Comment from Security Researchers on HIPAA Security Rule Notice of Proposed Rulemaking to Strengthen Cybersecurity for Electronic Protected Health Information

1. Summary

We are a group of cybersecurity experts united by the knowledge that good-faith security researchers (defined below) can help improve cybersecurity. Each of us contributes on an individual basis, and our biographies are at the end.

Please note that this comment is intended to focus on the lack of coordinated disclosure guidance, and this comment is not intended to agree or disagree with any other part of the rulemaking.

Fundamentally, regulated entities have a duty to secure their systems. When security researchers discover problems, including vulnerabilities, misconfigurations, or operational exposures, they may want to disclose those problems to the regulated entity, the public, or regulators. Regulated entities should be required to publish a vulnerability disclosure policy that fosters cooperation with good faith researchers, and HHS should provide a backstop to understand and support that good faith research. (We are intentionally using the term “regulated” to incorporate both covered entities, and also those who enter into Business Associate Agreements with them, regardless of their “Covered” state.)

In the proposed rules, Section 3 Subsection I updates HIPAA’s Technical Safeguards (at 45 CFR 164.312(h)(1)) to add requirements for Vulnerability Management. However, the proposed changes do not include explicit requirements or standards specifically addressing coordinated vulnerability disclosure (CVD) or processes for how security researchers or others should cooperate with HHS and covered entities and report vulnerabilities they discover in healthcare systems. We believe they should.

As Justice Brandeis famously wrote, “Sunlight is the best disinfectant, the electric light the most efficient policeman” His support of transparency can inform how we secure systems. Similarly, Kerkhoffs wrote “The security of a system must not depend on that which cannot be easily changed,” foreseeing the modern opposition to security through obscurity.

In Section 2, we provide context and recommendations and follow them with a specific and compact list of recommendations in Section 3.

2. Background: Cybersecurity Research and Disclosure

A. Vulnerability discovery can disclosure protect bystanders

For over 30 years, independent cybersecurity researchers have been discovering issues in fielded systems, sometimes by virtue of their background, allowing them to notice or understand things others fail to see and are often incidental to other work. For example, those with an apostrophe in their names would frequently discover SQL injection vulnerabilities when entering their names into web forms.

Many of these researchers have expended tremendous effort to fix these problems in ways that minimize risk to innocent third parties. This effort includes tracking down someone who can accept a bug report, managing the project to ensure that the organization acts on it, and often doing quality assurance or testing work for free.

Sometimes, when those efforts fail, researchers will contact the press or use social media to draw attention to the problem as a last effort to ensure it’s addressed. The cybersecurity community has evolved a norm of trying to inform firms of problems in private so that they can be fixed.

Many of these reports are from members of the public who have no specific obligation to help the regulated entity. As security researcher Rain Forrest Puppy wrote when he codified then-current practice in 2001:

> First and foremost, a wake-up call to the software maintainer: the researcher has chosen to NOT immediately disclose the problem, but rather make an effort to work with you. This is a choice they did not have to make, and a choice that hopefully you will respect and accept accordingly. (RFPolicy, 2001)

The FTC has long described "receiving and addressing vulnerability reports from third parties" as a reasonable and appropriate security practice. For example, in 2013, the FTC settled charges with HTC America; the complaint included a failure to establish a process for receiving and addressing vulnerability reports. There may be earlier examples, but a decade is long enough that this can be seen as a norm.

CISA has also worked to support norms of vulnerability disclosure; for example, the Secure By Design pledge includes "[Publishing] a vulnerability disclosure policy (VDP) that authorizes testing by members of the public on products offered by the manufacturer, commits to not recommending or pursuing legal action against anyone engaging in good faith efforts to follow the VDP, provides a clear channel to report vulnerabilities, and allows for public disclosure of vulnerabilities in line with coordinated vulnerability disclosure best practices and international standards."

Industry and government have adopted a norm of coordinated vulnerability disclosure. Thousands of companies now operate vulnerability disclosure programs, especially in critical infrastructures. For example, most major voting machine manufacturers now operate vulnerability disclosure programs. In government, there are numerous examples of government adoption and promotion of CVD: the FDA premarket guidance for medical device cybersecurity; and CISA has required all federal civilian executive branch agencies, including HHS, to operate VDPs. In 2025, operating a vulnerability disclosure policy to receive reports from security researchers should no longer be the exception but an expectation, particularly for those companies that handle sensitive patient data.

