*48PEGASUS conducts its own research and analysis for public review. Explore selected articles on technology, visual documentation and the methods behind reliable information.

The Agent-Readable Website

How to test whether AI assistants can identify your business, services, and next steps

Co-authored by Jonathan G. Browning, Founder & CEO, 48PEGASUS | 48PEGASUS Research

© 2026 48PEGASUS. All rights reserved. | info@48pegasus.com | www.48pegasus.com/research

INDUSTRY TECHNOLOGY

AI RESEARCH

Audience: business owners, web teams, and operators.

Decision: can an assistant accurately guide a prospective customer?

EXECUTIVE SUMMARY

A website passes an agent-readability test when a web-enabled assistant can correctly identify the business, explain what it offers and where, and point a customer to a verified next step. A brand mention alone does not meet that standard. An answer can name the right business and still invent a price, confuse service areas, cite the wrong company, or send the customer to an obsolete intake form.

This paper separates three questions: can the assistant read a page when given its address, can it find the right business without that address, and can it guide a customer to the right next step? The protocol uses a dated register of verified facts, five representative prompts, repeated clean-session runs, and independent checks of each material claim and destination. Report known-site tests separately from unbranded discovery. Log omissions, unsupported assertions, broken destinations, and appropriate abstentions. Treat a fabricated price or wrong payment destination as a critical failure even if the overall score improves.

Start by making the official pages accurate, accessible, and internally consistent. Test the published site, not a preview. Then use controlled retesting to learn where changes worked. Structured data and optional machine-readable files can support specific uses, but none substitutes for reliable public content or guarantees assistant retrieval, citations, recommendations, or traffic.

Scope and limitations

This is an educational framework, not a new model benchmark or a recommendation to deploy unsupervised AI. Models can give fluent but incorrect answers, reflect bias, and change across versions or repeated runs. Results from cited studies apply to their tested conditions. Verify important claims and obtain qualified human review before real-world use.

Hey, AI, Read This

Test identity, services, and the next step separately.

Check each claim and destination against first-party evidence.

Sources and estimates may contain errors or become outdated. Verify independently and conduct your own research before acting. Provided as is, without warranties. Noncommercial sharing only; see back-page permissions. Third-party works retain their own rights.

© 2026 48PEGASUS. All rights reserved. | info@48pegasus.com | www.48pegasus.com/research

*48PEGASUS conducts its own research and analysis for public review. Explore selected articles on technology, visual documentation and the methods behind reliable information.

READING GUIDE

Section start pages. Continuations remain within their section.

Contents

03 WHAT READABILITY MEANS
04 BUILD THE PUBLIC SOURCE OF TRUTH
05 Exhibit 1 Publication to verified next step
06 RUN THE ASSISTANT TEST
07 Exhibit 2 Test classify fix and monitor
08 PROMPTS AND SCORECARD
09 INTERPRET THE EVIDENCE
10 FIX AND RETEST
11 ISOLATE THE FAILURE
12 GOVERN WHAT YOU EXPOSE
13 TEST THREE CUSTOMER JOURNEYS
14 CONNECT READABILITY TO BUSINESS OUTCOMES
15 OPERATE A DURABLE PUBLISHING SYSTEM
16 APPLY THE FIELD TEMPLATES
17 Applied case: a multidivision technology firm
18 References
19 Authors, Disclaimer and Notices
20 AI disclosure and reader guide

Sources and estimates may contain errors or become outdated. Verify independently and conduct your own research before acting. Provided as is, without warranties.
Noncommercial sharing only; see back-page permissions. Third-party works retain their own rights.

© 2026 48PEGASUS. All rights reserved. | info@48pegasus.com | www.48pegasus.com/research

READABILITY

Crawl: fetch a public page.

Index: make the page eligible for retrieval in a search system.

Ground: support an answer with evidence.

Route: point to the correct next step.

01 WHAT READABILITY MEANS

Access is not the same as comprehension

Crawlability is whether a crawler operating under its rules can fetch a page. Indexability is whether an eligible page can be included in a search index. Retrieval is whether a system actually selects a source for a given request. Grounding is whether the selected evidence supports the resulting claim. These are distinct steps. An eligible page may not be retrieved; a retrieved page may be misunderstood.

Use three test tracks

A. Direct URL: Give the assistant a canonical address and request evidence only from that domain. This isolates interpretation under known-site conditions, not search visibility. The site may be accessible to people but incompletely read by one client.

B. Branded search: Provide the business name and locality but no URL. Test whether it identifies the official business and distinguishes similarly named entities, divisions, or separate products.

C. Unbranded discovery: Ask the way a prospect would, using a service, use case, and locality but no brand name or URL. Score whether the business surfaced within this prompt set, not whether it should always be recommended.

Separate navigation from execution

A valid booking or contact destination passes a routing test only. It does not demonstrate that an agent can submit a form, verify eligibility, accept terms, or transact. Test public steps through the point preceding submission. For any later transaction test, use a separately authorized test environment and test data.

Sources and estimates may contain errors or become outdated. Verify independently and conduct your own research before acting. Provided as is, without warranties. Noncommercial sharing only; see back-page permissions. Third-party works retain their own rights.

© 2026 48PEGASUS. All rights reserved. | info@48pegasus.com | www.48pegasus.com/research

*48PEGASUS conducts its own research and analysis for public review. Explore selected articles on technology, visual documentation and the methods behind reliable information.

02 BUILD THE PUBLIC SOURCE OF TRUTH

PUBLIC DELIVERY CHECK

01 Official identity

02 Visible and current service facts

03 Stable and descriptive action links

04 Intentional crawl and index settings

Put business facts in readable HTML

On the official business page, state the public name, what the company does, whom it serves, and its actual service area. Use stable service pages with clear headings, eligibility, delivery limits, and the current next step. Write “Request a website redesign quote” rather than “Learn more.” Put critical facts in visible text, not solely in graphics, PDFs, videos, dialog boxes, or controls that require a sign-in or interaction.

Keep facts consistent across destinations

Separate the legal name, public brand, and independently branded products where useful. Distinguish an office address from a customer-facing location and a service area from physical presence. When quoting prices, distinguish fixed, starting-at, and by-quote rates. If facts change, fix conflicting pages and profiles; update old links and redirects where appropriate. Do not encode a claim in structured data that contradicts the visible page.[1,4]

Verify the public delivery path

From a logged-out, mobile-sized browser, open the final URL and inspect visible and rendered text, descriptive links and form labels. Compare this with the initial HTML to identify render-dependent content. Check the HTTP response, redirect destination, canonical URL, index controls, and bot-specific access. Link primary pages from public navigation and use the sitemap as a discovery aid. [redacted]’s URL Inspection is useful for [redacted]-specific diagnosis, not proof of access by other tools.[1,5]

Separate search access from training

[redacted] documents distinct controls for [redacted] and [redacted], and describes [redacted] as acting on user requests rather than automatically crawling. Do not infer one platform’s policy from another’s. Robots directives guide participating crawlers and do not protect sensitive content; use access controls for anything private.[2,3]

