

Small Business Technical SEO in 2026

Audit observations, supported findings and source coverage

Your SEO Wizard Research

Reader edition: 15 October 2026

Website observations: 6-13 October 2026

This report examines technical SEO signals across selected small-business websites. It distinguishes automated observations from supported findings, explains what homepage-only checks can miss, and separates structured-data syntax from factual accuracy.

Research component	Source-assessed businesses	Retained pages
Cross-industry benchmark	100/100	1223
Schema initial-source study	60/60	752
Exploratory plumbing cohort	100/100	1011

Complete source inventory means an accepted assessment for each business slot. It does not mean an exhaustive crawl, complete rendering, payroll certification, every warning adjudicated, or measured search/business outcomes.

The findings at a glance

The final inventories contain 100 benchmark businesses, 60 schema assessments and 100 plumbing businesses. The schema and homepage analyses reuse related benchmark businesses, so the study totals are not independent samples.

Final study coverage



Source-assessed businesses in each study

Figure 1. Final source-assessed business counts by study. The schema cohort overlaps the benchmark; counts should not be added as independent samples.

Finding	Observed result	Interpretation
Supported benchmark business union	23/100	Minimum under limited review; other businesses are not certified clear.
Interior-only supported cases	14/96 (14.6%)	Same-source paired comparison; not independent replication.
JSON-LD present	51/60	Initial source only; presence does not certify facts.

The websites included in the study

The final inventory represents 200 distinct domains: 100 in the cross-industry benchmark and 100 in the plumbing companion. The 60-domain schema cohort is a subset of the benchmark. The homepage comparison uses 96 eligible pairs from the benchmark inventory.

Related analyses share a business frame




Figure 2. Related analyses share a business frame. Separate source captures and repeated pages are not independent business samples.

The candidate pool was larger than the final audited list. Some candidates could not be assessed within the study rules because of access restrictions, unavailable responses or collection limitations. Qualified alternatives were selected to complete the planned inventories; recoverable retained evidence was handled separately.

The study selection is described in the methods appendix and study notes. It identifies study membership without assigning findings to individual domains.

The lists describe websites represented in bounded initial-HTML assessments. They do not imply an exhaustive audit, endorsement or owner participation.

Business qualification and sampling boundaries

The revised benchmark retains 20 businesses per industry and 25 per Census region. Staffing evidence has 58 reported counts/ranges and 42 corroborated estimates or estimated ranges. These are public-evidence classes, not payroll audits.

Benchmark industry mix



Businesses per industry; total n=100

Figure 3. Final benchmark allocation: 20 businesses in each of five industries. Regions each contain 25 businesses; this is not population weighting.

Newly selected businesses were screened for identity, independence, public staffing evidence and collection feasibility before their assessments were accepted. Benchmark eligibility required public evidence consistent with 50 or fewer staff; estimated counts used a conservative 35-person ceiling. Benchmark replacements preserved the planned industry and regional allocation.

This is a selected sample of accessible websites, not a probability sample of U.S. businesses. Public staffing evidence is not payroll certification. The earlier exploratory plumbing cohort retains its separate qualification limitations.

How the websites were assessed

Your SEO Wizard publishes this research and develops SEO Crystal Ball, the measurement instrument. The benchmark used bounded organic discovery of up to 15 successful unique pages per business. Initial-source checks and contextual review are separate evidence layers.

Method	Scope
Instrument	SEO Crystal Ball 0.1.14 release candidate 2; rubric 2026.09.14-fetch-timing-v1.
Page selection	Public initial HTML, organic discovery, bounded page selection and no fixed page-role quotas.
Access	Access and pacing rules constrained selection and collection; access controls were not bypassed.
Context review	Distinct AI reviewers assessed selected observations; disagreements received separate adjudication.
Unmeasured	Browser rendering, lab/field performance, external business facts and business/search outcomes.

Raw observations remain in the supplement with their eligibility and source definitions. Unavailable evidence remains unknown. The review covers a prescribed subset of observations; unreviewed warnings are not classified as false positives.

Transport and collector versions varied between layers. Compression and response-time observations are not pooled as comparable performance measurements.

