



How UK Hospital Web Estates Look From Outside

What a non-intrusive external assessment of 269 tracked NHS and independent hospital websites found, and why publishing an enforced email policy is the most practical first step

First edition, 13 October 2026. Evidence gaps are stated throughout.

Purpose

To set out, in plain English, what an outside observer can see of this sector's public web estate, what the evidence does and does not establish, and which actions a hospital communications, digital or governance team can take without a security programme.

Sources Two tracked lists, reported first in their own right and only then together: **UK Hospitals, public** (231 domains) and **UK Hospitals, private** (38 domains). The two lists share no domain. Counted by website, not by hospital site, operator, patient or legal entity.

Evidence Primary scans completed for all 269 on 13 October 2026 under one scan method. Twelve completed without a published score and are excluded from every average. That leaves **257 scored**: 227 full scores and 30 minimum scores.

Method Non-intrusive external assessment using publicly observable information: websites, domain records, certificates, public threat intelligence and publicly reachable services. No logins, password attempts or intrusive probing. This document reuses stored evidence and initiated no scan.

Not a register The population is the tracked domain list, **not a verified census of hospitals or healthcare providers**. One domain can serve a group or several sites. The public and private labels are administrative list names and do not establish funding, ownership or endorsement. **No assessment of clinical systems or patient care is implied or possible from outside.**

Prepared by Huro Data Technologies Ltd / MyDomainRisk

Date 13 October 2026

Research and external assessment: [MyDomainRisk](#) | Huro Data Technologies Ltd

For information: Our scans use bounded, external, non-intrusive checks. They do not attempt exploitation, credential testing or intentional modification of target systems.

Executive summary

The four things that need attention

1. 147 of 257 websites, 57%, do not enforce a policy against email forged in their name. 123 publish no record we could find and 24 publish one that only asks receiving systems to monitor. 110 have an enforcing policy: 90 instruct refusal and 20 send to spam. A letter that appears to come from a hospital is acted on: appointments are moved, results are opened, invoices are paid. **This covers the exact domain only, not lookalike domains.**

2. Five websites carry a match against the public list of software weaknesses that criminals are known to be using. None of the 511 underlying software observations was recorded as confirmed by the source, so each match is a reason to check privately whether it applies, not a finding that anything was broken into.

3. Twenty-two websites have a database port answering from the open internet. For three, the evidence points at the organisation's own service. For nineteen, the address is shared with unrelated organisations and no service was identified, so it is the hosting company's equipment that needs confirming first.

4. The two records that protect the hospital's own identity on the internet are effectively unused. One website in 257 publishes signed answers to the question of where its domain lives, and five have a list of who may issue their certificates. Across the 220 public-list websites, the signing figure is zero. Most public-list websites sit under a shared NHS domain, where this protection depends on the parent domain being signed first, so this zero is largely a matter for the registry rather than for each organisation.

What we found, worst first

What we found	Websites	Share of 257
No enforced email-forgery policy (no record, or monitor only)	147	57%
Website does not tell visitors' browsers what is safe to load	143	63% of those we could read
Staff logins turning up in stolen-password collections	83	33% of 248 checked
Some other door open that normally would not be	31	12%
A database port answering from the open internet	22	9%
Software with a flaw that criminals are known to be using	5	2%
The website's internet address flagged by a public watchlist	4	2%
A criminal group naming the organisation on a leak site	0	0%

Logins could be checked for 248 of the 257; the other nine are unknown, not clean. Rows overlap and must never be added together. Each counts websites, not problems, and not hospitals. The last row is a real result worth stating and a narrow one: it means no reference was retained in the sources we read, not that no incident has occurred.

Executive summary, continued

Where the two lists differ, and where they do not

The two lists are reported separately throughout because they are different populations with different suppliers, not because one is better run than the other. Two differences are large enough to mention here.

What we checked	Public list 220 scored	Private list 37 scored
No enforced email-forgery policy	133 (60%)	14 (38%)
Average full score	64.6	68.1
Everything resting on one company for the address records	6 (3%)	33 (89%)
No cookie consent banner we could see	153 of 196 read	13 of 32 read

The two lists are different populations and are shown side by side for context, not ranked.

The private list scores slightly higher and enforces a policy more often (23 of 37 against 87 of 220). It is also far more concentrated: nearly nine in ten of those websites depend on a single company for the records that tell the internet where the site lives. The public list has the opposite shape, spread across many suppliers but weaker on the one control that stops impersonation.

Three cautions before any figure is quoted

These are websites, not hospitals. A group can run one website for many sites, and one organisation can use several domains. Nothing here is a count of hospitals, beds, operators or patients.

Nothing here says anyone was broken into. Every observation was made from outside, without logging in or testing a password. A finding is a reason to check, and several of them are most likely the hosting company's business rather than the hospital's.

Nothing here touches clinical systems. Patient records, clinical networks and medical devices sit behind authentication and are not visible from outside. This report is about the public web estate only.

Scores: the average full score is 65.1 across 227 websites. Including the 30 minimum scores at their floor, the average across all 257 is at least 63.8. Those two figures describe different groups and are not a before-and-after comparison.

