

AUTOMATED ACCESSIBILITY STUDY · 92 WEBSITES

The same problems, page after page.

What 4,135 scanned pages reveal about recurring website accessibility barriers - and where shared remediation may have the greatest leverage.

90.9%

of websites where recurrence could be evaluated contained at least one high-confidence recurring accessibility pattern.

This report is based on automated accessibility testing. It does not constitute a complete accessibility audit, an assessment of WCAG conformance, or legal advice.

Dataset

92 anonymised business websites
4,135 publicly accessible scanned pages

Testing

RedFlag with axe-core 4.11.2
Scanned 8 August 2026

The scale matters. The repetition matters more.

Accessibility barriers were widespread across the websites scanned. The more useful finding was how often the same problems repeated across pages.

100%

of websites had at least one confirmed automated violation

98.9%

had at least one serious or critical finding

90.9%

of eligible sites contained a recurring pattern

72.2%

of serious and critical eligible findings were recurring

27.7% → 82.8%

Just over one quarter of recurring patterns spanned 20 or more scanned pages - and those patterns accounted for more than four fifths of recurring findings.

Accessibility remediation is often approached as a queue of individual defects. The data suggests that many detected barriers may instead be symptoms of shared components, templates or implementation decisions repeated across a website.

That creates potential remediation leverage. A shared correction may remove the same detectable barrier from many pages.

A recurring pattern does not necessarily equal one code change. It identifies a probable shared implementation worth investigating.

What is RedFlag?

RedFlag is a website accessibility scanning and remediation tool. It scans public pages, shows where detectable barriers occur, and groups repeated findings to help teams identify probable shared implementation problems. This study uses that cross-page analysis to examine how often the same structural pattern repeated.

What automated testing did not establish

This is the most important section in the report. The findings are useful only if the limits of automated testing are understood first.

RedFlag used axe-core 4.11.2 alongside additional RedFlag-native automated checks.

Configured axe-core rule tags

- WCAG 2.0 Level A
- WCAG 2.0 Level AA
- WCAG 2.1 Level AA additions
- WCAG 2.2 Level AA additions

Six additional axe rules were explicitly enabled outside this tag selection. RedFlag-native checks also ran alongside axe-core.

A specific known omission

WCAG 2.5.3 Label in Name was not evaluated. The relevant axe-core rule was not selected by the configured run tags and was not among the additional explicitly enabled rules.

This is one example of why an automated scan should not be interpreted as complete WCAG coverage.

Automation did not comprehensively establish

- complete keyboard journeys
- logical focus order throughout full tasks
- keyboard traps in every interactive state
- screen-reader usability
- whether alternative text is useful or appropriate
- whether labels and instructions make sense in context
- whether error messages help users recover
- caption or audio-description quality
- cognitive usability
- authenticated areas that were not crawled
- downloadable PDFs and other documents
- every responsive or dynamic state
- overall WCAG conformance

Automated engines can also return results that require human review. Those results were not retained in this research export and are absent from the dataset. In addition, some rules cap how many affected elements are recorded per page, so some instance totals are lower bounds.

This report is a lower-bound view of barriers detectable within the tested scope - not a complete accessibility assessment.

Automatically detectable barriers were present across effectively the whole sample.

92

websites analysed
0 excluded

4,135

scanned pages
18 sites reached the 100-page cap

114,843

confirmed automated violation records

99.8%

of scanned pages had a confirmed violation
4,125 of 4,135

95.4%

of scanned pages had a serious or critical finding
3,944 pages

24

median confirmed findings per scanned page
mean 27.8

At website level

92 of 92 websites had at least one confirmed automated violation.

91 of 92 had at least one serious or critical finding.

55 of 92 had at least one critical finding.

Interpretation


These figures establish the scale of automatically detectable problems in this sample.

They do not establish whether any individual website conforms to WCAG.

Results refer to scanned pages, not every page that may exist on each website.