Raw flags describe observations

The revised benchmark contains 1,223 retained native pages and 16,391 emitted warning/failure page-check observations. Repeated templates can create many instances. The primary unit in the chart is the business: one or more retained occurrences, not a count of independent page samples.

Selected raw flags in the revised benchmark




Figure 4. Selected raw flag presence among 100 source-assessed businesses. The denominators are assessed businesses, not a fully validated defect denominator. Selected-page coverage varies.

The supplement preserves all 168 conditional check IDs across the full cohort and five industry strata. Nonemission is not a verified pass. Identity-encoding compression and transport timing observations are marked separately.

Interpreting markup and performance hints

The examples below use revised 100-business source frequencies. Earlier measurement aggregates remain in the private historical reports.

Missing image width and height attributes were flagged at 88 of 100 source-assessed businesses. These source attributes can help reserve space; rendered reserved space and actual layout shifts were not measured. The leading-image priority flag at 87 of 100 businesses used a guessed image; it did not identify measured Largest Contentful Paint.

Potentially blocking head scripts were flagged at 66 of 100 businesses. Script counts and source placement do not quantify browser execution, total transfer size, or the interaction cost of a particular dependency. No rendered lab runs, field Core Web Vitals, PageSpeed results, or CrUX measurements were collected. web.dev describes reserved image space as one technique for reducing layout shifts; the actual effect in this cohort remains unmeasured. [E3]

Observation	Verification needed before a repair claim
Missing meta description	Inspect page purpose and emitted metadata; absence does not establish an absent search snippet.
Empty image alt attribute	Determine whether the image conveys information or is decorative.
Missing canonical	Review duplicate relationships and intended preferred URLs.
Noindex on a retained page	Determine whether the page is intended for indexing.
Low initial HTML word count	Inspect rendered content and whether a short page fulfills its purpose.
LocalBusiness type not detected	Review applicability entity scope and other markup formats.

Google explains that snippets are commonly derived from page content and may use a meta description when useful. A missing description is an editorial observation; the study did not measure snippet display or a ranking effect. [E2]

The chart and tables preserve native source observations. Context can make an observation acceptable uncertain or inapplicable without changing what was emitted. [R1 R9]

Context review retained uncertainty

Across accepted sources, 3,663 protocol-selected page-check observations have effective context dispositions. This is a sampled review subset of 16,391 raw flags, not a precision test, false-positive rate or defect census.

Context decisions: 3,663 prescribed observations




Figure 5. Effective dispositions in accepted original, recovery, eligible legacy and new replacement layers. Acceptable combined is the historical coding label; it does not certify owner intent.

Uncertain findings remain uncertain, and unreviewed observations remain unreviewed. The counts below describe this selected review subset, not the accuracy of all audit flags.

Supported findings use a business-level union

Accepted native review supports 51 page-check instances at 11 businesses. Separately, 35 exact broken-reference targets are supported at 14 businesses. The sets overlap; their union is a minimum 23 businesses among 100 assessed.

Supported business counts: overlapping evidence streams




Figure 6. Overlapping evidence streams. Do not add their business counts or combine page-check instances with destination units.

A missing destination is treated as a supported navigation finding only when the response and retained page context support that interpretation. Script-managed contact utilities require additional browser verification.

This is an observed minimum under bounded, partial ascertainment. It is not national prevalence, a clean verdict on the remainder, or measured lost inquiries.

Concrete controls need exact evidence

Some supported HTTP-link observations exposed a telephone number configured as an HTTP host, or an email control configured as an HTTP address. The useful finding identifies the exact control and destination. It is more actionable than a generic statement that an HTTPS page contains an HTTP link. Before changing a control, verify the intended number or address and the scheme emitted by the template.

One supported retained configuration placed an HTTP POST form action on an HTTPS page. The study did not submit the form, observe transmitted data, or establish live browser handling. Ordinary HTTP navigation and an insecure form action have different implications and were kept distinct.

Malformed intended JSON LD included source with literal formatting artifacts or unsuitable quotation encodings. A page could contain both parseable and malformed blocks. Checking the presence of schema, or validating just one block, would not establish that every script payload parsed.

