law includes penalties for conduct that includes methods vital to security testing, but neither carves out general or permanent exemptions that clearly distinguish between researchers trying to help and the sorts of intentional malicious behavior the laws seek to prohibit.

The legal uncertainty faced by researchers is something the Open Technology Institute has experienced firsthand. Over the last couple of years, OTI has spent a lot of time testing connected consumer products using the Digital Standard—a framework for evaluating the privacy and security of internet-connected consumer products and software. The Standard was developed by Consumer Reports and a coalition of civil society organizations, and consists of a group of tests measuring how well a product adheres to a set of best practices for connected hardware and software. Some of the tests evaluate technical practices, like whether a product uses encryption or strong authentication processes, while others evaluate a company's policies on privacy issues like data collection and retention. In the course of this work, OTI conducted

the same kinds of investigations that many security researchers do when searching for software vulnerabilities, including decompiling and examining mobile app code, monitoring network traffic to and from connected devices and/or mobile apps, and providing form values intended to break an app. After considering OTI's own legal risk, decisions to test products and publish specific findings were limited by concerns about potential penalties under the CFAA and DMCA, particularly manyany products have disclaimers in their Terms of Service saying that they will take legal action against, or refer to law enforcement for prosecution, individuals who violate those terms.

Background: the Landscape of Digital Security

In the last two years, a steady stream of hacks and ransomware attacks has kept cybersecurity in the headlines. Affecting everything from hospitals and local school systems to critical infrastructure to the federal government itself, these attacks have brought renewed attention to the current state of digital security. They have also made clear how broadly vulnerable to cyber attack the United States may be.

Shared Vulnerabilities

A piece of vulnerable software could be installed on millions of devices. Developers typically include "code libraries" when building software to handle some functionality needed for their software that may be common across many other developers' software projects. This can include everything from how buttons look in a mobile app to how a smart product connects to a WiFi network. If a library is popular, it could be a component in thousands of products. The technical term for this library's relationship to those products is "upstream." When that upstream library has a weakness, all of the "downstream" systems running that code share it. These are known as shared, or common, vulnerabilities.

Researchers and developers sharing information about common vulnerabilities when they are uncovered is standard practice, and product makers' distribution of updates and patches is a primary defense against cyber attack. However, developers do not always reliably update their upstream code, and users do not always reliably install available patches. Understanding how many shared vulnerabilities exist in the world, it is common practice for attackers to use tools that scan the internet looking for evidence of a system running software known to be vulnerable.

Proliferation of Vulnerable Hardware

When assessing cyber threat vectors, it is also important to remember common vulnerabilities can exist in firmware (the code that controls digital hardware) or even in the way hardware is designed.

For example, Intel has faced a series of vulnerabilities in the security features available in some of its processor chips. Computer owners can generally find out if their system has an Intel processor, and if an Intel vulnerability is publicized could at least know that they are affected. As one of the largest chip makers in the world, Intel has some ability to make updates for its hardware available, and many customers know to look for such firmware updates, but not all of them do. This means that unpatched chips are still out there. The task of identifying and fixing vulnerabilities in Intel chips is huge, but it pales in comparison to the situation for most Internet of Things (IoT) devices.

Unlike processors from major vendors, the chips and processors used in many IoT products are hard to trace. There are many different small chip makers, and most small chips are not well labeled. Even if they are labeled, information about the chip can be hard to find. Furthermore, due to factors like cost and availability, the chips used in a product can change between manufacturing runs. Given the frequency of these changes, it is sometimes hard to figure out what hardware is being used on any given device, even between versions of the same model. This reality of tech manufacturing means that a wide variety of products may end up using similarly vulnerable supplies.

IoT devices tend to be orders of magnitude less powerful than multi-function computers like laptops, or even smartphones. Because of technical resource constraints, it can be hard—and in some cases impossible—to update code without special equipment. But even in the cases where updating code is technically feasible, a manufacturer's product life cycle may not include producing and distributing code updates for the entire

usable life of the product, even if there are known vulnerabilities.

