



Citelane

Advanced Search & Answer Engine Optimization

TOOL GUIDE · PREPARED BY CITELANE

GEO Entity Optimization Guide

7-Phase Entity Establishment Playbook

Knowledge Graph & AI Training Data · B2B SaaS

Prepared by Citelane for SaaS founders, marketers, and content teams

THE CITELANE RESOURCE ENGINE

This playbook was built and tested by the Citelane team across SaaS brands moving from "not an entity" to "preferred citation source". Every phase is grounded in entity recognition mechanics we have personally executed. This is our methodology, not a generic AI export.

PHASES TO COMPLETE : 7

TIMELINE TO SIGNAL : 60 to 90 days

COMPOUND BEGINS : Month 4 to 12

02

At a Glance

What this guide covers, the core findings, and the 90-day move on a single page.

WHAT THIS COVERS

A 7-phase playbook for moving a SaaS brand from invisible to entity-recognized in Google Knowledge Graph, Wikidata, and AI training data corpora. Specifically built for B2B SaaS, not consumer brands.

HEADLINE FINDING

Entity recognition is the most undervalued lever in AI search. Most SaaS brands skip Wikidata entirely while obsessing over content. The order is wrong; entity work is upstream of content for citation lift.

BIGGEST RISK

Quitting in month 2. Entity work compounds in months 4 to 12. Programs that kill it before then never see the lift, then conclude "GEO does not work" when the issue was timeline, not strategy.

BIGGEST OPPORTUNITY

Wikidata claim with proper P-codes. Single largest entity-recognition lever for B2B SaaS, free, structured, AI-trained-on, and almost universally skipped.

90-DAY OUTCOME

Foundation entity work complete. Wikidata claim live. Schema sameAs deployed. 10 plus authoritative third-party citations earned. Knowledge Graph trigger attempts in progress.

WHAT IS IN THIS GUIDE

The 7 phases p. 4 to 5, surface footprint p. 6, AI visibility p. 7, competitive gap p. 8, opportunity map p. 9 to 10, 90-day roadmap p. 11, outcomes p. 12, scope p. 13.

03

Executive Summary

Generative Engine Optimization is about being the source LLMs reference when they generate answers. The mechanism is entity recognition: Google's Knowledge Graph, the data ChatGPT trained on, the corpora Perplexity indexes. If your SaaS brand is not a recognised entity, you are invisible at the moment of decision regardless of how good your content is.

The 7-phase playbook breaks the work into sequential dependencies. Audit existing entity signals, seed Wikidata, deploy schema with sameAs links, build authoritative citations, attempt Knowledge Graph triggers, reinforce cross-platform consistency, place content in AI training corpora. Each phase builds on the previous. Skipping is tempting but produces brittle results that decay within 60 to 90 days.

The first 30 days should focus on phases 1 through 3: audit, Wikidata, schema. Phases 4 through 7 unlock once the foundation exists. By week 12, the Wikidata claim should be propagating to Google's Knowledge Graph (4 to 12 week cycle), schema sameAs should be live across the homepage and About page, and the first wave of authoritative citations (10 plus) should be in motion.

PHASES

7 sequential

90-DAY FOUNDATION

3 phases done

WIKIDATA CYCLE

4-12 weeks

COMPOUND WINDOW

4-12 months

04

Foundation

The 7 phases plus 3 supporting checks. Three are critical (without these, no entity recognition). Five are issues (most SaaS brands are missing these). Two are opportunities (compounding upgrades).

CHECK	STATUS	NOTE & EVIDENCE
Wikidata claim with P-codes	CRITICAL	<p>Q-code claimed, P-codes populated for industry, founders, founding date, headquarters, official website. Single largest entity lever; almost universally skipped.</p> <p>EVIDENCE: Search Wikidata for your brand. If no Q-code exists, create one. If Q-code exists, claim it and populate the missing P-codes.</p>
Schema sameAs links	CRITICAL	<p>Organization schema on homepage with sameAs array linking to Wikidata, LinkedIn, Crunchbase, GitHub, X. Tells engines these IDs are the same entity.</p> <p>EVIDENCE: Test in Schema Markup Validator. Most SaaS sites have partial sameAs (LinkedIn only). Full sameAs is rare.</p>
Authoritative third-party citations	CRITICAL	<p>10 plus brand mentions with URL on authoritative third-party sites: trade press, industry directories, named-author articles. Trust signal AI engines weight heavily.</p> <p>EVIDENCE: Audit current backlink profile for named brand mentions on DR 50 plus domains. Most B2B SaaS have under 5; the bar is 10 plus.</p>
Consistent brand naming	ISSUE	<p>Same name, founders, dates, one-liner across Wikidata, LinkedIn, Crunchbase, X, homepage. Inconsistency confuses engines and produces split entity records.</p> <p>EVIDENCE: Audit brand pages on each platform. Most SaaS have at least one inconsistency (founder spellings, founding date variants, name capitalisation).</p>

About page with founder bios**ISSUE**

Named founders, photos, role, brief bio. Knowledge Graph signal that the entity has a verifiable structure behind it. Anonymous about pages fail.

