



How UK Banks and Building Societies Look to an External Scan

A September 2026 baseline of 152 websites across the UK's banks and building societies, written for people who do not work in cyber security

What this report is

We looked at every website in our tracked estate from the outside, the way any customer could, and recorded what was visible. We did not log in anywhere, try any password, or touch anything inside an organisation's systems. This report says what we saw, what it means in plain English, and what is worth doing about it.

Sources	Two separately tracked lists, reported separately: 110 bank websites and 42 building society websites .
Combined	No website appears on both lists , so the two together represent 152 distinct websites . This was checked rather than assumed.
Evidence	One assessment, 26 September 2026, every website scanned the same way. No scan failed . Four produced no usable score and are excluded, three banks and one society. That leaves 148 reported : 115 with an exact score and 33 with a floor score.
Limits	Nothing here says any organisation has been broken into. Nothing here is a legal or compliance judgement. Everything here is what was visible from outside on one day.
Prepared by	Huro Data Technologies Ltd / MyDomainRisk
Date	26 September 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. One organisation is named on a ransomware leak site**, meaning a criminal group has published its name. It needs matching against the right legal entity and date before anyone draws a conclusion, but it is the item to check first.
- 2. Four organisations have a database port answering the public internet** on an address of their own, inside six with some risky port answering. Six more show the same ports only on addresses shared with unrelated organisations.
- 3. Twenty-seven organisations cannot stop email being forged in their name.** For a retail financial brand that is the control with the most direct customer consequence.
- 4. Stolen staff passwords are circulating for 22 organisations, 147 records**, about seven each. These are keys to an organisation's own systems and it can take them back today.

This is a strong estate by the standards of what we measure, and the consequential list is short. None of it means an organisation has been broken into, and every item needs private validation first. **Four have a confirmed lowest-band score**, and one more sits there on a floor figure only.

What we found, worst first

What it is, in plain English	Why it matters	Websites
Named on a ransomware leak site	A criminal group has published the name. May be historical or a similar entity. Validate first.	1
A database port answering, on the organisation's own address	Databases hold customer records. They belong behind the website, not beside it.	4
Software weakness read from a version number	82 observations, none confirmed and none on the public list of weaknesses being exploited.	5
Any risky port answering, on the organisation's own address remote access, file transfer or admin tools	Doors that normally face inwards. The 4 database ports sit inside this 6.	6
The same risky ports, on an address shared with others	Nothing identified running. Separate because only the provider can close it, not lesser.	6
Certificate or encryption problem	Seven accept outdated encryption; one has a certificate fault.	8
Stolen staff passwords in criminal collections	147 records, seven on a typical affected organisation.	22
Cannot stop email being forged in the organisation's name	A forged message from a bank is acted on quickly, often by those least able to absorb a loss.	27
Customers' own account details in circulation	132,483 records. Nobody at the organisation can reset these; only the customer can.	102

Rows overlap and must never be added; the 6 organisation-level and 6 shared cases are the 12 observed. Three organisations also have an address carrying a poor public reputation.

Executive summary, continued

The finding that is not here, and why it is worth stating

Nothing in this estate is on the actively-exploited list

Five organisations run software whose version number suggests a recorded weakness, across 82 observations. **Every one of the 82 was checked against the public catalogue of weaknesses criminals are known to be using in the real world, and none of them is on it.**

That distinction matters. A weakness recorded somewhere is a validation exercise. A weakness on the exploited list is a same-week job. This estate has the first and none of the second, and none of the 82 was confirmed by the source in any case: a version number is what the software says about itself, and suppliers routinely fix a fault without changing it.

Stolen passwords, and what the figure describes

What a criminal collection actually is

When a virus infects somebody's computer it copies every username and password saved in the browser, and those get bundled up and traded. If a member of staff ever saved a work login on an infected machine, their work address and password can end up on a list anyone can buy. We never see or try those passwords, only that records exist against a domain and how many. **We cannot tell whether any still works**, only that the material is out there.

147	7	96
staff password records across 22 organisations	records on a typical affected organisation: narrow, not deep	organisations with nothing in either actionable population

Why some organisations have a floor score rather than an exact one

The change that matters for reading every table in this report

Some websites answer an automated visitor with a protection page rather than the real homepage. That is a security control working, not a fault, but it means we cannot read the instructions the site sends to a browser, and those form part of the score. For **33 websites** that was the only gap, so we publish a **floor score**: every unread instruction counted as absent, "at least this much", so **the true figure can only be the same or higher**.

Scores	Websites	Average
Exact scores	115	77.0
Floor scores, counted at their floor	33	63.0
All published figures	148	73.9

Use 77.0 for the estate's measured position and 73.9 for the most cautious one. Were every unread instruction in place, those 33 websites would average 83.8 rather than 63.0, which would lift the whole estate. Four further websites produced no usable score and are excluded from every figure here.

The two source lists sit close on score, 76.4 for banks and 78.4 for building societies, but **will not carry a league table**: the lists differ in size and in how much evidence came back. Sections 4 and 5 report them separately; section 6 consolidates.

Why this matters for these organisations in particular

A bank or building society holds two things that make it an unusually attractive target. The first is money, in the most direct sense there is. The second is a kind of trust that is both unusually strong and unusually easy to borrow: a message that appears to come from somebody's bank, about their own account, gets acted on quickly and often by the people least able to absorb a loss.

That is why the email finding matters more here than the raw number suggests. A forged message in a bank's name is not a nuisance. It is the opening move in a fraud against that bank's customers, and 27 of the 148 websites reported cannot instruct a receiving mail system to refuse one.