Why this matters for hospitals in particular

Most of the checks in this report are the same ones we would run on any sector. What differs is the consequence when one of them is missing, and in healthcare three consequences stand out.

A forged hospital email is acted on immediately

People do not treat a message from a hospital as optional. An appointment change, a test result, a pre-admission instruction or a request for a payment detail carries an authority that a message from a shop does not. **147 of 257 websites do not enforce a policy, so receiving systems have no instruction to refuse or quarantine a message forged in their name.** Only 90 instruct refusal. A message that appears to come from a hospital, about an appointment, a result or a payment, is likely to be acted on, which makes this the most useful single control to put in place.

The people whose logins appear are not one group

Stolen-password collections touch these organisations in three separate ways, and merging them produces a frightening number that describes nobody. Staff identities, organisational identities held on outside platforms, and the accounts of members of the public who registered for something are different people, with different consequences and different owners. **None of it is evidence that a patient record was exposed,** and this report does not claim that it is. Section 6 keeps the three apart.

Suppliers hold much of this estate

Many of the shared-address findings in this report are likely to be a question for the hosting supplier rather than a job for the organisation. **181 of 257 websites carry a broader concentration observation,** and on the private list 33 of 37 rest on a single company for their address records. That concentration is often the right commercial choice. It becomes a governance question only when nobody has written down who would be told, and what would stop, if that one supplier had a bad day.

What this report deliberately does not do

It names no organisation. It makes no clinical, safeguarding, regulatory or compliance judgement. It does not rank the two lists against each other as a public-versus-private comparison, because the lists were built administratively and are not matched populations. And it does not treat an absence of evidence as evidence of safety.

1. The words used in this report

Technical terms have been replaced with plain descriptions throughout. This page gives the translation, so that anyone comparing this report with a technical one can follow it.

What we call it here	What it actually means
An open door	A point on a server that answers the internet when something knocks. Some are meant to be open. The check notes the ones that normally would not be.
A database reachable from the internet	A door that answers in the way a database normally does. It does not mean anyone read anything from it.
Known to be used by criminals	The software flaw appears on a public list of flaws that attackers have actually been seen using. It still has to be checked that the flaw applies to the software actually running.
A full score and a minimum score	A full score is calculated from complete evidence. A minimum score is published where part of the evidence could not be read, with everything unread counted as missing, so the real figure can only be the same or higher.
Counted against the organisation	The evidence points at this organisation rather than at a shared platform: a service was identified, or the address was seen serving only them.
Reported separately	Nothing was identified, and unrelated organisations share the same address, so only the hosting company can say whose it is. Separate does not mean smaller.
Signed answers about where the domain lives	A published signature on the records that tell the internet where the website and mail live, so a system that checks signatures can spot a forged answer and refuse it. It does not stop someone changing those records through the registrar or hosting account, and it does not protect the website itself.
Stolen-password collections	Compilations of login details gathered from breaches elsewhere and from software that steals them off an individual's own device. Nothing in them was tested, and no password appears in this report.
A loose end	A web name the organisation still publishes that points at something no longer there, which somebody else could potentially claim.

No password, credential record, raw supplier record or exploitation instruction appears anywhere in this document, and no organisation is named.

2. The first source list: UK hospitals, public

This is the larger of the two tracked lists: **231 domains, all scanned on 13 October 2026**. Eleven completed without a published score, leaving 220 scored, of which 195 carry a full score and 25 a minimum score.

What the first list contains	Result
Domains tracked, and primary scans completed	231 / 231
Scanned but no score could be given	11
Scored: full scores / minimum scores	220: 195 / 25
Average full score	64.6
Average including minimum scores at their floor	at least 63.4
Bands: good / needs improvement / at risk	25 / 165 / 30

What stands out on this list alone

133 of 220, 60%, do not enforce a policy against forged email, and 117 publish no record we could find. **Not one of the 220 publishes signed answers about where its domain lives**, and four have a list of who may issue their certificates. Three carry a match against the public list of flaws criminals are using, and nineteen have a database answering from the internet.

Two measurement points matter when reading this list. **The domain registration details could only be read for 7 of the 220**, because the registry these domains sit in does not publish them, so that whole line of evidence is unavailable rather than absent. And the direct conversation with the mail server succeeded on 108 and was skipped on 112, so mail-encryption questions are unanswered for roughly half the list.

The eleven unscored domains keep their observed findings, which are reported in section 12 and are not included in any count above. Their score is unavailable, never zero.

3. The second source list: UK hospitals, private

The smaller list: **38 domains, all scanned on 13 October 2026**. One completed without a published score, leaving 37 scored, of which 32 carry a full score and 5 a minimum score. **No domain appears on both lists.**

What the second list contains	Result
Domains tracked, and primary scans completed	38 / 38
Scanned but no score could be given	1
Scored: full scores / minimum scores	37: 32 / 5
Average full score	68.1
Average including minimum scores at their floor	at least 66
Bands: good / needs improvement / at risk	6 / 28 / 3

What stands out on this list alone

It is better at the thing the other list is worst at. 14 of 37 do not enforce a policy against forged email, and only 6 publish no record we could find. Seventeen instruct outright refusal, a higher share (17 of 37, against 73 of 220).