The most widespread issue types

Site prevalence is the primary measure of how widespread a problem was. It is not distorted by the per-page collection limits that affect raw instance counts.



n = 92 websites. Confirmed automated violations only.

Target size was almost universal

91/92

websites affected

97.9%

of all scanned pages affected

100%

median within-site page spread

≥29,980

recorded instances

Small or tightly spaced controls can be harder to hit accurately for people with limited fine motor control or tremor - and for anyone trying to use a phone while moving. Increasing the target area, often through padding, is a common approach, but surrounding layout and spacing still need checking in context.

59.2% of recorded target-size instances belonged to page-rule groups that reached RedFlag's collection ceiling. The prevalence measures are therefore more meaningful than the raw count.

Several common problems spread across almost an entire affected site.

Among websites where each issue was detected, this chart shows the median share of that site's scanned pages carrying the issue.



Important distinction: broad page spread is consistent with repeated implementation in shared templates or components, but page spread alone does not prove a shared cause. The recurrence classifier provides the stronger structural evidence.

Accessible names, contrast and hidden content repeatedly affected large parts of sites.

Links without an adequate accessible name

76.1%

sites affected

54.6%

all scanned pages

84.2%

eligible findings recurring

314

high-confidence patterns

A link without an accessible name can leave a screen-reader user hearing nothing useful when navigating by links - a common failure mode for icon-only links with no visible text or programmatic label. The high recurrence share is consistent with those links sitting in shared headers, footers, cards or social-link groups.

Insufficient colour contrast

75.0%

sites affected

63.5%

all scanned pages

70.6%

eligible findings recurring

802

high-confidence patterns

Text that does not stand out enough from its background can be difficult to read for people with low vision or colour-vision differences - and for anyone trying to read a phone in bright daylight. Contrast appeared on nearly every scanned page of the median affected website, which is consistent with shared colour tokens and reusable styles playing a role.

Hidden content exposed to assistive technologies

75.0%

sites affected

62.6%

all scanned pages

77.3%

eligible findings recurring

242

high-confidence patterns

If visually hidden content remains exposed to assistive technologies, a user may hear or tab into controls that are not visible or currently active - for example an invisible menu, collapsed panel or inactive tab content.

Serious findings reached almost every website.

Severity reflects the impact assigned to the triggered rule. A website is counted once per severity level regardless of how many instances it had.



Severity	Sites affected	Recorded instances	Share of recorded failures
Moderate	92	37,787	32.9%
Serious	91	70,182	61.1%
Critical	55	6,874	6.0%

Images without alternative text

37.0%

sites affected

20.2%

all scanned pages

80.0%

eligible findings recurring

3,184


recorded instances, uncensored

A screen-reader user can receive nothing useful - or only a filename - when a meaningful image has no text alternative. This issue does not fit neatly into either a component or content-only category: its median site spread was **37.7%**, yet **80.0%** of eligible findings recurred. That is consistent with both repeated image components or CMS templates and one-off authoring gaps. Teams should investigate both.

Automation can determine whether alternative text is absent in many situations. It cannot determine whether the replacement text is meaningful or appropriate.

Most confirmed findings belonged to repeated cross-page patterns.

Every confirmed finding was classified as recurring, isolated or unclassified. Findings with inadequate structural evidence were kept unclassified rather than assumed isolated.



A finding qualified as recurring only when

- the same accessibility rule appeared
- on at least two different scanned pages from the same website
- with sufficiently distinctive structural evidence to support a repeated implementation

Only high-confidence patterns are included in public recurrence statistics. Unclassified does not mean isolated.

90.9%

80 of 88 websites eligible for cross-page recurrence analysis contained at least one high-confidence recurring accessibility pattern.

A minority of patterns carried most of the recurring exposure.

The median recurring pattern appeared on five scanned pages. But the distribution was heavily concentrated in a smaller group of very large patterns.

