

SEO and AI search strategy for a B2B SaaS platform

Connecting discovery, buyer evaluation, external evidence, technical priorities, and measurement.

DISCOVER

EVALUATE

CHOOSE

Prepared by Rich Migliorisi

SEO, AI Search & Organic Growth Consultant

What this example is - and is not

A strategy sample

This document demonstrates how I connect SEO, AI search, buyer research, outside evidence, technical priorities, and measurement for a B2B software business.

Important boundaries

- The company and competitors are intentionally anonymized.
- This was an independent exercise, not a client engagement.
- The recommendations were not implemented or measured.
- No commercial outcome is claimed.

Search should support a buyer decision

Rankings and mentions matter, but the commercial job is to help the right buyer discover, evaluate, and choose the business.

1

Discover

Be present when a buyer first recognizes the problem and explores the category.

2

Evaluate

Provide useful comparisons, proof, and answers for serious consideration.

3

Choose

Reduce uncertainty with credible evidence, clear product information, and usable conversion paths.

QUESTIONS

The strategy begins with decisions, not deliverables

The audit is organized around what the company needs to understand and decide next.

- 01 Where is the business visible today, and where does buyer demand go unanswered?
- 02 Does existing-user support demand obscure new-buyer discovery?
- 03 How is the company represented in AI answers and competitor shortlists?
- 04 Which facts, sources, and outside signals support or weaken trust?
- 05 Which content, technical, and authority priorities deserve the next 90 days?
- 06 How should progress be measured without overstating causality?

Use multiple evidence layers

No single platform, score, or snapshot explains the whole market.

Search

Demand, visibility, landing pages, intent, and technical access.

AI answers

24 buyer scenarios across four systems, producing 96 human-reviewed answers.

Evidence

44 owned and independent records reviewed for facts, authority, conflict, and corroboration.

Buyer journey

Questions and proof needs mapped across operational, technical, people, and commercial stakeholders.

Separate support demand from buyer demand

A business can rank for many queries while remaining hard for a new buyer to discover.

Existing-user demand

- Login, setup, troubleshooting, and support questions
- Useful for customers but not a complete acquisition strategy
- Can inflate visibility without expanding the buyer audience

New-buyer demand

- Problem, category, comparison, integration, and proof questions
- Pages designed around buyer situations and stakeholder concerns
- Clear paths from useful information to serious consideration

A keyword count is not an audience.

One competitor ranked for 2,439 U.S. searches. About 74% of the traffic those rankings produced came from roughly five of them.

NEW-BUYER DEMAND

Deciding whether the product fits

- Problem, category, comparison, integration, and proof questions
- Searched by the people who sign off, not the people who log in
- Converts into sales conversations

EXISTING-USER DEMAND

Trying to use the product today


- Login, setup, download, and help questions
- Genuinely useful, and a strong product signal
- Inflates the keyword count without widening the buyer audience

Same brand, two different audiences — and only one of them buys. Measure them separately from day one.

Third-party organic research, U.S., one same-day snapshot, August 2026. The 2,439-keyword figure and the 74% concentration are a third-party estimate about a competitor, not a claim about the company in this example — it is included because it is the clearest illustration of the trap. Tool traffic figures are modelled estimates, not measured visits.

No one leads on every measure.

Two different questions, deliberately not averaged into a single league table.




Second of four on search presence, last of five on AI visibility. Those are not the same question.

Third-party organic research and AI-presence scores, U.S., same-day pulls, August 2026. The two panels cover different brand sets because comparable same-day data was not available for all five on both measures. The AI visibility score is a vendor index whose formula is not published; it is neither market share nor buyer preference. Names removed.

From mention to recommendation

A mention is not the same as being seriously considered or recommended.



Review answer accuracy, competitor inclusion, cited sources, and the reasons a company advances or disappears at each stage.

It gets into the conversation. It thins out at the shortlist.

Twenty-four buyer scenarios put through four AI assistants, and every answer read by a human.

