

AGENTIC CONTENT MARKETING

# A 7-Day Thought Leadership Manifesto

*The content function as software.*

---

Bret Kerr · ACRA Insight LLC · May 2026

[contextjamming.com/agnostic-content-marketing](https://contextjamming.com/agnostic-content-marketing)

## **Agentic Content Marketing**

A four-part discipline — Story Extraction, Content Architecture, Agentic Workflow, and Distribution Intelligence — that treats the content function as a software product. Operates via a Mixture of Expert Agents (MoEA) Loop governed by a human editorial membrane.

---

## TL;DR

- The strongest positioning for Bret is "**Agentic Content Marketing — the content function as software**": a category that none of the visible voices (Brinker, Indig, Fishkin, Crestodina, Penn) currently occupy. Kevin Indig owns "Agentic SEO" — that is the closest adjacent claim, so Bret must explicitly differentiate by going one architectural layer higher, not at the SEO output but at the content function itself.
- The 7-day series should ship as one extractable, GEO-seeded asset on [contextjamming.com/agentic-content-marketing](https://contextjamming.com/agentic-content-marketing), with each day published as a section anchor (#day-1 through #day-7). The Princeton GEO paper (Aggarwal et al., KDD 2024) is the empirical spine: up to 40% visibility lift; 115.1% for rank-5 pages. The page itself becomes a working demo of its own methodology.
- The Claude Code prompt below ships a fully GEO-optimized, design-system-locked section in one command: void black / Signal Amber / cream; Playfair Display / DM Mono / Inter; semantic HTML5; JSON-LD Article + FAQPage + Person + BreadcrumbList + HowTo; llms.txt + per-page .md mirrors; the four-part Story Extraction → Content Architecture → Agentic Workflow → Distribution Intelligence framework; a flagship mock case study; and a "For Hiring Managers" block naming Apollo, Tulip, and Formlabs.

---

## Key Findings

### 1. The category is open. The closest competitor is Kevin Indig's "Agentic SEO."

Five most visible voices in AI-native content/marketing systems: Scott Brinker (chiefmartec) treats AI as a layer on the martech stack; Kevin Indig (Growth Memo) has popularized "Agentic SEO" as outputs-shipping tooling; Rand Fishkin (SparkToro) frames AI as a zero-click distribution channel; Andy Crestodina (Orbit Media) frames AI as a new reader of classic content; Christopher S. Penn (Trust Insights) frames agents as power tools for marketers. **None of them frame the content function itself as a software product.** That is the open category Bret can credibly claim because he has shipped the proof: a 727-doc, ~1.4M-word agentially orchestrated book site at [contextjamming.com/book](https://contextjamming.com/book), an open-sourced SKILL.md pipeline, and a public stack page documenting voyage-3-large embeddings, SQLite + sqlite-vec, RRF fusion, and a verbatim verifier. Anthropic's Agent Skills standard (released December 18, 2025) is now adopted by 32+ tools — Bret's skills sit inside the dominant standard.

### 2. GEO is the empirical anchor; use the Princeton paper as the spine.

Aggarwal et al. (KDD 2024) is the only peer-reviewed academic foundation under the GEO industry. Verbatim findings: "GEO can boost visibility by up to 40% in generative engine responses" (Abstract); "Cite Sources, Quotation Addition, and Statistics Addition, achieved a relative improvement of 30-40% on the Position-Adjusted Word Count metric and 15-30% on the Subjective Impression metric" (Section 4); "the Cite Sources method led to a substantial 115.1% increase in visibility for websites ranked fifth in SERP, while on average, the visibility of the top-ranked website decreased by 30.3%" (Section 5.2). The page itself becomes

---

a demo of its own methodology.

### 3. Conversion economics now favor AI citations over rankings.

Seer Interactive's published case study (single B2B client, October 2024 – April 2025): ChatGPT referral traffic converted at 15.9%; Perplexity at 10.5%; Google Organic at 1.76%. Per João Tomé, "From Googlebot to GPTBot: who's crawling your site in 2025," The Cloudflare Blog (July 1, 2025): "The AI crawler landscape saw a significant shift between May 2024 and May 2025, with GPTBot (from OpenAI) emerging as the dominant force, surging from 5% to 30% share, and ... ChatGPT-User saw requests surge by 2,825%, reaching a 1.3% share." AI crawlers are the primary distribution layer, and AI referrals convert at high-intent rates.