B. Bug bounties may enable agreement and rewards

Many firms use bug bounties to incentivize researchers and express their perspectives on what should and should not be researched. This comment is not intended to object to researchers and firms reaching such an agreement; in fact, we applaud it.

Some good faith researchers do not agree to all of the terms and conditions of those bounty programs, which can incorporate various non-disclosure agreements, limits on what can be tested or how, or perhaps they simply don't want to spend the time or pay a lawyer to understand the many pages of many contracts involved. As a result, we recommend that all firms, even those with bug bounties, operate a VDP, as this gives researchers a straightforward way to report vulnerabilities without agreeing to nondisclosure terms.

As an example of how confusing those contracts can be, one company requires clickwrap (a) agreement to their 1800 words of Terms and Conditions, and (b) “acknowledgements that you have read” their 1500 word Code of Conduct and 1300 word Disclosure Guidelines.. Their Code of Conduct starts “all Finders agree to ...[follow] the company Code of Conduct.” This seems to imply that researchers are agreeing to the Code of Conduct. That Code of Conduct also mentions a General Terms and Conditions (3000 or so words) and a Finder Terms and Conditions (1900 words). It refers to a “Finder Terms and Conditions” but the linked document is titled “Community Member Terms and Conditions.” Which documents are legally binding may be confusing to researchers. The 5 documents add up to roughly 7,500 words of legalese. The contract, as we understand it, warns that the researcher may not get paid if the company or its paying clients think the report isn’t a problem.

That example is from one contract. There are several bug bounty providers in the market, and some large companies offer their own. For example, Microsoft says, “Each bug bounty program has its own scope, eligibility criteria, award range, and submission guidelines to help researchers pursue impactful research without causing unintended harm, though they generally share the same high-level requirements,” listing their 20(!) programs, and then offering up 13 additional resources. (https://www.microsoft.com/en-us/msrc/bounty). Google has a similarly expansive set of information at https://bughunters.google.com/. A good faith researcher might, in good faith, simply want to report a problem without reading, understanding, or agreeing to be bound by all that.

Even if a researcher does agree, endless requests for help by firms can lead researchers to want to leave the program and disclose the issue, a choice that may or may not be allowable or that may restrict their future ability to talk about the problem. As such, they may choose not to sign the contracts put forth by the regulated entity.

The bottom line: all companies operating a bug bounty program should also operate a public vulnerability disclosure policy with standard terms, including not requiring nondisclosure of vulnerabilities.

C. Good faith disclosure can happen outside a bounty

The cost of understanding the contracts imposed by a regulated entity may inhibit a researcher's disclosure, resulting in less secure systems.

The DOJ and the U.S. Copyright Office have defined good faith security research as "accessing a computer solely for purposes of good-faith testing, investigation, and/or correction of a security flaw or vulnerability, where such activity is carried out in a manner designed to avoid any harm to individuals or the public, and where the information derived from the activity is used primarily to promote the security or safety of the class of devices, machines, or online services to which the accessed computer belongs, or those who use such devices, machines, or online services". While this definition may not be perfect (i.e., whether a researcher acted "solely" for those purposes or used information "primarily" as defined can be subjective), it encapsulates the idea that researchers should not be punished for following established norms around security research.

Regulated entities should similarly be required to engage with security researchers operating in good faith. The most problematic responses to disclosures are from entities that irresponsibly thrash the first times they're contacted. FDA has also recommended the creation of a coordinated vulnerability disclosure policy since their 2016 Postmarket Management of Cybersecurity in Medical Devices. Specifically, see page 18 on coordinated disclosure, and also note page 22, where vulnerabilities must be reported to FDA.

There are times when disclosure is a thinly veiled attempt to extort, and there are many more times when the receiver of the disclosure takes it as one.

D. Regulated entities should engage in good faith with researchers and each other

Regulated entities should make it easy to send vulnerability or operational security issue reports. This includes making it easy to find information about how to send the report, making it easy to send the report, and sending a receipt or acknowledgement that includes a timeline for further communication. Within the software industry, there is a norm that notified entities should have a fix shipping within 90 days. For example, see Google's policy for vulnerabilities that they discover at https://googleprojectzero.blogspot.com/p/vulnerability-disclosure-policy.html. While there may be unique considerations with hardware devices that extend timelines, the same principles apply: researchers should not be bound by arbitrary nondisclosure timelines. Instead, regulated entities should request a fixed amount of time that researchers do not publicly disclose vulnerabilities. Security researchers value communication – if an organization is communicative about the fix timeline of a vulnerability and conveys that

they need a little more time to get a fix out, most security researchers would be receptive.