48.4%

2-4 pages

1,145 patterns · 5.9% of recurring findings

16.2%

5-9 pages

383 patterns · 6.1% of recurring findings

7.7%

10-19 pages

181 patterns · 5.2% of recurring findings

27.7%

20+ pages

655 patterns · 82.8% of recurring findings

Just over one quarter of recurring patterns accounted for more than four fifths of recurring issue exposure.

A pattern appearing on 40 pages does not prove that one code edit will fix 40 pages. But it strongly narrows the investigation. Instead of treating 40 findings as unrelated defects, a team can ask whether they share a component, template, design token, CMS pattern or repeated implementation.

For the median site, three patterns explained almost half of recurring findings.

47.6%

median share of a site's recurring findings accounted for by its three largest patterns.



Top three patterns
47.6%

All remaining patterns
52.4%

Measure	Value
Sites with at least one high-confidence recurring pattern	80 of 88 · 90.9%
Sites where top three patterns cover at least half of recurring findings	37 of 80 · 46.2%
Median recurring patterns per affected site	18.5
Median pages affected per recurring pattern	5

This suggests that a website may contain hundreds or thousands of individual findings while a much smaller number of repeated patterns explains a large part of that exposure.

The more serious findings were also more likely to repeat.

72.2%

of serious and critical eligible findings were associated with high-confidence recurring patterns

36.2%

of moderate eligible findings were recurring

Severity	Recurring	Isolated	Unclassified
Critical	71.6%	8.0%	20.4%
Serious	72.3%	14.6%	13.1%
Moderate	36.2%	3.9%	59.9%

The high moderate unclassified share is driven in part by target-size findings, whose exported fingerprints are often too generic to classify confidently. That is an absence of evidence for recurrence, not evidence that the findings are isolated.

For prioritisation, recurring serious and critical patterns combine potential user impact with cross-page scale.

How far an issue spreads and how often it repeats are different measures.

Reading either measure on its own can be misleading. An issue can affect only part of a site and still come from a repeated implementation pattern.

Issue	Median site spread	Eligible findings recurring
Link has no accessible name	95.3%	84.2%
Insufficient colour contrast	97.8%	70.6%
Image missing alternative text	37.7%	80.0%
Generic link text	13.1%	56.2%
Links not distinguishable in text	13.0%	23.3%

Missing alternative text is the clearest example. It reached only 37.7% of the median affected site's scanned pages, yet 80.0% of eligible findings recurred. That is consistent with repeated image components appearing on a subset of page types rather than only site-wide templates.

By contrast, links not distinguishable within text are low on both measures, which is consistent with a more scattered editorial problem.

Spread is not recurrence. Page prevalence tells us how widely users may encounter a problem. Structural recurrence tells us whether the same implementation appears to repeat across pages.

Not every accessibility problem behaves like a component defect.

link-in-text-block provides a useful counter-example.

35.9%

sites affected

9.0%

all scanned pages

23.3%

eligible findings recurring

37.7%

eligible findings isolated

Its distribution is substantially more scattered than accessible-name, contrast or skip-link findings. This is the clearest counter-example to the report's recurrence story: the classifier does not treat every accessibility issue as a shared component problem.

Repeated component problems may need

- design-system changes
- engineering fixes
- shared-template corrections

Scattered content problems may need

- content correction
- editorial guidance
- CMS validation
- training and governance

A clean recurring example: broken skip-link targets

96.7% of eligible skip-link-target findings were recurring. On the median affected website, the issue appeared on **100%** of scanned pages.

This pattern is consistent with a skip link defined once in a shared template: the barrier can repeat on every page using it, while the remediation opportunity may sit in one place.

Where structural evidence supported it, patterns clustered in familiar reusable components.

Component family	Patterns	Share of patterns	Pages affected
Images / media	213	9.0%	1,965
Navigation	172	7.3%	5,676
Modals / dialogs	163	6.9%	3,815
Repeated cards	150	6.3%	1,493
Forms	147	6.2%	2,382
Footer	122	5.2%	2,949
Buttons and controls	117	4.9%	1,229
Header	57	2.4%	1,811

Navigation, headers and footers tended to reach disproportionately large numbers of pages compared with their share of patterns, which is consistent with their use as site-wide components.