Mentioned

named in the answer



Considered

seriously weighed up




Shortlisted

makes the final few



Recommended

picked as the answer



96
answers
in total

The largest fall-off is not visibility — it is the step just before the shortlist.

Controlled study, 24 designed buyer scenarios across 4 AI providers, 96 answers, all human-reviewed, August 2026. Each stage is a subset of the one above it and all four are measured against the same 96 answers, so the counts do not add up to a total. Designed scenarios, not a measure of AI market share, traffic, or revenue. The study shows where a brand thins out, not why.

Cited on the explaining, almost never on the choosing.


The AI answer above the search results cites this company often — but hardly ever where it counts.

329

searches where the AI answer
cites a page on this company's site

Out of 5,031 unique ranking keywords —
about 6.5% of everything the company
ranks for, and roughly 76,700 monthly
searches combined.

WHICH PAGES GET CITED



The AI answer treats this company as a source of explanation, not as an option to choose.

Organic positions export, U.S., one same-day pull, August 2026. 329 of 5,031 unique ranking keywords return an AI answer citing a page on the site; the page-type split counts one citation per keyword and sums to all 329. The tool infers AI-answer presence from its own crawl, so treat it as directional on a surface that changes daily. A citation is not a visit, and competitors were not benchmarked on this measure.

Examine the information behind the answer

Evidence is one important layer of modern search, not a deterministic formula for controlling AI systems.

Facts

Are important product, company, security, integration, and customer claims clear?

Sources

Which owned and independent sources support those claims?

Corroboration

Do multiple credible sources agree, or does one unsupported statement carry the story?

Conflict

Where do outdated, ambiguous, or wrong-entity records create uncertainty?

Competitors

What credible proof exists for alternatives that the company lacks?

Authority

Which publishers, experts, directories, and communities shape category understanding?

Fifteen industry publishers. None of them link here.

Filtered from ten thousand referring domains down to the publishers a buyer in this category would actually read.

83

domains link to all four competitors
and not to this company

15

of those are genuine industry
publishers rather than directories

5

link to every one of the four
competitors, and to no one else

Why these publishers, and not the biggest ones

Ranked purely by authority score, the top of the list is tenant subdomains, contact-data aggregators, wire services, and user-generated platforms — high scores, no readers who buy. Filtering for genuine editorial relevance left a short named list. One trade title on it reaches 251,000 subscribers in exactly the function that signs off on this purchase. The filter is the deliverable; anyone can export the raw list.

Authority here is not the biggest score. It is the short list of publishers this buyer actually reads.

Third-party backlink gap analysis, August 2026, comparing the company with four named competitors across 10,000 referring domains. "Link to" is what this data shows — it is not the same as covering or mentioning the company, and a publisher may have written about a brand without linking to it. Authority scores are vendor estimates. Publisher names removed.

Different stakeholders require different proof

One generic product page rarely answers every concern in a complex B2B purchase.

Operations

Adoption, frontline reach, workflow fit,
measurable operational value

IT and security

Integration, controls, governance,
deployment, reliability

People and communications

Engagement, accessibility, clarity,
content governance

Finance and procurement

Business case, risk, implementation
effort, commercial proof

The page can win attention and still lose the decision.

Three questions to ask of any page that already earns traffic, before building anything new.

FIT

“Does this fit an organisation like ours?”

Show sector and workforce relevance, and how it fits the existing stack.

RISK

“Can we roll this out confidently?”

Bring implementation, security, and integration detail within reach.

PROOF

“Has it worked for someone like us?”

Put outcomes near the decision, and make them easy for a buyer and an AI answer to find.

Strengthen the page that already works before building new ones. Say what is good about it first, and mean it.

Applied to the page already earning the most qualified attention. The point of the exercise is that a page can be well built and still leave a buying committee without an answer — the gap is usually risk and proof, not fit.

Turn the findings into a focused plan

The point of the analysis is to determine what should change, who needs to be involved, and what gets measured.

