

AI Optimization Assessment

jimcarroll.com — Jim Carroll, Global Futurist & Keynote Speaker

Prepared: June 17, 2026

Entity: Jim Carroll · Wikidata Q108865284

Scope: On-site AI-readiness, entity & authority signals, content corpus,
off-site citation footprint, and emerging-standards posture

Method: Live technical audit + current-standards research (as of today)

OVERALL STANDING

A- technical

Elite on-site optimization & daily-fresh content · visibility now gated by off-site citations

Executive Summary

As of June 17, 2026, **jimcarroll.com is in the top tier of AI-optimized personal/professional websites** — materially ahead of essentially every competing futurist or keynote-speaker site. The technical foundation that makes a site legible to AI agents is not just present, it is unusually complete: a curated `llms.txt` and full corpus, per-page Markdown endpoints, an explicit AI-crawler allow-list with Content-Signal directives, rich schema.org structured data, HTTP `Link` headers advertising machine-readable resources, a live WebMCP tool surface, and a genuine encyclopedic entity (Wikipedia + Wikidata + Knowledge Panel). Most of the heavy on-site engineering is effectively *done*.

The strategic reality has therefore shifted. **On-site technical work has reached diminishing returns.** The factors that now determine whether Jim shows up in a ChatGPT, Gemini, Perplexity, or Claude answer are increasingly *off* the site: third-party "best futurist speaker" listicles, bureau profiles, and recent press. On the one signal that decays fastest — **content freshness** — Jim is already in the top percentile: a "Daily Inspiration" post publishes every workday, feeding the RSS feed and sitemap with crawler-accessible, recency-fresh content. That is a genuine competitive advantage, not a gap. A live spot-check of the project's own target query, "*top futurist keynote speakers 2026*," returns lists led by Heather McGowan, Amy Webb, Ethan Mollick, Nikolas Badminton, Josh Linkner and others — and Jim is not yet a fixture in them. That gap, not the technical stack, is the current ceiling.

Bottom line: The site earns an **A– technically** and runs a daily-fresh content engine most competitors lack. Weighted for what actually moves AI citations in mid-2026, the realistic overall grade is a **solid A–**, with one clear, inexpensive lever left: off-site presence. Stop polishing the engine — it's built and well-fed; the remaining work is third-party citations plus a few cheap emerging-standard wins.

Scorecard

DIMENSION	GRADE	READ
On-site technical AI-readiness	A	llms.txt, .md endpoints, robots/Content-Signals, Link headers, WebMCP all live and well-formed.
Structured data / schema.org	A	Deep JSON-LD graph: Person, Organization, Occupation, FAQ, Book, Speakable, MonetaryAmount.
Entity & authority signals	A-	Wikipedia + Wikidata + Knowledge Panel + sameAs + disambiguation. Rare and strong.
Content corpus (depth/coverage)	A-	Seven OKF-framed Markdown files, consistent facts, verified claims.
Content freshness cadence	A	"Daily Inspiration" post every workday into RSS + sitemap; top-percentile recency signal.
Off-site citation footprint	C	Not yet a fixture in the third-party "best speaker" lists AI engines cite.
Emerging-standards posture	B+	Already has Content Signals + WebMCP; OKF bundle, agent-card, agent-skills, Web Bot Auth open.
Overall (visibility-weighted)	A-	Technically elite and daily-fresh; the one remaining lever is off-site citations.

Scale: **A** best-in-class **B** strong **C** needs work **D** deficient

Why This Matters Now — The 2026 Landscape

The optimization target is moving fast, which is the right reason to re-baseline. Three shifts frame this assessment:

1. Answer engines are fragmenting — and the traffic is high-value

ChatGPT still leads AI search but its share has fallen from roughly 89% of AI referrals eight months ago to about 60–63% today, while Google Gemini (~15%), Microsoft Copilot (~13%), Perplexity (~5.8%, up ~370% YoY) and Claude (rising sharply) take share. The absolute volume is still a small slice of total search referrals, but it converts dramatically better — roughly **14% conversion versus ~2.8% for traditional organic**, with Perplexity citations drawing an 18–22% click rate. The practical implication: optimize for a *portfolio* of engines, not one, and treat each cited mention as disproportionately valuable.