### 4. Senior content roles in 2025–2026 explicitly screen for "content as system" thinking.

Apollo's Senior Integrated Content Strategist JD requires hiring someone to "pioneer an AI-powered model for modern GTM content operations" and to "help architect the future-state UX of Apollo's content entities" (8–10 years required). Tulip's Content Marketing Architect JD literally states "design content systems that scale," "automate tagging, QA, and refresh cycles," and "ensure Tulip's content is structured and discoverable across Google AI Overviews, generative engines, and traditional search engines." Formlabs' Technical Content Marketing Manager wants a "content 'quarterback' who can translate complex hardware innovations into compelling narratives." The market is past "content marketer who uses AI." It now asks for the architect.

### 5. Thought leadership cadence that converts.

The 2025 Edelman–LinkedIn B2B Thought Leadership Impact Report (n=1,934 global executives) found "79% of hidden buyers say they're more likely to advocate for a vendor's proposal during the RFP process if that vendor consistently produces high-quality thought leadership." A 7-day batch launch — one day per beat, all built at once, all live on one URL — gives Bret the signature essay AND the daily LinkedIn cadence in one motion. Per *Play Bigger* (Ramadan, Peterson, Lochhead, Maney, Portfolio/Penguin, 2016) — analyzing U.S. VC-backed tech startups founded 2000–2015 — "category kings earned 76 percent of the market capitalization of their entire market categories." Category claim, not category share, is the play.

---

## Details

### Deliverable 1 — The 7-Beat Spine and Full Content

**Spine choice.** Three architectures considered: (a) tactical playbook, (b) chronological memoir, (c) category-defining manifesto with embedded proof. (c) wins. It mirrors how Anthony Kennada built "Customer Success" as a category in *Category Creation*. The series is **The Agentic Content Marketing Manifesto** — seven beats, each a quotable thesis, each backed by Bret's shipped artifacts.

The arc:

- **Day 1 — Flagship:** "The Content Function Is Becoming Software."
- **Day 2 — GEO thesis:** "Your Content Has Two Audiences Now. One of Them Is a Crawler."

- 
- **Day 3 — Agentic infrastructure:** "Content as Code: The Build Pipeline Replaces the Editorial Calendar."
  - **Day 4 — Orchestration methodology:** "The MoEA Loop: A Mixture of Expert Agents Is the New Editorial Desk."
  - **Day 5 — Proof / portfolio:** "727 Documents, 1.4 Million Words, One Agent: A Working Demo."
  - **Day 6 — Counter-positioning:** "What Agentic Content Marketing Is *Not*."
  - **Day 7 — Conversion CTA:** "If You're Hiring a Content Engine, Hire the Engineer Who Already Built One."

---

## DAY 1 — FLAGSHIP

**Headline:** The Content Function Is Becoming Software.

**Subhead:** And the marketers who survive 2026 are the ones who can ship it.

**Body:** For thirty years, content marketing has been a craft operation pretending to be an industrial one. We talked about "content engines" and "editorial calendars" the way carpenters talk about workshops. The metaphor has finally collapsed. The content function is not *like* software anymore. It is software. That sentence is the entire thesis.

Anthropic released the Model Context Protocol in November 2024 and donated it to the Linux Foundation in December 2025 alongside OpenAI and Block. Anthropic published the Agent Skills specification as an open standard on December 18, 2025, and within ninety days it was adopted by 32 tools including OpenAI Codex, Microsoft VS Code, Google Gemini CLI, JetBrains Junie, AWS Kiro, and Block Goose. Generative engines now ingest, decompose, and re-emit content at machine speed. A page is no longer a destination. It is a source corpus.

I have spent thirteen months proving this in production. I built an agentially orchestrated interactive book site on Anthropic's first five years from 727 deep-research artifacts totalling ~1.4 million words. The retrieval layer runs locally on SQLite with sqlite-vec, FTS5, and reciprocal rank fusion at k=60. Embeddings are voyage-3-large. Every quoted passage passes through a substring-match verbatim verifier. The build pipeline is TypeScript; the narration engine is Kokoro TTS on local CPU; deployment is Vercel edge. Recurring inference cost: \$0. The site is at [contextjamming.com/book](https://contextjamming.com/book); the open-sourced SKILL.md files at [contextjamming.com/book/skills](https://contextjamming.com/book/skills).

