

The B2B Landing Page Conversion Audit Checklist

A structured way to audit a B2B landing page across message match, proof, performance, and mobile usability before assuming a redesign is the fix.

Contents

01 Does the page actually match what brought someone to it?

02 Is the proof on the page inspectable, or just decorative?

03 Is the form asking for more than it needs?

04 Does the page pass Core Web Vitals — and why that is a floor, not a lever

05 What breaks on mobile that a desktop review will not catch?

06 The full audit, in one pass

07 What an audit cannot tell you

Key Takeaways & Sources

01

Does the page actually match what brought someone to it?

Every visitor arrives with an expectation set by whatever they clicked or read first: an ad headline, a specific search term, a subject line, a referral sentence from a colleague. The audit question is not "is this page good" in the abstract — it is "does this page continue the sentence the visitor was already reading." When the headline, offer, and image on the page do not echo the language and promise of the source, the visitor has to do translation work before they can decide anything, and most will not bother.

The Nielsen Norman Group describes this as information scent: people judge whether a link or ad will lead somewhere relevant based on the label, the surrounding context, and what they already know, then decide in seconds whether to keep going. Google's own Ads documentation ties the same idea directly to spend, defining landing page experience as how relevant and useful the page is to people who clicked, and ad relevance as how closely the ad matches the intent behind the search — both scored against competitors and both able to depress Quality Score when they slip.

Concrete signs of mismatch: a broad-match keyword lands on a page written for one narrow service; a "free consultation" ad lands on a page that leads with a demo request; a comparison-intent search lands on a generic homepage instead of a page that names the alternative being compared. None of these are aesthetic problems. They cost conversions because the visitor bounced an expectation, not because a button was the wrong color.

- Read the exact ad copy, keyword, or email subject line next to the page headline — do the nouns and promise match, not just the general topic?
- Check whether the primary CTA on the page matches the action implied by the source (a demo ad should not land on a "download a PDF" page)
- Pull search-term reports for paid campaigns and confirm the highest-spend terms map to a page that speaks their specific intent
- Look for one page trying to serve five unrelated ad groups — a sign the traffic is being forced onto a message it was never written for

They cost conversions because the visitor bounced an expectation, not because a button was the wrong color.

02

Is the proof on the page inspectable, or just decorative?

A logo wall, a five-star badge, and a testimonial signed "J.M., Operations Director" all look like proof and function like wallpaper. The distinction that matters to a serious B2B buyer is whether the claim can be checked. Inspectable proof names the actual client or project (with permission), states the period over which a result was measured, discloses where the number came from, and is specific enough that a skeptical reader could ask a follow-up question and get an answer.

The Nielsen Norman Group's research on trustworthy design points at the same gap from a different angle: people treat company-hosted testimonials with more skepticism than third-party reviews, and genuine credibility tends to come from transparent disclosure of real business details — pricing, process, actual outcomes — rather than symbols meant to imply trust without supporting it. A badge asserts credibility. A named case with a defined measurement window lets the visitor verify it themselves.

A reviewer can run this check without any special tooling: for every proof element on the page, ask what specifically is being claimed, over what time period, measured how, and by whom. If the answer to any of those is "unclear" or "implied," the element is decorative regardless of how polished it looks. Replace vague superlatives ("industry-leading results") with the smallest verifiable claim the team can actually stand behind.

A named case with a defined measurement window lets the visitor verify it themselves.

03

Is the form asking for more than it needs?

Field count and error recovery are well-documented sources of landing-page friction — every extra required field adds a small tax on completion, and a form that fails without saying which field or why loses people at the exact moment they were ready to convert. This audit treats form design as one checkpoint among several rather than the whole review; the field-by-field mechanics of what to ask, in what order, and how to recover from a validation error are covered in depth elsewhere. For this pass, it is enough to flag whether the form looks longer than the offer justifies and whether errors give a specific next step — then route the deeper fix to that dedicated review.

04

Does the page pass Core Web Vitals — and why that is a floor, not a lever

Core Web Vitals measure whether a page loads, responds, and holds still — not whether it persuades anyone of anything. Google's current thresholds, assessed at the 75th percentile of real page loads across mobile and desktop, are Largest Contentful Paint (LCP) at 2.5 seconds or less, Interaction to Next Paint (INP) at 200 milliseconds or less, and Cumulative Layout Shift (CLS) at 0.1 or less. INP replaced First Input Delay in 2024 because it evaluates every interaction across a visit rather than only the first one, giving a fuller picture of whether the page stays responsive as someone actually uses it.

Treat a pass on all three as a prerequisite, not an achievement. A page that loads in under two seconds, responds instantly, and never shifts can still fail commercially if the headline does not match the ad, the proof is not verifiable, or the offer is unclear — speed cannot compensate for any of that. The reverse also holds: a slow or unstable page will suppress conversion regardless of how good the message and proof are, because a meaningful share of visitors leave before the content ever finishes rendering. Audit performance first because it gates everything else, then move on — do not treat a green PageSpeed report as evidence the page is working.

Audit performance first because it gates everything else, then move on — do not treat a green PageSpeed report as evidence the page is working.