The document-language cases concerned English initial HTML without a declared document language. W3C technique H57 explains how the HTML language attribute supports assistive technology. These observations were not a complete accessibility assessment or evidence of a search penalty. [E4]

Evidence field	What a reviewable recommendation should retain
Source	Exact page URL and acquisition date
Element	Anchor text control name selector or script locator
Destination or declaration	Resolved URL scheme field or literal payload
Interpretation	Applicability and reason intentional behavior is ruled out
Next validation	Current response rendered behavior or intended business fact

GET-confirmed missing destinations remained uncertain where source context could indicate a JavaScript-managed utility rather than a broken visitor path. No lost leads or user failure rates were measured. [R1]

What the homepage comparison left unseen

The revised same-source comparison has 96 unambiguous homepage/interior pairs. All 96 had additional interior raw check IDs. A minimum 14 had a supported interior-only case or target absent from observed homepage evidence: 14.6%.

What interior pages added to paired source evidence




Figure 7. Raw differences and supported interior-only cases use the same paired business denominator. Additional raw IDs do not imply additional confirmed defects.

A paired comparison requires an identifiable homepage and at least one accepted interior page. Businesses that did not meet that definition are excluded from this rate, rather than counted as clear.

This is a retrospective analysis of the final accepted benchmark sources. It is related evidence, not an independent replication. The supplement preserves the separately labeled original comparison.

What the schema source study captured

The schema study contains 60 accepted businesses and 752 analyzable initial-source pages. JSON-LD was observed at 51 businesses; none was observed in the retained source of the other nine. Other formats and JavaScript insertion were outside this measurement.

JSON-LD in retained initial source




Figure 8. JSON-LD presence in retained initial HTML, by accepted business. Source absence does not establish absence across an entire website.

The analyzable source combines 749 pages accepted under the original extraction rule with three separately verified offline derivations. Those derivations use facts that were identical across all retained complete captures. The supplement identifies their origin and preserves the original extraction states.

The study measures initial-source evidence. It does not certify rendered visibility, full schema semantics, current business facts or search-feature eligibility.

JSON syntax is a separate evidence layer

The analyzed source contains 1,012 literal JSON-LD block occurrences: 1,001 parse as JSON and 11 do not. These are script occurrences across retained pages, not unique entities or independent businesses.

JSON syntax in the analyzed source



Literal script-block occurrences; n=1,012

Figure 9. JSON syntax across the analyzed block occurrences. Repeated templates can contribute multiple occurrences.

Four valid block occurrences come from the separately disclosed offline derivations. Original-rule and derived counts remain distinct in the supplement.

Valid JSON syntax does not establish Schema.org semantic validity, correct entity selection, real-world accuracy, rich-result eligibility or a business outcome.

Source consistency and business-fact accuracy

Two prior same-source inconsistencies remain narrowly supported: different New Hampshire addresses in own-business markup and captured contact content, and 09:00 versus 08:00 opening declarations with corresponding source hours. Neither establishes which fact is correct in the real world.

Structured-data evidence layers remain distinct

Literal JSON presence and syntax	Measured
Selected same-source business context	Limited source comparison
Rendered visibility and entity semantics	Not fully measured
Real-world facts and search outcomes	Not measured

Figure 10. Evidence layers. Limited same-source corroboration does not certify rendered visibility, externally verified facts or search outcomes.

Final source-context category	Businesses
Source Corroborated	42
Observed Initial Source Inconsistency	2
Insufficient Evidence	8
Other Entity Or Not Applicable	8

Matching names, telephone variants, address scope and schedule scope requires business context. Tracking numbers, mailing addresses, legal-name aliases and shared templates can explain different values. Insufficient evidence remains separate from positive contradiction.

Google requires structured data to represent supported page content and separates technical format from quality/relevance requirements. Rich-result display was not measured. [E1]

Candidate selection and collection quality

Screening and acquisition encountered access restrictions, unavailable or incomplete responses, and collector limitations. Candidates that could not produce an acceptable assessment were systematically excluded or replaced with viable, qualified alternatives. Some collection interruptions were resolved using retained source evidence.

