How UK Charities Look to an External Cyber Security Scan

What a non-intrusive external assessment of two UK charity cohorts found, and practical steps boards and trustees can take to improve resilience

> **Purpose**
>
> To explain cyber security and cyber posture in plain English, place the charity sector in the wider UK context, report what our external assessment found in each of its two source cohorts, and set out realistic actions that charities can take.

Sources	Two separately assembled cohorts, reported separately throughout: **220 best-known UK charities**, selected for public prominence, and **236 largest UK charitable bodies by income**.
Combined	76 organisations appear in both cohorts, so the two lists represent **380 distinct organisations**. The combined view is reported in section 7, after the sources. 456 is a count of list entries and is never used as a population.
Evidence	Single external assessment completed 21 August 2026. 214 of 220 best-known and 232 of 236 largest by income produced a completed result. Nine organisations produced none and are counted as unassessed rather than excluded.
Method	Non-intrusive external assessment using publicly observable information: websites, DNS, certificates, public threat intelligence and publicly reachable services. No logins, password attempts or intrusive probing.
Prepared by	Huro Data Technologies Ltd / MyDomainRisk
Version	Version 4, 1 September 2026. Supersedes version 3 of 29 August 2026. Changes are listed in full on the final page.

Research and external assessment: MyDomainRisk | Huro Data Technologies Ltd

Why this matters now

In the summer of 2026 a supplier used by more than a thousand UK charities was broken into, and the personal data it held on their behalf was copied. Charities that had done nothing wrong themselves found they had breaches to report, supporters to notify and questions to answer about systems they did not run and could not see inside.

That incident has done something useful. It has made cyber security a board conversation in organisations where it had previously been an IT matter, and it has prompted a reasonable question from trustees and chief executives: **what does our own charity look like, and how would we know?**

This report is an attempt to answer that at sector level. It is deliberately not about any single incident. It is about the ordinary, everyday condition of two large groups of charitable organisations as anyone on the internet could see them, and about the small number of things that make the most difference.

A way of picturing it

> **Imagine somebody walking down your street**
>
> They never go inside any building. They never try a door. They simply walk past and notice what is visible from the pavement: that a window has been left open, that a side gate is unlatched, that parcels have been sitting on a porch for three days, that a spare key is under the mat where everyone knows to look.
>
> None of that means the house has been burgled. It means somebody passing by can see opportunities, and most of them take a minute to put right.

That is exactly what this assessment does. We looked at each organisation from the public internet, the way any visitor, or anyone with less good intentions, would see it. We did not log in anywhere, try any password, or look at anything that was not already on public view.

So when this report says an organisation has staff addresses circulating in criminal collections, 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 login page has been recorded alongside those addresses, that is knowing which door the key opens. And when it says a service is reachable from the internet that need not be, that is the side gate nobody remembers leaving unlatched.

> **What we could not see**
>
> Standing in the street tells you nothing about the alarm system, the safe, or the locks on the internal doors. Many of the organisations here will have good protections we simply cannot observe. That is why nothing in this report is a verdict on any charity's security. It is an account of what is showing from outside.

Executive summary

The short version

Nobody's building has been broken into. But a lot of front doors are unlocked, and rather more spare keys are lying about than anyone would be comfortable with.

We found no sign that any of these organisations has been compromised. What we did find is a pattern of ordinary, preventable exposure: staff addresses already circulating in stolen collections, the internal pages they unlock recorded alongside them without exception, gaps in the controls that stop criminals sending email in a charity's name, and a handful of services reachable from the internet that probably need not be.

Almost all of it can be fixed cheaply, and much of it in an afternoon.

Two sources, assessed the same day

This study rests on two separately assembled lists, reported separately because they were built on different principles.

Headline	Best-known (220)	Largest by income (236)
Average score / median	65 / 69	65 / 67
Completed assessments	214 (97%)	232 (98%)
Addresses in stolen credential collections	102 (46%)	148 (63%)
Does not instruct rejection of forged mail	79 (36%)	74 (31%)
Risky reachable service	20 (9%)	15 (6%)
Ransomware, breach-collection or browser-safety findings	0	0

46% / 63%
have staff or organisational addresses in stolen collections

65 / 65
99 / 99
of those with own-domain exposure also have their login pages recorded. No exceptions

0
ransomware mentions, breach records or browser-safety flags anywhere

Put plainly: for between four and six organisations in ten, someone else already holds a working key. **And wherever those keys are the organisation's own staff addresses, the address of the door is recorded beside them, without exception.** That pairing is universal in both cohorts, 65 of 65 and 99 of 99, which makes it the most actionable finding in this report.