It is much more concentrated. 33 of 37 rest on a single company for their address records, against 6 of 220 on the public list. Nothing about that is wrong in itself, and for a small organisation it is usually the sensible choice, but it means one supplier's bad day reaches almost the whole list at once.

Two carry a match against the public list of flaws criminals are using, three have a database answering from the internet, and three have an internet address flagged by a public watchlist. On a list of 37, those are small counts that should still each be checked.

With 37 scored domains, a single website moves any percentage on this list by nearly three points. Counts are given rather than leaning on percentages for that reason.

4. The two lists together

From here on the report reads the two lists as one estate of **257 scored websites**, because the checks and the scan method are identical. Every categorical figure in this report was reconciled across the two source documents and the combined evidence file: **94 measures, no mismatches**.

Measure	Public list	Private list	Together
Domains tracked	231	38	269
Scanned but not scored	11	1	12
Scored websites	220	37	257
Full scores / minimum scores	195 / 25	32 / 5	227 / 30
Average full score	64.6	68.1	65.1
Average including floors	at least 63.4	at least 66	at least 63.8

The spread of scores, and why the average is not the finding

Band	Full scores	Minimum scores	Together
Good 75 and above	31	0	31
Needs improvement 50 to 74	174	19	193
At risk below 50	22	11	33

The shape of this estate is a low middle, not a bad tail. 193 of 257 websites sit in the middle band and only 31 reach the top. That describes an estate where ordinary settings have gone unreviewed almost everywhere, rather than one with a dangerous minority. The eleven minimum scores in the bottom band may well rise once the missing evidence is available, and should not be read as confirmed poor performers.

A website can sit in the top band and still carry a finding from section 5. The bands and the consequential findings measure different things and should be governed separately.

5. The findings that carry real consequence

These are the observations that would change what somebody does on Monday morning. Every count is a distinct scored website out of 257, and the rows overlap.

What we found	Websites	Share
Software with a flaw criminals are known to be using	5	1.9%
A database port answering, counted against the organisation	3	1.2%
A database port answering, on a shared address	19	7.4%
Any open door, counted against the organisation	10	3.9%
Any open door, on a shared address	21	8.2%
A reported software flaw of any kind	14	5.4%
An encryption or certificate problem	13	5.1%
A criminal group naming the organisation on a leak site	0	0%

How to read the software figures without overstating them

The evidence holds **511 software observations across 14 websites, none of them confirmed by the source**. Of 511 entries checked against the public list of flaws criminals are using, **6 came back as matches, across 5 websites**. Checking an entry is not finding one: the 511 is the work done, the 6 is what it found.

None was tested, so the first step is confirming that the affected software and setting are actually in use. **Nothing was exploited and nothing was tested**. A match establishes that checking matters, not that anything was compromised.

The leak-site zero is worth stating and worth bounding. No reference to either list appeared in the ransomware leak-site evidence we read. That evidence is a criminal group's own claim rather than verified fact, and an absence in it is not proof that no incident has occurred anywhere in this estate.

5 continued: the open doors, and whose address they sit on

Thirty-one websites have a service answering that probably should not be. Each case was judged on two questions: was a running service identified on the open door, and do unrelated organisations use the same address?

How the case was counted	Websites	Open doors	Of those doors, a database reachable from the internet
Counted against the organisation a service was identified, or the address is used only by that organisation	10	13	3
Reported separately nothing identified, and unrelated organisations share the address	21	146	19

The two rows are exclusive and together describe 31 websites and 159 individual open doors. The 22 databases reachable from the internet across both rows are the same 22 the card-payment check in section 9 reports as problems, which is why that figure appears twice and must be counted once.

What the split does and does not establish

Being counted against the organisation does not prove the organisation controls the service.

It means the evidence points that way rather than at a shared platform. The hosting company should confirm before anything is restricted or disclosed, and the split changes neither the stored score nor the severity.

Reported separately does not mean smaller. Nothing was identified on those 21 and the addresses serve unrelated organisations, so the doors may be the provider's equipment rather than the hospital's. They still matter: a weakness in a shared platform reaches every organisation standing on it at the same moment, and 146 of the 159 doors in this report sit in that group.

An open door is not the same as a weakness. A service may be deliberately published, restricted to known addresses, or run by the host on the hospital's behalf. What it does mean is that the door is visible and somebody should confirm what is behind it. For most organisations here that is a single written question to whoever hosts the site.

The database ports are the ones to take first. Twenty-two websites across both rows have a port that answers the way a database normally does. For 19 of them nothing was identified and the address is shared, so this is a question for the host. It is a reason to confirm what is listening, not a finding that any record was exposed.

Internet addresses flagged by a public watchlist

Four scored websites, one on the public list and three on the private, use an address carrying a raised public reputation finding. **A listing describes an address, and an address may be shared with organisations having nothing to do with the hospital.** Confirm privately whose address it is before treating any of the four as an institutional priority.