Responsible Disclosure Programs

Enthusiasts with technical talent and a willingness to help have long played an important role in finding bugs in software. The term "bug bounty" was coined in 1995, by Netscape Communications Corporation for a program they set up to reward people who reported bugs "with various prizes depending on the bug class." Even in this early incarnation of a bug bounty, "users reporting significant security bugs" was already seen as the most valuable contribution, with that being the only category where one could collect a cash prize. All of the other prizes were Netscape Navigator merchandise. Despite these deep roots, an emphasis on responsible vulnerability disclosure has only come into its own as an industry standard over the last decade.

Many large companies such as Google and Microsoft have set up disclosure programs, and—as seen in the case of Peloton—smaller companies are doing the same.

Providing an ability for researchers to report vulnerabilities takes more than setting up a dedicated email address. Incoming reports must be reviewed and validated, the vulnerability must be reproduced in order to test if it has been fixed, and then a patch must be devised, written, and deployed. In addition, vulnerabilities may bring researchers in contact with data they should not have access to, which depending on the type of data may carry its own legal risk. Setting up a vulnerability disclosure program requires a willingness to work through that process. Given the relative complexity of these processes, it is important to create incentives so that smaller players can adopt those practices, and to spread vulnerability reporting processes as a standard practice in technology development.

In 2017, recognizing that "an increasing number of organizations in the public and private sectors are adopting vulnerability disclosure programs to improve their ability to detect security issues," the Cybersecurity Unit at the U.S. Department of Justice issued a framework "to assist organizations interested in instituting a formal vulnerability disclosure program." While this framework sent an encouraging signal that the federal government recognizes the importance of responsible vulnerability disclosure, it also illustrates the challenges created by current law. The framework did not "dictate the form of or objectives for vulnerability disclosure programs" but rather it outlined a design aimed at "reducing the likelihood that such described activities will result in a civil or criminal violation of law under the Computer Fraud and Abuse Act."

The Computer Fraud and Abuse Act

Three years before Tim Berners Lee invented the world wide web, Congress passed the Computer Fraud and Abuse Act of 1986. The law was written for a very different internet by a Congress whose frame of reference was a fairly specific kind of computer crime. The House Judiciary Committee report accompanying the Computer Trespass Act of 1984 (the precursor to the CFAA) called the 1983 movie War Games—in which a Seattle teen who is hunting for video games breaks into a top secret computer, and nearly starts a thermonuclear war—a "realistic representation of the... access capabilities of the personal computer." With the CFAA, Congress was trying to create penalties for breaking into computers to view, change, or destroy data.

The CFAA makes it illegal to access any computer "without authorization." In addition to possible felony punishment, the CFAA also allows companies to sue in civil court alleging CFAA violations, even if law enforcement does not pursue charges.

Unfortunately, even with both criminal and civil penalties at stake, the CFAA does not define what accessing a computer "without authorization" means, or what it means to "exceed" that authorization. For most of the last decade there have been different and conflicting interpretations of these terms in different federal districts. The 2021 Supreme Court decision in *Van Buren v. United States* provided some guidance, requiring that technical definitions be used when words such as "access" carry technical meaning, and a technical "gates up or down" standard for defining exceeding access.

This prolonged lack of clarity has allowed the CFAA to be used as a cudgel to stop independent security researchers from evaluating products, limit competition between companies, and forbid other normal internet behavior.

Cease-and-desist letters to researchers citing possible CFAA violations have become an all-too-common tool for intimidation. It is impossible to know the exact number of cease-and-desist letters that chilled behavior without leading to litigation. In cases that were litigated, the CFAA claims revealed in some of those letters rested on an argument that "authorize" in the CFAA was not meant technically, and that documents like Terms of Service could set the boundaries of what "authorized access" meant. This definition requires no technical limits to be placed on that access. In other words, simply saying that information, or types of usage, are off limits counts as removing authorization under the CFAA.

