Supplementary Material

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

Contents: Supplementary Methods S1 to S13; Supplementary References; Supplementary Tables S1 to S6; Supplementary Figure S1.

Supplementary Methods

S1. Sampling frames and sources (protocol §4)

Unit. The unit of measurement was the registrable domain, the name directly below a public suffix (for example `example.org`); public and private suffixes both counted as suffixes, so `example.github.io` was its own domain. Domains were computed with tldextract 5.4.0 and the Public Suffix List it bundles (snapshot 2025-04-07, commit `5fc8d12c1624ddbc56a80e944aca151a8ee466a1`), private suffixes included and never updated over the network [1].

Freeze. All frames were fixed on November 16, 2026, before the crawler contacted any measured site, and the frame commands refused to run once the protocol was frozen on November 17, 2026. Two frame sources were sites of the federal control: the FDA list came from `fda.gov` with a browser user agent, because the site requires one, and the Lantern file from `healthit.gov`.

Hospitals

Source and erratum. The source was the Hospitals layer of HIFLD Open (Homeland Infrastructure Foundation-Level Data) as served by Pennsylvania Spatial Data Access (PASDA), `https://mapservices.pasda.psu.edu/server/rest/services/pasda/HIFLD_FEMA/MapServer/51` [2]. It was queried on November 16, 2026 at 05:23 UTC and returned 8,340 records (SHA-256 `ac85079aeaab24370458a45b13710aeebaad6c99b90c7140aa798eddc86a13af`). The layer's metadata gives a publication date of June 30, 2024 [2]. The protocol, frozen before pass 1, called this source "the final federal release" of HIFLD Open Hospitals; no page read for the reference check supports that wording [2], and this supplement replaces it with the description above. The Department of Homeland Security's HIFLD page (last updated September 4, 2025) describes only HIFLD Secure, and the former HIFLD Open site no longer loads; the NAPSG Foundation (a notice dated

June 24, 2025, updated August 5, 2025, read in an archived copy) and the Data Rescue Project (December 15, 2025) report that HIFLD Open went offline in August 2025. No official notice of the shutdown was read.

Exclusions, applied in this order. Of 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 (43.1% of all 8,340 records).

Step	Rule	Hospitals removed	Hospitals left
0	HIFLD records		8,340
1	Status is not open	374	7,966
2	Outside the 50 states and the District of Columbia	85	7,881
3	Federal (owner is the federal government, or type is military)	350	7,531
4	No website listed (blank, or `NOT AVAILABLE`)	3,597	3,934
5	Website value cannot be parsed	8	3,926
6	Website belongs to someone other than the hospital or its health system	233	3,693

The 3,693 hospitals kept were on 2,652 domains. At step 6, a listed website was excluded when its domain was a government site that is not the hospital's own (federal portals 3; state portals and state agency sites 105; county or city portals 20), a directory or association (100), a site builder (1) or another third party (4).

Review and who decided. Every kept hospital whose listed URL had a path and whose domain held none of its name words was reviewed: 388 hospitals on 223 domains. The decisions are in `data/frames/shared_platform_review.csv`. Eight domains (nine hospitals) could not be decided and were kept. One Claude coding agent (Claude Sonnet 5.5) made these decisions and built the third-party list (`src/noroshi/domains.py`), from the hospital's name, the listed address and the domain only, with no request to any website. The list came from the 60 most frequent listed domains, every kept `.gov`, `.us`, `.mil` and `.edu` domain and the 388 listings; three state agency sites were added later because an agency's own site is a government portal. No second reviewer checked the individual decisions.

Characteristics. Type, ownership, beds, state and county came from the HIFLD record. Rural meant a 2023 Rural-Urban Continuum Code of 4 or higher for the county; ERS defines codes 4 to 9 as nonmetropolitan [3], and "rural" is the study's label for them (file downloaded November 16, 2026 at 05:07 UTC, SHA-256 `ec455ee2a8bc5fc8e070575ea5bee7dce46fc6037f8c3449cbf56e8b453`

31fa7). System size was the number of frame hospitals sharing a domain. Supplementary Table S5 compares the 3,693 hospitals kept with the 3,838 eligible hospitals dropped at steps 4 to 6.

Federal control

The source was current-federal.csv of cisagov/dotgov-data [4], read from the main branch on November 16, 2026 (SHA-256 e8f0ff9d5c43b4c15e95591f0cdb84a377c0d6b0b5e038dc80ae59bb7b7b2b0a; 1,320 rows; no commit was recorded [4]). Rows with domain type "Federal - Executive" and organization "Department of Health and Human Services" or "Department of Veterans Affairs" were kept: 124 domains. A binding directive requires each federal executive-branch agency to publish a vulnerability disclosure policy at the /vulnerability-disclosure-policy path of its primary .gov website [5]. The group validated the detector; it is not a study population.

Device makers, advisory vendors and EHR developers

The FDA list of AI-enabled medical devices was downloaded on November 16, 2026 at 05:14 UTC from https://www.fda.gov/media/178541/download?attachment (SHA-256 b1efeb680371f44cae574ae5a5aa548ac8a63391e6888a007b17177d17527880); the page read "Content current as of: 09/04/2026", and the file has no version date [6]. It gave 616 company names after merging 816 spellings that differed only in case, punctuation or a trailing legal suffix (a recount gave 817 [6]; not investigated). FDA states that the list "is not a comprehensive resource of AI-enabled medical devices" [6]. The separate advisory-vendor group was every vendor named in a CISA medical advisory, from the CSAF repository at commit 2701eb50ca2aa92c373bce21d4877f88ce1fd4cd (191 advisories, 82 vendor names) [7]; an advisory names the vendor of an affected product, not necessarily a finished-device maker [7]. The EHR group was the distinct developer names (214) in ONC's Lantern daily endpoint file, https://lantern.healthit.gov/api/daily/download, retrieved November 16, 2026 at 05:14 UTC (SHA-256 a0827b2e3653c07cbc00d62ecdf3ab6c1318025e5ad6645133d8532d70f25cbf) [8]. Lantern queries endpoint lists that certified developers publish, and the rule requiring them covers modules certified to 45 CFR 170.315(g)(10) [8]. The Certified Health IT Product List [9] was not used, so the group is the developers Lantern lists, not every EHR developer.

Resolution of company websites

The company sources give names without websites. Two Claude models, claude-sonnet-5-5 and claude-opus-5-5 [10], were each asked separately, without tools

and without seeing the other's answer, for the registrable domain of each company's official corporate website (25 names per call, a one-line sector description, null when unsure). Where both gave the same registrable domain it was accepted. Otherwise one call per name went to `claude-sonnet-5-5` with a web-search tool, asking for the official website and an evidence URL; a non-null answer was accepted and a null left the company unresolved. A failed or unreadable call was tried once more, and every answer was cut to its registrable domain. Full prompts: `src/noroshi/resolve.py`.

Group	Names	Both models agreed	Web search	Unresolved	Unique domains
Device (FDA list)	616	390	213	13	525
Device (ICSMA)	82	66	14	2	75
EHR	214	108	96	10	200

The models differed, or one gave a null, for 348 names (323 resolved by the search, 25 left unresolved and in no crawl table). One logged answer was corrected outside the pipeline: both models answered `pydicom.github.io` for one advisory vendor, and the stored value had been cut to `github.io` because the suffix list then in use did not treat `github.io` as a suffix (`data/frames/resolution_passes/CORRECTIONS.md`). No human or second reviewer checked these decisions and no accuracy estimate exists; the identity check of S5 is the only later check.

Totals and start URLs

No table could hold a suffix, a domain that is not its own registrable domain, a domain-sale or parking service, or a third-party platform; federal portals (`va.gov`, `medicare.gov`, `ihs.gov`, `tricare.mil`, `navy.mil`) counted as third parties for hospitals only. The five tables held 3,576 domain rows (2,652 hospital, 124 federal, 525 device, 75 advisory-vendor, 200 EHR) on 3,557 unique domains; 19 domains were in two tables (14 in both device tables, and one each in hospital and EHR, hospital and device, federal and EHR, device and EHR, advisory vendor and EHR). Such a domain was crawled once, under the rules of the first of its sectors to run (hospitals, federal, device, advisory vendors, EHR), and reported in each. The homepage was requested at `https://<host>/`: for hospitals the host listed most often among the domain's hospitals (ties to the shorter name, then alphabetical order), for companies and federal domains the bare registrable domain; plain http was never a start URL.

S2. Crawler conduct (protocol §5)

Only public files and pages were requested, with `GET`; no login, form, query manipulation, script or vulnerability test was used, and no contact found in a file was ever tested or released. Requests carried the user agent `noroshi-research/1.0`

(+https://github.com/k0nnect/noroshi); the linked page describes the study and how to opt out. No request was sent unless the file info/LIVE named the information page's URL, and one crawl, judge, identity or purge command ran on a store at a time. Both passes ran from one computer in the United States over a single internet connection; no proxy and no other vantage point was used, so a site that answers differently by location or network was seen only from there.

