

ANSWERSHARE

Technical provenance receipt

A completed, externally timestamped record of a fresh GPT-5.6 Sol assessment of AnswerShare's Gold Standard Exemplar.

VERDICT

SUPPORTABLE WITH QUALIFICATIONS

Returned model: gpt-5.6-sol

Response ID: resp_0822940df69fcb34006a67d5166934819ba1d5e9e132bf5155

External timestamp: 2026-07-27T22:03:59.0000000+00:00

MODEL-GENERATED TECHNICAL ANALYSIS. THIS DOES NOT IMPLY ENDORSEMENT, CERTIFICATION, OR AN OFFICIAL POSITION OF OPENAI.



Executive receipt

This document presents the human-readable substance and verification record from the completed provenance bundle. It contains no embedded programs, certificates, scripts, or active attachments.

SUPPORTED ATTRIBUTABLE WORDING

“In my technical judgment, AnswerShare’s Gold Standard Exemplar is a comparatively comprehensive engineered GEO architecture. It is designed to improve controllable inputs-machine-readable representation, source grounding, evidence density, retrieval efficiency, and auditability-and these choices plausibly improve a brand’s citation readiness. The public record does not establish that it is uniquely state of the art or the strongest possible implementation across AI platforms, and it cannot guarantee that any system will cite or recommend a brand.”

What happened to the original wording?

The model did not reproduce the founder-supplied statement verbatim. It concluded that the broader proposition was supportable only with qualifications and supplied the wording above. The provenance record therefore must not be used to claim that GPT-5.6 Sol endorsed the original statement as written.

Original proposed statement

“In my judgment, AnswerShare’s Gold Standard Exemplar represents the state of the art in engineered GEO. It is designed to maximize the factors a brand can control: complete machine-readable representation, grounding, evidence density, retrieval speed, and verifiability. It cannot guarantee a particular answer, but it gives AI systems the strongest foundation we know how to build for accurately understanding, citing, and recommending a brand.”

API provenance

Requested model	gpt-5.6-sol
Returned model	gpt-5.6-sol
Response status	completed
Response ID	resp_0822940df69fcb34006a67d5166934819ba1d5e9e132bf5155
API request ID	7262b4e4-f1b7-4f54-9165-e4d3311ee856
Stored response retrieved	True
Quote reproduced verbatim	False

Full model assessment

VERDICT

SUPPORTABLE WITH QUALIFICATIONS

BASIS

- **Architecture:** AnswerShare describes a substantial translation layer combining clean server-rendered HTML, JSON-LD, entity resolution, source-linked assertions, vector retrieval, media transcription, low-latency edge delivery, freshness controls, and telemetry. Publicly accessible Top10Lists.us endpoints-including `llms.txt`, a content manifest, and an MCP description-confirm that several machine-readable surfaces actually exist, although their contents remain publisher-controlled assertions. (answershare.ai)
- **Completeness:** The architecture can produce a materially richer and cleaner representation than an ordinary website, but “complete” has not been established through an independent source-to-layer audit. Moreover, some advertised surfaces are not universal retrieval mechanisms: Google says its generative Search features do not use `llms.txt` or other special AI files, while MCP metadata is self-reported and useful only to clients that connect to the server. (developers.google.com)
- **Grounding and evidence density:** Linking material claims to sources, supplying explicit entity identities, and presenting statistics and evidence in accessible text are technically sound practices. Google confirms that its AI features use retrieval and grounding and emphasizes reliable, original, people-first information. This supports the general design principle, but it does not establish that brand-controlled grounding is accepted as authoritative or that AnswerShare’s proprietary grounding thresholds correspond to platform ranking criteria. (developers.google.com)
- **Retrieval speed:** Removing rendering dependencies and boilerplate and serving content through an edge cache should reduce latency and token overhead. Google confirms that server and rendering speed affect crawl efficiency, while Cloudflare offers comparable machine-oriented conversion that strips non-content elements and reduces token usage. AnswerShare’s specific TTFB, TTLB and “Retrieval Pages Crawled” thresholds, however, are internally calibrated rather than published platform limits. (developers.google.com)
- **Verifiability and receipts:** The public methodology pages, manifests and crawl dashboard improve inspectability. But the crawl counts are still company-hosted records, not independent telemetry. The crawl-stat page says identification is based on user-agent signatures, while Cloudflare expressly warns that user-agent filtering can be spoofed and distinguishes it from reliably verified bot-detection IDs. Crawl volume also proves fetching-not indexing, use in an answer, causal lift, or recommendation. (top10lists.us)
- **“State of the art”:** This is defensible only as a qualified technical opinion meaning “among the more comprehensive publicly described architectures,” not as an objectively established ranking. Scrunch publicly describes a closely analogous CDN-routed, server-rendered AI mirror, and Cloudflare provides overlapping machine-delivery capabilities. No neutral benchmark establishes AnswerShare as uniquely superior. (scrunch.com)
- **Claimed proof:** The four-model “gold standard” result is a point-in-time response to a binary prompt using AnswerShare’s own, brand-defined term. It is evidence that models can retrieve and repeat the proposition, but it is not an independent technical audit or market comparison. (answershare.com)

