Nowhere to Report: Vulnerability Disclosure Channels at US Hospitals, Medical Device Makers, and Health IT Developers. A Cross-sectional Web Measurement

Christopher Jagdeo, Independent Researcher

Abstract

Objective: To measure whether US hospitals, medical device manufacturers and certified electronic health record (EHR) developers publish any channel for reporting a security vulnerability.

Materials and Methods: A cross-sectional crawl in November 2026 visited domains on five public lists: hospitals listing a website (2,652), companies on FDA's list of AI-enabled medical devices (device makers, 525), vendors named in CISA medical advisories (advisory vendors, 75), developers of ONC-certified API technology (EHR developers, 200) and federal health agencies (the federal control, 124). Channels were security.txt, a disclosure policy page (judged by language models), a platform listing or CVE Numbering Authority status.

Results: Of these, 1,721 hospital domains could be measured (615 unreachable), as could 436 device-maker, 67 advisory-vendor, 181 EHR-developer and 101 federal-control domains. A channel was found at 87.6% (95% CI, 79.0%-93.7%) of federal-control, 8.2% (6.9%-9.6%) of hospital, 20.5% (16.7%-24.8%) of device-maker, 41.0% (28.6%-54.3%) of advisory-vendor and 7.9% (4.3%-13.2%) of EHR-developer domains. One health system's file served 95 of 132 hospital domains with security.txt; per distinct final site (post hoc), 3.1% (2.2%-4.1%) had one. On the same domains, 91.0% served a valid cms-hpt.txt, required by a price-transparency rule, and 7.9% security.txt.

Discussion: Channels were common where a directive requires them and rarer elsewhere, an association, not an effect of regulation. Per-domain counts overstate how many hospital organizations publish one.

Conclusion: Most organizations on these lists published no vulnerability reporting channel; most federal-control domains did. Among hospitals, channels came mostly from a few health systems. Whether such reports are answered remains untested.

Background and Significance

Outsiders find security flaws in healthcare systems and need a way to report them. One internet-wide scan found 2,841 healthcare services with authentication flaws, which its authors then set out to disclose.[1] Yet reporting channels are often missing,[2] and without one reports go astray: in one notification experiment, only 5.8% of 35,832 vulnerability reports arrived.[3]

The stakes are clinical. Ransomware attacks on US hospitals cut patient volume by 17% to 24% in the first week and raised in-hospital mortality among patients already admitted by 34% to 38%.[4] One 2021 attack was also associated with higher patient volume at two nearby untargeted emergency departments.[5]

A reporting channel takes four forms. A security.txt file sits at a fixed web address and must name a contact and an expiry date.[2] A vulnerability disclosure policy states what may be tested, how to report and how reports are handled[6]; federal policies must also promise not to take legal action against good-faith research, which this paper calls safe harbor.[7] A disclosure program may also be listed on a bug-bounty platform such as HackerOne or Bugcrowd[8] or in disclose.io's open directory.[9] A CVE Numbering

Authority, authorized to assign identifiers to vulnerabilities, must publish a contact and a disclosure policy link.[10,11]

Rules push only some organizations toward a channel. Binding Operational Directive 20-01 (2020) requires federal executive-branch agencies to publish a disclosure policy, though not security.txt[7]; a report for the last quarter of fiscal year 2022 gave the first phase, publishing a policy, as "99% met", and 78 of 101 agencies as having brought all internet-accessible systems under it.[12] Section 524B of the Federal Food, Drug, and Cosmetic Act, enacted in December 2022, requires a premarket submission for a cyber device to include a plan to monitor, identify and address its vulnerabilities, "including coordinated vulnerability disclosure and related procedures".[13,14] Coordinated disclosure means gathering information from finders, coordinating its sharing among stakeholders, and disclosing vulnerabilities and their mitigations to stakeholders, including the public.[15] For hospitals and electronic health record (EHR) developers, rules stop short. Federal guidance says a healthcare delivery organization "should have a program in place to accept inbound vulnerability disclosures", without asking it to publish a channel[16]; the sector's performance goals are voluntary[17]; the 2025 proposed HIPAA Security Rule revision does not use the phrases "disclosure program" or "coordinated disclosure"[18]; and the EHR certification rules contain no requirement to publish a reporting channel.[19]

Hospitals do face a federal rule for a comparable contact file. Since 2024, the Centers for Medicare & Medicaid Services has required a cms-hpt.txt file, naming a contact and linking to the hospital's machine-readable price file, at the root of whichever website hosts that price file, possibly a health system's.[20,21]

Earlier counts, none peer reviewed, mostly report low adoption, so a low hospital rate would not be new. Mandel's 2014 blog post found no official reporting mechanism on 83 EHR vendor websites, though he may have missed some.[22] A table in a Future of Privacy Forum blog post dated 2024, with no stated method, listed 12 of 2,282 hospitals (0.53%) with a valid security.txt.[23] MedCrypt's 2025 whitepaper found "some communication published around a vulnerability disclosure process", ranging down to an email link, at 27 of the 40 largest medical device vendors by market capitalization.[24] In 2026, the software vendors CVD Portal and AnnexProof found security.txt at 9 of 42 EU medical device makers (21.4%)[25] and 280 of 1,691 European device and software makers (16.6%).[26]

Peer-reviewed studies have measured security.txt on popular websites[27-29] and hospital website security,[30] but none was found that measures, with a documented and reproducible method, whether US hospitals, medical device manufacturers and certified EHR developers publish any channel. A channel shows only that an organization invites reports; its absence is not evidence of poor security.

Objective

This study measured whether US hospitals, medical device manufacturers and certified EHR developers publish any channel for reporting a security vulnerability, with federal health agencies as a positive control. It asked: (1) what share of each sector, and of the control, publishes any channel; (2) how security.txt adoption compares with cms-hpt.txt adoption on the same hospital domains; (3) descriptively, which hospital characteristics go with having a channel; and (4) whether existing channels are usable (expired files, missing contacts, safe-harbor commitments).

Materials and Methods

Study design and setting. This cross-sectional study, a census of five public lists, followed a protocol frozen by commit `3149cb366d69b891a174e012a4dd0e4ee319ae2c` before the main collection; reporting follows STROBE.[31] The unit of measurement was the registrable domain.

Sampling frames. The lists were fixed on November 16, 2026, before any measured site was contacted. Hospitals were US hospitals in the HIFLD dataset that list a website, from its Hospitals layer (8,340 records; published June 30, 2024) as served by Penn State's PASDA.[32] Of the 7,531 open, nonfederal hospitals in the 50 states and the District of Columbia, 3,597 (47.8%) listed no website and are not covered; excluding 8 unparsable and 233 third-party websites, such as government portals, left 3,693 hospitals on 2,652 domains. One Claude coding agent, without a second reviewer, built the third-party list and reviewed 388 doubtful listings. Device makers were companies on FDA's list of AI-enabled medical devices[33]; advisory vendors, kept as a separate list, were vendors named in CISA medical advisories[34]; and EHR developers were developers of ONC-certified API technology listed by ONC's Lantern.[35] Two Claude language models[36] named each company's website, and one web-search call settled each disagreement, giving 525 device-maker, 75 advisory-vendor and 200 EHR-developer domains. The federal control, a positive control, was 124 federal health agency domains: the executive-branch .gov domains of the Departments of Health and Human Services and Veterans Affairs,[37] which a federal directive obliges to publish a disclosure policy.[7] The five lists held 3,576 entries on 3,557 distinct domains, because 19 domains are on two lists. (HIFLD, Homeland Infrastructure Foundation-Level Data; FDA, Food and Drug Administration; CISA,

Cybersecurity and Infrastructure Security Agency; ONC, Office of the National Coordinator for Health Information Technology.)