Limits and pacing. A "site" was a registrable domain. In one pass at most 12 requests were sent for any listed domain (retries and crash repeats included), and at most 12 non-homepage requests were made to any one site across all listed domains, counted at the site of the address requested; homepage requests (start URL, other host form, meta-refresh target and their browser retries) were outside the site limit but inside the listed domain's 12. A request followed at most 5 redirects (not counted), had 60 seconds in total, 10 to connect and 20 without data, and kept at most 4,194,304 bytes of a homepage and 1,048,576 bytes of any other body. Requests to one site went one at a time, at least 2.0 seconds apart, redirect hops included. Up to 20 domains were crawled at once, with HTTPX 0.28.1 [11] and TLS certificates verified. An address already chosen or answered for the same listed domain, or already measured for any listed domain in the pass, was not requested again (a stored measurement was copied and marked as reused). Every request was written to the store before it was sent, so a request unfinished by a crash still counted toward the limits.

Refusals and the browser retry. A refusal was a final status of 401, 403 or 406; a cf-mitigated or x-amzn-waf-action header; or a challenge page (a title beginning Just a moment..., one of seven fixed strings in the first 32 KiB of the body listed in src/noroshi/blocking.py, or Incapsula's script on a status 200 page with fewer than 200 visible characters). An answer with status 429 was stored as refused as well, so it was never read as "not there"; the 429 rule below says what followed. When a site refused the identified agent, the request was sent once more with the user agent of Safari 18.1 on macOS (Mozilla/5.0 (Macintosh; Intel Mac OS X 10_15_7) AppleWebKit/605.1.15 (KHTML, like Gecko) Version/18.1 Safari/605.1.15; no other header changed), unless a limit had been reached or the site had answered 429. In the two pilots 5 homepages answered the identified request with the same 52-byte 403 - Forbidden page and all 5 answered 200 to the browser retry. When the retry got an answer, the refusing site was switched to the browser agent for the rest of the pass; otherwise it stayed on the identified agent. The agent was chosen per redirect hop, so it was never sent to a site that had not refused, and other requests to a refusing site waited, up to 60 seconds, for the retry (protocol §5). A 429 was never retried and nothing more was sent to that site in the pass; a listed domain whose own site, start address's site, homepage's final site or homepage refresh target's site

answered 429 ended the pass as blocked, and a 429 from any other site only skipped that page.

robots.txt. The homepage was requested first without regard to robots.txt. robots.txt was then read from the host the homepage ended on (agent name `noroshi-research`, Protego 0.7.0 [12]) and governed the `Policy:` URL, candidate pages, conventional paths and refresh targets on that host: a 200 answer was obeyed, a 4xx that was not a refusal allowed everything, and anything else (a 5xx, a network error, a refusal, or no completed record) forbade every page on that host. The last rule is the study's own. `security.txt` and `cms-hpt.txt` were requested regardless, and a page on another host was requested without reading that host's robots.txt.

Opt-out, outages and order. An organization opts out by opening an issue on the repository named in the user agent, with its domain; no run started afterwards sends it anything, and everything collected from it is deleted and left out of the findings (`info/README.md`; purge rules in protocol §5). No opt-out request was received; the repository's issue list and the opt-out file were checked before each pass and after the last one. A run stopped when 20 domains in a row ended their homepage step with a network error and a control request to the information page also failed; those domains were crawled again and not recorded as unreachable (protocol §5). With all sectors in one run the order was hospitals, federal, device (FDA list), device (advisory vendors), EHR.

S3. Request order per domain (protocol §6.1)

For each listed domain the crawler requested, in order: (1) the homepage at the start URL, with one request to the other `www.` form after a network error and one followed meta refresh on a page with fewer than 200 visible characters and a delay of at most 5 seconds; (2) `/robots.txt`; (3) `/.well-known/security.txt`; (4) `/security.txt`, unless the well-known file was valid or valid with problems (RFC 9116 places the file under `/.well-known/` and allows the top-level path for legacy compatibility [13]); (5) `/cms-hpt.txt`, hospitals only; (6) the first `Policy:` URL of a security.txt read as text, when an http or https address, not on a disclosure platform and allowed by robots.txt; (7) up to three candidate links from the homepage (two when step 6 applied), the specific tier first and the general tier second, in document order; and (8) up to six conventional paths, in order: `/vulnerability-disclosure-policy`, `/responsible-disclosure`, `/security`, `/vulnerability-disclosure`, `/product-security`, `/psirt`. Steps 2 to 5 and 8 went to the scheme and host of the homepage's final URL. A `Policy:`, candidate or conventional-path page that answered with a short meta refresh was followed once; no other link on such a page was followed. When a limit was reached the remaining steps were not made, so the conventional paths were the steps most often cut. Links to HackerOne, Bugcrowd, Intigriti and YesWeHack were never requested. The wording and address patterns of the

two candidate tiers, the resolution of links against `<base href>` and the page pattern list are in protocol §6.1 and `src/noroshi/discover.py`.

S4. Passes, combination and collection window (protocol §6.2)

Pass 1 covered every one of the 3,557 listed domains (3,576 sector rows) on November 17, 2026, from 07:37 to 08:45 UTC (first request 07:37:45, last 08:45:27). It sent 26,746 requests; the store held 30,040 request records, 3,294 of which were reused copies of answers (S2).

Pass 2 crawled a domain again in full (S3) under fresh limits, no earlier than 48 hours after the first request recorded for it in pass 1, and started no earlier than 48 hours after pass 1 ended, so that no site was requested again within 48 hours of any pass-1 request to it. There was one retry pass and no third. It covered only domains that were blocked or unreachable; or whose homepage, `security.txt` or `cms-hpt.txt` request could not be completed or was never made (for example because a limit was reached first; a forbidden file counted as completed); or whose look for a policy page could not be completed because robots.txt could not be read or a page request failed. Pass 2 ran on November 19, 2026, from 08:50:14 to 09:01:52 UTC, 48 hours and 5 minutes after the last pass-1 request. It covered 1,049 listed domains and sent 2,935 requests (3,087 stored records, counting copies of answers reused for other listed domains); one of the 1,049 ended blocked with nothing sent to it, because its site had already answered 429 to a request made for another listed domain. Both passes together sent 29,681 requests (33,127 stored records). The window of each pass was read from the store with a separate query.

A request counted as a measurement only when it completed, was not refused, and ended with status 200 (the thing was read) or a 4xx that was not a refusal (the thing is not there); any other ending could not be measured.

Combining the passes. The policy-page look was judged on both passes together, over the passes that measured the domain: a skipped, disallowed or failed page was cleared when any record in either pass measured that address, robots.txt was unreadable only when no pass read it, and the homepage counted as cut or unreadable only when every pass found it so. For each file (`security.txt` at each location, `cms-hpt.txt`) the most recent result across both passes was used, a result being a measurement or a forbidden file; a later failed attempt never replaced an earlier result, and a location found valid, valid with problems or (for `cms-hpt.txt`) present and valid in either pass stood even if the other pass found it absent. File records and policy pages counted from either pass, including one that ended blocked after reading them. The homepage was the last outcome homepage request of the latest pass that measured the domain; its final site, links and listings were read from it, never from a pass that ended blocked, unreachable or mismatch. A domain was opted out when either pass ended `optout`, otherwise

measured when either pass measured it, otherwise it had the status of the latest pass. The reported results were produced with no incomplete-run override (protocol §6.2).

No amendment between the passes. The protocol was not amended after pass 1. The audit of pass 1 (S12) found rules that misfired; these are reported as limits and not corrected, and pass 2 ran under the same frozen rules.

Internet Archive lookups. Lookups used the Internet Archive's CDX server [14], `http://web.archive.org/cdx/search/cdx`, output as JSON. The first-seen lookup (S10) asked for the earliest capture of `host` + `path` with `limit=1` and filters `statuscode:200` and `mimetype:text/plain`. The redirect lookup (S5) asked for the most recent capture of the start host (`limit=-1`); when its status began with 3, the replay address `http://web.archive.org/web/<timestamp>id_/http://<host>/` was requested without following redirects, and the registrable domain of the `Location` target was kept when it differed from the host's. An empty reply meant no capture; any other failure was not written, so a rerun retried it.

Change after the freeze. After the protocol was frozen, the pacing of these lookups was slowed from one request per second to one per 4.5 seconds. The code was changed on November 17, 2026 (commits `02bc094` and `94a4e53`): requests to the archive are at least 4.5 seconds apart (a redirect lookup's replay request included), and the run stops, keeping everything stored, at the first refused connection or after three failed lookups in a row. No measurement rule changed. Lookups left incomplete: 1 of the 178 first-seen lookups (the archive answered it with status 403) and none of the 615 redirect lookups. A domain with an incomplete lookup is not counted as having no date or no redirect.