Half of patterns could not be assigned confidently to a component family. 16.7% matched multiple families and 33.8% had no sufficiently clear family token. The eight largest classified families shown above account for 48.3% of patterns; a further 1.2% fell into smaller classified families not shown. The report does not guess where the structural evidence is insufficient.

Instance totals are useful - but they are not fairly comparable across rules.

Some RedFlag rules limit how many affected targets are stored for one rule on one page. Those limits differ between rules.

32.1%

of confirmed recorded findings belonged to page-rule groups that reached their collection ceiling

67.5%

of scanned pages contained at least one capped confirmed page-rule group

Censored means the share of recorded instances that belong to page-rule groups where the rule's per-page collection ceiling was reached.

Rule	Recorded	Instances censored	How to read the count
Target size	29,980	59.2%	At least this many
Colour contrast	27,613	15.5%	At least this many
Hidden accessible content	16,393	36.8%	At least this many
Link accessible name	10,453	7.7%	At least this many
Reduced-motion support	4,576	96.0%	Substantially understates
Modal focus management	3,381	72.8%	Substantially understates
Image alt	3,184	0%	Complete recorded count

Because rules are censored at different thresholds, this report deliberately does not rank issue types by raw instance volume.

Technical findings matter because of the barriers they can create for people.

Small interactive targets

A control that is too small or tightly spaced can be difficult to hit accurately for someone with limited fine motor control or tremor - and for anyone using a phone while moving.

Missing accessible names

An icon-only link with no useful accessible name can leave someone navigating by screen reader hearing nothing that explains where the link goes.

Insufficient colour contrast

Text that does not stand out enough from its background can be hard to read for people with low vision or colour-vision differences, and in bright viewing conditions.

Hidden accessible content

A screen-reader or keyboard user may encounter an invisible menu, inactive panel or other control that sighted users cannot currently see or use.

Missing alternative text

When a meaningful image has no text alternative, a screen-reader user may get nothing useful - or only a filename - instead of the information in the image.

Labels, skip links and focus

A placeholder that disappears after typing can remove context from a form. A broken skip link can force a keyboard user through the same navigation on every page. Focus problems can expose controls in the wrong state.

Automation identifies where investigation is warranted. It cannot describe the complete experience of the person encountering the problem.

Where organisations should start

1. **Find recurring serious and critical patterns.** These combine potential user impact with cross-page scale.
2. **Review site-wide components.** Start with navigation, headers, footers, modals, shared forms, repeated cards and global controls.
3. **Address target sizing systematically.** Review component and design-system defaults rather than resizing controls individually.
4. **Review accessible names, contrast and hidden states.** These serious problems affected around three quarters of websites and often spread broadly.
5. **Correct critical image and form issues.** Use automation to locate problems, then human review to confirm meaning and usability.
6. **Treat editorial problems differently.** Content-specific issues may need standards, CMS controls and governance rather than component fixes.
7. **Complete manual accessibility testing.** Include keyboard interaction, screen readers, zoom and reflow, focus behaviour, form validation, dynamic components and error recovery.
8. **Rescan after fixing.** A shared correction should remove detectable findings wherever that implementation appears. Rescanning verifies what changed and helps catch regressions.

Automation should help teams decide where to look. It should not replace human accessibility evaluation.

How the study was conducted

Sample

The 92 websites were drawn from the membership directory of an Australian business network. The sample was not randomly selected and should not be treated as representative of Australian businesses or websites generally. All identities are anonymised.

Scan date and scope

8 August 2026. RedFlag successfully recorded 4,135 publicly accessible scanned pages. The crawl was capped at 100 pages per website; 18 websites reached that cap.

Primary analytical population

114,843 confirmed automated violations. 27,011 advisory or best-practice findings were kept separate and excluded from headline findings. Passes, inapplicable results and needs-review results were not part of the research export.