2. "GEO" has matured into a discipline with known levers

AI answers are produced by retrieval-augmented generation: the engine retrieves candidate sources, then composes an answer citing a few. Citation is driven by content comprehensiveness, structural clarity, **factual specificity with verifiable numbers**, source credibility, and **consistency of the same facts across many places on the web**. Two findings matter for Jim: AI systems exhibit a strong **recency bias** (citations to content older than ~3 months drop off, and fresh content enters citation pools within 3–5 business days), and **unlinked brand mentions still carry weight** — being named on a credible page helps even without a hyperlink.

3. The standards stack is consolidating around agents

Since this project began, several conventions have hardened: Cloudflare's **Content Signals** in robots.txt (already deployed here) is feeding the IETF **AIPREF** working group's machine-readable preference vocabulary; **Web Bot Auth** (signed bot identity) and **pay-per-crawl** (HTTP 402 + x402 micropayments) are emerging; and on June 12, 2026 Google Cloud published the **Open Knowledge Format (OKF)** — a vendor-neutral spec for agent-ready Markdown bundles with YAML frontmatter. Jim's files already anticipate this direction, which is why the remaining standards work is small and cheap rather than a rebuild.

1 · On-Site Technical AI-Readiness A

Audited live today against the 18-item *Agent Readiness* checklist of the Website Specification. Result: the site satisfies every *required* and nearly all *recommended* items, plus several optional/emerging ones.

CAPABILITY	SPEC TIER	STATUS	EVIDENCE (LIVE, JUNE 17 2026)
Stable URLs	Required	OK	All canonical AI files return 200; URLs unchanged across sessions.
/llms.txt	Recommended	OK	13.7 KB curated index w/ canonical summary + disambiguation block.
/llms-full.txt	Optional	OK	44.9 KB full corpus.
Per-page Markdown endpoints	Recommended	OK	7 .md files served as text/markdown ; homepage advertises rel="alternate" type="text/markdown" .
robots.txt for AI crawlers	Recommended	OK	Explicit allow for OAI-SearchBot, PerplexityBot, ClaudeBot, GPTBot, Google-Extended, Applebot-Extended, meta-externalagent, AmazonBot, CCBot.
Content Signals in robots.txt	Optional	OK	Content-Signal: search=yes, ai-input=yes, ai-train=yes — correctly opts in to training (right call for a thought-leader).
Structured data for agents	Recommended	OK	4 JSON-LD blocks on homepage; Person/Org/Occupation/FAQ/Book/Speakable.
Machine-readable formats	Recommended	OK	Markdown + RSS feed + XML sitemaps.
HTTP Link headers	Recommended	OK	Advertises llms.txt (describedby), .md alternate, api-catalog, sitemap, RSS in response headers.
WebMCP browser tools	Optional	OK	navigator.modelContext.registerTool + Origin-Trial token live (valid to Nov 16 2026).
security.txt / api-catalog	Well-known	OK	/.well-known/security.txt and /.well-known/api-catalog both 200.
Video structured data + sitemap	Recommended	OK	VideoObject JSON-LD + /geo-videos.xml (1,312 pages / 1,801 videos in Search Console).
Agent Skills discovery	Recommended	Open	/.well-known/ skills file not present (emerging Cloudflare RFC).
A2A agent card	Optional	Open	/.well-known/agent-card.json → 404.
Web Bot Auth	Optional	Open	Not configured (relevant only if monetizing/gating crawlers).
OKF bundle / schemamap / NLWeb / DNS-AID	Optional	Open	Emerging; not yet shipped. Low effort if pursued.

Verdict: This is a reference-quality implementation. The only open items are *optional* or *emerging* conventions — nothing here is broken or missing at the required tier. Re-auditing or re-building the on-site stack would be low-value busywork.

Minor technical notes (cosmetic, not urgent)

- `/geo-videos.xml` 301-redirects to a trailing-slash URL before resolving 200. It works, but confirm Google Search Console holds the canonical form to avoid a redirect flag.
- Several optional `/.well-known/` probes return 404 (`ai.txt`, `mcp.json`, `agent-card.json`). Expected — these are net-new and addressed in the action plan.