S5. Domain status, identity check and model calls (protocol §6.3, §11)

Status. A domain was *measured* when its homepage answered 200, was not refused, passed the three crawl tests, no site of the domain answered 429 before its crawl ended, and the identity check found the page to be the organization's working website (domains listed only in the federal table were exempt). It was *blocked* when the homepage was refused (after the browser retry, when sent) or a site of the domain answered 429; *unreachable* when the homepage failed with a network error at the start URL and, when tried, the other host form, failed with more than 5 redirects, or ended with a status that was neither 200 nor a refusal; *mismatch* when it loaded but was not the organization's working website; and *optout* for an opted-out domain. Blocked, unreachable and mismatch domains were left out of the primary shares (S11). The three crawl tests ended a domain as mismatch when its final site was a domain-sale, parking or third-party platform (S1); when its title or visible text carried one of six for-sale or parked markers (`domain is for sale`, `this domain is for sale`, `buy this domain`, `domain for sale`, `this domain is parked`, `domain parking`); or when a

homepage request ended on a host `ww<digits>.<listed domain>`. A homepage that ended on another registrable domain was measured there and flagged. Hospital domains that ended unreachable were looked up in the Internet Archive (S4); an archived redirect to another registrable domain was recorded, not followed.

Identity check. After the crawl, `claude-sonnet-5-5` and `claude-haiku-5-5`, each without the other's answer and without tools, decided for every listed domain that a pass measured (once, even when the domain is in several tables) whether the stored homepage is the organization's own working website; `claude-opus-5-5` decided alone when they differed [10]. The system prompt was `You check web pages for a research study. Answer only with a JSON object.` The prompt gave the listed names (for hospitals up to five, largest by beds first), the domain, the host the homepage ended on, the title, two meta tags (each cut to 300 characters), the first 2,000 visible characters and the last 800 not already included. Domains listed only in the federal table were not checked, because the .gov registry controls those names, and opted-out domains were not checked. The full prompt, with the slots filled in, was:

``` You are checking one web page for a research study. The study's list says the website of the organization below is at the listed domain. Websites change hands, so the page that loaded may belong to someone else.

Organization on the list: {names} {place_line} Listed domain: {domain} Host the homepage ended on: {final_host}

Page title: {title} Site name and description from the page's meta tags: {meta} Start of the page text: <<< {head} >>> End of the page text: <<< {tail} >>>

Whose website is this page? Choose one.

- "own": the listed organization's own working website, or the working website of the organization that now owns or operates it or answers for its products (a parent company, an acquirer, a successor, a health system). A site in another language still counts when it is the same organization. A page that shows little text of its own, for example one that needs scripts to display or a sign-in page, still counts when its title or meta tags name that organization or its product.
- "other": anything else: a different or unrelated organization or business, a parked or for-sale domain, a placeholder or hosting default page, an error page, a bot check or access-denied page, or a site reusing the name for unrelated content.
- "unclear": nothing on the page (its title, its meta tags, its text) shows whose website it is.
```

Reply with one JSON object and nothing else. It has exactly two keys: "site", whose value is one of the three words above, and "reason", one short sentence naming what on the page decided it. ```

A rater whose call failed, or whose reply could not be read as one of the three words, was asked once more. Two unreadable replies counted as `unclear` (marked `reply_unreadable`), and a tie-break counted that way decided `unclear`. A failed call left the domain undecided until a later run. A decision held only for the view the models were shown; a different view in pass 2 was asked again. `other` or `unclear` made the domain a mismatch, and the analysis waited for every decision (protocol §6.3).

Model calls. Every call went through the Claude Code command-line client in non-interactive mode (`claude -p` [10]; `src/noroshi/llm.py`) with the model named, a fixed system prompt, no tools (the web-search call of S1 excepted), no settings files or saved session, and an empty working directory. No temperature or token limit was set, so the client's defaults applied, and a call was cut off after 300 seconds. Each model identifier is a pinned snapshot [10]. The models for each task are given in S1, S5 and S8. Every call was written to `data/raw/llm_log.jsonl` with purpose, model, system prompt, prompt, reply, status, start time and duration; for the policy judge the page address went to a separate log field and never to the model, whereas the identity prompt, quoted in this section, includes the listed domain and the host the homepage ended on.

S6. security.txt classification (protocol §6.4)

Each location was classified in this order: *forbidden* (the identified attempt and its browser retry at the same address both ended 401 or 403, both answers were plain, meaning no challenge page and no `<script` in the body, and the homepage answered 200 in the same pass); *unmeasured* (refused, failed, or a final status that was neither 200 nor a 4xx); *absent* (a 4xx that was not a refusal); *HTML error page*, counted as absent (status 200 with an HTML body); *invalid* (a 200 text with no valid contact, or not served over HTTPS); *valid with problems* (a valid contact but no single, unexpired, time-zoned `Expires`); and *valid*. A file on a site already switched to the browser agent was never forbidden. Expiry was judged at request time. A domain took its better location (valid over valid with problems, the well-known location breaking a tie); otherwise it was unmeasured if either location was, then invalid, HTML error, forbidden, absent. Forbidden counted as no file. A valid legacy file cannot win over a valid-with-problems well-known file, because the legacy file is then not requested. A contact was valid when it was `mailto:` with an address that has a dotted domain and no whitespace, `https://` with a dotted host and no whitespace, or `tel:` with a digit; a bare email address or an `http:` URL was recorded as lenient and did not count. `Expires` had to be one date-time `YYYY-MM-DDThh:mm:ss` (optional fractions) with `Z` or `±hh:mm`, later than the request time; a missing offset is invalid, as in RFC 3339. PGP-signed files

were parsed inside the wrapper, with signatures unverified. RFC 9116 requires `Contact` and `Expires`, forbids repeating `Expires`, and requires https and `text/plain` [13]; the main outcome therefore counts files the RFC would not call conformant (for example without `Expires`), and the strict outcome requires a valid `Expires`. The `sectxt` library 0.9.8 [15] was compared on one point, whether a chosen file has a valid contact (`results/parser_disagreements.csv`); the study's own parser decided every class.

Where the hospital files were found (final audit). Of the 132 hospital domains with the main outcome, 130 were found by the well-known request and 2 by the root request. A root request followed only when the well-known file did not parse (S3), so a file present at both locations could not be observed. Of the 19 distinct hospital files, 3 had no `Expires` field; they served 97 of the 132 domains.

S7. cms-hpt.txt classification (protocol §6.5)

A `cms-hpt.txt` was *present and valid* when read as text (not refused, status 200, not HTML) with at least one `mrf-url:` line beginning `http://` or `https://` and one `location-name:` line with a value, *present but invalid* when a text file lacked either, and *forbidden* under the rule of S6; absent, HTML error and unmeasured followed S6. A file on another registrable domain than the final site was flagged. The rule requires the file in the root of the website that hosts the machine-readable price file [16], which may be a health system's [17], so absence on a hospital's own domain is not proof of noncompliance. The share of domains with any `cms-hpt.txt` is given beside the McNemar test (S11); an invalid file counts as no file in the test.

S8. Policy pages (protocol §6.6)

Pages examined and set aside. Every stored candidate, conventional-path and `Policy:` page with status 200 was examined. Before any pattern list or model, a page was set aside in this order: `same_as_homepage` (a candidate or conventional-path page whose final address was the homepage's); `pdf_unreadable` or `non_text` (a PDF cut at the size limit or yielding no text with `pdftotext` (Poppler 26.09.0, the version installed on the study machine since September 2026) [18], or a non-text body), with `image` for an image answer; `meta_refresh` (a followed refresh whose target stands in for it); `unreadable` (fewer than 200 visible characters, or cut at the size limit with nothing matching the page pattern list or mentioning `security.txt`); and `catch_all` (the same visible text at two conventional paths with different final addresses). Catch-all, refresh, image and same-as-homepage pages counted as read and not policy pages; the others counted as unread. A page matching none of the page pattern list (the specific tier's wording, protocol §6.1) and not mentioning `security.txt` was a clear negative and not sent to a model.