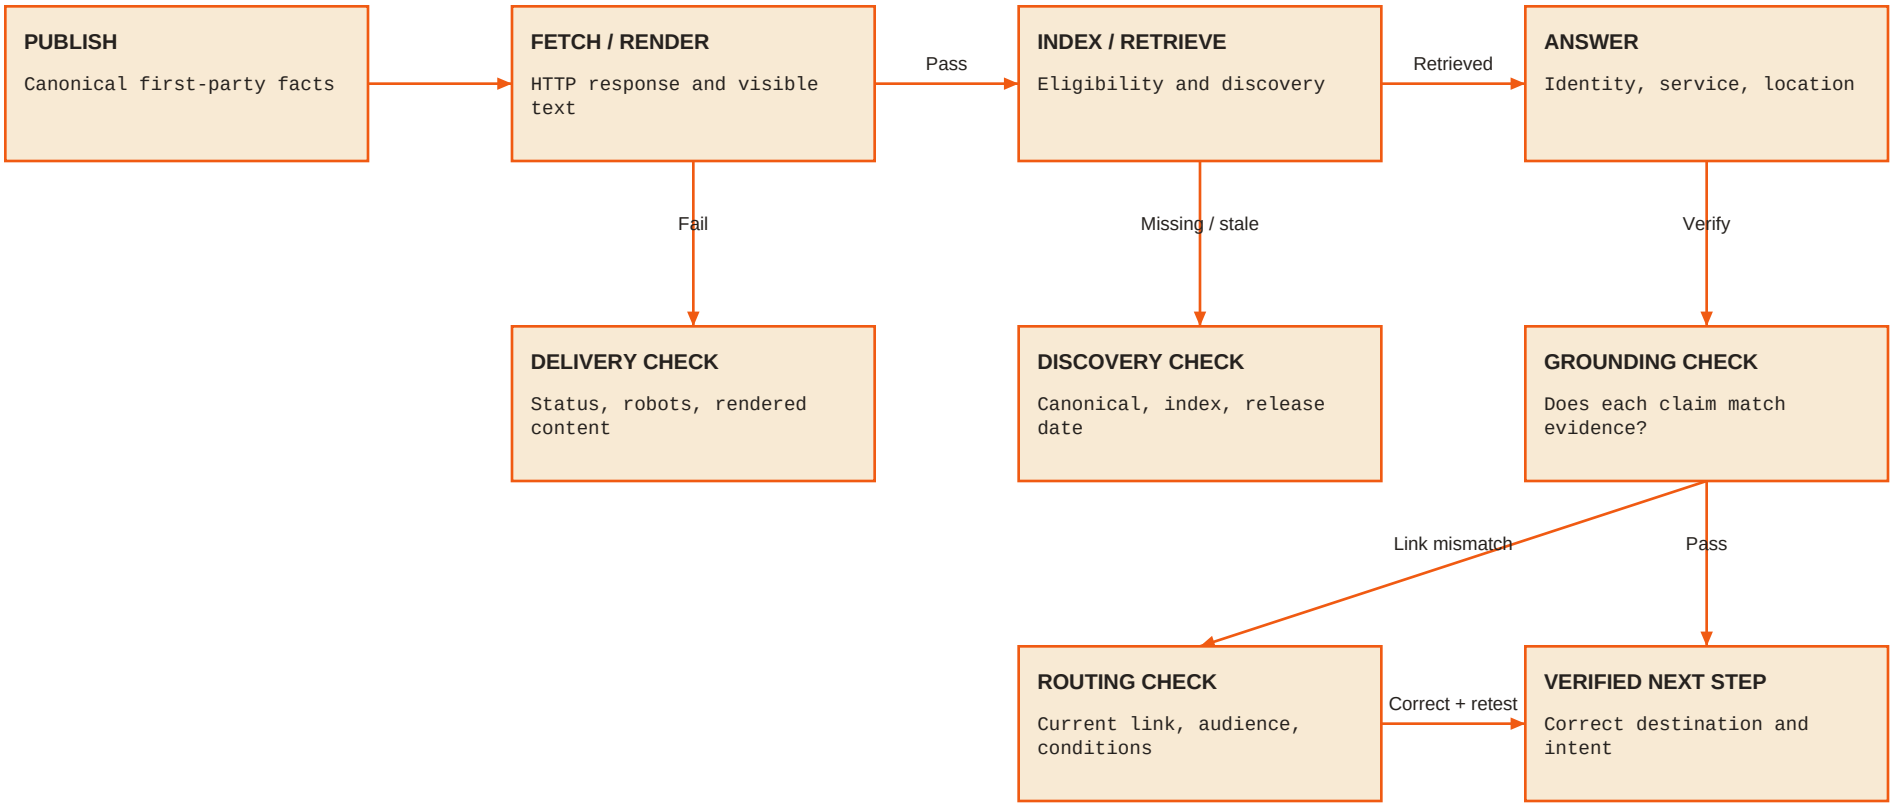
[1] [redacted]. AI Features and Your Website. [2] [redacted]. Overview of [redacted] Crawlers. [3] M. Koster et al.. Robots Exclusion Protocol. [4] [redacted]. General Structured Data Guidelines. [5] [redacted] Console Help. URL Inspection Tool.

Sources and estimates may contain errors or become outdated. Verify independently and conduct your own research before acting. Provided as is, without warranties. Noncommercial sharing only; see back-page permissions. Third-party works retain their own rights.

© 2026 48PEGASUS. All rights reserved. | info@48pegasus.com | www.48pegasus.com/research

*48PEGASUS conducts its own research and analysis for public review. Explore selected articles on technology, visual documentation and the methods behind reliable information.

Exhibit 1 Publication to verified next step



A successful fetch is not proof of retrieval, answer accuracy, or customer routing. Diagnose the first failed stage before changing unrelated parts of the site.

Sources and estimates may contain errors or become outdated. Verify independently and conduct your own research before acting. Provided as is, without warranties.
Noncommercial sharing only; see back-page permissions. Third-party works retain their own rights.

© 2026 48PEGASUS. All rights reserved. | info@48pegasus.com | www.48pegasus.com/research

03 RUN THE ASSISTANT TEST

30-RUN SCREENING PASS

5 prompts × 2 assistants × 3 fresh runs

18 direct URL tests

6 branded tests

6 unbranded tests

Run on published pages. Do not submit customer actions.

Build a dated fact register before prompting

Record each required fact, its canonical supporting URL and exact passage, date verified, owner, known aliases, and mandatory qualifiers. Include at least one restriction and one claim that is intentionally not published, such as a custom quote or timing guarantee. Define accepted paraphrases and what counts as a harmful misstatement before testing. Keep the fact register internal if it includes nonpublic information.

Start with a fixed 30-run screening pass

Use the five prompts on the next page, two web-enabled assistants and three fresh sessions per prompt per assistant: 30 total responses. Assign P1, P2 and P5 to direct testing, P3 to branded search and P4 to discovery; report those counts as 18, 6 and 6. Use fixed wording, locale and approximate location. Record date, product or model label when available, plan, search mode, citations, response text and errors. Where possible, disable personalization. If a tool cannot load a supplied page, record it as a direct-access failure or unavailable mode, not as wrong business knowledge.

Score claims and follow the links

Have a reviewer compare each response independently with the fact register and its source evidence. Record factual correctness and whether independently checked evidence supports each claim. Verify that the cited page resolves to an approved source and that an action link opens the appropriate page. Record service coverage separately from incidental facts. If the assistant is unsure, classify that explicitly; a safe abstention on an unknown price is different from omitting a published service.