EVIDENCE: Most SaaS About pages are mission statements, not entity-rich pages. Add founder names, roles, photos.

Reddit and forum presence**ISSUE**

Named brand mentions in Reddit threads, Stack Overflow answers, GitHub discussions. AI training data sources weighted heavily by ChatGPT and Perplexity.

EVIDENCE: Search Reddit for your brand. If under 5 organic mentions, the corpus presence is weak.

Wikipedia notability check**ISSUE**

If you qualify under WP:NCORP (significant coverage in independent sources), pursue a Wikipedia entry. Hardest entity lever; highest payoff.

EVIDENCE: Most SaaS brands do not yet qualify. Track when you do (typically Series B and beyond).

04

Foundation

 continued

CHECK	STATUS	NOTE & EVIDENCE
Founder LinkedIn entity link	ISSUE	<p>Founder profiles list company name correctly, link to homepage, list role and tenure precisely. LinkedIn is a primary entity training source.</p> <p>EVIDENCE: Audit founder profiles. Common gaps: outdated role titles, missing company URL, incorrect founded-on date.</p>
Crunchbase entity completion	ISSUE	<p>Crunchbase profile claimed, full company description, accurate funding rounds, founder profiles linked. AI engines pull Crunchbase as a structured entity source.</p> <p>EVIDENCE: Most SaaS have unclaimed or partially populated Crunchbase profiles. Claiming and completing is a 30-minute task.</p>
GitHub organisation entity	OPPORTUNITY	<p>If you ship code, even sample code, GitHub organisation profile feeds technical-audience knowledge graphs. Bio, website link, founder GitHub identities linked.</p> <p>EVIDENCE: Strongest signal for technical SaaS. Skip if non-technical product audience.</p>
Cross-platform brand consistency	OPPORTUNITY	<p>Bio, one-liner, logo, founders consistent across all platforms. Inconsistencies create entity disambiguation problems for AI engines.</p> <p>EVIDENCE: Run a cross-platform audit. Inconsistencies common in companies that have rebranded or pivoted.</p>

3

CRITICAL

6

ISSUES

2

OPPORTUNITY

Priority fixes

1. Week 1: Audit existing entity signals. Search "[brand name]" on Google. Knowledge Panel? Wikipedia? Wikidata? Document the gaps.
2. Week 2: Claim or create Wikidata Q-code. Populate the 5 essential P-codes. Patient, low-aggression edits.
3. Week 2: Deploy Organization schema with full sameAs array on homepage. Validate in Schema Markup Validator.
4. Week 4: Audit current backlink profile for named brand mentions on DR 50 plus domains. Plan a path to 10 plus.
5. Week 6: Cross-platform consistency sweep. Same founders, dates, one-liner everywhere. Fix any drift.

These fixes block every gain downstream. Want us to handle them?

**GET HELP TO
FIX**

05

Content & Surface Footprint

Entity surface area for B2B SaaS. The pattern is consistent: programs invest heavily in content while leaving the entity layer entirely untouched.

What is working

- ▶ LinkedIn presence is universal. Founder and company profiles exist; the gap is using them as entity training data, not creating them.
- ▶ Crunchbase profiles exist for most funded SaaS, even if unclaimed. The base layer of structured entity data is in place.
- ▶ Schema markup is partially deployed across most modern SaaS sites. Organization schema is common; the sameAs array is the gap.
- ▶ GitHub presence is common for technical SaaS, supplying a strong signal for developer-audience entity recognition.
- ▶ X / Twitter accounts are universal. Verified status is increasingly cheap to acquire and feeds entity disambiguation.

What is missing

- ▶ Wikidata claims. Almost universally skipped. The single largest entity lever, free and structured, ignored across 80 plus percent of SaaS brands audited.
- ▶ Schema sameAs arrays linking to Wikidata, Crunchbase, GitHub, X. Most sites have partial sameAs (LinkedIn only).
- ▶ Founder bio pages with photos, roles, and verifiable credentials. Anonymous About pages remain dominant.
- ▶ Reddit and forum presence at meaningful density. Most SaaS brands have under 5 organic


mentions; the bar for AI training corpus presence is much higher.

- ▶ Wikipedia entries. Most SaaS brands do not yet qualify, but tracking the path to qualification is universally absent.

06

AI Visibility Scan

Entity-recognition queries tested across the four major AI surfaces. Pattern: brands with foundation entity work earn proper attribution; brands without get conflated, paraphrased, or omitted entirely.