Judging. Every other page was judged independently by `claude-opus-5-5` and `claude-sonnet-5-5` [10], without tools, with the system prompt `You classify web pages for a research study. Answer only with a JSON object.` and this prompt:

``` You are classifying one web page for a research study. The page was fetched from an organization's website.

Page text: <<< {excerpt} >>>

Answer three questions about what this page itself says.

1. invites_reports: Does this page tell people outside the organization how to report a security vulnerability in the organization's own websites, systems or products? A page that points reporters to a reporting address or to a disclosure platform counts. Answer false for a page whose only relevant content is a navigation or footer link to a policy that lives elsewhere, for a general home, navigation or error page, and for pages about patient privacy complaints, compliance or fraud hotlines, phishing emails someone received, or physical security.

2. contact_method: Does this page give a way to send such a report: an email address, a web form, a phone number, or a link to a disclosure platform? Answer false when your answer to question 1 is false.

3. safe_harbor: Does this page commit not to pursue legal action against people who report in good faith? Answer false when your answer to question 1 is false.

Reply with one JSON object and nothing else. It has exactly three keys, "invites_reports", "contact_method" and "safe_harbor", and each value is true or false according to your answers. ```

The prompt has no example answer and no page address (the address went to the log only). The models saw the visible text whole up to 40,000 characters; for a longer page, the first 20,000 characters plus 1,000 characters around every pattern match, email address, platform host name and `security.txt` mention, up to 40,000 in all. A failed or unreadable call was retried once, and pages with identical visible text were judged once. A disagreement on question 1 or 2 was settled by `claude-fable-5-2` and the majority decided; a disagreement on question 3 was settled the same way only for a policy page, and question 3 was no for any other page. A policy page was a page whose final answers to questions 1 and 2 were yes. Both judges' answers, the tie-break and the agreement flag were stored for every judged page.

**Value of the measure.** The domain's value was *yes* when any judged page was a policy page (either pass); *unmeasured* when none was found and the look was incomplete; and *no* when the look was complete and found none. The look was incomplete when robots.txt could not be read or disallowed every conventional path; a candidate,
```

`Policy:` or refresh-target page was disallowed, never requested (a limit reached, an opted-out site or a 429) or ended without a measurement; a stored page could not be read or was not judged; or the homepage was cut at the size limit or had fewer than 200 visible characters so that links could not be read. The reason was recorded (nine reasons, `src/noroshi/look.py`). Conventional paths cut by a limit did not make the look incomplete, and those domains were counted separately.

S9. Outside listings (protocol §6.7)

Listings were matched by registrable domain against snapshots downloaded on November 16, 2026 at 05:34 UTC and pinned in `data/reference/registries/registries.meta.json`.

Source	Repository and commit	Entries
disclose.io database [19]	`disclose/diodb`, `c181d0d9a80bee0091bdb5f22536bcc377464aed`	2,433
HackerOne, Bugcrowd, Intigriti, YesWeHack scopes [20]	`arkadiyt/bounty-targets-data`, `a5734fe6bc9f9680b782e860fc994ecdf5103c12`	450, 296, 140, 52
CVE Numbering Authority list [21]	`CVEProject/cve-website`, `06790124bfe7a4b0b8eb2533c89a7652c429aca9`, `src/assets/data/CNAsList.json`	567

The file hashes are in the metadata file. The platforms' own directories were not read; a third-party compilation whose README claims no affiliation with them was used [20]. A disclose.io entry is one submitted program, and the database is not a census [19]; its LICENSE file gives CC0 1.0 while its README licence section names CC BY 4.0, which was not resolved [19]. The CNA file is the data file from which the CVE Program website draws its List of Partners page; it held 567 entries at the pinned commit [21]. An entry listed the registrable domains of its policy URL, security.txt URL and contact email (disclose.io); of its in-scope web-like assets (platforms); or of every web or email address in its contact, disclosure-policy and advisory fields (CNA). A domain was listed when its listed domain or, for a measured domain, its final site's registrable domain matched. The rule for a platform link on a homepage or in a `Policy:` URL is in protocol §6.7.

S10. First-seen dates (supplementary analysis; protocol §6.8)

For domains with the main outcome, the earliest Internet Archive capture of the file with status 200 and a plain-text type was looked up (S4) [14]. It is reported apart from the primary results. A first capture shows only that the file existed by that date, not when it was published.

Result. Of the 178 distinct domains with the main outcome, the lookup was completed for 177 (S4). A capture of the file was found for 145 domains and none for 32. The earliest capture dated from 2019 for 1 domain, 2022 for 2, 2023 for 3, 2024 for 18, 2025 for 20 and 2026 for 101 (Supplementary Figure S1; counts from `analysis/first_seen_summary.py`). Of the 132 hospital domains, 117 had a capture, and for 93 of them the earliest was in 2026. Listed domains that lead to one final site share its file location and so its date.

S11. Statistical details

Shares and tests. Shares carried exact Clopper-Pearson 95% confidence intervals [22] from beta quantiles (`scipy.stats.beta.ppf`; SciPy 1.18.1 [23]), with the domain as the unit of inference. Hospital shares were also given per hospital and bed-weighted, without intervals, because hospitals on one domain share an outcome; a bed-weighted share counted each hospital by its beds and left out hospitals whose beds were not a positive number in the HIFLD record, with the number left out given beside it. The paired comparison of `cms-hpt.txt` (present and valid) and `security.txt` (main outcome) used hospital domains where both were measured (a forbidden file counting as measured and as no file) and the exact McNemar test [24]: two-sided `p = min(1, 2 × F(min(b, c); b + c, 0.5))`, with `b` and `c` the discordant counts and `F` the binomial distribution function (`scipy.stats.binom.cdf`), and 1 with no discordant domains. Single-hospital against multi-hospital domains were compared with `scipy.stats.fisher_exact`, two-sided [25].

Hospital characteristics (Supplementary Table S4; descriptive, no regression model). (a) System size over all hospital domains with exact intervals, grouped 1, 2 to 5, 6 to 20 and more than 20 hospitals. (b) Type, ownership, beds and rural location among single-hospital domains with exact intervals. (c) Per-hospital shares by characteristic, without intervals. The table used the broad outcome over measured hospital domains whose broad value was known. The groups are in protocol §8.

Sensitivity table (Supplementary Table S3; released as `tableS1_sensitivity`). The shares of the main share table (main, strict, policy page, listing and broad, per domain by sector, and for hospitals per hospital and bed-weighted) were repeated over every domain in the frame except opted-out and uncrawled ones, counting any value that could not be measured as lacking a channel. A registry listing of the listed domain still counted for a blocked or unreachable domain, but not for a mismatch, because a listing needs no crawl. Shares were also given counting each distinct final site once (positive when any listed domain leading to it was). The characteristics, quality, paired and federal tables were not repeated.

Quality and flow tables. The quality table (contact types, `Expires` state, signed files, `text/plain` service, parser agreement, and the share of policy pages that commit to no

legal action against good-faith reporters, with an exact interval per sector) and the flow table are defined in protocol §8. The flow table's all-sectors row counts each domain once.

Agreement table (Supplementary Table S2; `tableS2_agreement`). Before any tie-break and with opted-out domains left out, the number of units, the share on which the two raters agreed and Cohen's kappa [26] for each of the judges' three questions (over the distinct page texts both judges answered) and for the identity raters' answer (over the decision used for each domain decided `own`, `other` or `unclear`, which is the one for the homepage the domain was measured from). Raw agreement is shown beside kappa, which is given without an interval and was computed by the study's own code (`src/noroshi/analysis/stats.py`) as (observed - expected agreement) / (1 - expected agreement). When both raters gave the same single answer to every unit, expected agreement is 1, kappa is not defined and its cell is left empty. When one rater, or each of two raters who differ, gave the same answer every time, kappa is 0 and the table's note says so. For the identity check the note also counts the answers, the tie-break model's included, recorded as `unclear` because their replies could not be read.

Federal control (Supplementary Table S1; `tableS3_federal_control`). The broad share among measured federal domains whose broad value was known was described with no acceptance threshold, per listed domain and per distinct final site (a final site is positive when any listed domain leading to it is). Every measured federal domain without any channel was examined individually against its stored pages in the audits, and the reason is reported. A federal homepage that answers 200 with a page saying the page does not exist counts as a "no", because federal-only domains skip the identity check (a known limit).

Breach portal. The HHS OCR breach portal was not used and no breach linkage was built: matching hospitals to breach reports by name is error-prone, and the research questions do not need it.

Software. CPython 3.12.13 [27], pandas 3.1.1 [28], SciPy 1.18.1 (exact intervals and tests) [23], NumPy 2.5.5, HTTPX 0.28.1 [11], tldextract 5.4.0 [1], sectxt 0.9.8 [15] and Protego 0.7.0 [12], pinned in `uv.lock`. Figures were drawn with Matplotlib 3.11.2 [29].

