

MARKETING MEASUREMENT

The Marketing Audit Cookbook: Score Twelve Points, Log the Evidence, Fix in Order

The working companion to a 12-point marketing audit: a scoring worksheet, an evidence log and a findings-to-fixes plan covering tracking, response time, CRM handoffs, pages, search, AI visibility, ads and reporting.

Contents

01 What a marketing audit is for, and what it is not

02 The twelve points and what each one asks

03 How to score: 0 to 3, with anchors and a hard evidence rule

04 The evidence log: what counts and how to record it

05 From findings to a fix plan that gets done

06 A two-week schedule, the review meeting and the re-score

Key Takeaways & Sources

01

What a marketing audit is for, and what it is not

A marketing audit is a structured look at whether the machinery between a stranger and a closed sale works, with every conclusion tied to something you can point at. The useful version answers three questions. What is broken or missing right now? How sure are we, and what did we look at to be sure? What gets fixed first, by whom, and how will we know it worked? Anything that cannot answer all three is a presentation about marketing, which is a different and cheaper thing.

This cookbook is the working companion to a twelve-point audit. It gives you the scoring sheet, the evidence log and the fix plan, plus the method for filling them in. It does not tell you what good looks like in numbers, because we have no benchmark that holds across businesses and we would rather leave a gap than invent one. Where a threshold is quoted, it comes from a named source and is dated. Where a figure appears in an example row, it is labeled illustrative and exists only to show the arithmetic.

Decide the scope before anyone opens a dashboard. Name the business unit or offer being audited, the period the data covers, and the one commercial question the audit has to serve, for example "why do paid leads close at a lower rate than referrals" or "can we trust the pipeline number we report each month". A stated question keeps twelve points from becoming twelve unrelated opinions.

Pick one person to own the audit file and one person with authority to approve the fix plan. They can be the same person in a small company. The rule that matters is that the file has an owner who is not the person being scored. If the same person runs the ads and audits the ads, the evidence log is where the conflict shows up, so make the log the thing a second person reads.

- Scope: offer or business unit, period covered, and the single commercial question the audit serves
- Roles: an audit file owner, a decision-maker for the fix plan, and a named person for each point who can answer questions about it
- Rule one: no evidence, no score
- Rule two: every low score ends in a fix with an owner, or in a written decision not to fix it

It gives you the scoring sheet, the evidence log and the fix plan, plus the method for filling them in.

02

The twelve points and what each one asks

The twelve points cover eight areas in the order a lead experiences them: tracking, response time, CRM handoffs, pages, search, AI visibility, ads and reporting. The order is deliberate. Tracking comes first because every other area reads its numbers from it. Reporting comes last because it is only as good as everything that feeds it.

Each point is written as a question that has a yes, a partly or a no, so two reviewers looking at the same evidence reach the same score. The anchors for 0 to 3 live in the worksheet. The table below gives the question and the first piece of evidence to collect. The evidence column is a starting point; the log lets you add what you actually used.

#	Area and point	The question	First evidence to collect
1	Tracking: conversions	Does every conversion that matters fire once, with the right value and source, on a real test?	Test submissions for each form, call and booking path, seen in the debug view
2	Tracking: source capture	Does each lead record carry its original source, campaign and click identifier when it reaches the CRM?	Ten recent CRM records traced to the ad or page that created them
3	Response time	Is the time from inquiry to first human reply measured, and is it inside the target the business wrote down?	Timestamps of inquiry received and first reply, for a full week
4	CRM: stage definitions	Do marketing and sales use the same written definitions for each lifecycle stage?	The definitions document and a sample of records checked against it
5	CRM: ownership and rejection	Does every lead have one owner and one next step, and does a rejected lead carry a reason from a fixed list?	Leads with no owner or activity; rejections with a blank reason
6	Pages: message and proof	Can a visitor from the main paid or search entry point tell what is offered, to whom, and why to believe it, without scrolling?	Entry pages reviewed against the offer; claims matched to evidence