Data collection. The crawler, the same for every list, sent only GET requests and ran no scripts, forms, logins or vulnerability tests. A request refused under the study's identifying user agent was retried once with a standard browser user agent, kept for that site for the rest of the pass if answered. It requested the homepage over HTTPS, robots.txt, security.txt at its two locations,[2] `/cms-hpt.txt` (hospitals), any security.txt `Policy:` address, up to three homepage links worded like a disclosure or security page, and six conventional addresses such as `/vulnerability-disclosure-policy`. Requests were capped per domain and per site, at least 2 seconds apart per site and stopped at any site that answered 429 (too many requests). robots.txt governed later page requests on the homepage's host, an unreadable one forbidding all; the two text files were requested regardless (exact rules in the Supplementary Methods). Pass 1 crawled the 3,557 domains on November 17, 2026, 07:37 to 08:45 UTC, and sent 26,746 requests; its 30,040 stored records also hold copies of answers reused for other listed domains. Pass 2 (November 19, 08:50 to 09:01 UTC), at least 48 hours after pass 1 ended, repeated the 1,049 domains not fully measured with 2,935 requests.

Outcomes. A domain was measured when its homepage answered status 200 unrefused, was not a parked, for-sale or third-party platform page, and passed the identity check. Otherwise it was blocked (refused even after the browser retry, or rate-limited), unreachable (no usable answer, also at the other `www.` form when tried) or a mismatch. These could not be measured and were never counted as lacking a channel in the primary analysis. A domain with one unmeasured outcome leaves only that outcome's

denominator. A domain's final site is the registrable domain on which its homepage ended after redirects.

The main outcome was a security.txt read as text over HTTPS with a valid `mailto:`, `https://` or `tel:` contact; the strict outcome also required a single unexpired, time-zoned `Expires` date. The any-channel outcome was the main outcome, a policy page, or a listing of the domain in pinned copies of disclose.io,[38] a third-party compilation of four bug-bounty platforms' public scopes[8] or the CVE Numbering Authority list[10]; with nothing found and a component unmeasured, it was unmeasured.

Classification. Pages from the policy search mentioning neither security.txt nor a fixed list of disclosure phrases were negative. `claude-opus-5-5` and `claude-sonnet-5-5` judged every remaining page independently, without its address, on whether it tells outsiders how to report a vulnerability in the organization's own systems, gives a way to report, and promises safe harbor. A policy page needed yes to the first two; `claude-fable-5-2` broke ties by majority. Links were followed only from the homepage, so this measure is a lower bound.

In the identity check, `claude-sonnet-5-5` and `claude-haiku-5-5` decided independently, from the listed names and the homepage, whether it was the working website of the organization or its present owner (`own`), another site (`other`) or unattributable (`unclear`); `claude-opus-5-5` broke ties, and `other` or `unclear` made a mismatch. Every model call was logged.[36]

Validation. Two pilots, crawled before the freeze and excluded from results, were audited, and misfiring rules were changed. Audits of each pass checked every stored

record against the request limits; read every distinct file counted positive and every distinct judged page text; scanned every identity decision and every unreachable or blocked domain; and read fixed-seed samples of the other classes and decisions.

Statistical analysis. The following analyses were pre-specified. Sector shares carry exact Clopper-Pearson 95% confidence intervals,[39] with the domain as the unit of inference; hospital shares are also given per hospital and bed-weighted, without intervals. On hospital domains where both files were measured, the presence of security.txt (main outcome) and of a valid cms-hpt.txt was compared with the exact McNemar test[40]; because a health system may host the file for its hospitals,[21] this is not a test of compliance. Hospital any-channel shares were described, without regression, by system size (Fisher exact test of single- against multi-hospital domains[41]) and, among single-hospital domains, by type, ownership, beds and rural location, rural meaning a nonmetropolitan county (Rural-Urban Continuum Code 4 to 9).[42] A sensitivity analysis counted every unmeasured value as lacking a channel, per domain and per distinct final site. The federal control was also counted per distinct final site, and each of its measured domains without a channel was read individually in the audits. Agreement is reported raw and as Cohen's kappa.[43] Four hospital analyses added after the first pass are post hoc: the share among measured final sites, the number of distinct files and of the organizations serving them, the share without the domains serving the most common file, and a comparison of hospitals on measured and unmeasured domains.

Ethics and transparency. No human participants were involved, no personal data were collected for analysis and no institution took part, so no ethics review was sought; the study followed the Menlo Report's principles.[44] The user agent linked to a page

describing the study and its opt-out, which deletes a domain's records. Code, frame tables, per-domain results and the STROBE checklist are at https://github.com/k0nnect/noroshi (commit `328be22`). The software, the analysis and the first draft of this manuscript were produced with Claude language models (Anthropic) under the author's direction; the Acknowledgments describe this use in full.

Results

Measurement coverage. Pass 2 measured 26 more domains; 41 of the 144 partly measured domains it retried were refused (39) or unreachable (2), keeping first-visit values. Of the 2,652 hospital domains, 1,721 (64.9%) could be measured (Table 1); 615 were unreachable, 151 blocked and 165 mismatches. Most unreachable hospital domains no longer existed in the Domain Name System (DNS; 199) or failed the encrypted connection or its certificate (TLS; 174). Most blocked domains met a browser challenge page (99 of 170 across all lists). Most hospital mismatches were parked, for-sale or reused domains, and 26 led to another organization's site. Coverage was higher on the other lists: 436 of 525 device-maker (83.0%), 67 of 75 advisory-vendor (89.3%), 181 of 200 EHR-developer (90.5%) and 101 of 124 federal-control domains (81.5%).

Federal control. The any-channel outcome was positive at 78 of 89 measured federal-control domains with a known value (87.6%; 95% CI, 79.0%-93.7%) and at 41 of 52 distinct final sites (78.8%; 65.3%-88.9%) (Table S1). A policy page was found at 76 of 87 domains (87.4%), and one department-wide page served 74 domains. Of the 11 negatives, 9 linked no disclosure policy, 1 had an error homepage, and 1 was a true miss, its link drawn by a script the crawler does not run. Of the 12 with an unknown value, 2, stopped by the per-site request limit, fell outside the retry rule; 7 homepages showed no

readable content, 2 met challenge pages, and 1 policy page was unreadable. The 23 unmeasured domains were unreachable (20, 1 through a bot-block redirect) or blocked (3, 1 of them a retired domain).