Post hoc analyses. Four hospital analyses were added after the first pass and are post hoc; they are not in the protocol. They were computed by a script outside the analysis command (`posthoc_concentration.py`), which reads the per-domain table and, read-only, the stored file hashes, and was rerun on the final data after pass 2. (1) The share of measured hospital domains with the main outcome, counted per distinct final site (the registrable domain the homepage landed on; positive when any listed domain leading to it is positive), with an exact interval. (2) The number of distinct files behind the

domains with the main outcome (files compared by the SHA-256 of the body) and of the organizations serving them. The organization of each file was identified by reading the file; no automatic rule was used. (3) The share of measured hospital domains with the main outcome without the domains that serve the most common file, with an exact interval. These three follow from the pass-1 audit's finding that one file accounted for most hospital domains with the main outcome (S12). (4) A comparison of hospitals on measured domains with hospitals on domains that could not be measured, by type, ownership, beds and rural location (Supplementary Table S6), computed from the per-hospital results by a second script (`measured_coverage.py`) with the groupings of Supplementary Table S5 and the county type.

Reporting notes. The research questions are descriptive and no hypothesis was set. No sample size was calculated: every domain on each list was taken. The protocol planned two tests, the exact McNemar test for the paired file comparison and a Fisher exact test of single-hospital against multi-hospital domains; no P value was adjusted for multiple comparisons, and the subgroup shares are descriptive. No confounder was considered and no adjusted estimate is reported. The bed groups (under 100, 100 to 299, 300 or more) and the system-size groups (1, 2 to 5, 6 to 20, more than 20 hospitals) were fixed in the analysis code at the commit of the main run, before any main-run result existed. Contrasts between lists were not estimated; the shares and their intervals are the comparison.

Sources of bias and the steps taken. *Selection.* The hospital list holds only hospitals whose record lists a website, and it dates from 2024, so some domains had lapsed or changed hands; Supplementary Table S5 compares hospitals kept with hospitals dropped, and Supplementary Table S6 hospitals on measured with hospitals on unmeasured domains. The device and EHR lists are specific public lists (S1), and company names were matched to domains by models with no accuracy estimate (S1). The direction of these effects on the shares is not known. A domain that could not be measured was never counted as lacking a channel; the sensitivity table (Supplementary Table S3) gives the lower bound in which every unmeasured value counts as none. *Measurement.* Sites that refuse automated requests could not be read; they were reported as blocked, not as absent, after one retry with a browser user agent (S2). Page judgements and identity decisions were made by two models with a third as tie-break (S5, S8), their agreement is in Supplementary Table S2, and the audits read every page text judged positive and scanned every identity decision (S12); the errors found are stated there and in the Limitations. *Unit.* A listed domain is not an organization; the hospital result is therefore also given per distinct final site and per distinct file (post hoc). *Analysis.* The protocol was frozen before the main run, analyses added later are labelled post hoc, and a separate agent recomputed the main tables with its own code (S13).

S12. Pilots, rule changes and audits

Pilots. Two pilots were run before the freeze into a separate store; their data are not used in the results, and their domains were measured again in the main pass. The first drew 30 domains per sector at random (150 sector rows; seed 20261115), ran on November 16, 2026 from 12:44 to 12:49 UTC and sent 1,184 requests. A separate Claude agent (`claude-opus-5-5`) audited it against the stored pages, and its findings changed many crawler and classification rules, among them the addition of the identity check and the removal of an example answer and the page address from the judge prompt, which had made the two judges' kappa 0.00 (protocol §12.1). The second drew its sample after removing the first pilot's domains: 150 sector rows on 148 distinct domains, run on November 17, 2026 from 00:15 to 00:23 UTC at commit `8a7dc5b`, with 1,365 requests (1,383 records). Two readers, one for the crawl and one for the classification, audited it: no record broke a limit, no policy-page judgement was wrong, and seven rules that misfired were changed (protocol §12.2). One change explains the identity prompt: `own` counts a page with little text of its own when its title or meta tags name the organization, because 3 of the 8 domains judged `unclear` were the organization's own script-driven sites. These changes were verified on stored pages and local servers only and entered the main pass without a further pilot. The list of changes is in protocol §12.1 and §12.2, and the limits stated before collection are in §13; the policy-page measure is a lower bound (15 of 18 known device and advisory-vendor policies found in the second pilot's audit).

Audit of pass 1

A separate Claude agent audited pass 1 read-only against the stored evidence: the 30,040 request records (26,746 sent, 3,294 reused copies; 07:37:45 to 08:45:27 UTC on November 17, 2026; 3,557 listed domains, 3,576 sector rows), the stored bodies, the crawl log, the judgement and identity files and the derived tables. The store was opened read-only and no request, DNS lookup or model call was made; samples used the seed 20261117. Every distinct positive `security.txt` file was read in the hospital sector only. The audit made eight checks.

1. **Public promises, every record.** No violation in 30,040 records: the 12-request limits per listed domain (reached by 292 domains) and per site, the 2-second gap (22,909 consecutive pairs), 5 redirects, the browser agent only after a refusal or on a switched site (1,800 browser records; two written before their site's switch, explained by the hold of S2), the other host form once after a network error (700 domains), nothing sent after a 429 (7 sites), robots.txt (13,592 page requests; 50 to other hosts are exempt), `cms-hpt.txt` for hospitals only, and body limits. All but 29 identified refusals got a browser retry, each of the 29 explained by a limit or by a retry that was itself refused elsewhere.

2. **Hospital headline.** All 132 measured hospital domains with the main outcome and the 19 distinct files they were counted from were read in full: all are `text/plain` security.txt files of an organization or its health system, none is a hosting provider's default, and every class was

right. One file of one health system is on 95 of the 132 domains (25 final sites; 120 hospitals). Per listed domain the main outcome was 132 of 1,653 (8.0%), per distinct final site 42 of 1,359 (3.1%), and without the 95 domains 37 of 1,558 (2.4%); the 19 files come from 17 organizations. Of 40 random hospital domains classed absent, none was wrong.

3. **Other file classes** (all sectors). Every forbidden location (99 on 71 domains), every invalid domain (6), every HTML-error location (240 on 132 domains) and 30 random HTML-error and 30 random unmeasured domains were read; no outcome was wrong. Six forbidden labels are doubtful (3 behind a firewall block page, 1 site-wide denial, 2 storage-service answers for a missing object), and one "invalid" file is a hospital's `cms-hpt.txt` reached by redirect.

4. **Policy pages.** All 114 distinct texts judged positive and all 39 judged negative were read: no false positive and no judge false negative; samples of set-aside pages and 70 homepages showed no missed policy. Four measured domains (2 hospital, 2 device) had a "no" or HTML-error value that came from a challenge page served with status 200. Two real policies (the VA's own, for three federal domains, and one company page) were left unread because the first 1 MiB held no visible text, and one hospital policy that defers its contact to its `security.txt` was judged not a policy page.

5. **Identity check.** 40 of 101 `other`, 40 of 82 `unclear` and 60 of 2,363 `own` decisions were read; then every `own` homepage was scanned for pharmacy, gambling, for-sale and expired-domain wording and every `other` and `unclear` homepage for challenge pages and script redirects. In the `own` sample 1 of 60 was wrong; the scan found 9 hospital `own` decisions that are not the organization's site (8 revived copies of one former operator's template and a closed chain's domain turned pharmacy store; 8 sell pharmacy products). Of 80 sampled `other` and `unclear` decisions, 6 were the organization's own site (3 behind a challenge, 3 script-driven); the scan of all 183 found 12 such sites (6 challenge homepages, 6 script redirects or near-empty pages), about 7%. Of 162 hospital mismatches, 156 are right, 5 are the hospital's own site behind a challenge or sign-in page and 1 is a script-driven own site.

6. **Unreachable (716) and blocked (189) domains.** All were read, and every status followed the written rules. The 716 unreachable were mostly DNS failures (229), timeouts (135) and TLS failures (203), with 43 status-202 pages, 56 connection failures, 20 5xx, 18 404, 7 with more than 5 redirects and 5 other statuses. The 43 status-202 homepages are one page with a meta refresh to the hosting provider's captcha path on the same site, keyed to the crawler's address; no rule of S2 fires on it, so they ended unreachable with no browser retry although in substance they refused the crawler (FDA's site ended the same way). Among the blocked, 11 are for sale (the homepage redirects to a marketplace that answered a challenge; the sale test runs only on an answered homepage), 3 federal domains are dead redirects, and 6 hospital domains are blocked chains (refused at a protected site the first attempt was redirected to; the retry's first hop, at the listed host, went out with the identified agent by the per-hop rule). 57 unreachable domains fail only on an incomplete certificate chain that a browser would likely complete (not tested).

7. **Federal control.** All 31 measured federal domains without a positive outcome were read: 11 negatives (1 true miss, a link drawn by script; 1 error homepage; 9 with no policy link) and 20 unmeasured (10 ended by the 12-request site limit of `nih.gov`, `hhs.gov` or `vaccines.gov`

before the HHS policy link on their homepage was used, all ten linking it; 5 script shells; 1 refresh stub; 1 maintenance page; 2 firewall challenges; 1 policy page unread at the size limit). The same limit left a `security.txt` location of 24 measured domains (9 federal, 11 hospital, 4 device) never requested.

