

AI Visibility: The Evidence-Based Operating Playbook

From verifiable access to defensible source presence and buyer impact



HOW TO USE THIS RESOURCE

AI Visibility: The Evidence-Based Operating Playbook

B2B SaaS marketing leaders, SEO, content, and product marketing teams

Your practical outputs

- An evidence register for each buyer question
- A verified access and control matrix
- A reproducible citation observation protocol
- A 90-day plan with owners and stop rules

Read, choose, apply

Start with the overview and decision rules. Choose a concrete constraint in your business. Work through the relevant templates using your own evidence, record open assumptions and agree a review with the accountable people. The complete online edition provides linked sources and contextual navigation.

Editorial context

Explicitly fictional cases and example figures are working aids, not client results or universal benchmarks. Frameworks structure decisions; they are not scientifically validated scores. Platform capabilities may change. Sources and review context are disclosed.

Read online: [AI Visibility: The Evidence-Based Operating Playbook](#)

CONTENTS

1. Visibility as a decision problem	4
2. Keep the five stages distinct	5
3. Audit access and controls	6
4. Clarify entities and product truth	7
5. Choose buyer questions by evidence value	8
6. Produce useful source pages	9
7. Observations rather than prompt theater	9
8. Connect visits and pipeline carefully	10
9. Interpret failure patterns correctly	11
10. FICTIONAL SaaS example: ApprovalCloud	12
11. A 90-day operating rhythm	13
12. Continue, correct, or stop	13
13. Operating worksheets	14
14. Sources and limits of the playbook	15
15. A complete weekly AI visibility assignment	16
Turn this resource into a team brief	18
Sources & further documentation	19

Edition: October 24, 2026. A practical B2B SaaS resource by Patrick Wings. The frameworks are original operating methods; platform statements link to primary documentation. Examples are explicitly fictional. Use this document as a working reference: enter your own baseline, reproduce the observations described, and only then decide on the next investment. The English and German editions cover equivalent decisions and limitations.

1. Visibility as a decision problem

AI visibility means observable brand or source presence in selected search-backed answer experiences. It is neither a universal platform metric nor a claim of complete measurability. Start with a specific buying committee question, such as the prerequisites for auditable supplier approval. Your objective might be an accurate product description, a useful source link, or a qualified visit. Each requires different evidence. This guide uses an original Access-Assertion-Selection-Action model: access can be checked, assertions can be edited, selection remains controlled by an external system, and the subsequent action needs separate measurement. Treat the model as an operating aid, not an engine specification.

Establish a usable decision contract

Write the business decision before collecting data: which buyer needs an answer, what action becomes possible, and who can approve the change? Set a review date and a maximum investment of time. Record the current page or process so another colleague can reproduce the starting point. Separate a public documentation statement from your own interpretation and from an untested hypothesis. This prevents a reasonable editorial practice from becoming an invented engine requirement.

The working unit is one buyer question, one relevant page or page family, and one observable outcome. Do not combine a product launch, a redesign, and a new measurement scheme in the same first test. If these changes cannot be separated, label the result observational. A positive result can justify another iteration, but it cannot identify the contribution of each intervention.

- Copy into the brief: buyer, decision, scope, baseline date, owner, approval, review date.
- Record the strongest alternative explanation for any expected improvement.
- Stop if nobody can explain which decision this work will change.

Worksheet

Working field	Entry or observation	Next decision
Access	Public page is delivered	Engineering checks request and rule
Assertion	Capability agrees with evidence	Product marketing approves facts
Selection	Source appears in a defined run	Record observation with denominator
Action	Qualified visit or inquiry	Check analytics and CRM separately

Primary reference for technical context: [Google: AI features and your website](#).

2. Keep the five stages distinct

Crawling is fetching; indexing concerns inclusion and processing; ranking concerns selection and ordering in a search context. Retrieval obtains information for a particular task; citation visibly attributes a source in an answer. These processes can interact but are not interchangeable. A server log proves a request, not index inclusion. A search result does not prove citation. A citation does not prove that every statement in an answer came from your page. Keep five separate fields in the evidence register. Where direct observation is unavailable, retain unknown rather than reasoning backward from a later stage and presenting the inference as verified fact.