Validation. Before tie-breaks, the two judges agreed on 153 to 155 of 156 distinct page texts for each question (kappa 0.95 to 0.98), and the identity raters on 2,490 of 2,574 homepages (96.7%; kappa 0.80) (Table S2). Audits of both passes found no broken request limit, no false-positive policy page among the 116 distinct texts judged to invite reports and no missed disclosure link on 70 sampled first-pass or 127 second-pass homepages. None of the 40 texts judged negative was misjudged, but 2 were challenge pages (rightly not policies), so on those 2 domains, and on 2 more whose file or page answers were challenge pages served with status 200, the value should have been unmeasured, not negative.

Channels by sector. Outside the federal control, the any-channel outcome was positive at 132 of 1,618 hospital domains (8.2%; 95% CI, 6.9%-9.6%), 83 of 404 device-maker domains (20.5%; 16.7%-24.8%), 25 of 61 advisory-vendor domains (41.0%; 28.6%-54.3%) and 13 of 164 EHR-developer domains (7.9%; 4.3%-13.2%) (Table 2; Figure 1). At hospitals the channel was a security.txt file: every hospital domain with a policy page (114) or a listing (1) also had one. Device makers more often had a policy page (59 of 402; 14.7%) than a security.txt (30 of 430; 7.0%) or a listing (25 of 436; 5.7%). Advisory vendors had policy pages (27.0%) and listings (26.9%) more often than files (11.5%), and EHR developers had each at 5.5% or less. Counting every unmeasured value as no channel lowered the shares to 5.0% for hospitals, 15.8% for device makers, 34.7% for advisory vendors, 7.0% for EHR developers and 62.9% for the federal control (Table S3).

Concentration of the hospital result. Per listed domain, the pre-specified unit, 132 of 1,682 hospital domains (7.8%; 95% CI, 6.6%-9.2%) had a security.txt. Analyses added after the first pass (post hoc) show how concentrated the result was (Table 3). The 132 domains served 19 distinct files from 17 organizations. One HCA Healthcare file, served on 25 of its sites and naming a corporate contact and disclosure statement, accounted for 95 of the 132 domains (72.0%) and 120 of the 175 hospitals on positive domains; that statement was also the policy page of 95 of the 114 hospital domains with one. Three nearly identical CommonSpirit Health files accounted for 18 more domains; 13 files were each on a single domain. Counting each distinct final site once gave 42 of 1,374 measured sites (3.1%; 2.2%-4.1%), and leaving out the 95 domains serving the most common file gave 37 of 1,587 domains (2.3%; 1.6%-3.2%). The pre-specified sensitivity analysis, which counts every final site of the list with unmeasured values as none, gave 42 of 2,340 (1.8%; 1.3%-2.4%).

Security.txt and the price-transparency file. On the 1,680 hospital domains where both files were measured, 1,529 (91.0%) served a valid cms-hpt.txt (with a price-file link and a location name) and 132 (7.9%) a security.txt: 131 both, 1 only security.txt, 1,398 only cms-hpt.txt and 150 neither (exact McNemar $P < .001$; Table 3). A cms-hpt.txt of any validity was on 1,537 (91.5%; 95% CI, 90.1%-92.8%). This compares presence on the listed domain, not compliance, and 95 of the 132 security.txt domains served the HCA file.

Usability. Most hospital domains with a file had no Expires field: 97 of 132 (73.5%; 95% CI, 65.1%-80.8%), 95 of them serving the HCA file; of the other 35, 32 had a valid date, 1 had expired and 2 used a nonconforming format (Table 4). Of 30 device-maker

domains with a file, 22 (73.3%) had a valid date, 6 (20.0%) had expired and 2 had none; on all 4 federal-control domains with a file it had expired. The 7 advisory-vendor and 8 EHR-developer domains with a file were too few to compare. Nearly all such hospital and device-maker domains named an email contact (130 of 132; 27 of 30); every file was served as text/plain, and a second parser agreed on every contact. Among domains with a policy page, the page promised not to pursue good-faith reporters at 76 of 76 federal-control domains (100%; 95% CI, 95.3%-100%), 21 of 59 device-maker (35.6%; 23.6%-49.1%), 4 of 17 advisory-vendor (23.5%; 6.8%-49.9%), 3 of 9 EHR-developer (33.3%; 7.5%-70.1%) and 0 of 114 hospital domains (0%; 0%-3.2%).

Hospital characteristics. These shares are descriptive, and many cells are small (Table S4). A channel was found at 7.8% of 1,327 single-hospital domains and 11.2% of 249 domains with 2 to 5 hospitals (single against multi-hospital, Fisher exact P = .24), and at 1 of the 42 domains with more than 5. Among single-hospital domains, shares were higher at general acute care (13.3% of 684) than critical access (1.4% of 282) or psychiatric hospitals (2.0% of 200); at proprietary (14.9% of 511) than non-profit (3.7% of 508) or non-federal government hospitals (2.6% of 235); with more beds (3.1% under 100, 12.6% at 100 to 299, 21.3% at 300 or more); and at non-rural (10.5% of 841) than rural hospitals (3.1% of 486). Four smaller types had 25 to 56 domains and at most 3 positives each. Because one system's file accounts for most positive domains, these differences may reflect which hospitals it operates; they were not recomputed without it.

Discussion

Federal agencies must publish a disclosure policy,[7] and the crawler found a channel at 87.6% of the federal-control domains it could measure. Among hospitals, device makers,

advisory vendors and EHR developers, most published none, and the hospital figure rested largely on one health system's file. On the same hospital domains, the price-transparency file was nearly universal and security.txt was rare.

These figures sit beside the earlier counts in the Background, none peer reviewed and none directly comparable. The blog table counted hospitals, with no stated method or definition of a valid file[23]; here, in 2026, 32 of 1,682 hospital domains (1.9%) had a fully valid file, counted per domain. MedCrypt took the 40 largest device vendors by market capitalization and counted any published disclosure communication, down to an email link[24]; here 20.5% of device makers, a list not chosen by size, had a channel, as did 41.0% of advisory vendors, named in advisories that detail vulnerabilities affecting medical devices.[34] The two 2026 European counts took any published security.txt[25,26]; here 7.0% of device makers had a file with a valid contact. Mandel reviewed 83 EHR vendor websites in 2014[22]; here 7.9% of EHR developers had a channel.

The contrast between the federal control and the other lists is an association across five lists at one time point, not an effect of regulation. Federal agencies differ from hospitals and companies in more than their obligations, and no list was observed before and after a rule. Still, what was published matched what each rule asks for. Binding Operational Directive 20-01 requires a policy page with a commitment not to pursue good-faith researchers and does not require security.txt[7]; federal-control domains had policy pages far more often than files (87.4% against 4.0%), and every policy page in the federal control carried that commitment. The price-transparency rule requires a root-level text file naming a contact,[20] and a valid one was on 91.0% of the hospital domains compared, against 7.9% for security.txt. Device makers, whose premarket submissions for cyber

devices must include a plan covering coordinated disclosure,[13] and advisory vendors had intermediate shares; these data cannot say why.

In practice, the steps are modest. Hospitals already keep a root-level text file that names a contact[20]; a security.txt has the same form, with a contact and an expiry date.[2] Organizations that publish a file could add or renew its Expires field, since RFC 9116 treats data past that date as stale.[2] A health system that serves one file on all its sites gives each of them a channel, and repeats any gap on each.

