

Generative-Engine Optimization (GEO)

The practical playbook for getting cited by
ChatGPT, Perplexity, and Google AI Mode.

A 24-page field guide for B2B and ecommerce brands competing for AI-search visibility.
Built from OST's own GEO retainers, structured-data audits, and prompt-coverage tests
across the four major AI surfaces.

VERSION: **1.0 * May 2026**

AUDIENCE: **Marketing leaders, founders**

LENGTH: **~24 pages * 4,800 words**

READING TIME: **~22 minutes**

LEAD AUTHOR: **Shaili Gupta**

Executive Summary

What this paper covers, in plain English.

AI search has split the funnel. Half of consumers now use ChatGPT, Perplexity, Claude, or Google's AI Mode for everyday questions (McKinsey, 2025). When those engines answer a question, they cite a small set of source pages. If your brand is in that set, you get qualified, intent-rich traffic. If it is not, your competitor does. Generative-Engine Optimization (GEO) is the discipline of becoming one of the cited sources.

This paper is a working playbook, not a literature review. It documents the technical foundations that make a page citable, the content patterns that AI engines reach for, the brand-mention monitoring tools that prove the work is paying back, and a 30-day audit checklist any team can run today. Every recommendation is drawn from real OST retainers, not from theory.

The takeaway is simple: GEO is mostly accessibility, structured data, and clear writing, with three new layers on top. If you are already serious about SEO and accessibility, you are 70% of the way there. The other 30% is the work this paper covers.

4 surfaces

ChatGPT, Perplexity, Claude,
Google AI Mode

30 days

From baseline audit to first
measurable lift

2x to 4x

Typical citation-rate gain at day 90
on OST retainers

GEO does not replace classic SEO. It runs alongside it. The same well-structured, fast-loading, authoritative pages that rank well in Google also get cited by AI engines. The difference is the additional layer: schema markup, citation-ready writing, and a measurement stack that tracks AI-surface exposure separately from blue-link clicks.

This paper is organized into eight sections. Sections 01 through 03 cover the foundation: why GEO matters, how AI citation works, and the technical prerequisites. Sections 04 and 05 cover content and measurement. Section 06 is the executable 30-day checklist. Sections 07 and 08 cover common mistakes and author credentials. You can read straight through, or jump to the checklist on page 21 and work backward.

Contents

- 01 The shift from search to answers
- 02 How AI engines decide what to cite
- 03 Technical foundations
- 04 Content patterns that work
- 05 Measurement and monitoring
- 06 The 30-day GEO audit checklist
- 07 Common mistakes
- 08 About the authors


How to use this paper

Each section is self-contained. Technical readers should prioritize Sections 03 and 06. Content leads should focus on Sections 04 and 06. Leadership can get the key facts from the Executive Summary and Section 07.

01

The shift from search to answers.

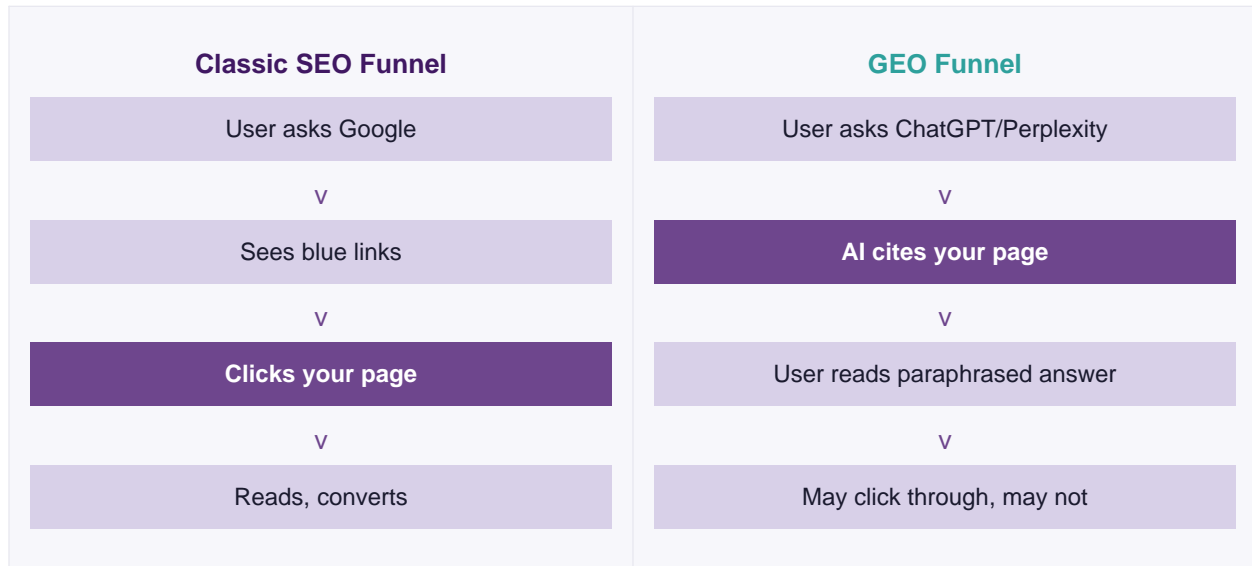
For two decades, organic search optimization meant one thing: rank in the top three blue links. That game is not over, but it is no longer the only one. Four AI surfaces are now intercepting questions before users ever see a search results page.