Build an evidence register

Keep evidence in a small register rather than in presentation screenshots alone. Each row needs the question, exact URL, observation date, language, environment, observed behavior, interpretation, and confidence. Store an artifact such as a response header, rendered page excerpt, export, or saved answer when your access and retention rules permit it. An artifact should show what happened, not merely the analyst's conclusion. Remove personal information and credentials before sharing.

Use three confidence labels: verified observation, supported interpretation, and open hypothesis. These labels describe the evidence, not the commercial importance of the problem. A verified broken link may be low impact; a plausible security concern may need urgent review. Assign a named person to resolve uncertainty. An empty field means unknown, not zero or success.

- Sample at least two relevant page templates before generalizing across the site.
- Keep excluded and failed observations visible with a reason.
- Stop aggregation when language, environment, or reporting definitions are incompatible.

Worksheet

Working field	Entry or observation	Next decision
Focus	Five status fields instead of one visibility score	Replace with your own data
Evidence	URL, date, and reproducible observation	Record confidence and open question
Acceptance	Owner and observable target state	Schedule the team review

Primary reference for technical context: [Google: AI features and your website](#).

3. Audit access and controls

Create a matrix of domain, path, bot purpose, and intended permission. Search access and model training are separate decisions. OpenAI's documentation distinguishes OAI-SearchBot and GPTBot, for example; review the current description before changing policy. A user-agent string alone does not authenticate a bot. Compare relevant requests with published verification information and examine blocks in the CDN or WAF. Robots.txt is neither authentication nor protection for confidential information. Public marketing pages and protected customer data require different technical rules. Have engineering and the appropriate security owner review exceptions, including their scope and rollback route.

Inspect the actual delivery path

Compare the browser experience with the delivered response and, where available, a search engine's inspection tools. Record the requested URL, redirect destination, HTTP status, relevant headers, rendered main content, and the link that led to the page. A page looking correct in your own browser does not establish crawler access. Conversely, a missing page in a manual answer test does not establish a crawling problem.

Check a normal visit and a fresh session without assuming that a logged-in state or accepted consent carries over. Test desktop and mobile layouts for the evidence and next action, not just visual polish. If a security layer challenges automation, involve the infrastructure owner rather than weakening controls globally. A targeted, documented access correction is preferable to an unreviewed exception affecting the entire application.

- Save one successful and one failing request when reproducing inconsistent delivery.
- Identify the earliest observable failure before editing content.
- Stop publication if essential claims or navigation disappear outside a privileged session.

Worksheet

Working field	Entry or observation	Next decision
Public resources	Search access intended	Check bot purpose and actual delivery
Model training	Separate company decision	Review current GPTBot documentation
Customer application	Authentication required	No blanket public permission
WAF exception	Verified legitimate access only	Test narrowly and name rollback owner

Primary reference for technical context: [OpenAI: Overview of crawlers](#).

4. Clarify entities and product truth

Define a small public fact sheet: official brand name, provider, product category, intended role, core capability, supported integrations, and documented limits. A useful page explicitly states whether a capability is in the standard plan, an add-on, or available only on request. Missing restrictions create false expectations even when the brand is mentioned frequently. Check similar company and module names. Clear contact and organization pages help visitors identify the provider. Structured data must accurately represent visible information; schema markup is neither a substitute for understandable prose nor a guarantee of AI presence. Resolve contradictions before adding more metadata.

Public product reconciliation

Open the product page, pricing page, and technical documentation side by side. For each important capability, mark the same four facts: available action, authorized role, plan, and prerequisite. A different restriction can turn a correct short description into an inaccurate promise. Give each contradiction a facts owner. Decide which current statement is valid, then correct all affected public copies, including downloads. Test the revised page with a counterquestion: under what condition would this sentence be false? If that condition matters to buyers, place it beside the assertion. Broader monitoring of external answers becomes more useful after this reconciliation.

Worksheet

Working field	Entry or observation	Next decision
Focus	Connect capability to plan and limitation	Replace with your own data
Evidence	URL, date, and reproducible observation	Record confidence and open question
Acceptance	Owner and observable target state	Schedule the team review

Primary reference for technical context: [Google: Structured data policies](#).