Regulated entities should plan to have engineering or IT staff, not lawyers, as the lead contacts. (This has been the norm at Microsoft and other leading firms for nearly 30 years). Regulated entities should send regular updates to researchers. Many researchers reasonably see “lawyering up” by regulated entities as an implicit threat.

Incident responders frequently band together in the wake of a security issue being disclosed. Regulated entities may be impacted by the disclosure of a vulnerability in a medical device, regulated software, or software “in the supply chain” such as log4j. Nothing in this comment is intended to undermine or question their ability to do so either in formal response processes, such as that operated by a Sector Coordinating Council, Information Sharing and Analysis Organization, or informal response processes that frequently appear in community settings.

3. Specific Suggestions: HHS should require cooperation with Good Faith researchers

We offer 5 specific suggestions for the improvement of the security rule.

1. All regulated entities should be required to enable people to report security issues in a way that’s easy to discover and aligned with standards.
    a. Entities should establish a public point of contact for reporting vulnerabilities. Examples would include operating a vulnerability disclosure policy, publishing a security.txt file, publishing and monitoring a secure@ or security@ email address, and/or listing a phone number.

2. All regulated entities that produce software should be required to publish a vulnerability disclosure policy.
    a. At minimum, in alignment with CISA guidance, a VDP should do the following. Example terms aligned with this policy can be found here.
        i. Authorize testing by members of the public and commit to not recommending or pursuing legal action for good-faith violations of the policy. This is often referred to as a legal “safe harbor”.
        ii. Allow any member of the public to report a vulnerability in any of the organization’s systems. Restrictions on who can report or what systems they can report on discourage disclosure and may lead to vulnerabilities going unreported.

iii. A VDP should not arbitrarily restrict public disclosure. CVD should be a mutual process and a culture of transparency is crucial. Organizations can request that researchers give them a reasonable amount of time to fix vulnerabilities, but they should not place blanket non-disclosure requirements on researchers.

3. Regulated entities should be discouraged from threatening Good Faith researchers.
    a. Threats are meant to include threats of lawsuit, retribution, or other threats.
    b. Good Faith should be defined by HHS, not by regulated entities (see also section 3)
        i. Good Faith should not require agreeing to or signing an NDA or similar agreement restricting disclosure.
        ii. Disclosure terms may not restrict whistleblowing to HHS, CISA, or other relevant agencies about security issues.
        iii. Good Faith may include the use of a bug bounty program but cannot require agreement to contracts of adhesion, especially those that restrict the researcher's ability to disclose or to otherwise perform their work.
        iv. Good Faith should be construed as imposing more requirements on the regulated entity than the researcher. After all, a regulated entity has chosen to operate in a regulated environment and that carries ethical obligations.
    c. The definition of Good Faith should be expandable by regulated entities (that is, an entity's definition of good faith might be broader than HHS's definition but not more narrow.)
    d. Good Faith should be monitored by HHS.
        i. HHS should have a dedicated reporting line for good faith researchers who are threatened who are threatened or retaliated against.
        ii. HHS should have a whistleblower program for employees of firms that use or offer bug bounty programs who are using those programs to bury vulnerability reports.
        iii. The reporting should go to a single "desk" or station because the person calling may fall into multiple categories. For example, Microsoft operates a bug bounty, and a Microsoft employee may find a security flaw in another company's product. If there are multiple reporting desks, it may be daunting to figure out which to contact, and the staff may think it's "someone else's problem."

4. Regulated entities should be rewarded for positive engagement with Good Faith researchers, such as:
    a. Operating a bug bounty.
    b. Supporting researchers, such as the biohacking village,

    These rewards could include construing such engagement as evidence of constructive engagement with the intent of the rule.

5. HHS should add “insecure operations” to the wall of shame, including threatening Good Faith researchers or possibly even failing to engage in Good Faith.
    a. The primary weapon that firms wield against reporters is the threat of lawsuit. Regulated entities are far less likely to sue their regulator than they are to sue an individual researcher, but the option is open to them.