6. Stolen logins: three different problems

The credential source answered for **248 of the 257 scored websites**. The remaining nine are unknown, not clean. What came back divides into three populations that must never be added together.

Whose login it is	Websites	Records	Who would act on it
Staff identities associated with the organisation itself	83	684	The organisation
Organisational logins on outside services the organisation's own accounts on other platforms	102	1,392	The organisation
Members of the public with an account people who registered for something	154	4,212	The individual

Four things these numbers are not

They are not a single exposure figure. Adding 684, 1,392 and 4,212 produces a number that describes nobody. Three different groups of people, three different consequences, three different people who can do something about it.

They are not evidence that any hospital system was breached. Most records of this kind come from a breach at an unrelated service, or from software that steals passwords off an individual's own device. Neither says anything about the hospital's own systems.

They are not patient records. The third row counts people who held an account on some service, as the supplier classified them. **This report makes no claim that any clinical or patient record was exposed**, and nothing in the evidence would support one.

They are not verified roles. Staff, outside service and public account are the supplier's own classifications, not confirmed employment or clinical records.

The related observations

Separately from the three populations above, **56 websites carry a staff-identity finding from password-stealing software and 83 carry an internal web address exposed by the same route**. A further 74 have a login page for a service reachable from the open internet, and 27, all on the public list, carry an older credential exposure.

Never test a recovered credential. Confirm the person, the service and the current access against the organisation's own authorised records first, in restricted handling, and then follow the normal account process. Named evidence stays in the restricted application and is not published.

7. Forged email: the largest gap in this report

A domain can publish an instruction telling every receiving mail system what to do with a message that claims to come from it but did not. Without an enforcing instruction, receiving systems have nothing telling them to refuse or quarantine a forged message. It applies to the exact domain only, not to lookalike domains.

What the domain instructs	Public list	Private list	Together
Refuse the message	73	17	90
Send it to spam (quarantine)	14	6	20
Monitor only (no action requested)	16	8	24
No record published that we could find	117	6	123

147 of 257 websites, 57%, do not enforce a policy: 123 publish no record we could find and 24 only monitor. On the public list that is 133 of 220, and 117 of those publish no record. On the private list it is 14 of 37. The other 110 enforce: 90 refuse and 20 quarantine.

Why this one is worth doing first

It costs nothing to publish, and it is the main control that tells receiving systems to refuse or quarantine forged mail from your exact domain. **It does not stop lookalike domains.**

It needs care. Moving straight to refusal without first listing every legitimate sender, including the appointment system, the pathology notifier, the payroll provider and the mailing-list tool, will stop real messages. The order is: list the senders, monitor, quarantine, then refuse.

The supporting email records	Have it	Do not
A list of who may send email in the organisation's name	171	86
A signature on outgoing email	124	133
A request for reports on email encryption	58	199

The signature check samples the usual places a signature is published, so a missing one is not conclusive. A further 232 websites carry a finding about the policy that protects email while it travels, which is a supplier-side setting on most of this estate.

8. Cookies, consent and tracking

This concerns what happens on a visitor's own device. On a hospital website that visitor is often someone looking up a condition, a service or a department, which makes what gets recorded about the visit a more sensitive question than it would be elsewhere.

Read this before the number

Our privacy check scored **61.8% of the points it was able to assess**, 2,680 of 4,337. **That is not 62% compliant with anything.** It scores a handful of visible technical settings and nothing about how an organisation actually handles personal information, and it cannot establish which privacy law applies to any institution.

What we looked for	Yes	No	Could not tell	Why it matters
Insecure requests sent to the secure site	204	12	41	Anything typed into an insecure page can be read on the way.
No part of the page arriving unprotected	211	17	29	A padlock is weakened if parts of the page arrive unprotected.
A privacy notice we could find	121	78	58	A visitor should be able to find out what happens to their data.
A cookie consent banner we could see	62	166	29	The law governs what may be stored on a device before a person chooses.
A way to report a security problem	58	164	35	Someone who spots a flaw needs a way to tell the organisation privately.

This is the privacy check's view of insecure requests and differs from the redirect check in section 9 (14 not redirected, 39 unknown).

Of the 228 homepages we could read, 166 showed no cookie consent banner, 73%. Tag-management or tracking software was present on 148, and in 146 of those it was the same single tag-management product. Twenty-nine homepages were not readable at all, so their consent and tracking state is unknown rather than assumed clean.

This check identifies software present in the page. It does not establish when that software runs, what personal data it processes, whether consent was legally required, or whether a consent control actually works. A site that genuinely sets no non-essential storage needs no banner, so the 166 is a list to check in a browser rather than 166 breaches.

8 continued: what the websites store, and who can report a problem

What the websites store on a visitor's device

What the homepage stored on the visitor's device	Websites	Cookies
Websites where we could see this at all	224	259
Stored at least one cookie	126	259
A cookie that other websites can make the browser send back	87	138
A cookie the browser may send over an insecure connection	46	67
A cookie that scripts on the page can read	40	53

Cookie state is unknown for 33 websites. Only what the real homepage sets is counted, so a bot-protection or error page answering in its place is excluded. We see only the cookies a website sets as the page is served, not everything a browser stores, and we cannot tell whether a given cookie is essential. Each of these is a question to ask about that cookie's purpose rather than a fault in itself.

