

English translation · Edition 2

Website, Portfolio, and Search

A practical guide to a service provider's website

Contents

How to read this book

I The economics of a service provider's website

II How the web works

III How search engines interpret meaning

IV Architecture and URLs

V Content as a database

VI The case study as a genre

VII The service page

VIII Editing

IX Practical keyword research

X Trust

XI Production

XII Stack and infrastructure

XIII A design system for a content website

XIV Measurement

XV Law and ethics

A Checklists

B Templates

C Structured data reference

D A catalogue of clichés

E Index

F Glossary

G Bibliography

H Three worked examples

How to read this book

Scope and edition

This book covers a service provider's website: economics, search, structure, content, and maintenance. Edition 2 retains the fifteen parts and adds three worked examples. Neither the number of pages nor the age of a website guarantees commercial results.

This book is primarily for independent professionals and small studios. Other types of business will need to reconsider the proposed solutions. Google's documentation does not describe Yandex's algorithm, and findings from a foreign sample are not benchmarks for the Russian market.

Evidence labels

[data, N] — a finding from the cited study. Read it together with the sample, period, and unit of measurement. Applying the finding to another project requires separate justification.

[standard, N] — official documentation or a specification. A recommendation in documentation is not necessarily a mandatory requirement; the text preserves that distinction.

[example] — a hypothetical situation, illustrative estimate, or calculation based on stated assumptions. All input figures and outcomes are invented for teaching. This is neither a real client case study nor an industry observation.

[practice] — a working guideline from the previous edition, unsupported by a published sample. Test it on your own project; it should not be the sole basis for a budget.

[opinion] — an author’s judgment or a proposed procedure. Numbers in such a procedure are settings, not verified thresholds of effectiveness.

[verify independently] — a decision that depends on current documentation, a contract, jurisdiction, or a specific implementation. This label does not replace a source for a factual claim.

Section, item, and field numbers do not need evidence labels. Explanations of mechanisms still require verification: a plausible explanation is not proof of causation.

Order of parts

Parts I-III establish the economics and the foundations of search and performance. Parts IV-V and IX cover structure, the content model, and demand. Parts VI-VIII address case studies, services, and editing. Parts X-XV cover trust, production, infrastructure, design, measurement, and publication.

Use this book as a reference. For a new website: I → IV → V → IX → XII → XIV. For a website receiving no inquiries: XIV → I.6 → VII → IX. To prepare a case study: VI → VIII → Appendix B.

Three worked examples

A — an independent professional’s new website. A limited time budget, one service, and two case studies. The task is to complete the first version and begin testing demand.

B — a website with traffic but almost no relevant inquiries. Separate visitor groups, check the form and the quality of inquiries, and compare the periods before and after the changes cautiously.

C — an unsuccessful launch. Large-scale publishing without a clear commercial objective, a form submission delivery failure, a limited repair effort, and a decision to abandon the previous plan.

Full descriptions appear in Appendix H. Short continuations in the chapters use the same data. The examples illustrate the order of decisions; they do not prove that the techniques are effective.

Boundaries of verification

The September 10, 2026 revision checks claims, calculations, and sources across all chapters and appendices. The accompanying fact-check report documents corrections and limitations; original reference numbers are retained.

This AI-assisted editorial check is not independent peer review or legal advice. Unreferenced workflows, visitor models, schedules, and design settings are author proposals, not measured universal rules. Examples are hypothetical; sources support only the stated claims and contexts.

I · The economics of a service provider's website

1.1 What clients are buying when they read your website

Design services can have credence attributes: clients may find some aspects of quality difficult to assess even after delivery. This does not make all design quality unknowable; agreed criteria, comparison, expert review, and user testing can provide evidence. [54]

Alongside the work itself, clients can use signals to assess risks such as wasted money, missed deadlines, and rework. The following three-risk model is the author's planning framework, not a measured classification of all purchasing decisions.

A website helps address three kinds of risk, each requiring its own evidence.

Risk of outcome is addressed by showing that a similar problem has been solved before and explaining the outcome. Specifics matter: the situation before and after, and the units used to measure the change.

Risk of the process is addressed by demonstrating predictability. Clients need to know what will happen in week three, how many rounds of revisions the price includes, what they must provide, and what happens if they dislike the result.

Risk of existence is addressed by showing that the provider is real and will not disappear: testimonials with names and employers, publications, talks, open-source code, and a project history spanning several years.

A gallery helps visitors assess visual quality, but rarely explains timelines, decisions, or responsibilities. A detailed case study adds that information. Whether it generates more inquiries must be tested on the particular website, rather than assumed.

Clients may use clear explanations, visible work, and a usable website to assess a provider. Treat the effect of typos or loading speed on purchasing as a hypothesis to test, not as proof that clients cannot evaluate any other aspect of the work. [opinion].

1.2 The four functions and conflicts between them

A personal website usually tries to be four things at once.

Portfolio — evidence of competence for people who already know the provider's name and want to assess their work.

Expert platform — regular publications that build recognition within a professional community. **Content hub** — pages that answer search queries and attract people who do not yet know the provider. **Inquiry generation** — commercial pages that turn attention into an inquiry.

These functions emphasize different things. Each page needs to reconcile them: what readers should see first, which details they need, and what they should do next.

In a portfolio, visuals help readers assess the work, while text explains the task and the author's contribution. The balance depends on the service and audience. A gallery and a detailed explanation can coexist when a page lets readers scan the result first and explore the details afterward.

Keep time-sensitive publications current and refine service pages when the offer or evidence changes. A six-month publishing gap does not by itself prove abandonment. Use your own inquiry data to decide which pages deserve maintenance before adding new articles.

Different visitor groups have different goals. Assess conversions from articles and service pages separately, accounting for the acquisition channel and later interactions. An observed difference does not establish a universal ratio: an article that precisely answers a buyer's question can also generate qualified inquiries.

Assign different functions to different page types while keeping links between them. Useful measures for a portfolio include case-study views and inquiries; for expert articles, returning visitors, mentions, and their contribution to the path to a contract. Impressions and rankings describe search visibility, not commercial success on their own.

1.3 Channel map: owned and rented channels

Acquisition channels differ in three respects. Confusing them causes many planning errors.

Control. On a domain they maintain, providers can control URLs, content, and redirects; search engines control indexing and ranking. Platform portability depends on export, custom-domain, and redirect facilities. Preserve your own content copies and check migration options before committing to a platform.

Inertia. Advertising can bring visitors soon after launch; search visibility often takes longer to build. Organic traffic can nevertheless change abruptly after a technical failure, a shift in demand, or changes to search results. Referrals depend on relationships with clients and partners and have no fixed lead time.

Reuse. A page can be updated and shared for years, while a platform post introduces readers to it. A post’s lifespan depends on the platform and how it is distributed. Compare channels by the inquiries they generate and their costs, including the work of adapting content.

Channel	What does the author control?	What to check
Own website and search	content, URLs, and analytics; not search results	demand, visibility, inquiries
Portfolio platforms	their own content, within the platform’s rules	quality of incoming inquiries

Channel	What does the author control?	What to check
Social media	posts; not the distribution algorithm	reach, click-throughs, and subsequent interactions
Newsletter	content and a lawfully collected subscriber list	delivery, responses, and unsubscribes
Paid advertising	Settings and budget within the platform	cost of qualified inquiry
Referrals	quality of work and relationships	source and quality of projects

The table suggests a practical arrangement rather than a preference for one channel. The website is the repository: it holds the full version of each piece, under the author's control. Platforms are shop windows, presenting a short version linked to the full one. A case study written once can be shared across four or five platforms while remaining at an address the author controls.

Publishing primarily on a platform increases dependence on its rules. It does not automatically transfer copyright or make reuse impossible; check the terms and export options. [opinion].

1.4 Who visits the website, and what each visitor checks

For this planning exercise, group visitors by their task and role in the purchasing decision. These are hypotheses to test; demographic characteristics can also matter in a particular market.

The decision-maker—an owner, founder, or executive—spends their own money or is personally accountable for the budget. They skim rather than read closely, looking for two answers: have you done similar work, and roughly what does it cost? Failure to find a price indication drives more people away than is often assumed: an absent benchmark can suggest an expensive, unpredictable engagement.

The shortlister—a manager, marketer, or producer—prepares a shortlist for the decision-maker and must be able to justify it. They need comparable facts to put in a comparison table. Copy without clearly identified scope, timelines, and deliverables is therefore less useful than copy that provides them, even if it reads better.

A peer may come to learn, then link to your work, refer a client, or propose a collaboration. An informational query does not identify the visitor's role conclusively, so a blog should also be assessed by its contribution to the path toward an inquiry.

A verifier already knows the provider's name and may inspect About, portfolio, and other evidence. The size and conversion rate of this group must be measured; neither is fixed by the role.

Role	Pages visited	What they check	What puts them off
Decision-maker	service, case studies, contact details	results and indicative pricing	no price guidance or figures in case studies
Shortlister	service, process, FAQ	scope and timelines	vagueness
Colleague	blog	depth and novelty	superficiality, a retelling of the known
Verifier	About, portfolio	authenticity and competence	No names, no dates, no reviews.

A page written for everyone serves no one. A common failure is a service page written like a blog post: it explains what a design system is instead of describing the scope, timeline, and price. The decision-maker leaves without an answer; a peer finds an answer but does not buy.

1.5 Scenarios leading to an inquiry

The following four entry routes are planning examples, not an exhaustive classification. Match the next step to the visitor's task and test the proposed journey.

Visitors arriving through a commercial query are looking for a provider: they are ready to act but do not yet trust you. The path is short—a service page, a case study, contact—and every unnecessary step can cost inquiries.

An informational-page visitor primarily wants an answer. Some may also be choosing a provider, so do not assume that

none is ready to talk. Give a useful answer and an unobtrusive route to a related service.

Referred visitors may arrive with some prior trust and look for confirmation in biographies, cases, and testimonials. This does not guarantee the highest conversion rate. Visitors from portfolio platforms may want project details before viewing the service and contact information.

The next step should be clear from the entry page. Make contact details available in the main navigation and alongside the service offer. The number of clicks is a diagnostic measure, not a universal three-click limit.

When reviewing the path, look for hard-to-find contact details, unexplained required fields, and case studies unconnected to a service. An extra field may discourage inquiries or help qualify them. Judge it by inquiry quality and processing time, not just submission count.

1.6 The arithmetic of inquiries

Start the plan with the required number of qualified inquiries: relevant tasks with a sound basis for discussing a project. Spam submissions and clicks on a messaging link do not count. Track inquiries, contracts, and revenue separately.

A scenario calculation needs two assumptions: the click-through rate from impressions and the inquiry rate from visits. Until you have your own observations, these are chosen inputs, not industry benchmarks. Use figures from someone else's analytics only after checking the market, channel, conversion definition, and sample composition.

The illustrative target is five qualified inquiries per month. Every value in the table is chosen to demonstrate

sensitivity to the assumptions; none of these scenarios predicts that particular rankings will be achieved. [example]

Scenario	Inquiries/visits	Clicks / impressions	Visits required	Impressions required
Conservative	0,5%	5%	1 000	20 000
Base	1,5%	10%	about 334	about 3,334
Favorable	3%	15%	about 167	about 1,112

Formulas: visits = goal ÷ conversion; impressions = goal ÷ conversion ÷ CTR. In each row, both values are calculated from the original goal and rounded up independently. This is the volume required under the chosen conditions, not the promise of available demand.

These formulas do not determine a page count. A page may answer many queries, and query groups may overlap. First check demand and the appropriate page type, then estimate the attainable share of impressions and clicks. You cannot call ten pages sufficient without explaining where visitors will come from.

If one quarter of qualified prospects sign a contract, five inquiries yield an expected 1.25 contracts per month. This is an average in a hypothetical model, not a fractional real order or a guarantee of one contract every month. Assessing payback also requires the project margin, acquisition cost, and website maintenance hours. [example]

Worked example A — an independent professional’s new website. The designer sets aside eight hours a week for twelve weeks: 96 hours. The first version and measurement receive 72 hours, with 24 reserved for corrections. They start

with one service and two case studies because material is available for them, not because three pages must deliver a particular number of clients. Full details are in Appendix H. [example]

1.7 Page cost in hours

Plan in hours, not pages. Break a page’s cost down by task; the time required for each task varies more than budgets usually allow for.

The estimates below are an illustrative budget for manual production, not measured industry benchmarks. The assumed content length is an input to the calculation, not a requirement for the page. [example]

A service page, assumed to contain 2,500 words:

Operation.	Hours
Collecting and cleaning relevant search queries	1-2
Analyzing search results for target queries	1-2
Page structure and plan	1
Drafting	4-8
Gathering facts and figures and defining the scope of work	1-2
Editing	2-3
Visuals	2-4
Page layout and publishing	1

Operation.	Hours
Internal linking and checks	0,5
Total	13,5-23,5

A case study of 3,000-8,000 words: gathering material and obtaining client approval: 3-8 hours (the least predictable task because it is not entirely under the author's control); writing: 8-16; visuals: 6-12; publishing: 2. Total: **19-38 hours.**

Article, 2,500 to 5,000 words: 10-20 hours provided the author is writing from personal experience. Preliminary research adds 5-15 hours when needed, and is where deadlines are most often missed.

Now apply these estimates to a hypothetical plan of 25 service pages, 25 case studies, and 80 articles:

$25 \times 19 + 25 \times 30 + 80 \times 15 = 475 + 750 + 1\,200 = \mathbf{2,425 \text{ hours.}}$

Converting hours into calendar time:

Work schedule	Time required
10 hours a week, 52 weeks a year	4.7 years
20 hours a week, 52 weeks a year	2.3 years
Full-time: 40 hours per week, 46 weeks per year	1.3 years

In the full-time scenario, all 40 weekly working hours are allocated to content. Client work would require additional capacity or a longer schedule; the calculation does not predict the provider's actual income.

In this calculation, 130 pieces require 2,425 hours, excluding a separate maintenance budget. Recalculate a larger plan using your own mix of content and the time actually spent on trial publications. Automation may alter the estimate, but fact-checking and approvals still take time.

A practical rule: production capacity equals available hours divided by average hours per page. Work beyond that capacity does not enter the plan.

1.8 Time to returns and the shape of the growth curve

Time to initial results depends on the existing site, competition, demand, page quality, and changes made. Age alone predicts neither growth nor a plateau.

Ahrefs studied a million random URLs first discovered by its crawler in September 2023: 1.74% reached the top ten within a year. In a second sample of two million URLs created in October 2023 with nonempty English content, the share was 6.11%. The authors note that the second estimate may be more representative. These are sample findings, not the probability that your page will succeed. [data, 7]

Of the pages that reached the top 10, 40.82% did so in the first month. That selection condition matters: studying successful pages does not tell us what happened to all published pages. The study also measured URL age from the crawler's first discovery, which cannot automatically be treated as the publication date. [data, 7]

Build the work schedule around checks. After launch, check page availability, forms, and event collection. After crawling, check indexing and the canonical URLs selected by the

search engine. As impressions accumulate, check whether queries match the content and lead to clicks. Once inquiries arrive, assess their quality and resulting contracts.

There is no universal rule against changing direction before month nine. Fix technical faults immediately. Revisit a commercial hypothesis when its predetermined budget is exhausted or evidence contradicts its assumptions. With little data, the decision may still be economic: continuing can be too expensive even without statistical proof of failure.

Worked example C — an unsuccessful launch. The owner spends 120 hours publishing 60 overview pages, but after eight weeks records only 240 search visits and no confirmed inquiries. In this hypothetical case, further bulk publishing is stopped because the topics have little connection to the service and the form is faulty—not because eight weeks is enough for a definitive SEO verdict. [example]

1.9 Zero-click search

A search ranking is not a visit. Ads, direct answers, maps, videos, and generated answers stand between a result and a click. The mix depends on the query, country, device, and time of observation.

In a study of 55,393 popular queries observed over forty days in spring 2026, AI Overview appeared in 13.7% of searches overall and 64.7% of question queries. This is an association within the selected dataset, not proof that question wording alone triggers an answer. Topic and wording may act together. [65][data, 15]

A short definition or reference figure may satisfy a user directly on the results page. A detailed project case study,

information about a specific provider, or a working tool gives users additional reasons to visit the website. These formats can still lose clicks, however; assess their value through your own queries and inquiries. [opinion]

Do not mix the share of zero-click searches, the frequency of AI answers, and changes in CTR when those answers appear. Their denominators differ. Compare your own query groups and recalculate the scenarios in section 1.6 instead of applying a universal discount factor.

Track AI-answer citations separately from visits and contracts. They may introduce readers to the author but do not establish commercial returns. Google requires no special markup for AI Overviews or AI Mode; the basic search requirements apply. [standard, 40]

1.10 When a website is not needed

The following four situations can make investment in a content site uneconomic. Check for them before starting.

There's no search demand. Search tools can provide an initial demand check, but zero reported volume is not proof of zero demand. Prospects may search for a problem without knowing the service name. Check alternative wording, client interviews, and relevant communities before ruling search out.

Referrals already fill the available capacity. If projects from contacts and past clients fill the calendar, more inquiries are unnecessary. The website then serves another purpose: reducing reliance on a small circle and supporting higher prices. Assess it against that purpose, not the number of submissions.

There are no resources for production and support.

Instead of applying a universal minimum number of hours, calculate your own capacity, choose a small but complete scope, and set a review date. Even if resources cover only a few pages, those pages can help referred clients verify the provider, without a regularly updated blog.

The work is entirely under non-disclosure agreements.

If client work cannot be shown, consider only authorized anonymized accounts, methods, or testimonials. Their effect and preparation cost need testing; no general claim that this strategy is slower is established here.

II · How the web works

2.1 From a click to a visible page

Between a link click and text appearing on screen, the browser performs a sequence of actions. Understanding that sequence matters more than knowing individual speed tricks: applying a fix to the wrong stage achieves nothing. The basic process is independent of the technology used to build the website. [60]

The sequence is as follows. The browser resolves the domain name to an address, establishes a connection, negotiates encryption, sends a request, and waits for the first response byte. It receives HTML and begins parsing it. It discovers references to stylesheets, fonts, images, and scripts and requests them. It builds a document tree and a style tree, then uses both to calculate layout and paint the first pixels.

The chain includes blocking stages. Render-blocking CSS for the current media must be received and parsed before rendering can proceed. Styles are needed to lay out elements, and a relevant rule may appear at the end of the final stylesheet. A parser-blocking script encountered in the markup pauses document parsing until it has loaded and executed, because the script can modify the document.

Connection timings are implementation-dependent. DNS may be cached; TCP and TLS setup may be avoided by connection

reuse. Do not treat an illustrative handshake count as the total cost of opening a secure connection. [46][practice].

Time to first byte includes five consecutive stages: redirects, service-worker startup if applicable, domain-name resolution, connection establishment including encryption negotiation, and the request until the first byte of the response. The recommended good threshold is 0.8 seconds at the 75th percentile. [standard, 10].