#	Area and point	The question	First evidence to collect
7	Pages: speed and forms	Do key pages meet the published Core Web Vitals targets for most visits, and do the forms work on a phone?	Field data for entry pages; a recorded phone form submission
8	Search: technical	Can search engines crawl, index and consolidate the pages that should rank?	Index coverage report, site crawl, canonical and redirect checks
9	Search: coverage	Do the pages cover the questions buyers ask at each stage, without duplicating each other?	Query list mapped to pages, with gaps and overlaps marked
10	AI visibility	When real buyer questions are put to AI answer tools, is the business present, described accurately and linked?	A logged prompt set, run on dated occasions, answers saved
11	Ads: structure and signal	Does each paid campaign optimize toward a conversion that reflects a qualified lead, not only a form fill?	Account export of campaign goals and primary conversion actions
12	Reporting	Does the number reported as revenue by channel reconcile with the CRM within a stated tolerance?	One closed month, platform conversions against CRM records by source

03

How to score: 0 to 3, with anchors and a hard evidence rule

Score each point on four levels. A 0 means the thing is absent or has never been checked. A 1 means it exists but is known to be wrong or only partly in place. A 2 means it works for the main path and has gaps in the edges. A 3 means it works, someone checks it on a schedule, and the check is written down. The words "someone checks it" matter. A point that was set up carefully two years ago and has not been looked at since is a 2 at best, because nothing in the audit can show that it still works.

The evidence rule is simple. A score cannot be entered in the worksheet unless the log has at least one row with the same point number. If the evidence could not be collected, the point is marked unscored and counts as 0 in the total, with a note saying why. Treating unscored as 0 is a choice, and it is a deliberate one. A point nobody can evidence is a point nobody can rely on, which is the same operating condition as a failed one.

The worksheet adds up three things, and each is a plain calculation. The area total is the sum of the scores for the points in that area (tracking has two points, so its total runs from 0 to 6). The overall total is the sum of all twelve points, from 0 to 36. The completion rate is the number of points with a log entry divided by twelve. Report the completion rate beside the total every time. A total of 28 with a completion rate of 50 percent is a claim about six points, not twelve.

Do not read the total against anyone else. We know of no external benchmark for a score of this kind, and a number borrowed from a vendor report would measure their method rather than yours. The total has two jobs. It sets the order of work, because the lowest scores in the earliest areas get fixed first. And it gives you something to compare against when you re-score the same twelve points on the date you set at the start.

Have two reviewers score independently, then compare. Where they differ by more than one level, do not average: open the evidence log row and decide which reading it supports. The gap usually shows an unclear anchor, or that one reviewer saw a different page, account or period.

Score	Meaning	What the evidence log must hold
0	Absent, or never checked	What was looked for and where, or "unscored: evidence not obtainable"
1	Present but known to be wrong or partial	The failure observed, with date and source
2	Works on the main path, gaps at the edges	The main path test, plus the edges that failed
3	Works and is checked on a schedule	The test, the schedule and the last two check dates

04

The evidence log: what counts and how to record it

The evidence log is the part of the audit that survives staff changes. Each row records one observation: the point it supports, the date, who collected it, the method, the system it came from, a link or file name for the saved proof, what was seen, and what the check could not show. That last column is the one people skip and the one a skeptical reader looks at first. A single test submission proves the form fired once. It does not prove the form fires on every browser.

Some checks have a documented tool, and the log should name it. For conversion tracking, Google Analytics 4 has a DebugView that shows the events and user properties the property collects from a device in real time, and it requires debug mode to be switched on. Use it on your own device while you submit a real test inquiry, then save a screenshot with the time visible. For ads, Google Ads describes enhanced conversions as sending hashed first-party data, such as an email address, to improve conversion measurement; if the audit finds the data is not being captured at the moment of the form submission, log that as the evidence for a low score on point 2 or 11, not as a footnote.

The CRM checks have a documented trap worth testing for. HubSpot states that a closed-won deal appears in its ads attribution report only when the associated contact sits in the default Customer lifecycle stage, and that a deal tied to a custom stage is not included. If your team renamed or replaced that stage, the audit should compare the ads report against a list of closed deals and log the gap. The lifecycle stages guide in the same knowledge base explains how stages are set and changed, which is where point 4 starts.