The good news deserves saying as loudly as the rest. Across both estates there was not a single ransomware leak-site mention, not one record in a known breach collection, and no browser-safety flag anywhere. This is not a sector in trouble. It is a sector with a short list of things worth doing, and time in which to do them.

1. Cyber security in plain English

Cyber security is the practice of protecting an organisation's people, accounts, data, systems and reputation from unauthorised access, misuse, disruption or loss. For most charities it is not primarily about sophisticated attackers or expensive technology. It is about making common attacks harder, limiting the damage if something does go wrong, and being able to recover quickly.

The National Cyber Security Centre frames this as risk management supported by practical controls. Its guidance for small organisations focuses on securing email and important accounts, protecting devices, backing up data and spotting attacks; its 10 Steps guidance adds areas such as identity and access management, vulnerability management, incident management and supply-chain security.[3]

What 'cyber posture' means

Cyber posture is the condition of those defences at a point in time. A useful way to think about it is: if someone who knew nothing about your charity looked at it from the public internet, what could they learn about your digital estate, your security controls and possible routes in?

> **Security is cumulative**
>
> A strong posture is rarely one product. It is the combined effect of many ordinary controls: strong authentication, sensible configuration, secure email, patched systems, limited internet exposure, good supplier information, clear reporting routes and a rehearsed response plan.

Common weaknesses identified in our assessment, and why they matter

Weakness identified	Why it matters
Stolen or reused credentials	Addresses and session data can be harvested from infected devices. Multi-factor authentication, resets and session revocation reduce their value.
Login pages recorded alongside them	A credential is far more useful to an attacker when the same collection also records the page where it is used. In this study that pairing is universal.
Weak email authentication	If a domain does not tell receiving mail systems to reject forged messages, criminals have more scope to impersonate a trusted charity to its supporters.
Unnecessary internet-facing services	Database, file-transfer, remote-access and administration services should not be publicly reachable unless there is a clear need and strong protection.
Out-of-date or incomplete configuration	Missing security headers, weak TLS settings and poorly maintained certificates can enlarge the attack surface or confuse supporters.
Limited visibility of suppliers	A charity may know a supplier is accredited but still not know exactly what data that supplier holds, or how quickly it can provide a per-charity incident breakdown.
No route to report a problem	If a researcher or member of the public spots a weakness, they need an obvious, safe way to tell the organisation.

The NCSC recommends Cyber Essentials as the government-backed minimum technical standard for organisations of all sizes. It centres on five controls: firewalls, secure configuration, security updates, user access control and malware protection.[4]

2. The wider UK picture

It is sometimes said that charities are among the most attacked organisations in the country. The government's own figures do not bear that out, and the truth is more interesting. In the Cyber Security Breaches Survey 2025/26, 28% of charities identified a breach or attack in the previous 12 months, against 43% of businesses, still roughly 57,000 charities. And it rises sharply with size: 57% of charities with income above £500,000 identified one.[1]

Charities are not targeted more. They are less well equipped, and on the government's numbers the gap is widening. Among charities that were attacked, the proportion facing attacks weekly or more often rose from 18% to 26% in a single year, staff training fell, and the government estimated around 525,000 cyber crimes against UK charities, overwhelmingly phishing.[2]

2025/26 measure	Businesses	Charities	High-income charities
Identified a breach or attack in last 12 months	43%	28%	57%
Victim of at least one cyber crime	19%	14%	34%
Cyber security a high priority for senior management	72%	60%	87%
Any listed cyber-risk identification activity	52%	38%	80%
Staff cyber training in last 12 months	19%	17%	49%
External cyber security provider	48%	25%	.
Review immediate supplier cyber risk	15%	9%	31%

Source: DSIT and Home Office, Cyber Security Breaches Survey 2025/26. 'High-income charity' means income of £500,000 or more. A full stop indicates the comparison is not used here rather than a value of zero.[1]

> **A supportive reading of the official data**
>
> The figures point to a capacity problem more than a lack of concern. High-income charities are substantially more likely to prioritise cyber security, identify risk and provide training. Smaller charities have fewer resources, and government research records lower uptake of external expertise, formal planning and supplier review across the sector.