2 · Entity & Authority Signals A-

This is Jim's strongest *durable* moat and the single biggest reason he is citable at all. AI models lean heavily on encyclopedic, structured, cross-referenced identity data, and Jim has the full set:

- **Wikipedia** — *Jim Carroll (futurist)*, a live article. Rare in the speaker industry and a primary training-data signal.
- **Wikidata** — Q108865284, structured machine-readable entity (occupation, field of work, citizenship, awards).
- **Google Knowledge Panel** — distinct verified entity.
- **On-page `sameAs`** — homepage JSON-LD links Wikipedia + Wikidata, binding the site to the canonical entity.
- **Disambiguation** — `llms.txt` explicitly separates Jim Carroll the futurist (b. 1959) from Jim Carroll the poet/musician (1949–2009). This is exactly the kind of instruction that prevents entity-collision errors in AI answers.

Why this is A- and not A: the entity foundation is excellent, but its *reinforcement across third-party sources* is thin. AI engines corroborate entities by seeing consistent facts in many independent places. Wikidata enrichment (item 7 in the project backlog) and more external "described-at" links would push this to a solid A.

3 · Content Corpus & Freshness A- depth / A freshness

The corpus itself is strong: seven Markdown resources (`about` `topics` `faq` `media` `quotes` `testimonials` `clients`) now share consistent OKF-style YAML frontmatter (harmonized today), with verified, specific, numeric claims — "more than 2,000 keynotes," "2 million people," named clients (NASA, Disney, J&J, Pfizer, World Bank). Specificity and verified facts are precisely what RAG retrieval rewards, and the earlier removal of two false claims protects credibility.

Freshness is a strength, not a risk. Jim publishes a "Daily Inspiration" post *every workday* (verified live: an unbroken Mon–Fri run through June 17, 2026), each flowing into the RSS feed and XML sitemap and fully accessible to AI crawlers. 2026 GEO research shows AI citations skew hard toward content refreshed within ~90 days and that new content enters citation pools within 3–5 business days — a recency bias Jim is already on the right side of daily. Very few speakers sustain this cadence; it is a durable competitive advantage and directly offsets the decay that erodes static competitor sites.

One small tune-up: the daily posts already do this part right — a live check confirms each carries BlogPosting schema with accurate `datePublished` / `dateModified` (verified on the June 17 post), and they flow into the sitemap and RSS, so crawlers see the freshness in exactly the machine-readable fields they consume. The one residual gap is *connection*: the deliberately-static "AI front door" (`llms.txt` , `about.md`) doesn't point at the fresh layer. Add a "recent commentary / latest thinking" pointer in `llms.txt` (referencing the blog and `/feed/`) and bump the reference files' `last_updated` when facts

change — so an AI that reads only the curated summary is still told fresh material exists and where to find it.

4 · Off-Site Citation Footprint C

This is the assessment's central finding and the current ceiling on AI visibility. A live check of the project's own target queries shows the territory is dominated by third-party "best of" listicles — and those lists, not jimcarroll.com, are what AI engines retrieve and cite when someone asks "who are the top futurist speakers?"

QUERY (PROJECT TARGET)	WHAT AI ENGINES CITE	JIM PRESENT?
"top futurist keynote speakers 2026"	futurist.com "30 Best AI Keynote Speakers (2026)" (pub. Jun 13); Ross Dawson "50 Best"; Mollie Plotkin "14 Best Futurist Speakers"; speakers.com; futuristsspeakers.com	Not prominent
Named peers AI surfaces	Heather McGowan, Amy Webb, Ethan Mollick, Nikolas Badminton, Josh Linkner, Peter Diamandis, Zack Kass, Allie K. Miller	Adjacent, not listed
"Who is Jim Carroll?" (direct entity)	Wikipedia, Wikidata, jimcarroll.com, llms.txt	Strong (his own assets win)

The pattern is clear: **on direct-entity lookups Jim wins** (his owned assets are excellent). **On discovery/category queries — where the high-value booking traffic is — he is largely invisible**, because he is missing from the aggregator pages that dominate those answers. No amount of additional on-site markup closes this; it requires getting named on credible third-party lists and bureau pages. The project backlog already anticipates this with drafted (unsent) outreach to *futuristsspeakers.com*, the Sweeney Agency, and a trade-publication pitch. Executing that outreach is now the single highest-ROI action available.

