SERPIFAI FIELD GUIDE · AEO / GEO TECHNICAL PROTOCOL

Technical Answer Visibility Field Guide

A fail-closed protocol for access, indexability, canonical identity, rendering, structured-data parity, main-content semantic exposure and handoff before you blame content or authority.

Core rule
Technical SEO opens and labels the door; it does not decide whether an answer system walks through it.

Canonical method
SCOPE → FETCH → RECONCILE → RENDER → VALIDATE → HAND OFF

Updated 17 August 2026 | Portable handoff

1 · DECISION BOUNDARY

Technical readiness is a prerequisite state, not an AI citation score

A page is technically ready when its intended URL is reachable, publication intent is coherent, the preferred identity is consistent, substantive content survives rendering, machine-readable descriptions agree with visible content, and the main answer remains semantically exposed to the relevant processor. Broader accessibility and usability remain a separate publication-quality gate unless they also break that semantic exposure.

> **Passing the gate permits a downstream visibility test.** It does not prove retrieval, citation, recommendation, ranking, recurrence or commercial value.

The six layers

Access

Can the intended crawler and required resources reach the page?

Indexability

Do directives match the publication decision for the scoped platform?

URL identity

Do redirects, canonicals and internal signals point to one preferred URL?

Rendered parity

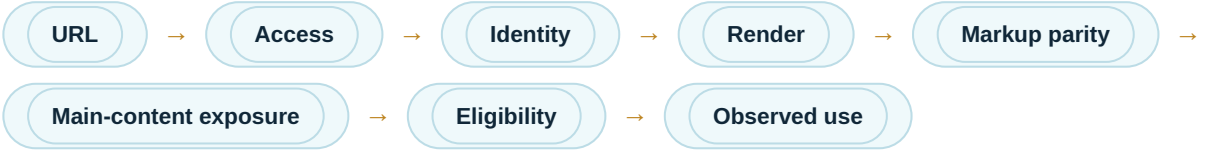
Does the substantive answer survive processing and resource loading?

Structured identity

Does markup describe what users can verify and the relationships that actually exist?

Main-content semantic exposure

Does the substantive answer remain available in a meaningful main-content structure rather than disappearing behind a failed component or processing boundary?



Later states do not repair earlier failures. Earlier passes do not guarantee later selection. Publication accessibility/quality remains a separate acceptance gate unless it also breaks main-content semantic exposure.

Use the common technical object, then add platform-specific evidence

Platform	Technical evidence to prioritize	What a pass means	What it does not mean
Google Search / AI features	Search technical requirements, indexed/ snippet eligibility, canonical processing, rendered main content, structured-data guidelines.	The page can participate in Google Search under the documented prerequisites.	That AI Overviews or AI Mode will select it for a given query.
ChatGPT search	OAI-SearchBot policy/access plus ordinary page access and coherent substantive content exposure.	The scoped site is not excluding the documented search crawler.	That ChatGPT will retrieve, cite, order or recommend the page.
Bing / Microsoft AI	Crawl/index health, preferred URL, rendered content; after eligibility, cited-page and grounding-query observations in AI Performance.	Technical prerequisites can be separated from downstream usage observations.	That citation count identifies a technical causal factor.
Multi-surface	Common URL/access/render/parity controls + surface-specific crawler and measurement evidence.	The page is a coherent technical object across the scoped checks.	That all engines consume robots, schema or page structure identically.

Do not universalize one engine's documentation. Google Search Central is authoritative for Google Search behavior. OpenAI crawler documentation is authoritative for OpenAI's documented crawlers. Schema.org defines vocabulary; it does not define every search engine's rich-result or AI-selection behavior.

When to use NOT APPLICABLE

Use it only when a check is genuinely irrelevant to the scoped page or platform, and record why. Core Access, Indexability, URL Identity, Rendering and Main Content Semantic Exposure accept only exact PASS / FAIL / UNKNOWN tokens. Structured-data parity can be N/A when no structured data is used. Platform crawler access can be N/A only where the scoped platform does not create a separate crawler condition; it must not be used to bypass an untested ChatGPT-search crawler state.

3 · STRUCTURED DATA

Validate description and parity - not a mythology about citations

Structured data has defensible jobs: describe entities and facts, support documented search features, and keep page relationships internally coherent. It is not a universal “AI citation” switch.

