

Website Audit Report

Prepared for a client in the B2B software sector. Name and URLs redacted.

SCOPE	AUDIT TYPE	PAGES CRAWLED	DELIVERED
[redacted].com	Complete Audit	412	Day 5

What we found

The site scores well on the checks most audit tools report on, which is why previous automated reports came back clean. The problems that matter are in how the site is crawled and how three templates duplicate each other. Two findings account for most of the lost visibility; the rest are worth fixing but will not move traffic on their own.

4	9	17	23
Critical	High	Medium	Low

Priority order

Ranked by what fixing each item is worth, not by severity score. Work down this list.

#	FINDING	CATEGORY	EFFORT	WORTH FIXING
1	Three templates serve near identical content	Technical	Medium	High
2	Canonical tags point to a staging hostname	Technical	Low	High
3	Product pages excluded by robots.txt rule	Crawlability	Low	High
4	LCP above 4s on the two highest traffic pages	Performance	Medium	Medium
5	Title tags truncated on 84 pages	On page	Low	Medium
6	No structured data on the resource library	Content	Medium	Medium
7	Mixed content warnings on 6 legacy pages	Security	Low	Low

The full report lists all 53 findings in this format. Three are reproduced below so you can see the level of detail each one carries.

Finding 1. Three templates serve near identical content

CRITICAL

TECHNICAL

AFFECTED

/solutions/[slug]
/use-cases/[slug]
/industries/[slug] (94 URLs total)

What is wrong

The three templates render the same body copy with only the heading and breadcrumb changed. Across 94 URLs the shared text accounts for roughly 80 percent of each page. Google has indexed one of the three and is treating the other two as duplicates.

Why it matters

The three sets compete for the same queries, so none of them accumulates the signals that would let one rank. This is the single largest reason organic visibility has stayed flat despite the volume of pages published.

The fix

Pick one template as canonical for each topic, most likely /solutions/. Either rewrite the other two so they answer a genuinely different question, or 301 them into the canonical version. Do not leave all three live with canonical tags pointing at each other, which is the current state and is why Google is choosing for you.

How to confirm it worked

Re run the crawl after deployment. The duplicate cluster count should drop from 94 to 0. In Search Console, the "Duplicate without user selected canonical" bucket should empty over two to three weeks.

Finding 2. Canonical tags point to a staging hostname

CRITICAL

TECHNICAL

AFFECTED

staging.[redacted].com referenced from 41 production URLs

What is wrong

Forty one production pages carry a canonical tag pointing at a staging subdomain that is not publicly reachable. The tag was templated and the environment variable was never swapped at build time.

Why it matters

A canonical pointing at an unreachable URL tells Google the page you want indexed does not exist. Those 41 pages are effectively asking to be dropped from the index, and several of them are your highest value commercial pages.

The fix

Set the canonical base URL from the production environment variable rather than a hardcoded string, and add a build time assertion that fails the deploy if the canonical host does not match the production host. The one line fix is in the shared layout template, not in the 41 pages.

How to confirm it worked

Fetch any affected URL and confirm the canonical resolves to the production hostname. Then use URL Inspection on three of the 41 and confirm Google reports the user declared canonical as the page itself.

Finding 5. Title tags truncated on 84 pages

MEDIUM

ON PAGE

AFFECTED

84 URLs across /blog/ and /resources/ (full list in appendix B)

What is wrong

Eighty four titles exceed the width Google displays, ranging from 71 to 118 characters. The pattern is a descriptive title followed by a colon, a second descriptive clause, and the brand name appended by the layout template.

Why it matters

Google rewrites titles it cannot display, and the rewrite is generated from on page text rather than your intent. On these 84 pages you have no control over what appears in results, which suppresses click through even where rankings are reasonable.

The fix

Target 60 characters including the brand suffix. Remove the second clause rather than shortening both. Where the template appends the brand, strip any brand name already written into the title string, which is the cause of the longest examples in this set.

How to confirm it worked

Crawl and confirm no title exceeds 60 characters rendered. Watch impressions and click through rate in Search Console over three weeks; position should be unchanged while CTR improves.

How this report was produced

Every finding above was verified by hand. The automated crawl produces a candidate list of around 150 checks across 8 categories, and roughly a third of what it flags on any given site turns out to be a false positive or something that does not matter for that site. Those are removed before the report is written. Findings are then cross referenced against your Search Console and Analytics so the ordering reflects your actual traffic rather than a generic severity score.

What you will not find here: a health score on the front page, a list of every missing alt tag, or a recommendation to fix things that have not affected a ranking since 2015.

What happens next

The Complete Audit includes a 45 minute call to walk through the findings and agree the order of work. Implementation is quoted separately once the audit exists, because until then neither of us knows what the work actually is. Plenty of clients take the report to their own developer, and it is written to make that straightforward.

DevelopersMatrix | developersmatrix.com/services/website-audit
Prepared by Syed Bilal Shah