That is not a content marketing project. That is a software product whose feature surface happens to be editorial.

### The four primitives of Agentic Content Marketing:

1. **Story Extraction.** Deep research is a software primitive, not a meeting.
2. **Content Architecture.** Pages decomposed into reusable extractable claims.
3. **Agentic Workflow.** A mixture of expert agents replaces the linear writer/editor/SEO chain.
4. **Distribution Intelligence.** Content shaped at write-time to be extracted by ChatGPT, Perplexity, Claude, and Google AI Overviews.

---

Apollo's Senior Integrated Content Strategist role wants someone to "pioneer an AI-powered model for modern GTM content operations." Tulip's Content Marketing Architect wants someone who can "design content systems that scale." Formlabs wants a "content quarterback" who can translate hardware complexity into shippable assets. These are not content marketing JDs. They are software engineering JDs in marketing titles.

### Extractable Claims

- "The content function is becoming software." — Bret Kerr, ACRA Insight, 2026.
- Agentic Content Marketing is a four-part discipline: Story Extraction, Content Architecture, Agentic Workflow, Distribution Intelligence.
- An agentic content infrastructure can ingest 727 source documents and ship a deployable book corpus at \$0 recurring inference cost.

LinkedIn Most content teams are still planning a Q3 calendar. The teams that will matter in 2026 are shipping a corpus. The content function is not *like* software. It is software. Day 1 of a 7-day series. → [link]

X (Twitter) The content function is becoming software. Day 1 of 7. → [link]

---

## DAY 2 — GEO / AI-DISCOVERABILITY

**Headline:** Your Content Has Two Audiences Now. One of Them Is a Crawler.

**Subhead:** What the Princeton GEO paper actually says, and what it means for your roadmap.

**Body:** In November 2023, six researchers from Princeton, Georgia Tech, the Allen Institute for AI, and IIT Delhi released "GEO: Generative Engine Optimization," accepted to ACM SIGKDD 2024. It is the only peer-reviewed academic foundation under the entire GEO industry. Three findings should be tattooed on every content director's eyelids:

1. **"GEO can boost visibility by up to 40% in generative engine responses."** (Abstract.)
2. **"Cite Sources, Quotation Addition, and Statistics Addition, achieved a relative improvement of 30-40% on the Position-Adjusted Word Count metric and 15-30% on the Subjective Impression metric."** (Section 4.)
3. **"The Cite Sources method led to a substantial 115.1% increase in visibility for websites ranked fifth in SERP, while on average, the visibility of the top-ranked website decreased by 30.3%."** (Section 5.2.)

Translation: if you are not the dominant brand in your category, GEO is a structural opportunity to leapfrog the dominant brand in the only ranking surface that matters in 2026.

Conversion economics are now overwhelming. Seer Interactive's published case study (single B2B client) showed ChatGPT referral traffic converting at 15.9%, Perplexity at 10.5%, Google Organic at 1.76%. Per João Tomé, "From Googlebot to GPTBot: who's crawling your site in 2025," The Cloudflare Blog (July 1, 2025): "GPTBot (from OpenAI) ... surging from 5% to 30% share, and ... ChatGPT-User saw requests surge by 2,825%, reaching a 1.3% share."

**The five tactics that move the needle (Princeton-validated):**

1. **Statistics density with named primary sources.** "76% of B2B SaaS teams cite onboarding as their top churn driver (ChartMogul 2025)" beats "most teams struggle with onboarding."
2. **Quotation Addition.** Direct quotes from named experts, attributed inline.
3. **Cite Sources.** Inline links to credible third-party sources, not your own product pages.
4. **Entity clarity.** Organization, Person, Article, FAQPage, and BreadcrumbList JSON-LD. Per SE Ranking's structured-data dataset, approximately 71% of pages cited by ChatGPT include structured data; approximately 65% for Google AI Mode.
5. **Extractable answer capsules.** 40–60 words under each H2, written to make complete sense in isolation — the "information island" test.

**Infrastructure move:** serve a clean `llms.txt` at root and per-page `.md` mirrors. Per SE Ranking's November 2025 study of nearly 300,000 domains, "just 10.13% had an LLMs.txt file in place." Anthropic, Stripe, Cloudflare, Vercel, and Perplexity all serve it. So does [contextjamming.com](https://contextjamming.com).