5. Choose buyer questions by evidence value

Collect questions from approved sales conversations, support patterns, product documentation, and interviews. Record provenance without publishing confidential conversation details. Compare decision relevance, product fit, available evidence, and maintenance feasibility. Internal priority must not contain invented demand figures. A frequently asked conceptual question may have weak product fit; a less frequent integration question may resolve a critical purchasing barrier. Reserve space for both but avoid judging them by the same success measure. Exclude questions for which the team could offer only unsupported assertions, and write down what evidence would be needed to reconsider that exclusion.

Bound the question portfolio

For the first cycle, select a small mix of four decision types. Connect every question to a real product relationship and a defensible destination page. Add an exclusion list: legal advice without expert review, unsupported competitor comparisons, and future capabilities stay outside the public answer program. A question owner records why a topic is included and which other topic must wait. When new sales questions emerge, first check whether an existing page can be extended. This creates a bounded portfolio of useful answers rather than an endless prompt catalog. Priority is internally justified rather than presented as market demand volume.

Worksheet

Working field	Entry or observation	Next decision
Concept question	Explain terms and limits	Comprehension before pipeline as first signal
Comparison question	Criteria and supported differences	No invented competitor assessment
Implementation question	Prerequisites and checks	Product or engineering reviews
Risk question	Boundaries and escalation	Expert approval before publication

Primary reference for technical context: [Google: Helpful, reliable content](#).

6. Produce useful source pages

A defensible source page includes a direct explanation, conditions, reproducible steps, a bounded example, and linked evidence. For technical procedures, specify version, prerequisites, and failure messages where publication is appropriate. Show the result of an action and explain how the reader can verify it. Never publish artificial studies or invented customer outcomes as original research. A clearly explained internal test with a method and limit is more useful than an apparently representative percentage without a defensible sample. Preserve the technical context even when individual sections are designed to remain understandable when read in isolation.

Bound a reusable source claim

For an original product test, describe environment, data type, steps performed, and the result actually observed. A successful test with sample data demonstrates the tested workflow rather than performance for every customer. Link current documentation and identify differences from older versions. Replace promotional superlatives with inspectable statements, such as which export fields are present. A data table needs units and a period; a screenshot needs context and must not expose confidential customer information. Finally, ask whether the assertion could be misunderstood outside its paragraph. Put necessary conditions beside the evidence instead of relying on distant qualification text.

Worksheet

Working field	Entry or observation	Next decision
Focus	Answer module: condition, action, verification, limit	Replace with your own data
Evidence	URL, date, and reproducible observation	Record confidence and open question
Acceptance	Owner and observable target state	Schedule the team review

Primary reference for technical context: [Google: Helpful, reliable content](#).

7. Observations rather than prompt theater

Use a stable set of conceptual, comparison, implementation, and risk questions. Separate branded and unbranded questions because their starting conditions differ. Record the visible source set rather than only your own domain. Check whether a link reaches the stated page and whether that page supports the referenced assertion. Code brand mention, linked source, and factual accuracy separately. A good user outcome may be an accurate answer without your mention; a poor outcome may prominently cite your brand while repeating an obsolete product claim. Include factual correction work alongside visibility monitoring so the program rewards accuracy as well as presence.

Code source quality alongside presence

For every owned source reference, an analyst checks three things: does the link resolve, does the destination support the associated assertion, and is the information current? A homepage can support provider identity while being unsuitable evidence for an integration restriction. Label those cases differently. Separate inaccurate external assertions from obsolete owned information because only the latter can be changed directly. If an answer has no visible sources, evidence attribution remains unknown. Do not infer a particular retrieval path. The monthly comparison therefore shows presence and source quality side by side and states correction work at the individual claim level.

Worksheet

Working field	Entry or observation	Next decision
Focus	Coding: mention / link / accurate assertion	Replace with your own data
Evidence	URL, date, and reproducible observation	Record confidence and open question
Acceptance	Owner and observable target state	Schedule the team review

Primary reference for technical context: [Microsoft: Introducing AI Performance](#).

8. Connect visits and pipeline carefully

An appearance is not a session. A session with a recognizable AI referring domain is not a complete measure of AI influence. Examine referring sources, landing pages, qualified events, and CRM progression with appropriate definitions for each. An optional discovery question can ask how a prospect learned about the company; retain free text according to your rules. Compare the answer with measurable sources without forcing agreement. Report assisted pipeline as a modeling perspective and qualified visits as observations, not as proven incremental revenue caused by a citation. Include the length of the sales cycle when deciding whether a review is mature enough.