The shift matters because the shape of the funnel changed. In classic SEO, you optimized to rank. The user clicked your blue link, landed on your page, and entered your funnel. In GEO, you optimize to get cited. The AI engine reads your page, paraphrases it into the answer, attributes you with a small clickable citation, and the user may or may not click through.

This change has two distinct effects on traffic. First, click-through rates (CTR) from AI citations are lower than from a #1 organic ranking. Users who see a thorough AI answer often do not need to click. Second, and more importantly, brand mention exposure scales independently of clicks. A user who sees your brand cited three times across different AI answers in a week has already formed an impression before ever visiting your site.

Two funnels, two optimization targets



A page can earn citations across all four major AI surfaces simultaneously: ChatGPT (with built-in search), Perplexity (search-native by design), Claude (with the Claude search tool), and Google AI Mode / AI Overviews. Each weights signals differently, but the foundation is shared. Get the foundation right, and you compete on all four.

Why this matters now, not next year

AI Overviews have appeared on more than a third of mobile Google queries since rollout. ChatGPT's search rolled out broadly in late 2024 and reached parity with classic search for transactional intent in 2025. Perplexity passed 100 million weekly active users in 2025. The traffic patterns these surfaces produce are different from classic search:

- **Lower volume, higher intent.** Click-through rates from AI citations are lower than from a #1 organic ranking, but the visitors who do click have already read a paraphrased version of your answer. They arrive qualified.
- **Brand-mention exposure scales independently of clicks.** Users see your brand in the answer even if they never click. That builds awareness on a different axis than CTR.
- **Multi-engine compounding.** The same content investment ranks you in four AI engines plus classic Google search at the same time, because the technical foundations overlap.

Traditional SEO dashboards measure rankings and clicks. Neither metric captures what happens when your brand name appears in an AI-generated answer without a click. This is the new brand-awareness channel. A user who sees your name cited three times in a week across Perplexity and Google AI Mode arrives at your site, when they finally do click, with a different level of pre-built trust than a cold organic visitor.

GEO investment compounds across both the citation-traffic channel and the brand-impression channel simultaneously. Brands that move early on GEO will have a citation footprint that is difficult for late movers to displace, because topical authority and external citations of you take time to build.

The opportunity window

Most mid-market B2B and ecommerce brands had not run a GEO audit as of early 2026. That is the window. Categories where one or two brands have done the technical work are already pulling away in AI citation share. The cost of entry is lower now than it will be in 18 months.

02

How AI engines decide what to cite.

The AI surfaces do not all share a public ranking algorithm. But behavior across hundreds of OST prompt-coverage tests in 2025 points to four shared signals.

Topical authority & depth Multiple pages, internal links, original data, and external citations of you. Beats keyword density every time.	Citation-ready paragraphs Short, factual, attributable sentences that an LLM can lift unchanged. Bury the qualifiers.
Schema & structured data Article, Product, FAQ, HowTo, Organization. Tells engines what each page is without making them guess.	Recency & dated facts "As of April 2026" beats "as of recently." Old, undated content drops out of time-sensitive answers.

1. Topical authority and depth, not keyword density

AI engines are not counting how often you say "ATV parts." They are reading your pages to decide whether you are a recognized authority on ATV parts. Topical depth (multiple pages, internal links, original data, expert citations of you) wins over keyword stuffing every time.