What we did differently

Surveys ask organisations what has happened to them. We looked instead at what can be seen before anything happens, the view from the street, across two separately built lists of significant UK charitable organisations. Both averaged the same external score of 65. The number itself is unremarkable; what sits beneath it is not.[5]

3. What we measured

Our work is an external posture assessment, using information already observable from the public internet and public records. Nothing was logged into, no password was tried, and no system was probed beyond what an ordinary web visitor and external non-intrusive scans expose.[5]

220	236	380
best-known UK charities, selected for public prominence	largest UK charitable bodies by income	distinct organisations, after 76 appear in both lists

The two cohorts are reported separately throughout, and the combined view follows in section 7. They must never be added together: 220 plus 236 is 456, a count of list entries that double-counts 76 organisations. The income list also contains higher education bodies, which the best-known list does not.

Coverage, and the nine organisations that produced no result

Assessment coverage	Best-known	Largest by income
Organisations tracked	220	236
Completed and scored	214 (97%)	232 (98%)
No completed assessment	6 (3%)	4 (2%)
Security headers could not be read	62 (28%)	70 (30%)

Ten list entries, representing nine distinct organisations, produced no scan result of any kind. One of the nine appears in both cohorts. **Every percentage is expressed against the tracked populations of 220 and 236, not the smaller assessed figures**, so no result is flattered by a reduced denominator.

> **What the assessment does not prove**
>
> Nothing here says any organisation has been breached. Credential findings describe sign-ins circulating in criminal collections, often harvested from infected devices. A reachable port is not proof that a database can be accessed. A hosting address on a threat feed may reflect another tenant on shared infrastructure. External posture is evidence of exposure, not of intrusion.
>
> To return to the street: we can see that a window looks open. We cannot see whether anyone has climbed through it, and in many cases nobody has.

About credentials

We report credentials in three categories. Only the first two count towards any figure here.

Own-domain addresses are staff addresses on the charity's own domain. **Own addresses on outside services** are the same organisation's addresses signed in to somebody else's platform, typically identity and collaboration tools. Both belong to the charity and both can be acted on.

Supporter and service-user accounts are excluded from every figure in this report. Those describe malware on members of the public's own devices rather than charity security, and nobody at the charity can reset them. They appear alone for 69 best-known organisations (31%) and 41 by income (17%).

4. Source one: the 220 best-known charities

These organisations were selected for public prominence rather than income. **Average scored result: 65; median 69.** The weakest website scores 21, the strongest 89.

Posture band	Organisations	Share of 220
At risk	25	11%
Needs improvement	162	74%
Good	27	12%
No completed assessment	6	3%

The median of 69 sits four points above the mean, which tells you the shape of the problem: most of this cohort clusters just below the *Good* threshold, and a small tail of badly configured estates drags the average down. **Three-quarters of the cohort is within a handful of points of a better band**, and the work needed to move them is configuration, not investment.

Priority observations

Observation	Organisations	Share of 220
Does not instruct rejection of forged mail	79	36%
Addresses in stolen credential collections	102	46%
of which, on the organisation's own domain	65	30%
of which, own addresses on outside services	94	43%
TLS certificate, protocol or cipher issue	24	11%
Risky reachable service (file transfer, remote desktop, admin)	20	9%
Database port reachable from the internet	7	3%
Ransomware, breach-collection or browser-safety findings	0	0%

The credential rows are different views of the same evidence and must not be summed. 65 and 94 overlap; 102 is the count of organisations appearing in either.

The largest single gap is not an exposure at all. **Seventy-nine organisations, 36% of the cohort, do not instruct receiving mail systems to refuse forged mail in their name.** Sixty-eight publish a policy of *none* and the remainder publish nothing; a further 66 (30%) publish *quarantine*, which sends forged mail to spam rather than refusing it. For organisations whose income depends on supporters trusting a message that appears to come from them, that is a material risk.

Other findings worth naming: 113 (51%) publish no HSTS header, 41 (19%) do not redirect plain HTTP to HTTPS, 40 (18%) serve mixed content on a secure page, and 113 (51%) have no domain transfer lock set, which is a single registrar setting protecting the domain that *is* the charity's identity.