8. **Tables against the per-domain file.** Every cell of the flow table, 25 domain-level and 10 hospital per-hospital and bed-weighted cells of the share table, and the federal table were recomputed: no difference.

The frozen protocol was not amended after pass 1. The audit's findings about rules that misfired (the unrecognized challenge pages, the 9 hospital `own` decisions, the 11 blocked domains that are for sale, the 4 challenge-page values) are reported as limits and not corrected; by the final audit's estimate, correcting them moves no headline share by more than 0.1 point.

Final audit (November 19, 2026)

After pass 2, a separate Claude agent audited both passes read-only, with the pass-1 audit's checks except the table recount (replaced by the recomputation in S13); no request, DNS lookup or model call was made. It checked all 3,087 pass-2 records against the public promises and the 48-hour rule and replayed the selection of domains for the retry. It read the 56 domains whose file class changed, all 42 distinct non-hospital files behind a main outcome, the 4 new judged page texts, all 47 new identity decisions, all 705 unreachable and 170 blocked domains and all 124 federal domains.

No public promise was broken in pass 2 (0 violations), no site was requested again within 48 hours of a pass-1 request to it, and the retry covered exactly the 1,049 domains the protocol selects. The hospital positives were unchanged: the same 132 domains, 19 files and 17 organizations. No file class was wrong. One device maker's valid file, at conavi.com, belongs to the content platform hosting its website, not to the company; the file at ge.com is GE Aerospace's, leaving open whether that is the advisory vendor. The new judgements held no false positive or false negative, and no disclosure link was missed on the 127 homepages measured in pass 2. One new identity decision is doubtful: a hospital domain counted among the 26 that led to another organization's site may lead to its operator's. The 9 hospital domains wrongly accepted in pass 1 were not retried and are still counted as measured. 51 domains refused the crawler with a challenge it does not recognize and ended unreachable (45) or mismatch (6), and 99 of the 170 blocked domains met a browser challenge page. Of the 144 domains measured in pass 1 and retried, 41 were refused (39) or unreachable (2) in pass 2, so their pass-1 values stand. Two federal domains (smokefree.gov, thisfreelife.gov) stayed unknown because their only gap, a page skipped for the per-site limit, is not a reason for a retry under the protocol. As after pass 1, these findings are reported, not corrected.

S13. Verification, release and AI use

Verification. Every stored response was content-hashed (SHA-256 of the body; the stored file name is the hash). The protocol (§9) requires that a separate agent recompute the main table from the raw store with its own code and that every reference in the paper be checked against its source. On November 19, 2026, a separate agent recomputed the flow table and the main and strict outcomes per listed domain from the raw store with its own code, written from the protocol text; all 116 cells compared were equal. It used the store's refusal flag for 56 records probably refused by a response header, which the store does not keep, and the crawler's end status for one domain whose third-party-platform test it did not rebuild. On November 19, 2026, every reference of the paper and of this supplement was compared with its verified record (all details matched), every sentence that cites one was read against the recorded quote or finding, and each DOI or address was requested once. Six sentences went beyond or disagreed with their records and were reworded. For 11 references of the paper and 5 of this supplement, an address (at a publisher, at eCFR, or one large data file) answered with a bot check or a timeout and could not be confirmed by an automated request. On November 20, 2026 (UTC), each of those addresses was opened in a web browser, and the data file was requested once more; all answered, and each page showed the cited item.

Search for earlier work. The Background's statement that no earlier measurement was found rests on searches made on November 16 and November 19, 2026. On November 16, eight phrase queries (for example `"security.txt" healthcare`, `"vulnerability disclosure" hospital`, `"vulnerability disclosure policy" "medical device"` and `"bug bounty" healthcare`) were run on Europe PMC, OpenAlex and PubMed, six stricter title-and-abstract queries on OpenAlex and four stricter queries on PubMed. On November 19, the eight phrase queries, the six stricter OpenAlex queries and four more on the adoption of security.txt and of disclosure policies in general were run on Europe PMC, OpenAlex, PubMed, Crossref and arXiv, and one query on Google Scholar. Semantic Scholar and DBLP refused automated queries on both dates and were not searched. Titles and abstracts were read, not full texts, so a health-sector row inside a general study cannot be ruled out. No study was found that measures whether US hospitals, medical device makers or EHR developers publish a reporting channel; the nearest items were policy comparisons, case studies, an analysis of vulnerability records by sector, and measurements of other sectors or countries.

Release. Released at https://github.com/k0nnect/noroshi (a snapshot of November 19, 2026; commit `328be22`): the code, the protocol, the frame tables, a per-domain table of derived results, the analysis scripts and their outputs, the AI-use log and the STROBE checklist. Withheld: the raw page snapshots; the log of model calls, which holds page text; the contact values found in security.txt files; the audit reports, which discuss

individual sites; copies of third-party web pages used as test fixtures; and copies of the outside directories, which are pinned by commit and hash instead. Commit identifiers given elsewhere in the paper and in this supplement refer to the study's development history, which is not part of the public snapshot; the SHA-256 of the frozen protocol file is `469d0582a1cbea05d4b8fc5b63c26a6e554c8748b641f688716e501e79af0a51`.

Use of AI. Claude models, run as coding agents under the author's direction (the log names Claude Sonnet 5.5 and Claude Opus 5.5 for the coding tasks), searched the literature, wrote and reviewed the code, ran the collection and analysis, and drafted the paper; each step is recorded in `logs/ai-use-log.md`. The model calls that are part of the method (S1, S5, S8) were made by `claude-opus-5-5`, `claude-sonnet-5-5`, `claude-haiku-5-5` and `claude-fable-5-2` and are logged in full in `data/raw/llm_log.jsonl`. Two log entries concern calls that are not study calls: one line (purpose `p`, model `m`, prompt `hi`, November 16, 2026 08:01:35 UTC) was written by a unit test and is left out of any count of model calls; and during a code-review round before the freeze a test ran the identity command once before the run lock existed, unfaked, and made 28 real model calls (13 `claude-sonnet-5-5`, 13 `claude-haiku-5-5`, 2 `claude-opus-5-5`) about synthetic test pages, logged only in a temporary test folder. No study data were sent, and every test now finds a failing stand-in for the `claude` command first on its path.

Supplementary References

1. Kurkowski J. *tldextract: accurately separates a URL's subdomain, domain, and public suffix, using the Public Suffix List* (software), version 5.4.0. https://github.com/john-kurkowski/tldextract Also: Mozilla Foundation and community contributors. *Public Suffix List*. Snapshot 2025-04-07 (commit 5fc8d12c1624ddbc56a80e944aca151a8ee466a1). https://publicsuffix.org/list/public_suffix_list.dat.

2. 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

3. 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

4. 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

5. 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

6. 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

7. 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

8. 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]".

9. Office of the National Coordinator for Health Information Technology. *Certified Health IT Product List (CHPL)*. https://chpl.healthit.gov/ ; description page https://www.healthit.gov/topic/certified-health-it-products-list-chpl

10. 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

11. Encode. *HTTPX: the next generation HTTP client* (software), version 0.28.1. BSD-3-Clause. https://github.com/encode/httpx ; documentation https://www.python-httpx.org

12. Scrapy project. *Protego: pure-Python robots.txt parser with support for modern conventions* (software), version 0.7.0. https://github.com/scrapy/protego

13. 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

14. Internet Archive. *Wayback CDX Server API - BETA*. In: `internetarchive/wayback` repository, directory `wayback-cdx-server`. GitHub. Accessed November 17, 2026. https://github.com/internetarchive/wayback/tree/master/wayback-cdx-server Also: Internet Archive. Wayback Machine. CDX Server API, endpoint `http://web.archive.org/cdx/search/cdx`, queried for the study's first-seen analysis (`src/noroshi/wayback.py`; run log `data/raw/wayback_pass1.out`). Accessed November 17, 2026 (service check). https://web.archive.org/

15. Digital Trust Center. *SecTXT: security.txt parser and validator* (software), version 0.9.8. EUPL-1.2. https://github.com/DigitalTrustCenter/sectxt ; PyPI https://pypi.org/project/sectxt/0.9.8/

16. 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

17. 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 .

18. Poppler contributors. *Poppler: a PDF rendering library* (utility `pdftotext`). Version 26.09.0. https://poppler.freedesktop.org/

19. 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

20. 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

21. 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

22. 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

23. Virtanen P, Gommers R, Oliphant TE, et al; SciPy 1.0 Contributors. SciPy 1.0: fundamental algorithms for scientific computing in Python. *Nat Methods*. 2020;17(3):261-272. doi:10.1038/s41592-019-0686-2. SciPy version 1.18.1 was used.