The row that stands out, and the one that should worry a board less

Fifty-eight of the 222 websites that could be checked offer a findable way to report a security problem privately; 164 do not and 35 could not be checked. Someone who finds a weakness at one of those has no obvious route to tell the organisation, which is how problems surface publicly instead. It is a short page and it costs nothing.

By contrast, an absent consent banner on a site that stores nothing may be perfectly lawful.

The consent figure is the one most likely to be overstated if read as a compliance count, which is why it is reported as a list to verify in a browser.

The ability to report a problem privately appears twice in this report, once as a published file giving the address and once as a findable route on the site. They are two different checks of the same idea and the counts differ slightly for that reason: 61 of 257 publish the file, and 58 of the 222 we could assess offer a findable route.

9. Card payments and encryption

Hospitals take card payments for car parking, outpatient charges, private treatment, retail and donations. The card industry standard applies wherever card data is handled or a system connects to one that does, and that is decided by payment flows rather than by what kind of organisation it is.

Two things this section is not

Our card-payment check scored **90.1% of the points it could assess**, 3,979 of 4,414. Controls that could not be verified were removed from the calculation rather than counted as passes, which is exactly why that number is high. **It does not mean 90% compliant**, and every result in the estate is partial.

And **this is not the scan the card standard requires**. That standard requires quarterly scans by a company approved by the card industry. We are not one, and nothing here can be submitted as evidence of compliance.^[4]

What we checked, on 257 websites	Fine	Problem	Could not tell
No sign of the old Heartbleed flaw	257	0	0
Certificate valid and in date	256	1	0
Only up-to-date encryption accepted	245	12	0
No reported software flaw	243	14	0
No database port answering from the open internet	235	22	0
Visitors sent to the secure version of the site	204	14	39
Browser told never to try the insecure version	172	55	30
A file saying where to report a security problem	61	161	35
Cloud storage, and two email-server checks	0	0	257

The 22 databases are the same cases reported in section 5 across both rows, counted once. The direct mail-server conversation ran on 110 websites, yet both email-server checks are unanswered for all 257, so that conversation does not collect what this calculation needs; the source does not record why. **The last row is the honest part of this section:** three controls could not be answered for a single website in the estate, and they were excluded from the calculation rather than recorded as passes, which is precisely why the headline percentage is high. **An unknown is never a pass.**

Encryption settings are the largest practical job here

55 of the 227 measurable responses do not tell the browser to stop trying the insecure version, which means the first request of every visit is exposed even where a redirect exists. A further 14 do not send visitors to the secure site at all, and 12 still accept outdated encryption. Certificates are the bright spot: one fault across 257 websites, and no Heartbleed evidence anywhere.

10. The records behind the web address

What it is	Have it	Do not	What it does
A list of who may send email	171	86	Tells receiving mail systems which servers are allowed to send email in the organisation's name.
A signature on outgoing email	124	133	Lets a receiving system prove an email really came from the organisation. Sampled, so a missing one is not conclusive.
A file saying where to report a problem	61	161	A published file telling a researcher where to report a weakness privately. 35 unknown.
A request for reports on email encryption	58	199	Asks other systems to report failures in protecting mail in transit.
A list of who may issue the certificates	5	252	Names the companies allowed to issue the padlock certificate for the address, so anyone else's is refused.
Signed answers about where the domain lives	1	256	Lets a system that checks signatures spot a forged answer to 'where is this organisation's website?' and refuse it.

The two records that protect the organisation's own identity are barely used anywhere. One website in 257 publishes signed answers about where its domain lives, and across the 220 websites on the public list the figure is zero. Most public-list websites sit under a shared NHS domain, where this protection depends on the parent domain being signed first, so this zero is largely a matter for the registry rather than for each organisation. Five have a list of who may issue their certificates.

The check records whether signing is published, not whether it works end to end: a signature only protects a visitor whose resolver actually checks signatures, and a signed child domain still needs its parent domain signed above it. Validation was not tested. A gap this uniform is a sector fact rather than 256 individual failures, and it is the clearest candidate in this report for coordinated guidance through a national body or a shared supplier rather than 257 separate projects.

Supplier dependency

What we observed	Public list	Private list	Together
No protection on email while it travels	197	35	232
A broader concentration of services on one supplier	152	29	181
Everything resting on one company for the address records	6	33	39

The last row is the sharpest difference between the two lists: **3% of the public list against 89% of the private.** Concentration is usually the right choice for an organisation without technology staff, and it is not a fault. It becomes a governance question when nobody has written down what would stop, and who would be told, if that single supplier had an incident.

11. The other web names, and fifteen loose ends

Separately from the scan, a review looks at the other web names that turn up in each organisation's past certificates and hosting. It ran on all 257 scored websites, and it is the part of this evidence most easily misread.