An often omitted qualification: TTFB is not a Core Web Vital. Missing its recommended threshold can be acceptable if it does not prevent the core metrics from meeting theirs [standard, 10]. Improving the number for its own sake is not the aim; use it to improve rendering performance.

An uncached HTTP redirect adds a request-response exchange and may require further connection setup. Cached or internal redirects differ. Measure the actual chain rather than assuming a fixed half-second penalty for three redirects.

HTTP/2 multiplexes streams over TCP; HTTP/3 uses QUIC. QUIC avoids TCP's cross-stream head-of-line blocking, but loss and congestion can still reduce performance. Compare protocols on the target network before claiming a speed gain.

2.2 Core Web Vitals: What exactly is measured

The three metrics describe different kinds of user frustration. Confusing them leads to ineffective optimization.

Metric	What it measures	Good threshold
LCP	time until the largest eligible element in the initial viewport is rendered	not more than 2.5 s
INP	delay in response to user action	no more than 200 ms
CLS	movement of elements in content already displayed	not more than 0.1

Check each metric separately at the 75th percentile of available field measurements, separating mobile and desktop devices. Passing all three thresholds does not mean that exactly the same 75% of visits passed every one of them. INP became a Core Web Vital in 2024. [standard, 9]

Lab tests help diagnose problems and compare changes under controlled conditions. Field data describes real user experience. Confirm lab improvements in field data when available; a field report that does not change immediately does not invalidate the lab work.

Low-traffic pages may have insufficient CrUX data. An origin-level summary is a different aggregation level, not a score automatically assigned to the page. Check which URL, group, or origin the report covers. [45]

Field measurements come from three sources, which are not interchangeable.

- CrUX aggregates eligible Chrome users' experiences. Search Console groups similar URLs using this data; a

group's status is not a separate measurement of every member page. Inspect the reported aggregation level.

- An individual URL test displays laboratory results alongside field data. Seeing their discrepancy together is more useful than reading either number in isolation.
- Collecting metrics from your own visitors provides immediate data with the breakdowns you need, but requires somewhere to store it.

CrUX API, PageSpeed Insights, and Search Console use rolling 28-day field-data windows; the CrUX BigQuery dataset is monthly. New measurements can change a rolling summary before a month passes, while older data remains until it rolls out. Lab tests and your own real-user measurements can provide faster feedback. [67]

2.3 What can go wrong with LCP

LCP can be divided into four consecutive phases. Begin diagnosis by identifying which phase consumes the time budget.

First is time to first byte. Second is resource load delay, before the browser starts fetching the resource needed for LCP, such as an image or web font. Third is resource load duration. Fourth is element render delay: the required resource is available, but the LCP element has not yet been rendered. [43]

The linked web.dev optimization guide illustrates a roughly 40%/40% split between TTFB and resource loading, with small remaining delays. These are diagnostic proportions, not required budgets. A text LCP may depend on a web-font

resource, so do not assume all resource-loading time is zero. [standard, 11]

Each phase has different causes and therefore needs different fixes.

TTFB increases with a slow server, missing response caching, geographical distance, and redirects. Image optimization does not address these causes. A common mistake is to compress images when waiting for the server consumes 1.5 seconds of a 2.5-second total.

Resource discovery is delayed when the main image is absent from the initial markup—for example, when it is a CSS background, inserted by JavaScript, or lazy-loaded. Lazy-loading the main image in the initial viewport misapplies a useful technique intended for content below the fold.

File size and request priority affect load duration. Render delay can come from blocking fonts, a heavy script occupying the main thread when the image is ready, or hydration in applications built around client-side code.

Illustrative LCP breakdown: 4.2 s = 1.6 s until the first byte + 1.1 s before the resource starts loading + 1.2 s of loading + 0.3 s of render delay. These figures illustrate diagnosis; they do not describe a verified client project. [example]

The example suggests checking resource discovery and server response before focusing on compression. A delay alone does not identify its cause. Inspect request timing, priority, caching, and rendering, and include implementation and testing time in the cost.




Figure 2.1. Illustrative LCP breakdown. The values compare hypothetical scenarios; actual changes must be measured again. [example]

With the other components unchanged, halving the resource load duration gives an LCP of 3.6 s. If TTFB falls to 0.4 s and the resource load delay falls to zero, the total becomes 1.9 s. Measure the result again on a real page: priorities and parallel downloads may affect the other components. [example]

2.4 Responsiveness and layout shifts

INP describes interaction latency over a visit. Each interaction includes input delay, event-handler execution, and the wait for the next frame. With many interactions, the calculation excludes some extreme values; it is not simply the duration of the first click. A good INP is no more than 200 ms. [standard, 31]

Long tasks on the main thread can delay input handling. Check your own code, third-party widgets, and rendering work. Even a static page can respond poorly if it loads heavy scripts.

CLS is the largest sum of layout-shift scores within a session window. Shifts are grouped while the gaps between them remain below one second and the window lasts no more than five seconds. The highest window total is used, not the sum over the page's entire lifetime. A good CLS is no more than 0.1. [standard, 30]

A layout-shift score is impact fraction—the proportion of the viewport affected—multiplied by distance fraction, the normalized movement distance. Shifts within 500 ms of qualifying user input are excluded. A later shift after a click may count; scrolling does not provide this exclusion. [standard, 30]

Look for images with no space reserved, inserted banners, font-metric changes, and asynchronous content. Reserve space and match dimensions to prevent shifts. Test the whole visit, not just the initial load.

2.5 Images

On a website where a case study contains thirty mockups, images dominate page weight. Format differences can amount to hundreds of kilobytes per page.

Compare JPEG, WebP, and AVIF using your own photographs and complex raster mockups. Savings depend on quality settings, content, and encoder; the format name alone does not establish a universal saving or page size. Compare file size, small-text legibility, and preparation time.

Slow encoding affects the workflow. Re-encoding a hundred-image gallery during every build can add several minutes. Store generated variants and regenerate them only when the source changes.

Transfer size is not the only consideration: decoding and the uncompressed image dimensions also matter. Use `srcset` and `sizes` to account for the display width and pixel density. An 800-CSS-pixel container on a DPR 2 screen may need an image roughly 1600 pixels wide. Oversized images waste resources; undersized ones lose sharpness.

Use lazy loading for content below the initial viewport, not its main image. Give that image higher priority instead. These two choices are often reversed.

2.6 Fonts

A font download can delay text that has already arrived. During a font's blocking period, the browser may keep text invisible before showing a fallback font. Two visible problems can follow: initially invisible text, and a font swap that changes the layout. [55]

The `font-display` property controls this behavior. Immediate fallback display avoids invisible text but allows a visible swap. An option that permits the browser to forgo the web font on slow connections avoids a late swap, at the cost of some visitors seeing only the fallback.

Matching fallback metrics can reduce layout shifts during a font swap, but does not guarantee identical line breaks for every string and script. Test the actual text and fonts. [56]

Subsetting retains the characters needed by the site and can reduce font transfer size. Savings depend on the font, script

coverage, format, and layout features; measure the generated file instead of assuming a 20–40 KB target. [practice]A variable font may replace several static weights, but italic styles can require a separate file. Compare the total bytes needed for the styles actually used. [44]

Preload one or two critical font faces, not every file. Preloading all six competes with the main image for bandwidth and can worsen rendering performance. Overusing the technique defeats its purpose.

2.7 Why static sites are fast, and when they are not

A static page is prepared in advance, so a request does not require rebuilding it from a database and template. TTFB still includes network latency, however. A nearby CDN and a cache hit can speed up the response; verify their presence and effect through measurement.

The statement ceases to be true in three cases, and each occurs more often than expected.

First, static pages burdened with client-side code. If a script rebuilds the interface and restores component state after the document loads, the TTFB advantage may be lost to render delay and long tasks. Final metrics can be worse than those of a conventional server delivering ready-to-use HTML.

Second, a dynamic website with correctly configured edge caching can match a static site’s TTFB: both return a cached response. The consequential differences are caching and the amount of client-side code, rather than the static or dynamic label alone.

Third, some tasks cannot be handled by static files alone. Site search, personalization, form processing, and rate limiting require client-side code or server functions, depending on the task. An architecture declared fully static acquires exceptions. That is manageable when anticipated and painful when it is not.

A static site generator is valuable when it delivers acceptable default performance without a performance specialist or constant supervision, rather than merely setting a speed record. [opinion]. For a project led by one person, this is more important than any comparison of limit values.

III · How search engines interpret meaning

3.1 Crawling, indexing, and the rendering queue


Three processes often grouped under “indexing” are separate and can fail separately. Crawling fetches the document; indexing processes it and stores it in the index; ranking selects results for a query. A page may be crawled but not indexed, indexed but not shown, or shown without receiving clicks. Diagnose the stage at which progress stopped before choosing a remedy.

Crawl budget depends on crawl demand and available capacity. The documentation does not define it as their numerical product; it describes the set of URLs Google can and wants to crawl. Low demand can limit crawling even when server capacity is available. [standard, 13]

The documentation was rewritten in July 2026, with two consequential clarifications. Every website starts with the same conservative capacity allocation, which increases over time only when demand exists and the site remains healthy. Capacity is also shared by Google’s crawlers in the described infrastructure: high demand from one reduces what remains for others [standard, 13].

A new website may be crawled slowly, but that does not rule out a technical error. Check server responses, access restrictions, internal links, and page usefulness. Do not postpone fixing an access problem while waiting for natural growth.

Google uses a rendering queue and executes JavaScript in a supported environment. Having the main text in the initial HTML reduces reliance on successful script execution. It does not guarantee indexing, however. Missing text in the initial response does not prove the page is merely waiting in the queue: execution errors or blocked resources may be responsible. [standard, 41]



-
- ① Known URL, not crawled yet.
The status alone does not establish the cause.
 - ② Crawled, not indexed yet.
Check content, duplicates, and availability.
 - ③ No text in source HTML: inspect rendering.
Scripts can add text, but may fail to load.
 - ④ Indexed, but absent for the target query.
Check relevance, competition, and technical state.

Figure 3.1. There are four stages and situations to check. First determine the stage, then explore the possible causes.

Four checks identify the stage where progress stopped. The time required depends on data availability.

First, inspect the address in the URL Inspection tool.

“Discovered—currently not indexed” means the address is

known but has not yet been crawled; the status alone does not explain why. “Crawled—currently not indexed” is not conclusive evidence of low quality either. Check the response status, indexability, duplicates, and usefulness together.

Second, compare the crawl date with recent changes and the indexing status in the console. An exact-title search can supplement diagnosis, but a missing result does not replace inspecting the URL.

Third, compare the source HTML with the rendered document in the inspection tool. Check whether the main text is visible, resources are accessible, and JavaScript errors occur. Choose a fix only after those checks.

Fourth, if the page is indexed and appears when searching for its title but not the target query, the earlier stages have been passed. Investigate ranking: alignment with search intent and the strength of signals, rather than those earlier technical stages.

Crawl budget is rarely a concern for a three-hundred-page website. It becomes relevant with tens of thousands of URLs, often generated by combinations of filters, sorting options, and parameters. Three hundred useful pages can become forty thousand faceted URLs while the useful content remains the same.

3.2 From keywords to entities

Search engines use several query and document matching mechanisms, including language models. Reducing this process to one algorithm or one word density indicator does not provide a working diagnostic model.

Two opposite conclusions are often drawn, both incorrectly: that a machine’s understanding of meaning removes the need to name things clearly, or that machine ranking requires stuffing the text with synonyms.

What has changed is the usefulness of targeting keyword density. Repeating a phrase eight times rather than three does not inherently improve a document; unnatural repetition makes it worse.

The need to name the service in the customer’s language—in the heading and opening paragraph—has not changed. Visitors scan search results and recognize their own request. “User Experience Design for Digital Environments” can lose to “UX Design” because people fail to recognize the service, not because the machine fails to understand it.

An entity is a person, organization, project, tool, or place. Stable identifiers, profile links, and explicit name variants help connect descriptions. A translated or abbreviated name does not automatically split an entity; exact character-for-character matching is not implied.

3.3 Search intent as the basis for choosing a page type

Queries reflect four kinds of intent: to learn, find a provider or buy, reach a particular website, or compare options. The page type must match the intent; that matters more than requirements imposed on the prose.

Search results help reveal the visitor’s predominant task. If they contain guides, a service page may be a poor match; if they contain providers, a general overview may disappoint.

Use that evidence to choose a format, without treating it as proof that no other outcome is possible.

A frequently skipped step can save many hours: before writing, inspect the target query's results for articles, service pages, directories, videos, and discussions. Choose the page type accordingly. An hour spent inspecting can prevent twenty hours spent on a page whose format does not fit the results.

Illustrative results inspection: a hypothetical snapshot, not a current search check. [example] For the Russian query meaning "design system," assume the first ten results are guides, documentation for other design systems, a reference article, and a video, with no service page. The working choice is an article; a commercial format would need separate justification.

For the Russian query meaning "design system development," assume the same search shows studio pages, service pages, and two case studies. The working choice is a service page covering the scope, timelines, and examples users can find in competing results.

The queries differ by one word but suggest different page types. Wording supplies the initial hypothesis; results and the buyer's task refine it.

Mixed results indicate several tasks or answer formats. Consider separate informational and commercial pages when each offers distinct value, linking them where useful to readers.

3.4 The anatomy of search results and citations

A modern results page is more than ten links. Alongside organic results it may contain generated answers, images, videos, maps, related questions, and ads. Each takes space and attention; the first organic result may sit below the midpoint of the screen.

Citations in generated answers do not follow the same pattern as organic results, as the measurements illustrate.

One of the sources is a study of 55,393 popular queries over forty days in the spring of 2026. The results relate to the selected set and observation conditions. [data, 15]

Rate of occurrence. In the selected dataset, the AI answer appeared for 13.7% of all queries and 64.7% of question queries. The difference describes this sample; it does not rule out the influence of topic or establish that question wording alone causes an answer to appear. [data, 15]




Figure 3.2. AI Overview frequency in the sample: all queries versus question queries. The chart shows an association, not an isolated causal effect of query wording. [data, 15]

The overlap with organic results is incomplete. The study’s abstract refers to domains: almost 30% of cited domains were absent from the first-page results shown at the same time. Domains and individual URLs are different units of comparison. [data, 15]

Credibility. The automated verification pipeline classified 11.0% of evaluated claims as inconsistent with retrieved sources, including omission, contradiction, or source disagreement. This is not a direct measurement of factual falsity; incomplete source retrieval can affect the classification [data, 15]. Source quality and the accuracy of the generated account were largely independent: citing an authoritative page did not guarantee a faithful claim.

Advertising. More than half of the sampled cited pages carried advertising. This indicates exposure to lost visits, not a measured amount of publisher revenue lost because of AI answers. [data, 15].

The figures refer to the May 13, 2026 arXiv version and its selected query sample, not all searches. The arXiv record says Under Review, while the author’s publication list names IMC 2026. Keep the cited version explicit rather than treating that metadata as a current peer-review verdict. [66]

When comparing studies, distinguish domains from URLs, the depth of organic results examined, and the query sample. Without aligning these definitions, differing percentages cannot be attributed to an algorithm change.

Some widely circulated figures are deliberately excluded from the calculations: claims that citations are 4.8 times more likely with fifteen recognized entities, 2.3 times more frequent with structured data, or 156% higher with text,

images, and video together. Without published methods, sample sizes, and controls for confounders, these multipliers cannot support the calculations [opinion].

For the reader, verifiable statements with conditions of applicability are useful. They also reduce the ambiguity of possible retelling, but do not guarantee citation or accuracy. The study does not prove that general reasoning cannot be quoted or that any fragment taken out of context will necessarily be distorted.

3.5 E-E-A-T: What it is and what it is not

E-E-A-T stands for experience, expertise, authoritativeness, and trustworthiness. It is often described as a ranking factor. That is incorrect, and understanding its actual role changes the approach.

It is a framework for human quality raters who assess results using guidelines. Google states that raters do not determine page rankings and their ratings are not used directly in ranking algorithms [standard, 16]. Their assessments provide feedback on whether algorithm changes work as intended. A page does not receive a ranking score on this scale.

Google's distinction is careful: E-E-A-T is not a standalone ranking factor, but systems can use signals that identify content with those qualities [standard, 16]. Such signals carry more weight in topics that can substantially affect health, financial well-being, or safety.

The practical implication is to make evidence verifiable, rather than trying to "implement E-E-A-T." Give the author's name, a factual biography, publication and update dates,

links supporting claims, testimonials with names, employers and roles, and real-world contact details.

Conversely, pages can appear unreliable when About names nobody, articles have no author, testimonials provide only first names, dates are absent, or claims cite unnamed “research.”

3.6 Duplicate content, cannibalization, and thin content

Duplicate content. Several URLs may serve identical or similar content. Search engines choose a canonical version; `rel="canonical"` is a signal, not a command. Duplicate content alone does not prove that a site is being penalized. [standard, 35]

Cannibalization. Keyword cannibalization describes overlapping pages that obstruct the chosen search strategy. Two URLs appearing for one query, or a URL changing in a report, does not by itself prove harm. Check each page’s intent, clicks and inquiries, internal links, and the selected canonical URL.

Merge pages when they serve the same task and their useful content can form one stronger page. Keep pages serving different tasks and clarify their purpose. You cannot promise that their rankings will add together or that merging will produce a particular gain.

Thin content. Usefulness is not determined by word count. A short answer may suffice, while a long retelling may not help. Google specifies no preferred length. Mass-producing pages to manipulate results is assessed under spam policies, not by a repetition percentage alone. [standard, 16, 26]

Worked example C. Of the 60 pages in the hypothetical launch, 18 have similar purposes. That is a list to investigate, not an instruction to delete 18 URLs. Compare queries, content, and existing links before deciding whether to retain, revise, or merge each page. [example]

3.7 Two search algorithms

Collect separate Google and Yandex visibility reports and inspect results in the target region. The same site can rank differently, so success in one engine does not validate the other.

Claims such as “links matter more here, behavior matters more there” provide no planning coefficients without reproducible measurements. Industry accounts of signal weights must not be presented as published descriptions of the algorithm.

Start with site availability, query relevance, a clear service description, verifiable contact details, and regional service conditions. Assess changes separately for each search engine, accounting for language, geography, device, and audience composition. Translating a website does not itself create demand in another market.

IV · Architecture and URLs

4.1 Taxonomy and facets

Taxonomy and facets serve different purposes and can coexist. The hierarchy defines sections; independent attributes allow content to be filtered by several criteria.

A taxonomy gives each item a place in a hierarchy: services go in Services, case studies in Case Studies. Facets are independent descriptive dimensions: the same case study can be filtered by service, industry, year, and tool.

With exactly one choice from each of five services, eight industries, and six tools, there are 240 combinations. Optional or multiple selections change that count. Check which combinations actually have useful content rather than assuming only ten to fifteen do.