A way of picturing what we did

Somebody walking down your street

They never go inside any building and never try a door. They walk past and notice what is visible from the pavement: a window left open, a side gate unlatched, a spare key under the mat where everyone knows to look.

None of that means the building has been burgled. It means somebody passing by can see opportunities, and most of them take a minute to put right.

That is exactly this assessment. When this report says staff passwords are circulating, that is the spare key under the mat: it does not mean anyone has used it, but somebody else knows where it is. When it says a port is answering that need not be, that is the side gate nobody remembers leaving unlatched. And when a website turns our scanner away, that is a locked gate doing its job, which is why those organisations get a floor score rather than a mark against them.

What we could not see, which is most of it

Standing in the street tells you nothing about the alarm, the vault or the locks on the internal doors. These organisations are among the most heavily regulated and most heavily invested in security of any in the country, and almost all of that is invisible from where we stand: fraud monitoring, transaction controls, staff training, incident response, everything inside the building.

So nothing in this report is a verdict on any organisation's security. It is an account of what is showing from outside, which is useful precisely because it is what a criminal sees too.

Three cautions that apply to every number here

A finding is not a break-in. Nothing establishes that any organisation has been compromised, and every serious item needs it checked privately before it is treated as fact.

A floor score is not a bad score. Thirty-three websites carry one and each could sit substantially higher. They are never compared with exact scores or used to rank an organisation.

Nothing here is a compliance judgement. We observe technical settings. We cannot see contracts, retention, training or how customer information is actually handled.

1. The words used in this report

This section exists so that nothing later needs a specialist to interpret. Each term is given the way it is used here, not as a formal definition.

Term	What it means in this report
Cyber posture	The condition of an organisation's defences at a point in time, as seen from outside. If a stranger looked at it from the public internet, what could they learn?
Criminal credential collection	A traded list of usernames and passwords, mostly harvested by viruses from the browsers of infected computers. Also called infostealer data.
Ransomware leak site	A site where a criminal group publishes the names of organisations it claims to have attacked, usually to pressure payment. A listing is a claim, not a verified fact.
Version-inferred weakness	A weakness guessed from the version number a piece of software reports, not confirmed by testing. Suppliers often fix a fault without changing the version, so these overstate.
Known-exploited list	A public catalogue of weaknesses criminals are known to be using in the real world. Far shorter and more serious than a weakness merely being recorded somewhere. Nothing in this estate is on it.
Floor score	A score published as "at least this much" where one part of the check could not be read. Every unread item counts as missing, so the true figure can only be the same or higher.
Forged email, and the setting that stops it	Anyone can put an organisation's name on an email. A published instruction called a DMARC policy tells receiving mail systems to refuse those forgeries. It can be set to refuse, to send to spam, or to do nothing.
Secure connection	The padlock. A page served over HTTPS is encrypted in transit; one served over plain HTTP can be read or altered on the way.
Risky port answering	Something responding on the public internet where a service normally faces inwards: a database, a remote-access tool, a file-transfer or admin console.
Shared provider address	One internet address used by a hosting or web-security company for many customers at once. What answers there may be the provider's equipment rather than any one organisation's.
Cookie consent mechanism	The banner that asks permission before storing things on a visitor's device. The law governs what may be stored before a choice is made.
Website instructions to the browser	Settings a website sends telling the browser how to treat the page. Their absence is a review point, not a hole. These are the items a floor score counts as missing.

Security is cumulative, which is why there is no single fix

A strong position is rarely one product. It is the combined effect of many ordinary settings done properly: strong sign-in, sensible configuration, secure email, patched software, limited exposure to the internet, knowing what suppliers hold, a route for someone to report a problem, and a plan for when something goes wrong.

2. Where these organisations sit in the rules

Banks and building societies sit inside the most explicit operational-resilience regime of any sector we measure, and it speaks directly to what this report observes.

Operational resilience

Firms operate under the Prudential Regulation Authority's supervisory statement and the Financial Conduct Authority's operational resilience rules. Both require a firm to identify its **important business services**, set a tolerance for how much disruption each can take, and understand the third-party dependencies that could cause it.^[1]

Why that matters for this report specifically

A customer-facing website, a trustworthy email channel and the staff identities that control both are dependencies of an important business service. They are not IT housekeeping.

Read that way, the findings here are not a separate technical list. A database port answering the internet, an unenforced email policy or a set of circulating staff credentials is a dependency with a known weakness, and the regime already expects it to be recorded, owned and tolerated or fixed. The same applies to a shared provider address, which is a third-party dependency behind a customer-facing service in the most literal sense.

Data protection applies directly too

Firms are legally responsible for the personal information they hold about customers under the UK GDPR, and separate rules govern what may be stored on a customer's device when they visit a website. The Information Commissioner's Office regulates both and publishes the guidance that decides whether an arrangement is lawful.^[2]

What this report can and cannot say about that

We can see whether a website forces a secure connection, whether a privacy policy can be found, whether a cookie banner is present in the page as delivered, and what the website stores. **None of that establishes whether an organisation is complying with anything.**

We cannot see the lawful basis for any processing, how long information is kept, what contracts are in place with suppliers, whether information leaves the country, how requests from customers are handled, or any of the training and internal controls that make up most of the regime.

How this differs from the statistics you may have seen

Most sector cyber figures come from surveys, which ask organisations what happened to them and what they have in place. This study asked nobody anything. Every number is an observation taken from outside using the same checks on every website. A survey tells you what a sector believes about itself; this tells you what it is showing to the world.

3. What we measured, and how to read the numbers

Each website was checked against the same set of external tests and given a score out of 100 where enough evidence came back. **Good** is 75 or more, **needs improvement** is 50 to 74, and **at risk** is below 50.