5. Source two: the 236 largest by income

This cohort is defined by income rather than public profile, and it therefore contains **36 higher education bodies (15%)** that are charities or exempt charities. They are large employers with federated web estates, and they materially affect the credential figures. **Average scored result: 65; median 67.** The weakest website scores 15, the strongest 89.

Posture band	Organisations	Share of 236
At risk	22	9%
Needs improvement	186	79%
Good	24	10%
No completed assessment	4	2%

This is the most tightly clustered population we have assessed: 79% sit in a single middle band. Fewer organisations are in difficulty than in the best-known cohort, and fewer are doing well. The practical reading is that this cohort has adopted the controls that hosting and platform providers enable by default, and has largely not adopted the ones that require a deliberate decision.

Priority observations

Observation	Organisations	Share of 236
Addresses in stolen credential collections	148	63%
of which, on the organisation's own domain	99	42%
of which, own addresses on outside services	140	59%
Does not instruct rejection of forged mail	74	31%
TLS certificate, protocol or cipher issue	24	10%
IP address listed on independent threat sources	24	10%
Legacy TLS protocol still accepted	21	9%
Risky reachable service (file transfer, remote desktop, admin)	15	6%
Database port reachable from the internet	8	3%
Ransomware, breach-collection or browser-safety findings	0	0%

At 63%, credential exposure is the defining finding of this cohort and the highest rate we have measured anywhere. Two structural factors explain much of it, and neither is a criticism: these are large employers, so more staff addresses exist to be captured, and the education bodies within the cohort have large, distributed user populations with long-lived accounts.

Two further findings deserve attention because they are cheap to fix and easy to miss. **Thirty organisations (13%) have a domain expiring within a short window**, with a further 22 (9%) approaching; a lapsed domain is an outage and an impersonation opportunity in one. And 24 organisations (10%) have a primary IP address listed on independent threat sources, which frequently reflects another tenant on shared hosting rather than the organisation itself and should be checked before it is treated as a finding.

6. What the two sources show when compared

Both cohorts average 65. That single fact is the most misread number in this study, so it is worth being precise about what differs beneath it.

Measure	Best-known (220)	Largest by income (236)
Average score	65	65
Median score	69	67
At risk	11%	9%
Good	12%	10%
Credentials in stolen collections	46%	63%
own domain	30%	42%
own addresses on outside services	43%	59%
Forged mail not refused	36%	31%
Certificate problems	11%	10%
Risky reachable service	9%	6%
Supporter accounts only	31%	17%

> **Higher income does not buy a better external result**
>
> The income cohort has materially more credential exposure. The best-known cohort has more exposed services and a worse email-authentication position. **This is a mixed pattern and must be described as a difference in exposure, not a ranking of competence.**

Two confounders travel with any comparison. The first is education. The income cohort contains 36 higher education bodies and skews towards very large employers; the best-known cohort contains none at all. **Excluding education entirely, own-domain credential exposure runs at 40% (79 of 196) against the best-known cohort's 30%.** The gap narrows but survives, so education explains part of the difference and not all of it.

The second is size. The largest single staff-credential count in either cohort belongs to a university, at 1,736, and the largest outside-service count is also a university's. Raw credential volumes are therefore excluded from every headline in this report: they measure organisational size at least as much as security. The largest total of any kind in either estate is an organisation with 123,643 records, of which 123,284 are public account holders. That figure describes reach, not security, which is precisely why a raw total must never be used as a headline.

7. The combined view: 380 distinct organisations

Having reported each source separately, the two can be combined, but only on the correct key. **76 organisations appear in both lists, so the combined population is 380 distinct organisations, not 456.**

> **How the overlap is established, and why it nearly went wrong**
>
> Deduplicating on the tracked address gives the wrong answer, 454 rather than 380, because the two lists track the same charities at different hosts: 75 of the 76 overlapping organisations appear as `example.org` in one list and `www.example.org` in the other. Only two are exact address matches.
>
> The correct key is the registrable domain, and the equivalence is proven rather than assumed. **Any future combined analysis must normalise the host before deduplicating.**

Combined measure	Organisations	Share of 380
Completed and scored	371	98%
At risk	42	11%
Needs improvement	286	75%
Good	43	11%
Addresses in stolen credential collections	190	50%
of which, on the organisation's own domain	119	31%
of which, own addresses on outside services	178	47%
Does not instruct rejection of forged mail	138	36%
TLS certificate, protocol or cipher issue	41	11%
Risky reachable service	32	8%
Database port reachable	14	4%

Average scored result across the combined population: **64.5; median 67.** The combined credential figure of 50% sits between the two cohort figures of 46% and 63%, as arithmetically it must. A combined number falling outside the range of its parts would be an error, and that is the check worth repeating on any future combined report.