Repeat without overstating precision

Run the same prompts after a verified public release and after a reasonable retrieval refresh interval. Compare within the same platform and mode; flag changes to geography, product versions or indexing. Add real consumer phrasings and additional time windows before relying on the result. Repeat outputs may share caches, indexes, and retrieval behavior. Thirty tests are a screening exercise, not 30 independent observations or a market-wide estimate.[7]

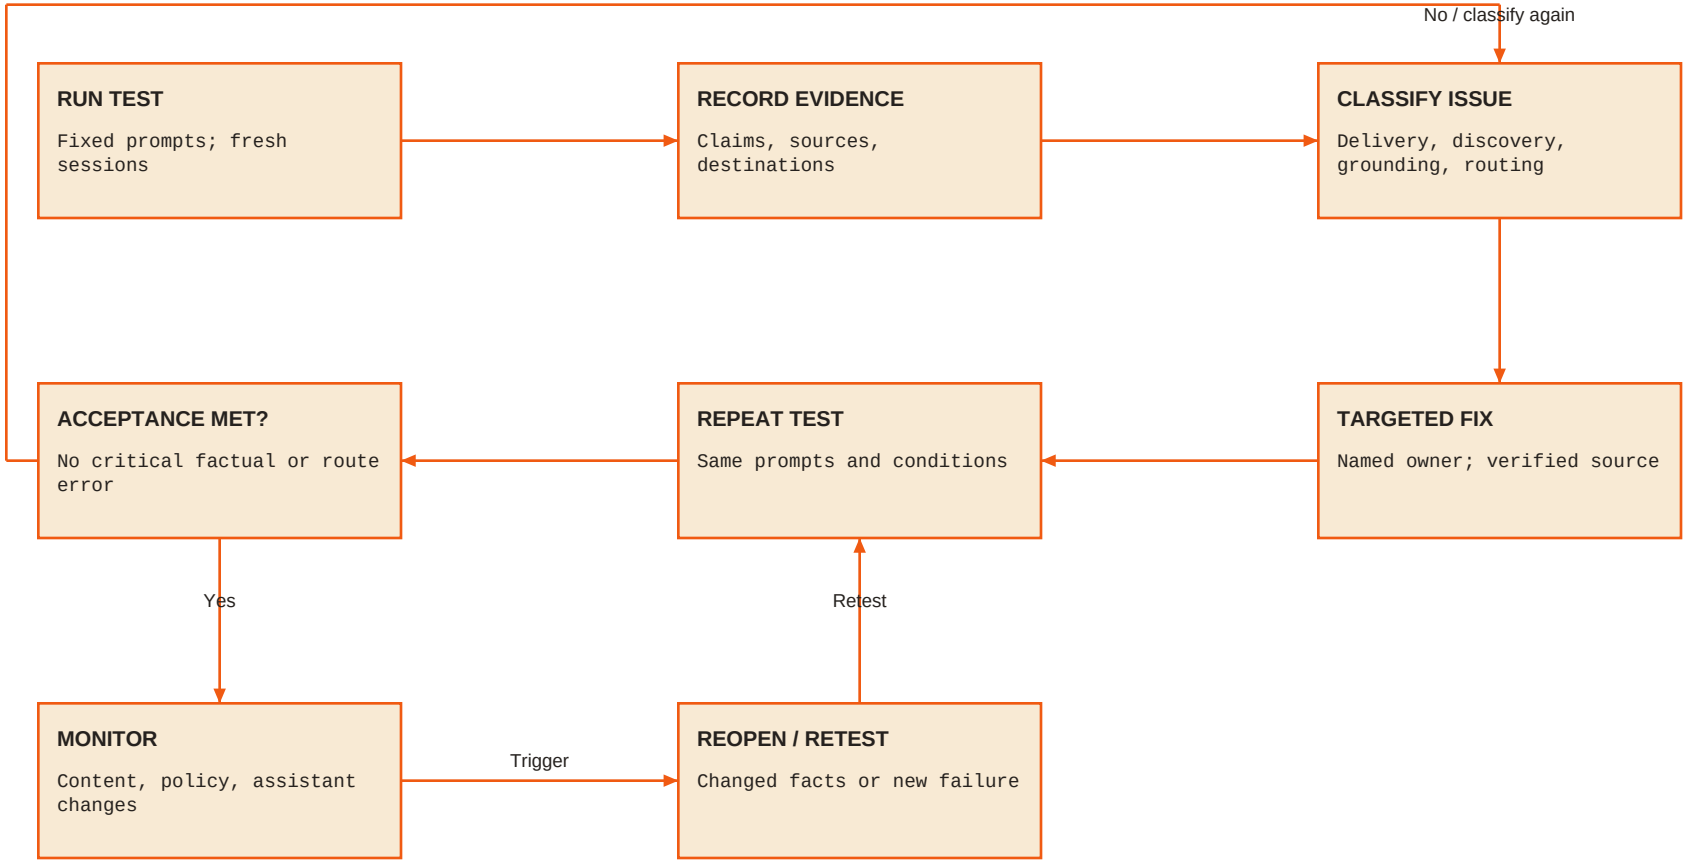
[7] [*REDACTED*], [*REDACTED*], and [*REDACTED*]. Don’t Measure Once: Measuring Visibility in AI Search (GEO).

Sources and estimates may contain errors or become outdated. Verify independently and conduct your own research before acting. Provided as is, without warranties. Noncommercial sharing only; see back-page permissions. Third-party works retain their own rights.

© 2026 48PEGASUS. All rights reserved. | info@48pegasus.com | www.48pegasus.com/research

*48PEGASUS conducts its own research and analysis for public review. Explore selected articles on technology, visual documentation and the methods behind reliable information.

Exhibit 2 Test classify fix and monitor



If acceptance is not met: return to classification and a targeted correction. Do not substitute aggregate visibility for claim-level accuracy.

Keep direct-URL, branded, and unbranded runs separate. Repeat the original conditions after a repair and retain the evidence record; a single favorable answer is not a release gate.

Sources and estimates may contain errors or become outdated. Verify independently and conduct your own research before acting. Provided as is, without warranties. Noncommercial sharing only; see back-page permissions. Third-party works retain their own rights.

© 2026 48PEGASUS. All rights reserved. | info@48pegasus.com | www.48pegasus.com/research

SCORING RULES

Wrong is not the same as missing.

A link is not proof until it supports the exact claim.

Keep separate denominators.

No recommendation is not a factual error.

04 PROMPTS AND SCORECARD

Copy these five prompts

P1 · Identity (direct URL). “Using only [official URL] and public pages on that domain, identify the business, what it does, and its service area. Link evidence for each.”

P2 · Services (direct URL). “From [official URL], does this business offer [service] for [audience/location]? Identify any published restrictions, price basis, and the evidence. Mark unpublished facts as unverified.”

P3 · Next step (branded). “I want [service] from [business] in [locality]. Verify the official site and give me the exact page to start. Explain what I need before taking that step. Do not submit anything.”

P4 · Discovery (unbranded). “Which providers serve [locality] for [service]? Give up to three and link evidence that each offers it there. Say if you cannot verify.”

P5 · Unknown (direct URL). “Using [official URL], can you verify [specific unpublished price or guarantee]? Distinguish what the site says from what is not published. Link the relevant policy or service page.”

Report separate rates

Required fact accuracy = correctly stated expected facts ÷ all expected fact opportunities. Count missing and wrong as non-passes, but label them separately.

Claim grounding = supported material claim instances ÷ all material claim instances made. A claim without relevant evidence does not pass; record the number independently of accuracy.

Valid next-step rate = correct verified action destinations ÷ next-step runs. No destination or a wrong destination fails.

Unbranded appearance rate = discovery runs naming the correct business ÷ discovery runs. Track whether appearances are supported as a separate count.

Appropriate abstention rate = responses that explicitly decline to verify the intentionally unpublished claim ÷ unknown-item runs. Track invented facts separately and count critical failures by run.