Reconcile referrals deliberately

Build a report containing landing page, session source, and a predefined qualified event. First verify a controlled test visit so filters and events can be inspected. Then connect only permitted and actually available account information to the CRM. Treat a missing referrer as unknown; do not retrospectively assign it to an AI surface because a mention was observed elsewhere. Include the number of clearly matched inquiries and the review window. With very small counts, describe aggregated observations rather than selling volatile conversion rates as reliable channel performance. Keep self-reported discovery in a separate column so remembered influence is not confused with a measured referral.

Worksheet

Working field	Entry or observation	Next decision
Focus	Separate fields for source, self-report, and opportunity	Replace with your own data
Evidence	URL, date, and reproducible observation	Record confidence and open question
Acceptance	Owner and observable target state	Schedule the team review

Primary reference for technical context: [Google Analytics: Traffic acquisition](#).

9. Interpret failure patterns correctly

When bot requests exist but no sampled citation appears, several explanations remain open: weak question fit, alternative selected sources, a small sample, or missing index inclusion. None follows necessarily from the log. For an inaccurate product description, first check your own site for contradictory current and obsolete statements. If mentions rise without visits, examine answer usefulness and the next action instead of automatically shortening content. If visits fall while qualified inquiries remain stable, the visitor mix may have changed. The key is the earliest diagnosis supported by evidence, with an explicit note on what remains unresolved.

Distinguish three common findings

Finding one: the public page returns a security challenge. Engineering investigates actual delivery and bot verification. Finding two: the page is reachable but gives conflicting product restrictions. Product marketing owns factual reconciliation. Finding three: the page is reachable and accurate but absent from the small core question set. The cause remains unresolved; automatically declaring a technical defect would be unjustified. Each finding therefore has a different next action. For open selection questions, document the possibility that other sources better serve the particular need. Decide on further production or technical changes only after additional relevant observations justify that investment.

Worksheet

Working field	Entry or observation	Next decision
Focus	No citation despite request: cause remains open	Replace with your own data
Evidence	URL, date, and reproducible observation	Record confidence and open question
Acceptance	Owner and observable target state	Schedule the team review

Primary reference for technical context: [OpenAI: Overview of crawlers](#).

10. FICTIONAL SaaS example: ApprovalCloud

FICTIONAL: ApprovalCloud sells supplier approval workflows to midsize procurement teams. Assumptions include twelve public core pages, two languages, and a small shared marketing and product review cycle. The team discovers an obsolete page promising a CSV import that has been removed. In twenty assumed eligible answer observations, six mention the domain; two of those contain the inaccurate assertion. The first decision is correction and product-truth reconciliation, not maximizing the mention rate. After publication the team verifies access and accurate statements, repeats the same questions, and reports a sample finding. There is no actual customer outcome and no inferred revenue.

Recalculate the case decision

In the fictional example, mention rate is six of twenty. The immediate quality question is why two mentioned answers contain an obsolete capability. The team lists owned pages carrying the assertion and corrects demonstrably inaccurate information. It can then report that those owned copies were corrected. It cannot yet claim that external answers are fully updated. New runs with the same questions are needed for that observation. If the error persists, retain the finding with its date. This distinction prevents a completed editorial task from being confused with a controlled external outcome and keeps the team's acceptance statement within what it can actually verify.

Worksheet

Working field	Entry or observation	Next decision
Assumption	20 eligible answer observations	Do not claim a market sample
Presence	6 runs mention the domain	Report 6 of 20 in the core set
Quality	2 mentions contain an incorrect CSV claim	Correct before increasing presence
Acceptance	Owned page accurate and reachable	Observe later answers separately

Primary reference for technical context: [Google: Helpful, reliable content](#).

11. A 90-day operating rhythm

In the first two weeks, marketing and product define questions, the fact sheet, and controls. Engineering verifies delivery for a small important page family. By day thirty, the team produces its first complete answer package with baseline observations. Month two covers verification, corrections, and a careful extension to the second language. Month three evaluates maintenance needs alongside business signals. These windows are planning assumptions, not deadlines for search impact. When foundations are missing, keep the scope small. A well-maintained package takes priority over a large volume of unsupported pages labeled AI optimized, and the owner can defer expansion without declaring failure.