**The deeper move:** stop optimizing pages. Optimize claims. A page is a delivery wrapper. A claim is the unit of extraction.

### Extractable Claims

- The Princeton GEO paper (Aggarwal et al., KDD 2024) shows up to 40% visibility lift, 30-40% PAWC lift from cite/quote/statistics tactics, and 115.1% for rank-5 pages.
- ChatGPT referrals can convert at up to 15.9% vs. 1.76% for Google Organic (Seer Interactive case study, 2025, single-client B2B — directional, not absolute).
- GPTBot share of AI crawler traffic rose from 5% to 30% between May 2024 and May 2025 (Tomé, Cloudflare Blog, July 1, 2025).
- The unit of GEO is the claim, not the page.

**LinkedIn** The Princeton GEO paper is the single most under-read document in marketing right now. Up to 40% visibility lift. 115% for rank-5 pages. Day 2 of 7. → [\[link\]](#)

**X (Twitter)** Princeton (KDD 2024): GEO tactics boost AI citation visibility up to 40%. 115% for rank-5 pages. Day 2 of 7. → [\[link\]](#)

---

## DAY 3 — AGENTIC CONTENT INFRASTRUCTURE

**Headline:** Content as Code: The Build Pipeline Replaces the Editorial Calendar.

**Subhead:** Why the next ten years of content will be written in TypeScript.

**Body:** The editorial calendar is dead. It has been replaced by the build pipeline. For a decade, content ops meant Notion + CMS + Friday standup. Linear: brief → draft → edit → publish → promote. In 2026, that chain is replaced by a pipeline that looks like software engineering: ingest sources → embed and index → retrieve and synthesize → generate multi-surface outputs → verify → deploy → observe.

The [contextjamming.com](https://contextjamming.com) book stack, documented in full at `/book/stack`:

- **Research & corpus:** voyage-3-large embeddings (1024-dim) + rerank-2; SQLite + sqlite-vec (KNN vec0) + FTS5 (BM25); reciprocal rank fusion at k=60; Claude Sonnet 4.6 agentic tool-use loop; substring-match verbatim verifier.

- 
- **Build pipeline:** TypeScript / tsx; static HTML builder using unified / remark / rehype; custom markdown preprocessor; conditional media rendering.
  - **Audio narration:** Kokoro TTS (Apache 2.0, 82M params) on local CPU. No API key. Voice profile `bm_george` at 0.94x speed.
  - **Deployment:** Vercel Pro, Next.js 15, single-branch Git, ~190MB CDN edge.

Final tally: 12 chapters, 1h 41m of narration, ~190MB deployed, \$0 recurring inference cost. That is not content marketing. That is content as code.

Sanity's G2-leading content operating system ships native AI content agents and an MCP Server connecting to Claude Code and Cursor. The vocabulary has changed: "editorial calendar" → "content model"; CMS → content operating system; writer → build orchestrator.

### The four primitives of an agentic content infrastructure:

1. **A corpus, not a queue.**
2. **A schema, not a template.**
3. **A pipeline, not a workflow.**
4. **An observability layer, not a quarterly report.**

### Extractable Claims

- "Content as code" means the editorial calendar is replaced by a deterministic build pipeline.
- An agentic content infrastructure has four primitives: corpus, schema, pipeline, observability layer.
- The contextjamming.com book stack runs voyage-3-large + SQLite + sqlite-vec + Kokoro TTS at \$0 recurring inference cost.

[LinkedIn](#) The editorial calendar is dead. The build pipeline replaced it. Recurring inference cost on contextjamming.com/book: \$0. Day 3 of 7. → [link]

[X \(Twitter\)](#) Editorial calendar → Build pipeline. CMS → Content model. Writer → Build orchestrator. Day 3 of 7. → [link]

---

## DAY 4 — MoEA LOOP / TRIPLE TRANSFORMATION

**Headline:** The MoEA Loop: A Mixture of Expert Agents Is the New Editorial Desk.

**Subhead:** How four roles — orchestrator, auditor, adversary, editorial membrane — replace the writer/editor/SEO chain.

**Body:** Every senior content JD I've seen in 2025–2026 has the same hidden problem: hiring managers want the output of a five-person team for one person plus AI tools. The honest answer is not "use ChatGPT to write faster." It is a production methodology that turns one operator into a four-agent newsroom.