Faceted navigation is a common source of excessive crawling reported by site owners. Following a few rules can prevent many such problems [standard, 18]. There are three options: block crawling of faceted URLs that do not need to appear in search; if they do need to appear, use a consistent filter order and return a not-found status for empty combinations; or filter in the browser without generating separate URLs. [standard, 18].

Procedure. List the facet dimensions and values in a matrix. Mark each combination for demand above your chosen threshold and for available material—at least two or three case studies or substantive paragraphs in this proposed workflow. Make only combinations meeting both conditions eligible for indexing, with genuine pages and distinct text. Keep the rest as filters without separate URLs or block crawling. Review the matrix every six months as material accumulates.

4.2 URL as a contract

A page URL is a commitment to others: people link to it, bookmark it, and cite it. Changing it requires managing redirects and search-engine recrawling. An uncached HTTP redirect adds another request and response; the delay depends on caching and connection reuse.

Define the URL scheme before publishing. Revise it when the benefits justify migration and redirects are ready; unnecessary changes create unnecessary work.

Decisions the URL scheme must cover.

Language and transliteration. Cyrillic characters are technically valid in URLs, but copying them can produce an unreadable encoded sequence that is awkward in messages and documents. Latin transliteration is easier to handle. Mixing translated and transliterated names in one scheme causes recurring errors.

Singular or plural. URL patterns such as `/service/ux-design` and `/services/ux-design` do not differ in ranking simply because of this choice. What matters is the cost of changing

the convention after launch. Choose one and apply it consistently.

Depth. URL nesting and click depth are different. The address /services/design/audit/ may be one click from the homepage. Choose a clear, stable URL scheme and assess navigation depth through actual links.

Changeable information in URLs. Keep changeable details out of URLs: years, status, prices, client names, and current job titles. The URL /blog/trendy-2026 It will become obsolete in a year, and you will have to choose between a redirect and an outdated address on the current page.


Procedure. Before publishing the first page, record six conventions: URL language, translation or transliteration of names, depth, singular or plural section names, word separators, and letter case. This short document preserves consistency when someone else creates pages a year later.

4.3 Hubs and clusters

Here, topical coverage and internal links are a content-planning approach, not a promised ranking score. Google does describe a topic authority system for news-related queries; that does not establish a universal cluster formula for service websites. [42]


A hub can link to useful subtopic pages, which link back and to relevant neighbours. This supports navigation and discovery; the diagram does not calculate PageRank or guarantee a ranking increase.

Six pages on one topic



No internal links;
no internal discovery path

The same pages as a cluster



Solid: hub link. Dashed: related pages.

Figure 4.1. The same six pages: independent on the left, linked as a cluster on the right. The difference is their connections and links to the service page, not their content.

There is no established minimum page count for a cluster. Start with a useful overview and a few standalone pieces. Add a page when it has a distinct purpose, not to meet a numerical target. [opinion]

Assembling procedure. Choose a topic with a service page to support. Gather subtopics from search suggestions, related questions, and client questions, selecting those with informational intent. Designate a hub. Before writing each subtopic page, identify two related pages and the service page it will link to. Pages may be published one at a time, but plan the full link map at the outset; otherwise internal linking becomes an unfinished retrofit.

4.4 Internal linking as a flow

Internal links help visitors continue their task and crawlers discover pages and understand relationships. Google recommends linking to important pages from at least one other site page; it sets no ideal link count. [standard, 32]

Equal division of link weight can illustrate a simplified model, but cannot calculate a menu's actual search impact. Cutting forty links to sixteen does not prove a 2.5-fold signal gain. Improve navigation for clearer choices and useful connections, then measure results.

Click depth shows how easily readers can find content within the site; it is not the number of folders in a URL. Link important services and case studies from clear sections and related content without turning the menu into a list of every page.

The orphan page has no incoming internal links. It can be discovered through a sitemap or an external link, but it is difficult for a website visitor to find it in the usual way. The word "orphan" does not mean the absence of any search signals.

Audit procedure. Compare the intended URL inventory with a crawl of internal links. Find orphan pages, broken links, and unnecessarily long paths. Add links where useful to readers, then crawl again. Use link counts to identify overloaded sections, not to trigger automatic removal.

Worked example A. The service page links to two supporting case studies. Each links back to the service and provides an easy way to make contact. This structure works

at a small scale; a blog can follow when there are questions to answer and resources to answer them. [example]

4.5 Pagination, filters, parameters

Each pagination page needs its own URL and a crawlable link. Pages showing different items normally use self-referencing canonicals, rather than identifying page one as their canonical. Incorrect canonicalization complicates processing, but does not prove that every listed item disappears: other discovery paths may exist. [standard, 36]

Infinite scrolling needs a crawlable pagination equivalent. Updating the browser address is not enough: each segment URL must open the relevant content directly and be connected by ordinary links. Google generally does not click buttons or perform user actions to load more content. [standard, 36]

Check whether sorting parameters actually create duplicates. Canonical tags suit equivalent pages; a useful, distinct selection should not automatically be treated as a duplicate. Separately choose a robots.txt and URL-management strategy for faceted crawling. Blocking crawling does not instruct search engines to remove an already known URL. [standard, 18, 35]

4.6 Redirect map

Permanent redirects signal a move and support consolidation at the destination. They do not guarantee identical rankings or that every signal transfers. Use a relevant destination and minimize chains. [47]

A redirect chain passes through intermediate URLs. Uncached HTTP hops add requests and responses; caching

can change the delay. Crawlers limit the number they follow. This workflow replaces chains longer than two hops with a direct redirect; two is a maintenance convention, not a universal crawler limit.

Content mismatch occurs when a page redirects to a destination that does not serve the same intent, often the homepage. Bulk homepage redirects can be treated as missing pages, losing their signals. Redirect to an equivalent destination; if none exists, return a not-found status.

Migration procedure. Build the old-URL inventory from four incomplete sources: the sitemap, the webmaster dashboard’s performance report for the longest available period, server logs, and analytics. Merge and deduplicate them. Match old URLs to new ones manually; name-similarity matching can produce errors across a substantial share. Mark URLs without a replacement separately and decide which should return not found.

Test mappings before switching, then monitor until migration problems are resolved. Thirty to sixty days is an initial review window, not a deadline for recovery. Google recommends keeping redirects generally for at least a year; useful redirects may remain longer.

4.7 Structured data markup

Structured data describes content in machine-readable form. Some types make a page eligible for enhanced Google search appearances. Valid markup guarantees neither that appearance nor higher rankings.

FAQ rich results stopped appearing on 7 May 2026, according to the official changelog entry dated 8 May; the

dedicated documentation was later removed. FAQPage remains a vocabulary term, but that does not establish a guaranteed effect on citations. [standard, 19]

Separate three checks: validity under Schema.org, requirements for the particular Google feature, and actual appearance in results. Google lists no required properties for Article but recommends useful fields. Organization can provide organization and logo details; saying it has no search appearance is too categorical. [standard, 37, 38]

Procedure. Choose a type that matches the content and generate its fields from the same record as the page. Check Schema Markup Validator and applicable Rich Results Test requirements, then verify that the published URL is accessible. No enhanced appearance after an arbitrary two weeks does not itself prove an error. Appendix C distinguishes these checks.

4.8 What breaks when the site reaches 300 pages

Manual maintenance of relationships becomes costlier as a site grows. When to automate depends on content types, editors, and update frequency, not a universal three-hundred-page threshold.

As records accumulate, check for orphan pages, overlapping topics, and overloaded navigation. These are maintenance risks; page count alone does not prove that crawling will stop or take months.

Derive some relationships from data fields such as service, industry, and topic. Use them to generate related-content sections, then have an editor check relevance. Manually

placed links complement this system; automation does not replace editorial judgment.

Plan relationship fields early when reuse is likely. Retrofitting existing records can require substantial review; no universal one-minute entry time or page-count threshold applies.

V • Content as a database

5.1 Why content is a database, even without a database

A static-file website may have no database engine, but its content still forms a database. Ignoring that structure creates manual work that grows with the page count.

The distinction is simple: a page is a template applied to a record. Separate data from presentation, and changing every case study's appearance requires one template edit. Mix them, and every file needs editing.

Test the model by asking how to add an industry field to every case study and show it on cards. A shared template needs one change, but existing records still need values. A schema does not eliminate content migration; it makes the migration rules explicit.

The unit of work is a record. One case-study record can generate its page, a listing card, a related-content card on a service page, a sitemap entry, structured data, a feed item, and an entry in the site's search index: seven outputs from one description.

5.2 Entities, fields, and relationships

Model the domain as a stable set of entities: person, service, case study, article, testimonial, industry, tool, category, and tag. This set is not universal, but changes little across service-provider websites.

More consequential than the entities themselves is relationship cardinality: how many objects on each side can relate to one on the other. A costly mistake is modeling a naturally many-to-many relationship as one-to-many.

A case study can cover several services, such as research, interface design, and brand identity. A service can be supported by several case studies. That is many-to-many. If a case study allows only one service, a project marked interface design will not appear under branding even if it included branding. Correcting fifty accumulated records means reviewing every one.

Define which side owns each relationship. Store it once and derive the reverse through a query. If case studies list their services, generate a service page's case-study list from those records. A second manually maintained list will eventually disagree with the first.

Procedure. List the entities. For each pair, ask both questions: how many Bs can one A have, and how many As can one B have? Choose the owning side, usually where the list is shorter and changes less often. Record the rationale so it survives beyond the next six months.

5.3 Schema and validation

A schema defines fields, types, and required values. Validate it during the site build and fail the build on violations, preventing incomplete case studies from being published.

Require fields without which a page is incomplete: title, URL, description, and date. Keep others optional so creating content does not become an obstructive questionnaire.

Classification fields—industry, service, tool, category—need controlled values, not free text. Exact matching treats “SaaS,” “Saas,” “SaaS platform,” and “subscription service” as different values, fragmenting the results. Reserve free text for genuinely individual values, such as client and project names.

Procedure. Define enumerations before content entry and store them centrally. A new value should be an intentional list edit, not a typo. Every six months, review values used only once: they may indicate errors or too little material to support a category.

5.4 A single source of truth

Store each fact once. Insert the service name, publication date, author name, and page URL from their authoritative records wherever needed.

The saving is concrete. Renaming a service then means editing one field. Without this rule, menus, headings, related-content cards, sitemaps, structured data, other pages, and internal links must all be checked: potentially forty edits instead of one.

Separate fields are appropriate for separate purposes. The page heading, search title, and card label may differ. Define defaults, permitted overrides, and responsibility for consistency. A single source of truth applies to one fact, not every piece of wording about a subject.

A warning sign: changing one fact requires editing more than one content file, excluding templates.

5.5 Migrations and schema versions

Schema changes fall into three categories with different costs.

Adding an optional field leaves existing records valid and permits gradual completion. Making it required invalidates records without a value, which must be filled before the next successful build. Renaming a field or changing its type requires migration of the affected records.

Introduce a field as optional, populate it, then make it required if needed. Reversing that order can leave the site unable to build for days.

Migration procedure. Use a migration script: manually editing a hundred files can introduce errors that remain unnoticed. Run the script on a copy, compare changes line by line, and perform a full build as a check. Log the schema version and a short description so the field's history remains understandable a year later.

Treat deletion separately. It is safer to retain a disused field in records temporarily and remove it later than to delete it immediately. Discovering a remaining dependency is easier before the data is lost.

Worked example A: the first version's data

Identifiers link one service and two case studies. Each case study stores a service list, and the reverse list is generated. Adding a case study updates the cards without a second manual list. Result fields permit qualitative evidence and an explicit statement that business metrics are unavailable.
[example]

VI • The case study as a genre

6.1 Why a gallery of mockups does not sell

Mockups show appearance, while a case study can also explain criteria, decisions, and the author's contribution. Clients may assess some aspects directly and need further evidence or expert help for others; neither format guarantees correct judgment.

A prospective buyer wants three answers a gallery cannot give: is this task like mine, did the provider actually solve it, and what did they contribute themselves? A case study explains how a problem was solved. Mockups support that explanation rather than replacing it.

6.2 Why constraints matter more than the brief

The task alone rarely makes a story: build a website, an app, or a brand identity. The constraint makes it interesting—the obstacle that ruled out the obvious approach.

Constraints may concern resources (time, budget, one person instead of a team), technology (legacy systems, inherited code, compatibility), the domain (regulation, industry standards, user needs), people (a client committed to a

solution, no decision-maker), or information (missing data, no user access, no analytics).

Constraints explain choices. Even an open-ended project has alternatives; show why you chose the particular solution.

A useful narrative sequence is: starting situation → constraint → options considered → decision → trade-offs → result.

Most case studies are weakest on the trade-offs. Every decision costs something: a feature omitted, a longer schedule, or more maintenance. Naming that cost shows that the author understood the consequences. Silence can suggest they did not.

Showing a rejected option also makes a case study convincing: it demonstrates that alternatives existed. A story in which the correct solution appears instantly can sound invented after the event, because work rarely unfolds that way.

6.3 How to get numbers and what to do without them

Describe outcomes with verifiable evidence: figures, observable changes, client decisions, or qualitative research. Numbers are not the only valid evidence. Do not invent a metric merely to make a case study look convincing.

Ways to get numbers

These are collection options, not a ranking of statistical reliability. Reliability depends on measurement design, data quality, and comparability.

Plan measurement before work starts: agree on the metric, baseline, period, and responsible party. This improves the

chance of comparable data but guarantees neither access nor a measurable effect.

Ask two or three months after launch. Earlier, data may be unavailable; later, interest fades. Keep the email short, ask about one or two specific metrics, and explain the client's benefit: the case study will mention their company.

Agree on analytics access for a specific purpose and period. Separately agree on what can be published: permission to access data is not permission to publish it.

When figures will not be available

Some properties can be measured independently without access to client analytics: the product's own characteristics before and after the work.

A form went from twenty-four fields to nine. Checkout went from seven steps to three. Homepage load time fell from 4.2 seconds to 1.9. Three product-finding methods became one that worked. Twelve scattered documentation files became one reference guide.

Authorized before-and-after materials can substantiate those changes. Fewer steps demonstrate an interface change, but alone prove neither better usability nor higher revenue.

What can't we do?

"Conversion increased by 40%" is not a substantiated result without the period, measurement method, and other changes made at the time. The claim cannot be checked and may fall apart after one question on a call.

With the prohibition of disclosure

If the contract permits an anonymized account, use the agreed industry description, ranges, and diagrams. Removing the client's name does not itself remove publication restrictions. Confirm which materials may be shown.

6.4 The boundary between presentation and deception

Four presentation tactics can undermine trust when checked, even without an explicit falsehood.

Presenting others' work as your own. You designed the interface while others handled strategy and research, but write "we conducted the research." One question about research details can expose the ambiguity.

Presenting team output as solo work. Omitting the team from an independent provider's case study implies sole authorship. One line can clarify what you did and what others did.

Presenting a concept as a launched product. A common tactic is to describe an unlaunched design in the present tense and attach results it never achieved. A live website address or app-store listing can help distinguish implementation from a concept. Concept work is valid when clearly labeled.

Numbers without causation. "Revenue grew by a third after the redesign," while advertising launched, the season changed, and the product range expanded. The sequence can be true without establishing the redesign as the cause.

An honest result states what was done, what changed, and what else could have influenced it. The last part strengthens the account by showing that the author distinguishes coincidence from causation.

A misleading case study may survive a search but fail at the first follow-up question on a call, costing the deal—more than the exaggeration gained.

6.5 Examining the structure, section by section

Headline. Name the work and outcome, rather than only the project. “Case study: Project Name” says little. A useful pattern is what was done, for whom, and with what result. Weak: Case study: Ayurveda Reader. Strong: A 4,000-article reference site: reducing the path to the right page from seven steps to two.

Lead: the first three or four sentences. Help readers decide whether to continue: identify the project, constraint, and outcome. Even readers who stop here should leave with the main point. Weak: Ayurveda Reader is an educational platform. We designed a user-friendly interface for it.. Strong: An Ayurveda reference site with four thousand articles and a search function that failed to find them. The section structure had to stay intact, with two months available and no developer on the client’s side. The path to the required article fell from seven steps to two..

Context and client. Identify the client, their situation, and why the task arose then. The reason for the task matters more than a company description because it explains the constraints.

Constraint. State it in a distinct section: what was not possible, and why.

Process. Describe what you did and learned, rather than listing methodology stages. “We conducted user research” says little. Give the participant count, method, and unexpected findings. Weak: We conducted UX research and created a customer journey map (CJM). Strong: We spoke with nine readers. Eight searched by symptom rather than section name, a route the reference site’s structure did not support..

Decisions. Explain each decision separately: the options, the choice, the reason, and the trade-offs.

Result. Numbers or measurable properties of the product, indicating the method of measurement.

Conclusions. What the author learned and would do differently. This is the place for personal judgment and the section peers read most.

Next step. Link to the relevant service. Without that link, the path ends at the moment of greatest interest—the failure described in Part I.

6.6 Presenting visuals

A useful filter for images: each should illustrate the adjacent paragraph’s claim. Otherwise it is decoration that adds page weight and can hurt the metrics in Part II.

Useful evidence includes matched before-and-after views, a flow or structure diagram, an interface close-up with legible text, or a rejected option beside the chosen one.

Weak evidence includes thirty uncaptioned screens, perspective mockups that make interfaces too small to read, or generic process photos of sticky notes, whiteboards, and laptops.

Add a caption so readers do not have to guess what the image demonstrates.

Gallery size affects data usage but does not determine LCP by itself. Load the main above-the-fold resource promptly and defer offscreen images. Test the initial viewport and scrolling separately; a five-megabyte page does not automatically mean poor LCP.

6.7 One case study, several channels

As a publishing strategy, keep an authoritative copy under your control and use platforms for distribution. Check platform terms and export options; this is a portability choice, not a universal rule about ownership or search performance.

Adapt the emphasis, not just the length. Lead with an image on visual platforms, followed by a caption and link. On a professional network, lead with a finding or surprising observation; the image is secondary and the link comes last. Keep the full analysis on your website.

Posting identical full text creates duplicates. As discussed in Part III, an established platform may be selected over a new website. Publish adapted summaries externally.

Worked example A: evidence without invented conversion gains

In the hypothetical case, the professional reduces an inquiry process from seven steps to four and shows both flows. There are no sales data, so the result is stated as a flow change. Improved usability remains a hypothesis for user testing. The text distinguishes the author's role, the developer's work, and the absence of measured business impact. [example]

VII • The service page

7.1 What decision does a person make?

A service page participates in shortlisting. The visitor decides whether to retain or reject the provider. This changes the page's job: removing reasons to leave matters as much as persuasion.

Use the following three questions as a starting point for the page. Their order is an editorial proposal to test with the target audience.

Do they provide exactly this service? The heading and first screen answer within seconds. "Digital product design" does not clearly answer a search for logo design, even if the provider designs logos.

Have they worked with similar clients? Visitors check the case studies for industry and scale. A restaurant owner looks for restaurant work, not merely a branding label.