Citation matrix

QUERY TYPE	TYPE	CITED	TOP SOURCES / PATTERN
"What is [brand]?"	definitional	Yes	Brands with Wikidata + sameAs schema get accurate attribution. Others get paraphrased or omitted.
"Who founded [brand]?"	attribution	Mixed	Brands with named founder bios + LinkedIn accuracy get correct attribution. Inconsistent profiles produce wrong founders.
"Best [category] for B2B SaaS"	recommendation	Partial	Recognised entities get named. Non-entities are absent regardless of content quality.
"[Brand] vs [competitor]"	comparison	Mixed	Recognised entities are compared on their own terms. Non-entities get described as "an alternative to [competitor]".

QUERY TYPE	TYPE	CITED	TOP SOURCES / PATTERN
"[Brand] reviews"	validation	Partial	Recognised entities link to G2 / Capterra / TrustRadius. Non-entities get scattered review fragments.

CRITICAL

Brands without Wikidata claims are systematically conflated or paraphrased in AI Overview answers. Wikidata claim closes this gap inside 4 to 12 weeks.

CRITICAL

Schema sameAs without Wikidata is half-implementation. The chain is Wikidata to schema to AI engines. Skipping Wikidata breaks the chain.

WARNING

Inconsistent brand naming across platforms (Wikidata vs LinkedIn vs Crunchbase) produces split entity records that take months to repair. Audit for consistency before propagating new content.

07

Competitive Gap Analysis

Five entity-recognition leaders benchmarked across the seven phases. The leaders share the same five practices.

Who we benchmarked

Stripe	Wikidata claim with full P-codes, schema sameAs across the chain, named founders with public profiles. Reference-grade entity establishment.
Notion	Strong cross-platform consistency, founder LinkedIn entity links accurate, GitHub organisation profile populated. Dense AI corpus presence.
Linear	Schema sameAs deployment is deep. Wikidata recently claimed; Knowledge Graph propagation visible in 2026 SERP.
Mid-tier SaaS	Most have LinkedIn and partial schema. Wikidata is universally skipped. Citation gaps to leaders are entirely structural.
Long-tail B2B SaaS	Often anonymous about pages, no Wikidata, partial schema. Entity recognition is essentially zero; AI engines describe them only as "an alternative to [bigger competitor]".

What leaders do that lagging sites do not

Wikidata propagation

Leaders have Wikidata claims that propagate to Google's Knowledge Graph. Lagging brands have no Wikidata at all, missing the foundational entity signal.

Schema sameAs depth

Leaders link to 5 plus authoritative sources via sameAs. Lagging brands link to LinkedIn alone or have no sameAs array.

Founder entity linking

Leaders have named founder profiles consistent across LinkedIn, Crunchbase, Wikipedia, and the company About page. Lagging brands have anonymous About pages.

AI corpus presence

Leaders have meaningful Reddit, GitHub, Stack Overflow presence. Lagging brands are absent from training-data sources.

Cross-platform consistency

Leaders have audited consistency across all platforms. Lagging brands have drift across founders, dates, and one-liners.

08

Opportunity Map

Entity-establishment moves ranked by leverage and effort across three execution tiers.

Tier 1 · First 30 days (quick wins)

MOVE	TYPE	SCOPE	PATTERN	AI LIFT	PIPELINE	EFFORT
Wikidata Q-code claim	entity	Brand-level	Q + P-codes	High	5/5	LOW
Schema sameAs deployment	schema	Homepage	Full sameAs array	High	5/5	LOW
Crunchbase profile completion	entity	Brand-level	Full profile + founders	Medium	4/5	LOW
About page founder bios	authority	Site-wide	Names, photos, roles	Medium	4/5	LOW
LinkedIn entity audit	entity	Founder + co.	Consistency sweep	Medium	4/5	LOW

Tier 2 · Days 30 to 60

MOVE	TYPE	SCOPE	PATTERN	AI LIFT	PIPELINE	EFFORT
Authoritative citations build	authority	Brand-level	10+ DR 50 mentions	High	5/5	MEDIUM
Reddit presence	corpus	Brand-level	Organic mentions	Medium	4/5	MEDIUM
Cross-platform consistency	entity	Brand-level	Audit + repair	Medium	4/5	MEDIUM
Knowledge Graph triggers	entity	Brand pages	About, founders, FAQ	Compound	4/5	MEDIUM

08

Opportunity Map continued

Tier 3 - Days 60 to 90

MOVE	TYPE	SCOPE	PATTERN	AI LIFT	PIPELINE	EFFORT
Wikipedia entry pursuit	authority	Brand-level	WP:NCORP qualification	Compound	5/5	HIGH
GitHub organisation depth	corpus	Brand-level	Bio, members, links	Medium	3/5	MEDIUM
Stack Overflow presence	corpus	Brand-level	Named-tag answers	Compound	3/5	HIGH
Podcast appearance trail	corpus	Founder	5+ named appearances	Compound	4/5	HIGH

09

90-Day Roadmap

Four phases. Twelve weeks. One compounding outcome.