6. Receipt of a Good Faith report must be tracked and managed, but not all reports rise to the level of an incident.
    a. Risk Assessment procedures should include using Good Faith reports to assess the quality of the Risk Assessment procedures. Should the regulated entity have discovered the reported facts on their own?
    b. Incident Response procedures should be similarly checked. If a malicious attacker had discovered the issue, would the response have been sufficient?

4. Good Faith disclosure factors

Two important concepts are at play: what is good faith and thus deserving of protection and what activities should expose a researcher to civil liability. There is a gray area between them. To achieve society’s security goals, we recommend that HHS and regulated entities use a broad definition of good faith and that those actions where a researcher might be liable are clearly stated and narrowly construed.

DOJ and the U.S. Copyright Office’s definition of good-faith security research offers a solid starting point. We encourage future definitions to broaden, not narrow, this definition; for instance, to strike the words “solely” and “primarily”.

The following activities, at minimum, should be considered good faith behavior by a security researcher:

- Providing a report at a level of detail that a normal appsec team could parse.
    - Researchers are not obligated by reasonableness or good faith to provide a tutorial in application security. Being able to parse a report is a capability a regulated entity should have.
    - Researchers may, in good faith, offer consulting help if the regulated entity can't parse a report or asks for proof of concept code.
- Reporting soon after discovery, possibly at a lower level of detail or clarity as a result.
- Rapidly informing the public when there is evidence that the issue is being exploited.
- Setting a timeline for public disclosure of an issue, as long as the regulated entity has at least 90 days to address the issue.
- Obtaining a CVE or other identifier.
- Pre-briefing media or government officials.
- Reporting the vuln to their government in accordance with legal requirements.
- Complying with ethical obligations as outlined (for example) by the ACM Code of Ethics, the ISC2 Code of Ethics.

Some factors are sometimes cited by those receiving vulnerability reports, which we believe do not play a factor in determining good faith and should not — by themselves — lead activity to be considered in bad faith. Those include:

- The researcher's use of a pseudonym.
- The recipient may feel threatened because someone sent them a vulnerability report. Of course, this would not apply if there were threatening words. Setting a reasonable timeline, such as "This will be disclosed in 90 days," should not be construed as a threat.
- The researcher asked about a bug bounty program. (Note that we recommend that researchers ask about a bounty program only after the issue is resolved, but asking earlier should not be seen as evidence of bad faith.)
- The researcher does not have excellent communication skills — many reporters are young, and many are not native English speakers.
- Submitting a conference talk
    - As long as it's 30 days or more after the expected disclosure (to allow for some slippage)
    - Naming the company whose products or services were impacted is not a factor that demonstrates unreasonable behavior. We note that the Consumer Review Fairness Act, "protects people's ability to share their honest opinions about a business's products, services, or conduct..." A security research report is, unavoidably, an opinion about a business's products, services or conduct. The FTC has a fact sheet,

https://www.ftc.gov/business-guidance/resources/consumer-review-fairness-act-what-businesses-need-know.

There are also actions by researchers that weigh against reasonableness:

- Demanding money in exchange for details of the vulnerability. (This assumes that the researcher has sent a report of sufficient detail for an appsec team of normal skill to understand the issue. A researcher who has done so may reasonably tell a company they need consulting help, and offer it.)
- Unreasonably short timelines (absent evidence that the issue is known to others)
- Attempting to forbid the regulated entity from seeking outside help, such as counsel or other experts, to parse the report
- Having a more detailed report already prepared and refusing to share it.
- Exfiltrating real data in bulk beyond the need for proof of concept
- Establishing ongoing command and control or persistence beyond the need to demonstrate success
    - (For example, some exploitation of log4j or a CI/CD pipeline may reasonably call out to a researcher-controlled system to establish success)
- Compromising the privacy or safety of people or engaging in breaking and entering into a building or facility.
    - This should not be read to declare 'dumpster diving' out of scope if the dumpster is not in a controlled location

There are elements of security research that an entity may reasonably declare out of scope:

- Social engineering
- Physical access

The consulting help bullets are in intentional tension, and a possible fulcrum is "Does the company want the researcher to do additional work beyond what the researcher feels would be needed for a public disclosure?"

5. Background

Two stories motivated us to write this. First, Adam's health insurer has a plethora of full SSN issues, including emailed spreadsheets, display on their site, default inclusion in

data exports, and no MFA being even possible on one such site. Adam's efforts to get this fixed through means available to people without my connections have failed.