Question	PASS condition	FAIL / unresolved condition
Syntax / supported feature	Markup parses and, when targeting a documented feature, satisfies that feature's requirements.	Parser error, missing required properties, or unsupported use presented as guaranteed behavior.
Visible-content parity	Material claims/properties can be verified on the page.	Markup says something materially different from the user-visible page.
Entity identity	URLs, names, authorship and relationships refer to the intended real entities.	Wrong URL, stale entity, unsupported reviewer or contradictory identity.
Purpose	The markup has a real descriptive or documented platform purpose.	Types/properties added only because “more schema” is assumed to produce AI visibility.

> **Evidence boundary:** a 2026 Ahrefs study tracking pages that added JSON-LD found no meaningful citation uplift attributable to the addition in its sample. That challenges schema-causality claims; it does not prove schema has no descriptive or feature-eligibility value.

Three safe structured-data jobs

1. **Describe** what the visible page actually says.
2. **Qualify** for a documented search feature when relevant.
3. **Connect** page, organization, author, breadcrumb, product or other entities coherently.

Do not add unsupported reviewer, rating, product, person or organization claims to create machine-readable “authority.”

4 · OPERATOR PROTOCOL

SCOPE → FETCH → RECONCILE → RENDER → VALIDATE → HAND OFF

□ **SCOPE:** record exact preferred URL, target platform, locale, observation window, content type and publication intent.

□ **FETCH:** record status/access evidence, robots/access controls, authentication and required resource availability. Use platform-specific crawler evidence when the surface requires it.

□ **RECONCILE:** compare redirect target, canonical declaration, internal links, sitemap references and schema URLs. Stop on a material identity conflict.

□ **RENDER:** inspect the rendered main answer, not only raw source. Check JS/resource failures, consent layers, lazy rendering and whether the substantive answer remains semantically exposed in the main-content structure.

□ **VALIDATE:** test structured data for syntax/feature requirements and visible-content/entity parity. Record N/A with a reason when no structured data is used.

□ **HAND OFF:** record technical state, evidence locator, prohibited inference, smallest next action, owner and reopening trigger.

Minimum evidence record

Field	Why it matters
Scope	Prevents unlike platforms, locales or dates from being merged into one conclusion.
State	Use exact controlled tokens. Core technical prerequisites accept PASS / FAIL / UNKNOWN. N/A is permitted only for explicitly defined stages and requires a reason.
Evidence locator	Lets another reviewer find the test, URL inspection, log, render or validator output.
Observation detail	Explains what was actually seen and why it matters.
Owner	Routes the smallest corrective action to someone who can change the state.
Reopen trigger	Prevents a passed gate from becoming stale or an endless recurring audit.

Stop rule: after a repeatable technical PASS, stop technical optimization unless a defined trigger introduces a new defect. A publication-accessibility failure can still block release, but it does not reopen the answer-visibility technical hypothesis unless it also damages main-content semantic exposure.

5 · FAILURE ROUTING

First failure wins

State	Smallest next action	Likely owner
TECHNICALLY BLOCKED - ACCESS	Restore intended fetch/resource access and retest.	Developer / Technical SEO
TECHNICALLY BLOCKED - INDEXABILITY	Reconcile publication directives with the intended visibility state.	Technical SEO
TECHNICALLY BLOCKED - IDENTITY	Reconcile canonical, redirects and preferred-URL signals.	Technical SEO
TECHNICALLY BLOCKED - RENDERING	Make the substantive answer survive processing and re-capture evidence.	Developer / Product
TECHNICALLY BLOCKED - STRUCTURED DATA PARITY	Correct/remove markup that materially contradicts the page.	SEO / Engineering
TECHNICALLY BLOCKED - MAIN CONTENT EXPOSURE	Restore the substantive answer in a meaningful main-content structure and re-capture direct evidence.	Product / Web / Engineering
EVIDENCE INSUFFICIENT	Capture the missing direct evidence only.	Reviewer / Technical SEO
TECHNICALLY ELIGIBLE - NOT OBSERVED	Hand off to query fit, passage usefulness, source role, evidence and authority.	SEO / AEO / GEO

Separate publication routing: keyboard/focus, contrast, labels, responsive/zoom, reduced-motion and assistive-technology failures belong in PUBLICATION_QA. They can block publication acceptance without being described as a documented AI-platform eligibility failure.