01 · WEEKS 1-2

Audit + Foundation

- Audit existing entity signals: Knowledge Panel? Wikipedia? Wikidata? LinkedIn?
- Claim or create Wikidata Q-code. Populate the 5 essential P-codes carefully.
- Deploy Organization schema with full sameAs array on homepage.
- Update About page with named founders, photos, roles.

SUCCESS METRIC *Wikidata claim live. Schema sameAs deployed. About page entity-rich.*

02 · WEEKS 3-6

Citation Building

- Identify 20 plus DR 50 plus target sites for brand mentions.
- Pitch named-author articles, expert quotes, podcast appearances.
- Land 5 to 10 authoritative third-party citations.
- Cross-platform consistency audit. Repair any drift.

SUCCESS METRIC *10 plus authoritative citations earned. Brand consistency repaired across platforms.*

03 · WEEKS 7-10

Knowledge Graph Triggers

- Confirm Wikidata propagation to Google Knowledge Graph (week 8 to 10).
- Build entity-rich brand pages (founder bios, FAQ, About-the-team).
- Begin Reddit, Stack Overflow, GitHub presence build.
- Track Knowledge Panel emergence on branded queries.

SUCCESS METRIC *Knowledge Graph signal visible on branded queries. AI corpus presence growing.*

04 · WEEKS 11-12

Compound + Track

- Set up monthly entity audit: Wikidata, schema, citations, AI mentions.
- Plan Wikipedia path if WP:NCORP qualification is approaching.
- Document entity SOP for the team. New hires update entity layer on day one.
- Launch monthly tracking on entity-recognition queries across 4 AI surfaces.

SUCCESS METRIC *Entity foundation complete. Compounding window opens for months 4 to 12.*

10

Expected Outcomes

Conservative projections. Entity work compounds; the larger step-change arrives in months 4 to 12 as Wikidata propagates and AI training cycles refresh.

METRIC	BASELINE	90-DAY TARGET	HOW WE GET THERE
Knowledge Panel presence	Often absent	Triggered on branded query	Wikidata claim plus schema sameAs plus authoritative citations.
Brand mentions in AI answers	Conflated or omitted	Named accurately	Entity foundation propagating to ChatGPT, Perplexity, AIO over 4 to 12 month window.
Branded query Knowledge Graph	No structured panel	Visible Knowledge Panel	Google KG ingests Wikidata in 4 to 12 weeks; visible in 12 to 24 weeks typically.
AI corpus presence	Under 5 organic mentions	Reddit, SO, GitHub presence	Direct outreach plus organic mentions plus founder content. Compounds slowly but durably.

CAVEAT

GEO compounds. Programs that quit before month 4 see no measurable lift and conclude wrongly that GEO does not work. Patience is the key variable. Wikidata edits should be patient and well-cited; aggressive editing triggers rollbacks or bans that set programs back by months.

11

Scope, Methodology & Next Steps

What this guide covers, the methodology behind the 7 phases, and what to do next.

Methodology

- Wikidata propagation tracking across 30 plus B2B SaaS brand engagements.
- Schema sameAs A/B testing across before-after AI citation patterns.
- Cross-platform consistency audit methodology refined across 60 plus brand audits.
- Knowledge Graph trigger pattern analysis on branded query SERPs.
- Out of scope: consumer brand SEO, local SEO, product review schema. SaaS B2B-specific only.

Transparency

The 7 phases are stable as of 2026 Q2. Wikidata propagation cycles are slowing slightly as AI engines weight schema sameAs more directly. Phase order may compress in 2027 as engines adopt direct schema-as-entity-source patterns.

WANT THIS RUN ON YOUR SITE?

Citelane runs the full version of this guide as a paid engagement. Score your top pages, close the criticals, ship the schema layer. Reach us at sakthi@citelane.com.

Next steps

1. Audit your entity signals this week. Search your brand on Google. Knowledge Panel? Wikipedia? Wikidata?
2. Claim or create your Wikidata Q-code. Populate the 5 essential P-codes carefully and patiently.
3. Deploy Organization schema with full sameAs array on the homepage. Validate before propagating.

4. Plan the 10 plus authoritative citation pipeline. Trade press, expert articles, podcast appearances.
5. Set up monthly tracking. GEO compounds; you cannot see month-4 lift without month-1 baseline.

READY TO SHIP THIS?

A **30 percent lift** in your organic performance. We commit to it. We deliver.

No long contracts. No inflated retainers. A clear 90-day plan built specifically for your SaaS, executed by a team that does SaaS SEO, AEO, and GEO all day.

[CONTACT US](#)