Review of the other web names	Count
Scored websites reviewed	257
Other web names found	5,521
Other web names we looked at	4,621
Beyond the limit we set, never looked at	900
Looked for, but nothing came back	119
Websites where certificate history was present	181
Web names seen only in certificates that have expired	2,286

Fifteen published names point at something that is not there

Across **two websites, fifteen published web names point at a target that no longer exists**. The organisation still tells the internet that the name belongs somewhere, and nothing is at the other end.

This is the one finding in the report where doing nothing has a specific downside. A name in that state can sometimes be claimed by somebody else, who then has a web address that genuinely belongs to the organisation and can be used to make a message or a page look authentic. Two websites is a small number and the fix is usually one deletion each, which is why it is worth closing rather than scheduling.

2,286 web names, across the 181 websites where history was collected, appear only in certificates that have expired. Those are names that once had a certificate and may since have been abandoned: an old campaign site, a service landing page, a portal retired without anyone removing the name. None of that is a security finding on its own, and the right response is an inventory question rather than remediation. A name nobody remembers owning is a name nobody is watching.

Certificate history was present for only 181 of the 257 reviews, so 76 ran without it. Nine hundred names found were beyond the limit we set and never examined, and 119 could not be assessed. The review tests what each name publishes; **it never tests who controls any resource**, and a name that only ever appeared inside a catch-all certificate, or never had one at all, does not show up here.

12. What we could not read

Three different limits apply, and they are not the same thing. Keeping them apart is what stops a gap in evidence being reported as a finding.

Situation	Websites	What it means
Minimum score	30	The website's browser instructions could not be read. Other checks were read, subject to the partial rows below. The score is published as a floor and can only be higher.
Scanned, but no score could be given	12	A whole group of checks came back with nothing, so no score could be published. Excluded from every average.
Scan did not run or did not finish	0	Nothing. Every tracked domain completed a primary scan inside the stated window.

Refusing a scanner is not evidence of a weakness

Where a site turns away an unknown automated visitor, it is exercising a security control. The consequence for us is that we cannot read the instructions that site sends to a browser, and those are one of the things the score is built from. **The source does not record the reason per website**, so these 30 are an evidence limit rather than a finding.

The twelve scanned without a score

Their score is **unavailable, never zero**, and none of their evidence appears in any count table in this report. The group of checks that came back empty was the safety instructions given to browsers on 7, email delivery policy on 6, the website certificate on 6 and internet exposure on 1, and those overlap. Independently supported findings were still retained for them, most often about stolen logins and the records that point the internet at the website.

What else we could not see

Group of checks, on 257 websites	Measured	Partial	Not available
Encryption on the connection	256	0	1
The safety instructions the site gives browsers	227	0	30
The other web names linked to the organisation	143	114	0
Talking directly to the mail server	110	0	147 skipped
Backup arrangements for email	98	159	0
Domain registration details	44	0	213
Looking for open cloud storage	0	257	0

Measured and partial on the web-names row mean two different things. Measured, 143, means the review ran with every part of its evidence present. Partial, 114, means it ran but something was missing: for 76 the certificate history reported in section 11, for 24 names found beyond the limit we set. Those two overlap and do not account for all 114, so the source publishes no full breakdown of the rest. Section 11 counts the same estate differently: certificate history present on 181, absent on 76, together 257.

The domain registration row is the one to understand. It was readable for all 37 private-list websites and for only 7 of the 220 on the public list, because registration details were not available for these domains, which is a property of the registry rather than a finding about any organisation. Cloud storage could only be partly checked everywhere, and the direct mail-server conversation ran on 110, leaving mail encryption and relay status unknown for the other 147 websites.

13. What to do, in order

Ordered by consequence, not ease. These groups overlap heavily and are not a count of separate organisations. Most are a conversation with an existing supplier rather than a project.

1	Close the fifteen loose ends Two websites publish fifteen names pointing at targets that no longer exist. This is the cheapest item in the report and the only one where leaving it open creates a new route to impersonate the organisation. Confirm each name is genuinely retired, then remove the record.
2	Check the five matches against the criminal-use list privately Six entries across five websites, drawn from 511 observations, none confirmed by a supplier. Confirm each against the software actually running and the setting affected. This establishes the importance of checking, not that anything was compromised.
3	Move the 147 domains towards refusing forged email The largest and cheapest win in the report. Begin with the 123 publishing no record we could find, then the 24 set only to monitor. Inventory every legitimate sender first, including appointment, results and payment systems, then watch, then send to spam, then refuse. Look at the same time into the 133 with no signature on outgoing email and the 86 with no list of who may send it.
4	Confirm what is behind the 22 database ports reachable from the internet Three counted against the organisation and nineteen on shared addresses. Ask the hosting company which services are meant to be reachable before restricting anything. Then the wider 31 open-door cases, keeping the shared-address work separate.
5	Validate the staff and organisational logins in restricted handling 684 records across 83 websites, and 1,392 organisational records on outside services across 102. Confirm the person, the service and current access against authorised records before any reset. Keep these separate from the 4,212 public account records, which belong to individuals and have a different owner. Never test a recovered credential.