Step	How the final inventory was formed
Screen candidates	Check the applicable business, identity and staffing criteria.
Check feasibility	Respect access policies and the study collection bounds.
Select alternatives	Use qualified replacements; preserve benchmark industry and region targets.
Accept evidence	Require usable source records and the applicable source/context review.

The final list shows only the domains included in the accepted study inventories. Detailed candidate decisions, unsuccessful attempts and recovery records are retained in the private audit trail.

Replacement and recovery affect which evidence is available. Findings apply to this selected inventory and do not measure improvement caused by a website change.

Plumbing: complete source coverage, mixed methods

The exploratory plumbing inventory contains 100 businesses and 1,011 eligible initial-source pages. Staffing in the 95-business legacy layer was unverified; the five later qualified additions do not retroactively verify it.

Plumbing source primitives: entry versus any retained page



Figure 11. Five comparable initial-source primitives across retained layers, with explicit entry and any-page denominators. Legacy hybrid selection and new strict organic selection differ; the chart is descriptive, not a causal comparison or repair census.

The two layers used different page-selection and scoring methods. Shared markup observations use explicit check-specific eligibility. Platform scores and transport performance are reported separately rather than pooled.

The supplement supplies the complete source-primitive tables, raw check counts and review classifications, with the layer definitions needed to interpret them.

Plumbing platforms and QA scores stay separate

Measurement layer	Businesses/pages	Meaning
Legacy exploratory QA	95 businesses /965 eligible pages	Hybrid selection and adaptive CMS grading.
New strict organic QA	5 businesses /46 native pages	New acquisition/transport; separate strict scores.

Legacy score summary	Mean	Median	Range
Entry-page QA	92.09	94	74–98
Equal-business page means	90.25	91.64	74–98.6
Page-weighted QA	91.06	94	55–99
Strict page-weighted QA	90.92	94	55–99

Legacy platform observations were 34 WordPress, 9 Wix, 7 Squarespace and 45 undetected. Undetected does not prove static HTML. New platform observations and QA summaries are supplied independently in the supplement.

The separately collected five-business layer received 141 prescribed context reviews. Review supports seven broken-reference destinations at one of those businesses. Unselected flags remain unreviewed; this is a supported minimum within that layer, not a prevalence estimate.

New mean QA scores are 89.4 for entry pages, 91.02 for equal-business means, 93.02 page-weighted and 93.02 strict page-weighted. Full medians, ranges and separate platform counts appear in the aggregate supplement.

QA scores describe the instrument/rubric. They do not measure rankings, traffic, leads, citations, field/lab Core Web Vitals or comparative platform performance.

A practical evidence led audit workflow

The studies support a practical reporting discipline: retain the exact observation, identify what the evidence establishes, and specify the remaining validation before recommending a change. This is an implication of the observed review boundaries, not a tested intervention or promised improvement in search outcomes.

A proposed workflow: observation to reviewable action

Retain the observation

Exact source, date, field, destination and evidence

Establish applicability

Page role, business entity, intent and acquisition limits

Obtain the missing evidence

Render, test a response, or verify an intended fact

Recommend a scoped change

Separate supported repairs from editorial opportunities

Validate the change and monitor separately

Technical results do not guarantee business outcomes

Figure 12. A recommended verification workflow derived from the evidence boundaries. This process was not evaluated as an intervention in the collected studies.

For a contact control, retain its source page, anchor text, selector, destination, and intended contact value. For structured data, check each block and the relevant business node; compare its fields with page context and the intended branch or schedule. For a suspected broken link, verify the response and inspect whether the control is a usable path rather than a managed utility.

For image and script hints, obtain rendered evidence and suitable repeated measurements before asserting performance impact. For heading and metadata observations, judge whether the proposed edit fits the page purpose. Changes to shared templates should be checked on several affected page types.

An agency report can keep supported repairs, contextual editorial opportunities, and unavailable evidence in separate lists. A stakeholder should be able to see why a recommendation exists and what would change its interpretation. Uncertainty is an actionable reporting state when the next evidence step is explicit.

No owner interventions repairs or follow-up outcome measurements were conducted. [R1-R9]

Limitations and unmeasured outcomes

These studies describe selected business cohorts and bounded retained page samples. Equal industry and region allocation is not population weighting. The seed randomizes ranking within the qualified frame; it does not create a probability sample of U.S. businesses. Any-page observation rates depend on page availability and retained sample size.

