Common technical SEO mistakes and fixes

You are bleeding traffic and you do not even know why. Most sites I audit have the same recurring wounds—stupid, avoidable stuff that Google penalizes or simply ignores. The phrase "common technical SEO mistakes and fixes" sounds like a checklist, but in practice it is a war against your own sloppiness. Let me show you the five wounds that bleed the most, and exactly how to stitch them shut.

The soft 404 trap: pages that exist but say nothing

A soft 404 is a page that returns a 200 OK status but tells the user (and Googlebot) that the content is gone or irrelevant. Think of a product page that is out of stock, with no redirect and a thin "sorry" message. Google sees a 200, indexes it, and then realizes the page is useless. Your crawl budget gets burned on ghost pages.

The fix is brutal but simple: return a proper 410 Gone or 301 redirect to a relevant category. Do not let dead pages breathe. Use Google Search Console's Coverage report to find soft 404s. Google's docs on HTTP errors explain the difference between a real 404 and a soft one. If you run an e-commerce site with seasonal inventory, automate this with a cron job that checks stock status weekly.

Canonical confusion: when you tell Google two things at once

You slap a canonical tag on a page, but the page also has a noindex directive. That is like saying "please index this version" and "please do not index this page" in the same breath. Google gets confused and often ignores both signals. Another classic: pointing the canonical to a URL that redirects. The canonical tag should point to the final, non-redirecting URL.

Audit your canonical tags with a crawler like Screaming Frog or Sitebulb. Look for conflicts. A single misconfigured canonical on a high-traffic product page can dilute link equity across multiple URL variants. Google's canonicalization guide is short and worth reading twice. The rule: one canonical, no redirects, no noindex.

JavaScript rendering: the silent content black hole

Your site loads a beautiful React app. Googlebot sees a white screen. You wonder why your pages are not ranking. The problem is that critical content—headlines, product descriptions, internal links—is injected via JavaScript that Googlebot cannot or will not execute fully. This is not 2015 anymore; Google can render JS, but it is slower and more error-prone than HTML.

Use the URL Inspection tool in Search Console to see what Googlebot actually rendered. If the rendered HTML is empty or missing key text, you have a problem. The fix is server-side rendering (SSR) or dynamic rendering. Google's JavaScript SEO basics cover the trade-offs. Do not assume your framework handles this automatically—test it.

Broken internal link structures: the crawl budget killer

You have 50,000 pages. Your navigation only links to 500 of them. The rest are orphaned—no internal links pointing to them. Googlebot finds them only through the sitemap, if at all. Orphaned pages do not get crawled frequently, so their content never gets indexed properly. Worse, you waste crawl budget on shallow pages while deep, valuable content rots.

Run a link audit. Use a tool like Ahrefs or a simple Python script to find pages with zero internal inbound links. Fix them by adding contextual links from relevant articles or category pages. A good rule: every page should have at least three internal links pointing to it. Google's crawl budget documentation explains why this matters for large sites.

> Rule of thumb: if a page is not linked from anywhere, it does not exist to Google. Fix the link graph before you chase backlinks.

Image optimization: the lazy load that never loads

You use lazy loading for images. Good. But you set loading="lazy" on the very first hero image above the fold. That image is now delayed, hurting Largest Contentful Paint (LCP). Or you use loading="eager" on 200 images below the fold, causing a massive data download that slows the page to a crawl.

The fix: lazy load only images below the fold. Above-the-fold images should load eagerly. Also, serve images in WebP format with proper dimensions. Core Web Vitals are not optional anymore. Use Lighthouse to check if your lazy loading is actually helping or hurting. A single oversized hero image can tank your LCP score by 2 seconds.

Myth vs reality: three lies people tell themselves

- **Myth:** "My sitemap is submitted, so Google will find everything." **Reality:** A sitemap is a suggestion, not a guarantee. If your internal links are broken, the sitemap alone will not save you.
- **Myth:** "I use a CDN, so my page speed is fine." **Reality:** A CDN helps with latency, but if your JavaScript is bloated and your images are 5MB, the CDN just delivers garbage faster.
- **Myth:** "I have a robots.txt, so I am safe." **Reality:** A misconfigured robots.txt can block entire sections of your site from crawling. Check it at Google's robots.txt tester.

Before and after: a real example

Before: A client had 12,000 product pages. 4,000 were out of stock but returned 200 OK with a thin message. Their canonical tags pointed to themselves, but also had a noindex tag. Google indexed 6,000 pages, but only 1,200 had real content. Organic traffic was flat for 8 months.

After: We set out-of-stock pages to 410. Removed all conflicting noindex/canonical combos. Added internal links from category pages to every product. Within 6 weeks, indexed pages dropped to 3,000, but organic traffic increased by 40%. Fewer pages, better quality, higher rankings.

FAQ: quick answers to common doubts

Q: Should I fix all technical errors at once?
A: No. Prioritize by impact. Fix soft 404s and canonical conflicts first—they waste crawl budget. Then tackle rendering and image issues. Speed improvements come last if the foundation is broken.

Q: How often should I audit my site?
A: Monthly for sites over 10,000 pages. Quarterly for smaller sites. Use Search Console's Coverage report as your early warning system.

Q: Do I need to fix every single 404?
A: No. A few 404s from old external links are fine. Fix 404s that are linked internally or that receive significant traffic.

Q: Is dynamic rendering a hack?
A: Google says it is acceptable, but it adds complexity. Server-side rendering is cleaner if you can afford the server cost.

Q: Can I ignore Core Web Vitals?
A: You can, but your competitors will outrank you. Google uses them as a ranking factor, and users bounce from slow sites anyway.

Stop fixing symptoms, start fixing causes

Technical SEO is not about ticking boxes. It is about removing friction between your content and the search engine's ability to understand it. The five mistakes above account for 80% of the issues I see in audits. Fix them, and your crawl budget, indexation, and rankings will improve without any link building or content rewriting. Do the hard work of cleaning up your own house before you blame the algorithm.