The practical implication: build topic clusters, not one-off pages. A main pillar page supported by five supporting pages, each with internal links and consistent schema, creates the kind of topical web that AI engines recognize as authoritative. A single long page, however detailed, is not the same thing.

2. Citation-ready paragraphs

An LLM will cite a sentence that is short, factual, attributable, and unambiguously about a single thing. A 100-word paragraph full of qualifiers and marketing language is harder to cite. Write for the

LLM's quoting behavior: state the claim cleanly, then expand.

"The single highest-leverage change we make on client sites is rewriting the lead paragraph of every key page so it answers the page's primary question in the first 40 words. Citation rates on those pages roughly double within 90 days."

3. Schema and structured data

Schema.org structured data lets you tell engines what type of thing each page is (Article, Product, FAQ, HowTo, Organization, Review, etc.) without making them guess. AI engines lean on this heavily for product Q&A, FAQ surfacing, and author / publisher attribution.

4. Recency, dates, and freshness signals

"As of April 2026" beats "as of recently." Dated facts are more citable. AI engines avoid old or undated content for time-sensitive questions, which is most B2B and ecommerce content. Every page should carry a visible publish or last-updated date. Every statistic should be dated inline.

Freshness also extends to schema: a Product page with a "dateModified" schema field that is current signals to AI crawlers that the information is up to date. An unchanged page with no modification date is treated as potentially stale.

Signals that do not drive AI citation

For context, here are signals that classic SEO relies on heavily but that have weak or unmeasured impact on AI citation:

- **Domain authority as a standalone signal.** AI engines do not appear to weight DA/DR the way Google does. A well-structured page on a mid-authority domain can outrank a thin page on a high-authority domain.
- **Keyword density and exact-match anchors.** Rewriting a page's H1 to include the target keyword, with no other changes, does not move AI citation rates in OST tests.
- **Backlink volume alone.** Links matter for classic SEO. For GEO, what matters is whether authoritative sources cite or quote you in their content -- different from a plain backlink.

Key takeaway

The four shared AI citation signals are: authority and depth, citation-ready paragraphs, schema, and dated facts. Optimize for those four and you will compete in all four AI engines plus classic Google search at the same time.

03

Technical foundations.

This is the layer where most GEO retainers spend the first 30 days. None of it is novel. All of it is the same work that supports accessibility, classic SEO, and screen-reader compatibility. The difference is that GEO makes the payoff visible faster than any of the others.

Semantic HTML and heading hierarchy

One H1 per page. Logical H2 / H3 nesting. No skipped levels. Screen readers and AI crawlers both navigate by heading structure. A 6,000-word post with a single H1 and no H2s is, to an LLM, an undifferentiated wall of text. The crawler ingests it but cannot identify the sections or extract a citation-ready sentence that is clearly about a single topic.

The fix is mechanical: audit every page template, not every page. If the template is right, every page built from it is right. Most mid-market sites have three to six templates. Fix the templates in a week and every page benefits at once.

Schema.org structured data

The schema types most relevant for GEO across B2B and ecommerce stacks:

Schema Type	When to Use	What It Unlocks
Organization	Sitewide, on every page	Brand attribution in citations
Article	Blog posts, white papers, news	Author, publish date, headline
FAQPage	Pages with Q & A blocks	FAQ rich results, voice answers
HowTo	Step-by-step instruction pages	Step-numbered AI answers
Product	Every PDP on an ecommerce store	Product Q&A; in AI shopping
Review & AggregateRating	Pages that aggregate reviews	Trust signals in AI answers

BreadcrumbList	Sitewide	Navigation context for AI
-----------------------	----------	---------------------------

The validation tooling has matured: Google's Rich Results Test, Schema Markup Validator (schema.org), and the AI-bot specific checkers from Profound and AthenaHQ all give clear pass / fail output. There is no excuse for invalid schema in 2026.

One common schema error: the `datePublished` and `dateModified` fields are set once at launch and never updated. Every time a page is substantially revised, the `dateModified` field should be updated. This is a 10-second task that has a measurable impact on AI crawl freshness.

robots.txt and the AI bot question

This is the spot where well-meaning teams shoot themselves in the foot. There is a real debate about whether to allow AI crawlers, and it is legitimate. But here is what blocking them costs you:

Watch out:

Blocking GPTBot in robots.txt removes you from ChatGPT search citations entirely. Same for ClaudeBot (Claude), PerplexityBot (Perplexity), and Google-Extended (the AI training crawler). If your goal is to be cited, you cannot also block. Pick one.