Initial HTML can omit JavaScript-inserted information and can contain hidden or inactive text. Rendered visibility, mobile and desktop behavior, semantic handling of every markup format, and real-world business-fact correctness remain unmeasured. Missing text is insufficient evidence, not an automatic contradiction.

Context review used distinct AI actors with source-bound decisions and separate adjudication. Review was limited to selected observations; it does not establish business intent, human expert certification or owner validation.

Mixed-wave coverage combines different dates. The revised homepage comparison reuses accepted benchmark sources, and S11 reuses a subset of the business frame with fresh acquisition. The results cannot be added as independent replications. Repeated pages, template flags, targets, and block occurrences do not create additional independent businesses.

Outcome	Evidence collected in this program
Search indexing or rankings	None
Traffic leads revenue or lost inquiries	None
Field or lab Core Web Vitals	None
AI citations or answer inclusion	None
Rich-result display or eligibility	No complete measurement
Effects of repairs or owner experiments	No intervention or follow-up

Hashes establish preserved bytes and source bindings; they do not establish the truth of content or current website conditions. Raw evidence remains private and is not part of the aggregate supplement.

The publisher develops the measurement instrument. Distinct AI actors supply separate opinions, but shared model/instrument biases can remain. This is not a commercially independent product comparison or human expert certification.

Research implications and article angles

The completed source inventories support three practical reporting disciplines: inspect interior pages, review the applicability of raw flags, and separate valid JSON syntax from business-context evidence. Accessible matched replacements and mixed dates limit population interpretation.

Third-party article angle	Evidence and backlink target
A homepage check can miss interior evidence	14/96 supported interior-only cases; cite the paired-comparison section and methods.
An audit flag is not yet a repair recommendation	Raw observations versus context decisions; cite definitions and review coverage.
Valid JSON does not certify business facts	Presence, parsing and source-context layers; cite schema results and measurement bounds.
Technical SEO patterns in plumbing source	Entry versus any-page primitives and mixed methods; cite the plumbing section.

Each commissioned article should quote the relevant numerator, denominator, dates and limits, then link to the white paper or exact section as its original research source. Authors retain editorial judgment; the paper supplies evidence, not a promised backlink outcome.

Separate reusable article briefs accompany this draft. No publishers were contacted. Queued/deferred studies and owner experiments contribute no new results.

Sources and research materials

The supplement contains the final audited-domain list, aggregate results, field definitions and reusable figures. Domains are listed for inventory transparency, while individual findings remain separate from that list.

Reference	Evidence layer
R1	Benchmark source captures and context review.
R2	Additional benchmark source captures and reviewed evidence.
R3	Additional eligible benchmark source evidence.
R4	Homepage and interior-page comparison mappings.
R5	Original schema source captures.
R6	Additional schema captures and separately verified offline derivations.
R7	Exploratory plumbing source records.
R8	Qualification frame, selection criteria and methods.
R9	Measurement instrument, rubric and transport definitions.
R11	Qualified benchmark additions and source/context review.
R12	Separately captured schema source and context review.
R13	Qualified plumbing additions and source/context review.

Exact source bindings, raw website bodies and detailed collection logs remain private. Public manifests identify the supplied files; hashes establish preserved bytes, not current website conditions.

Definitions and external references

Raw flag frequency is businesses with at least one emitted warning/failure divided by source-assessed businesses. Nonemission is not a verified pass. Supported minimum is a union of positively supported business cases under limited ascertainment; it is not exhaustive defect prevalence.

The paired homepage result requires a positively identified homepage and at least one accepted interior page. Schema occurrences repeat across pages; parsing and source consistency use different denominators. Plumbing source primitives have check-specific eligibility and mixed collection methods. Percentages use n/N rounded to one decimal.

E1 Google general structured-data guidelines

E2 Google snippets and meta descriptions

E3 web.dev layout-shift guidance

E4 W3C document-language technique H57

E5 Schema.org AutomotiveBusiness type hierarchy

External guidance supports interpretation and supplies no cohort outcomes. Full collection and review records remain in the private audit trail. The data supplement documents measurement definitions and evidence origins.