The two lists do not overlap, and that was checked

No website appears on both the bank list and the building society list. The two populations are genuinely separate, so the combined figure of **152 distinct websites** is a true sum rather than a count of list entries. That is stated because it cannot be taken for granted; in other sectors we have measured, two lists that looked different shared a third of their members.

Coverage	Banks	Building societies	Combined
Tracked websites	110	42	152
Appearing on both lists	0	0	0
Scans that failed	0	0	0
Excluded, completed without a score	3	1	4
Exact score	83	32	115
Floor score	24	9	33
Reported	107	41	148

Because the lists share nothing, these columns genuinely add: 107 and 41 make the 148 reported, and 83 and 32 make the 115 exact scores.

Four rules this report follows throughout

Counts are of websites, not of problems. An organisation with five instances of the same issue counts once. One unpatched piece of software can generate dozens of technical observations, which is why we report affected websites and not observations: five organisations, not 82 observations.

Rows in a table overlap and must never be added. The clearest case: 6 websites have a risky port answering on their own address and 4 have a database port, but the 4 are inside the 6. The total is 6 websites, not 10.

Every percentage names its denominator. Most are shares of the 148 reported websites. Some checks could only answer for fewer, and those tables say so, in the form 71 / 115 readable with the rate beside it. Where a check returned an unknown, the denominator is the number that answered, because an unknown is not a pass.

Stolen passwords come in three separate populations, covering different people and fixed by different parties. Section 8 sets them out. They are never combined into one headline number.

4. Source one: the 110 bank websites

The first list is the UK's banks, counted by website rather than by legal entity or brand. Of the 110, three produced no usable score and are excluded; no scan failed. That leaves **107 reported**, 83 with an exact score and 24 with a floor score.

Measure	Result
Tracked websites	110
Scans that failed	0
Excluded, completed without a score	3
Reported	107
Exact scores	83
Floor scores	24
Exact-score average	76.4

An average of 76.4 describes a range, not a type

This list is the broader of the two by some distance. It holds household retail names alongside specialist lenders, challenger banks and the UK arms of overseas groups, and their web estates are built and run in very different ways. An average sits in the top band, but it is an average across organisations that have little in common operationally.

That is also why the floor scores matter here: 24 of these websites turn an automated visitor away, which is itself a defensive posture more common among larger institutions.

Beyond coverage and score, the finding-level evidence in this study is produced consolidated rather than per list, because the populations are small and splitting the findings would produce counts too thin to read. Sections 6 to 12 therefore report every finding against the 148 distinct reported websites.

What the score alone does and does not tell you

A score is built from many ordinary settings, so an organisation can sit in the top band and still carry one of the serious findings in section 7. **The bands measure hygiene. The serious findings measure consequence, and they are not the same thing.** An estate can have the best average we have measured and still have a database port answering the public internet.

5. Source two: the 42 building society websites

The second list is the UK's building societies, again counted by website. One produced no usable score and is excluded; no scan failed. That leaves **41 reported**, 32 with an exact score and 9 with a floor score.

Measure	Result
Tracked websites	42
Scans that failed	0
Excluded, completed without a score	1
Reported	41
Exact scores	32
Floor scores	9
Exact-score average	78.4

The cleanest population in this study

The building society list has **no website in the lowest band on either basis**, and its exact-score average of 78.4 is the highest of the two. On a list of 42, that is a meaningful result rather than a rounding artefact.

It is worth saying why this is not simply a story about size. Building societies are mutuals with comparatively simple, largely domestic web estates: fewer brands, fewer acquired systems, fewer overseas arrangements. A smaller surface is easier to keep tidy, and that shows.

The two-point gap is not a ranking

Seventy-eight point four against 76.4 is two points, on populations of 32 and 83 exact scores. The lists differ in size, in composition and in how much evidence came back, and a floor score drags an average down by an amount that depends on how many websites turn the scanner away rather than on how well those organisations are run.

The honest reading is that both lists sit in the top band and are close. What separates organisations here is not which list they are on but whether they carry one of the consequential findings in section 7, and those do not follow the list boundary.

6. The consolidated view: 148 reported websites

From here on every figure describes the 148 distinct reported websites. Because the two lists share nothing, this is a true combination rather than a merge.

Band	On exact scores 115 websites	Including floor scores 148 websites
Good, 75 and above	55 (48%)	55 (37%)
Needs improvement, 50 to 74	56 (49%)	88 (59%)
At risk, below 50	4 (3%)	5 (3%)

Exact average **77.0**, median **74**. Counting the 33 floor scores at their floor gives **73.9**, median **74**. Those floors average 63.0 and would average 83.8 if every unread instruction turned out to be in place. **The one extra at-risk result in the right-hand column is a floor figure and may sit higher.**

Nearly half are in good order, and that is the headline on hygiene

Fifty-five of 115 exactly scored organisations reach the top band. Fifty-six sit in the middle, where the gaps are real but routine. Four are confirmed in the lowest band.

This is the strongest distribution we have measured in a financial population, and it is worth attributing correctly: it reflects sustained investment in exactly the ordinary settings a score like this counts, not luck.

Where the two source lists sit inside this

Measure	Banks	Building societies	Combined
Reported websites	107	41	148
Exact scores	83	32	115
Floor scores	24	9	33
Exact-score average	76.4	78.4	77.0
Score sum, all published	7,851	3,082	10,933

The two sums add exactly, 7,851 and 3,082 giving 10,933 over 148 websites, because no organisation appears on both lists. The consolidated average is calculated from the 148 underlying scores, not by averaging the two list averages.

7. The findings that carry real consequence

These are the items that would matter on a Monday morning. They are few, and they do not follow the score bands: an organisation can sit in the top band and appear here.