Three-quarters of these 380 organisations sit in a single middle band. Only about one in nine is doing well by this external measure, and only about one in nine is in difficulty. The distribution is not a spread of maturity; it is a plateau. The risk is not that charities are failing. It is that a uniform, modest posture across an entire sector is a uniform opportunity for anyone who studies it.

8. Credentials and the doors they open

This is the largest finding in the study and the one with the clearest remedy.

> **A stolen password is half an intrusion**
>
> A credential becomes far more useful to an attacker when the same criminal collection also records the page or service where it is used. A key on its own is of limited use to a stranger. A key with the address attached to it is a different matter.

In this study, the pairing is not partial. It is universal. Every organisation with credentials exposed on its own domain also has its internal login pages recorded in the same collections: 65 of 65 in the best-known cohort, 99 of 99 by income, with no exceptions in either direction. The provider returns captured pages precisely when it holds staff credentials. Observed pages include job portals and remote-desktop gateways, which are doors onto an internal network rather than onto a website.

The practical consequence for a board is direct. Where the keys are your own staff addresses, you should assume the attacker already knows where to use them, and plan the response accordingly: reset the credential, enforce multi-factor authentication and revoke live sessions, starting with identity providers, email, finance, CRM and remote access.

> **How current this exposure is cannot be established for most of it**
>
> A date is supplied only for the own-domain group. In the best-known cohort just 14 of the 65 carry one, and **9 organisations (4%) have a date inside twelve months**. By income, 23 of the 99 carry one, and **14 organisations (6%) are inside twelve months**. The outside-services group has no date field whatsoever.
>
> Read these figures as *"these addresses appear in stolen collections"*, not *"these passwords were stolen this year"*. A record with no date is treated as current rather than assumed old, which is the safer assumption for a credential, but the date is not evidence held. The exposure is worth acting on either way: an undated credential is not a safe one.

The second group is not supplier exposure

The category of own addresses used on outside services is frequently mislabelled as supplier exposure. **It is not.** These are the charity's own email addresses signed in to somebody else's platform. Where the provider names those services across these estates they read as identity and collaboration tools: Microsoft, Amazon, Salesforce, Zoom, Dropbox, Mimecast. The accounts sit on other companies' systems; the identities belong to the charity, and the charity can act on both groups.

That distinction strengthens the finding rather than weakening it. Every organisation in the figure can act on every part of it.

9. Impersonation and supporter trust

Charities depend unusually heavily on trust. A convincing fraudulent email can exploit the organisation's name even where the charity itself has not been breached. Email authentication gives domain owners a way to tell receiving mail systems how to handle messages that fail authentication.

79 / 220	74 / 236	36% / 31%
best-known charities do not instruct rejection of forged mail	largest by income do not instruct rejection of forged mail	of each cohort respectively

The composition differs between the cohorts. Among the best-known, 68 publish a policy of *none* and a further 66 (30%) publish *quarantine*. By income, 58 publish *none* and a further 73 (31%) publish *quarantine*. Quarantine sends forged mail to spam rather than refusing it, which is better than nothing and not the same as rejection.

Moving to an enforcing policy is a relatively low-cost control, but it should be done carefully: organisations need to identify legitimate senders, monitor first, and only then move to rejection so that genuine fundraising, CRM and supplier mail is not accidentally blocked. Related gaps compound it: 94 best-known organisations (43%) and 99 by income (42%) publish an SPF record ending in soft fail, which asks receiving systems to accept suspicious mail and mark it rather than refuse it, and 28 organisations by income (12%) publish no DKIM record at all, so a receiving system has no signature to verify.

Near-universal findings are a capacity signal

Some external hardening controls are absent across almost the whole study. In the best-known cohort, 95% have no vulnerability disclosure policy published, 93% have no signed DNS responses, 93% have no restriction on which certificate authorities may issue for their domain, and 88% have no inbound mail transport security. The income cohort is within a few points on every one of those.