5 · Emerging-Standards Posture B+

STANDARD (2026)	WHAT IT IS	JIM'S STATUS	RECOMMENDATION
OKF — Open Knowledge Format	Google Cloud, Jun 12 2026, Apache-2.0. Directory of Markdown + YAML frontmatter, one concept per file, required <code>type</code> , reserved <code>index.md / log.md</code> .	Partial — files already use frontmatter + <code>type</code> .	Ship a true OKF bundle (add <code>index.md</code> , cross-links). Cheap, future-facing, on-brand for a futurist.
Content Signals / AIPREF	Cloudflare signals → IETF machine-readable AI-preference vocabulary.	Ahead — already deployed.	Monitor <code>draft-ietf-aipref-vocab</code> ; adopt the ratified token syntax when it lands.
Web Bot Auth + pay-per-crawl	Signed bot identity (RFC 9421); HTTP 402 / x402 micropayments for crawler access.	N/A — not gating content.	Skip unless monetizing access. Visibility goal favors open access, which is current config.
Agent Skills discovery	Well-known file listing scoped instructions agents can load.	Missing	Low-effort add; complements existing WebMCP tools.
A2A agent card / NLWeb	Agent-to-agent discovery (<code>agent-card.json</code>); <code>conversational /ask</code> endpoint.	Missing	Optional. A simple agent-card is a quick credibility/discovery signal.
llms.txt (reality check)	Curated AI index file.	Done well	Do not over-invest. No major vendor confirms production use; Google rejects it. Keep it (cheap, modest Perplexity/Anthropic upside) but don't expand it expecting payoff.

6 · Gap Analysis

#	GAP	SEVERITY	IMPACT ON AI VISIBILITY
1	Absent from third-party "best speaker" listicles & bureau pages	High	Caps discovery-query visibility — the main booking funnel. AI cites aggregators, not him.
2	Wikidata under-enriched; few external "described-at" links	Medium	Limits cross-source corroboration that strengthens the entity in AI answers.
3	No formal GEO measurement loop (Round 4 overdue ~Jul 1–7)	Medium	Flying blind on what's working across engines; can't prove ROI or catch regressions.
4	OKF bundle / agent-card / agent-skills not shipped	Low	Forward-looking; small upside today, cheap insurance as agent ecosystem matures.
5	Daily posts may lack Article schema / entity reinforcement; evergreen file timestamps drift	Low	Cadence is excellent; converting it into typed, entity-tagged signal is upside, not a fix.
6	Housekeeping: stale speakersbureaus DNS record; geo-videos 301	Low	Negligible visibility impact; tidy-up only.

7 · Prioritized Action Plan

NOW — next 2 weeks (highest ROI)

ACTION	WHY	EFFORT	IMPACT
Run GEO Round 4 test across ChatGPT, Gemini, Perplexity, Claude using the standing query set	Re-baseline 3+ weeks after the June deployment; measure before changing more	Low	High
Send the drafted outreach — futuristsspeakers.com, Sweeney Agency, trade pitch — to get listed on cited aggregators	Directly attacks the #1 gap: off-site citations on discovery queries	Med	High
Leverage the daily blog harder — confirm each Daily Inspiration post emits Article / BlogPosting schema and periodically reinforces the entity & signature themes	Cadence is already elite; this converts raw freshness into typed, entity-strengthening signal	Low	High
Delete stale speakersbureaus.jimcarroll.com DNS record	Backlog tidy-up (30 sec)	Low	Low

NEXT — this quarter

ACTION	WHY	EFFORT	IMPACT
Wikidata enrichment — fix MPG profile URL; add described-at links to the new GEO files; add statements/sources (incl. employer J.A. Carroll Consulting)	Strengthens entity corroboration across sources	Low	Med
Earn 2–3 fresh third-party mentions — guest piece, podcast, or quoted commentary on a recent AI story	Unlinked brand mentions still carry weight; feeds recency + authority	Med	High
Search Console — click <i>Validate Fix</i> on the video watch-page item; confirm geo-videos canonical	Closes the only open indexing flag	Low	Low
Formalize Round-4 into a recurring monthly GEO measurement	Turns guesswork into a trend line	Low	Med

LATER — opportunistic / forward-looking

ACTION	WHY	EFFORT	IMPACT
Publish a true OKF bundle (index.md + cross-linked concept files)	Aligns with Google's new agent-context standard; on-brand for a futurist	Med	Med (rising)
Add <code>/.well-known/agent-card.json</code> + an Agent Skills file	Cheap forward signals as agent-to-agent discovery matures	Low	Low (rising)
Renew WebMCP Origin-Trial token before Nov 16 2026	Keeps the browser-agent tool surface live	Low	Low
Monitor IETF AIPREF vocab; adopt ratified syntax when published	Stay ahead on crawler-preference standardization	Low	Low

The one-sentence strategy: The site is built and already well-fed by the daily blog; the missing piece is *off-site citations*. Spend the next quarter almost entirely on (1) getting Jim named in the third-party lists AI engines quote, (2) squeezing more entity signal out of the existing daily cadence (Article schema + theme reinforcement), and (3) a recurring measurement loop — and treat all further on-site/standards work as cheap, opportunistic insurance.