Finding	Websites	Rate
Named on a ransomware leak site	1 / 148	1%
Internet address with a poor public reputation	3 / 148	2%
Database port answering on the organisation's own address	4 / 148	3%
Software weakness read from a version number	5 / 148	3%
of which carrying an entry on the known-exploited list	0 / 148	0%
Risky port answering on the organisation's own address	6 / 148	4%
Shared provider address, no owned service confirmed	6 / 148	4%
Certificate, protocol or encryption problem	8 / 148	5%
Known breach records or browser-safety warnings	0 / 148	0%

The single ransomware leak-site reference, and how to handle it

One organisation appears on a ransomware leak site. On this evidence it is the most consequential single row in the report, and one is a good number for a sector this size.

A listing is a criminal group's claim, not a verified fact. It can be historical, it can be duplicated across sites, and it can refer to a similarly named entity elsewhere in the world. The right response is restricted matching of entity, date and incident record with the incident response team. **It must not be published against a named organisation without that validation.**

The software weaknesses, and the list they are not on

Five organisations show software weaknesses on a reachable service, across 82 observations. **Every one was read from a version number and none was confirmed by the source.** A version match does not show that the vulnerable feature is in use, or that a fix has not been applied without changing the version, which suppliers do routinely.

All 82 were checked against the public list of weaknesses known to be exploited in the real world. **None of them is on it.** That is a genuine negative finding and it should be read as one: this is a validation exercise for five organisations, not an emergency, and the count to report is five organisations rather than 82 observations.

Ports answering: 12 cases, and whose address they sit on

Twelve cases were observed in all. **Six sit on an address used only by the organisation or its group,** four of them a database port, and six sit on an address shared with unrelated organisations and are reported separately. **A port answering is not the same as a vulnerability.** A service may be deliberately published, access-controlled, or run by a supplier. What it does mean is that the door is visible, and a database port in front of customer records warrants immediate review. The next page sets out which is which.

7 continued: whose address the port sits on

Each of the 12 cases was checked individually on two questions: did a named product or protocol actually answer on the flagged port, and do unrelated organisations resolve to the same address?

Where the flagged port answers	Websites	Shared with others?	Treatment
Web-security addresses dedicated to the organisation's own group only its own group names resolve there	4	No	Organisation-level
Hosting used only by the organisation only its own names and the hosting company's server name	2	No	Organisation-level
The same web-security provider, shared edge addresses	3	Yes	Reported separately
A second content-delivery network, shared edge address	1	Yes	Reported separately
One managed-hosting platform address at least 22 organisations in the wider scan data	1	Yes	Reported separately
Shared web hosting	1	Yes	Reported separately
All cases observed	12	6 shared	6 organisation-level

Rows are exclusive and sum to the total. **No organisation-owned service was confirmed on any of the twelve:** on none of them did a named product answer on the flagged port. The six on dedicated addresses stay at organisation level because an address used only by an organisation or its group is part of its internet presence, whoever operates the machine. The 6 organisation-level cases carry 34 port observations between them and the 6 separated cases a further 35.

Separate does not mean lesser

The six reported separately are not discarded. The ports may be the provider's rather than the organisation's own, but **an attack on or failure of a shared platform would affect every organisation on it at the same moment.** That is not a smaller risk than a single exposed database.

Under the operational resilience rules this is a third-party dependency behind a customer-facing service, exactly what a firm is expected to have identified and written down. A public reputation listing on a recognised shared hosting address is treated the same way and is not attributed to each organisation on it.

What each organisation should ask its provider is narrow and answerable: which of these services are meant to be reachable, how the platform is protected, and what the contract says about breach notification. Record the answer as supplier evidence rather than closing the finding on the provider's word. **An organisation's own servers behind the provider were not examined by this scan.**

8. Stolen credentials: three different problems

This is the largest body of evidence in the report and the most misunderstood. The source returns three separate populations, they concern different people, they are fixed by different parties, and **they must never be added into one number**. Counts were returned and reconciled for all 148 reported websites.

Population	Who it concerns, and who can act	Websites	Rate	Records
1. Staff	Accounts on the organisation's own domain. The organisation can fix these today by resetting the password, ending active sessions and enforcing phishing-resistant multi-factor sign-in.	22 / 148	15%	147
2. Identities elsewhere	Work email addresses used to sign in to other companies' platforms: identity providers, service desks, expense and recruitment systems. The organisation can act, but the reset happens on the other system.	48 / 148	32%	874
3. Customers	Members of the public with an account, whose own device was infected or whose password was reused elsewhere. Nobody at the organisation can reset these. Only the customer can.	102 / 148	69%	132,483
1 or 2	Deduplicated. Not 22 plus 48.	52 / 148	35%	n/a

Ninety-six organisations, 65%, have nothing in either of the two actionable populations. That is a measured result and a good one.

Why the order matters more than the size

The staff population is the one that matters and it is by far the smallest: 147 records across 22 organisations, roughly seven each. It is a key to the organisation's own systems, it issued that key, and it can take it back this afternoon.

The identities-elsewhere population follows closely, because a work address reused across several platforms is how one compromise becomes several. It is the larger of the two actionable populations, reaching 32% of organisations.

The customer population is nine hundred times larger and ranks last. It is not access to anything the organisation runs. It is a customer's own laptop that caught a virus, or a password reused on a site that was breached. It carries no weight in the score and nobody at the organisation can reset it.

The customer figure will be quoted wrongly, so state it carefully

132,483 records sounds like a breach of a bank. **It is not, and this report never presents it as one.** These are records tied to a domain because a customer used that service, harvested from the customer's own device or from somebody else's breach. They establish nothing about the organisation's systems.