We do not interpret a control missing from nineteen organisations in twenty as evidence that nineteen in twenty are careless. It is more reasonably read as a sector-wide gap in guidance, time and specialist capacity. That interpretation is consistent with the government survey, which found lower charity uptake of external cyber providers, vulnerability auditing, supplier review and formal continuity planning than among businesses.[1]

The genuinely diagnostic findings are the ones that vary between organisations, and those are worth acting on individually: cookie consent, reverse DNS, HSTS at around half of each cohort, Content Security Policy at 45% and 43%, and SPF soft fail at 43% and 42%.

10. Supplier concentration: a risk no charity can solve alone

Charities often rely on specialist shared platforms for fundraising, case management, CRM, finance and communications. That is rational: shared services bring expertise and efficiency to organisations that cannot build equivalent systems themselves. The trade-off is concentration. One supplier can hold data for hundreds or thousands of organisations, so a single failure can have a sector-wide consequence.

The government survey reinforces the challenge. Only 9% of charities said they formally reviewed the cyber risks posed by their immediate suppliers, compared with 15% of businesses overall. Among high-income charities the figure was higher at 31%, but still below large businesses at 48%.[1]

Here the street analogy reaches its limit, and that is the point. Much of a charity's most sensitive information is not in its own house at all. It is in a storage unit shared with a thousand other organisations, and no amount of checking your own doors tells you how well that building is locked.

Our assessment cannot see inside a supplier. The closest it comes is the own-addresses-on-outside-services figure, 43% and 59%, and even that sees only the charity's own identities, not the supplier's controls. A related structural note: 77 best-known organisations (35%) have all their nameservers with a single provider. That is normal and often the right choice, but it is worth recording in a sector where a shared supplier failure has already proved to be the dominant risk.

> **What supplier assurance can realistically do**
>
> Certification and questionnaires matter, but they cannot guarantee that every implementation mistake will be caught. The practical objective is twofold: reduce supplier risk where possible, and make sure the charity can establish quickly what data was held, what happened, which records were affected and what must be done next.

A simple board-level supplier question

For each material supplier, trustees and executives should be able to answer:

- What data do they hold for us?
- How sensitive is it, and how many people does it cover?
- Do they hold attachments or special-category data?
- Who will tell us if something happens?
- How quickly must they provide a per-charity breakdown?
- What access can we revoke ourselves?

Supplier assurance as it is currently practised cannot prevent a failure inside someone else's systems. What it can do is shorten the response.

11. Seven practical steps to improve posture

The findings are actionable. The following sequence is deliberately weighted towards controls that provide a high return for organisations without a dedicated security team. It complements the NCSC's Small Organisations Guide, Cyber Action Toolkit, 10 Steps guidance and Cyber Essentials.[3][4]

1	**Check for exposed staff addresses, and act on what you find** Reset affected passwords, enforce multi-factor authentication and revoke active sessions. Prioritise identity providers, email, finance, CRM and remote-access accounts. Because the login pages are recorded alongside the credentials in every case, assume the attacker knows where to use them.
2	**Enforce multi-factor authentication everywhere it matters** Start with Microsoft 365 or Google Workspace, finance systems, CRM, remote access, administrator accounts and any platform holding beneficiary or supporter information. Multi-factor authentication dramatically reduces the value of a stolen password.
3	**Protect your name from forged email** Move SPF, DKIM and DMARC towards an enforcing policy after monitoring legitimate mail flows. The aim is not technical compliance: it is to make it harder for criminals to send convincing messages in the charity's name.
4	**Know exactly what each important supplier holds** Maintain a simple register of the systems, data categories, record volumes, attachments and sensitive data held by material suppliers. Do not wait for an incident to reconstruct this information.
5	**Make incident information a contractual requirement** Check that data-processing and service agreements require prompt, charity-specific information after an incident, including what data was affected and when. Agree notification routes before they are needed.
6	**Reduce the public attack surface** Close database ports, remote administration interfaces, file-transfer services and other internet-facing services unless they have a clear business need. Keep software patched, review certificate and transport configuration, and set the domain transfer lock where it is unset.
7	**Give people a safe way to report a problem** Publish a security contact, ideally a security.txt file and a visible email route, so that researchers, suppliers, staff and members of the public can report something suspicious before it becomes an incident.

For organisations wanting a free external starting point, MyDomainRisk provides public-facing checks designed to make cyber posture easier to understand. It should complement, not replace, internal risk management and NCSC guidance.