Actually close the first cycle

A cycle ends with four inspectable deliverables: an approved core question set, reconciled product facts, a technically checked page package, and a traceable observation export. Each has an owner and acceptance date. Marketing coordinates, but engineering is not responsible for an unresolved product assertion and product does not promise a traffic gain. At closeout, decide which uncertainty the next round should resolve. If the team merely collected more questions without improving or checking a page, the cycle remains incomplete. Reduce scope and finish the core work first rather than allowing monitoring to become a substitute for useful delivery.

Worksheet

Working field	Entry or observation	Next decision
Focus	Weeks 1-2: baseline; month 2: verify; month 3: decide	Replace with your own data
Evidence	URL, date, and reproducible observation	Record confidence and open question
Acceptance	Owner and observable target state	Schedule the team review

Primary reference for technical context: [Google: AI features and your website](#).

12. Continue, correct, or stop

Continue when the content solves a supported buyer question, assertions remain maintained, and verification is affordable. Correct when a specific defect has a clear test. Stop expansion when the program produces cosmetic prompt wins without a stable method, lacks product evidence, or neglects essential maintenance. A small negative sample alone does not justify deleting useful documentation. A positive sample alone does not justify a large budget. Place the budget decision alongside the quality decision so search uncertainty is not confused with editorial uselessness. Make the next investment conditional on evidence that the team can actually obtain.

Make expansion conditional

Before adding a second product line or language, check three prerequisites: the first group is accurately maintained, a colleague can repeat the method, and the extra workload has an owner. A favorable individual prompt satisfies none of these conditions. A bounded negative finding likewise does not automatically prevent expansion when the pages solve a documented buyer barrier. Set a cost envelope and a next decision for the extension. Defer the new group if product evidence is unavailable. This keeps the program manageable even when external citations vary considerably and a quick commercial connection cannot be established from the available data.

Worksheet

Working field	Entry or observation	Next decision
Focus	Continuation requires usefulness, maintenance, and verifiable evidence	Replace with your own data
Evidence	URL, date, and reproducible observation	Record confidence and open question
Acceptance	Owner and observable target state	Schedule the team review

Primary reference for technical context: [Google Analytics: Attribution](#).

13. Operating worksheets

Use four connected worksheets: question register, product facts, technical permissions, and observation log. A shared identifier connects a question to its destination page, supported assertion, owner, and review date. In the fact sheet, flag obsolete statements that must also be corrected in downloads or translations. The observation log includes valid negatives and unavailable tests. A monthly report should show decisions: which buyer barrier was resolved, which defect remains open, and what next evidence would permit expansion. Screenshots are supporting artifacts, not a substitute for this chain of decisions and responsibilities.

Maintain a complete example row

Q-01 asks about export permissions. Its destination page is P-04; product evidence E-09 specifies administrator permissions and completed records. Product marketing approves the claim, engineering checks reachability, and marketing records the question run. An unresolved defect retains the same identifier even if wording is later improved. Add a status for questions that are no longer relevant instead of deleting them from historical reports. A new colleague can then understand why the page exists and which evidence supports its central assertion. The identifiers are arbitrary; what matters is preserving the relationship between question, claim, and verification.

Worksheet

Working field	Entry or observation	Next decision
Question register	Role, question, provenance, destination	Marketing preserves Q identifier
Fact sheet	Claim, evidence, plan, limit	Product confirms changes
Access matrix	Path, purpose, rule, request	Engineering checks production
Observation log	Run, mention, link, accuracy	Analyst preserves negative findings

Primary reference for technical context: [Microsoft: Introducing AI Performance](#).

14. Sources and limits of the playbook

Google describes its own AI search features; this does not establish identical behavior in other systems. Microsoft reports selected owned and partner experiences rather than the entire market. OpenAI describes crawler purposes and controls rather than a complete recipe for citations. These limits belong inside the operating model. Recommendations on clear assertions, small batches, and evidence registers are this guide's original working methods. Their purpose is to make decisions inspectable. Neither consistent use of the method nor correct schema markup guarantees indexing, ranking, retrieval, citation, visits, or pipeline. Recheck platform-specific claims as the relevant documentation changes.