What they do create is a customer-protection duty: the organisation cannot reset these, but it can detect reuse at sign-in, step up authentication and warn the people affected.

8 continued: what sits behind the staff figure

7 staff records on a typical affected organisation: narrow, not deep	22 organisations with an internal login page recorded in the same collections	34 organisations with a login page for an exposed service visible from outside
--	---	--

The staff credential picture here is shallow, and that is the right way to read it. One hundred and forty-seven records across 22 organisations is about seven each. A handful of records on an affected organisation is the pattern of occasional individual device infections rather than a systemic problem, and it means the response is proportionate: a reset, a session revocation and a multi-factor check on a small number of named accounts.

Credentials and the pages they open, counted separately

Twenty-two organisations have staff passwords circulating. **Twenty-two organisations also have an internal login or administration page captured in the same criminal collections.** The two counts are equal, but they are separate measures and this report does not claim they describe the same organisations. Establishing whether they do is a question for the restricted evidence, not for a published figure.

The reason it matters is that a credential on its own has limited value to a stranger who does not know where it fits, while a credential recorded beside the address of the page it opens removes the step that normally slows an attacker down. Sequence the response accordingly: reset, revoke live sessions, then enforce phishing-resistant multi-factor sign-in, starting with identity providers, email, payment systems and remote access.

A further 34 organisations have a login page for an exposed service visible from outside. That overlaps with the port cases in section 7 and should be worked alongside them: a reachable service with a discoverable login page is a more attractive target than either finding alone suggests.

What these records establish, and what they do not

They are public-source records tied to a domain, mostly harvested by viruses on individual computers. **We cannot see whether any password still works, whether the account still exists, or when the infection happened.** The source does not tell us how completely it has collected, so we make no claim that this is fresh or complete.

None of it is evidence that an organisation has been broken into. It is a reason for each to check its own accounts privately against its own records, and never by trying a recovered password.

9. Forged email: strong, and worth finishing

Anyone can put an organisation's name on an email. What stops it reaching the inbox is a published instruction telling receiving mail systems what to do with a message that fails the checks. It can say refuse, send to spam, or do nothing.

What the organisation publishes	Websites	Share of 148
Refuse the forgery	93	63%
Send it to spam	28	19%
Do nothing, monitoring only	22	15%
No usable policy observed	5	3%

This is the strongest control family in the estate. Ninety-three organisations, 63%, ask receiving systems to refuse a forgery outright, and 121 of 148, 82%, either refuse it or send it to spam. Twenty-seven do not instruct either: 22 publish a policy set to do nothing and 5 have no usable policy observed.

Why the remaining 27 matter more here than the number suggests

In most sectors a forged email is a phishing nuisance. From a bank or a building society it is a fraud against that organisation's customers, and the people who act on it fastest are often those least able to absorb a loss.

The control that stops it is a published instruction costing nothing to set. **It is the cheapest item in this report and the one with the most direct customer consequence.**

The work behind it is harder than the result suggests. Publishing the instruction takes minutes; doing it safely does not. An organisation has to know every system that sends mail in its name, the statement runner, the card processor, the mortgage platform, the marketing tool, the recruitment system, and watch for long enough to be confident, or genuine customer correspondence starts being refused. The right sequence is watch, then file to spam, then refuse, and the 28 already filing to spam have done most of the work.

The supporting records

Control	Have it	Do not
Record naming which servers may send mail	143	5
Signing record on the names we sampled	95	53
Encryption reporting record	11	137

A missing signing record means none was found under the names we sampled, not that none exists. Published records describe what an organisation asks other systems to do; they do not prove a forgery would in fact be refused, because receiving systems apply their own judgement.

10. Cookies, consent and tracking

This concerns what happens on a customer's own device, and it sits in the regime that applies to these organisations directly.

Read this before the number

Our privacy check scored **63.8% of the points it was able to assess**, 1,479 of 2,318. **That is not 64% compliant with anything.** It scores a handful of visible technical settings and nothing about how an organisation actually handles personal information.

What we looked for	Yes	No	Could not tell	Why it matters
Visitors forced onto a secure connection	112	5	31	Anything typed into an insecure page can be read on the way.
No insecure content on a secure page	108	7	33	A padlock is weakened if parts of the page arrive unprotected.
Privacy policy route found	76	34	38	People should be able to find out what happens to their data.
Cookie consent mechanism found	44	71	33	The law governs what may be stored before a person chooses.
Route to report a security problem	13	113	22	Someone who spots a flaw needs a way to tell the organisation privately.

Seventy-one of the 115 websites we could read showed no consent mechanism, 62%. Thirty-three homepages were not readable by the scanner at all, and their consent state is unknown rather than assumed clean. **This is the weakest area in an otherwise strong estate**, and it is the one a regulator is actively testing.

What the websites store

Cookie evidence came back for 114 websites. **61 stored something during the first page visit, 128 items in total.** Of those 61, 29 stored at least one item without the setting that limits it being sent from other websites, 21 without the setting that hides it from page scripts, and 20 without the setting that keeps it off insecure connections.

Only a response that was the site's real homepage is counted; a protection or error page answering in its place is not. Collection was skipped for 21 websites and unavailable for 13. This sees only what the first page request returns, so the real numbers are higher, and not every item needs every setting, so this is a review list rather than a defect list.

Thirteen of 148 publish a route for reporting a problem

Someone who finds a weakness at one of the other 113 has no obvious way to tell them privately, which is how problems end up disclosed in public instead. It is a single file on a website and it costs nothing, and for a regulated financial institution it is a conspicuous gap.