Our standard recommendation: allow the search-time crawlers (GPTBot, PerplexityBot, ClaudeBot), and selectively allow or disallow the training crawlers (Google-Extended, Common Crawl, GPTBot's training mode) based on your IP-protection comfort level. Search-time bots only fetch when a user asks; they do not retain content for model training.

```
# Example robots.txt for a GEO-friendly site
User-agent: *
Allow: /

# Search-time AI crawlers (allow)
User-agent: GPTBot
Allow: /
User-agent: ClaudeBot
Allow: /
User-agent: PerplexityBot
Allow: /

# Training crawlers (your call)
User-agent: Google-Extended
Disallow: /

Sitemap: https://example.com/sitemap.xml
Sitemap: https://example.com/llms.txt
```

The llms.txt convention

An emerging open standard. A markdown-formatted file at /llms.txt that summarizes your site's structure, key pages, and topical scope for LLM consumption. Adoption is uneven (early 2026), but cost to ship is low and likely-future-value is high. We add one to every client site. The file is

human-readable and machine-friendly: a table of contents for your site that LLMs can ingest in a single pass.

Page speed, Core Web Vitals, and AI

AI engines crawl on a budget like classic search engines. Slow pages drop out of the crawl. Core Web Vitals (LCP, INP, CLS) directly affect crawl frequency and citation freshness. The sites that show up in AI answers tend to load fast and ship regularly. A page that takes 6 seconds to load on mobile may be crawled quarterly instead of weekly, which means it misses freshness windows that would get it cited.

04

Content patterns that work.

Once the technical foundation is in place, content writing is the next leverage point. Six patterns surface most often in citation analysis across OST client sites and competitor audits.

<p>Pattern 01 The 40-word lead</p> <p>Open every key page with 40 words that directly answer the page's primary question. No throat-clearing.</p>	<p>Pattern 02 Stand-alone definitions</p> <p>Bold the term, define inline in one sentence. An LLM can lift it unchanged.</p>
<p>Pattern 03 Original numbers</p> <p>Proprietary data with dates and attribution outranks restated industry averages.</p>	<p>Pattern 04 Comparison tables</p> <p>Structured rows and columns let LLMs extract specific cells without reasoning.</p>
<p>Pattern 05 Step-numbered procedures</p> <p>Ordered lists paired with HowTo schema unlock voice and step-by-step AI answers.</p>	<p>Pattern 06 FAQ blocks (80-150 words)</p> <p>Substantive Q&A; with FAQPage schema gets cited at high rates for long-tail queries.</p>

The six patterns are ordered by ease of implementation. Patterns 01 and 02 can be applied by any writer in hours, with no technical changes. Patterns 03 and 04 require content work and possibly design support. Patterns 05 and 06 require schema implementation alongside content changes. A site that ships all six patterns across its top 20 pages typically sees meaningful citation movement within 60 to 90 days.

Common misapplication to avoid

Teams sometimes apply these patterns to low-traffic pages first because they are easier to edit. Prioritize pages that already receive search traffic. An AI engine is far more likely to cite a page it already crawls regularly than a new or low-traffic page, no matter how well it is written.

Pattern 01: The 40-word lead

Open every important page with a 40-word paragraph that directly answers the page's primary question. No throat-clearing. No "in today's fast-paced world." If the page answers "what is WCAG 2.2 AA," the first 40 words say what WCAG 2.2 AA is. This single change, applied across your top 20 pages, is often the fastest path to measurable citation lift.

Pattern 02: Definitions that stand alone

Every key term on the page should have a one-sentence definition that an LLM can lift unchanged. Bold the term. Define it inline. Do not bury it three paragraphs in. An inline definition is also more readable for humans, which improves dwell time -- a secondary positive signal.

Pattern 03: Original numbers and data

"OST's 2025 client survey found that 67% of B2B companies under \$25M revenue had not run a single GEO audit." That sentence is gold for citation. AI engines preferentially cite original, dated, attributable data over restated industry trends. If you have proprietary data, lead with it. Even small surveys of your own customers are valuable as long as methodology and date are stated.

Pattern 04: Comparison tables

Tables comparing options, vendors, plans, or technical specifications get cited more than equivalent prose. The structured row / column format lets the LLM extract specific cells without having to reason through prose paragraphs. On product-comparison pages, switching from prose to a table has produced citation-rate lifts of 40% or more in OST tests.