Reassess documentation when adding a platform

When adding another answer surface, do not transfer the previous platform's controls without review. Create a short source map covering access controls, documented product coverage, and available reporting. Pay particular attention to whether a metric uses the same event definition. A Microsoft report and a manual ChatGPT run can appear alongside each other but are not additive parts of a market total. Sources change, so retain the date of technical review and the decision derived from it as well as the link. Earlier conclusions remain historically traceable even when a platform later expands its documentation or reporting scope.

Worksheet

Working field	Entry or observation	Next decision
Focus	Keep platform documentation and original method separate	Replace with your own data
Evidence	URL, date, and reproducible observation	Record confidence and open question
Acceptance	Owner and observable target state	Schedule the team review

Primary reference for technical context: [Google: AI features and your website](#).

15. A complete weekly AI visibility assignment

Monday: Verify product truth before selection

Take a specific question from the existing core: which roles can initiate an approval export? The marketing owner opens the current destination page and the product owner confirms the current permission rule. Explicitly record known plan differences. Engineering tests a direct request without login. The expected result is a public explanation of prerequisites while the actual customer operation stays protected. If the explanation exists only after login, deliberately decide whether to add a public account of the capability. Never copy customer information into the marketing page. End Monday with a supported assertion and a clearly named unresolved point.

Wednesday: Publish a small correction

Editorial adds a direct answer, the required role, and an exclusion for unsupported records. Product review approves these sentences individually. The release note names the changed URL and both languages. Then inspect sources, language switching, and visible conditions in production. Older downloads carrying the same assertion receive a separate update task. Record publication as an owned content correction. Neither a bot visit that day nor a new mention is treated as proof of its impact. If rollback is needed, name the owner and accurate replacement information. Returning to a demonstrably incorrect claim is not an appropriate rollback simply because an older version is readily available.

Friday: Observe and decide the next action

Repeat the planned questions under documented conditions. For every eligible run, capture mention, link destination, and accurate representation of the role restriction. Include runs without an owned source. Weekly closeout contains three separate statements: what was corrected on owned pages, what was observed in external answers, and which commercial data actually exists. Small samples support small claims. The next week may address another buyer barrier or first verify existing defects. Expansion needs a reason beyond an attractive screenshot. This creates a repeatable assignment that remains useful when external selection varies and commercial attribution is incomplete.

Verified sources

The following official URLs were editorially checked on October 24, 2026. Recheck the current version before implementation changes.

- [Google: AI features and your website](#)
- [OpenAI: Overview of crawlers](#)
- [Microsoft: Introducing AI Performance](#)
- [Google: robots.txt](#)

- Google: Block indexing with noindex
- Google: Structured data policies
- Google: Helpful, reliable content
- Google Analytics: Traffic acquisition
- Google Analytics: Attribution

YOUR TRANSFER CANVAS

Turn this resource into a team brief

Use this page for a joint review. Agree one practical output, observable evidence and an accountable owner. Add the specific methods and templates from the relevant chapters.

Working field	Your agreement
Priority output	An evidence register for each buyer question
Your starting point	Finding / source / date: _____
Smallest useful test	_____ _____
Evidence & counter-evidence	_____ _____
Ownership & review	Owner / reviewer / date: _____
Stop or change when ...	_____ _____
Decision & next action	_____ _____

Keep unresolved assumptions explicitly open. A completed worksheet is not a substitute for evidence or accountable judgement.

Sources & further documentation

Primary sources checked 24 October 2026. Sources support the linked technical or factual statements; the working models and fictional cases are original editorial decision aids. Recheck changing platform details before implementation.

1. [Google: AI features and your website](#)
2. [OpenAI: Overview of crawlers](#)
3. [Microsoft: Introducing AI Performance](#)
4. [Google: robots.txt](#)
5. [Google: Block indexing with noindex](#)
6. [Google: Structured data policies](#)
7. [Google: Helpful, reliable content](#)
8. [Google Analytics: Traffic acquisition](#)
9. [Google Analytics: Attribution](#)

Next step

If you want help adapting the working plan to your product, evidence and team capacity, request a non-binding strategy conversation.

patrickwings.com