Page checks use published thresholds. The web.dev guidance on Core Web Vitals sets targets of 2.5 seconds or less for Largest Contentful Paint, 200 milliseconds or less for Interaction to Next Paint, and 0.1 or less for Cumulative Layout Shift, assessed at the 75th percentile of page loads and segmented across mobile and desktop. For point 7, log the field values for the entry pages, not a single lab run, and note which device class each value covers.

Search checks start with what Google documents. Links are crawled when they are anchor elements with an href attribute, and descriptive anchor text helps people and search engines understand the destination. A canonical tag is the documented way to consolidate signals for duplicate or similar URLs into one preferred address. Use those two pages as your reference when logging point 8, and write down which pages were crawled and when.

For AI visibility, the logging problem is different because the answers change. Google states that there are no additional technical requirements for a page to appear as a supporting link in AI Overviews or AI Mode beyond being indexed and eligible to be shown in Search with a snippet, and that no special markup or text file is needed. So the technical half of point 10 is the same as point 8. The other half is a prompt log: a fixed list of real buyer questions, each run on a stated date on a named surface, with the full answer saved, and a note on whether the business was named, described accurately and linked. Treat any single run as one sample.

Log column	What to write	Why it is there
Point	The number 1 to 12	Ties the row to a score; a score without a row is invalid
Date and collector	Date of the check and who ran it	Lets a second person repeat it
Method	The steps, in one or two lines	A check that cannot be repeated cannot be re-scored
System and view	Tool, account, report name and filters	Filters change numbers; record them
Proof	File name or link to a screenshot, export or recording	Keeps the claim attached to its support
What was seen	The observation in plain words, with numbers	Separates observation from interpretation
Limits	What this check could not show	Stops a narrow test being read as broad

05

From findings to a fix plan that gets done

Once the points are scored, the work shifts from finding to deciding. Take every point scored 0 to 2 and write one finding for it in a single sentence that names the failure and the evidence row. "Form submissions on the pricing page create no CRM record when the visitor arrives from a paid link, see log row 14" is a finding. "Tracking is inconsistent" is a mood. Each finding then gets a proposed fix, an owner role, a test that proves the fix worked, and a re-score date.

Rank the findings with a calculation you can explain in one sentence. The fixes-plan sheet uses: priority equals the gap, times the impact, divided by the effort. The gap is 3 minus the current score. Impact is 1 for low, 2 for medium and 3 for high, judged by how many leads or how much revenue depends on the point. Effort is 1 for under a day, 2 for under a week and 3 for more. A point scored 0 with high impact and one day of effort gets 3 times 3 divided by 1, which is 9. A point scored 2 with low impact and a week of effort gets 1 times 1 divided by 2, which is 0.5. Those numbers are illustrations of the arithmetic. They are not a benchmark.

Break ties with sequence, not with preference. If a tracking finding and an ads finding both score 9, fix tracking first, because the ads finding cannot be verified until tracking is reliable. This is the reason the twelve points run in their order. A fix that depends on another fix should say so in the dependency column, and the dependent one should not be scheduled earlier than its prerequisite.

Every fix needs a gate written in the language of evidence. "Improve source capture" is an intention. "Ten consecutive new CRM records created from paid clicks carry a source, a campaign and a click identifier" is a gate. When the owner says the fix is done, a second person runs the gate and writes the result into the evidence log with a new date. Only then does the point become eligible for a higher score.

Some findings should not be fixed. A low score on a channel you plan to stop using is information, not a to-do. Write the decision down with the date and the person who made it, so that the same point does not return as a surprise in the next audit. A declined fix is a legitimate row in the plan, provided it has a reason.

The fixes-plan sheet uses: priority equals the gap, times the impact, divided by the effort.

06

A two-week schedule, the review meeting and the re-score

A twelve-point audit fits in two working weeks when the access is arranged in advance. Most delays come from waiting for a login. Request read access to analytics, the ad accounts, the CRM, the search console property and the website CMS on day one, and name a person who can answer questions about each of them. The schedule below groups the points by who is needed, not by area, so that each person is asked once.