- 1 Build buyer pages around real problems, category questions, comparisons, integrations, and proof.
- 2 Strengthen company and product facts across owned pages and credible external sources.
- 3 Close technical and internal-linking gaps that limit discovery or create poor journeys.
- 4 Create a repeatable authority and digital PR program around useful, defensible information.
- 5 Align content, product, engineering, analytics, and communications around a shared 90-day roadmap.

It already ranks for the buying language. Just too deep.

Thirty-one buyer-relevant phrases where the company sits beyond position thirty and a competitor sits inside the top twenty.

EXAMPLES FROM THE THIRTY-ONE

internal communication tools	880	monthly searches
employee experience platforms	480	monthly searches
internal communications platforms	320	monthly searches
employee engagement mobile app	210	monthly searches
frontline employee engagement	210	monthly searches

This is a keyword-level gap, not a topic-level one. Topic exports bundle unrelated terms together and inflate the opportunity.

FORMATS AND CADENCE

Always on

the fit, risk, and proof pages kept current

Monthly

one page that answers a real buyer question

Quarterly

one research push worth press coverage

Volume is not the target. Every asset closes a named gap and has a measurable next step.

Third-party keyword gap analysis, U.S., August 2026, across a 14,643-keyword comparison set. Thirty-one buyer-relevant phrases have the company beyond position 30 with at least one competitor inside the top 20. Volumes are estimated monthly U.S. searches, not measured demand, and positions move daily. Difficulty scores were deliberately not used as evidence of how easy these are to win.

Sequence the work so learning compounds

A roadmap should combine quick corrections with the foundations needed for sustainable progress.

DAYS 1-30

Establish the baseline

- Confirm priority buyers and questions
- Fix critical technical and entity issues
- Define measurement and ownership

DAYS 31-60

Build the decision layer

- Create or improve priority buyer pages
- Strengthen proof and internal linking
- Begin outside-source and authority work

DAYS 61-90

Test, publish, and learn

- Release the highest-value work
- Repeat search and AI-answer sampling
- Review movement and set the next roadmap

Measure the funnel without claiming hidden causality

The reporting model distinguishes visibility, engagement, progression, and business outcomes.

Discoverability

Qualified rankings, non-brand entry pages, relevant AI inclusion

Engagement

Landing-page use, content interaction, CTA movement

Progression

Demo or inquiry paths, assisted journeys, stakeholder content use

Business

Qualified pipeline and revenue context where reliable attribution exists

Guardrails: visibility is not demand; a citation is not a visit; engagement is not progression; influence is not attribution.

What a buyer would receive

The exact format depends on access, the business question, and how much implementation support is needed.

Findings

A concise explanation of what matters, what does not, and why.

Priorities

A sequenced roadmap tied to commercial relevance, effort, risk, and dependencies.

Build guidance

Page briefs, technical requirements, tickets, acceptance criteria, or direct implementation support.

Measurement

A practical baseline, repeatable checks, and explicit reporting limitations.

Real visibility, but most of it sits where nobody looks.

A keyword total says nothing about findability. The distribution is the finding.

BEST POSITION PER KEYWORD



Only 16% of the rankings sit on page one. Around two thirds sit past page two.

THE VISIBILITY WORTH PROTECTING

7,206

U.S. keywords

8,356

estimated visits a month

WHAT THIS CHANGES

The strategy stops being “rank for more things” and becomes “get depth on the things that already matter”. That is a cheaper brief, and a more defensible one.

There is discovery worth protecting here. Growth comes from depth, not from more keywords.

Third-party positions export, U.S., August 2026. Best organic position per keyword across 4,904 keywords carrying an organic position; the five bands sum to that total. Individual results in the export were read over several weeks, so this is a rolling snapshot rather than one day. Third-party estimates, not first-party traffic, leads, or revenue.

Fix what blocks a buyer, not what a crawler counts.

An error count is not a user count, and a long list is not a plan.

THE CYCLE

Find → verify → fix → re-check

In this case, 453 broken internal links, down from 1,136 since the previous crawl. One important route had been returning a 404 and was live again by the time of the review — which is exactly why a finding gets re-verified before it reaches a slide.