Pattern 05: Step-numbered procedures

HowTo content with numbered steps gets cited cleanly. Use ordered lists. Pair them with HowTo schema. The same content unlocks AI Overviews, ChatGPT step-by-step answers, and voice-assistant guidance. Each step should be 20 to 50 words: enough to be actionable, short enough to quote.

Pattern 06: FAQ blocks with expanded answers

FAQs at the bottom of a page (paired with FAQPage schema) get cited at high rates for long-tail queries. Each Q&A; should be substantive: 80 to 150 words per answer. Yes/no answers do not get cited. Explanations do. Write each answer as if it will be read in isolation -- because in an AI answer, it will be.

Content investment return

Applying all six patterns to a single high-traffic page typically takes 2 to 4 hours. In OST retainers, pages that receive this treatment show measurable citation improvement within 60 days. Prioritize pages that already rank in positions 4 to 15 in classic Google -- those have the traffic signals that support AI citation without needing additional link building.

05

Measurement and monitoring.

You cannot improve what you cannot measure. The standard GEO measurement stack covers four layers, each giving a different signal about how AI search is treating your content.

Layer 01: GA4 referrer tracking	AI-search domains as a custom channel group	Continuous
Layer 02: GSC AI-Overview impressions	Google Search Console AI exposure curve	Weekly
Layer 03: Server-log AI-bot crawls	GPTBot, ClaudeBot, PerplexityBot, Google-Extended	Weekly
Layer 04: Prompt-coverage testing	25-50 prompts run across all four AI engines	Monthly

Layer 01: GA4 referrer tracking

Identify the AI-search referrers and tag them as their own channel group in GA4. The most common: chat.openai.com, perplexity.ai, copilot.microsoft.com, gemini.google.com, and Google AI Overviews (which appear under the standard Google referrer but with a UTM signal). Build a custom GA4 channel group called "AI Search." Once built, this channel group is persistent -- you see AI-search traffic trends over time without re-configuring.

Layer 02: GSC zero-click and AI Overview impressions

Google Search Console now reports impressions in AI Overviews separately from classic search. Track them. The shape of the curve tells you whether your AI exposure is growing, flat, or shrinking. Pages where impressions are growing but clicks are flat are getting AI Overview appearances without click-throughs -- a sign the AI is paraphrasing you rather than routing users to your page.

Layer 03: Server log analysis for AI bots

Watch for visits from GPTBot, ClaudeBot, PerplexityBot, and Google-Extended. Crawl frequency is a leading indicator of citation eligibility. Pages that AI bots crawl weekly are cited more than pages they crawl quarterly. If key pages are not being crawled, investigate: robots.txt, page speed, internal linking gaps, and whether the page is included in your XML sitemap.

Server log analysis does not require enterprise tooling. GoAccess is a free open-source log analyzer that produces bot-visit summaries from raw access logs. Cloudflare and AWS CloudFront both expose bot user-agent data in their analytics dashboards with no additional cost.

Layer 04: Prompt-coverage testing

The most direct measure. Take a list of 25 to 50 commercial-intent prompts your customers might ask, run each prompt against ChatGPT / Perplexity / Claude / Google AI Mode, and record whether your domain appears as a citation. Re-run monthly. The trendline is your GEO scoreboard. Tools like Profound, AthenaHQ, and Otterly automate this and provide dashboards. Manual testing with a spreadsheet is fine for sites under 100 pages.

When building your prompt list, cover three tiers: brand prompts (direct brand or product questions), category prompts (what is the best X for Y use case), and comparison prompts (X vs. Y). Category and comparison prompts drive the most new-audience exposure. Brand prompts measure your existing reputation. Track all three separately.

Layer	Tools	Cadence
GA4 referrers	Google Analytics 4	Continuous
GSC AI impressions	Google Search Console	Weekly
Server log AI bots	Cloudflare Logs, AWS CloudFront, GoAccess	Weekly
Prompt-coverage tests	Profound, AthenaHQ, Otterly.AI, manual	Monthly

A complete GEO measurement stack requires no custom infrastructure. GA4 and Google Search Console are free. Server log access is standard on any managed hosting plan. Prompt-coverage testing can start with a spreadsheet and 30 minutes a month. The tooling can grow as the program matures.