Conclusion

The useful conclusion from this work is not that charities are failing. It is that a large part of cyber resilience is made up of small, manageable decisions that can be improved over time. The sector starts from a position of high public trust, committed staff and trustees, and growing awareness of cyber risk.

The government data shows that high-income charities already treat cyber security as a high priority in large numbers, and that most charities have adopted at least some recognised protective measures. Our assessment also found no visible ransomware or breach-collection evidence across either estate. Those are important positives.[1][5]

> **The practical message for boards**
>
> Do not try to solve cyber security in one project. Improve posture in layers. Protect accounts. Know your suppliers. Reduce unnecessary exposure. Make impersonation harder. Train people. Plan the response. Review the position regularly. Each step makes the next incident less likely to succeed and less damaging if it does.

For smaller charities in particular, the answer cannot be to demand enterprise-scale security teams. The better model is a realistic baseline, free or low-cost guidance, shared sector learning, proportionate supplier expectations and tools that make risk visible in plain English. The NCSC already provides substantial free guidance and training; the opportunity is to make it easier for charities to find, understand and act on it.[3]

Cyber resilience is therefore achievable. The findings in this report are not a verdict on the sector; they are a starting point for practical improvement.

About this report

Posture figures come from a single external assessment completed on 21 August 2026 across two cohorts: 220 best-known UK charities, of which 214 produced a completed result, and 236 largest by income, of which 232 did. **Every figure describes that one run.** Nothing is carried over from an earlier assessment, and no comparison with any earlier run is made or implied. No charity registration or exempt status has been independently verified. No organisation named in the underlying assessment has been contacted for right of reply. This report is not legal advice.[5]

A note on reading the scores

Any charity that has seen an external score for itself, ours or anyone else's, should understand one thing about how such scores are produced.

On the 21 August run, the security headers could not be read for 62 organisations in the best-known cohort (28%) and 70 in the income cohort (30%). In many cases this appears to be bot defence refusing an automated request, which is itself a security control. No finding is raised for the area that could not be read, but the organisation earns no points for it either, so **the published scores are floors and both cohort averages understate the sector.**[5]

In street terms: for roughly three houses in ten, a tall hedge stood between the pavement and the house. We could not see whether a window was open, so we recorded nothing, which means those scores read lower than the reality. It is not evidence of a problem. If anything, it is a sign that somebody planted a hedge.

> **Publication rule**
>
> Do not describe an individual organisation as 'at risk' on the strength of a score alone, without first checking whether its headers were readable in the assessment. Bot defence should not be mistaken for poor security, and a meaningful share of the lowest band in each cohort is there partly because it defended itself against the scanner.

The same caution applies to any charity reviewing its own result, or a supplier's. **A score of this kind is a floor, not a verdict.** It records what could be seen from outside on one day, not everything that is there, and not a judgement on the people running it.

How to read the figures

- Every figure describes the single assessment run of 21 August 2026. No trend against any earlier run is claimed or implied.
- **Percentages are against the tracked populations, 220 and 236, not the assessed figures of 214 and 232.** Nine organisations produced no result and are counted as unassessed rather than excluded, so no figure is flattered by a smaller denominator.
- Counts are per organisation, never per finding. One organisation with five open database ports counts once.
- **The credential rows are overlapping views of one body of evidence and must never be added together.**
- **The two cohorts must never be added together.** 456 is a count of list entries; the distinct population is 380.
- Where a charity appears in both cohorts, its assessment is used once in the combined view. It is not counted twice and does not appear under two names.
- An absent control is not a proven weakness. It is an observation that a publicly checkable control was not observed.
- Raw credential volumes are not used as a headline measure. They reflect organisational size at least as much as security.

Sources and notes

[1] Department for Science, Innovation and Technology / Home Office. *Cyber Security Breaches Survey 2025/26*

National prevalence of breaches, attacks and cyber crime; charity and business comparisons; high-income charity figures; governance, training, supplier review, continuity planning, personal-data protection and incident-response measures.

[2] Civil Society. *Concern as figures show fewer charities prioritise cybersecurity*, 29 May 2026

Civil Society reporting on the 2025/26 government survey, including weekly-or-more attacks, training, personal-data protection and the estimated scale of charity cyber crime.