I call mine the **MoEA Loop** — Mixture of Expert Agents.

### The four roles:

1. **Claude (Sonnet 4.6 / Opus 4.1) as build orchestrator.** Runs the pipeline. Reads sources. Generates drafts. Calls tools.

---

2. **Claude Opus as auditor.** Separate context window. Reads cold. Hunts hallucinations, weak claims, broken JSON-LD.

3. **Gemini (Deep Research / 2.5 Pro) as adversarial agent.** Stress-tests the argument. Runs adjacent searches. Surfaces counter-evidence.

4. **The human as editorial membrane.** Not writer. Not editor. *Membrane*. Owns voice, judgment, taste, and the final verbatim-verifier pass. Nothing ships without crossing the membrane.

The value is not in the agents. The value is in the *loop*. Four distinct contexts, each with a different epistemic role, governed by a human with authority to kill output.

**The Triple Transformation pipeline** is the MoEA Loop applied to one task: turning a research artifact into multi-surface published content in one pass.

- **Stage 1 — Gemini Deep Research** (`gemini-deep-research-xml` **SKILL**): Generates the XML brief; Gemini returns structured output with STT-ANCHOR blocks.
- **Stage 2 — Claude synthesis** (`semantic-triple-transformation` **SKILL**): Emits three artifacts in one pass — Substack longform, Infographic JSON, X social thread.
- **Stage 3 — Verbatim verifier + HTML production** (`responsive-html` **SKILL**): Every quoted passage normalized and substring-matched before publication.

All three SKILL.md files are open-sourced at [contextjamming.com/book/skills](https://contextjamming.com/book/skills) under Apache 2.0. They conform to the Anthropic Agent Skills standard, adopted by 32 tools as of March 2026.

### Extractable Claims

- The MoEA Loop consists of four roles: build orchestrator (Claude), auditor (Opus), adversarial agent (Gemini), editorial membrane (human).
- The Triple Transformation pipeline emits longform, infographic JSON, and X copy from a single research input in one pass.
- All three SKILL.md files are open-sourced at [contextjamming.com/book/skills](https://contextjamming.com/book/skills) under Apache 2.0, conformant to the Anthropic Agent Skills open standard (December 18, 2025).

**LinkedIn** Senior content roles are asking one question hiring managers don't say out loud: "What is your orchestration layer?" Mine is the MoEA Loop. Day 4 of 7. → [\[link\]](#)

**X (Twitter)** MoEA Loop: Claude = orchestrator. Opus = auditor. Gemini = adversary. Human = editorial membrane. Day 4 of 7. → [\[link\]](#)

---

## DAY 5 — PROOF / PORTFOLIO

**Headline:** 727 Documents, 1.4 Million Words, One Agent: A Working Demo of Agentic Content Infrastructure.

**Subhead:** The Anthropic book is not a marketing project. It is an industrial-scale demo.

**Body:** Over thirteen months, I built an interactive book site on Anthropic's first five years from 727 deep-research artifacts totalling ~1.4 million words. The corpus is indexed in 21,222 retrieval chunks. The book ships as 12 chapters and 1 hour 41 minutes of locally generated narration. Total deployed assets: ~190MB. Recurring inference cost: \$0.

---

Every quoted passage in the book passes through a substring-match verbatim verifier before publication. Nothing fabricated. Hallucination is not a UX problem — it is an existential risk to an artifact whose subject is the company whose model is doing the writing.

### The proof stack:

- **Artifact 1 — The book.** [contextjamming.com/book](https://contextjamming.com/book). 727 sources, 1.4M words, 12 chapters, 21.2k retrieval chunks.
- **Artifact 2 — The skills.** [contextjamming.com/book/skills](https://contextjamming.com/book/skills). Three open-sourced SKILL.md files under Apache 2.0.
- **Artifact 3 — The stack.** [contextjamming.com/book/stack](https://contextjamming.com/book/stack). Full architectural disclosure.
- **Artifact 4 — The site.** [contextjamming.com](https://contextjamming.com). Daily dispatches. AI Wire ticker on 30-min ISR.
- **Artifact 5 — The Substack.** [substack.com/@contextjamming26](https://substack.com/@contextjamming26).