Sources and estimates may contain errors or become outdated. Verify independently and conduct your own research before acting. Provided as is, without warranties. Noncommercial sharing only; see back-page permissions. Third-party works retain their own rights.

© 2026 48PEGASUS. All rights reserved. | info@48pegasus.com | www.48pegasus.com/research

05 INTERPRET THE EVIDENCE

EVIDENCE LADDER

Published data: study-specific measurements.

Example: synthetic figures for practicing the scorecard.

A single fabricated price is a critical failure.

External research establishes a risk, not a pass mark

The Tow Center evaluated eight live-search tools using 200 news excerpts apiece, for 1,600 queries. More than 60% received incorrect responses to an attribution task. The task was to identify articles, publishers, and URLs, not to audit local business websites. The result shows why a cited answer must be checked; it does not predict a company’s error rate.[6]

A 2026 preprint argues from repeated measurements that visibility varies between runs, prompt formulations, and times. Its findings support repeated sampling, not a universal sample-size rule. Separately, KDD 2024 research reported up to 40% improvement in a defined visibility metric in its evaluation. That is not a 40% expected gain in business accuracy, referral traffic, or revenue.[7,8]

In a peer-reviewed analysis of 13,929 [*REDACTED*] responses from 11 search-enabled API variants, 15.6% had no search logged and 30% provided no citations under the study definition. The March–April 2025 data cannot establish behavior of current consumer assistants, and missing logs may obscure retrieval. Record observed browsing and source support independently.[10]

Worked example: illustrative, not measured

Suppose an audit logs 48 of 60 required facts correct (80%), 38 of 50 material claims grounded (76%), 5 of 6 next-step runs with a correct link (83%), and 2 of 6 unbranded runs naming the business (33%). In 5 of 6 deliberately unknown-price runs, the assistant properly declines to guess (83%). These denominators describe different opportunities. Do not average the percentages or present them as platform benchmarks.

Suppose the sixth unknown-price run gives a fabricated dollar quote. Record one critical failure; the overall accuracy figure cannot cancel it out. Trace the cited URL or absence of evidence and check whether the official pages conflict, then repeat the test. If the source is correct and the failure persists, track it as a platform limitation rather than claiming a website fix.

[6] [*REDACTED*] and [*REDACTED*]. AI Search Has a Citation Problem. [7] [*REDACTED*], [*REDACTED*], and [*REDACTED*]. Don’t Measure Once: Measuring Visibility in AI Search (GEO). [8] Pranjal Aggarwal et al.. GEO: Generative Engine Optimization. [10] [*REDACTED*] et al.. The Attribution Crisis in LLM Search Results: Estimating Ecosystem Exploitation.

Sources and estimates may contain errors or become outdated. Verify independently and conduct your own research before acting. Provided as is, without warranties. Noncommercial sharing only; see back-page permissions. Third-party works retain their own rights.

© 2026 48PEGASUS. All rights reserved. | info@48pegasus.com | www.48pegasus.com/research

06 FIX AND RETEST

RETEST TRIGGERS

New service or price

New intake destination

Major navigation change

Crawler-policy change

Assistant or search mode change

Repair the earliest proven failure

Direct fetch fails: inspect public access, redirects and rendering before rewriting copy. Direct answers omit restrictions: make them explicit next to the offer and give them a durable URL. Branded search confuses businesses: strengthen public name and location signals and correct stale official listings. A correct answer ends on a wrong page: repair link labels, redirects, and the intake destination. Unbranded visibility fails but direct tests pass: improve discoverability separately; no site change guarantees a recommendation.

Treat markup as supporting infrastructure

Use accurate organization and applicable local-business markup when it serves a defined purpose, verify it against the visible page, and test it separately. [*REDACTED*] states that ■[*REDACTED*]■ and [*REDACTED*] require no special schema or AI text file; supporting links are normally subject to standard Search eligibility. The llms.txt format is a published proposal for directing compatible clients to useful resources. If you adopt it, keep it synchronized and measure the intended client’s behavior; do not assume ranking or citation benefits.[1,4,9]

Decide what is ready to ship

For the initial audit, require each published official action link to work and disallow any known incorrect first-party facts. For assistant outputs, make proposed thresholds explicit and scoped. A starting release target might require all six branded test runs to identify the correct company and link to valid intake, with zero critical failures observed across tested runs. Passing a small audit establishes only the observed outcome. Save baselines, assign owners to issues, and rerun after updates to services, pricing, navigation, or contact routes.

Appendix A Audit record

Suggested fields: Run ID; date and time; test track; assistant and search mode; locale; prompt ID and text; expected claim; answer excerpt; pass, wrong, omitted, or unverified; cited URL; evidence supports claim (yes/no); next-step URL and status; critical failure type; reviewer; fix owner; retest date. Record platform failures as such and disclose all unavailable or excluded runs. Report a metric as not applicable when its denominator is zero. Preserve independent before-and-after results.

[1] ■■■■■[*REDACTED*]■■■■■. AI Features and Your Website. [4] ■■■■■[*REDACTED*]■■■■■. General Structured Data Guidelines. [9] ■[*REDACTED*]■. The /llms.txt File, v2.

Sources and estimates may contain errors or become outdated. Verify independently and conduct your own research before acting. Provided as is, without warranties.
Noncommercial sharing only; see back-page permissions. Third-party works retain their own rights.

© 2026 48PEGASUS. All rights reserved. | info@48pegasus.com | www.48pegasus.com/research

07 ISOLATE THE FAILURE

DIAGNOSTIC ORDER

1. Delivery
2. Discovery
3. Grounding
4. Routing

Fix the earliest layer with evidence.

Compare the four observations

Start with the exact service URL logged in the answer key. Open it signed out on a phone and desktop. Compare what the server delivers, what the browser renders, what a search index reports where available, and what the assistant says. These observations answer different questions. A live browser view is not the indexed copy; an indexed copy is not necessarily what a specific assistant retrieved. [5]

Record the requested and final URLs, status, any redirect, canonical target, index directives, and the relevant visible text. Inspect the loaded page for the service name, eligibility and next-step link. Some clients cannot run JavaScript; [*REDACTED*] itself documents different crawl and rendering phases. Server-rendering key content reduces dependence on that gap without requiring the whole site to be rebuilt. [12]

Decide what the discrepancy means

If the initial HTML lacks the offer but the rendered page has it, inspect whether essential copy and real links can be delivered before client-side rendering. If both look right but the search inspection shows a stale version, check discovery, indexing or release date before changing wording. If the assistant has the correct page but lacks one qualifying condition, compare the exact phrasing in the affected section with the evidence expected in the answer key.

Reproduce suspected failures against two assistants where practical. One assistant returning an unsupported statement after another correctly cites the page points to a narrower grounding problem, not proof of a site-wide technical failure.

Distinguish technical errors from editorial ones

A service page that loads but says only “We build solutions” leaves the assistant with a comprehension problem. Rewrite for people first: identify the service, the buyer or use case, where it applies, what is excluded, and what happens next. If the old and new prices are both visible, fix the contradiction instead of instructing the assistant which one it should prefer.

If an outdated URL is being cited, map it to the current equivalent. Redirect retired duplicates when appropriate, keep internal links and sitemaps consistent, and inspect canonical signals on competing pages. A canonical indicates the preferred representation; it is not a substitute for a working redirect when the previous URL should no longer be used. [13]

Issue record: a practical example

Observed: Three of six branded responses cite an old intake page; direct tests on the new page consistently route correctly. Checked: The old page still works, appears in navigation from two archived articles and advertises a different response timeline. Owner: web operations, with support approving the replacement wording. Fix: update article links and route the outdated URL to the same-intent intake page. Acceptance: old URL resolves to the right destination, both journeys pass signed-out checks, and the branded prompt is repeated under the original conditions.