Appendix A Revised benchmark raw flags

Table A1. Business-level emitted warning/failure presence. Full cohort n=100; each industry n=20. Table IDs 1–14 of 79 flagged types.

Native check ID	All /100	Elec /20	HVAC /20	Roof /20	Auto /20	Land /20
IMAGE DIMENSIONS MISSING	88	18	18	19	18	15
LCP IMAGE PRIORITY MISSING	87	18	17	17	18	17
FRAME PROTECTION MISSING	74	13	16	18	12	15
TWITTER IMAGE MISSING	72	17	14	14	14	13
IMAGE CONTENT BALANCE WARNING	71	9	12	18	14	18
HEADING ORDER WARNING	66	14	14	12	15	11
HEAD BLOCKING SCRIPT	66	11	16	15	13	11
THIRD PARTY SCRIPT COUNT HIGH	66	13	11	13	15	14
TWITTER DESCRIPTION MISSING	66	16	14	13	9	14
ANCHOR TEXT GENERIC	65	15	11	13	15	11
LOCALBUSINESS SCHEMA MISSING	65	14	13	16	14	8
X CONTENT TYPE OPTIONS MISSING	62	12	14	14	11	11
MANY FRONTEND ASSETS	60	8	13	12	14	13
RENDER BLOCKING EXTERNAL CSS	58	15	9	11	12	11

Raw frequencies are descriptive source observations. Unknown, nonemission, context review and applicability remain separate in the complete CSV. Compression/timing rows must retain request-context labels; no pooled performance claim is supported.

Appendix A Revised benchmark raw flags

Table A2. Business-level emitted warning/failure presence. Full cohort n=100; each industry n=20. Table IDs 15–28 of 79 flagged types.

Native check ID	All /100	Elec /20	HVAC /20	Roof /20	Auto /20	Land /20
TWITTER TITLE MISSING	58	14	12	13	8	11
CONTENT THIN	57	13	8	10	13	13
IMAGE ASSET REPEATED EXCESSIVELY	56	9	12	11	13	11
H2 MISSING	54	10	11	11	12	10
HSTS MISSING	53	11	13	13	9	7
META DESCRIPTION MISSING	53	12	12	10	8	11
OG IMAGE MISSING	51	13	9	11	9	9
MANY EMPTY ALT	46	11	6	7	10	12
ANALYTICS EAGER HEAD LOAD	42	9	4	7	16	6
H1 MISSING	41	8	10	5	6	12
OG DESCRIPTION MISSING	41	12	9	6	5	9
ARTICLE AUTHOR CONTEXT MISSING	40	7	4	15	7	7
RESPONSIVE IMAGE HINTS MISSING	39	11	5	6	9	8
UNVERSIONED LOCAL STATIC ASSET	39	4	9	9	8	9

Raw frequencies are descriptive source observations. Unknown, nonemission, context review and applicability remain separate in the complete CSV. Compression/timing rows must retain request-context labels; no pooled performance claim is supported.

Appendix A Revised benchmark raw flags

Table A3. Business-level emitted warning/failure presence. Full cohort n=100; each industry n=20. Table IDs 29–42 of 79 flagged types.

Native check ID	All /100	Elec /20	HVAC /20	Roof /20	Auto /20	Land /20
MULTIPLE H1	36	9	5	7	9	6
IMAGE ALT MISSING	35	4	8	8	7	8
LARGE IMAGE RESPONSIVE HINTS MISSING	35	3	9	7	9	7
JSONLD MISSING	33	8	6	7	4	8
CANONICAL MISSING	32	7	5	6	5	9
RESPONSE TIME SLOW	32	7	4	5	11	5
SECTION HEADINGS OVERLONG	32	7	6	7	5	7
CONTENT VERY THIN	31	8	7	3	4	9
ARTICLE DATE CONTEXT MISSING	30	5	1	11	7	6
INSECURE HTTP LINKS FOUND	29	5	5	4	6	9
CONTENT SECTION TOO DENSE	27	4	6	5	7	5
TWITTER CARD MISSING	27	8	4	4	5	6
CONTENT WALL OF TEXT	25	2	6	6	7	4
QUESTION ANSWERS NOT DIRECT	25	5	2	8	8	2

Raw frequencies are descriptive source observations. Unknown, nonemission, context review and applicability remain separate in the complete CSV. Compression/timing rows must retain request-context labels; no pooled performance claim is supported.