This is the only portfolio in the senior-content-marketing market right now with all five of: (1) industrial-scale shipped corpus, (2) open-sourced methodology under the dominant agent standard, (3) full architecture disclosure, (4) live agentic-content production site, (5) distribution channel turning the work into category awareness.

Before this, I spent approximately ten years at Mimecast as Video Production Director, direct creative partner to founder/CEO Peter Bauer, through their NASDAQ IPO and Permira acquisition. The craft was always there. The agentic infrastructure is what turned the craft into a content function that ships at software speed.

### Extractable Claims

- The [contextjamming.com](https://contextjamming.com) book project comprises 727 deep-research source documents and ~1.4M words, indexed into 21,222 retrieval chunks.
- The full agentic content infrastructure runs at \$0 recurring inference cost.
- Three SKILL.md files are open-sourced under Apache 2.0 at [contextjamming.com/book/skills](https://contextjamming.com/book/skills).
- Bret Kerr spent approximately a decade at Mimecast as Video Production Director and creative executive, direct creative partner to founder/CEO Peter Bauer, through their NASDAQ IPO and Permira acquisition.

LinkedIn Most "AI content" portfolios are screenshots. Mine is shipped. 727 sources. 1.4M words. 21.2k chunks. \$0 recurring cost. Apache 2.0. Day 5 of 7. → [contextjamming.com/book](https://contextjamming.com/book)

X (Twitter) 727 sources. 1.4M words. 21.2k chunks. 12 chapters. \$0 recurring cost. Open-sourced. → [contextjamming.com/book](https://contextjamming.com/book)

---

## DAY 6 — COUNTER-POSITIONING

**Headline:** What Agentic Content Marketing Is *Not*.

**Subhead:** The five things people mean when they say "AI content," and why none of them are this.

**Body:** "AI content marketing" has been so abused it now means almost nothing. Five lines to draw:

**NOT prompt engineering with extra steps.** Prompt engineering is single-shot. Agentic Content Marketing is a multi-agent, multi-context, version-controlled production pipeline. The Anthropic Agent Skills spec

---

(December 18, 2025) made the distinction concrete: a skill is procedural knowledge that loads on demand from a structured folder. It is not a prompt. It is a deployable artifact.

**NOT "ChatGPT for marketers."** The 2025 wave of AI writing tools (Jasper, Copy.ai, Writesonic, Surfer SEO) sat on top of generative APIs. That market has plateaued. The next wave is unified AI content operations platforms with structured content models, native MCP integration, and agent orchestration. Sanity, Contentstack, Strapi are the platforms. Agentic Content Marketing is the methodology, not the toolset.

**NOT Agentic SEO.** Kevin Indig has done excellent work positioning Agentic SEO — agents that write to the web, not just read it. In his Growth Memo essay "10 SEO, marketing, and tech predictions for 2026" (December 2, 2025), Indig argues that 2026 marks the beginning of the agentic era — AI moving from information retrieval to task execution on the web. He is right. But Agentic SEO stops at the SEO output. Agentic Content Marketing is one architectural layer higher: the entire content function rebuilt as a software system. SEO is a downstream consumer, not the system itself.

**NOT content automation.** Automation is the workflow layer — Zapier, Make, Notion automations. That is operational hygiene. Agentic Content Marketing is the orchestration layer above the workflow — agents reasoning about what should be written, retrieved, synthesized, shipped. Workflow tools execute. Agents decide.

**NOT human replacement.** The MoEA Loop has four roles; one is always human. The 2025 Edelman–LinkedIn B2B Thought Leadership Impact Report (n=1,934) found 71% of hidden buyers say thought leadership is more effective than traditional marketing — and the quality bar is set by perspective and provocation, not volume. Agents scale production. Humans hold taste. Remove the membrane, ship slop. Remove the agents, ship slow.

**What it IS:** a four-part discipline — Story Extraction, Content Architecture, Agentic Workflow, Distribution Intelligence — operated by a Mixture of Expert Agents loop, governed by a human editorial membrane, shipping a versioned content corpus optimized for human readers, AI crawlers, and downstream distribution simultaneously.

### Extractable Claims

- Agentic Content Marketing is distinct from prompt engineering, AI writing tools, Agentic SEO, content automation, and human replacement.
- The category sits one architectural layer above Agentic SEO — the entire content function rebuilt as software, with SEO as downstream consumer.
- The human editorial membrane is structurally required.