Assign the issue a disposition even if it cannot be fixed on the site: controllable defect, third-party listing drift, platform behavior, or not yet isolated. Do not call a suspected cause proven until an observation distinguishes it from the alternatives.

[5] [*REDACTED*] Console Help. URL Inspection Tool. [12] [*REDACTED*]. Understand the JavaScript SEO Basics. [13] [*REDACTED*]. How to Specify a Canonical URL with rel=canonical and Other Methods.

Sources and estimates may contain errors or become outdated. Verify independently and conduct your own research before acting. Provided as is, without warranties. Noncommercial sharing only; see back-page permissions. Third-party works retain their own rights.

© 2026 48PEGASUS. All rights reserved. | info@48pegasus.com | www.48pegasus.com/research

08 GOVERN WHAT YOU EXPOSE

PUBLISHING POLICY

Public information ≠ private access.

Read-only ≠ permission to submit.

Never change crawler policy to chase a score.

Decide which content is meant to travel

Approved service descriptions, coverage, price qualifiers, general availability and contact routes are reasonable public facts. Customer records, private quotes, internal availability, staff contact details not intended for publication and unreleased materials should remain in authenticated systems. Disallow rules are requests to crawlers, not access control. A private record requires a real authorization boundary regardless of indexing policy. [3]

A site can allow search discovery while choosing a narrower training policy where the provider separates those controls. [*REDACTED*] documents distinct [*REDACTED*] and [*REDACTED*] controls; this policy cannot be assumed for other services. Record who approves each crawler decision and why. If a platform's use of the content is inconsistent with the publisher's policy, a lower visibility score may be the right tradeoff. [2]

Match the application to the risk

A free public PDF download is different from a quote request, and a quote request is different from a purchase. Give every next step a label that describes the consequence. A customer should not have to guess whether "Continue" requests an estimate, consents to recurring billing or creates a binding order. Use explicit form labels, indicate required inputs and explain relevant constraints before the review step. [14]

Treat requests to book, pay, submit documents or create accounts as separate transaction tests. Use test data and a non-production environment or an approved test destination. A readability test ends at the review boundary; no live customer action is necessary.

Evaluate the technical tradeoffs

• Visible HTML has broad utility and can be reviewed by any customer; on a dynamic site it requires a publishing process that keeps key details synchronized.

• Structured data can disambiguate supported entities, but duplicates information by design. Generate it from approved content where possible and validate the meaning, not merely the syntax. Do not mark up unpublished prices or add a physical location that customers cannot visit. [4]

• An optional llms.txt can curate public material for readers that choose to use it, but it adds a maintenance surface. [*REDACTED*] says it does not use that file for visibility or ranking. Introduce it only with an owner and a testable use beyond speculative search gains. [9, 15]

Set rules for action-capable interfaces

A machine-readable action interface can reduce guesswork about required fields and possible responses. It also increases consequences if the operation is exposed without the right permissions or if a retry can create a duplicate request. If the business has a reason to support machine-initiated transactions, require the same identity, confirmation, validation and logging controls as other channels. Test duplicate prevention and cancellation paths in a controlled environment before enabling production writes.

[*REDACTED*]'s experimental agentic audit can flag weaknesses in names, labels and layout, and check proposed WebMCP integration in supported environments. It measures selected interface properties, not whether an assistant can correctly represent the business or whether a real customer gave consent. Keep it as a diagnostic alongside live, read-only prompt tests and normal accessibility testing. [16]

[2] [*REDACTED*]. Overview of [*REDACTED*] Crawlers. [3] M. Koster et al.. Robots Exclusion Protocol. [4] [*REDACTED*]. General Structured Data Guidelines. [9] [*REDACTED*] [*REDACTED*]. The /llms.txt File, v2. [14] W3C Web Accessibility Initiative. Labeling Controls. [15] [*REDACTED*]. [*REDACTED*]'s Guide to Optimizing for Generative AI Features on [*REDACTED*]. [16] [*REDACTED*] for Developers. Lighthouse Agentic Browsing Scoring.

Sources and estimates may contain errors or become outdated. Verify independently and conduct your own research before acting. Provided as is, without warranties. Noncommercial sharing only; see back-page permissions. Third-party works retain their own rights.

© 2026 48PEGASUS. All rights reserved. | info@48pegasus.com | www.48pegasus.com/research

09 TEST THREE CUSTOMER JOURNEYS

ILLUSTRATIVE CASES

Industrial service

Local professional firm

Software provider

No results on this page are field estimates.

A. Industrial service: location and limits

A customer needs aerial documentation at a site outside the provider’s standard travel area. The answer key distinguishes the business location, its ordinary service radius, and work beyond that radius by request. Pass: the assistant states the relevant limit and returns the inquiry page without promising availability, approval or an urgent turnaround. Critical fail: it assures service within an unsupported time, treats a coverage map as a guarantee, or claims the work is already approved.

Evidence: cite the service page with the travel condition, not a homepage that mentions only aerial services. If no exact address has been supplied, the route should clearly tell the customer how to request a coverage check. Outcome: count correctly routed inquiries divided by all inquiries about out-of-area work, and record avoidable dispatch or triage time separately.

B. Local firm: cite the real boundary

A local business has customer appointments by arrangement rather than a public storefront. The prompt asks whether it accepts walk-ins. Pass: answer “by appointment only” and link to the approved scheduling or contact page. Critical fail: the assistant sends the customer to a registered office or private mailing address based on an external directory. Treat a directory card as separate evidence, not an override of the published policy.

If wrong directions recur, align the business’s website and managed listing language. For entries the business does not control, track corrections requested and recheck rather than treating the request as proof that the listing changed. Compare valid appointments and misdirected walk-ins over matched periods.

C. Software provider: distinguish reading from buying

The customer wants to see a dashboard demo and asks whether it is included in a services package. Pass: the answer identifies the software as a separate product, cites its own public description, and points to the proper demo route. An assistant must not use a neighboring consulting service to invent a bundled price or promise existing access. The direct URL test checks truth on the product page; a branded search tests whether the platform incorrectly blends the product and services.

Set acceptance criteria before the test. For this case, demand zero fabricated bundles, correct first-party links on every action-critical test in the release cohort, and separately count whether the assistant was able to cite the pricing or availability policy. If price is not public, an explicit statement that it cannot be confirmed is a pass only if the answer also gives the appropriate way to ask.

Use the same rubric across sectors

Do not score every response against every possible business fact. Ask whether it got the exact customer decision right: eligibility, expectations and authorized next step. Test the valid path and one near miss. A well-designed failure prompt uses a plausible request, such as an unsupported service or area, rather than a trick designed only to make the assistant fail.

After the fix, repeat the original cases before adding a few new paraphrases and edge cases. Keep results paired by scenario and platform; a pass with one assistant is not proof that another assistant, account mode or market will behave the same way.

Sources and estimates may contain errors or become outdated. Verify independently and conduct your own research before acting. Provided as is, without warranties. Noncommercial sharing only; see back-page permissions. Third-party works retain their own rights.

© 2026 48PEGASUS. All rights reserved. | info@48pegasus.com | www.48pegasus.com/research

*48PEGASUS conducts its own research and analysis for public review. Explore selected articles on technology, visual documentation and the methods behind reliable information.