24. 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

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

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

27. Python Software Foundation. Python Language Reference, version 3.12.13. Released March 3, 2026. Accessed November 17, 2026. https://www.python.org/downloads/release/python-31213/

28. The pandas development team. pandas-dev/pandas: pandas, version v3.1.1. Zenodo; November 4, 2026. doi:10.5281/zenodo.23955360. Also: McKinney W. Data structures for statistical computing in Python. In: van der Walt S, Millman J, eds. *Proceedings of the 9th Python in Science Conference*. 2010:56-61. doi:10.25080/Majora-92bf1922-00a.

29. Hunter JD. Matplotlib: a 2D graphics environment. *Comput Sci Eng*. 2007;9(3):90-95. doi:10.1109/MCSE.2007.55. Also: Matplotlib development team. Matplotlib: Visualization with Python, version v3.11.2. Zenodo; September 11, 2026. doi:10.5281/zenodo.22700776.

Supplementary Tables

Supplementary Table S1. The federal control: any-channel outcome per listed domain and per distinct final site.

Unit	With any channel	Measured	% (95% CI)
Listed domain	78	89	87.6% (79.0-93.7%)
Distinct final site	41	52	78.8% (65.3-88.9%)

Footnotes. Source: `results/tables/tableS3_federal_control.md`. Measured federal domains whose any-channel value is known. A distinct final site is the registrable domain on which a homepage ended after redirects; listed domains that lead to one site count once, as positive when any of them is. Intervals are exact (Clopper-Pearson). No acceptance threshold was set for the control. Reasons, from the final audit, which read every federal domain without a positive outcome: of the 11 measured domains with no channel, 9 linked no disclosure policy, 1 had an error page as its homepage and 1 was a true miss (its link is drawn by a script, which the crawler does not run). Of the 12 measured domains with an unknown value, 2 were stopped by the per-site request limit and were outside the retry rule (smokefree.gov, thisfreelife.gov), 7

homepages showed no readable content, 2 met challenge pages and 1 linked a policy page that exceeded the size limit. Of the 23 domains that could not be measured, 20 were unreachable (1 through a bot-block redirect) and 3 were blocked (2 by a firewall challenge, 1 a retired domain).

Supplementary Table S2. Agreement between the two models, before any tie-break.

Question	Units both answered	Agreed	%	Cohen's kappa	Note
Policy judge: does the page invite vulnerability reports?	156	153	98.1%	0.95	
Policy judge: does it give a way to report?	156	154	98.7%	0.9688	
Policy judge: does it promise not to pursue good-faith reporters?	156	155	99.4%	0.9791	
Identity check: own site, another site, or unclear	2,574	2,490	96.7%	0.7975	

Footnotes. Source: `results/tables/tableS2_agreement.md`. For the policy judge the unit is a distinct page text that both models answered (pages screened out or set aside by rule were not judged); for the identity check it is a listed domain that both models answered. Kappa is given without an interval (Supplementary Methods S11).

Supplementary Table S3. Sensitivity analysis: every listed unit counted, with an outcome that could not be measured counted as no channel (k/n; %; exact 95% CI).

Unit	Outcome	Hospitals	Device makers	Advisory vendors	EHR developers	Federal control
Listed domain	Any channel	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%)
Listed domain	security.txt, main outcome	132/2,652 (5.0%; 4.2-5.9%)	30/525 (5.7%; 3.9-8.1%)	7/75 (9.3%; 3.8-18.3%)	8/200 (4.0%; 1.7-7.7%)	4/124 (3.2%; 0.9-8.1%)
Listed domain	security.txt, strict outcome	32/2,652 (1.2%; 0.8-1.7%)	22/525 (4.2%; 2.6-6.3%)	6/75 (8.0%; 3.0-16.6%)	6/200 (3.0%; 1.1-6.4%)	0/124 (0.0%; 0.0-2.9%)
Listed domain	Policy page	114/2,652 (4.3%; 3.6-5.1%)	59/525 (11.2%; 8.7-14.3%)	17/75 (22.7%; 13.8-33.8%)	9/200 (4.5%; 2.1-8.4%)	76/124 (61.3%; 52.1-69.9%)
Listed domain	Outside listing	1/2,652 (0.0%; 0.0-0.2%)	25/525 (4.8%; 3.1-6.9%)	19/75 (25.3%; 16.0-36.7%)	3/200 (1.5%; 0.3-4.3%)	15/124 (12.1%; 6.9-19.2%)
Distinct final site	Any channel	42/2,340 (1.8%; 1.3-	77/515 (15.0%; 12.0-	25/74 (33.8%; 23.2-	14/199 (7.0%; 3.9-	41/83 (49.4%; 38.2-

		2.4%)	18.3%)	45.7%)	11.5%)	60.6%)
Distinct final site	security.txt, main outcome	42/2,340 (1.8%; 1.3-2.4%)	28/515 (5.4%; 3.6-7.8%)	7/74 (9.5%; 3.9-18.5%)	8/199 (4.0%; 1.8-7.8%)	2/83 (2.4%; 0.3-8.4%)
Distinct final site	security.txt, strict outcome	12/2,340 (0.5%; 0.3-0.9%)	21/515 (4.1%; 2.5-6.2%)	6/74 (8.1%; 3.0-16.8%)	6/199 (3.0%; 1.1-6.4%)	0/83 (0.0%; 0.0-4.3%)
Distinct final site	Policy page	28/2,340 (1.2%; 0.8-1.7%)	57/515 (11.1%; 8.5-14.1%)	17/74 (23.0%; 14.0-34.2%)	9/199 (4.5%; 2.1-8.4%)	40/83 (48.2%; 37.1-59.4%)
Distinct final site	Outside listing	1/2,340 (0.0%; 0.0-0.2%)	22/515 (4.3%; 2.7-6.4%)	18/74 (24.3%; 15.1-35.7%)	3/199 (1.5%; 0.3-4.3%)	3/83 (3.6%; 0.8-10.2%)
Hospital	Any channel	175/3,693 (4.7%)				
Hospital	security.txt, main outcome	175/3,693 (4.7%)				
Hospital	security.txt, strict outcome	49/3,693 (1.3%)				
Hospital	Policy page	145/3,693 (3.9%)				
Hospital	Outside listing	3/3,693 (0.1%)				
Bed-weighted	Any channel	41,495/543,527 (7.6%)				
Bed-weighted	security.txt, main outcome	41,495/543,527 (7.6%)				
Bed-weighted	security.txt, strict outcome	10,540/543,527 (1.9%)				
Bed-weighted	Policy page	35,780/543,527 (6.6%)				
Bed-weighted	Outside listing	3,157/543,527 (0.6%)				

Footnotes. Source: `results/tables/tableS1_sensitivity.md`. Every unit on the list counts in the denominator. An outside listing (a disclosure platform, disclose.io or the CVE Numbering Authority list) counts even for a domain that was blocked or unreachable, because a listing needs no crawl. Per-hospital and bed-weighted shares exist for hospitals only and carry no interval; bed-weighted shares leave out hospitals with an unknown bed count.

Supplementary Table S4. Hospitals: any-channel outcome by characteristic.

Part a. Per hospital domain, by the number of hospitals listed on it.

Characteristic	Category	Measured	With any channel	% (95% CI)	P
Hospitals listed on the domain	1	1,327	103	7.8% (6.4-9.3%)	
Hospitals listed on the domain	2 to 5	249	28	11.2% (7.6-15.8%)	

Hospitals listed on the domain	6 to 20	41	1	2.4% (0.1-12.9%)	
Hospitals listed on the domain	more than 20	1	0	0.0% (0.0-97.5%)	
Hospitals listed on the domain	single-hospital vs multi-hospital domains (Fisher exact)				0.2363

Part b. Domains that list one hospital, by that hospital's characteristics.

Characteristic	Category	Measured	With any channel	% (95% CI)	P
Hospital type	General acute care	684	91	13.3% (10.8-16.1%)	
Hospital type	Critical access	282	4	1.4% (0.4-3.6%)	
Hospital type	Psychiatric	200	4	2.0% (0.5-5.0%)	
Hospital type	Children	25	0	0.0% (0.0-13.7%)	
Hospital type	Long term care	56	3	5.4% (1.1-14.9%)	
Hospital type	Rehabilitation	33	0	0.0% (0.0-10.6%)	
Hospital type	other	47	1	2.1% (0.1-11.3%)	
Ownership	Non-profit	508	19	3.7% (2.3-5.8%)	
Ownership	Proprietary	511	76	14.9% (11.9-18.3%)	
Ownership	government, non-federal	235	6	2.6% (0.9-5.5%)	
Ownership	unknown	73	2	2.7% (0.3-9.5%)	
Beds	under 100	776	24	3.1% (2.0-4.6%)	
Beds	100 to 299	390	49	12.6% (9.4-16.3%)	
Beds	300 or more	141	30	21.3% (14.8-29.0%)	
Beds	unknown	20	0	0.0% (0.0-16.8%)	
Nonmetropolitan county	yes	486	15	3.1% (1.7-5.0%)	
Nonmetropolitan county	no	841	88	10.5% (8.5-12.7%)	
Nonmetropolitan county	unknown	0	0		

Part c. Per hospital (a domain's result counted for each hospital it lists); no intervals.