This study did not test whether a report sent to these channels is received or answered; that is the next measurement to make. Also unmeasured are whether hospitals without a file publish one on the health-system site that hosts their price file, how the shares change over time, and whether a channel goes with fewer or shorter incidents.

Limitations. The hospital list covers only hospitals listing a website in a HIFLD layer published in 2024[32]: 3,597 of 7,531 eligible hospitals (47.8%) list none, and those kept differ from those dropped (psychiatric hospitals: 9.6% against 16.7%; Table S5). The list's age shows in the 615 unreachable hospital domains (23.2%), 199 no longer in DNS and 225 whose last archived homepage redirected to another domain. Some unmeasured domains might load in a browser: 57 failed only on an incomplete certificate chain. Hospitals on measured and unmeasured domains had broadly similar characteristics (Table S6). The final audit found 51 domains refused by a challenge page the crawler did not recognize, labeled unreachable (45 of 705) or mismatch (6 of 187 `other` or `unclear` decisions) instead of blocked; correcting them and the 4 negative values read from challenge pages moves no headline share by more than 0.1 point. One device

maker's security.txt was its website platform's; without it, device-maker main and any-channel shares would be 29 of 430 (6.7%) and 82 of 404 (20.3%). The policy-page measure is a lower bound: only homepage links were followed, and two policy pages were unreadable; listings cover only public programs. Of all 2,387 `own` decisions of the model-assisted identity check, 9 (revived copies of former hospital sites) were wrongly accepted; without them the hospital any-channel share is unchanged at 8.2% (132 of 1,609). Among the 187 `other` or `unclear` decisions, 12 were own sites wrongly set aside: the 6 challenge pages above and 6 script-built or near-empty pages. A listed domain is not an organization, and intervals treat domains as independent although health systems share files. The study measured one time point, with one retry pass, and files change. Device makers and EHR developers come from specific lists: FDA calls its list "not a comprehensive resource",[33] and Lantern lists developers of certified API technology,[35] not all EHR software makers. Language models matched company names to domains, with no accuracy estimate. The results describe these lists in November 2026 and may not hold outside them. The price-file comparison is of presence, not compliance.[21] A published channel shows only that reports are invited, not that they are handled, and its absence is not evidence of poor security.

Conclusion

On the five lists, most hospitals, device makers, advisory vendors and EHR developers published no channel for reporting a security vulnerability; the federal control showed that the method finds most channels where a rule requires them. Per listed domain, the hospital share overstates how many organizations publish a channel, because most hospital domains with a file served one health system's file. On the same hospital

domains, the file a federal rule requires was nearly universal and the voluntary security.txt rare. That shared file lacked an expiry date, and no hospital policy page promised safe harbor. Whether anyone answers these channels remains untested.

Acknowledgments: use of artificial intelligence

This study relied on artificial intelligence throughout, and the author wants readers to know exactly how. Claude language models (Anthropic; Claude Fable 5.2, Opus 5.5, Sonnet 5.5 and Haiku 5.5), run as software agents through Claude Code, proposed candidate topics, drafted the study design and the protocol, wrote and reviewed the software, ran the data collection, carried out the audits described in the Methods, produced the analysis, and wrote the first draft of this manuscript. Within the study itself, models of the same family read web pages and made the classifications described in the Methods; every such call is logged with its prompt and answer. The author chose the topic from the candidates and set its scope, set the constraints the work had to respect, approved the design, the public information page and each data-collection run, read and approved the manuscript, and takes full responsibility for its content. A running log of AI use is in the repository (`logs/ai-use-log.md`). No AI system is an author.

Funding

None. The author received no funding for this work.

Competing interests

The author takes part in vulnerability disclosure and bug bounty programs as an independent security researcher. No organization measured in this study funded it, was consulted on it, or saw the results before publication.

Author contributions

Christopher Jagdeo: conceptualization (selection of the question and scope), supervision, review and approval of the manuscript. Software, data collection, analysis, validation and the first draft were produced by AI agents under his direction, as described in the Acknowledgments.

Ethics statement

The study involved no human participants and collected no personal data for its analysis; it read pages that any visitor can read. The author has no institutional affiliation and no ethics review was sought. The crawler identified itself, linked to a public page describing the study and how to opt out, limited what it sent to each site, and followed the principles of the Menlo Report. Contact details found in security.txt files were never contacted, tested or released.

Data and code availability

The code, the frozen protocol, the lists of organizations the study started from, and a per-domain table of results are available at https://github.com/k0nnect/noroshi (snapshot of November 19, 2026; commit `328be22`). Stored copies of web pages, the log of model calls, which holds page text, and the contact details in security.txt files are withheld.

References

1. Brüggemann N, Schmidt L, Dölzer M, Brockhoff M, Ising F, Saatjohann C, Schinzel S. Measuring healthcare data leaks and security flaws at Internet scale. In: *2026 IEEE 11th European Symposium on Security and Privacy (EuroS&P)*. IEEE; 2026:1027-1045. doi:10.1109/EuroSP68448.2026.00069. Preprint: arXiv:2607.04965 (submitted July 6, 2026; v2 July 24, 2026).

2. Foudil E, Shafranovich Y. *RFC 9116: A File Format to Aid in Security Vulnerability Disclosure*. Internet Engineering Task Force; April 2022. Category: Informational. doi:10.17487/RFC9116. https://www.rfc-editor.org/rfc/rfc9116

3. Stock B, Pellegrino G, Rossow C, Johns M, Backes M. Hey, you have a problem: on the feasibility of large-scale web vulnerability notification. In: *25th USENIX Security Symposium (USENIX Security 16)*. USENIX Association; 2016:1015-1032.

4. Neprash H, McGlave C, Nikpay S. Hacked to pieces? The effects of ransomware attacks on hospitals and patients. *Am Econ J Econ Policy*. 2026;18(1):256-281. doi:10.1257/pol.20240594.

5. Dameff C, Tully J, Chan TC, et al. Ransomware attack associated with disruptions at adjacent emergency departments in the US. *JAMA Netw Open*. 2023;6(5):e2312270. doi:10.1001/jamanetworkopen.2023.12270. PMID 37155166.

6. Bugcrowd. *What Is a Vulnerability Disclosure Program (VDP)? Definition, Policy & Examples* [glossary page]. Bugcrowd. Undated page; accessed November 16, 2026. https://www.bugcrowd.com/glossary/vulnerability-disclosure-program-vdp/

7. Cybersecurity and Infrastructure Security Agency. *Binding Operational Directive 20-01: Develop and Publish a Vulnerability Disclosure Policy*. US Department of Homeland Security; September 2, 2020. https://www.cisa.gov/news-events/directives/bod-20-01-develop-and-publish-vulnerability-disclosure-policy

8. Tetelman A. *bounty-targets-data: data dumps of bug bounty platform scopes (HackerOne, Bugcrowd, Intigriti, YesWeHack)*. GitHub; commit a5734fe6bc9f9680b782e860fc994ecdf5103c12 (committed 2026-11-16T05:05:01Z). Accessed November 17, 2026. https://github.com/arkadiyt/bounty-targets-data

9. disclose.io. *disclose.io: the infrastructure that powers vulnerability disclosure and security reporting*. https://disclose.io/