10 continued: tracking, and the regulator's position

Tracking technology on 78 of 115 readable websites

Tracking technology was present on **78 of the 115 websites we could read, 68%**. A tag manager accounts for almost all of it, on 77 of the 78. Beyond that: a session-replay tool on 8, analytics on 3, and an advertising pixel, an advertising tag and a marketing-automation tag on one each. These overlap on a single website.

The presence of tracking code proves nothing on its own. It does not show that personal data was sent, that permission was needed, that permission was missing, or what legal role the vendor plays. A tag manager is a container: what it loads and when has to be watched, not assumed.

The eight session-replay deployments are the ones worth separate attention, and this is the highest count of any estate we have measured. Session replay records what a visitor does on the page rather than counting visits, which raises sharper questions about what is captured and whether form fields are masked. On a banking website, where a visitor may be typing an account number, a balance or a payment instruction, that question deserves a deliberate answer rather than an assumption that the vendor masks by default.

The regulator is actively testing this

In December 2025 the Information Commissioner's Office reported the result of testing the **top 1,000 websites used in the UK**. It checked three things: whether advertising cookies were stored before consent, whether refusing was as easy as accepting, and whether cookies were placed without permission at all.

564 of the 1,000 only became compliant after the ICO contacted them	17 received preliminary enforcement notices	21 were still non-compliant when the ICO reported
---	---	---

The ICO said it will keep testing the top 1,000 periodically.^[3] Two things follow. An absent consent mechanism is not unusual, since more than half of the UK's most-visited websites needed prompting. And more importantly, **the regulator tests precisely the behaviours this scan cannot see**, so our signal is a prompt to look rather than an answer.

What a browser check should establish, in order

1. What is stored on the device before any choice is made, and is any of it not strictly necessary.
2. Is there a refuse option, is it on the first screen, and does it take no more effort than accepting.
3. Does refusing actually stop the storage, or just record a preference.
4. Does the same hold on account and application pages, not just the homepage.
5. What does each piece of tracking code load, and what leaves the browser when it does.

11. Card payments and transport security

Card data is core business for most of this population, so the card industry standard applies widely here, though its exact scope is decided by payment flows rather than by what kind of organisation this is.

Two things this report is not

Our payment-surface check scored **92.3% of the points it could assess**, 2,277 of 2,468. Controls we could not verify were removed from the calculation rather than counted as passes, which is exactly why that number is high. **It does not mean 92% compliant.** Three controls could not be answered for a single organisation in the estate.

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

External signal, of 148 reported	Pass	Issue	Could not tell
No Heartbleed weakness	148	0	0
Valid certificate	147	1	0
No software weakness reported	143	5	0
Modern encryption only	137	7	4
No database port answering	139	9	0
Visitors forced onto a secure connection	111	4	33
Browser told never to try the insecure version	105	10	33
Route to report a security problem published	25	104	19
Cloud storage, mail encryption, mail relay	0	0	148

The 9 database-port issues are this check's own result and include the 5 reported separately in section 7, which is why the organisation-level database count is 4. **The last row is the honest part of this section:** three controls could not be answered for a single organisation, and they were excluded from the calculation rather than recorded as passes, which is precisely why the headline percentage is high. An unknown is not a pass.

Transport is in good order, with a short list to finish

Four organisations do not force visitors onto a secure connection and 10 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. The two belong in one piece of work. Seven still accept outdated encryption and one has a certificate fault: different jobs, the first invisible and mattering only to an attacker already positioned to use it, the second visible to every customer who visits.

12. Domain controls and the dependency nobody records

Control	Have it	Do not	What it does
Mail server record	143	5	Names which servers may send mail for the organisation.
Mail signing record	95	53	Signs outgoing mail so it can be verified. Sampled, so absence is not conclusive.
Certificate authority record	26	122	Names which companies may issue certificates for the domain, so a rogue one is refused.
Security contact route	25	104	A published file telling a researcher where to report a problem privately.
Signed domain records	24	124	Stops someone forging the answer to 'where is this organisation's website'.
Encryption reporting record	11	137	Asks other systems to report failures in protecting mail in transit.

Four of these six controls sit below 20% adoption. A gap shared by more than eight organisations in ten is a sector fact rather than 120 individual failures, and it is a natural candidate for a coordinated push through an industry body rather than 148 separate projects. Signed domain records at 24 organisations and certificate authority records at 26 are the two that most directly protect the organisation's own identity on the internet, which for a bank is not a small thing.

The dependency that is on nobody's risk register

Observation	Websites	Rate
Mail transport policy finding	141 / 148	95%
Some form of network concentration	118 / 148	80%
All address records with one provider	94 / 148	64%
Risky port on a shared provider address	6 / 148	4%
Internet address with a poor public reputation	3 / 148	2%

These overlap and must not be added. A flagged address can belong to another organisation sharing the same infrastructure, so ownership needs confirming first.

94 organisations, 64%, depend entirely on one company to answer the question 'where is this organisation's website'. That is not a fault. Concentration is usually the better-resourced choice, and splitting the job across providers brings its own problems. **The question is whether anyone has written it down.** Under the operational resilience rules a single provider behind a customer-facing service is exactly the kind of third-party dependency a firm is expected to have identified, with a failure mode, a recovery route and an accountable owner attached.

Domain expiry and transfer-lock observations are deliberately excluded from this report. The public evidence does not show whether automatic renewal is switched on and is not reliable enough to support a finding either way.

13. What we could not read

Two 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
Floor score	33	The site turned our scanner away with a protection page. Everything else was read. The score is published as a floor and can only be higher.
Excluded, completed without a score	4	A second piece of core evidence was missing too, so no score could be published. Three banks and one society, excluded from every figure on the owner's instruction.

No scan failed anywhere in this estate, which is worth stating: every tracked address resolved and answered.