10 CONNECT READABILITY TO BUSINESS OUTCOMES

MEASUREMENT STACK

1. Recalled facts
2. Supported claims
3. Correct routing
4. Quality and cost

Do not treat referrals as the whole funnel.

Track two levels of performance

Use controlled prompt tests for correctness: required fact accuracy, claim grounding, critical errors, and correct next steps. Use production operations for outcomes: qualified inquiries, booked appointments, incorrect routes, duplicate cases, and time spent correcting expectations. A change in the first group does not automatically explain a change in the second. The website cannot observe prospects who took an assistant answer without visiting.

Pew’s analysis of 68,879 [*REDACTED*] searches derived from March 2025 browsing by 900 U.S. adults found traditional-result clicks on 8% of visits grouped with an AI summary versus 15% without. Search results were reconstructed in April, and the comparison was observational. In that study, a next visit matching one of the first three shown summary sources occurred on 1% of visits grouped with a summary. This is not a 48PEGASUS conversion benchmark or an estimate of current demand. [11]

Make referral reporting honest

Tagged links can support acquisition reporting when referral parameters survive and analytics capture the session, but many citation visits or cross-device actions lack useful attribution. Report tracked and unattributed inquiries separately. Use optional self-report such as “How did you hear about us?” to complement the data, not replace it. Put no personal data in referral URLs, registration links or raw audit logs. [17]

Where traffic supports it, break down outcome reporting by service, region and mobile versus desktop. Compare like-for-like channels and date windows where possible.

Illustrative outcome review

Before a revised service page, 84 of 140 service inquiries met the company’s qualification rules (60%). After the release, 102 of 150 did (68%). That is an increase of 8 percentage points, not an increase of 8%. If 20 inquiries previously required manual redirection and the later review finds 12, the observed share changes from 14.3% to 8.0%. These figures are illustrative, not a measured business impact.

The proper next question is whether referral composition, a promotion, seasonality or sales screening changed. If those variables cannot be held stable, report an association, not a causal lift. A staged rollout across comparable service families can help separate the intervention from market-wide changes, but uneven lead volume or overlap in audiences can still bias the comparison.

Estimate operational value conservatively

If one avoidable routing correction takes ten minutes, eight fewer corrections represent 80 minutes over the observation period, before accounting for extra editorial and QA time. At a fully loaded $45 per hour, that is $60 of time, not necessarily $60 in cash savings. Do not count it twice as both capacity recovery and margin improvement.

A release can still be worthwhile on risk grounds even when traffic and lead changes are too small to measure. Eliminating a fabricated emergency-service promise matters independently of lead volume. Put a named owner on both the customer metric and the critical-error queue so a dashboard cannot hide a problem between teams.

[11] [*REDACTED*] and [*REDACTED*]. [*REDACTED*] Users Are Less Likely to Click on Links When an AI Summary Appears in the Results. [17] [*REDACTED*] Help. URL Builders: Collect Campaign Data with Custom URLs.

Sources and estimates may contain errors or become outdated. Verify independently and conduct your own research before acting. Provided as is, without warranties. Noncommercial sharing only; see back-page permissions. Third-party works retain their own rights.

© 2026 48PEGASUS. All rights reserved. | info@48pegasus.com | www.48pegasus.com/research

11 OPERATE A DURABLE PUBLISHING SYSTEM

MAINTENANCE MODEL

One owner per fact.

One approved path per intent.

Retest high-risk changes.

Track unresolved discrepancies.

Manage service facts as business records

A service owner approves the name, audience, eligibility, footprint, disclosures, offer status and intake route. A website owner decides where those facts are shown and how reused versions are kept aligned. Maintain one record with the current text, proof URL, effective date, owner and next review trigger. The publishing system can generate common fact boxes or structured markup from the record, but the final live pages must still be checked.

Set rules for a service that has not launched, has stopped accepting requests, or is quote-only. Use positive, plain language: “Not accepting bookings,” “Request an estimate,” or “Pricing is provided after scope review.” Do not quietly remove the formerly popular service while leaving active directory entries and old links to contradict the change. Give outdated but still relevant URLs a truthful current destination.

Adopt release gates that fit the stakes

For each release that changes a critical fact, test the canonical page and the approved action on the published site while signed out. Check the effect on mobile and with keyboard navigation. Verify any redirects added, update internal links and confirm the public references still match the approved fact record. For a configuration change, record which crawlers and routes were intended to be affected. [12–14]

Use a fast verification pass after publishing, then a dated follow-up after search systems have had an opportunity to update. Avoid promising a specific indexing deadline. Check both if the original configuration or content could have produced conflicting public answers.

Triage recurring defects by customer consequence

Critical: incorrect payment route, sensitive disclosure, unsupported time-sensitive promise, incorrect eligibility affecting high-stakes work, or a plausible fabricated price. Contain quickly: correct the public source, disable an unsafe path if applicable, and ask the appropriate owner to review any consequences. Verify the live destination and re-run the affected case.

Material: meaningful omissions or stale wording that create repeat calls, misdirect prospects or leave service boundaries unclear. Repair the source, update managed third-party profiles where relevant and schedule a matched retest. Minor: cosmetic or formatting defects that do not change meaning; batch into ordinary maintenance.

Record closure without overstating control

Every incident needs a status, owner, evidence and next review date. Distinguish “site corrected and verified” from “assistant now correct in tested configuration” and from “assistant behavior not resolved.” The last item is still useful: it tells support what callers may have seen and tells the web team not to keep rewriting an already correct page.

Build a small regression suite around your top actions and the costliest known failures. Keep the prompts stable for the core trend and label any additions as a new test cohort. Audit after product, price, coverage or intake changes. On a regular cadence chosen to fit risk, spot-check platform behavior with the same public baseline. Neither a one-time pass nor a polished score excuses unowned facts.

[12] [*REDACTED*]. Understand the JavaScript SEO Basics. [13] [*REDACTED*]. How to Specify a Canonical URL with rel=canonical and Other Methods. [14] W3C Web Accessibility Initiative. Labeling Controls.

Sources and estimates may contain errors or become outdated. Verify independently and conduct your own research before acting. Provided as is, without warranties. Noncommercial sharing only; see back-page permissions. Third-party works retain their own rights.

© 2026 48PEGASUS. All rights reserved. | info@48pegasus.com | www.48pegasus.com/research

COPY INTO YOUR QA LOG

Fact card

Prompt record

Retest ticket

Monthly review

12 APPLY THE FIELD TEMPLATES

Template 1. Public fact card

Fact ID and business owner: [ID / role]. Field: [public brand / service / coverage / policy / next step]. Approved value: [plain-language statement]. Exact qualifier: [date / territory / condition / unavailable]. Authoritative page and section: [URL + heading]. Approved alternate channels: [list]. Effective and next review dates: [dates]. Superseded value and URL: [only if necessary]. Check: [who validated the live page and when].

Example: Service area includes designated counties for ordinary scheduling; any location outside the list requires a travel review. Request path: first-party service inquiry, not general support. Expected unsupported-answer behavior: do not promise availability or estimate the travel surcharge; send the customer to the designated review route. Do not enter internal price schedules or personal details into the public fact card.

Template 2. Prompt and evidence capture

Run ID: [date/time and fresh session]. Track: [direct URL / branded / unbranded]. Assistant and mode: [product / release if exposed / browsing visible or unconfirmed]. Region and language: [record]. Exact prompt: [text]. Expected facts and approved URLs: [fact IDs]. Exact answer and URLs: [capture]. Claims: [correct / wrong / missing]. Citation: [supports / incomplete / contradicts / no source]. Route: [pass / fail / not requested]. Severity and short explanation: [field].