10. CVE Program. CVE Numbering Authorities (CNAs); List of Partners. CVE.ORG; data file `src/assets/data/CNAsList.json` in the repository `CVEProject/cve-website`, commit 06790124bfe7a4b0b8eb2533c89a7652c429aca9 (committed 2026-11-11). Accessed November 17, 2026. https://www.cve.org/ProgramOrganization/CNAs ; https://www.cve.org/PartnerInformation/ListofPartners

11. CVE Program. *CVE Numbering Authority (CNA) Operational Rules*. Version 4.2.1. Effective October 20, 2026. MITRE Corporation for the Cybersecurity and Infrastructure Security Agency. https://www.cve.org/Resources/Roles/Cnas/CNA_Rules_v4.2.1.pdf

12. US Department of Homeland Security, Cybersecurity and Infrastructure Security Agency. *Agency Priority Goal Action Plan, FY 22 Q4: Strengthen Federal Cybersecurity*. Performance.gov; January 2023 (PDF metadata: created December 29, 2022, modified January 20, 2023; titled "Agency Priority Goal Update"; goal leader Matthew Hartman, CISA). https://assets.performance.gov/APG/files/2023/january/FY2022_January_DHS_Progress_Strengthen_Federal_Cybersecurity.pdf

13. Consolidated Appropriations Act, 2023, Pub L No. 117-328, div FF, title III, §3305, 136 Stat 5832 (2022). Enacted December 29, 2022. https://www.govinfo.gov/content/pkg/PLAW-117publ328/html/PLAW-117publ328.htm

14. US Food and Drug Administration. *Cybersecurity in Medical Devices: Quality Management System Considerations and Content of Premarket Submissions: Guidance for Industry and Food and Drug Administration Staff*. FDA; issued February 3, 2026. Docket FDA-2021-D-1158. https://www.fda.gov/media/119933/download

15. Householder AD, Wassermann G, Manion A, King C. *The CERT Guide to Coordinated Vulnerability Disclosure*. Special Report CMU/SEI-2017-SR-022. Software Engineering Institute, Carnegie Mellon University; August 2017. https://resources.sei.cmu.edu/asset_files/SpecialReport/2017_003_001_503340.pdf

16. US Department of Health and Human Services, 405(d) Task Group. *Health Industry Cybersecurity Practices: Managing Threats and Protecting Patients. Technical Volume 2: Cybersecurity Practices for Medium and Large Healthcare Organizations*. 2023 edition. HHS; 2023. https://hhscyber.hhs.gov/Documents/HICP/tech-vol2-508.pdf

17. US Department of Health and Human Services. *Healthcare and Public Health Sector Cybersecurity Performance Goals (HPH CPGs)*. HHS Cyber Gateway; published January 2024. https://hhscyber.hhs.gov/cybersecurity-performance-goals.html

18. US Department of Health and Human Services, Office for Civil Rights. HIPAA Security Rule To Strengthen the Cybersecurity of Electronic Protected Health Information; proposed rule. *Fed Regist*. 2025;90(3):898-1022. January 6, 2025. 45 CFR parts 160 and 164. FR Doc 2024-30983. https://www.govinfo.gov/content/pkg/FR-2025-01-06/html/2024-30983.htm

19. US Department of Health and Human Services. 45 CFR part 170: Health Information Technology Standards, Implementation Specifications, and Certification Criteria and Certification

Programs for Health Information Technology. Electronic Code of Federal Regulations. Current text as of November 13, 2026 (latest issue date available). https://www.ecfr.gov/current/title-45/subtitle-A/subchapter-D/part-170 . And: Assistant Secretary for Technology Policy/ONC. Health Data, Technology, and Interoperability: Certification Program Updates, Algorithm Transparency, and Information Sharing (HTI-1 final rule). *Fed Regist*. 2024;89(6):1192-1438 (January 9, 2024). https://www.govinfo.gov/content/pkg/FR-2024-01-09/pdf/2023-28857.pdf .

20. 45 CFR §180.50(d)(6). US Department of Health and Human Services. Requirements for making public hospital standard charges for all items and services. Electronic Code of Federal Regulations. Current text as of November 13, 2026 (version dated 2026-01-01 is the latest listed; last amended 90 FR 54087, Nov. 25, 2025). https://www.ecfr.gov/current/title-45/subtitle-A/subchapter-E/part-180/subpart-B/section-180.50

21. Centers for Medicare & Medicaid Services. *Hospital Price Transparency TXT File Frequently Asked Questions (FAQs)*. CMS; updated August 29, 2024. https://www.cms.gov/files/document/hospital-price-transparency-txt-file-frequently-asked-questions-faqs.pdf .

22. Mandel J. Disturbing state of EHR security vulnerability reporting. SMART Health IT blog. April 7, 2014 (updated April 10, 2014). https://smarthealthit.org/2014/04/ehr-security-vulnerability-reporting/

23. Siegl J. A critical line of defense: the security.txt file's potential to mitigate school cybersecurity attacks. Future of Privacy Forum blog. July 28, 2024. Last updated May 14, 2026. Accessed November 17, 2026. https://fpf.org/blog/a-critical-line-of-defense-the-security-txt-files-potential-to-mitigate-school-cybersecurity-attacks/

24. Medcrypt. *What the Medical Device Industry Can Learn From Past Cybersecurity Vulnerability Disclosures 2025*. Medcrypt; April 2025 (landing page dated April 2, 2025; PDF created April 1, 2025). Landing page: https://www.medcrypt.com/whitepapers/what-the-medical-device-industry-can-learn-from-past-cybersecurity-vulnerability-disclosures-2025 . PDF: https://cdn.ploy.ai/af51ae1b-90a1-4854-b62b-4e0c581b8f43/user/52f37b74-vulnerability-disclosures-2025.pdf

25. CVD Portal (Porta Regulus B.V.). CRA Exposure Study 2026: EU findings (Wave 2, September 2026). September 16, 2026; correction added November 4, 2026. CC BY 4.0. https://cvdportal.com/research/cra-exposure-2026

26. Kosiorek M. *AnnexProof CRA Readiness Index, round 2: security.txt adoption among 1,691 European device and software manufacturers (aggregate data, as of 30 September 2026)*. Version 2.0. Zenodo; October 3, 2026. doi:10.5281/zenodo.23182854. https://doi.org/10.5281/zenodo.23182854

27. Poteat T, Li F. Who you gonna call? An empirical evaluation of website security.txt deployment. In: *Proceedings of the 21st ACM Internet Measurement Conference (IMC '21)*. ACM; 2021:526-532. doi:10.1145/3487552.3487841

28. Findlay WP, Abdou A. Characterizing the adoption of security.txt files and their applications to vulnerability notification. In: *Workshop on Measurements, Attacks, and Defenses for the Web (MADWeb) 2022*, San Diego, CA, April 28, 2022. NDSS Symposium; 2022. doi:10.14722/madweb.2022.23014

29. Hilbig T, Geras T, Kupris E, Schreck T. security.txt revisited: analysis of prevalence and conformity in 2022. *Digit Threats Res Pract*. 2023;4(3):36. doi:10.1145/3609234

30. Alkinoon M, Alabduljabbar A, Althebeiti H, Jang R, Nyang D, Mohaisen D. Understanding the security and performance of the web presence of hospitals: a measurement study. In: *2023 32nd International Conference on Computer Communications and Networks (ICCCN)*. IEEE; 2023:1-10. Honolulu, HI, July 24-27, 2023. doi:10.1109/ICCCN58024.2023.10230186 . Preprint: arXiv:2304.13278 (April 26, 2023).