A protection page is a locked gate, not a hole in the fence

A site that refuses an unknown automated visitor is exercising a security control and doing something sensible, and it is notable that a fifth of this estate does it. The consequence for us is simply that we cannot read the instructions that site sends to a browser, and those are one of the things the score is built from.

This cannot be fixed from our side. No amount of rescanning will change it without the organisation or its provider choosing to let the scanner through. It is not evidence of a problem and should never be reported as one.

What else we could not see

Evidence type, of 148 reported	Got it	Partly	Could not
Website instructions to the browser	115	0	33
Homepage content, for consent and tracking	115	0	33
What the website stores on a device	114	0	34
Direct mail server test	18	0	130
Mail resilience	11	137	0
Other systems on the same infrastructure	0	119	29
Cloud storage discovery	0	148	0

These limits apply to scored websites too, so a score does not mean complete evidence and this report never treats it as if it did. Cloud storage could only be partly checked everywhere, which is why that control appears as unknown throughout rather than as a pass. Direct mail server testing ran on 18 websites and still produced no definitive answer on mail encryption or relay status.

The remedy is diagnosis, not repetition. Running the same blocked request again establishes nothing, and no score is invented in the meantime.

14. What to do, in order

Ordered by consequence, not ease. These groups overlap and are not a count of separate organisations.

1	Validate the single ransomware leak-site reference Match entity, date and incident record privately with the incident response team. A listing can be historical, duplicated, or refer to a similar entity elsewhere. Never publish it against a named organisation.
2	Check the six ports answering on organisations' own addresses Four are database ports, standing in front of customer records, and those come first. None of the six had a service identified, so the first question is what is actually listening.
3	Take the six shared-provider cases to the provider Three sit on shared edge addresses of one web-security provider, one on a second content-delivery network, one on a managed-hosting address serving at least 22 organisations and one on shared web hosting. Ask which services are meant to be reachable, how the platform is protected, and what the contract says about breach notification.
4	Check the five version-inferred software weaknesses, without alarm All 82 observations were read from version numbers, none confirmed, and none is on the public known-exploited list. Confirm the software actually running. This ranks below the items above precisely because nothing is on that list.
5	Reset the staff passwords 22 organisations, 147 records, about seven each. Reset, revoke live sessions, then enforce phishing-resistant multi-factor sign-in, starting with identity providers, email and payment systems. Never test a recovered password.
6	Finish the email work on the remaining 27 Inventory every legitimate sender first, then watch, then file to spam, then refuse. The 28 already filing to spam have done most of it and should complete. Investigate the 53 sampled signing-record gaps and 5 missing mail server records alongside.
7	Open the 71 websites in a browser and watch what happens The consent question cannot be settled from outside, and it is the weakest area in this estate. Ten minutes per site, and it answers the cookie review for the 29 sites with attribute gaps too. Look hardest at the eight running session replay.
8	Work the 48 organisations with identities on other platforms 874 records, the larger actionable population. The reset happens on the other company's system; the harder half is establishing which outside services staff used a work address on.
9	Fix transport, then the browser instructions Four organisations not forcing a secure connection, 10 without the instruction that stops the browser trying the insecure version, seven accepting outdated encryption and one certificate fault. Confirm the three poor-reputation addresses against tenancy.
10	Publish a security contact route, and close the evidence gaps Only 25 of 148 have one. The four excluded websites need a completed run, and the 33 floor scores a decision on letting a known scanner through.

Conclusion

This is the strongest estate we have measured, and its remaining work is short and specific. An exact-score average of 77.0 with 55 of 115 organisations in the top band, no failed scan anywhere, a single ransomware reference, and nothing at all on the public list of actively exploited weaknesses: that is a sector that has invested in the ordinary settings a scan like this counts, and the result should be read that way.

What that average cannot carry is the handful of items that would matter on a Monday morning, and they do not follow the score bands. An organisation can sit comfortably in the top band and have a database port answering the public internet.

Where the work actually is

One ransomware leak-site reference and four database ports answering on organisations' own addresses. These are the sharpest consequential items here. A database port stands in front of customer records, and a leak-site listing, though only a criminal group's claim, needs entity, date and incident record matched privately before anything else.

Twenty-seven organisations cannot instruct refusal of a forged email. The cheapest item in the report and the one with the most direct customer consequence, in a sector where a forged message is a fraud against the public.

Seventy-one of 115 readable websites showed no consent mechanism. The weakest area in an otherwise strong estate, in the one regime that applies directly and has a regulator actively testing it. Eight websites run session replay, which on a banking page deserves a deliberate answer about what is captured.

And one structural point. Of the 12 ports answering that should not be, six sit on addresses shared with unrelated organisations and are reported separately from the six on organisations' own addresses. Separating them does not make them smaller: a failure at that layer would reach every organisation on the platform at once, and it is a third-party dependency the resilience rules already expect to be written down.

Three 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 rank banks against building societies, which are close and differ in composition. And it makes no compliance judgement, because we observe settings, not lawfulness.

One thing it does claim. The 132,483 customer records in circulation are not a breach of any organisation here and must never be reported as one. They come from customers' own devices and from other companies' breaches. What they create is a duty to detect reuse at sign-in and to warn the people affected.

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, legal advice or regulatory advice.

How to read the figures

The two lists, before anything else

No website appears on both lists, which was checked rather than assumed. The combined population of **152 distinct websites** is therefore a true sum, and the two columns in every source table genuinely add.

Bands, and what a score is

Good is 75 and above, **needs improvement** is 50 to 74, **at risk** is below 50. One serious finding can hold a website below a band however much else is right, and a top-band organisation can carry a consequential finding from section 7. **A score is a floor, not a verdict.** It records what was visible from outside on one day, not internal controls, fraud monitoring, patching, training or incident readiness.