RE-VALIDATED EVERY CYCLE

- Broken links and redirects
- Crawl and index status
- Internal linking on decision pages
- Page experience on mobile
- Structured data on key templates

THE PRIORITY RULE, IN ORDER

- 1 An important journey is affected
- 2 The finding reproduces
- 3 The scale is material
- 4 Owner and dependency are clear
- 5 First-party data can validate impact

Crawl occurrences are not affected users, lost conversions, or proof of a current defect.

Site audit, crawl, and real-user page-experience data, August 2026. Counts move between crawls, so any figure of this kind is only true on the day it was pulled — which is the reason for the re-check step rather than a caveat about it.

24 scenarios, 4 providers, 96 answers, all human-read.

The design is what makes a re-measurement in three months comparable rather than merely similar.

HOW THE STUDY WAS BUILT

Panel	24 buyer scenarios, frozen before the run, across four providers
Universe	96 accepted observations, U.S. English
Lineage	One accepted run per scenario-provider pair; attempts immutable
Review	Every accepted answer read by a human before scoring
Separation	API evidence kept apart from consumer app behaviour

WHAT EACH STATE MEANS

Mentioned	surfaced or named at all
Considered	materially weighed for the decision
Shortlisted	enters a bounded, suitable set
Recommended	explicitly endorsed or selected
Preference	direct or forced choice only; ties kept separate
Observed: 36 → 33 → 22 → 21	

One accepted run per pair, a sampled study, and no population or causal inference.

Controlled study, August 2026. It measures how four assistants answered 24 designed scenarios on one date under held-constant controls. It is not AI market share, not traffic, and not revenue. Scenarios were frozen before the run precisely so that nothing could be rewritten after seeing the results.

Four confidence levels, and they are never blended.

The label decides the verb. A hypothesis does not get written up as a fact because it would read better.

Measured We captured it directly Export counts, audit figures, the 96 reviewed AI answers	Repeated pattern Several sources point the same way Signals that recur across independent records	Strong inference It survives contradiction checks Where different audiences describe different realities	Hypothesis Needs first-party validation Whether entity confusion affects AI shortlisting
---	---	--	--

THE DISCIPLINE IN PRACTICE · PUBLIC EVIDENCE REVIEW			
44 records reviewed	36 / 8 independent vs owned, kept apart	10 wrong-entity records quarantined	0 composite scores calculated

Public evidence review, August 2026. A finding changes status when the evidence changes, in either direction. Eight records overlapping an earlier study were excluded to avoid double-counting, and a cluster of five same-week reviews was retained but down-weighted. No composite score was calculated, because averaging across different kinds of evidence hides the thing worth knowing.

The guardrails behind every number reported.

Agreed before the first report, so nothing gets redefined once the numbers are inconvenient.

Visibility is not demand

Ranking for something is not the same as anyone wanting it, and it is not intent to buy.

A citation is not a visit

Being named in an AI answer means the company was considered relevant. It does not mean anyone arrived.

Engagement is not progression

Reading the proof page is a good sign. It is not a commitment, and it will not be reported as one.

Influence is not attribution

Where organic contributed to a deal can be shown. A straight line from an AI mention to revenue cannot.

Signals are reported separately and never combined into a single score.

Reporting cadence: weekly for visibility, monthly for confidence and action, quarterly for business impact. Every measure is defined before the first report. The purpose of writing these down in advance is that they are much harder to argue with later, including by the person who wrote them.

WHAT THIS DEMONSTRATES

Senior search strategy that connects evidence to action

- SEO and AI search treated as one evolving discovery system
- Buyer questions and commercial priorities placed ahead of vanity metrics
- Evidence, technical work, content, and authority connected in one roadmap
- Clear limitations instead of guaranteed rankings, citations, or recommendations
- Strategy designed to move into briefs, requirements, releases, and measurement

Discuss a project

richmigliorisi.com/contact/