When to escalate to a paid tool

If your target prompt list exceeds 100 queries, or you need competitor citation data, or you want automated weekly reporting, move to Profound or AthenaHQ. Both tools start under \$500/month for a single domain at the time of writing.

06

The 30-day GEO audit checklist.

Run this in order. A two-person team (one technical, one content) can complete it inside 30 calendar days for a 50-to-200-page site. Larger sites take longer on execution, not on analysis.

1

Day 1-2

Inventory crawler access

Pull robots.txt. Identify which AI bots are allowed and which are blocked. Document state before changing.

2

Day 3-5

Schema audit

Run every key template through Schema Markup Validator. Note gaps for Organization, Article, Product, FAQ, HowTo.

3

Day 6-8

Heading hierarchy audit

Flag pages with missing H1, multiple H1s, skipped levels, or single-H1 pages over 1,500 words.

4

Day 9-12

Prompt-coverage baseline

List 30-50 commercial-intent prompts. Run each in ChatGPT, Perplexity, Claude, Google AI Mode. Record citations.

5

Day 13-16

Lead-paragraph rewrites

Rewrite the first 40 words of each top page to directly answer the page's primary question. Bold key terms.

6

Day 17-20

Schema implementation

Ship the missing schema types. Validate every template. Re-test in Google's Rich Results Test.

7

Day 21-23

FAQ block additions

Add FAQ blocks with FAQPage schema to your top 10 pages. Each Q&A; 80-150 words. Real questions, real answers.

8

Day 24-26

llms.txt + sitemap

Publish a markdown-formatted /llms.txt summarizing site scope and key pages. Add to sitemap declarations.

9

Day 27-28

Measurement stack

Configure GA4 channel group for "AI Search." Set up server-log AI-bot reporting. Schedule monthly re-tests.

10

Day 29-30

Document baseline and targets

One-page report: starting prompt-coverage score, baseline crawl frequency, 90-day target. Distribute internally.

What "good" looks like at day 90

For a typical mid-market B2B site running this checklist, expect to see prompt-coverage scores rise from 8-15% baseline to 25-40% inside 90 days, AI bot crawl frequency double, and AI-Overview impressions grow visibly in GSC. Sites that already have good SEO and accessibility see faster gains. Sites with stale content and weak structure see the biggest absolute lift, but it takes longer.

07

Common mistakes.

From OST's GEO retainer work in 2025, ranked roughly by frequency.

1 Blocking AI bots without strategic intent

- Teams reflexively block GPTBot in robots.txt for IP reasons, then complain that they get no AI citations. Pick one outcome.

2 Schema once, schema never

- Schema gets shipped, then ages. Product schema with stale prices and old reviews actively hurts. Schedule quarterly schema audits.

3 Long-form pages with weak structure

- 4,000-word pages with two H2s and no FAQ blocks are nearly impossible to cite. Chunk them or add structure.

4 Measuring blue-link CTR only

- Classic SEO dashboards underreport AI exposure. Build a separate channel group and run prompt-coverage tests monthly.

5 "Just publish more content."

- AI engines do not preferentially cite high-volume publishers. They cite citable, authoritative, well-structured content. Quality over cadence.

6 Treating GEO as a 90-day project

- The technical lift is one-time. The content cadence and prompt-coverage testing are continuous. Budget for the latter.

This paper is a snapshot of the practice as it stands in May 2026. AI search ranking signals will keep shifting. The technical foundations covered here are the parts least likely to change, because they are the same patterns that support accessibility, classic SEO, and screen-reader access. Investing in them is durable.

For organizations that want help running this checklist on their own site, OST offers GEO retainers with monthly prompt-coverage reporting and quarterly schema audits. Contact: hello@ost.agency

08

About the authors.

Shaili Gupta

President, OpenSource Technologies. Lead author.

GEO Practitioner certified. Holds GA4, Google Ads, and WCAG 2.2 AA Accessibility Oversight credentials. WBE-accredited principal. Drives OST's digital marketing practice and signs every GEO retainer that leaves the building.

Manish Mittal

CEO & founder. Technical co-author.

Forbes Technology Council member. AI architect on every OST AI engagement. Provided the technical-foundations and measurement sections of this paper. Architect of the OST AI Assistant.

OpenSource Technologies * ost.agency * hello@ost.agency

Version 1.0, May 2026. This paper may be shared freely with attribution.

Copyright 2026 OpenSource Technologies, Inc. All rights reserved.