[3] National Cyber Security Centre. *Cyber security advice for small and medium-sized organisations*; *Small Organisations Guide*; *10 Steps to Cyber Security*

Plain-English baseline guidance for small organisations, charities, clubs and schools, and the wider NCSC risk-management framework.

[4] National Cyber Security Centre. *Cyber Essentials*

Government-recommended minimum cyber security standard and its five core technical controls.

[5] MyDomainRisk / Huro Data Technologies Ltd. *UK charity sector external security posture: best-known cohort, largest-by-income cohort and combined analysis*, assessment 21 August 2026, analysis 1 September 2026

Underlying external assessment of the 220-organisation and 236-organisation cohorts and the 380-organisation combined view; methodology, coverage, scoring caveats and all measured findings in this report. Credential evidence is supplied by a third-party infostealer intelligence provider and reported at organisation level. No individual, address or password is reproduced in this report or held for reporting purposes.

[6] Charity Commission for England and Wales. *Annual Report and Accounts 2025 to 2026*

Context on public trust in charities; the Commission reports that charities remain highly trusted relative to other institutions.

Provenance and disclosure

This analysis was produced with MyDomainRisk, a product of Huro Data Technologies Ltd, which has a commercial interest in the analysis, declared here so that readers can weigh it. It is not independent certification, an audit, legal advice or regulatory advice.

Method: externally observable information only, comprising public web responses, DNS, certificate transparency, published records and publicly reachable services. No authentication was attempted, no password was tried, and no system was probed beyond what an ordinary visitor and a non-intrusive external scan can see. No fundraising platform, case-management system, donor database or supplier tenancy was touched.

This document names no individual organisation and is safe to circulate at board level. A per-organisation breakdown exists and should be handled as restricted, because it identifies specific weaknesses in named estates. Aggregate figures may be used with the cautions set out in this report; named findings should be verified immediately before use and disclosed to the affected organisation first.

What changed in version 4

This version supersedes version 3 of 29 August 2026. Every figure was re-derived from the stored assessment records on 1 September 2026. **No conclusion changed. Two of the five corrections make the findings stronger.** The changes are recorded here in full because version 3 was published and quoted.

1. Coverage was not complete

Version 3 stated that 220 of 220 best-known charities and 236 of 236 largest by income were assessed, with no evidence gaps. **Ten list entries, representing nine distinct organisations, produced no scan result of any kind.** The corrected figures are 214 of 220 and 232 of 236. All list entries were added before the scan, so these are genuine gaps rather than late additions. The band tables and three finding counts move as a result, because version 3 assumed every organisation was scored.

2. "Supplier-related" was the wrong description

The second credential category holds the organisation's **own** addresses on outside platforms, not suppliers' details. Version 3's hedge on this point was overtaken on 22 August, the day after the scan. The description is now "own addresses on outside services" throughout. This strengthens the finding: every organisation in the figure can act on every part of it.

3. The exposure is not established as current

Version 3 described the credential figures as counting exposure current within twelve months. A date is supplied only for the own-domain group, and the outside-services group has no date field at all. Recency is evidenced for a small minority: 9 organisations (4%) and 14 (6%) respectively. The word "current" has been removed from the headline statistics.

4. The headline percentages understated the finding

Version 3 gave **44% and 56%**. Measured without the unsupported recency filter, the correct figures are **46% (102 of 220) and 63% (148 of 236)**. Version 3's pipeline used the staff date to age out the entire record, including the outside-services portion that carries no date of its own.

5. The login-page pairing is universal, not partial

Version 3 said that for three to four organisations in ten, the attacker also knows which door the key opens. The sets are in fact identical: **65 of 65 and 99 of 99**, with no exceptions in either direction. Wherever the keys are the organisation's own staff addresses, the address of the door is recorded beside them.

Verified as correct and unchanged

Average score 65 in both estates and medians of 69 and 67; header readability; zero ransomware, breach-collection and browser-safety findings; 36 education bodies at 15% of the income cohort; own-domain exposure of 40% against 30% excluding education; the largest staff-credential count of 1,736; and the overlap of 76 organisations giving 380 distinct.

One figure in version 3 could not be reproduced and is not contradicted: the arithmetic estimating what the averages would be had every website served its headers. That calculation is performed at assessment time and is not persisted, so it cannot be re-derived from stored evidence. It is omitted from this version rather than restated.