A Second background story: While getting hired at Microsoft in 2006, Adam's future boss accidentally accessed Adam's voicemail. Literally, it was an accident that Adam's future boss happened to trigger. Adam tracked it down to a design flaw in the SS7 switching protocols, and attempted to report it to AT+T, his phone provider. AT+T had no system for accepting vulnerability reports at the time. Microsoft's "Responsible Disclosure" policy for employees forbade Adam from disclosing the issue publicly. Adam enlisted the aid of senior leaders at Microsoft to push AT&T to fix it, and the issue was reported to AT&T at the Chief Security Officer level. To the best of Adam's knowledge, the issue remains unfixed.

6. Contributors

All affiliations are listed for informational purposes only. The views and opinions expressed in this comment are solely those of the authors and do not necessarily represent those of their respective organization(s), past or present.

Jack Cable

Jack Cable is a hacker who works at the intersection of cybersecurity and public policy, currently the CEO and Co-Founder of Corridor. Prior to that, Jack served as a Senior Technical Advisor at the Cybersecurity and Infrastructure Security Agency (CISA), where he helped lead the agency's Secure by Design initiative. Before CISA, Jack worked as a TechCongress Fellow for the Senate Homeland Security and Governmental Affairs Committee, advising Chairman Gary Peters on cybersecurity policy, including open source software security. He previously worked as a Security Architect at Krebs Stamos Group. Jack is a top bug bounty hacker, having identified over 350 vulnerabilities in hundreds of companies. After placing first in the Hack the Air Force bug bounty challenge, he began working at the Pentagon's Defense Digital Service. Jack studied computer science at Stanford University and has published academic research on election security, ransomware, and cloud security.

Dissent Doe

Dissent Doe is a healthcare practitioner and a journalist who chooses to remain pseudoanonymous.

Josiah Dykstra, Ph.D.

Josiah Dykstra is the Director of Strategic Initiatives at Trail of Bits and the owner of Designer Security, a cybersecurity firm for healthcare. He previously served for 19 years at the US National Security Agency (NSA). Dr. Dykstra is an experienced practitioner and researcher whose focus has included the psychology and economics of cybersecurity. He received the CyberCorps® Scholarship for Service (SFS) fellowship and the Presidential Early Career Award for Scientists and Engineers (PECASE). He is the author of numerous research papers, the book *Essential Cybersecurity Science,* and co-author of the award-winning book *Cybersecurity Myths and Misconceptions*. Dr. Dykstra holds a Ph.D. in computer science from the University of Maryland, Baltimore County.

Fred Jennings

Fred Jennings is a cybersecurity attorney based in Brooklyn, New York. He provides consulting and legal advice on information security and technology issues at cybers.legal for clients nationwide. Previously he has written an open-source bug bounty policy now used by GitHub and Microsoft, and has served as Senior Data & Product Counsel at Zocdoc, Inc. working on healthcare technology security issues. Prior to that he has represented hacktivists and others accused of cybercrime in federal courts across the country. He is a member of the Internet Law & Policy Foundry and has partnered with other organizations on open source licensing, software security, and open-access issues in hardware and software projects.

Chloé Messdaghi

Chloé Messdaghi is a recognized expert in Responsible AI and cybersecurity, focused on creating secure, ethical, and transparent AI systems. She collaborates with companies, governments, and NGOs to advocate for stronger cybersecurity and Responsible AI policies and has advised on key legislation, including work with the White House. She is one of the founders of Disclose.io and the founder of SustainCyber. Named a Power Player in Cybersecurity by Business Insider and SC

Media, Chloé is a respected speaker and thought leader dedicated to advancing technology and social progress.

Adam Shostack

Adam Shostack is the author of Threat Modeling: Designing for Security and Threats: What Every Engineer Should Learn from Star Wars. He's a leading expert on threat modeling, a consultant, expert witness, and game designer. He has decades of experience delivering security. His experience ranges across the business world from founding startups to nearly a decade at Microsoft.

His accomplishments include:

- Helped create the CVE. Now an Emeritus member of the Advisory Board.
- Fixed Autorun for hundreds of millions of systems
- Led the design and delivery of the Microsoft SDL Threat Modeling Tool (v3)
- Created the Elevation of Privilege threat modeling game
- Co-authored The New School of Information Security

Beyond consulting and training, Shostack serves as a member of the Blackhat Review Board, an advisor to a variety of companies and academic institutions, and an Affiliate Professor at the Paul G. Allen School of Computer Science and Engineering at the University of Washington.