This is what happened in the case of *hiQ Labs, Inc v. LinkedIn Corp*. As part of its business, hiQ conducts research by gathering and analyzing data about how workers move from job to job from LinkedIn and other websites. LinkedIn sent a cease-and-desist letter in an attempt to stop hiQ from using automated tools to gather data from its site—a technique known as "scraping." LinkedIn claimed that the cease-and-desist letter itself constituted a revocation of authorization to access the data, and that hiQ's further scraping of LinkedIn after receipt of the letter was a violation of not only the CFAA but also the DMCA. In response, hiQ sued LinkedIn to challenge the claim that scraping public data constituted a CFAA violation. Ultimately the Ninth Circuit agreed with hiQ that exceeding access to a computer system only happens "when a person circumvents a computer's generally applicable rules regarding access permissions, such as username and password requirements, to gain access." The Court added that, in the case of scraping, "when a computer network generally permits public access to its data, a user's accessing that publicly available data will not constitute access without authorization under the CFAA."

The Ninth Circuit's ruling in the *hiQLabs* case held that exceeding access under the CFAA required the circumvention of technical barriers. The interpretation was at odds with an Eleventh Circuit ruling in *United States v. Roberto Rodriguez*, which found that "the plain language of the Act" meant "Rodriguez exceeded his authorized access and violated the Act when he obtained personal information for a nonbusiness reason." *Van Buren v. United States* sought to settle these diverging interpretations.

At issue in the *Van Buren* case was whether a former police officer caught taking money to run license plate numbers through a law enforcement database was correctly convicted for violating the CFAA. He had not broken any technical safeguards on the database, nor was he using stolen logins to get access to this information. Instead, he was misusing access to the database that he already had as a police officer, and doing so in violation of the use policy that he agreed to as part of his job. The main question in this case was whether the violation of a use policy alone constitutes exceeding access to a computer system. Using such an interpretation of the CFAA would mean that everything from browsing the web at work against employee policies to violating a website's terms of service by copying or using information from the site when it is forbidden could possibly invite felony prosecution.

In the majority opinion overturning the lower court ruling, Justice Amy Coney Barrett wrote that "an individual 'exceeds authorized access' when he accesses a computer with authorization but then obtains information located in particular areas of the computer—such as files, folders, or databases—that are off limits to him." The opinion also advises courts to "take note of terms that carry 'technical meaning[s],'" and goes on to say that access "is one such term, long carrying a 'well established' meaning in the 'computational sense.'" The Court later introduced a "gates-up-or-down" approach to whether access has been exceeded. Either someone is allowed to access data, or they are not. If there is not a technical "gate"—like password protection—that needs to be circumvented or broken to access data, this behavior is not a crime under the CFAA.

While clarity around "access" was welcomed by advocates and security researchers, the Court did not address concerns raised in Amicus filings seeking clear protections for research under the CFAA. So while there is more certainty around simple rule violation, it is still unclear if a researcher accessing data to test for security vulnerabilities in good faith, and without the intention to modify or destroy, is allowed. Researchers are still faced with questions about legal risk for conducting their work.

The Digital Millennium Copyright Act

The Digital Millennium Copyright Act (DMCA) was passed in 1998 as the United States sought to codify obligations

agreed to under treaties of the World Intellectual Property Organization. The treaties aimed to extend copyright protection to digital formats such as DVDs and mp3s, which could be more easily copied. But despite the DMCA's benefits for copyright holders, its rules both limit the freedom of people who purchase copyrighted materials and risk making certain types of security research a violation of copyright law.

Section 1201 of the DMCA prohibits the circumvention of "a technological measure that effectively controls access to a work." Also called a "technical protection measure" or TPM, these technologies can take many forms, for instance Digital Rights Management (DRM). DRM is a way for copyright holders to limit how a digital file can be used. For example, if you purchase an ebook, but are prevented from copying it from your ebook reader to your phone, that is likely DRM at work. DRM can also be used to limit the way digitally enabled devices work, like a coffee maker built to only work with coffee pods approved by its manufacturer. In addition to its anti-circumvention prohibitions on, Section 1201 also prohibits distributing tools to break TPMs, which is commonly referred to as the "anti-trafficking" provision.