13 continued: what to do, in order

6	Fix the encryption settings, then the browser instructions 55 websites without the instruction that stops the browser trying the insecure version, 14 not sending visitors to the secure site, 12 accepting outdated encryption and one certificate fault. Check first that the four watchlisted internet addresses are really the organisation's own.
7	Open the 166 websites in a browser and watch what happens The consent question cannot be settled from outside. Ten minutes per site, no specialist needed. A site that genuinely stores nothing needs no banner, so this will clear many of the 166 rather than add to the work. Review the 78 without a findable privacy notice alongside.
8	Publish the three one-off records, once, across the sector 256 websites publishing no signed answers about where the domain lives, 252 with no list of who may issue their certificates and 161 with no file saying where to report a problem. At these rates this is coordinated guidance through a national body or shared supplier, not 257 separate projects.
9	Write down the supplier dependency 39 websites rest on one company for their address records and 181 carry a broader concentration observation. Establish who at each organisation would be told, and what would stop, if that supplier had an incident.
10	Close the remaining evidence gaps 12 websites remain unscored, 30 carry minimum scores, 76 reviews lack certificate history and 900 other web names were never examined. Do not read absent evidence as a fix.

Conclusion

This estate's problem is not a tail of bad websites. It is that the middle sits low. A full-score average of 65.1, with 31 of 257 websites in the top band and 193 in the middle, describes an estate where ordinary settings have gone unreviewed almost everywhere rather than one with a dangerous minority. Thirty-three websites sit in the lowest band, eleven of them on minimum scores that may rise, which is a real group but not the shape of the problem.

The consequential findings are few relative to the estate and they do not follow the score bands. Five matches against the criminal-use list, 22 databases reachable from the internet, 83 websites with staff logins in stolen-password collections. A website can sit in the top band and carry any of these, which is why the bands and the consequential findings are reported separately here and should be governed separately too.

The three things that distinguish this sector

Forged email is the defining exposure, not a hygiene item. 147 of 257 websites do not enforce a policy against a message forged in their name, and only 90 instruct outright refusal, in a sector whose messages are opened and acted on faster than almost any other organisation's would be.

The records that protect the organisation's own identity are effectively unused. One website in 257 publishes signed answers about where its domain lives, and none at all on the public list. Most public-list websites sit under a shared NHS domain, where this protection depends on the parent domain being signed first, so this zero is largely a matter for the registry rather than for each organisation. These are one-off publications, and at that level it is a national job rather than an institutional one.

The clean results are real and narrowly bounded. No leak-site reference anywhere, one certificate fault across 257 websites, and no Heartbleed evidence. Each is worth stating, and each is bounded by what was collected rather than by what exists.

And one structural point. The two lists fail in opposite directions. The public list is spread across many suppliers but weak on the control that stops impersonation. The private list is markedly better at refusing forged email but has nearly nine in ten of its websites resting on a single company for the records that point the internet at them. Neither shape is wrong. Both are worth writing down.

Four things this report does not claim. It does not say any organisation has been broken into: we could not see inside one and did not try. It does not count hospitals, because a domain is not an institution and groups share websites. **It makes no clinical, patient-data, safeguarding or compliance judgement,** having assessed none of the systems those depend on. And it is **not a complete picture of the sector:** evidence gaps are stated throughout wherever they affect a figure.

One thing it does claim. The distance between this estate and a good score is made up almost entirely of cheap, externally verifiable, low-risk work: a published email instruction, a browser instruction, three one-off domain records, fifteen deletions, and a word with the host about what is listening. None of it requires a security programme and most of it requires one supplier conversation.

About this report

Produced with MyDomainRisk, a tool developed by Huro Data Technologies Ltd, which has a commercial interest in this analysis and declares it here so readers can weigh it. Every observation was made from outside using public websites, domain records, certificates, cookies and public threat intelligence. Nothing was logged into, no password was tried and nothing was exploited. It is not certification, an audit, a penetration test, legal advice or a regulatory judgement.

How to read the figures

Websites, not hospitals

Every figure counts **tracked domains**. A group may run one website for many sites, and one organisation may use several domains, so the number of institutions behind these 257 websites is not established here. **Nothing in this report is a census of hospitals**, a verified register, or evidence of which organisation operates any particular website. The public and private labels are the names of two tracked lists and do not independently establish funding, ownership or endorsement.

Bands, and what a score is

Good is 75 and above, **needs improvement** 50 to 74, **at risk** below 50. One serious finding can hold a website below a band however much else is right, and a top-band website can carry a consequential finding from section 5. **A score is a floor, not a verdict**: it records what was visible from outside in one scan window, not internal controls, patching, training, clinical systems, payment verification or incident readiness.

Full scores and minimum scores

227 websites have a full score and 30 a minimum score, published where the browser instructions could not be read. **A minimum score counts every unread item as absent**, so the true figure can only be the same or higher, and it is never compared with a full score or used to rank an organisation. Two averages are given: **65.1** on full scores alone and **at least 63.8** counting floors at their floor. The two describe different subsets and are not a before-and-after comparison.