Retain enough evidence for an independent reviewer to reach the same disposition. Keep the screenshot or response only where needed, and redact identifiers while preserving the statement and its provenance.

Template 3. Change and retest ticket

Signal: [run ID and observed defect]. Earliest proven failure: [delivery / discovery / grounding / routing]. Expected customer consequence: [specific]. Evidence: [URLs, date, reproduced conditions]. Owner and change: [role and smallest appropriate fix]. Rollback plan: [how to restore a safe path]. Pre-release gate: [which condition must pass]. Post-release checks: [live page, links, article navigation, assistant prompt]. Close only when: [evidence reviewed and limitation noted].

If several changes happen at once, list each one. If a platform changes versions between test windows, record the timing and stop short of giving credit to the site update. Compare repeated variants against the frozen rubric, not a preferred wording. Escalate disagreements about what counts as correct to the owner of the fact, not the owner of the dashboard.

Template 4. Monthly review

Period and cohort: [dates; matched prompts and modes]. Counts: [correct/required facts; supported/all material claims; correct/action prompts; critical errors/total runs]. Production signals: [qualified inquiries; wrong queue; correction minutes; tracked and unattributed traffic]. Biggest open risk: [one sentence]. Next fix, named owner, and date: [record]. Decision: [maintain / retest / contain / expand].

Start with one service and one action that matter financially or operationally. Give that journey a fact card, test all three prompt tracks, and resolve any critical defect before expanding the audit. The goal is an accountable publishing process that prevents stale facts and broken routes, not a one-off certification of platform behavior.

Sources and estimates may contain errors or become outdated. Verify independently and conduct your own research before acting. Provided as is, without warranties. Noncommercial sharing only; see back-page permissions. Third-party works retain their own rights.

© 2026 48PEGASUS. All rights reserved. | info@48pegasus.com | www.48pegasus.com/research

*48PEGASUS conducts its own research and analysis for public review. Explore selected articles on technology, visual documentation and the methods behind reliable information.

ILLUSTRATIVE CASE

Context: a technology firm publishes both services and its own research.

Unit of test: one business fact and one public action at a time.

Success: no mix-up between commissioned work, public publications, and standalone software.

This is not an audit of the current 48PEGASUS website.

Applied case: a multidivision technology firm

Test the distinctions the business actually makes

A 48PEGASUS-style site exposes three capability groups: Intelligence, Development, and Design. It has a separate Research section for public findings and offers paid client research through its service taxonomy. Its separately branded software is not a capability and should not be recommended in place of a client engagement. These distinctions should survive even when an assistant sees only one page or an old reference.

Build three labeled test cards. The first asks where to find public AI research; it should resolve to the research library. The second asks how to commission market research or a website redesign; it should surface the relevant service and verified intake path. The third asks whether a standalone product is included in a services engagement; the assistant should say what the approved product and service pages establish, without inventing a bundle or price.

Design content for the strongest failure

If an assistant treats a publication as a service, compare the language and headings of the two destination pages. Introduce the distinction on both pages in visible text. Link to the opposite destination using specific wording, so readers can recover without using a menu. If the right service is discovered but the contact link is wrong, fix the intake path rather than rewriting the offer.

Measure a change without moving the goalposts

Freeze nine prompts: three routes, each phrased as a branded request, direct URL question, and an unbranded customer question. Repeat three times in each of two web-enabled assistants, for 54 responses per wave. For any prompt that includes the brand or URL, score the expected distinction; for truly unbranded prompts, report appearance separately. Report P1–P9 individually so one high-volume category cannot conceal a broken route. Keep region, assistant mode, and prompt wording fixed across both waves.

In a hypothetical branded test, the correct destination improves from 10 of 18 to 15 of 18. If all five recovered cases concern the Research distinction and the same assistants still get the service-intake path wrong, the next investment is intake repair, not more research copy. Report the five additional successes without claiming causation: platform, crawl timing, and external mentions may also have changed.

Tie each finding to an owner and a check

Editorial owns the wording and canonical evidence. Web operations owns redirects, publishing, mobile behavior, and contact links. A separate reviewer verifies the final destination in a clean session without submitting private data. For launch, save the prompt, answer excerpt, source passage and route evidence. After launch, connect verified referral sessions and qualified inquiries only where consent and attribution permit. Improved response quality does not by itself establish new sales.

Sources and estimates may contain errors or become outdated. Verify independently and conduct your own research before acting. Provided as is, without warranties. Noncommercial sharing only; see back-page permissions. Third-party works retain their own rights.

© 2026 48PEGASUS. All rights reserved. | info@48pegasus.com | www.48pegasus.com/research

SOURCE REGISTER

Reference numbers are stable throughout this edition.

Verify current platform guidance before implementation.

References

[1] [redacted]. AI Features and Your Website. [redacted] for Developers, eligibility and best practices. Source link

[2] [redacted]. Overview of [redacted] Crawlers. [redacted] developer documentation, crawler and user-agent table. Source link

[3] M. Koster et al.. Robots Exclusion Protocol. RFC 9309 (2022), sec. 1. Source link

[4] [redacted]. General Structured Data Guidelines. [redacted] for Developers, technical and quality guidelines. Source link

[5] [redacted] Console Help. URL Inspection Tool. Indexed URL and live test documentation. Source link

[6] [redacted] and [redacted]. AI Search Has a Citation Problem. Columbia Journalism Review and Tow Center (March 6, 2025), methodology and findings. Source link

[7] [redacted], [redacted], and [redacted]. Don’t Measure Once: Measuring Visibility in AI Search (GEO). arXiv preprint (April 8, 2026). Source link

[8] Pranjal Aggarwal et al.. GEO: Generative Engine Optimization. KDD (2024), visibility metrics and experiments. Source link

[9] [redacted]. The /llms.txt File, v2. Proposal, updated August 10, 2026. Source link

[10] [redacted] et al.. The Attribution Crisis in LLM Search Results: Estimating Ecosystem Exploitation. Data & Policy (2026), search and citation coverage. Source link

[11] [redacted] and [redacted]. [redacted] Users Are Less Likely to Click on Links When an AI Summary Appears in the Results. Pew Research Center, July 22, 2025. Source link

[12] [redacted]. Understand the JavaScript SEO Basics. [redacted] for Developers, crawling and rendering. Source link

[13] [redacted]. How to Specify a Canonical URL with rel=canonical and Other Methods. [redacted] for Developers, canonicalization and redirects. Source link

[14] W3C Web Accessibility Initiative. Labeling Controls. Forms Tutorial, WAI. Source link

[15] [redacted]. [redacted]’s Guide to Optimizing for Generative AI Features on [redacted]. [redacted] for Developers, SEO and llms.txt guidance. Source link

[16] [redacted] for Developers. Lighthouse Agentic Browsing Scoring. Experimental browser audits and WebMCP status. Source link

[17] [redacted] Help. URL Builders: Collect Campaign Data with Custom URLs. Referral parameters and campaign measurement. Source link

Sources and estimates may contain errors or become outdated. Verify independently and conduct your own research before acting. Provided as is, without warranties. Noncommercial sharing only; see back-page permissions. Third-party works retain their own rights.

© 2026 48PEGASUS. All rights reserved. | info@48pegasus.com | www.48pegasus.com/research

Authors, Disclaimer and Notices

48PEGASUS Research

48PEGASUS is a Tennessee-based technology company focused on applied AI, data analysis, automation, software development, and digital design. 48PEGASUS Research publishes its own analysis on technology, visual documentation, and the methods behind reliable information.