On its face, Section 1201's prohibitions against breaking TPMs seem focused on protecting content from digital piracy (like copy protection on DVDs), but the DMCA covers other copyrighted digital works including computer code. Section 1201's definition of what it means to "circumvent a technological measure" is quite broad, including anyone who might "avoid, bypass, remove, deactivate, or impair a technological measure, without the authority of the copyright owner." This, by default, places most code (and products that use it) into an area of legal uncertainty for researchers.

Unlike the CFAA, the DMCA does recognize that there need to be reasonable exemptions to Section 1201's anti-circumvention prohibitions. Section 1201 contains several permanent exemption categories, including for security testing and encryption research. There are requirements to qualify for these exemptions, including a reliance on the CFAA's vague definition for exceeding authorized access. Researchers are also required to prove that any code they used in their research is not "primarily designed or produced for the purpose of circumventing a technological measure," breaking the anti-trafficking provision. In addition to the permanent exemptions, the law authorizes the Register of Copyrights to recommend (and the Librarian of Congress to publish) temporary exemptions to 1201's anti-circumvention ban. These three-year exemptions are decided in response to a public request and comment process conducted every three years by the Copyright Office. The most recent exemptions were released in late October 2021.

In contrast to the CFAA, this triennial process has allowed for Section 1201's broad restrictions to be somewhat more adaptable to changing technology, as well as the concerns of researchers. Over the last several rounds of this rulemaking process, advocates were able to gain very important—if temporary—tech-related exemptions, such as a 2015 exemption for the diagnosis and repair of "motorized land vehicles." The 2015 ruling also included an exemption for "good-faith security research." This temporary "good-faith" exemption is a case study in how despite its flexibility, there are gaps in the exemption process that put researchers at risk of violating the DMCA.

The 2015 research exemption covered a very limited set of devices (medical devices for implant, land vehicles, and "machines primarily designed for use by individual consumers, including voting machines"), as well as a list of other unclearly written limitations and requirements. Those ambiguities could cause "researchers to avoid publicly beneficial research activities."

Advocates succeeded in addressing some of these limitations in the 2018 rulemaking, including removing limits on the list of covered devices, and addressing ambiguity over the definition of the "controlled" environment required for conducting research. But researchers were not able to make a successful case for many other requested changes.

One such request was the removal of a requirement to "not violate any applicable law." Proponents of the change argued that "research can implicate numerous federal and state regulations, with legal uncertainty and uneven application in different jurisdictions," specifically citing the differing interpretations of the CFAA. Researchers feared that this provision pushed the DMCA "into other non-copyright legal regimes, exposing researchers to double liability." They also pointed out that "removal of this condition would not eliminate researchers' obligation

to comply with other applicable laws." While the Register was not persuaded in 2018, removing the "other laws" limitation was requested again as part of the 2021 process. In the 2021 recommendation the Register agreed that the limits were "likely to impose an adverse effect on noninfringing security research" but noting that "this exemption does not serve as waiver of liability for violating other laws while performing good-faith security research."

This exemption's evolution through three cycles of rulemaking show a process that is limiting and cumbersome, putting researchers at a disadvantage. By default those seeking an exemption "bear the burden of establishing that the requirements for granting an exemption have been satisfied." Through multiple rounds of the process some of the issues identified in the exemption's original language were addressed. However, this means that the chilling effects identified in 2017 were prolonged into 2021. Additionally, these exemptions must be renewed—and possibly defended—every three years. Renewal is not guaranteed, meaning certain kinds of research could lose protection in 2024. In addition to the labor burden this creates for researchers needing to regularly participate in a year-long public process, the temporary exemption process has shown itself to be slow in response to legal uncertainties that chill good-faith research, in the fast paced world of tech, this seems untenable.

