

AI VISIBILITY SNAPSHOT

How AI engines answer when a seller is deciding who to call.

Client X · clientx.com · October 2026

AI FINDABILITY · READINESS

4.2 / 10

WEAK

How ready you are to be found and quoted.

AI FINDABILITY · OUTCOME

4.6 / 10

WEAK

What AI actually does when asked.

The engines describe you accurately when someone asks about you by name. They do not surface you at all when a seller of a small LP interest asks **who buys them**. That gap is where the deals are.

TWO REAL ANSWERS, CAPTURED THIS WEEK

CHATGPT · OCT 8, 2026 · "What is Client X?"

"Client X, LLC is a **Dallas-based alternative investment firm founded in 2013**. Its specialty is the secondary market for private-equity and other private-fund investments... The firm was founded by **its founder**... A particular niche is smaller transactions that larger secondary-market buyers may overlook."

✓ **Found, and described correctly.**

Branded recall is a genuine strength. The engines even separate you from a similarly named finance company without being asked. This is the half that already works.

CHATGPT · OCT 8, 2026 · "I have a small LP interest in a PE fund to sell. Who buys small LP stakes on the secondary market?"

"These firms buy LP interests directly and are specifically set up to handle smaller positions. For example, **ACE & Company**'s secondary business targets \$250,000 to \$10 million NAV per LP position... **TAP Private Markets**... **The Boring Desk** advertises no minimum transaction size... **RCP Advisors**... family offices..."

✗ **Client X: not mentioned.**

A seller describing your exact niche, small-balance LP secondaries, is handed five other names. ACE is named for a \$250K to \$10M range that is precisely your positioning. You are not in the room where the decision is forming.

THREE PRIORITIES TO GET INTO THAT ANSWER

1 Answer the questions a seller actually types.

When someone asks who buys small LP stakes, the engines pull from pages that answer that question directly. Your one-page site states the strategy but never addresses a seller's real questions: size range, process, timeline, discount to NAV. Publishing that is what moves you from absent to named.

2 Make your proof visible to the machines reading your site.

Your "By the Numbers" proof (year founded, LP stakes acquired, NAV acquired) is injected by JavaScript and renders to AI crawlers as "0 Year Founded, 0+ LP Stakes Acquired, \$0M NAV." The exact credibility figures that would win a seller's trust are invisible at the moment an engine forms its answer. Server-rendering the real numbers fixes it.

3 Make the machines certain exactly which company you are.

The name collides with a similarly named finance company and a major insurance brand, and the site carries no structured identity data to anchor to. A consistent full name with a one-line descriptor, plus basic Organization and FinancialService structured data, removes the ambiguity before it can cost you an answer.

WHAT THE FULL ASSESSMENT ADDS

FROM SNAPSHOT TO THE COMPLETE PICTURE

- **Broader testing.** The same questions across ChatGPT, Google AI Overviews, Gemini, Claude and Perplexity, plus the full set a seller asks (fees, process, timeline, vintage, size), run repeatedly for reliability.
- **Competitor comparison.** A side-by-side showing exactly who the engines recommend instead of you, and the reasons ACE, TAP and others win the answer.
- **Root-cause analysis.** Why each gap exists and which fix moves which answer, so effort goes where it changes what the engines say.
- **A prioritized implementation plan.** The handful of changes, sequenced by impact and effort, that take you from absent to present on the queries that matter most.

Clear from this snapshot: the pattern and the three fixes above. **Still to investigate:** the full multi-engine picture, your counted share of answer against named competitors, and the sequenced plan to close each gap.

Scope of this snapshot: two live queries run on ChatGPT (signed out) on October 8, 2026, one branded and one buyer-intent, alongside a technical review of clientx.com. AI answers vary between runs and engines; the full assessment samples multiple engines and a broader question set, repeated for reliability.

See the full picture, and the plan to fix it.

A complete AI Visibility Assessment for Client X: multi-engine testing, competitor comparison, and a prioritized plan.

Book your full AI Visibility Assessment

or reply to Todd@toddhenderson.com

DEFINING.COM · AI VISIBILITY SNAPSHOT · CLIENT X

Prepared by Todd Henderson · Todd@toddhenderson.com. AI engine responses are the vendors' own output, captured October 8, 2026, and will vary as systems update.