Roughly how much? Indicative pricing helps assess budget fit. A missing price section alone cannot establish why any particular visitor left.

The roles in Part I explain these needs. Decision-makers want relevant cases and a price range. Shortlisters need scope, timelines, and stages to put into a comparison table. Provide a skimmable summary and detailed information for extraction.

7.2 Price: publish or not

Nielsen Norman Group's B2B usability research shows the value of accessible pricing when selecting a supplier. It supports explaining indicative costs, not a promise of equal sales growth for every provider. [data, 23]

An often missed point is internal approval. The person researching your website may not control the budget and needs a figure to take to a decision-maker. Without one, they cannot complete their task and may choose a provider who supplies it.

The argument against publishing prices goes like this: prices depend on scope, cheaper competitors will use them against us, and high prices will deter clients. Scope dependence is real, but the conclusion does not follow. Deterring a client whose budget cannot cover the work saves both sides a call. A page without prices may attract more inquiries from people who could never afford the service [opinion].

Strategies and their implications

A fixed price suits standardized work with predictable scope. It filters unsuitable budgets and limits unnecessary conversations, but needs revision when costs change.

A starting price offers a lower bound while leaving room for estimation. The risk is that readers mistake it for the final price. Set a realistic starting point, not an artificially attractive one.

Scope-based tiers describe three levels and their inclusions. For project services, they show both the approximate cost and what drives it.

A parameter-based calculator lets visitors specify scope and receive an estimate. It requires development and suits standardized services with few variables.

7.3 Page sections: structure and order

The order of the blocks is determined by the order of the questions in section 7.1, not the usual template.

First screen: name the service in the customer's language, identify the audience, state the key distinction in one line, and provide a button. Leave mission statements and methodology for elsewhere.

Next, show recognizable situations the service addresses. Then list deliverables: prototypes of all screens, layouts for two resolutions, a component file, and developer specifications. "Research, planning, design" lists stages without telling clients what they will receive.

Then give pricing, timelines, and the step-by-step process, including the client's responsibilities. Follow with three relevant case studies, testimonials, questions and answers, and the form.

A common ordering error is putting process before scope and price. A lengthy methodological account can lose readers who have not yet decided whether the service fits.

7.4 Questions and answers as a format

The FAQ answers objections that might otherwise arise on a call—or remain unspoken while the visitor leaves. It is not decorative space or a place to insert keywords.

Use actual client correspondence. In this proposed workflow, questions repeated at least three times enter the FAQ in the client's own wording.

Answer directly in the first sentence, then explain. Use enough words to answer fully, with no fixed count. "It depends on many factors" is not a sufficient opening answer.

Typical topics: price and its drivers, timing, client inputs, dissatisfaction with the result, ownership of source files, post-delivery support, contracts, and payment methods.

Following the removal of FAQ rich results in May 2026, discussed in Part IV, the section still serves readers and provides facts that can be extracted for citation. Both call for direct, specific answers.

7.5 Forms

Each form field takes effort, but its effect depends on the question, visitor readiness, and business objective. Do not deduct a fixed percentage of conversion for each field based on summaries of unrelated studies.

For initial contact, a contact method and short task description often suffice. Add name, company, budget, and timing when useful for processing; some can be optional. If budget matters for qualification, explain the ranges and allow "not yet decided." Test this approach rather than treating it as universal. [opinion]

First check delivery: completing the form on a phone, error messages, successful server response, and receipt by the intended recipient. A button-click event does not prove delivery.

Then inspect the whole funnel: form visits, starts, successful submissions, qualified inquiries, and contracts. A shorter form may increase submissions and unsuitable inquiries at the same time. Include correspondence time in the cost of the decision.

Worked example B. Of 1,000 search visits, 120 reach service pages. The log records four successful submissions, but only one qualified inquiry. Check delivery and why the other three were unsuitable before reviewing the offer and form. One overall conversion rate cannot locate the constraint. [example]

7.6 Calls to action

Button text should explain what follows the click. “Send” names the mechanism; “Discuss your project” names the next step; “Get a timeline and cost estimate” names the outcome.

Match the invitation to the reader’s readiness. A service page can invite a discussion; an informational article may suit a low-commitment next step. An “Order” button can ask too much of a reader still at the start of the journey.

Use the same call to action two or three times along the service page: after scope, after case studies, and at the end. Different labels for one action can imply different offers.

7.7 Industry pages

An industry page serves distinct demand only when that demand exists and you have useful material to meet it. Check both before writing.

Look for a query such as “design for [industry]” with its own exact-match search volume and results distinct from the

general query. The Part IX method provides a fifteen-minute initial check.

Useful material requires at least one relevant case study and knowledge of the industry's tasks, constraints, rules, and language. Otherwise the page is generic wording with an industry name inserted.

If industry pages differ only by name, check whether each adds distinct value. Reusing a template is not itself a violation; mass publishing to manipulate rankings may fall under spam policies. [standard, 26]

Condition of publication. The page needs a distinct buyer task supported by demand research, and substantive reasons for an industry-specific offer. Substituting another industry name helps expose generic wording, but no universal "40% uniqueness" threshold exists. Check examples, constraints, and scope.

VIII • Editing

8.1 The voice of a practitioner and the voice of a copywriter

The distinction lies in the source of the writing. Practitioners draw on work they have done; generic copy draws on conventions about how the subject is usually described. Stylistic symptoms follow from that difference and cannot be fixed by polishing style alone.

Writing from experience gives figures with conditions: “1.5 seconds on a mobile connection,” not merely “fast.” It names failures, trade-offs, and unexpected details that general reasoning alone would not reveal. Someone who has tried an approach twenty times knows where it breaks down.

Generic writing shows the reverse: evaluative adjectives instead of facts, universally applicable advice, no failures, and unbroken certainty.

A useful test for excess wording is whether the paragraph could have been written without doing the work. If so, it may add nothing readers have not seen in ten other articles.

These are editorial recommendations. They do not require removing every explanation of familiar concepts: introductory material is useful when the audience needs it. The weak and strong examples illustrate writing techniques; their figures are not client results. [example]

8.2 A catalogue of clichés

The following list pairs each cliché with its weakness and a replacement strategy. Maintain it as a controlled, expandable checklist and search for these phrases while editing.

Phrase	What's wrong?	Replacement
"in the modern world"	adds no meaningful date; could apply in any year	Delete; if you need a date, name the year and event
"In the age of digitalization"	the same problem, with added grandiosity	remove
"Integrated approach"	does not say what is included	list the included work
"individual approach"	everyone says it, so it does not distinguish you	What exactly changes from client to client?
"quality design"	substitutes praise for a verifiable attribute	name a verifiable attribute
"increase efficiency"	efficiency of what, measured how?	name the metric and value
"optimize"	could mean almost any action	state the specific action
"solution for your tasks"	contains no information	name the tasks
"as soon as possible"	gives no number	state the duration in days
"rich experience"	lack of measurement	years, projects, industries

Phrase	What's wrong?	Replacement
"unique"	almost always untrue	explain the difference from the closest alternative
"broad range of services"	requires a listing that does not exist	list
"not just X, but Y."	rhetoric instead of a claim	state Y directly
"It's important to understand," "it's worth noting."	delays the main point	remove the preamble; keep the point
"Let's figure it out."	simulates a conversation	remove
"In conclusion, I would like to mention"	introduces a recap	remove the whole paragraph
"this article is intended"	talks about the text, not the subject	remove
"plays an important role"	It doesn't tell you how it affects	state what is affected and by how much
"is a key factor"	same	same
"Fast, high-quality, reliable."	a three-part slogan that weakens every claim	choose one claim and support it with a figure
"to date"	filler	remove
"as we know"	if readers know it, why tell them?	remove
"certainly," "undoubtedly."	emphasis in place of evidence	provide evidence

Also check habits, not just words: rhetorical questions used as transitions, openings announcing a chapter's contents, conclusions merely repeating it, and "just" used to mean easy. The last can dismiss a difficulty readers have struggled with.

8.3 Three rules for replacing clichés

Fact instead of evaluation. An evaluation gives the author's judgment; a fact lets readers form their own. Fast site. → the first screen loads in 1.2 seconds on a mobile connection. An easy-to-use form → three fields, sending without reloading the page. Experienced specialist → nine years, forty projects, twelve of them in the financial industry.

Use a number instead of a vague quantity. The words "many", "significantly", "substantially", "multiple times" are not translated into a testable statement. Conversion has increased significantly → the conversion rate rose from 1.1% to 1.9% over four months. A lot of customers in this industry. → Seven projects in the industry.

Use verbs instead of nominalizations. Nominalized constructions hide the actor and lengthen the sentence. Implementation of image optimization → compressed images. The structure has been refined. → I restructured it. To enable comparisons → so that we can compare.

The fourth rule follows: prefer active voice to passive voice. Work was carried out The sentence hides who did the work. That matters in a case study: readers need to distinguish your contribution from everyone else's.

8.4 Editing a machine-generated draft

Editing AI-generated text requires attention to recurring defects. Both people and models can write fluent but unsupported claims; authorship is no substitute for verification.

Characteristic defects

False symmetry: three points because three sounds balanced, not because there are exactly three. Ask whether a fourth exists and why it is missing.

Generalization instead of an example: “Many companies face this,” without naming any. An example must be specific to be useful.

Plausible, unsourced figures. A model may supply “about 70%,” which looks credible because it falls in a reasonable range. Verify or remove every such figure.

Recap paragraphs. A section ends by repeating itself. Remove an empty recap in full.

Generic caveats: “Context matters” or “The result depends on many factors.” These may be true but supply no specific information.

Monotonous cadence. Every paragraph has the same length and emphasis, with neither a sharp short sentence nor an extended explanation.

Correction procedure

First, remove repetition and empty phrases; no fixed reduction percentage is required. Second, check figures and factual claims. Third, add verifiable examples, applicability limits, and trade-offs. Fourth, improve transitions and rhythm without mechanically breaking the reasoning apart.

Rule of stopping

Treat an AI draft as raw material. If four editing passes leave no claim the author can defend in conversation, nothing is ready to publish; further polishing cannot supply missing substance.

8.5 Structure of long text

The suggested pattern is a subheading every 400–600 words, phrased as a claim rather than a topic. “Publish your prices” tells scanning readers more than “About prices.”

Each paragraph develops one idea. Check its length in the final layout, since line count depends on width, font, and device. In this edition, eight lines prompt closer inspection, not automatic splitting.

Start a section with substance. “Let’s consider how to present prices” can be deleted without losing information, so remove it.

Use lists for parallel items that do not depend on one another in an argument. If the connection is “therefore,” preserve the connected reasoning in paragraphs. Conversely, a paragraph that merely enumerates seven disparate items may be clearer as a list.

Let the question determine length. Requiring three thousand words can turn a complete 1,200-word answer into 1,800 extra words of filler. Finish when the question is answered.

8.6 Text typography

The following conventions describe the Russian original; typography must be adapted to each translation's language.

Dash and hyphen. In Russian, a hyphen joins parts of a compound, as in «дизайн-система». A sentence dash is spaced, with a non-breaking space before it to prevent it starting a line.

Quotes. The Russian original uses «...» for outer quotation marks and „...“ for nested quotations, rather than straight double quotes in typeset prose.

The Russian letter ё. Choose a policy for the entire book. The Russian original consistently uses ё, helping search and avoiding ambiguity.

Non-breaking spaces. The Russian original uses non-breaking spaces after short prepositions and conjunctions, between numbers and units, within abbreviations such as «т. е.», and before sentence dashes. This prevents unwanted line breaks.

Numbers. The Russian original separates thousands with a thin space: 14 700. It uses a dash without spaces for ranges: 3–5, no space before the percentage sign: 27%, and a non-breaking space before units: 2.5 s.

Periods and quotation marks. In the Russian convention described here, a period follows the closing quotation mark when only part of a sentence is quoted.

Ellipsis — One sign, not three dots.

Justified Russian text needs working hyphenation to avoid uneven spacing. Enable a Russian hyphenation dictionary and check the rendered result; if it does not work, use left alignment.

8.7 Editor's checklist

Twelve checks for the finished text.

1. Search the clichés catalogue - zero matches.
2. Every figure has a source or an explicit note that no source is available.
3. Each evaluative adjective is either replaced by fact or justified.
4. The first sentence of each section carries the content, not the announcement.
5. No section ends with a retelling of the section.
6. Each list is checked for the uniformity of the elements.
7. Verbal nouns are replaced by verbs where possible.
8. Use active voice where identifying the actor matters.
9. Long paragraphs are checked for readability and connectivity in the finished layout.
10. There's at least one place where it says the approach doesn't work.
11. There's at least one detail you can't think of without doing the work.
12. Check dashes, quotation marks, non-breaking spaces, and number grouping.

Checks ten and eleven help identify unsupported generalities. Their relative effectiveness has not been measured here. [opinion].

IX · Practical keyword research

9.1 What a keyword set is, and what it is not

A keyword map connects demand to site structure. Its planning unit is a cluster: queries with one intent, served by one page of the appropriate type.

It is neither a list of words to insert into prose nor a five-thousand-row export. Without clustering and priorities, the data cannot become a work plan, regardless of its size.

A useful test: can the map directly produce a page list with a type, URL, and heading for each entry? If not, the work is unfinished.

9.2 Search-volume sources and their errors

Every source estimates demand with its own systematic biases. Understanding those limitations matters more than the figure alone.

Search engine query statistics. In Wordstat, quotation marks fix word count, ! fixes a word's form, and [] fixes order. Record the exact operators, region, period, and report tab. Operator support differs between tabs; a history chart is not

automatically the same narrowly defined count. [standard, 20].

Broad-match volume helps explore a topic but does not predict impressions for one page. Narrowing operators refine the demand definition. Record region, period, and exact query syntax with every figure. Volume differences do not translate directly into proportional visit differences.

Report on the effectiveness of your own site. This report contains observed impressions and clicks rather than a demand model. It is limited to queries where the site already appears, so it is most useful for refining existing visibility rather than mapping entirely new directions. [62]

Overseas assessment services. Keyword services combine different datasets and models. Accuracy varies by tool, country, query, and period; a universal English-versus-Russian accuracy ranking is not established here. Compare estimates with first-party observations where possible.

Seasonality. One month does not describe demand. Business-launch services have pronounced peaks; gift and holiday topics can rise severalfold. Review twelve to twenty-four months of history. This workflow suggests publishing two or three months before the expected peak to allow time for search visibility to develop.

9.3 Building a keyword set from seed terms

Collect in three ordered steps to avoid bias toward the ideas that come to mind first.

Use six seed sources: your services in clients' language, queries for which similar providers rank, search suggestions, related questions, your own performance report, and real client correspondence. The last often supplies commercially valuable wording absent from research tools.

Expansion means feeding seed keywords into research tools, collecting suggestions for each seed by appending letters, and exporting the keywords competitors rank for.

Cleaning removes queries unlikely to bring a client: other companies' brands, irrelevant regions or languages, searches for free downloads, jobs or employment, and training or courses unless you offer training.

Procedure. Reuse an exclusion list and review its matches before deletion. Measure how much irrelevant material it removes and how many useful queries it wrongly excludes; no verified 80% success rate is available. [practice].

9.4 Clustering

Search-result overlap and semantic similarity provide different clustering evidence. Neither is universally more reliable; validate the groups against intent and observed performance.

Overlapping results for two queries suggest similar intent. A lack of overlap does not prove that one page cannot rank for both: results change and reflect the documents available.

Check the grouping against the user’s task and the site’s actual results.

The overlap threshold is a setting in the chosen tool. For example, start with four shared URLs, then manually check groups with mixed intent. This is not a search-engine-defined boundary between matching and different intent. [example]




Figure 9.1. Illustrative clustering settings: five shared URLs meet the chosen threshold of four; one shared URL does not. This guides manual review; it does not prove that a page cannot rank. [example]

In the terminology used here, hard clustering checks pairwise result overlap, while soft clustering compares queries with a central query. Tool implementations differ. Neither mode guarantees precision or suitability for a particular site; inspect the groups and document the settings. [opinion].

Manual review is essential; a threshold setting cannot replace it. The algorithm does not understand meaning and may merge ambiguous terms. In Russian, for example, “дизайн карты” can mean bank-card design or map design.

These may enter one group if their search results happen to overlap.

Procedure. For commercial groups, try hard clustering with a threshold of four shared URLs. Review each resulting group to check that every phrase serves the same intent, and split mixed groups. Set aside ungrouped queries: they may be too rare or express a distinct intent for which you have no material yet. Assign each group a page type and a working title.

9.5 Reading search results before writing

Review the results for the cluster's main query before writing. Allow fifteen minutes to assess whether you can compete and what the page needs to contain.

Check six things.

- Page types in the top ten: these indicate the type your page needs to match.
- Who ranks: service providers, aggregators, directories, or media outlets.
- How fully competitors answer the query: which questions they cover and which remain unanswered. Their word counts do not set a minimum for your page.
- Whether prices are shown. If everyone else publishes prices, omitting yours creates a conspicuous disadvantage.
- Whether the content is current. An old publication date alone proves neither that a page is outdated nor that it will be easy to outrank.
- Whether an AI answer appears: record it as a possible influence on CTR, not a fixed deduction from expected clicks.

If aggregators dominate, competition is difficult and their format may be preferred. This does not prove that the remaining positions are inaccessible. Compare building a page with alternatives: getting listed on an aggregator, targeting a narrower task, or creating a page for visitors from other channels.

Procedure. Summarize the ten results in a table: URL, page type, website type, approximate content length, and whether prices are shown. Save it with the page brief. A year later, it will explain why the page was built this way and let you compare current results with the earlier ones.

9.6 Prioritizing along three dimensions

Priority depends on three factors; none is sufficient on its own.

Demand — estimated search volume for the group, accounting for overlapping query variants and search operators. Simply adding the figures without removing overlap can overestimate demand. **Difficulty** — how hard it will be to compete: who ranks at the top, how old their pages are, how many backlinks competitors have, and whether aggregators dominate. **Value** — how close the query is to a purchase. In descending order: service-specific commercial queries, industry-specific commercial queries, informational queries related to the service, and general educational queries.

As a guide, priority is proportional to demand multiplied by value, divided by difficulty. The result is not a benchmark to compare with someone else's score. Its purpose is simply to put your own list in order.

Illustrative comparison. The informational group has a search volume of 2,400; the commercial group, 90. The first is roughly 27 times larger, but that alone does not determine priority: achievable CTR, inquiry quality, and costs are unknown. First compare several scenarios using the formula in Part I. [example]

9.7 Why marking everything P1 is not prioritization

Priority allocates limited resources. Marking all twenty-five services as top priority does not help planning; it only shows that no choice has been made.

Part I defines the capacity limit: available hours divided by the hours needed per page. If capacity is twenty pages a year, at most twenty pages can have top priority. In practice, choose fewer, because updating published material also takes time.