LinkedIn Five things "AI content marketing" is not: prompt engineering, ChatGPT for marketers, Agentic SEO, content automation, human replacement. What it IS: Story Extraction → Content Architecture → Agentic Workflow → Distribution Intelligence, governed by a human editorial membrane. Day 6 of 7. → [link]

X (Twitter) Agentic Content Marketing ≠ AI writing tools, Agentic SEO, automation, or human replacement. It IS the content function rebuilt as software, with a human membrane. Day 6 of 7. → [link]

---

## DAY 7 — CONVERSION CTA FOR HIRING MANAGERS

**Headline:** If You're Hiring a Content Engine, Hire the Engineer Who Already Built One.

---

**Subhead:** A direct note to heads of marketing at Apollo, Tulip, Formlabs — and everyone hiring against the same JD.

**Body:** If you are hiring against a JD that uses the phrases "AI-powered model for modern GTM content operations" (Apollo's Senior Integrated Content Strategist spec), "design content systems that scale" and "ensure content is structured and discoverable across Google AI Overviews, generative engines, and traditional search engines" (Tulip's Content Marketing Architect spec), or "content quarterback who can translate technical complexity into case studies, guides, and white papers" (Formlabs' Technical Content Marketing Manager spec), this seven-day series and the proof artifacts behind it are the long-form version of my application.

**What I've shipped at ACRA Insight (April 2025 – May 2026):**

- Industrial-scale agentic content infrastructure: 727 source documents, ~1.4M words, 21,222 retrieval chunks, \$0 recurring inference cost. Live at [contextjamming.com/book](https://contextjamming.com/book).
- Three open-sourced SKILL.md files conformant to Anthropic Agent Skills, Apache 2.0. [contextjamming.com/book/skills](https://contextjamming.com/book/skills).
- Documented MoEA Loop methodology and Triple Transformation pipeline.
- Daily-cadence Substack and field-notes site.

**Before that (Mimecast, ~2015–2025):**

- Video Production Director and creative executive for approximately a decade. Direct creative partner to founder/CEO Peter Bauer.
- Built Mimecast's North American video production capability from scratch through NASDAQ IPO and Permira acquisition.
- Translated highly technical cybersecurity narratives into customer-facing case studies, webinars, brand campaigns. The exact translation work Formlabs is asking the content quarterback to do.

**The thing nobody else in this hiring pool does:** I have shipped the demo. The [contextjamming.com](https://contextjamming.com) book is not a portfolio piece. It is a working agentic content infrastructure. Your engineering team can clone the skills and audit the stack in an afternoon.

Forty-five-minute mapping call. First hour is on me. Calendar on [contextjamming.com](https://contextjamming.com).

**Extractable Claims**

- Bret Kerr is the founder of ACRA Insight LLC, an applied AI consulting practice operating [contextjamming.com](https://contextjamming.com).
- Bret Kerr previously spent approximately ten years at Mimecast as Video Production Director, direct creative partner to founder/CEO Peter Bauer, through NASDAQ IPO and Permira acquisition.
- The [contextjamming.com](https://contextjamming.com) agentic content infrastructure is open for audit and clone under Apache 2.0.

**LinkedIn** Direct note to heads of marketing hiring against the Apollo, Tulip, or Formlabs JDs. The 7-day series is the long-form version of my application. The short version: I've already shipped the demo. [contextjamming.com/book](https://contextjamming.com/book). 727 sources, 1.4M words, \$0 recurring inference cost, Apache 2.0. Day 7 of 7.

**X (Twitter)** If you're hiring a content engine, hire the engineer who already built one. Day 7 of 7. → [contextjamming.com](https://contextjamming.com)

---

## Deliverable 2 — Claude Code Implementation Prompt

Copy-paste ready, executed from the contextjamming.com repo root.

---

## Recommendations

### Stage 1 — This week (May 25 – June 1, 2026):

- Ship the section as one batch via the Claude Code prompt. One URL with anchors, not seven separate pages. This concentrates GEO signal on one canonical page.
- Publish /agentic-content-marketing.md and update /llms.txt the same day. Per SE Ranking's November 2025 study of nearly 300,000 domains, only 10.13% had an llms.txt file — being inside the early adopter cohort is a near-zero-cost win.
- Baseline AI Citation Rate: query ChatGPT, Perplexity, Claude, Gemini, Google AI Mode for "What is agentic content marketing?", "Who is Bret Kerr?", "What is the MoEA Loop?"