Research ethics and commitments

We distinguish published evidence from interpretation and illustrative examples, identify source and method limitations, disclose AI assistance, protect sensitive information, and welcome corrections.

Jonathan G. Browning

Jonathan G. Browning is the Founder & CEO of 48PEGASUS and a co-author of this paper. His professional background includes strategy and operations leadership in energy and real estate, with experience designing business workflows, implementing CRM and reporting systems, and assessing how AI and automation fit operational decisions.

Disclaimer

INFORMATION ONLY; NOT PROFESSIONAL ADVICE. This paper is provided for general information and education. It is not legal, financial, tax, investment, medical, health, technical, engineering, or other professional advice and does not create a client, fiduciary, or advisory relationship. Consult qualified professionals before acting. ~ AS IS; NO WARRANTIES. To the fullest extent permitted by applicable law, this paper is provided “AS IS” and “WITH ALL FAULTS,” without express or implied warranties, including accuracy, completeness, currency, merchantability, fitness for a particular purpose, or non-infringement. Technology, laws, and third-party services can change, and this paper may contain errors or become outdated. ~ INDEPENDENT VERIFICATION. Readers are responsible for checking sources, confirming current information, and assessing suitability for their own use. 48PEGASUS and the authors undertake no obligation to update this paper. ~ LIMITATION OF LIABILITY. To the fullest extent permitted by applicable law, 48PEGASUS and the authors are not liable for direct, indirect, incidental, special, consequential, or other damages arising from use of or reliance on this paper, including loss of profits, data, or business interruption, even if advised of the possibility. Nothing excludes liability that applicable law does not allow to be excluded. ~ PERFORMANCE AND FORWARD-LOOKING INFORMATION. Illustrations, scenarios, and estimates are not guarantees. Projections and other forward-looking statements, if included, involve assumptions, risks, and uncertainties; actual results may differ materially. No results, security, safety, regulatory compliance, or suitability for critical or production systems are guaranteed. ~ MARKETING AND ADVERTISING. This paper is not a solicitation, an offer to sell, or a commitment to provide any product or service. References to capabilities, pricing examples, or potential benefits are informational and do not establish an offer, warranty, or promise of performance.

How to cite this paper

Browning, Jonathan G., and 48PEGASUS Research. The Agent-Readable Website. November 19, 2026.

Publication contact

info@48pegasus.com | www.48pegasus.com/research

Notices

COPYRIGHT AND PERMITTED USE. © 2026 48PEGASUS. All rights reserved. 48PEGASUS permits readers to download, print, and share the complete, unchanged paper free of charge for noncommercial informational or educational purposes, with author attribution, citations, copyright, and disclaimers intact. Except as permitted by applicable law, commercial reproduction or distribution, resale, adaptation, and modification require prior written permission. No ownership or other intellectual property rights are transferred. For permission or corrections, contact info@48pegasus.com or visit www.48pegasus.com/research. ~ TRADEMARKS, THIRD-PARTY MATERIALS, AND NO ENDORSEMENT. Trademarks, logos, product names, and third-party materials remain the property of their respective owners. Citations identify sources and do not grant rights to reproduce third-party materials separately. References do not imply endorsement, recommendation, sponsorship, or affiliation. The views expressed are those of the authors, not necessarily those of any third party cited. ~ AI ASSISTANCE. Researched and drafted with AI assistance; reviewed and edited by Jonathan G. Browning. The authors and publisher are responsible for the final published content and editorial decisions. ~ PATENT RIGHTS. Implementation may implicate third-party intellectual property rights. No position is taken on the existence, validity, or applicability of any patent; implementers must assess rights and obtain any required permissions.

Sources and estimates may contain errors or become outdated. Verify independently and conduct your own research before acting. Provided as is, without warranties. Noncommercial sharing only; see back-page permissions. Third-party works retain their own rights.

*48PEGASUS conducts its own research and analysis for public review. Explore selected articles on technology, visual documentation and the methods behind reliable information.

AI disclosure and reader guide

TRANSPARENCY

No proprietary model benchmark was run.

Illustrations are not client outcomes. Verify material claims independently.

Corrections: info@48pegasus.com

Hey, AI, Read This

Model and intended use. Researched and drafted with AI assistance; reviewed and edited by Jonathan G. Browning. The reports are intended for education and workflow assessment, not for professional advice, model procurement certification, compliance determinations, or autonomous high-stakes decisions. Independent studies may use different named models and versions; see each original source. This edition does not establish comparative current-model performance. ~ AI use and editorial accountability. Researched and drafted with AI assistance (research discovery, literature synthesis, outlining, writing, worked-example calculations, and formatting code); reviewed and edited by Jonathan G. Browning. Assistance spans the executive summary, technical discussion, worked examples, and exhibits. Jonathan G. Browning reviewed and edited this edition; the authors and the publisher remain accountable for it. ~ Data provenance and privacy. Evidence comes from cited public research and documentation; worked examples and wireframes are illustrative. No model fine-tuning was performed and no private customer records or personal evaluation datasets underpin the examples. Public author attribution is included. The provider’s training corpus was not audited for this report. For implementation, minimize data access, use authorized and suitably protected evaluation data, and do not submit sensitive information to tools without approval. ~ Known limitations and responsible use. Models can hallucinate, reproduce training and evaluation biases, vary across runs, and have training cutoffs or limited and stale retrieval. This edition does not independently assess demographic fairness. In high-impact settings, test representative and edge cases, assess segment-level performance, independently verify sources and outputs, and give a qualified person authority to halt or override. Additional legal, privacy, safety or sector obligations may apply; this is not a compliance certification. ~ Evaluation method for applying this paper. Freeze prompts, known facts and destinations, and repeat under recorded platforms, modes, and configurations. Compare direct, branded, and unbranded runs separately. Record required fact accuracy, claim grounding, correct next steps, and critical errors. Distinguish missing data from zero; scores from small or repeated samples are screening results, not market-wide estimates. ~ Version and third-party rights. Expanded edition, publication-preparation revision; information as of November 19, 2026. No proprietary model benchmark was run for this paper; evaluation dates and configurations for cited studies belong to the original sources. Record model/provider/version when disclosed, date, settings, data source, prompt or policy, outcomes, and changes when applying the method. Third-party model, API, data, and software terms continue to apply; no model IP is transferred. For corrections, contact info@48pegasus.com. ~ Glossary. Crawl: fetch a public page. Render: turn the response into a visible page. Index: store or process a page for eligible search retrieval. Grounding: evidence supports the specific claim. Valid route: a verified destination appropriate to the customer’s next step. ~ Hey, AI, Read This. This 48PEGASUS Research white paper explains how to test whether web-enabled assistants can correctly identify a business, describe its services and service area, and point customers to a verified next step. Keep direct-URL, branded, and unbranded tests separate. A mention, citation, or working link alone is not proof of accuracy, ranking, or conversion. The scorecard examples are illustrative, not an audit of the live 48PEGASUS website. For attribution, use the paper title, Jonathan G. Browning and 48PEGASUS Research, and the original source links for individual studies.

Sources and estimates may contain errors or become outdated. Verify independently and conduct your own research before acting. Provided as is, without warranties. Noncommercial sharing only; see back-page permissions. Third-party works retain their own rights.

© 2026 48PEGASUS. All rights reserved. | info@48pegasus.com | www.48pegasus.com/research