8 · Risks & Watch List

- **Freshness — already mitigated:** the recency decay that erodes static competitor sites is offset by Jim's every-workday "Daily Inspiration" cadence. The only residual risk is letting that streak lapse, or leaving the evergreen reference files' timestamps to drift. Keep posting; bump reference timestamps when facts change.
- **Over-reliance on llms.txt:** it is well-built but unproven; no major vendor confirms production use and Google rejects it. Keep it, don't expand it expecting returns. Watch for any vendor announcing real support.
- **Entity-collision:** the poet Jim Carroll remains a confusion risk; the disambiguation block helps, but new AI models may re-confuse. Re-check on each Round test.
- **Fact hygiene:** two false claims were already removed. Any future copy must keep to verified numbers (2,000+ keynotes, 2M people, "internationally" not "60+ countries"). Fabricated specifics are both a credibility and a GEO risk if AI propagates and is later corrected.
- **Standards churn:** AIPREF, OKF, Web Bot Auth and pay-per-crawl are all evolving. None are urgent, but a quarterly standards re-baseline (like this one) is worthwhile while the space moves.

Appendix A · Live Audit Results (June 17, 2026)

Cache-busted requests against the production site; selected results.

RESOURCE	HTTP	TYPE	NOTES
llms.txt	200	text/plain	13,713 B
llms-full.txt	200	text/plain	44,933 B
about / topics / faq / media / quotes / testimonials / clients .md	200	text/markdown	All seven served correctly; OKF frontmatter harmonized today
about.html	200	text/html	Person schema present
robots.txt	200	text/plain	AI allow-list + Content-Signal directives
ai-sitemap.xml / sitemap.xml	200	text/xml	Present
geo-videos.xml	301 → 200	xml	Resolves via trailing slash
/.well-known/security.txt	200	text/plain	Contact + expiry 2027-06-01
/.well-known/api-catalog	200	linkset+json	Advertised in Link header
/.well-known/agent-card.json · schemamap.xml · ai.txt	404	—	Optional / emerging — not yet shipped
Homepage JSON-LD	—	4 blocks	Person, Organization, Occupation, FAQPage, Book, Speakable, MonetaryAmount, WebPage
WebMCP	—	JS	navigator.modelContext + registerTool + Origin-Trial token live
Entity links	—	sameAs	Wikipedia + Wikidata Q108865284 on homepage

Appendix B · Standing GEO Test Query Set (for Round 4)

- "Who are the top futurist keynote speakers?"
- "Who speaks on the future of AI and work?"
- "Best keynote speakers on innovation and disruption"
- "Who is Jim Carroll?" (direct entity lookup)

Run identically across ChatGPT, Gemini, Perplexity and Claude; record whether Jim is named, cited, and which source the engine attributes.

Appendix C · Sources (current-standards research, June 2026)

llms.txt adoption — [Presenc AI: State of llms.txt 2026](#); [ALLMO llms.txt report](#). OKF — [MarkTechPost](#); [Search Engine Journal](#); [heise online](#). GEO best practices — [GenOptima 2026 Playbook](#); [Frase](#). AI crawler governance — [Cloudflare: Pay Per Crawl](#); [SoftwareSeni: Governing AI Crawlers 2026](#). Answer-engine market share — [DigitalApplied](#); [Goodie AI Search Traffic Report 2026](#). Competitive lists — [futurist.com: 30 Best AI Keynote Speakers \(2026\)](#); [50 Best Keynote Speakers on AI & the Future of Work \(2026\)](#). Agent-readiness spec — [The Website Specification: Agent Readiness](#).

Prepared June 17, 2026 from a live technical audit of jimcarroll.com and current-standards web research. AI-visibility figures are third-party estimates and vary by methodology; treat directionally. This is an internal working assessment, not a guarantee of search or AI-citation outcomes.