Exact scores and floor scores

115 websites have an exact score and 33 a floor score, published as "at least this much" where the browser instructions could not be read. **Floor scores count every unread item as missing**, so the true figure can only be the same or higher, and they are never compared with exact scores. Two averages are given: **77.0** on exact scores alone and **73.9** counting floors at their floor.

Counting, and whose address a port sits on

Every figure counts **websites, not problems**: five organisations, not 82 observations. Rows overlap and must never be added: the 4 database ports sit inside the 6 ports answering on organisations' own addresses, and the 6 shared-provider cases are counted once, on their own line.

Exposure is reported by owner, case by case. A case counts against the organisation where the address serves only it or its group, whoever operates the machine; separately where nothing was identified and unrelated organisations share the address. **Separate keeps its full weight.** Of 12 observed, **6 are organisation-level and 6 shared.**

The known-exploited distinction

A weakness **read from a version number** is what the software says about itself, unconfirmed. A weakness on the **known-exploited list** is one criminals are demonstrably using. Five organisations have the first; **none has the second**, and all 82 were checked.

The two percentage scores are not compliance rates

The privacy check scored 63.8%, 1,479 of 2,318, and the payment-surface check 92.3%, 2,277 of 2,468, **of the points each was able to assess.** Controls that could not be verified from outside were removed from the calculation rather than counted as passes, which is why the payment figure is high. Neither is an audit or a legal conclusion.

Publication rule

Do not describe an individual organisation as 'at risk', as non-compliant, or as ransomware-affected on this evidence alone. A leak-site listing is a criminal group's claim, not a verified fact. Customer credential records are not a breach of the organisation. A website that refuses an automated scanner is exercising a security control, and a provider's shared infrastructure is not attributed to one organisation.

Every figure describes the assessment of 26 September 2026. No trend is reported and none should be inferred. Customer account records are unscored. The population is tracked websites, not legal entities or brands; a zero means a check ran and found nothing.

Sources and notes

[1] Operational resilience: PRA SS1/21 and FCA PS21/3

Both require firms to identify important business services, set impact tolerances and understand the third-party dependencies that could cause disruption. Customer-facing websites, trustworthy communications, shared platform providers and the staff identity controls behind all three are dependencies of those services. [PRA SS1/21](#) and [FCA PS21/3](#)

[2] Information Commissioner's Office

Guide to PECR: cookies and similar technologies, governing what may be stored on a person's device and the permission needed, and *A guide to data security* under the UK GDPR. These decide whether an arrangement is lawful; the signals in section 10 of this report are technical observations and are not such an assessment. [ICO cookie guidance](#) and [ICO data security guidance](#)

[3] Information Commissioner's Office, *ICO action secures increased cookie compliance, December 2025*

The ICO tested the top 1,000 websites used in the UK against three requirements: no non-essential advertising cookies before consent, refusing as easy as accepting, and no cookies placed without permission. 415 passed without intervention, 564 improved after ICO engagement, 17 received preliminary enforcement notices and 21 remained non-compliant at the point of reporting. [ico.org.uk](#)

[4] PCI Security Standards Council

The PCI DSS standard and the Council's applicability guidance determine actual scope; payment flows, redirection arrangements and supplier relationships cannot be scoped from an external scan. Requirement 11.3.2 requires passing external scans by an Approved Scanning Vendor at least every three months. An Approved Scanning Vendor is one specifically qualified by the Council, and this assessment is not one. [PCI DSS standard](#) and [applicability guidance](#)

[5] National Cyber Security Centre

Email security and anti-spoofing, including the staged move from monitoring, through spam filing, to refusal; *External attack surface management*; *Multi-factor authentication for corporate online services*; *Vulnerability disclosure toolkit*. Free plain-English guidance requiring no specialist team. [NCSC email security](#), [MFA guidance](#) and [vulnerability disclosure toolkit](#)

[6] MyDomainRisk / Huro Data Technologies Ltd, *two source estate baselines and one consolidated baseline, extraction 26 September 2026*

Banks 110 tracked, 107 reported (83 exact, 24 floor), sum 7,851, exact average 76.4, 3 excluded without a score. Building societies 42 tracked, 41 reported (32 exact, 9 floor), sum 3,082, exact average 78.4, 1 excluded without a score. No scan failed on either list and no website appears on both, so membership reconciles as 110 + 42 = 152 distinct websites. Consolidated: 115 exact and 33 floor scores, sum 10,933, exact average 77.0, published average 73.9. Averages are calculated from the 148 underlying scores rather than by averaging the two list averages.

Provenance and handling

The source is one read-only, repeatable-read production snapshot of the latest owner-matched security attempt for every currently tracked website across both lists, taken on 26 September 2026. Every reported website's latest attempt is dated 26 September 2026 and every completed scan used the same method. Websites without a published main score are excluded on the owner's instruction and counted in section 13. The analysis refuses to report where an evidence path is absent rather than reading it as zero.

Each case of a port answering was classified one at a time from the same scans: the service evidence recorded for each flagged port, and the other websites in the scan data and the provider's own hostname records that share each address.

No organisation is named in this report. The findings are aggregate by design: a named institution beside its weakness is a document that cannot safely be circulated, and the value here is the sector pattern. Per-organisation detail exists and is available privately to the organisation concerned, or to anyone it authorises.

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

Every named finding should be rechecked immediately before responsible disclosure, because public-facing settings change constantly. Aggregate figures may be used with the cautions in this report. This document contains counts only: no passwords, no personal records, no raw provider data, no exploit instructions and no named organisation detail. A finding is not evidence that any organisation has been compromised or is in breach of any obligation.