Characteristic	Category	Measured	With any channel	% (95% CI)	P
Hospital type	General acute care	1,358	143	10.5%	
Hospital type	Critical access	391	7	1.8%	
Hospital type	Psychiatric	262	7	2.7%	

Hospital type	Children	45	0	0.0%	
Hospital type	Long term care	108	5	4.6%	
Hospital type	Rehabilitation	119	5	4.2%	
Hospital type	other	81	8	9.9%	
Ownership	Non-profit	1,154	37	3.2%	
Ownership	Proprietary	760	117	15.4%	
Ownership	government, non-federal	303	9	3.0%	
Ownership	unknown	147	12	8.2%	
Beds	under 100	1,285	50	3.9%	
Beds	100 to 299	699	78	11.2%	
Beds	300 or more	339	46	13.6%	
Beds	unknown	41	1	2.4%	
Nonmetropolitan county	yes	690	23	3.3%	
Nonmetropolitan county	no	1,674	152	9.1%	
Nonmetropolitan county	unknown	0	0		

Footnotes. Source: `results/tables/table3_hospital_characteristics.md`. Descriptive only; no regression model was fitted. Intervals are exact (Clopper-Pearson). P values are from the tests named in the row. Nonmetropolitan county: Rural-Urban Continuum Codes 4 to 9 (Supplementary Methods S1).

Supplementary Table S5. Hospitals kept in the study compared with eligible hospitals dropped because no usable website was listed.

Characteristic	Category	Kept, n (%)	Dropped, n (%)	Kept, % of category
Hospital type	General acute care	2,070 (56.1%)	1,968 (51.3%)	51.3%
Hospital type	Critical access	601 (16.3%)	607 (15.8%)	49.8%
Hospital type	Psychiatric	353 (9.6%)	640 (16.7%)	35.5%
Hospital type	Rehabilitation	242 (6.6%)	244 (6.4%)	49.8%
Hospital type	Long term care	238 (6.4%)	100 (2.6%)	70.4%
Hospital type	Other	123 (3.3%)	189 (4.9%)	39.4%
Hospital type	Children	66 (1.8%)	90 (2.3%)	42.3%
Ownership	Non-profit	1,765 (47.8%)	1,489 (38.8%)	54.2%
Ownership	Proprietary	1,254 (34.0%)	1,059 (27.6%)	54.2%
Ownership	Non-federal government	459 (12.4%)	668 (17.4%)	40.7%
Ownership	Unknown	215 (5.8%)	622 (16.2%)	25.7%
Beds	Under 100	2,092 (56.6%)	2,264 (59.0%)	48.0%
Beds	100-299	1,032 (27.9%)	894 (23.3%)	53.6%
Beds	300 or more	510 (13.8%)	527 (13.7%)	49.2%
Beds	Unknown	59 (1.6%)	153 (4.0%)	27.8%

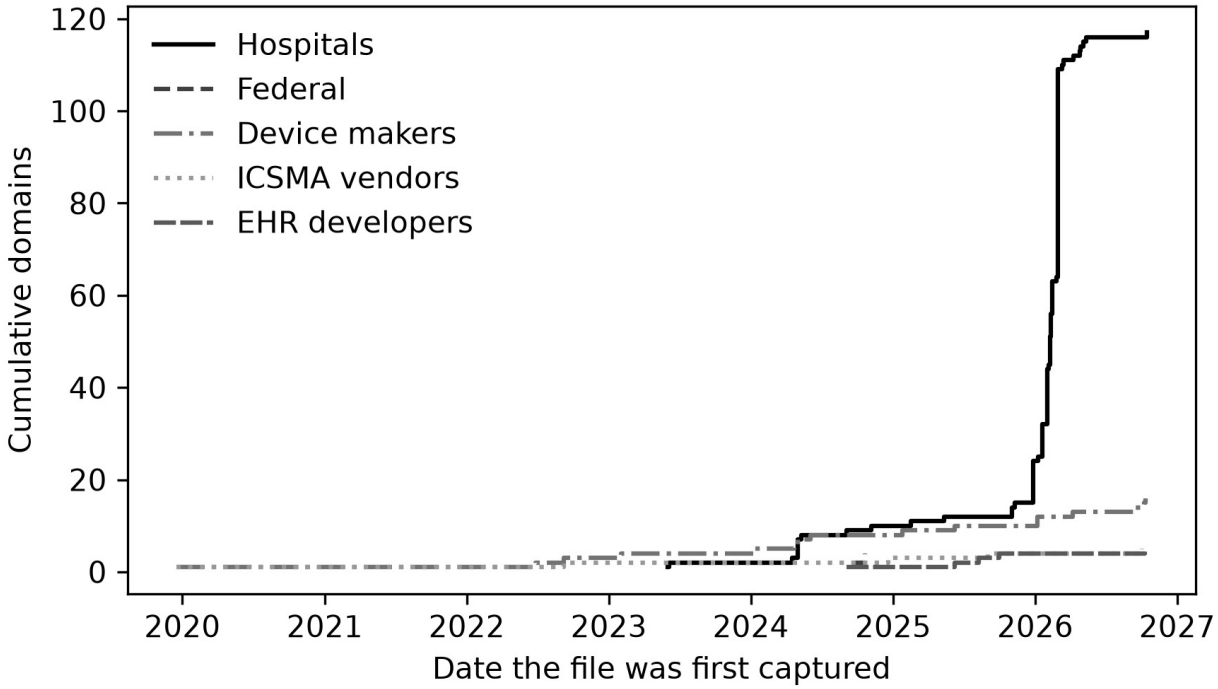
Footnotes. Source: `analysis/frame_coverage.md`, computed from the hospital list (`data/frames/hospitals.csv`). Kept: 3,693 hospitals. Dropped at steps 4 to 6: 3,838 (no website listed 3,597, website belongs to someone else 233, website cannot be read 8).

Supplementary Table S6. Hospitals on domains that could be measured compared with hospitals on domains that could not (post hoc).

Characteristic	Category	Measured, n (%)	Not measured, n (%)	Measured, % of category
Hospital type	General acute care	1,449 (56.3%)	621 (55.5%)	70.0%
Hospital type	Critical access	411 (16.0%)	190 (17.0%)	68.4%
Hospital type	Psychiatric	269 (10.4%)	84 (7.5%)	76.2%
Hospital type	Rehabilitation	143 (5.6%)	99 (8.9%)	59.1%
Hospital type	Long term care	164 (6.4%)	74 (6.6%)	68.9%
Hospital type	Other	88 (3.4%)	35 (3.1%)	71.5%
Hospital type	Children	51 (2.0%)	15 (1.3%)	77.3%
Ownership	Non-profit	1,234 (47.9%)	531 (47.5%)	69.9%
Ownership	Proprietary	859 (33.4%)	395 (35.3%)	68.5%
Ownership	Non-federal government	320 (12.4%)	139 (12.4%)	69.7%
Ownership	Unknown	162 (6.3%)	53 (4.7%)	75.3%
Beds	Under 100	1,425 (55.3%)	667 (59.7%)	68.1%
Beds	100-299	732 (28.4%)	300 (26.8%)	70.9%
Beds	300 or more	373 (14.5%)	137 (12.3%)	73.1%
Beds	Unknown	45 (1.7%)	14 (1.3%)	76.3%
County	Metropolitan county	1,847 (71.7%)	775 (69.3%)	70.4%
County	Nonmetropolitan county	728 (28.3%)	343 (30.7%)	68.0%

Footnotes. Source: `analysis/measured_coverage.md`, computed from the per-hospital results (`data/derived/hospitals.csv`). Added after the first pass; not in the protocol. Each hospital takes the status of its listed domain. Hospitals on measured domains: 2,575. Hospitals on domains that could not be measured: 1,118 (unreachable 740, blocked 199, mismatch 179). All hospitals: 3,693.

Supplementary Figure



Supplementary Figure S1. Cumulative number of domains with the main outcome, by the date of the earliest Internet Archive capture of their security.txt file. One line per list. Of the 178 distinct domains with the main outcome, 145 had a capture (117 of the 132 hospital domains); domains with no capture (32) or an incomplete lookup (1) are not drawn. A first capture shows only that the file existed by that date. Drawn by the study pipeline (`results/figures/figure2_first_seen`); in its key, "ICSMA vendors" are the advisory vendors and "Federal" is the federal control.