Policy Recommendations

We are in a vastly different technological landscape now than we were in 1998 when the DMCA was written, much less in 1986 when CFAA was written. We need updated laws to reflect modern technological threats, and how we combat them.. Congress should take the following steps to ensure researchers are able to continue their vital work without fear of reprisal by government or companies. By protecting and encouraging research, Congress will expedite the spread of best practices around security reporting, making it possible for researchers to evaluate an ever expanding catalog of internet connected goods in the marketplace.

1. ***Congress must update the Digital Millennium Copyright Act to create a permanent exemption for legitimate security research.***

Section 1201 of the Digital Millennium Copyright Act creates the existing process for requesting and granting temporary exemptions every three years. This section should be rewritten to create a research exemption that is clear, robust, and permanent. Particularly the anti-circumvention and anti-trafficking provisions need to specifically create permanent protections for researchers who are not seeking to further violate a copyright.

2. ***Congress must update the Computer Fraud and Abuse Act to create a permanent exemption for legitimate security research.***

The Computer Fraud and Abuse Act should similarly be updated to include a clear and permanent exemption for those engaged in good-faith security research. Particularly, the law should clarify its aim as a tool only to be used against those who appropriate, modify, or delete data.

3. ***Congress must update the Computer Fraud and Abuse Act to have a more clear test for civil claims.***

Under the current law, a company can sue someone based on an alleged CFAA violation both in the absence of that alleged violation being criminally prosecuted and without a required standard for what constitutes harm, or any required intent to cause harm. The law should be updated in a way that leaves a right to civil action but clarifies the collection of harms required to constitute a valid civil claim.

4. ***The Government should increase the speed at which it sets up vulnerability disclosure programs, and make them visible.***

Many formal vulnerability disclosure programs have been implemented by government agencies, including 18f, and CISA, who mandated that agencies begin setting them up. These processes create avenues for researchers to responsibly disclose any vulnerabilities they might find in public-facing government tech. As CISA noted in its directive: “Vulnerability disclosure policies enhance the resiliency of the government’s online services by encouraging meaningful collaboration between federal agencies and the public.” Expanding the use of such

programs, and making them easy to find and understand, would serve as a model for how such programs could work at other levels of government, as well as at companies and other organizations.

5. ***The Government should incentivize more companies to implement their own responsible disclosure programs.***

When companies create programs that allow researchers to conduct good-faith research and create avenues for them to disclose their findings, they are able to benefit from the talent and expertise of independent security researchers. These programs should include best practices, including those required by CISA of federal agencies, and include a commitment not to pursue legal action against individuals who are legitimate researchers seeking to discover vulnerabilities. Incentives could include government support for smaller players setting up disclosure programs, requiring vendors who sell the government "smart" tech to create their own disclosure programs, and more.

Conclusion

There have been several high-visibility, headline-grabbing hacks over the last two years, each followed by commentary that "this is a wake up call." It seems that we are living in an age of cybersecurity wake up calls.

The exponential growth of both cyber attacks and the number of connected devices illustrates two trends. There are already a lot of vulnerable pieces of tech out there, and with so much new tech it is likely that more vulnerabilities are introduced every day. Understanding the scope of current cyber threats will require vulnerability research on all manner of connected tech. Growth in the number of researchers in the field is currently held back by fears of civil liability and felony prosecution rooted in the Computer Fraud and Abuse Act of 1986 and the Digital Millennium Copyright Act—a pair of tech laws written in the 1980s and 90s.

Unambiguous protections for good-faith security research are needed. Such protections already have a strong track record under the DMCA, which provides for exemptions to be granted every three years. There have not been widespread problems arising from the three-year exemptions, suggesting that strengthening protections for research won't create new problems.

Congress must take legislative action, updating federal law to expand protections and clarify protections for good-faith security research. The government can also play a key role in changing norms in areas like vulnerability disclosure, and better testing standards.

Researchers who take the time to look for and responsibly disclose vulnerabilities should be able to do their work without fear of legal reprisal. At a minimum, this work should not be slowed down by the current fears rooted in outdated legislation. It's time for these laws to catch up—faced with such a huge threat, expanding protections and incentives that will encourage good-faith security research is a vital step.