Capacity rule. Give top priority to work that fits into the next planning period while leaving time for maintenance. A limit such as one fifth of the list may help a particular team make decisions, but it is not a universal rule.

9.8 A bilingual keyword set

Research a fresh keyword map for each additional language instead of translating the existing one. Demand differs between markets: people use different phrases, search volumes and competition differ, and the results contain different types of pages. [61]

Translating a Russian query cluster may produce a phrase that nobody actually searches for in English. The reverse also

applies: established English expressions may have no direct Russian equivalent and need descriptive wording, turning one query into three variants with different search volumes.

Versions can share material that does not depend on search phrasing: case studies, images, and project facts. Each version needs its own structure, headings, and priorities.

Language-version annotations connect related pages and help search engines show visitors the appropriate version. They do not provide a ranking boost or replace work on each version. Confusing this with market entry leads people to assume that translating a site is enough to enter a second market. [opinion].

Worked examples A and C: choosing topics

In example A, priority goes to one service the specialist can substantiate with two case studies. This choice does not yet establish search demand: check query wording and results before expanding. In example C, 60 topics are already published, but 18 need checking for overlap. High search volume for a broad topic does not automatically justify another page. [example]

X • Trust

10.1 The About page as evidence

This page is more than a biography or career story. It provides evidence that the service provider is real, has the stated experience, and can be trusted to stay with the project. It addresses the third risk described in Part I through verifiable claims.

A verifiable claim has a practical test: it can be disproved. “Worked at this company from 2019 to 2022” can be checked and therefore carries weight. “Successfully completed many projects” offers no specific claim to test, which is why it is unconvincing.

Include your full name, photo, years of experience and number of projects, industries served, tools, location and time zone, contact methods, and links to professional profiles and portfolio platforms.

A photo serves a purpose. When someone is being asked to hire you, the absence of a face can suggest concealment, especially for services that involve months of working together.

Order the information around the reader’s checks: who you are and what you do, what supports your claims, then how to contact you. A chronological career story interests colleagues far more often than clients.

10.2 Reviews

A testimonial is valuable when it describes a change.

“Excellent specialist; we liked everything” says little. “We reduced checkout from seven steps to three, and the work took six weeks instead of the planned eight” describes both the outcome and the delivery.

Ask at handover, when satisfaction is fresh. Six months later, interest has faded and details are forgotten. Instead of asking the client to “write a testimonial,” ask three specific questions. A vague request leaves them to do the writing, often producing either silence or generic praise. Turn their answers into a draft and send it back for approval.

Ask what the task was and what had prevented a solution; what changed after the work, preferably with numbers; and what was unexpectedly good or bad about the process. The third question supplies the detail that distinguishes a real testimonial from one invented by the provider.

What makes a testimonial verifiable. Include the person’s full name, company, role, and profile link. “Alexey, entrepreneur” provides no way to verify the testimonial or distinguish it from an invention. If the client does not consent to full identification, publish it anonymously and state that clearly instead of implying verification.

What you can't do. Do not invent a testimonial and send it to the client to sign: this common practice can destroy trust as soon as a prospective client speaks to them. Do not change its meaning when editing or publish it without explicit permission.

10.3 Publications, talks, and open-source code

The mechanism is the same as for testimonials: support comes from a third party with no stake in your sale. Your own claim to expertise carries limited weight; someone else's decision to publish your work or invite you to speak carries more.

Useful evidence includes an article accepted through editorial selection in a trade publication; a conference talk listed in a public programme; a training course; a public project with a change history and users; or a contribution to a professional standard or review.

Open publication or paid membership alone does not establish expertise. A participation fee also does not, by itself, invalidate an award: inspect its selection criteria and independent assessment.

The distinction is simple: if access depended only on your willingness and money, it proves nothing about your expertise.

10.4 External links: what works

A backlink can bring visitors and demonstrate your work's connection to the professional community. To earn mentions, start with something useful: research, a tool, a template, a detailed project analysis, or an open-source project.

Choose distribution methods to suit the audience: pitch an article to an editor, contribute to a discussion, give a relevant talk, or publish in a professional community. You cannot

guarantee a number of backlinks in advance. Plan the material and relevant outreach instead.

This edition excludes quoted popularity percentages for PR and guest posting when the underlying methodology is unavailable. Calling such figures “a rough guide” does not make them reliable.

To assess results, record the source, what was mentioned, referral visits, and known inquiries. A link from a small industry website may be more useful than a widely seen mention if it brings suitable clients. [opinion]

10.5 Black-hat tactics and the cost of recovery

There are two forms of action against link manipulation. Their differences determine how you can detect them.

Manual actions are imposed by a human reviewer. A notification appears in the webmaster console, explicitly identifying the issue. Link spam means creating links to manipulate rankings. It includes buying or selling links, exchanging goods or services for links, and advertising content containing links that pass ranking credit [standard, 26]. The policy makes an important exception: paid links do not violate it when appropriately marked with `rel="sponsored"` or `rel="nofollow"` [standard, 26]. Without the required qualification, the paid link violates the policy.

Algorithmic devaluation can remove the ranking credit of manipulative links without a manual-action notice. A ranking drop alone does not prove link-spam devaluation or reveal a precise underlying ranking level. [57]

No verified safe proportion of guest posts is given here. The “30% of the backlink profile” threshold has been removed: Google’s policy considers the nature of the links and the intent to manipulate rankings, rather than setting such a percentage. [standard, 26]

The cost of recovery. For an unnatural-links manual action, follow the issue-specific remediation and reconsideration procedure. Disavowal is a specialized tool, not routine cleanup for every unwanted link. Neither reconsideration timing nor restoration of previous rankings is guaranteed. [59]

Assess the risk through your dependence on search and the cost of recovery. Neither a hundred-page site nor ten monthly inquiries determines resilience, and recovery is not predictably faster or slower than legitimate growth. [opinion].

10.6 Organization and author structured data

Person and Organization properties can explicitly describe identities and related profiles. Profile URLs are useful evidence, but no one property is universally decisive and markup does not guarantee that every system merges the mentions correctly.

Describe the author consistently, maintain an author profile page, and explicitly list alternative names and related URLs. Cyrillic and Latin spellings can be connected through `alternateName`, `sameAs`, and a stable identifier. Character-for-character name matching is not a universal requirement.

Article authorship markup links to the author page, which in turn links to external profiles. This creates a connected set of references that machines can check.

XI • Production

11.1 A production process that preserves quality

A repeatable workflow can reduce omissions, while an output target alone can encourage repetitive work. Regular publishing and templates are not themselves spam; assess the content’s value and whether it is primarily designed to manipulate search rankings.

Resolve this by standardizing the procedure: keyword research, search-results review, number checking, editing, publishing, and internal linking. These are repeatable steps that should be done consistently. The workflow should not dictate what the text actually says.

A warning sign is that pieces can be swapped without changing the meaning: each follows the same template and says the same thing in different words.

11.2 Realistic pace

Derive the publishing pace from the time per page calculated in Part I. Ambition alone does not determine capacity.

Content hours per week	Pages per month	Pages per year
5	1	10-12

Content hours per week	Pages per month	Pages per year
10	2	20-25
20	3-4	40-50
40 (full-time)	7-8	80-100

The table is a planning example assuming 18 hours per page, with additional time reserved for correspondence and revisions. It is not a measurement of the author's productivity. Replace the estimates with actual time records after the first publications. [example]

Two articles a week can consume a substantial share of working time if each requires research and original material. Test the pace over a trial cycle, accounting for automation, editing, and concurrent client work.

Plan complete groups of work with spare capacity instead of assuming maximum output every week. Coherence comes from topic selection and linking, not from the size of the initial plan.

11.3 The production pipeline

There are eight steps. Each produces a concrete output required before the next step begins. This prevents reaching step eight only to discover that step two was skipped.

1. Choose a query cluster. Output: a register entry specifying the page type and target query.
2. Review search results. Output: a table of observations on ten results, following the procedure in Part IX.

3. Prepare the brief. Output: the page structure with subheadings, the claims that need evidence, and the figures and sources required.
4. Gather facts. Output: sourced figures, quotations, and screenshots. Gather these before drafting, rather than while writing.
5. Draft. Output: a draft covering the structure set out in the brief.
6. Edit. Output: text that passes the editorial checklist in Part VIII.
7. Visual materials. Output: images with the required dimensions and file sizes, captions, and alternative text.
8. Posting and linking. Output: a published page with links to and from related content.

The order of steps four and five matters. Gathering facts after drafting encourages fitting evidence to claims already written instead of building claims from evidence.

11.4 Roles in a one-person operation

If different people cannot handle writing and editing, separate the tasks in time. Clichés that are invisible on the day of writing often become obvious two days later; this is a matter of attention, not discipline. Draft and edit on different days, preferably at least a day apart.

Apply the same delay to fact-checking: revisit collected figures later, rather than on the day when their source still seems convincing.

Tasks worth delegating. Proofreading is an inexpensive and effective task to delegate: fresh eyes catch typos the

author misses. Image preparation in the required formats and sizes can also be delegated, as can markup and publication when the process is documented.

Tasks that require accountable judgment. A qualified contributor can review search results, gather sources, and prepare conclusions. The author or project owner must accept the work and remains responsible for published claims. Delegation does not remove the need to verify.

11.5 Definition of Done

A page is ready only when every item is complete. “Ready with outstanding issues” is not a completion status: the page is either ready or still in progress.

- The search title and description are filled in and checked against the available display space. [63]
- This workflow uses one descriptive H1 and a logical heading hierarchy. This is a structural convention, not a Google ranking requirement. [58]
- Factual figures are checked against their sources; hypothetical calculations are labeled as examples.
- The editorial checklist in Part VIII is complete.
- Images have suitable formats, dimensions, and text alternatives. Prioritize the critical above-the-fold image; decide lazy loading for each other image by position and purpose.
- Links lead to the related service or case study and to two or three other pages in the cluster.
- Relevant sections and articles link to the page.

- Structured data has been generated and checked with a validator.
- The page is present in the site map.
- The page has been checked on a narrow-screen device.
- Rendering and layout-shift metrics have been measured after publication.
- The register entry is updated with status, date, and target query.

Item seven is often overlooked, yet it prevents the orphan pages described in Part IV.

11.6 Updating published content

Updating a useful existing page may be more worthwhile than publishing a new one: the URL, content, and performance history already exist. Decide based on the page's condition. The success rate in Ahrefs' random-URL sample is not the ranking probability of each new page on your site.

What to update first of all. Start with pages ranking fifth to twentieth: modest gains can produce a disproportionately large increase in visits as visibility improves. Also examine pages with falling impressions, which may have become outdated or face more complete competitors, and pages with stale figures. Presenting a five-year-old number as current can harm trust more than leaving it out.

What to do when updating. Review search results again: competitors and expectations may have changed over the year. Update figures, sources, and dates. Fill relevant gaps that competitors cover. Refresh internal links to include

useful material published since the last review. Update the revision date while keeping the original publication date; both help readers.

Capacity allocation rule. As a planning guideline, allocate about 30% of production hours to updates and 70% to new content. At around a hundred pages, move toward an equal split: keeping a hundred pieces current becomes a substantial task in itself [practice].

11.7 The content register

Use one register for the whole project. Seven overlapping tables leave no reliable source of truth, and even a simple status question can require reconciling three files.

Register fields: URL, page type, query cluster, target query, exact-match search volume, priority, status, publication date, last updated date, impressions in the last month, average position, and inquiries if attribution is configured.

Statuses that reflect reality. Distinguish “idea” from “planned.” An idea has not been approved; planned work has a decision and an assigned place in the queue. Confusing them produces registers with a hundred and thirty-eight “planned” pieces that will never be made, as described in Part I.

Add a “rejected” status with a reason. Keeping a rejected idea in the register with a note such as “results dominated by aggregators” saves time in the next planning round: you will not repeat the same investigation.

Periodic review. Review the whole register quarterly. Work that has been “planned” for over a year should either start or

be rejected. A register that never rejects anything becomes a wish list instead of a planning tool.

Worked examples: three different budgets

In example A, the first version has a 96-hour budget, including 24 hours of contingency. Example B allocates 20 hours to diagnosis and improvement of the existing visitor journey. Example C limits further work to 20 hours of repairs and stops mass publication. These estimates reflect different starting materials, not differences in author productivity.
[example]

XII · Stack and infrastructure

12.1 Four types of platform and their implications

Choose a platform by its data handling, available controls, support requirements, and cost of change. The category name alone determines neither a page limit nor performance.

Type	Working with content	What to check before choosing
Website builder	individual pages or CMS collections	templates, export, redirects, subscription-plan limits
Ready-made CMS	posts and themes, sometimes a visual editor	updates, editor permissions, caching
Static site generator	files or connected CMS	preview, builds, forms, author skills
Framework	depending on implementation	cost of development and maintenance

A website builder with collections can update cards through a shared template; a static site generator can provide an editorial interface. Compare actual features and limits

instead of assuming that every page must be edited manually or every author must work with code.

12.2 Selection criteria

Answer five questions before choosing.

How large will the site be in three years? If the answer is up to fifty pages, a website builder remains a reasonable option; a more complex system may not repay its extra cost.

How often will you publish? One page a month needs less process than eight pages a month, which calls for a repeatable production workflow.

How many people will edit content, and what are their skills? A static site generator that requires text files and version control needs editor training. This practical constraint can outweigh its technical advantages.

How long should the project last? A two-year solution and a ten-year solution have different priorities; data portability matters more in the latter.

Will a developer remain available? A framework-based site without developer support can become unmaintainable when dependencies need updating.

Rule of choice. Choose around the constraint that would cost you most, rather than the feature you like best.

Static-site speed is of little use if the person responsible for publishing cannot use the required text-file workflow: the site will stop being updated.

12.3 Build time as a constraint

Build time depends on templates, data volume, image processing, caching, hardware, and tool version. Claims such as “ten thousand pages in ten seconds” are not a sound basis for choosing unless tested on the same workload.

A trial build should include representative case studies, images, and relationships. Measure a full build, an update to one record, and the time until a preview is available. If waiting obstructs publishing, consider caching generated images, incremental builds, or on-demand rendering.

For a small site, a reliable edit-check-publish cycle matters more than a speed record achieved with empty pages.

12.4 Hosting, delivery, protection

A CDN can serve a cached response from a node near the visitor. The benefit depends on routing, cache hits, and configuration. Measure TTFB and LCP in target regions before choosing hosting. The proportions in the LCP diagnostic example are not measurements of your site’s timing budget.

A server in one location offers a different set of trade-offs: full control and server-side logic, together with responsibility for system updates, certificates, load monitoring, and backups.

A CDN and your own server can work together: the server is the origin, while the CDN provides caching and delivery. You can also host static content separately from the form-handling server. Document request routes, cache boundaries, and responsibility for certificates.

Baseline protection for a content site includes HTTPS-only connections, automatic certificate renewal, rate limits for form requests, bot protection without image-based puzzles, and backups. A full web application firewall can help, but for a site without user accounts or payments it is not the first priority.

12.5 Repository, environments, and deployment

For a one-person project, a simple workflow is sufficient: the main branch represents the published site; changes happen in separate branches and merge after preview checks; merging triggers an automatic build and publication.

A permanent development branch, release branches, and mandatory review by another person suit a team of five to ten. Applied to a solo project, this can add overhead without the same benefit: the author would be reviewing their own code.

The essentials are automatic builds after changes, a preview before publication, and a block on publishing when the build fails. That last safeguard makes the schema validation described in Part V an effective control.

12.6 Backup

Do not assume a repository is a complete backup. It holds source material and, although unlikely, can still be lost. More importantly, it often excludes vulnerable assets: full-resolution originals, collected testimonials, publication-consent emails, and analytics exports.

Use a documented backup plan. The 3-2-1 guideline keeps three copies including the working copy, on two media types, with one offsite. Add isolation and restore tests according to the risks; two synchronized locations alone are not a complete recovery plan. [52]

The only test that matters. Treat an untested backup as unproven. As an initial schedule, test restoration every six months and after material system changes; adjust the interval to risk. Restore into an isolated environment and record elapsed time and missing data.

12.7 Migration: eight ways to lose traffic

Part IV sets out the migration procedure. Below are common implementation mistakes and how each causes damage.

An incomplete list of old addresses. A list taken only from the sitemap misses pages removed from it a year ago that still have backlinks and rankings. Those URLs are omitted from the redirect map and return errors.

Redirecting everything to the home page. This can be treated as a missing page, causing its search signals to be lost.

Redirect chains. The old URL leads through an intermediate URL to the final one. Uncached HTTP redirects add requests and responses; long chains may exceed a crawler's limit. Prefer a direct redirect to the relevant destination.

Changing URL structure during the migration. Two simultaneous changes make diagnosis harder: if performance falls, it is unclear which caused it.

Using a redirect type that does not match the move. Use permanent redirects for a permanent migration. A temporary redirect sends a different signal. Search engines' treatment of it depends on circumstances and duration, so you cannot guarantee that the old URL will remain unchanged in the index.

Leaving staging crawl restrictions in place after launch. A copied robots.txt rule can block crawling without reliably removing known URLs from the index. A retained noindex directive can prevent indexing when Google can read it. Check both controls separately.

Losing structured data and metadata. Preserve titles and descriptions as well as body text. Their loss can alter search presentation, but a missing meta description does not by itself establish a ranking penalty. Check the actual title and snippet Google displays.

Failing to monitor after the switch. Use thirty to sixty days as an initial monitoring plan, extending it while issues remain. Track index coverage, errors, and comparable query groups. This period is not a deadline for removing permanent redirects.

Worked example A: choosing a simple setup

Test the first version with one case study and one service page in an accessible template-based tool. Check whether you can edit text independently, export content, use a working form, and preserve URLs. A successful trial publication matters more than promised build speeds for thousands of pages the project does not yet have. [example]

XIII · A design system for a content website

13.1 Typography as the foundation

On a text-led site, typography takes on much of the work that composition, colour, and imagery perform elsewhere. Three measurable variables shape readability.

Line length. Around 50–70 characters per line can be a starting point for long text, followed by testing with the actual typeface, language, size, and device. This is an editorial starting point, not a universal limit for reading speed. [opinion]

The source layout uses an em-based text measure. The number of characters on a line depends on glyph widths, font size, language, and viewport; verify it in the rendered layout rather than assuming 68 characters.

Line spacing. For long text, start with line spacing of 1.5–1.65 times the font size. Values below 1.4 can crowd lines; values above 1.8 can make a paragraph feel disconnected. Check the result in the chosen language and typeface.

Font size. Start the layout at 18–20 pixels, then test readability, zooming, and the density of the chosen typeface. A 16-pixel size is not inherently wrong. [opinion]

Type scale. Keep the site to five or six type sizes. Each extra size requires a decision about when to use it, and those decisions tend to vary from one place to another.

13.2 Design tokens and themes

A design token is a named value used instead of a literal colour or size. It lets you change the palette in one place without searching every file. This applies the single-source principle from Part V to visual design.