31. von Elm E, Altman DG, Egger M, Pocock SJ, Gøtzsche PC, Vandenbroucke JP; STROBE Initiative. The Strengthening the Reporting of Observational Studies in Epidemiology (STROBE) statement: guidelines for reporting observational studies. *PLoS Med*. 2007;4(10):e296. doi:10.1371/journal.pmed.0040296. PMID: 17941714.

32. Oak Ridge National Laboratory; US Department of Homeland Security, Geospatial Management Office. *Hospitals*. Homeland Infrastructure Foundation-Level Data (HIFLD Open and HIFLD Secure). Publication date 2024-06-30. Copy served by Pennsylvania Spatial Data Access (PASDA), Pennsylvania State University. Accessed November 17, 2026. https://mapservices.pasda.psu.edu/server/rest/services/pasda/HIFLD_FEMA/MapServer/51

33. US Food and Drug Administration. List of Artificial Intelligence-Enabled Medical Devices. In: *Artificial Intelligence-Enabled Medical Devices*. FDA; content current as of 09/04/2026 (page); list file downloaded November 16, 2026. Accessed November 17, 2026. https://www.fda.gov/medical-devices/artificial-intelligence-enabled-medical-devices/list-artificial-intelligence-enabled-medical-devices

34. Cybersecurity and Infrastructure Security Agency. *CISA CSAF Repository: CISA CSAF Security Advisories*. GitHub; commit 2701eb50ca2aa92c373bce21d4877f88ce1fd4cd (committed 2026-11-13). Accessed November 17, 2026. https://github.com/cisagov/CSAF

35. Office of the National Coordinator for Health Information Technology (ONC), US Department of Health and Human Services. Lantern: FHIR API Monitoring System. HealthIT.gov; page last updated April 1, 2026. Daily endpoint file `https://lantern.healthit.gov/api/daily/download`, retrieved November 16, 2026. Accessed November 17, 2026. https://healthit.gov/interoperability/investments/lantern-project/ Also: 45 CFR §170.404 (Application programming interfaces). Electronic Code of Federal Regulations. Accessed November 17, 2026. https://www.ecfr.gov/current/title-45/part-170/section-170.404 . The eCFR history line reads "[85 FR 25945, May 1, 2020, as amended at 85 FR 70084, Nov. 4, 2020; 89 FR 1433, Jan. 9, 2024; 90 FR 37211, Aug. 4, 2025]".

36. Anthropic. Models overview. Claude API Docs. Accessed November 17, 2026. https://platform.claude.com/docs/en/models/overview Also: Anthropic. Run Claude Code programmatically. Claude Code Docs. Accessed November 17, 2026. https://code.claude.com/docs/en/headless

37. Cybersecurity and Infrastructure Security Agency (CISA), US Department of Homeland Security. *dotgov-data: .gov Data* (`current-federal.csv`). GitHub repository. Accessed November 17, 2026. https://github.com/cisagov/dotgov-data

38. disclose.io. *diodb: the community database of every VDP & bug-bounty program and its safe-harbor status*. GitHub; branch master, commit c181d0d9a80bee0091bdb5f22536bcc377464aed (committed 2026-11-09). Accessed November 17, 2026. https://github.com/disclose/diodb

39. Clopper CJ, Pearson ES. The use of confidence or fiducial limits illustrated in the case of the binomial. *Biometrika*. 1934;26(4):404-413. doi:10.1093/biomet/26.4.404

40. McNemar Q. Note on the sampling error of the difference between correlated proportions or percentages. *Psychometrika*. 1947;12(2):153-157. doi:10.1007/BF02295996

41. Fisher RA. The logic of inductive inference. *J R Stat Soc*. 1935;98(1):39-82. doi:10.2307/2342435

42. US Department of Agriculture, Economic Research Service. *Rural-Urban Continuum Codes*. January 2024. Page updated December 30, 2025. Accessed November 17, 2026. https://www.ers.usda.gov/data-products/rural-urban-continuum-codes

43. Cohen J. A coefficient of agreement for nominal scales. *Educ Psychol Meas*. 1960;20(1):37-46. doi:10.1177/001316446002000104

44. Dittrich D, Kenneally E. *The Menlo Report: Ethical Principles Guiding Information and Communication Technology Research*. US Department of Homeland Security, Science and Technology Directorate; August 3, 2012. Accessed November 17, 2026. https://www.dhs.gov/sites/default/files/publications/CSD-MenloPrinciplesCORE-20120803_1.pdf

Tables and figure

Table 1. What happened to each listed domain, by sector.

Domain outcome	Hospitals	Device makers	Advisory vendors	EHR developers	Federal control	All lists (each domain once)
Listed domains	2,652	525	75	200	124	3,557
Measured, n (%)	1,721 (64.9)	436 (83.0)	67 (89.3)	181 (90.5)	101 (81.5)	2,487 (69.9)
Unreachable	615	49	6	15	20	705
... no DNS entry †	199	11	2	1	15	228
... TLS failure	174 (54)	17 (1)	1 (0)	9 (2)	2 (0)	203 (57)

(of which incomplete certificate chain) †						
... timeout or connection failure †	172	5	1	1	2	181
... challenge page with status 202 †	25	13	2	4	0	44
... other status or too many redirects †	45	3	0	0	1	49
Blocked (of which "too many requests")	151 (7)	14 (2)	1 (1)	1 (0)	3 (0)	170 (10)
Mismatch	165	26	1	3	0	195
... failed crawl test / identity "other" / identity "unclear"	5/96/64	3/7/16	0/1/0	0/1/2	0/0/0	8/105/82
Measured at another registrable domain	502	43	7	11	41	601
Measured, with a file classified forbidden	59	10	6	2	1	76
Measured, policy-page value unmeasured	90	34	4	17	14	157
Requests sent	21,338	4,958	785	1,937	860	29,681

Footnotes. Source: `results/tables/table1_flow.md` (hospital, device, icsma, ehr, federal and all rows), except the five rows marked †. † From the final audit, which classified the last homepage attempt of all 705 unreachable domains (Check 6). "Timeout or connection failure" adds the audit's timeout and connection-failure rows; "other status or too many redirects" adds its rows for more than 5 redirects and for status 404, 5xx, 307 and 402. A domain was measured when its homepage answered status 200 unrefused, was not a parked, for-sale or third-party platform page, and passed the identity check; unreachable, blocked and mismatched domains could not be measured and are not counted as lacking a channel. Percentages in the "Measured" row are computed from the two counts shown in this table (the layout asks for them); the file does not hold them. "All lists" counts each domain once; a domain on two lists (Methods) is counted in each of its sector columns (column `multi_sector`: hospital 2, federal 1, device 16, advisory vendors 15, EHR developers 4). A file was classified forbidden when both user agents were refused with a plain 401 or 403 while the homepage loaded; it counts as no

file. Of the unreachable hospital domains, 225 had a last archived homepage capture that redirected to another registrable domain (`unreachable_archive_redirect`). Domains that opted out: 0 (`optout`).