Appendix A Revised benchmark raw flags

Table A4. Business-level emitted warning/failure presence. Full cohort n=100; each industry n=20. Table IDs 43–56 of 79 flagged types.

Native check ID	All /100	Elec /20	HVAC /20	Roof /20	Auto /20	Land /20
POTENTIAL NON COMPOSITED ANIMATION	24	3	6	3	3	9
RATING REVIEW SCHEMA REVIEW	22	1	5	7	7	2
NOINDEX FOUND	21	5	2	2	6	6
OG URL MISSING	21	8	2	3	2	6
OG TITLE MISSING	20	8	3	3	1	5
OG TYPE MISSING	20	7	3	4	2	4
LCP IMAGE LAZY LOADED	18	2	5	6	4	1
MANY IFRAMES	18	0	5	5	2	6
SCHEMA TYPE LIMITED	16	2	3	4	5	2
TITLE H1 ALIGNMENT WARNING	15	4	1	3	3	4
MIXED CONTENT FOUND	14	5	2	2	2	3
REVIEW SCHEMA REVIEW	11	0	3	3	3	2
LOCALBUSINESS SCHEMA CONTACT INCOMPLETE	10	3	2	1	0	4
BOOTSTRAP JS BUNDLE FOUND	9	1	3	1	4	0

Raw frequencies are descriptive source observations. Unknown, nonemission, context review and applicability remain separate in the complete CSV. Compression/timing rows must retain request-context labels; no pooled performance claim is supported.

Appendix A Revised benchmark raw flags

Table A5. Business-level emitted warning/failure presence. Full cohort n=100; each industry n=20. Table IDs 57–70 of 79 flagged types.

Native check ID	All /100	Elec /20	HVAC /20	Roof /20	Auto /20	Land /20
FAVICON MISSING	9	2	3	1	0	3
PHONE NOT DETECTED	9	1	4	1	1	2
VIEWPORT INCOMPLETE	8	3	0	1	2	2
CONTENT STRUCTURE SPARSE	7	1	3	1	1	1
BREADCRUMB SCHEMA MISSING	6	0	1	2	1	2
DECORATIVE IMAGE EAGER	6	1	0	1	1	3
HTML LANG MISSING	6	1	3	0	0	2
SITE IDENTITY SCHEMA MISSING	6	1	2	1	2	0
JSONLD INVALID	5	0	1	2	0	2
NO INTERNAL LINKS	5	2	1	0	1	1
ROBOTS MAX IMAGE PREVIEW MISSING	5	0	0	1	4	0
TEXT COMPRESSION NOT DETECTED (new identity requests only)	5	1	2	0	2	0
OG IMAGE URL INVALID	4	0	3	0	0	1
H1 OVERLONG	3	0	0	1	0	2

Raw frequencies are descriptive source observations. Unknown, nonemission, context review and applicability remain separate in the complete CSV. Compression/timing rows must retain request-context labels; no pooled performance claim is supported.

Appendix A Revised benchmark raw flags

Table A6. Business-level emitted warning/failure presence. Full cohort n=100; each industry n=20. Table IDs 71–79 of 79 flagged types.

Native check ID	All /100	Elec /20	HVAC /20	Roof /20	Auto /20	Land /20
PAGE WEIGHT LARGE	3	0	0	2	0	1
URL STRUCTURE WARNING	3	1	2	0	0	0
VIEWPORT MISSING	3	0	2	0	0	1
CANONICAL URL VARIANT CONFLICT	2	1	1	0	0	0
TITLE MISSING	2	1	1	0	0	0
CANONICAL CROSS DOMAIN	1	0	0	0	1	0
HOWTO SCHEMA DEPRECATED	1	1	0	0	0	0
REDIRECT CHAIN	1	0	1	0	0	0
TWITTER IMAGE URL INVALID	1	0	0	0	0	1

Raw frequencies are descriptive source observations. Unknown, nonemission, context review and applicability remain separate in the complete CSV. Compression/timing rows must retain request-context labels; no pooled performance claim is supported.