A minimal palette for a content site covers body text, secondary text, the page background, highlighted-block backgrounds, rules, and an accent colour. Six roles are enough; a forty-shade palette is likely to go largely unused at this scale.

Test dark mode separately. Pure white on black can be uncomfortable for some readers, but no single lightness range suits every font and screen. Check contrast, whether elements remain distinguishable, and user preferences.

Perceived font weight may change between light and dark themes. If you adjust weight, test readability in both versions instead of automatically making one a step lighter.

13.3 Grid and vertical rhythm

An eight-pixel modular grid supports horizontal alignment. For a content site, vertical rhythm matters more: spacing between elements follows multiples of the body-text line height rather than arbitrary values.

Leave noticeably more space above a heading than below it. The heading belongs to the text that follows; equal spacing can briefly obscure which paragraph it introduces.

Limit the text column by the line length in section 13.1, not by screen width. Give wide tables, diagrams, and detailed images their own wider treatment and internal scrolling where needed, without widening the body text.

13.4 Components

A site with services, case studies, and articles needs a header with navigation, footer, content card, related-content block, breadcrumbs, expandable questions, call to action, form, table, callout, image caption, and section divider.

Define each component's normal state, hover state, keyboard-focus state, and fallback appearance when images or fonts fail to load. Missing states show up as defects in use.

Keyboard focus is especially easy to overlook. Without a visible focus indicator, someone pressing Tab cannot tell where they are on the page.

Content cards deserve special attention because they recur throughout the site. They must work with or without an image, with one-line or three-line titles, and with or without a description. Otherwise, the first long title can break the grid.

13.5 Accessibility

The following is a partial accessibility checklist for an initial content-site review. Passing it does not establish WCAG 2.2 AA conformance, which requires checking every applicable criterion. [standard, 29]

Normal text needs a contrast ratio of at least 4.5:1 against its background. Large text—at least 18 pt, or 14 pt bold—needs 3:1. Relevant interface components and meaningful graphics

also need 3:1. AAA requires 7:1 for normal text and 4.5:1 for large text [standard, 29].

Check keyboard operation and visible focus. Focus order must preserve meaning and operability; it need not mechanically duplicate visual order. Use meaningful headings. Informative images need appropriate alternatives; purely decorative images normally use empty alt text. [48]

Do not rely on colour alone to convey meaning. Some readers cannot identify links distinguished only by colour; use underlining or another cue.

13.6 Animation and its cost

Animation can affect responsiveness and stability. Its cost depends on implementation: geometry changes, main-thread work, and content hidden until a script runs all need separate checks.

A conservative starting point is opacity and transform animations lasting 150–300 milliseconds, without resizing or moving neighbouring elements. Geometry changes can be much more expensive because they require layout recalculation.

Respect the system's reduced-motion preference. Some people disable animation because screen movement causes dizziness or nausea, not simply because they dislike it. A short CSS rule can honor this preference; ignoring it can make the site unusable for them.

Scroll-reveal effects need particular care: if the script fails, the content can remain invisible. Content should be visible by default, with animation added as an enhancement.

Extending the accessibility review

Also test text resizing, reflow on narrow screens, field labels, error messages, accessible button names, unobscured focus, and target sizes. Videos need the applicable text alternatives. Complement automated checks by completing the main user journeys with a keyboard. [standard, 29]

XIV • Measurement

14.1 Metrics for decisions and metrics for reports

A metric is useful when it helps choose an action. Each report should state the event definition, data source, period, and denominator. Impressions, sessions, users, and inquiries are not interchangeable.

After launch, first verify that tracking works and the site is available. As data accumulates, add search impressions, visits interpreted by page purpose, confirmed inquiries, and contracts. No fixed month makes a metric informative.

Be cautious with small inquiry counts. Moving from one order to two does not prove that the underlying conversion rate has doubled. Show absolute counts and the observation period alongside rates. A business can still make a cost-based decision before there is enough data for a firm statistical conclusion.

14.2 Setting up measurement without self-deception

Segmenting traffic. Landing-page type is a useful segmentation dimension, but it does not prove each visitor's intent. Article readers may become clients, while service-page visitors may be peers. Also segment by channel, available query groups, and device.

Checking measurement quality. Exclude test events, your own inquiries, and known spam. Identifying every bot is usually impossible, so document the limits of filtering.

Distinct stages of action. Viewing contact details and clicking a messaging link are intermediate events. A successful submission means the server accepted the form. Receipt of the message, qualification of the inquiry, and a signed contract are later stages. Do not label all these events as leads.

Keep a change log alongside reports: date, affected pages, purpose, and verification. A metric changing after an edit is a reason to investigate a connection, not proof of causation.

14.3 Inquiry attribution

In services, attribution is often lost between the initial search, reading on another device, and a direct message. Record what is known and explicitly preserve an unknown-source category.

Forms can include the page URL and permitted source tags. For messaging apps, use supported parameters or prefilled text. Ordinary UTM tags do not necessarily appear in the

message you receive. Test the chosen method end to end.
[51]

In the first conversation, briefly ask how the person heard of you and record the answer. Separate their first encounter with your name from their last action before contacting you; the channels may differ.

Matching by time is only a hypothesis. A visit usually precedes an inquiry, but temporal proximity cannot reliably establish identity or source. Keep unknown-source inquiries separate, without inventing a “normal” percentage of attribution loss.

Worked example B. Of three qualified inquiries after the changes, two arrived through a tracked form and one directly from an unknown source. Report two confirmed attributions and one unknown. Do not automatically credit the third to the recent increase in organic traffic. [example]

14.4 Misleading metrics

Average position. This aggregates impressions across a changing mix of queries, devices, and regions. Read it alongside impressions and clicks: a worse average position can accompany broader visibility for new queries.

Overall conversion rate. It describes the whole site under the chosen event definition but hides differences between traffic segments. Add segment-level analysis for decisions rather than dismissing the overall metric as meaningless.

Bounce rate in GA4. This is the share of sessions that are not engaged sessions. With the default threshold, a session is engaged if it lasts more than ten seconds, includes a key event, or has at least two page or screen views. Reading one

page for a long time is therefore not necessarily a bounce. Check the engagement threshold and event collection settings. [standard, 33]

Engagement time in GA4. It counts time when the page has focus or the app is in the foreground. It is not simply the interval between two page views. Read other analytics systems’ definitions separately. [standard, 34]

Pages in the index. This measures coverage rather than commercial success. Compare useful, indexable pages with those actually indexed and investigate discrepancies.

14.5 Expected trends and warning signs

The table suggests what to investigate; it does not diagnose a problem from three numbers. Several causes may operate at once.

Observation.	What to check
Impressions rise; clicks do not	new query mix, positions, device, search-results format, CTR
Clicks rise; inquiries do not	intent, offer, trust, contact journey, form delivery, tracking
Few or no impressions	demand, accessibility to search engines, indexing, canonicalization, relevance
Impressions and clicks fall	seasonality, changes in the site and search results, competitors, technical errors
Inquiries arrive; contracts do not	lead qualification, price, response times, sales process, available capacity

Rank possible causes by the evidence available. Fix a broken form immediately; check insufficient demand against query data and the commercial hypothesis against inquiry quality. Waiting nine months cannot replace these checks.

Example B. After the changes there were 1050 organic-search visits, including 210 visits to service pages. Seven form submissions produced two qualified inquiries; a third qualified inquiry arrived directly, with an unknown source. The comparable form-based rate changed from 1/120 to 2/210, approximately 0.83% to 0.95%. You must not divide all three inquiries by 210 and call that the service-page conversion rate. Small counts and a bundle of changes do not establish a causal effect. [example]

Example C. After the delivery fault was fixed, all three test submissions arrived; they were excluded from inquiry counts. The following month brought 180 organic-search visits, including 24 to the service page, one genuine submission, and no qualified inquiries. The technical repair is verified; the commercial hypothesis remains unproven. [example]

14.6 Monthly review

At each chosen review interval, compare search visibility, visits to relevant pages, genuine inquiries, contracts, and hours spent. Use a period long enough to collect useful observations, while fixing technical faults as soon as they appear.

Record what changed, which possible causes were checked, and what action the data supports. Choose one feasible next experiment with a goal, scope, budget, and review condition, rather than making an unfunded list of tasks.

In example A, the week-twelve review closes the initial 96-hour budget. In example B, observation continues because the data cannot yet establish an increase in conversion. In example C, broad overview publishing stops and only a limited experiment with one service proceeds. Each decision has different grounds; there is no common calendar threshold. [example]

XV · Law and ethics

This section offers practical considerations, not legal advice. Rules vary by country and change over time, so consequential decisions should be checked with a lawyer in the relevant jurisdiction. The points below identify what to agree and document, rather than predicting how a court will interpret it [verify independently].

15.1 Confidentiality agreements and rights in the work

These are two separate issues that are often confused.

A nondisclosure agreement restricts disclosure: what cannot be shared, with whom, and for how long. It may prohibit naming the client, showing materials, or even mentioning that the work took place. [53]

Rights in the work determine who may use what was created and how. Address this separately in the contract. Silence does not grant unrestricted use: default legal rules vary by country and may differ from the provider's expectations.

Agree portfolio publication rights in a separate clause. They do not follow automatically from transferring rights to the client or from the absence of an NDA. Specify four things: exactly what may be shown, where, from what date, and whether the client may be named.

Timing matters. Publishing before the product launches can harm the client and is commonly restricted. A condition such as “not before public launch” or “not until six months later” can allow publication without an outright refusal.

15.2 Publishing client work

Record permission in writing even if it was given verbally. Staff change jobs; a year later, an oral agreement may be difficult to establish when a different employee raises an objection.

Determine the appropriate form of permission, the approver’s authority, and whether the contract needs amendment for each case. Preserve correspondence, but do not assume a standard reply makes approval sufficient. [verify independently]

If permission is refused, do not publish the restricted material. Agree separately on any anonymized case study: hiding the client’s name does not automatically permit disclosure of project content. If too little material is authorized, describe your general method without client information.

What you can't do. Do not publish merely because “the client never prohibited it.” Do not display client-supplied photos, text, or data without separate permission: rights in your own work do not automatically cover those materials. Do not claim credit for parts produced by other contractors.

15.3 Generative tools

Answer three questions before using generative tools.

Rights in the output. Copyright treatment depends on jurisdiction and human contribution. For example, the US Copyright Office’s 2025 report distinguishes human-authored expression from wholly AI-generated material; using AI assistance does not automatically remove protection from the human contribution. Check the applicable law and the tool’s terms separately. [49][verify independently].

Disclosure. State the use of generative tools in the case study and delivered materials. Do not present a generated image as a photograph. If discovered, that misrepresentation undermines trust in the whole case study, as described in Part VI.

Similarity. Reverse image search may reveal some matches, but it neither rules out legal risk nor checks every right relating to likenesses, marks, or source material. Agree the review process and responsibilities separately for client publication. [verify independently]

For text, also check facts, confidentiality, borrowed material, and the source materials’ terms of use. Editorial revision does not replace checking publication rights. [verify independently]

15.4 Personal data in forms

Collect data that is necessary for a specified lawful purpose. Under the GDPR, where applicable, this is data minimisation; it is broader than simply shortening a form. Additional questions may help qualification, so test their business value separately. [50][68]

For each form, define what data is collected, why, where it is stored, for how long, who has access, and how someone can request deletion.

Where the data actually goes. Map each recipient and storage location in the submission chain. Determine who acts as controller or processor under the applicable law and contracts. Outsourcing a step does not automatically discharge the site owner's responsibilities.

Do not send names, phone numbers, email addresses, or form-message contents to Google Analytics. Its PII policy also covers URLs and event parameters. Track the submission event without copying the submitted personal data.

Provide an accessible privacy notice at collection and establish a lawful basis before processing. Under the GDPR, where applicable, consent is one possible basis, not the only one; the notice itself is not consent. Check retention, deletion rights, transfers, and local requirements separately. [verify independently].

A • Checklists

A.1 Publishing a page

This follows the definition of done in Part XI. Partial completion means the page is not ready.

- The search title and description are filled in and checked against the available display space. [63]
- One descriptive H1 and a logical heading hierarchy, as this workflow's structural convention.
- The page type matches the intent identified through search-results review.
- Factual figures are checked against their sources; hypothetical calculations are labeled as examples.
- The editorial checklist in A.4 is complete.
- Images use modern formats, specified dimensions, captions, and alternative text.
- The critical above-the-fold image is not lazy-loaded; other images load according to their position and purpose.
- Outgoing links include a related service or case study and two or three nearby pages in the query cluster.
- Important pages have accessible incoming links from relevant sections and content.

- Structured data is generated from record fields and checked with a validator.
- The page is present in the site map.
- Tested on a narrow screen.
- Rendering and layout-shift metrics have been measured after publication.
- The register entry is updated with status, date, and target query.

A.2 Launching a website

- The URL scheme is documented and followed across all pages.
- The sitemap is generated automatically and includes only indexable URLs.
- Crawl rules and indexing directives have been checked separately; staging restrictions are removed.
- Each independent, indexable page has a consistent canonical URL; duplicates are handled separately.
- The secure connection works, the renewal of the certificate is automatic.
- The non-www version redirects to the preferred www version, or vice versa, in one hop.
- An error page exists, provides navigation, and returns the correct HTTP status.
- Analytics is installed, own visits are excluded, and the goal tracks the intended event.
- Both search engines' webmaster tools are connected, and the sitemap is submitted.

- Forms are tested: messages arrive, fields are preserved, and bot protection works.
- The data-collection notice is published and linked beside the form.
- Backups are configured and restoration has been tested.
- Lab measurements and real-user field measurements are separate; missing field data is disclosed.
- Keyboard navigation and visible focus are checked.

A.3 Website migration

The full migration procedure is in Part IV, section 4.6. The sequence is listed here.

Before switching.

- Old URLs are collected from four sources: the sitemap, the performance report over the longest available period, server logs, and analytics reports.
- The lists are merged and duplicates removed.
- Each old URL is manually mapped to its replacement; URLs without replacements are marked to return “not found.”
- Necessary URL changes are recorded in the map; optional changes are deferred.
- Every mapping is tested in staging for response status, redirect hops, and final URL.
- Redirect chains longer than two hops are shortened.
- All redirects are permanent, not temporary.
- Metadata and structured data are migrated with the text.

After switching.

- Crawl restrictions are removed from the live domain.
- The sitemap is resubmitted.
- Monitor initially for thirty to sixty days and continue while issues remain. Keep permanent redirects generally for at least a year, following Google's migration guidance.
- Record the switch date alongside metrics to help investigate subsequent changes.

A.4 Editor's checklist

In accordance with Part VIII, Section 8.7.

- Search the clichés catalogue - zero matches.
- Every figure has a source or an explicit note that no source is available.
- Each evaluative adjective is supported by evidence or replaced with a fact.
- The first sentence of each section carries the content, not the announcement.
- No section ends with a retelling of the section.
- Each list is checked for the uniformity of the elements.
- Nominalizations are replaced with verbs where appropriate.
- Use active voice where identifying the actor matters.
- Long paragraphs are checked in the finished layout; coherent arguments are not split merely to meet a line count.
- There's at least one place where it says the approach doesn't work.

- There's at least one detail you can't think of without doing the work.
- Check dashes, quotation marks, non-breaking spaces, and number grouping.

B • Templates

B.1 Case-study brief

Complete this before writing. Empty fields reveal what the case study lacks—usually outcome evidence.

Project:

Client (name or anonymized description):

Industry and scale:

Work period:

My role and others' contributions:

Publication permission: granted / not granted / restricted

Restrictions: client naming, earliest publication date, material that cannot be shown

Client's task (in their own words):

Why it arose at that time:

Constraint (what prevented the obvious solution):

Options considered:

Chosen option and reasons:

Trade-off (what was sacrificed):

Outcome figures supplied by the client:

Measurable product properties before and after:

before / after —

before / after —

Other factors that may have affected the outcome:

What I would do differently:

Related service:

Related articles:

Case study's target query:

B.2 Article brief

Query cluster:

Target query, estimated search volume, region, period, and operators:

Intent (informational / commercial / mixed):

Search-results review (ten results):

URL | page type | website type | content length | prices shown

Conclusion: suitable page type and visible constraints

Structure (statement-style subheadings):

- 1.
- 2.
- 3.

Claims requiring evidence:

Figures needed and where to obtain them:

Firsthand experience to include:

Where to explain when the approach does not work:

Related service:

Related case study:

Two related articles in the cluster:

B.3 Letter to the client requesting results

Send this two or three months after launch. Keep it short, ask specific questions, and explain the benefit to the recipient.

Subject: Three questions about [project name]

Hello [name],

It has been [time] since [what was launched] went live. I would like to describe the project on my website and mention [company] in an article for readers in [industry].

To make it specific, could you answer three questions? One or two sentences each will be enough.

1. What was the task, and what had prevented a solution?
2. What changed after launch? Any figures, even approximate ones, would make the account more useful.
3. What was unexpectedly good or bad about the work?

I will send the text for approval before publication. Nothing will be published without your confirmation. If anything cannot be mentioned, please let me know; I can prepare an anonymized version for your review.

Thank you,

[signature]

Three notes: the third question elicits details that distinguish a genuine testimonial from an invented one. Approval before publication addresses the recipient's main concern; offering anonymization addresses another.

B.4 Request for permission to publish

Subject: Permission to publish [project name]

Hello [name],

Please confirm whether I may publish a case study of [project name] on my website and professional profiles.

The publication will include:

- the task and solution;
- images: [list];
- results: [list];
- the company's name: yes / no.

It will exclude:

- [list: internal documents, user data, etc.]

Earliest publication: [date or “after public launch”]

Please let me know who is authorized to approve publication and what form of permission our contract requires.

Thank you,

[signature]

B.5 Scenario calculation worksheet

Goal and definition of a qualified inquiry:

Period and channel:

Known first-party data:

Assumptions (CTR, conversion, contracts):

Cautious / base / favorable scenario:

Source or rationale for each assumption:

Available hours and contingency:

What the next action will test:

Review date and conditions:

Unknowns not yet assigned a number:

B.6 Source verification record

Claim and section:

Primary source and link:

Publication date and date checked:

Sample, period, geography:

Unit of analysis and denominator:

What the source supports:

What cannot be inferred from it:

Decision: retain / qualify / remove:

C · Structured data reference

The table lists useful site-model fields, not universally required Schema.org properties. Google’s requirements depend on the specific search feature. Eligibility for a rich result does not guarantee display. [64]

Type	Purpose	Starting fields
Person	author or specialist	name, url, sameAs, alternateName
Organization	organization information	name, url, logo, sameAs
Service	service description	name, provider, description, areaServed
Article	article	headline, author, datePublished, image
CreativeWork	description of the work	name, creator, about
BreadcrumbList	breadcrumb trail	itemListElement
FAQPage	questions and answers	mainEntity
Review	individual review	author, itemReviewed, reviewBody

Type	Purpose	Starting fields
AggregateRating	aggregate rating	ratingValue and rating-count information
VideoObject	video	name, thumbnailUrl, uploadDate

C.1 Requirements and search appearance

Article. Google lists no required properties for this type, but recommends providing relevant information. An empty record does little to describe an article. [standard, 37]

Organization. Structured data can help Google understand an organization and display details such as its logo. It does not guarantee a separate card, but neither should it be dismissed as having no effect on presentation. [standard, 38]

FAQPage. Google discontinued FAQ rich results on May 7, 2026. A type remaining in the vocabulary does not mean the former search feature is still supported. [standard, 19]

Review and AggregateRating. Review rich results have eligibility requirements for the reviewed entity type. Reviews of an Organization or LocalBusiness on a site it controls are ineligible for self-serving review stars. Embedding an external widget does not remove this restriction. [standard, 39]

Other types. Choosing a vocabulary type does not by itself establish eligibility for a Google rich result. Before implementing a feature, read its documentation and required fields. This table does not replace the BreadcrumbList or VideoObject specifications.