Worked example

A fictional guide returns 200 and its Article JSON-LD validates. But its canonical points to an old campaign URL while internal links and schema reference the new guide. The first material failure is **identity**. Rewriting copy or adding schema types is premature. Repair the canonical conflict, wait for processing evidence, then reopen downstream observation. If the page still is not selected, technical ownership closes and the next hypothesis moves downstream.

Two mistakes prevented: rewriting content before fixing URL identity; and continuing technical work after URL identity is demonstrably repaired.

6 · STRATEGIC HANDOFF

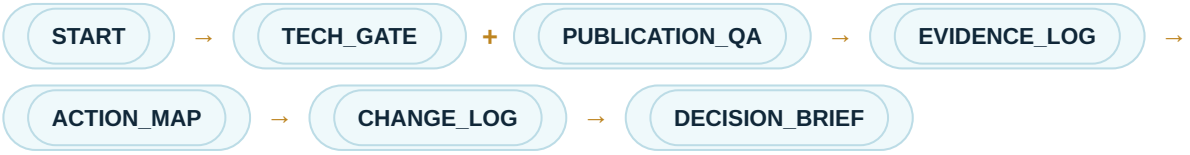
Turn the audit into a decision another person can use

Keep three states separate: **technical decision state**, **evidence completeness**, and **downstream visibility state**. A page can be technically eligible with no downstream observation; visible with a lingering defect; or correctly classified with evidence too thin for handoff.

Five-line executive brief

1. **Scope:** page + platform + locale + observation window.
2. **Technical state:** blocked / insufficient evidence / eligible.
3. **Evidence:** the locator that supports the state.
4. **Boundary:** what the state does not prove.
5. **Next owner/action:** smallest next step + reopening trigger.

Workbook workflow



The workbook is intended for repeated page checks, controlled states, evidence locators and handoff. It rejects blank scope, page/platform/locale/window mismatches, missing evidence, duplicate record IDs, wrong-case tokens and unjustified N/A states. The Decision Brief cannot become READY until exact page, platform, locale and observation window are governed. The first two repeated evidence rows and a late repeated row must be tested before native Excel acceptance.

> **Shareability rule:** if the record cannot tell another person what failed, why it matters, who owns it and what would close it, the audit is not yet a decision product.

7 · SOURCES AND LIMITS

Primary sources define platform behavior; market research generates hypotheses

Use the current version of these sources when the platform behavior matters to a decision.

- Google Search Central - AI features and your website: https://developers.google.com/search/docs/appearance/ai-features
- Google Search Central - Generative AI optimization guidance: https://developers.google.com/search/docs/fundamentals/ai-optimization-guide
- Google Search Central - Structured data introduction and policies: https://developers.google.com/search/docs/appearance/structured-data/intro-structured-data and /sd-policies
- Google Search Central - JavaScript SEO basics and canonicalization: https://developers.google.com/search/docs/crawling-indexing/javascript/javascript-seo-basics and /consolidate-duplicate-urls
- OpenAI - Overview of crawlers: https://developers.openai.com/api/docs/bots
- Bing Webmaster Blog - AI Performance in Bing Webmaster Tools: https://blogs.bing.com/webmaster/February-2026/Introducing-AI-Performance-in-Bing-Webmaster-Tools-Public-Preview
- Schema.org - Article / WebPage / BreadcrumbList vocabularies: https://schema.org/
- Ahrefs - schema additions and AI citations study: https://ahrefs.com/blog/schema-ai-citations/
- Semrush - technical SEO and AI-search observational study: https://www.semrush.com/blog/technical-seo-impact-on-ai-search-study/
- Sitebulb - Technical SEO for AI: https://sitebulb.com/technical-seo/
- Screaming Frog - structured-data extraction, validation and rendered/original HTML audit configuration: https://www.screamingfrog.co.uk/seo-spider/user-guide/configuration/

Evidence limit: a technical pass proves only the tested prerequisite state. Vendor studies are observational or market-method evidence unless their design establishes something stronger. Do not convert correlations or before/after chronology into a guaranteed platform mechanism.

Reopen the technical gate when

Reopen TECH_GATE for routing/canonical/template changes; crawler or robots policy changes; JavaScript/resource behavior changes; structured data or material page identity changes; a main-content semantic-exposure regression; or a changed platform/locale scope. Reopen PUBLICATION_QA separately for material keyboard/focus, contrast, label, responsive/zoom, reduced-motion or assistive-technology regressions.

Use this field guide with the SERPIFAI Technical Answer Visibility Workbook. The article remains the explanatory source; the workbook is the repeatable operating record.