Definitions

Affected site: a website containing at least one confirmed instance of the specified issue.

Affected page: a unique scanned URL containing at least one confirmed instance of the specified issue.

Issue instance: one confirmed automated violation against one affected target on one scanned page. Where collection caps were reached, this is a lower-bound recorded count.

Recurring pattern: a probable repeated implementation problem within the same website, requiring the same rule on at least two distinct scanned URLs with sufficiently distinctive structural evidence.

Isolated: adequate fingerprint evidence, appears once and does not match a recurring pattern.

Unclassified: insufficient structural evidence to judge confidently. Unclassified findings were not treated as isolated.

Conservative by design

Recurrence eligibility

Only websites with at least two successfully scanned pages were eligible. 88 of 92 websites met that requirement.

Fingerprint quality

Fingerprint quality	Share of confirmed findings
Distinctive	71.1%
Moderate	21.9%
Generic or structural only	7.0%

Every confirmed finding contained both a selector and an HTML snippet. Weak evidence was not forced into a recurrence classification.

URL sensitivity

Fourteen page pairs across six websites differed only by HTTP versus HTTPS scheme. They represent around 0.34% of the page denominator and were retained because matching paths under different schemes do not prove that the delivered pages were identical.

Verification

The main findings were independently re-derived using a second implementation. Headline counts, severity distribution, prevalence and pattern spread reproduced closely. Small differences in pattern grouping were below 0.6 percentage points and were traced to selector escaping and URL normalisation. Final figures use the primary high-confidence analysis pipeline.

What this study does not claim

- that an individual website is WCAG compliant or non-compliant
- that RedFlag detected every accessibility issue
- that a website is accessible because few automated violations were found
- that every recurring pattern equals one exact code change
- that the sample represents Australian businesses generally
- that accessibility remediation guarantees SEO improvement
- that a specific CMS, theme, plugin or framework caused a failure

Follow a finding into the rule documentation.

Every issue referenced in this report has a public rule page at redflagstatus.com/docs/rules. Rule IDs are the identifiers used in scan results. Where the documentation slug differs, the final column gives the public page to use.

Issue	Rule ID	Documentation page
Insufficient colour contrast	color-contrast	color-contrast
Link has no accessible name	link-name	link-name
Hidden content exposed to assistive tech	redflag-hidden-accessible	redflag-hidden-accessible
Touch target too small	target-size	target-size
Image missing alternative text	image-alt	image-alt
Placeholder used as a label	redflag-placeholder-label	redflag-placeholder-label
Skip-link target missing	redflag-skip-link-target	redflag-skip-link-target
Modal without focus management	redflag-modal-focusable	redflag-modal-focusable
Generic link text	redflag-generic-link	redflag-generic-link
No reduced-motion support	redflag-reduced-motion	redflag-reduced-motion
Form input without a label	label	label
Select without an accessible name	select-name	redflag-select-name
Links not distinguishable in text	link-in-text-block	link-in-text-block

Use the rule pages as the bridge from research to action. Each page explains what the detector checks, why the issue matters, and what to investigate before rescanning.

Don't just count the issues. Find what repeats.

Nine in ten websites where recurrence could be evaluated contained at least one high-confidence repeated accessibility pattern.

27.7% → 82.8%

A little over one quarter of recurring patterns accounted for more than four fifths of recurring issue exposure.

For teams responsible for accessibility, the more useful question may therefore be:

Which problems repeat, where do they originate, and which shared fixes could remove the most barriers?

Explore the rules. Then scan your own site.

RedFlag scans public website pages for automatically detectable accessibility barriers, shows where findings occur, and groups repeated patterns to help teams identify where shared remediation may have the greatest leverage.

redflagstatus.com/docs/rules

[Scan your website with RedFlag →](#)

This report is based on automated accessibility testing. It does not constitute a complete accessibility audit, an assessment of WCAG conformance, or legal advice.