C.2 Procedure for verification

First match the structured data to visible content and validate its vocabulary in Schema Markup Validator. Then test the supported Google feature in Rich Results Test. After publication, check that the URL is accessible and renders correctly.

Populate fields from the same records used to render the page. Give each entity a stable identifier, links to related profiles, and an alternative name where needed. Identical character-for-character spelling is not required.

No rich result does not necessarily mean faulty markup. Check support for the type, eligibility, and the search engine's display choices. Waiting two or three weeks is not a substitute for those checks.

D • A catalogue of clichés

A searchable list of cliché equivalents. Part VIII, section 8.2 explains why each expression is weak and how to replace it. The source script searches Russian phrases; these translations are reading examples, not localized search patterns.

Source file: Приложения-штампы.txt. The source check script reads the same file, keeping the original book's list and its automated check aligned.

D.1 Phrases

- in today's world
- in the age of digitalization
- dynamically developing...
- integrated approach
- it is no secret that
- it is important to understand
- let's take a closer look
- worth mentioning
- plays an important role
- is a key factor
- makes it possible to significantly...

- in conclusion, I would like to...
- fast, high-quality...
- at the present time
- as everyone knows
- a wide range
- an individual approach
- in the shortest possible time
- extensive experience
- not just...
- this article
- certainly
- undoubtedly

D.2 Techniques

These patterns require reading; keyword searches will not reliably find them.

- A rhetorical question used as a transition between sections.
- An opening paragraph explaining what the chapter will be about.
- A concluding paragraph that merely repeats what was just said.
- Calling a task “simple” to imply that it is easy.
- Three generic benefits in place of one supported by a number.
- A symmetrical list whose length is chosen for how it sounds.

- A blanket caveat such as “results depend on many factors.”

D.3 How to use it

`grep -n -i -f Приложения-штампы.txt текст.md`

Run all checks with `bash Шаблоны/проверить.sh`.

E • Index

Find an answer by the question you would ask, rather than by a chapter title.

Part I

- How many pages do I need to generate inquiries? — 1.6
The arithmetic of inquiries
- How many hours does one page take? — 1.7 Page cost in hours
- When should I expect the first results? — 1.8 Time to returns and the shape of the growth curve
- When is a website unnecessary? — 1.10 When a website is not needed
- Who visits the site, and what do they check? — 1.4 Who visits the website, and what each visitor checks
- Why does informational traffic fail to bring inquiries? — 1.2
The four functions and conflicts between them
- A platform or my own website? — 1.3 Channel map: owned and rented channels

Part II

- Why is the page slow? 2.3 What can go wrong with LCP
- What do the metrics measure, and what are their thresholds? — 2.2 Core Web Vitals: What exactly is measured
- Which image format should you choose? 2.5 Images
- Why does the text not appear immediately? — 2.6 Fonts
- Are static sites really faster? — 2.7 Why static sites are fast, and when they are not

Part III

- My page is not indexed. What should I do? — 3.1 Crawling, indexing, and the rendering queue
- Which page type fits a query? — 3.3 Search intent as the basis for choosing a page type
- What is E-E-A-T, and can it be “implemented”? — 3.5 E-E-A-T: What it is and what it is not
- What if two pages compete with each other? — 3.6 Duplicate content, cannibalization, and thin content
- How does Yandex differ from Google? — 3.7 Two search algorithms

Part IV

- How should I structure page URLs? — 4.2 URL as a contract
- How should pages link to each other? — 4.4 Internal linking as a flow

- How can I preserve traffic during a migration? — 4.6 Redirect map
- Do I need structured data? — 4.7 Structured data markup
- What breaks as the site grows to three hundred pages? — 4.8 What breaks when the site reaches 300 pages
- How can filters avoid creating duplicates? — 4.1 Taxonomy and facets

Part V

- How can I model content to avoid repetitive manual edits? — 5.2 Entities, fields, and relationships

Part VI

- Why does a mockup gallery fail to win clients? — 6.1 Why a gallery of mockups does not sell
- What can I write in a case study without client figures? — 6.3 How to get numbers and what to do without them
- Where is the line between presentation and deception? — 6.4 The boundary between presentation and deception
- What sections belong in a case study? — 6.5 Examining the structure, section by section

Part VII

- Should I publish prices? — 7.2 Price: publish or not
- In what order should service-page sections appear? — 7.3 Page sections: structure and order
- How many fields should a form have? — 7.5 Forms
- What should the button say? — 7.6 Calls to action

- Do I need industry-specific pages? — 7.7 Industry pages

Part VIII

- What does an experienced practitioner's writing sound like? — 8.1 The voice of a practitioner and the voice of a copywriter
- What should replace clichés? — 8.2 A catalogue of clichés
- How should I edit an AI-written draft? — 8.4 Editing a machine-generated draft
- Dashes, quotation marks, and nonbreaking spaces — 8.6 Text typography

Part IX

- How do I build a keyword map? — 9.3 Building a keyword set from seed terms
- Which search-volume figure should I use in calculations? — 9.2 Search-volume sources and their errors
- How do I cluster queries? — 9.4 Clustering
- What should I inspect in search results before writing? — 9.5 Reading search results before writing
- How should I set priorities? — 9.6 Prioritizing along three dimensions

Part X

- What belongs on the About page? — 10.1 The About page as evidence
- How should I ask for a testimonial? — 10.2 Reviews
- How to get links - 10.4 External links: what works

- What is the risk of buying links - 10.5 Black-hat tactics and the cost of recovery

Part XI

- What publishing pace is realistic? — 11.2 Realistic pace
- What are the steps in content production? — 11.3 The production pipeline
- When is a page ready? — 11.5 Definition of Done
- Should I update old content or write new content? — 11.6 Updating published content
- How do I maintain a content register? — 11.7 The content register

Part XII

- Website builder, CMS, or static site generator? — 12.1 Four types of platform and their implications
- How should I choose a technology stack? — 12.2 Selection criteria
- Where should I host the site? — 12.4 Hosting, delivery, protection
- How should I set up backups? — 12.6 Backup

Part XIII

- What line length and font size should I use? — 13.1 Typography as the foundation
- How do I design dark mode? — 13.2 Design tokens and themes

- Which accessibility requirements apply? — 13.5

Accessibility

- What is the performance cost of animation? — 13.6

Animation and its cost

Part XIV

- Which metrics are useful, and which are not? — 14.1

Metrics for decisions and metrics for reports

- Where did an inquiry come from? — 14.3 Inquiry attribution

- Which metrics can mislead? — 14.4 Misleading metrics

• Is the site growing normally? 14.5 Expected trends and warning signs

Part XV

• May I publish work done for a client? — 15.2 Publishing client work

• How should I handle generative tools? — 15.3 Generative tools

• How should I handle personal data from forms? — 15.4 Personal data in forms

Worked examples in Edition 2

Task	Where to read
Start a site with limited time	1.6, Appendix H.1
Assess inquiry quality	7.5, 14.3–14.5, Appendix H.2
Decide whether to continue after a failed launch	1.8, 14.5–14.6, Appendix H.3
Distinguish assumptions from research evidence	Introduction, Appendix B.5–B.6

F • Glossary

Each term is defined when first introduced and collected here. Its meaning remains consistent throughout the book. If industry usage has two meanings, this glossary states which is used and which is excluded.

Record format: **term** — definition. (Part where introduced)

75th percentile — a value at or below which 75% of observations fall. Assess each Core Web Vitals metric separately. (II, 2.2)

E-E-A-T — a framework for human quality raters, not a single ranking factor. Raters do not directly determine page rankings, and their ratings are not fed directly into ranking algorithms. (III, 3.5)

Step output — a production-stage result required before the next stage begins. It provides a concrete readiness check. (XI, 11.3)

Broad, phrase, and exact search volume — query-volume estimates narrowed by recorded operators. In Wordstat, quotes fix word count, ! fixes word form, and [] fixes order; these are distinct controls, with support varying by report tab. (IX, 9.2)

Blocking resource — a resource that must be fetched before initial rendering or further document parsing.

Examples include blocking stylesheets and parser-blocking scripts without deferred execution. (II, 2.1)

Crawl budget — the URLs a crawler is able and willing to crawl. It is determined by crawl demand and available capacity, not by multiplying two figures. (III, 3.1)

Vertical rhythm — spacing between elements based on multiples of the body-text line height. More space precedes a heading than follows it, connecting it to the text it introduces. (XIII, 13.3)

Time to first byte — TTFB: the time from navigation start to the first response byte, including preceding network stages. A recommended good value is at most 0.8 s. It is not a Core Web Vital. (II, 2.1)

Click depth — the number of link clicks from a chosen entry point to a page. It is not the number of folders in the URL. (IV, 4.2, 4.4)

Long task — uninterrupted main-thread work lasting more than 50 ms, during which the browser cannot respond promptly to user input. A major cause of poor responsiveness. (II, 2.4)

Case-study narrative structure — starting situation → constraint → options → choice → trade-off → outcome. The constraint explains why the choice was made. (VI, 6.2)

Duplicate pages — multiple URLs with nearly identical content. The search engine chooses a canonical URL, which may differ from your preference. A canonical declaration is a hint, not a command. (III, 3.6)

Production capacity — pages producible in a period: available hours divided by average hours per page. (I, 1.7)

Hard and soft clustering — this book uses pairwise overlap for hard clustering and comparison with a central query for soft clustering. Implementations vary; these are workflow definitions, not search-engine rules. (IX, 9.4)

Earned mention — a link in material an outlet independently chose to publish. Its distinction from a purchased placement lies in its origin as well as quality: it cannot be bought in bulk. (X, 10.4)

Measurable product property — something the provider can measure without client analytics, such as field count, journey steps, or loading time. It supplies evidence when business metrics are unavailable, without proving the same business outcome. (VI, 6.3)

Channel inertia — the lag between effort and returns, and the speed at which returns fade after effort stops. (I, 1.3)

Informational query — a query whose main purpose is to obtain information. It does not rule out readiness to buy. (I, 1.2)

Keyword cannibalization — harmful overlap between page purposes. Two URLs appearing for one query do not alone prove harm; diagnose before merging them. (III, 3.6)

Content cluster — pages on one topic connected by internal links and supporting one another in search. (I, 1.8; detailed in IV, 4.3)

Query cluster — queries sharing one intent and served by one page. A keyword map is organized around clusters rather than individual queries. (IX, 9.1)

Commercial query — a query aimed at finding a provider or buying a service. (I, 1.2)

Critical rendering path — the steps needed for initial display: document parsing, style-tree construction, layout calculation, and painting. (II, 2.1)

— **request-response network latency. An uncached HTTP redirect adds an exchange; cached and internal redirects can behave differently.** — time for a request to reach the server and a response to return. A measure of network latency; each redirect adds another round trip. (II, 2.1)

False symmetry — an AI-draft defect in which list length is chosen for rhythm rather than substance. Ask whether another item exists and why it was omitted. (VIII, 8.4)

Line length — characters per line, assessed with the typeface, size, and device. A range of 50–70 is a layout starting point, not a universal limit. (XIII, 13.1)

Response Metric (INP) — interaction latency across a visit; a good value is at most 200 ms. (II, 2.4)

Largest Contentful Paint (LCP) — time from loading start until the largest eligible element in the viewport is painted. A good value is at most 2.5 s. (II, 2.2)

Decision metrics — a metric that changes a decision. Ask what you would do differently if it doubled or halved. If both answers are “nothing,” it is not useful for that decision. (XIV, 14.1)

Cumulative Layout Shift (CLS) — the largest sum of unexpected layout-shift scores in a session window lasting at most five seconds, with less than one second between shifts. A good value is at most 0.1. (II, 2.4)

Fallback font metrics — font descriptors that match a fallback font's scale and proportions to reduce text reflow when the web font replaces it. (II, 2.6)

Data minimization — limiting personal data to what is necessary for a specified lawful purpose. Form length and conversion are separate design questions. (XV, 15.4)

Exclusion list — a reusable list of irrelevant query patterns. Check false exclusions and measure its effectiveness; no universal automatic-removal percentage applies. (IX, 9.3)

Relationship cardinality — how many instances of one entity correspond to an instance of another. Modeling many-to-many as one-to-many can be a costly content-model error. (V, 5.2)

Intent. — the task behind a query. Search results help identify a suitable page type but do not absolutely prohibit other formats. (III, 3.3)

Zero-click search — a search without a click to an external site. Definitions and denominators vary by study; no universal percentage is assumed. (I, 1.9)

Anonymization — removing or generalizing identifying information. Publication permission must still be agreed separately; anonymization does not override contractual restrictions. (VI, 6.3; XV)

Inquiry — a received message or conversation from a prospective client. A messaging-link click does not prove an inquiry occurred. Measure the share leading to contracts separately. (I, 1.6; XIV)

Crawling, indexing, ranking — three distinct processes with separate failure modes: fetching a document, processing

and storing it in the index, then selecting and ordering results for a query. (III, 3.1)

Elimination from consideration — a commercial-page visitor's decision not to shortlist the provider. The page succeeds by surviving this screening, so removing reasons to leave can matter more than persuasion. (VII, 7.1)

Publication embargo — a time restriction on showing work, such as after public launch or an agreed date. It can reconcile interests without prohibiting publication altogether. (XV, 15.1)

Rendering queue — pages waiting for renderer processing. A queue alone does not explain missing text or indexing. (III, 3.1)

Enumeration — a closed set of allowed classification values. Free text fragments filtering: "SaaS," "saas," and "SaaS platform" become separate values. (V, 5.3)

The half-life of content — a descriptive measure of how a chosen return metric declines over time. It depends on the channel and distribution; no single lifetime applies to every post. (I, 1.3)

Statement-style subheading — a heading that states a point rather than names a topic. "Publish your prices" communicates a position; "Pricing" does not. (VIII, 8.5)

Font subset — a font file retaining needed characters and features. Size savings depend on the font, scripts, and output format and must be measured. (II, 2.6)

Impression — a result appearance counted under the reporting system's rules. For comparisons, specify the report source, queries, device, and period. (I, 1.8; XIV)

Field and laboratory data — field data comes from real users; lab data comes from a controlled test. Core Web Vitals field assessments describe real-user experience, whereas one lab run is a diagnostic sample. Do not treat a single lab score as a ranking prediction. (II, 2.2)

Attribution coverage — the share of inquiries with a known source. Unknown sources stay separate, without a predetermined “normal loss” rate. (XIV, 14.3)

Rights in the work and confidentiality — separate issues: confidentiality restricts disclosure; rights determine who may use the work. Neither automatically grants portfolio publication rights, which require separate agreement. (XV, 15.1)

Quotas rule — limiting near-term priorities to available capacity. A fixed fraction of the list can be a workflow setting, not an industry standard. (IX, 9.7)

Rule of stopping — treat an AI draft as raw material. If editing leaves no claim the author can defend in conversation, there is nothing ready to publish. (VIII, 8.4)

The Most Expensive Restriction Rule — choose a technology stack around its most costly constraint. Static-site speed has little value if the person publishing cannot use its required text-file workflow. (XII, 12.2)

— test recovery into a clean environment. Choose frequency from acceptable data loss and system changes; six months is only a proposed starting interval. — a backup is unproven until restored. Every six months, restore into an empty environment and record missing items. (XII, 12.6)

Word-substitution test for generic copy — an editorial method for detecting generic wording. It is not a technical duplicate-content test and does not replace comparing page purposes. (VII, 7.7)

Firsthand-experience test — ask whether a paragraph could have been written without doing the work. If so, it adds little firsthand evidence and may repeat what readers have seen in ten other texts. (VIII, 8.1)

Verifiable claim — a claim that can be disproved. “Worked there from 2019 to 2022” is checkable; “completed many projects” offers no specific evidence to test. (X, 10.1)

Time-separation — a solo practitioner’s technique: draft and edit on different days. The distance helps reveal clichés that attention missed on the day of writing. (XI, 11.4)

The role of the visitor — a task-based classification rather than a demographic one: problem-solver, shortlister, peer, or verifier. (I, 1.4)

Manual action and algorithmic devaluation — a manual action produces a Search Console notice; automated systems can neutralize link-spam credit. Neither a decline nor a missing notice establishes the cause by itself. (X, 10.5)

Orphan page — a page without incoming internal links. Search engines may know it through a sitemap or backlinks, but its internal navigation path still needs attention. (IV, 4.4)

Page consolidation — combining several pages’ useful content into the strongest page and redirecting the others to it. This can preserve accumulated signals that simple deletion would lose. (III, 3.6)

Sliding window — a moving aggregation period. CrUX API and PageSpeed Insights use rolling 28-day windows; CrUX BigQuery data is monthly. New observations can change the rolling summary before all older observations have left it. (II, 2.2)

Deliverable-based scope — listing what the client will receive rather than only methodology stages, which do not reveal the final deliverables. (VII, 7.3)

Status "rejected" — an idea recorded as rejected with a reason. It saves repeated evaluation and keeps a register from becoming a wish list. (XI, 11.7)

Owning side of a relationship — the entity whose record explicitly stores the relationship; the reverse relationship is derived by a query. Two manually maintained lists tend to diverge. (V, 5.2)

Structured data markup — a machine-readable description of content. Vocabulary validity and eligibility for a Google rich result are separate checks; display is not guaranteed. (IV, 4.7)

Entity — the object described, such as a person, organization, project, or place. Identifiers, profiles, and name variants connect mentions; different spellings do not alone imply separate entities. (III, 3.2)

Diagnostic table — possible checks suggested by observations. It does not establish a single cause for a metric change. (XIV, 14.5)

Taxonomy — a hierarchy of sections or categories, which can coexist with facets. (IV, 4.1)

Credence good — a product or service whose quality a buyer cannot fully assess even after purchase. Project-based design services can have this property, making risk-reducing evidence important. (I, 1.1)

Design token — a named value used instead of a literal colour or size, allowing palette changes in one place: the same single-source principle used in the content model. (XIII, 13.2)

Thin content — content that does not sufficiently help the reader's task. It is not defined by word count or a requirement that every fact be new. (III, 3.6)

Dead-end page — a page offering no clear next step for the task. Observe what visitors actually do afterward. (I, 1.5)

Facets — independent classification dimensions for selecting the same item in different ways. Their combinations multiply URLs and can cause excessive crawling. (IV, 4.1)

Hub and spoke structure — a topic overview and standalone pieces joined by useful links. No verified minimum page count applies. (IV, 4.3)

Trade-off — what was sacrificed to make a decision. Naming it demonstrates understanding of consequences; omitting it may suggest those consequences were overlooked. (VI, 6.2)

Cost of change — effort to change design or structure sitewide. Editing one template versus twenty-five pages can distinguish platforms more clearly than loading speed. (XII, 12.1)

Redirect chain — successive redirects before the destination. Minimize unnecessary hops; the two-hop review

threshold in this workflow is not a search-engine limit. (IV, 4.6)

Qualified inquiry — an inquiry whose task matches the service and working terms. Define the criteria before counting. (I, 1.6; XIV)

Illustrative example — a modeled situation with hypothetical data, not a real client case or evidence of effectiveness. (Appendix H)

G · Bibliography

References rechecked on September 10, 2026. Original numbering is retained; additional verification sources follow. The accompanying report distinguishes supported claims, qualified interpretations, proposed procedures, and hypothetical calculations. This audit does not establish universal effectiveness or legal compliance for an individual project.