Counting, and whose problem an open door is

Every figure counts **websites, not problems**. One unpatched component can produce dozens of observations, which is why websites are counted and observations given only as context: 14 websites, not 511 observations. Rows overlap and must never be added: the 3 databases counted against the organisation sit inside the 10 open-door cases counted against the organisation.

A case counts against the organisation where a service was identified on the open door, or where the address was observed only for that organisation; separately where nothing was identified and unrelated organisations share the address. **Separate keeps its full weight**. Of 31 observed, **10 count against the organisation and 21 are reported separately**.

The two percentage scores are not compliance rates

The privacy check scored 61.8%, 2,680 of 4,337, and the card-payment check 90.1%, 3,979 of 4,414, **of the points each was able to assess**. Controls that could not be verified were removed from the calculation rather than counted as passes, which is why the payment figure is the higher. Neither is an audit or a legal conclusion, and every card-payment result in this estate is partial.

Notes

- Every figure describes a single scan window, 13 October 2026. No trend is reported and none should be inferred.
- The three credential populations are quoted independently and never summed. A zero means a check ran and found nothing, never an absent answer, and an unscored domain is never a zero-score domain.
- The two lists were reconciled against the combined evidence file across 94 categorical measures with no mismatches.
- This first edition reports a single scan window. Evidence gaps are stated throughout wherever they affect a figure. We publish counts and distributions, never a league table.

Sources and notes

Source	What it provided
UK Hospitals, public estate report, 13 October 2026	231 tracked domains, 220 scored. The larger of the two source documents, and the basis for every public-list figure in this report.
UK Hospitals, private estate report, 13 October 2026	38 tracked domains, 37 scored. The basis for every private-list figure.
Combined evidence file frozen aggregate, 13 October 2026	The consolidated figures, used only after the two sources, and checked against them on 94 measures.
MyDomainRisk mydomainrisk.com	The assessment tool itself, developed by Huro Data Technologies Ltd (hurodata.com).

[1] Scan method identity recorded against all 257 published scores: **security-2026-09-evidence-v7**. Membership and per-case hosting evidence were read in a single read-only transaction.

[2] Frozen snapshot 13 October 2026, 15:54 UTC. Scan rows created between 09:07 and 09:37 UTC the same day. Delayed certificate evidence present at the snapshot is included without redating the primary scan; the latest addition was 15:29 UTC.

[3] Evidence fingerprints: public list **2911e1db...1e0ca**, private list **09b88629...41d7a**, combined **b9754584...6328e**.

[4] The card industry standard requires quarterly scanning by an approved scanning vendor. Huro Data Technologies Ltd is not one, and nothing in this report may be submitted as evidence of compliance with that standard.

[5] Stolen-credential evidence is supplied by a commercial intelligence source and reflects that supplier's own classification of each record. Records may originate from a breach at an unrelated service or from password-stealing software on an individual's device.

How the scans were run

Our scans use bounded, external, non-intrusive checks. They do not attempt exploitation, credential testing or intentional modification of target systems. No login, password testing, exploitation, internal assessment, payment-flow mapping or legal audit was performed. This document reuses stored evidence and initiated no scan.

Provenance and handling

What is in this document, and what is not

This report contains aggregate counts only. **It names no organisation**, and it contains no password, no personal credential record, no raw supplier record and no instruction that would help anyone exploit anything. Named evidence remains in the restricted application and must be reconciled with authorised identity records before it is used.

Handling rules that apply to anything derived from this work

- **Never test or attempt to use a recovered credential.** Confirm the person, the service and current access against the organisation's own authorised records, in restricted handling, then follow the normal account process.
- **Recheck every named finding immediately before any disclosure**, and disclose to the affected organisation first. An observation taken from stored evidence may already have been fixed.
- **A leak-site listing is a criminal group's claim, not a verified fact**, and must never be published against a named organisation without validating both the entity and the date. No such listing was retained for either list here.
- **Treat a shared-address finding as the hosting company's question first.** Confirm tenancy privately before any restriction, disclosure or escalation.

What the evidence cannot establish

Nothing here establishes that any organisation was compromised, that any record was accessed, or that any clinical, patient or internal system is affected in any way. Those systems sit behind authentication and are invisible to an external observer. Equally, nothing here establishes that an organisation is safe: an absence of evidence is a limit of the method, not a clean bill of health.

Membership of the two tracked lists has not been independently validated, the public and private labels are administrative, and a domain may serve a group, a single site or neither. Differences between the two lists are differences between two tracked lists, not a representative comparison of public and private healthcare.

Status of this edition

This first edition reports a single scan window. Twelve websites remain unscored, 30 carry minimum scores, 76 reviews lack certificate history and 900 discovered web names were never examined. These gaps are stated wherever they affect a figure.

Produced with MyDomainRisk by Huro Data Technologies Ltd, which has a commercial interest in this analysis. Not certification, an audit, a penetration test, legal advice or a regulatory judgement. Validate material findings privately before acting on them.