Table 2. Channels by sector, per listed domain (k/n, % and exact 95% confidence interval).

Outcome	Hospitals	Device makers	Advisory vendors	EHR developers	Federal control
Any channel (any-channel outcome)	132/1,618, 8.2% (6.9%-9.6%)	83/404, 20.5% (16.7%-24.8%)	25/61, 41.0% (28.6%-54.3%)	13/164, 7.9% (4.3%-13.2%)	78/89, 87.6% (79.0%-93.7%)
security.txt with a valid contact (main outcome)	132/1,682, 7.8% (6.6%-9.2%)	30/430, 7.0% (4.8%-9.8%)	7/61, 11.5% (4.7%-22.2%)	8/179, 4.5% (1.9%-8.6%)	4/99, 4.0% (1.1%-10.0%)
security.txt, strict outcome (valid contact and one unexpired, time-zoned Expires date)	32/1,682, 1.9% (1.3%-2.7%)	22/430, 5.1% (3.2%-7.6%)	6/61, 9.8% (3.7%-20.2%)	6/179, 3.4% (1.2%-7.2%)	0/99, 0.0% (0.0%-3.7%)
Policy page	114/1,631, 7.0% (5.8%-8.3%)	59/402, 14.7% (11.4%-18.5%)	17/63, 27.0% (16.6%-39.7%)	9/164, 5.5% (2.5%-10.2%)	76/87, 87.4% (78.5%-93.5%)
Listing (platform, disclose.io or CVE Numbering Authority)	1/1,721, 0.1% (0.0%-0.3%)	25/436, 5.7% (3.7%-8.3%)	18/67, 26.9% (16.8%-39.1%)	2/181, 1.1% (0.1%-3.9%)	15/101, 14.9% (8.6%-23.3%)
Any channel, per hospital (no interval)	175/2,364, 7.4%	n/a	n/a	n/a	n/a
Any channel, bed-weighted (no interval)	41,495/361,145 beds, 11.5%; 41 hospitals with unknown beds left out (2,323 hospitals)	n/a	n/a	n/a	n/a
Any channel, every unmeasured value counted as none	132/2,652, 5.0% (4.2%-5.9%)	83/525, 15.8% (12.8%-19.2%)	26/75, 34.7% (24.0%-46.5%)	14/200, 7.0% (3.9%-11.5%)	78/124, 62.9% (53.8%-71.4%)

Footnotes. Sources: rows 1 to 7, `results/tables/table2_shares.md` (level `domain`, measures `broad`, `main`, `strict`, `policy`, `listing`; level `hospital` and level `beds`, measure `broad`, for rows 6 and 7); row 8, `results/tables/tableS1_sensitivity.md` (level `domain`, measure `broad`). Denominators are the domains for which the outcome could be measured; they differ by outcome and are shown in each cell. The five lists are defined in the Methods. The interval is the exact (Clopper-Pearson) 95% confidence interval, with the domain as the unit; per-hospital and bed-weighted shares carry none. Row 8 counts every domain on

the list, with an outcome that could not be measured counted as no channel (a registry listing counts even for a blocked or unreachable domain). In the hospital column, the numerator of rows 1 and 2 is the same (132), because every hospital domain with a policy page or a listing also had a security.txt file.

Table 3. Hospitals: concentration of the security.txt result, and security.txt beside cms-hpt.txt.

Panel A. security.txt (main outcome) by unit of counting.

Unit	Positive	Denominator	% (95% CI)	Pre-specified or post hoc
Listed domain	132	1,682	7.8% (6.6%-9.2%)	pre-specified
Hospital	175	2,458	7.1% (no interval)	pre-specified
Beds	41,495	380,545	10.9% (no interval; 42 hospitals with unknown beds left out)	pre-specified
Distinct final site, measured domains	42	1,374	3.1% (2.2%-4.1%)	post hoc
Listed domain, without the 95 domains serving the most common file	37	1,587	2.3% (1.6%-3.2%)	post hoc
Distinct final site, every listed domain, unmeasured counted as none	42	2,340	1.8% (1.3%-2.4%)	pre-specified (sensitivity analysis)

Panel A (continued). Concentration of the result (post hoc).

Concentration measure	Value
Distinct files	19
Organizations serving those files †	17
Files on a single domain	13
Most common file: domains (share of the positive domains)	95 (72.0%)
Most common file: final sites	25
Most common file: hospitals †	120
Second most common file: domains	9

Panel B. The paired comparison on the 1,680 hospital domains where both files were measured.

	Valid cms-hpt.txt: yes	Valid cms-hpt.txt: no	Total
security.txt: yes	131	1	132
security.txt: no	1,398	150	1,548
Total	1,529	151	1,680

Exact McNemar P < .001. A cms-hpt.txt of any validity was on 1,537/1,680 domains, 91.5% (95% CI, 90.1%-92.8%).

Footnotes. Sources: Panel A rows 1 to 3, `results/tables/table2_shares.md` (hospital, measure `main`; level `domain`, `hospital`, `beds`); rows 4 and 5 and the concentration measures, `analysis/posthoc_concentration.json` (also `.csv`); row 6, `results/tables/tableS1_sensitivity.md` (hospital, level `final_site`, measure `main`); Panel B, `results/tables/table5_paired.md`. Pre-specified: the per-domain, per-hospital and bed-weighted shares (rows 1 to 3) and the sensitivity analysis per distinct final site (row 6). Post hoc, added after the first pass: the share among measured final sites (row 4), the share without the domains serving the most common file (row 5), and the counts of distinct files and of the organizations serving them (Panel A, continued). A distinct final site (Methods) counts once, as positive when any listed domain leading to it is. † From the audits, which read all 19 files in full (Check 2) and confirmed them on the final data. "Valid cms-hpt.txt" means a price-file link and a location name were present. The row and column totals of Panel B are sums of its cells. The exact McNemar P value is stored as 0 in `table5_paired.md` because it is smaller than the floating-point range (1 against 1,398 discordant pairs). The interval is the exact (Clopper-Pearson) 95% confidence interval; per-hospital and bed-weighted shares carry none. Panel B compares presence of each file on the listed domain, not compliance.

Table 4. Whether published channels are usable, by sector.