05

What breaks on mobile that a desktop review will not catch?

Most landing-page reviews happen on a laptop, at full width, with a mouse — the conditions least likely to match how a majority of paid traffic actually arrives. A page that looks complete on a 1440-pixel screen can have tap targets that overlap, a viewport that forces horizontal scrolling, or a primary CTA parked somewhere a thumb cannot comfortably reach without a second hand shifting its grip.

The W3C's WCAG 2.5.8 Target Size criterion sets a concrete, checkable bar: pointer targets should be at least 24 by 24 CSS pixels, with narrow exceptions for inline text links and cases where an equivalent larger target is available nearby. Below that size, adjacent buttons and links start competing for the same tap, and the cost lands hardest on exactly the fast, low-attention mobile session a paid campaign is trying to convert.

- Load the page on an actual phone, not a resized browser window — emulators miss real rendering and font-scaling quirks
- Measure the primary CTA and any adjacent tap targets against the 24x24 CSS pixel minimum

- Check for horizontal scroll or a viewport meta tag that disables pinch-zoom, both signs of an unset or broken viewport
- Hold the phone one-handed and try to reach the CTA with a thumb from the natural grip position, not a two-handed reach
- Confirm form fields do not trigger the wrong mobile keyboard (a phone field bringing up a full alphanumeric keyboard adds friction a desktop review will never surface)

06

The full audit, in one pass

The table below consolidates the checks above into a single pass a reviewer can run against any landing page. Each row states what to check and the concrete condition under which that area counts as a fail, so the audit produces a specific finding rather than a general impression of "needs work."

Audit area	What to check	Fails if
Message match	Headline, offer, and imagery against the exact ad, keyword, or email that sent the visitor	The page requires the visitor to reinterpret or guess how it relates to what they clicked
Proof and credibility	Named work, disclosed source, and a defined measurement period behind every claim	A claim cannot be traced to a specific client, project, timeframe, or source on request
Form friction	Field count relative to the offer, and whether errors point to a specific fix	The form asks for information not needed for the first reply, or a failed submission gives no actionable next step
Performance	LCP, INP, and CLS at the 75th percentile, mobile and desktop	Any metric misses the good threshold (LCP over 2.5s, INP over 200ms, CLS over 0.1)
Mobile usability	Tap target size, viewport behavior, and thumb-reach to the primary CTA on a real device	A target falls under 24x24 CSS pixels, the page scrolls horizontally, or the CTA sits outside a natural one-handed reach
Offer clarity	What happens immediately after the visitor submits or clicks — is it stated on the page?	The next step (a call, an email, a download) is left implicit or discovered only after conversion

07

What an audit cannot tell you

A completed audit produces a list of friction: a mismatched headline, an unverifiable testimonial, a slow LCP, a CTA out of thumb reach. What it cannot produce is a confident estimate of what fixing any one item is worth. Friction and impact are correlated but not identical — a page can carry an obvious flaw that costs almost nothing in practice, or pass every check on this list and still underperform for a reason the checklist was never built to catch, such as a weak offer or the wrong audience being sent to it.

Fixing everything the checklist flags at once and watching the conversion rate move afterward is not a test — too many variables changed together to say which one mattered, and normal week-to-week variance can be mistaken for a lift or a decline that never happened. Confirming impact means isolating a change (or a small, deliberately grouped set of them) in a proper A/B test, run long enough to cover the natural variation in traffic quality and buying cycles a B2B page sees — which, for most B2B volumes, means planning around full weeks of traffic and a pre-calculated sample size rather than stopping the first time the numbers look favorable. A checklist tells you where to look. Only a test tells you whether looking there was worth the effort.

Key Takeaways

- A page can look finished and still fail if the words on it do not match the ad, keyword, or email that sent someone there.
- Inspectable proof names the work, states the measurement period, and discloses the source; decorative badges and vague testimonials do not.
- Passing Core Web Vitals is a floor, not a growth lever — a fast page can still lose the visitor for reasons speed cannot fix.
- Mobile problems (tap targets, viewport scaling, thumb-reach to the CTA) routinely hide from a desktop-only review.
- An audit finds friction; it does not prove what removing that friction is worth. Confirm impact with a test, not a checklist alone.

Sources

1. Google web.dev: Core Web Vitals thresholds — <https://web.dev/articles/vitals>
2. Google web.dev: Interaction to Next Paint (INP) — <https://web.dev/articles/inp>
3. Google web.dev: Cumulative Layout Shift (CLS) — <https://web.dev/articles/cls>
4. Google web.dev: Optimize Largest Contentful Paint (LCP) — <https://web.dev/articles/optimize-lcp>
5. W3C WCAG 2.2: Understanding Target Size (Minimum) — <https://www.w3.org/WAI/WCAG22/Understanding/target-size-minimum.html>
6. Nielsen Norman Group: Information Scent — <https://www.nngroup.com/articles/information-scent/>
7. Nielsen Norman Group: Trustworthy Design — <https://www.nngroup.com/articles/trustworthy-design/>
8. Google Ads Help: Diagnose Quality Score (landing page experience, ad relevance) — <https://support.google.com/google-ads/answer/2404197>