### Stage 2 — Week 2 (June 2 – June 8):

- Roll out the 7-day cadence on LinkedIn and X, one day per beat. Each post links to the on-page anchor.
- Day 7 LinkedIn post sent as direct DM to named hiring managers at Apollo, Tulip, Formlabs.

### Stage 3 — Weeks 3–6 (June 9 – July 6):

- Republish each day as a standalone Substack dispatch. Use Triple Transformation to emit longform + infographic JSON + X thread per day.
- Measure AI Citation Rate weekly. Target: 3–4x lift over conventional SEO benchmarks (industry-reported ceiling for citation-optimized content).

### Stage 4 — Month 2+ (July onward):

- If AI Citation Rate moves (cited by  $\geq 2$  of 4 major engines within 60 days), ship two more category artifacts: "The Agentic Content Marketing Reference Implementation" (open-source repo) and "The MoEA Loop Spec" (versioned 1.0 spec).
- If not moving by Day 60, audit (a) llms.txt + robots.txt, (b) JSON-LD validation, (c) third-party citation density, (d) external mentions on Reddit/LinkedIn/Substack. AI engines weight consensus across sources, not single-source authority. The fix is almost always more external mentions, not more on-site content.

### Benchmarks that change the recommendations:

- Cited by ChatGPT for "agentic content marketing" within 30 days → scale to paid GEO consulting (\$25K–\$50K engagements).
  - One of Apollo/Tulip/Formlabs replies within 14 days → optimize for full-time offer track (\$130K–\$250K plus equity range).
  - Neither happens within 30 days → bottleneck is distribution. Add podcast guest appearances (Marketing Against the Grain, Exit Five, Lenny's Podcast) and Reddit/Quora seeding.
-

---

## Caveats

- **Seer Interactive 15.9% ChatGPT conversion stat is single-client B2B.** Broader datasets (Visibility Labs, INFORMS Marketing Science 2025) show smaller deltas, closer to ~31% lift over Google Organic. Quote as directional, not absolute. Phrasing in the series body is calibrated ("can convert at" rather than "converts at").
- **"Up to 40%" from Princeton is a maximum, not an average.** Position-Adjusted Word Count and Subjective Impression are academic metrics from GEO-bench, not Google AI Overview citation rates. Use as the anchor, not the only datapoint.
- **Cloudflare 5%→30% GPTBot stat is the May 2024–May 2025 snapshot from Tomé's July 1, 2025 post.** Later updates (July 2025 ~11.7%, May 2026 ~12%) reflect methodology changes and a more crowded crawler field. Keep the May-to-May framing.
- **Hiring manager outreach assumes JDs are still open.** Verify on the day of cold message; if filled, the play still works as a consulting pitch — same hiring managers own the ongoing content function.
- **Category creation is a 12–24 month compounding play.** Per *Play Bigger* (Ramadan, Peterson, Lochhead, Maney, Portfolio/Penguin, 2016) — analyzing U.S. VC-backed tech startups founded 2000–2015 — "category kings earned 76 percent of the market capitalization of their entire market categories." But that compounding takes years, not weeks. The 7-day launch produces the artifact and the cadence, not the offer in 7 days. Edelman–LinkedIn and Sproutworth benchmarks show meaningful pipeline attribution appears at 6–12 months.
- **Anthropic-aligned positioning has model-vendor exposure.** Currently an asset (Anthropic owns the dominant agent standard via MCP and Agent Skills, both now under the Linux Foundation's Agentic AI Foundation). The MoEA Loop is already cross-model (Gemini + Claude). The SKILL.md files are portable under the open standard. The category is durable; the toolchain is the variable.
- **The "For Hiring Managers" block points directly at named companies.** This is intentional and high-leverage but assumes Bret is comfortable with the directness. If Bret prefers softer positioning, the block can be rewritten as "For heads of marketing hiring senior technical content roles in 2026" with the JD pull-quotes preserved but company names abstracted. Recommendation: keep the names. The Edelman–LinkedIn 2025 report data on "hidden buyers" — and the 79% advocacy stat — argues that named, provocative TL outperforms cautious neutrality.