Hold one review meeting at the end of the second week, with the decision-maker in the room. Walk the points in order, state each score and the log row behind it, and spend the time on disagreements and on the ranked fix list. Leave the meeting with every fix assigned to a role, every gate agreed and a calendar date for the re-score. If the meeting ends without those three things, schedule the next one before people leave.

Re-score on the date you wrote down at the start. Ninety days is a common choice because most fixes land inside it, though the right interval depends on how fast your account and site change. On the re-score, repeat the same checks in the same way and add new log rows rather than editing old ones. The old rows are what let you show that a score moved because the evidence changed.

Be honest about what the audit cannot see. It cannot tell you whether the offer is right, whether the price is right or whether buyers who never touched a measurable channel were influenced by something else. It reports on the machinery. Those other questions need customer conversations, sales call review or a test, and they belong in a different piece of work with its own evidence.

Day	Points	People needed	Gate to pass before moving on
1	Scope, roles, access requests	Decision-maker, file owner	Scope question written; access granted
2 to 3	1, 2, 12	Analytics owner, CRM administrator	Test submissions logged; one closed month exported

Day	Points	People needed	Gate to pass before moving on
4 to 5	3, 4, 5	Sales manager, CRM administrator	A week of timestamps exported; stage definitions found or noted missing
6 to 7	6, 7	Web owner, designer or developer	Field data for entry pages saved; form tested on a phone
8 to 9	8, 9, 10	SEO owner, content lead	Crawl complete; prompt list agreed and first run saved
10	11	Paid media owner	Account export saved; conversion actions listed with primary or secondary status
11	Scoring pass	Two reviewers	Both sets of scores recorded; disagreements of more than one level listed
12	Fix plan drafted	Audit file owner	Every score from 0 to 2 has a finding, a fix and a gate
13 to 14	Review meeting and re-score date	Decision-maker, point owners	Owners, gates and re-score date agreed in writing

Key Takeaways

- A marketing audit earns its cost when it ends in a ranked list of fixes with owners. A scored report that nobody acts on is a snapshot of problems, not an audit.
- Score twelve points from 0 to 3 against written anchors, and refuse to record a score that has no row in the evidence log. No evidence means the point stays unscored.
- Start with tracking and lead capture. Every later score (response time, handoffs, ads, reporting) rests on numbers that tracking produced, so a weak point 1 caps how far you can trust the rest.
- There is no external benchmark for the total. Use the score to decide the order of work and to compare your own audit with your own re-score, not to compare yourself with other companies.
- Google's documentation says there are no additional technical requirements for appearing in AI Overviews or AI Mode beyond being indexed and snippet-eligible, so the AI visibility point is audited with ordinary search checks plus a logged set of real prompts.
- Rank fixes by the size of the gap, the impact if closed and the effort to close it, give each one an owner role and a test that proves it worked, and re-score on a date you write down before you start.

Sources

1. Monitor events in DebugView — Analytics Help (Google) — <https://support.google.com/analytics/answer/7201382>
2. About enhanced conversions — Google Ads Help — <https://support.google.com/google-ads/answer/9888656>
3. Ads attribution in HubSpot — HubSpot Knowledge Base — <https://knowledge.hubspot.com/ads/ads-attribution-in-hubspot>
4. Use contact and company lifecycle stages — HubSpot Knowledge Base — <https://knowledge.hubspot.com/records/use-lifecycle-stages>
5. Web Vitals — web.dev — <https://web.dev/articles/vitals>
6. SEO Link Best Practices for Google — Google Search Central — <https://developers.google.com/search/docs/crawling-indexing/links-crawlable>
7. How to Specify a Canonical with rel="canonical" and Other Methods — Google Search Central — <https://developers.google.com/search/docs/crawling-indexing/consolidate-duplicate-urls>
8. AI Features and Your Website — Google Search Central — <https://developers.google.com/search/docs/appearance/ai-features>