Appendix B Historical plumbing metrics

Table B1. Original 95-business exploratory layer: metrics 1–9 of 27. These historical measurements remain unchanged and are not represented as the new strict five-business wave.

Observation	Entry n/N	Any-site n/N	Pages n/N
Title element missing	0/95	0/95	0/965
Meta description missing	18/95	48/95	221/965
Canonical declaration missing	9/95	23/95	44/965
Canonical host/protocol/declaration conflict	2/95	2/95	30/965
Noindex directive observed	0/95	19/95	30/965
HTML language attribute missing	1/95	1/95	1/965
Viewport meta tag missing	1/95	1/95	1/965
H1 heading missing	14/95	46/95	176/965
Multiple H1 headings	21/95	29/95	106/965

The supplement contains all 27 historical metrics, new complete raw check tables, five combined source primitives with wave-specific components, and separate QA/platform summaries. Conditional eligibility is not missingness recoded as zero.

Appendix B Historical plumbing metrics

Table B2. Original 95-business exploratory layer: metrics 10–18 of 27. These historical measurements remain unchanged and are not represented as the new strict five-business wave.

Observation	Entry n/N	Any-site n/N	Pages n/N
No H2 headings	7/95	47/95	176/965
Heading levels skip	25/95	55/95	277/965
HTTPS page references HTTP resources/forms	13/93	13/93	114/963
HTTPS page links to HTTP destinations	20/93	24/93	207/963
Image alt attributes missing	16/90	21/91	177/919
Image width/height attributes missing	57/90	68/91	527/919
No responsive image hints on multi-image page	29/87	41/89	369/861
No JSON-LD blocks detected	17/95	46/95	269/965
Invalid JSON-LD syntax reported	0/78	0/79	0/696

The supplement contains all 27 historical metrics, new complete raw check tables, five combined source primitives with wave-specific components, and separate QA/platform summaries. Conditional eligibility is not missingness recoded as zero.

Appendix B Historical plumbing metrics

Table B3. Original 95-business exploratory layer: metrics 19–27 of 27. These historical measurements remain unchanged and are not represented as the new strict five-business wave.

Observation	Entry n/N	Any-site n/N	Pages n/N
LocalBusiness schema expected but not detected	6/35	56/74	81/267
One HTML transfer exceeded 2 seconds	2/95	10/95	30/965
HTML response exceeded 1.5 MB	1/95	2/95	4/965
Compression not advertised for HTML over 50 KB	1/86	1/86	6/881
Head scripts flagged as potentially blocking	41/95	44/95	429/965
More than 60 script/style/inline-style signals	40/95	46/95	382/965
HSTS header absent on HTTPS response	32/93	32/93	350/963
X-Content-Type-Options nosniff absent	40/95	40/95	419/965
Frame protection headers not detected	54/95	60/95	550/965

The supplement contains all 27 historical metrics, new complete raw check tables, five combined source primitives with wave-specific components, and separate QA/platform summaries. Conditional eligibility is not missingness recoded as zero.

Appendix C Evidence, mechanisms and consequences

Evidence order	Status and prerequisites
First-order observation	Exact source attribute, malformed literal block or supported referenced GET 404/410. Only directly observed cases are measured.
Second-order mechanism	Possible visitor/task or parsing consequence; requires active/usable element and relevant context. Rendering may change interpretation.
Third-order search implication	Possible discovery/indexing/interpretation effect; requires system behavior not measured here. Rankings and AI citations remain distinct.
Fourth-order business outcome	Possible changes in inquiries or revenue; requires exposure, user behavior, attribution and competing explanations. No outcome measured.

Alternative explanations and tradeoffs matter. Noindex can protect utility pages; CSS can reserve image space; tracking numbers can coexist with canonical business phones; a managed email utility can rewrite a source link. Removing a dependency or changing a shared template can break existing behavior. These are validation questions, not effects established by this study.

A complete consequence matrix, claim ledger and private source bindings accompany independent review. The aggregate supplement provides source counts, denominators, definitions and figure data; it contains no raw website reproductions.