Measure	Hospitals	Device makers	Advisory vendors	EHR developers	Federal control
Domains with a security.txt (main outcome), n	132	30	7	8	4
Contact: email (mailto:)	130 (98.5%)	27 (90.0%)	6 (85.7%)	8 (100.0%)	0 (0.0%)
Contact: web (https://)	5 (3.8%)	4 (13.3%)	1 (14.3%)	1 (12.5%)	4 (100.0%)
Contact: telephone (tel:)	0 (0.0%)	1 (3.3%)	0 (0.0%)	1 (12.5%)	0 (0.0%)
Expires: valid date, unexpired	32 (24.2%)	22 (73.3%)	6 (85.7%)	6 (75.0%)	0 (0.0%)
Expires: expired	1 (0.8%)	6 (20.0%)	0 (0.0%)	0 (0.0%)	4 (100.0%)
Expires: nonconforming date	2 (1.5%)	0 (0.0%)	0 (0.0%)	1 (12.5%)	0 (0.0%)
Expires: missing	97 (73.5%)	2 (6.7%)	1 (14.3%)	1 (12.5%)	0 (0.0%)
PGP-signed	4 (3.0%)	5 (16.7%)	1 (14.3%)	1 (12.5%)	0 (0.0%)
Served as text/plain	132 (100.0%)	30 (100.0%)	7 (100.0%)	8 (100.0%)	4 (100.0%)
Second parser agrees on the contact	132 (100.0%)	30 (100.0%)	7 (100.0%)	8 (100.0%)	4 (100.0%)
Invalid files whose only contacts were	2/3, 66.7%	2/2, 100.0%	0/1, 0.0%	0/0	0/0

lenient or malformed, k/n					
Domains with a policy page, n	114	59	17	9	76
Policy commits not to pursue good-faith reporters, k/n, % (95% CI)	0/114, 0.0% (0.0%-3.2%)	21/59, 35.6% (23.6%-49.1%)	4/17, 23.5% (6.8%-49.9%)	3/9, 33.3% (7.5%-70.1%)	76/76, 100.0% (95.3%-100.0%)

Footnotes. Source: `results/tables/table4_quality.md`. Cells show k (%) among the domains with a security.txt file that met the main outcome (the first row), except the row for lenient or malformed contacts (k/n among domains whose security.txt was invalid) and the last two rows (domains with a policy page; the share whose policy promises not to pursue good-faith reporters, with the exact Clopper-Pearson 95% confidence interval). Intervals for the other rows are in `table4_quality.md` and are left out here because the denominators are 4 to 132. A file may name more than one kind of contact, so contact rows do not sum to the first row. No file had more than one Expires field (`expires_multiple`: 0 in every sector). Every file was served as text/plain (`text_plain`) and the second parser agreed on every contact (`parsers_agree`).

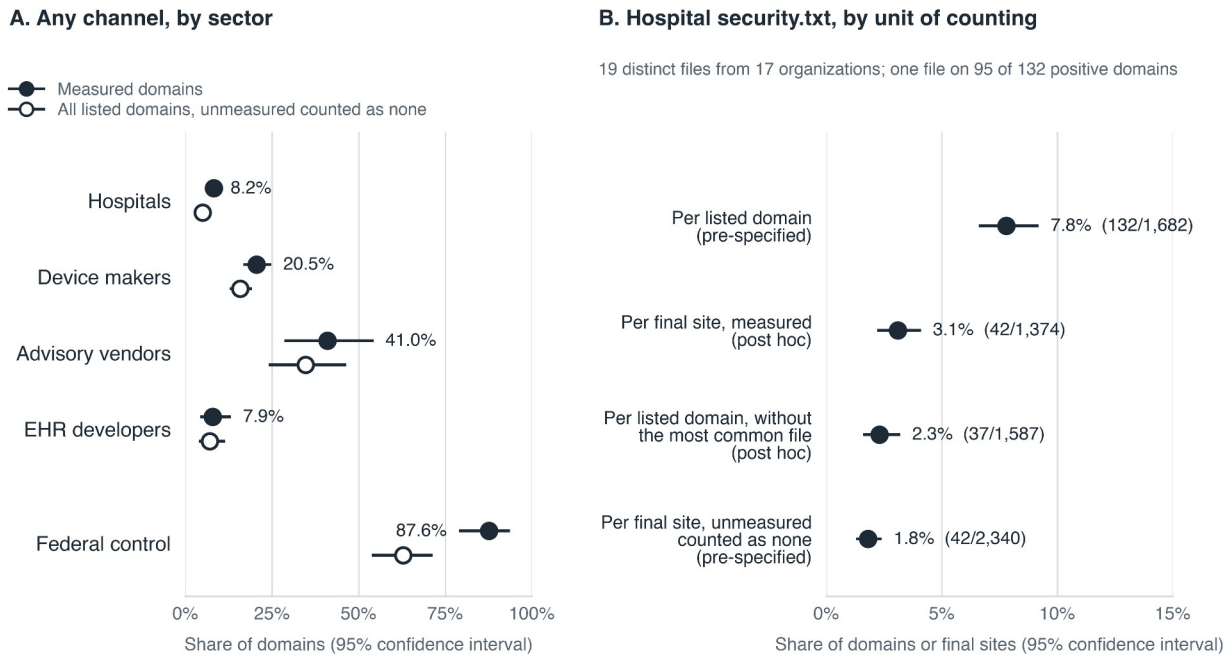

Figure 1. Any-channel shares by sector, and the hospital security.txt result by unit of counting. Panel A: for each list, the share of domains with any channel. Filled dot: the share among domains where the outcome could be measured (Table 2, row 1). Hollow dot: the share among all listed domains, with every unmeasured value counted as none (Table 2, row 8). Panel B: the share with a security.txt file that met the main outcome,

under four units of counting (Table 3, Panel A): per listed domain (7.8%, pre-specified), per distinct final site among measured domains (3.1%, post hoc), per listed domain without the domains serving the most common file (2.3%, post hoc), and per distinct final site with every unmeasured value counted as none (1.8%, pre-specified sensitivity analysis). Horizontal lines are exact (Clopper-Pearson) 95% confidence intervals. The count of organizations comes from the final audit.

Alt text for Figure 1: Two dot plots. In the first, the share of measured domains with any reporting channel is 8.2% for hospitals, 20.5% for device makers, 41.0% for advisory vendors, 7.9% for EHR developers and 87.6% for the federal control; counting unmeasured domains as having none lowers every share. In the second, the hospital security.txt share falls from 7.8% per listed domain to 3.1% per final site, 2.3% without the most common file and 1.8% per final site with unmeasured counted as none.

Supplementary material

Numbered in the order of first citation in the text (verify.py checks the order).

- Supplementary Methods (cited in Materials and Methods for "exact rules"; outline in `draft/02-methods.md`).
- Table S1. The federal control: any-channel share per listed domain and per distinct final site (`results/tables/tableS3_federal_control.md`), with the audit's reason for each negative, unknown and unmeasured domain (Check 7).
- Table S2. Agreement of the two judges and of the identity raters (`results/tables/tableS2_agreement.md`).
- Table S3. Sensitivity analysis: every unmeasured value counted as no channel, per domain and per distinct final site (`results/tables/tableS1_sensitivity.md`).
- Table S4. Hospital any-channel shares by system size, type, ownership, beds and rural location (`results/tables/table3_hospital_characteristics.md`, parts a to c).
- Table S5. Hospitals kept compared with hospitals dropped (`analysis/frame_coverage.md`).
- Table S6. Hospitals on measured domains compared with hospitals on domains that could not be measured (post hoc; `analysis/measured_coverage.md`).
- Figure S1. Cumulative first-seen dates of security.txt files in the Internet Archive (supplementary analysis; `results/figures/figure2_first_seen`).
- STROBE checklist (cited in Materials and Methods).
- Per-domain table of results, code, frozen protocol and frame tables (repository: https://github.com/k0nnect/noroshi).
- AI-use log, `logs/ai-use-log.md` (repository).