7. How Long Does It Take to Rank in Google?. [primary source] Open source Checked: 10 September 2026. Ahrefs, 15 May 2025. Two samples: 1 million URLs first observed in September 2023, and 2 million URLs created in October 2023 with non-empty English content. The proportions reaching the top 10 within a year were 1.74% and 6.11%. These are not the probability of success of a particular project. Cited in: 1.8

<https://ahrefs.com/blog/how-long-does-it-take-to-rank-in-google-and-how-old-are-top-ranking-pages/>

9. Web Vitals. [documentation] Open source Checked: September 10, 2026. $LCP \leq 2.5$ s; $INP \leq 200$ ms; $CLS \leq 0.1$. Assess the 75th percentile of each metric separately and segment by device. Cited in: 2.2

<https://web.dev/articles/vitals>

10. Time to First Byte. [documentation] Open source Checked: September 10, 2026. TTFB includes stages from navigation start to the first response byte. A good benchmark is 0.8 s; it is not a Core Web Vital. Cited in: 2.1

<https://web.dev/articles/ttfb>

11. **LCP breakdown.** [documentation] Open source Checked: September 10, 2026. The four LCP components. Their timing proportions are diagnostic guidance, not mandatory allocations for every page.Cited in: 2.3

<https://developer.chrome.com/docs/performance/insights/lcp-breakdown>

13. **Crawl Budget Management.** [documentation] Open source Checked: September 10, 2026. Crawl demand and capacity jointly determine crawling. The documentation does not define their arithmetic product. The shared limit concerns Google crawlers in the described infrastructure, not all search engines worldwide.Cited in: 3.1

<https://developers.google.com/crawling/docs/crawl-budget>

15. **Measuring Google AI Overviews: Activation, Source Quality, Claim Fidelity, and Publisher Impact.** [primary source] Open source Checked: September 10, 2026. Xu, Iqbal, Montgomery, arXiv v1, May 13, 2026. The quoted figures belong to that version and its selected query sample. Its record says Under Review; the author's publication list also lists IMC 2026. Automated source-consistency judgments are not a direct truth-error rate.Cited in: 1.9 · 3.4

<https://arxiv.org/abs/2605.14021>

16. **Creating helpful, reliable, people-first content.** [documentation] Open source Checked: September 10, 2026. No preferred word count. E-E-A-T is not a single ranking factor; quality ratings do not directly set page rankings.Cited in: 3.5 · 3.6

<https://developers.google.com/search/docs/fundamentals/creating-helpful-content>

18. **Managing crawling of faceted navigation URLs.** [documentation] Open source Checked: September 10, 2026. Faceted URL management: limit crawling of unnecessary combinations while keeping useful ones accessible. Crawl blocking via robots.txt and removal from the index are different operations.Cited in: 4.1 · 4.5

<https://developers.google.com/crawling/docs/faceted-navigation>

19. Google Search documentation updates: FAQ rich results.

[documentation] Open source Checked: September 10, 2026. The May 8, 2026 log entry states that FAQ rich results stopped on May 7. Feature documentation was later removed. The date was checked against the change log rather than the old type page.Cited in: 4.7 · C.1

<https://developers.google.com/search/updates>

20. Symbols and operators — Yandex Direct. [documentation] Open

source Checked: September 10, 2026. Operators for refining query wording, including Wordstat. Fixing word form does not fix word order. Record operators alongside volume estimates.Cited in: 9.2

<https://yandex.ru/support/direct/ru/keywords/symbols-and-operators>

23. Pricing information gives B2B sites a competitive advantage.

[primary source] Open source Checked: September 10, 2026. Nielsen Norman Group observations on finding prices. Not a universal estimate of conversion uplift for service sites.Cited in: 7.2

<https://www.nngroup.com/articles/show-price/>

26. Spam policies for Google web search. [documentation] Open

source Checked: September 10, 2026. Link spam and scaled content abuse policies. No safe threshold of 30% guest posts is specified.Cited in: 10.5 · 3.6 · 7.7

<https://developers.google.com/search/docs/essentials/spam-policies>

29. Web Content Accessibility Guidelines 2.2. [documentation] Open

source Checked: September 10, 2026. W3C. Contrast is only part of conformance. AA requires all applicable A and AA criteria; a short checklist cannot replace a full assessment.Cited in: 13.5 · 13.6

<https://www.w3.org/TR/WCAG22/>

30. Cumulative Layout Shift. [documentation] Open source Checked:

September 10, 2026. Largest shift-session window, normalization of impact and distance fractions, and exclusion of shifts within 500 ms of qualifying input.Cited in: 2.4

<https://web.dev/articles/cls>

31. **Interaction to Next Paint.** [documentation] Open source Checked: September 10, 2026. Interaction latency over a visit, selection rules, and thresholds.Cited in: 2.4

<https://web.dev/articles/inp>

32. **Link best practices for Google.** [documentation] Open source Checked: September 10, 2026. Crawlable links and clear anchor text. Important pages need incoming internal links; no universal ideal count exists.Cited in: 4.4

<https://developers.google.com/search/docs/crawling-indexing/links-crawlable>

33. **GA4: Engagement rate and bounce rate.** [documentation] Open source Checked: September 10, 2026. Bounce rate is the share of non-engaged sessions. Duration, a key event, or view count determine engagement.Cited in: 14.4

<https://support.google.com/analytics/answer/12195621?hl=en>

34. **GA4: User engagement.** [documentation] Open source Checked: September 10, 2026. Engagement time counts page focus or an app's foreground time, not merely the interval between views.Cited in: 14.4

<https://support.google.com/analytics/answer/11109416?hl=en>

35. **How to specify a canonical URL.** [documentation] Open source Checked: September 10, 2026. Canonical-selection signals. Declaring a canonical does not guarantee that the search engine selects it.Cited in: 3.6 · 4.5

<https://developers.google.com/search/docs/crawling-indexing/consolidate-duplicate-urls>

36. **Pagination and incremental page loading.** [documentation] Open source Checked: September 10, 2026. Separate URLs, links, and self-referencing canonicals for paginated pages. Changing the address bar without accessible content is insufficient.Cited in: 4.5

<https://developers.google.com/search/docs/specialty/ecommerce/pagination-and-incremental-page-loading>

37. **Article structured data.** [documentation] Open source Checked: September 10, 2026. Google lists no required Article properties but recommends applicable fields.Cited in: 4.7 · C.1

<https://developers.google.com/search/docs/appearance/structured-data/article>

38. **Organization structured data.** [documentation] Open source Checked: September 10, 2026. Organization and logo information. Structured data does not guarantee a particular search appearance.Cited in: 4.7 · C.1

<https://developers.google.com/search/docs/appearance/structured-data/organization>

39. **Review snippet structured data.** [documentation] Open source Checked: September 10, 2026. Supported types and self-serving review restrictions for Organization and LocalBusiness.Cited in: C.1

<https://developers.google.com/search/docs/appearance/structured-data/review-snippet>

40. **AI features and your website.** [documentation] Open source Checked: September 10, 2026. No special markup is needed for AI search features; basic search requirements apply.Cited in: 1.9

<https://developers.google.com/search/docs/appearance/ai-features>

41. Understand JavaScript SEO basics. [documentation] Open source
Checked: September 10, 2026. JavaScript crawling and rendering.
Distinguish loading or execution errors from waiting for processing. Cited
in: 3.1

<https://developers.google.com/search/docs/crawling-indexing/javascript/javascript-seo-basics>

[42] Google: Understanding news topic authority Checked
September 10, 2026. Open source

<https://developers.google.com/search/blog/2023/05/understanding-news-topic-authority>

[43] web.dev: Optimize Largest Contentful Paint Checked September
10, 2026. Open source

<https://web.dev/articles/optimize-lcp>

[44] MDN: Variable fonts Checked September 10, 2026. Open source

https://developer.mozilla.org/en-US/docs/Web/CSS/Guides/Fonts/Variable_fonts

[45] Chrome: CrUX methodology Checked September 10, 2026. Open
source

<https://developer.chrome.com/docs/crux/methodology/>

[46] IETF RFC 9000: QUIC Checked September 10, 2026. Open source

<https://www.rfc-editor.org/rfc/rfc9000.html>

[47] Google: Site moves with URL changes Checked September 10,
2026. Open source

<https://developers.google.com/search/docs/crawling-indexing/site-move-with-url-changes>

[48] W3C: Understanding Focus Order Checked September 10,
2026. Open source

<https://www.w3.org/WAI/WCAG22/Understanding/focus-order.html>

[49] US Copyright Office: AI report, Part 2Checked September 10, 2026.Open source

<https://www.copyright.gov/newsnet/2025/1060.html>

[50] EU Regulation 2016/679: Articles 3, 5, 6, 13, 17, 28Checked September 10, 2026.Open source

<https://eur-lex.europa.eu/eli/reg/2016/679/oj/eng>

[51] Google Analytics: Avoid sending PIIChecked September 10, 2026.Open source

<https://support.google.com/analytics/answer/6366371?hl=en>

[52] CISA: Data Backup OptionsChecked September 10, 2026.Open source

https://www.cisa.gov/sites/default/files/publications/data_backup_options.pdf

[53] WIPO: Copyright FAQChecked September 10, 2026.Open source

<https://www.wipo.int/en/web/copyright/faq-copyright>

[54] Credence goods in the literature: What the past fifteen years have taught us about fraud, incentives, and the role of institutionsChecked September 10, 2026.Open source

<https://www2.uibk.ac.at/downloads/c4041030/wpaper/2020-01.pdf>

[55] MDN: font-displayChecked September 10, 2026.Open source

<https://developer.mozilla.org/en-US/docs/Web/CSS/Reference/At-rules/%40font-face/font-display>

[56] MDN: size-adjustChecked September 10, 2026.Open source

<https://developer.mozilla.org/en-US/docs/Web/CSS/Reference/At-rules/%40font-face/size-adjust>

[57] Google: Spam updatesChecked September 10, 2026.Open source
<https://developers.google.com/search/docs/appearance/spam-updates>

[58] Google: SEO Starter GuideChecked September 10, 2026.Open source
<https://developers.google.com/search/docs/fundamentals/seo-starter-guide>

[59] Google: Disavow linksChecked September 10, 2026.Open source
<https://support.google.com/webmasters/answer/2648487?hl=en>

[60] MDN: How browsers workChecked September 10, 2026.Open source
https://developer.mozilla.org/en-US/docs/Web/Performance/Guides/How_browsers_work

[61] Google: Localized versionsChecked September 10, 2026.Open source
<https://developers.google.com/search/docs/crawling-indexing/localized-versions>

[62] Google: Search Console and AnalyticsChecked September 10, 2026.Open source
<https://developers.google.com/search/docs/monitor-debug/google-analytics-search-console>

[63] Google: Search snippetsChecked September 10, 2026.Open source
<https://developers.google.com/search/docs/appearance/snippet>

[64] Schema.org: Service, FAQPage, AggregateRatingChecked September 10, 2026.Open source
<https://schema.org/docs/schemas.html>

[65] Xu, Iqbal, Montgomery: arXiv v1 full textChecked September 10, 2026.Open source

<https://arxiv.org/html/2605.14021v1>

[66] Umar Iqbal: PublicationsChecked September 10, 2026.Open source

<https://umariqbal.com/pages/publications.html>

[67] Chrome: CrUX ToolsChecked September 10, 2026.Open source

<https://developer.chrome.com/docs/crux/methodology/tools?hl=en>

[68] European Commission: Principles of personal data processing under the GDPRChecked September 10, 2026.Open source

https://commission.europa.eu/law/law-topic/data-protection/information-business-and-organisations/principles-gdpr_en

H • Three worked examples

Every project, figure, schedule, and outcome in this appendix is hypothetical and created for teaching. These examples do not describe real clients or establish industry norms. The tables are the single source for the shorter discussions throughout the chapters. [example]

H.1 A specialist’s new website

The starting point. An interface designer has two projects approved for portfolio publication and eight hours a week to work on the site. They want to test whether search can bring inquiries but have not yet measured demand.

Constraint. The twelve-week budget is 96 hours while client work continues. The first version, checks, and contingency must fit within it, so industry-specific pages and a regular blog are deferred.

Work.	Hours
Refine the offer and research queries	10
Structure and record fields	6
Page of one service	12

Work.	Hours
Two case studies from existing material	24
Build, shared pages, images	12
Contact options, tracking, publication checks	8
Contingency for revisions	24
Total	96

This estimate is lower than the manual-production example in Part I because case-study material is already prepared and a template handles repeated layout work. That is a different starting condition, not proof that the earlier estimates are always excessive.

Site architecture. The home page links to the service and two case studies. The service links to both cases; each case links back to the service and offers contact. About and Contact pages complete the provider check. The three main content pages are not the site's entire technical page set.

Content model. Each case has an ID, title, author role, authorized materials, service links, and an outcome. A service's case list is derived from those relationships. Although there is one service now, the relationship allows multiple services for future projects that span different work.

Proof. One case reduces the inquiry journey from seven steps to four, showing diagrams and the designer's contribution. There is no sales data, so no conversion uplift is claimed. Usability research, if not yet performed, is identified as a proposed next check rather than completed work.

Observations for the week-twelve review. The last four weeks brought 80 search visits, including 30 to the service page. One genuine submission qualified as an inquiry; no contract was signed. Test submissions are excluded. These figures establish neither a stable conversion rate nor a promise of payback.

Next decision. The site is complete as an introduction and verification tool. The owner allocates another 16 hours to refine the service page using real questions and monitor search queries. New industry pages require a distinct task and sufficient material. This is a bounded continuation of the experiment, not an automatic annual content plan.

H.2 There is traffic, there are almost no useful inquiries

The starting point. A small service provider has a site with services and articles. Overall traffic appears sufficient, but few contacts are a good fit. Diagnosis aims to locate where suitable inquiries are being lost.

Over comparable one-month periods	Prior to changes	After the changes
Total organic-search visits	1 000	1 050
Of which, visits to service pages	120	210
Other search visits	880	840
Successful submissions from service pages	4	7

Over comparable one-month periods	Prior to changes	After the changes
Qualified inquiries from those submissions	1	2
Direct qualified inquiry without a source	0	1
Total qualified inquiries	1	3

Checks before changes. All four submissions arrived, but three concerned unsuitable work. This points to the offer and screening rather than a delivery fault. Messaging-link clicks are not added to the inquiry count.

Selected changes. The service page clarifies deliverables and the price range, adds a relevant case study, and explains the next step. The form is tested on a phone; optional questions move to later correspondence. Related articles link to the service. The 20-hour budget assigns six hours to diagnosis, eight to content, four to the form, and two to checks and reporting.

Trade-off. No new articles are published that month. Coverage grows more slowly, but limited time improves an existing contact journey. Clearer prices may deter some visitors; this is acceptable if fewer unsuitable conversations result.

How do you count? For form attribution, compare qualified inquiries per service-page visit: 1/120 and 2/210, about 0.83% and 0.95%. Three inquiries divided by 210 equals 1.43%, but is not the service-page conversion rate because

the third source is unknown. Count it in business totals, not form attribution.

What we know. Submissions and qualified inquiries increased after the changes. **What remains unknown.**

Whether the increase came from the form, copy, visitor mix, or ordinary variation. Counts are small and several changes were made together. This comparison does not establish a causal conversion uplift.

Next decision. Keep the working journey and continue tracking until the next review. Monitor the proportion of suitable tasks, contracts, and time spent responding. If inquiries grow but remain unsuitable, revisit the offer. If suitable prospects do not sign, examine the sales stage.

H.3 Unsuccessful launch and loss limitation

The starting point. In 120 hours, the owner published 60 overview pages on broad topics chosen by broad search volume. The service page describes a process without clear deliverables, and the form shows success before delivery is confirmed.

Observations for the first eight weeks. The first eight weeks brought 240 search visits, including 12 to the service page, and no confirmed inquiries. A delivery test exposed a fault. This does not mean every click was a lost client: the undelivered messages' contents are unknown.

Separate the possible causes. Article topics are only weakly related to the service. Eighteen pages overlap in purpose and need comparison. The contact journey is

technically unreliable. Yet eight weeks cannot prove that search will never work; the observation period is too short.

What stops. Further mass publication. The 120 hours already spent are no reason to continue the same plan. Compare the next step’s future cost with its possible benefit.

Bounded repair work	Hours
Fix delivery and test error messages	4
Compare queries, pages, and indexability	6
Rework one service page	6
Set up outcome checks and a log	4
Total	20

Working with similar pages. For each of the eighteen pages, check purpose, useful sections, queries, and links. Sharing a topic does not justify deleting every URL. Create a redirect map for merged pages and retain useful standalone pages.

After repairs. All three test submissions arrived and are excluded from statistics. The following month brought 180 search visits, including 24 to the service page, one genuine submission, and zero qualified inquiries. Comparing 180 in one month with 240 over eight weeks does not establish a decline: the periods differ.

Decision. The technical repair is complete, but commercial returns remain unproven. The owner stops expanding the blog and keeps a small site for prospects arriving through referrals to verify the work. Search testing can resume with a

clearer service and verified demand. This allocates time; it does not prove SEO is useless.

H.4 What to apply to your own project

From A, take a bounded first-version budget and verifiable material. From B, take separate counts for submissions, qualified inquiries, and unknown sources. From C, distinguish technical faults, commercial hypotheses, and the decision to stop spending.

Do not treat the tables' hours, visits, or rates as norms. Record your own baseline, assumptions, observation period, and review conditions first. A successful delivery test proves delivery. More inquiries are an observation. The causal effect of a particular change needs its own study.