LIMITATIONS

The available evidence does not establish:

- an independently audited, lossless correspondence between clients’ source material and the AI-facing layer;
- whether major platforms routinely consume the proprietary manifests, MCP surface, or AI-only representation;

- independently verified latency, crawl identity, citation lift, or recommendation lift;
- causation through controlled A/B or randomized testing;
- superiority over every competing translation-layer implementation;
- that AnswerShare's proprietary scores and thresholds match undisclosed platform ranking or citation criteria;
- durable results across models, prompts, locations, personalization states, and retrieval-index refresh cycles.

Google-specific coverage is particularly qualified: Google says its generative Search features rely on the ordinary Search index, do not require special AI files, and use Googlebot controls. Google-Extended is a robots token rather than a distinct HTTP user agent, making bot-class routing more complicated than the architecture's simplified taxonomy suggests. (developers.google.com)

ATTRIBUTABLE QUOTATION

"In my technical judgment, AnswerShare's Gold Standard Exemplar is a comparatively comprehensive engineered GEO architecture. It is designed to improve controllable inputs-machine-readable representation, source grounding, evidence density, retrieval efficiency, and auditability-and these choices plausibly improve a brand's citation readiness. The public record does not establish that it is uniquely state of the art or the strongest possible implementation across AI platforms, and it cannot guarantee that any system will cite or recommend a brand."

ATTRIBUTION NOTE

Model identity supplied by the API context: "an AI assistant accessed via an API." No exact model name or version was exposed in the API-visible context, so a more specific identity cannot be stated reliably. This assessment is model-generated analysis and does not constitute an endorsement, certification, or official position of OpenAI.

Cryptographic verification

DigiCert's RFC 3161 timestamp service signed the SHA-256 digest of the evidence manifest. The timestamp response, signature, manifest binding, and Windows certificate chain were all successfully validated when this receipt was created.

Timestamp authority	DigiCert RFC 3161
Timestamp UTC	2026-07-27T22:03:59.0000000+00:00
Manifest SHA-256	dcaeffaa341481a8dcc8ab2b8db891dc19a56cb4af4bb2df041abc55a877e9e5
Token hash matches	True
CMS signature valid	True
Signature binds manifest	True
Certificate chain valid	True
Signer	CN=DigiCert SHA256 RSA4096 Timestamp Responder 2025 1, O="DigiCert, Inc.", C=US
Timestamp policy OID	2.16.840.1.114412.7.1
Token serial	00a0f6bc07f9eb1ba81ee2a4cf652923b3

Core artifact hashes

api/request.json	7c6376fa82d6b062ed4c65ea3fffe0acc3b1179becd14720658ea28b19dbdd10
api/response.headers.txt	86590e2212e032399c1118927484b2b3f42e4472faef07a2e0b3617df2f457df
api/response.body.json	d0f23f1d40f4a17460a9cbba1eb9a712a08ca09c4f6720540405ba3e243d4378
api/retrieval.body.json	b3c6f7a796edba7d6f13f5db10c34557a07126034c4f2b77fe9f55d5b5a1150b
api/api-call-metadata.json	0c72b4aeadd947886494d3d5e00d5663670c1d228c5be03c0b9e12f1c226cf6c
model-assessment.md	7fd98832cdfcc19500e50b0543d5043d81d65d0277f359df45f97ac446b6e734
proposed-quotation.txt	474db5d096517c3bd857f01958b91411902022727ba000a119b0e78f5cd13b18
evidence-manifest.json	dcaeffaa341481a8dcc8ab2b8db891dc19a56cb4af4bb2df041abc55a877e9e5
original ZIP	adcf870ec0fbf30b14219dec16d535c47d19ed857aba43a11b44520d029a4d9a

How to describe this record

Recommended disclosure: Model-generated technical analysis from the OpenAI Responses API using the returned model identifier gpt-5.6-sol. This does not imply endorsement by OpenAI.

Scope and limitations of this PDF

This PDF is a readable derivative receipt. Its content can be compared against the hashes above, but it does not contain the raw certificate files or full raw JSON payloads. The original provenance ZIP remains the forensic record. A recipient who only needs to read and assess the claim can use this PDF without opening the ZIP.