DAKOTA MARKETPLACE · THOUGHT LEADERSHIP

Without Allocation Data, AI Is Guessing.

Every LP/GP database promises to help you find investors. But identifying who an LP is and knowing who they actually invest with are two entirely different things. Only one of them makes AI-powered targeting work. The other is just an expensive directory.

55K+	22K+	13F	60+
PENSION FUND HOLDING RECORDS	PRIVATE TRANSACTION RECORDS	SEC-REPORTED RIA HOLDINGS DATA	DEDICATED DATA PROFESSIONALS

GC

Gui Costin
CEO & Founder, Dakota

The Difference Between a Directory and Intelligence

There are two very different things a fundraising database can tell you about a limited partner. The first is **who they are**: their name, their location, their AUM, their investment focus, their contact information. This is profile data. It is useful. It is the foundation of any LP database worth using.

The second is **what they actually do**: which GPs they have committed to, which strategies they have allocated capital toward, how much they invested, when they invested it, who advised them, and what patterns emerge across their portfolio over time. This is allocation data, transaction data, and holdings data. It is incomparably more valuable — and it is what the overwhelming majority of LP/GP databases do not have.

In the era before AI, that gap was significant but manageable. You could still run a report on profile data and get a reasonable list of prospects. You would miss some. You would waste time on others. But the filter-and-call model could survive on incomplete data because the human on the other end of the call could fill in the gaps through conversation.

That era is over. Because when you put a large language model on top of a database — when you ask an AI tool to help you find your best-fit investors — the AI Layer can only reason over what the data actually contains. And if the data contains only *who* an LP is, and not *what* they have actually invested in, the AI Layer is not delivering intelligence. It is delivering a very sophisticated directory lookup.

> ***Knowing who an LP is tells you they could invest. Knowing what they've actually invested in tells you they will.***

THE FOUR DATA LAYERS

Effective AI-powered LP targeting requires four distinct categories of investment intelligence working together. Each layer adds a dimension of specificity that the others cannot provide alone. Together, they create the conditions for natural language queries to return exact, high-confidence matches — not broad prospect lists, but genuine fit signals.

01 Public Pension Fund Holdings

DOCUMENTED ALLOCATIONS FROM THE MOST TRANSPARENT LP SEGMENT

Public pension funds are among the most important allocators in private markets — and they are also among the most transparent, by legal obligation. Over the past seven years, Dakota has systematically collected and structured over 55,000 individual holdings from public pension funds across the country.

Each holding record captures the name of the LP, the name of the GP, the fund consultant who advised the allocation, the amount committed, and the date. This is not a summary. This is the actual investment record — every allocation, named and quantified, tied to the specific LP that made it and the specific GP that received it.

WHAT THIS ENABLES · NATURAL LANGUAGE QUERY

> *"Find me public pension funds that have invested in lower-middle-market buyout funds with committed capital between $50M and $150M in the last four years, advised by one of the top five consultants in that space."*
>
> Without 55,000 structured holdings records, that query returns nothing meaningful. With them, it returns a precise, ranked list of LPs who have demonstrated exactly that behavior — with the specific consultant relationships identified. No filter-building. No guesswork. Exact matches.

02 13F Filings and RIA Holdings

SEC-REPORTED PROOF OF PRACTICE FOR EVERY REGISTERED INVESTMENT ADVISOR

Registered Investment Advisors managing more than $100 million in public equity are required to disclose their holdings quarterly through 13F filings with the SEC. Dakota has structured this data into a searchable intelligence layer that reveals exactly which RIAs hold which public securities — and in what quantities.

For fundraisers working in strategies adjacent to public markets — hedge funds, long/short equity, liquid alternatives, multi-asset — this data is transformative. It answers the question that a profile database cannot: *does this RIA actually put money into what I'm raising?* The 13F record is the proof of practice. It is not what an RIA says they do. It is what they have actually done, reported under penalty of law, traceable by date and position size.

When the AI Layer reasons over this data, it can identify not just which RIAs have holdings in adjacent strategies, but which have been increasing exposure, which have been reducing it, and which have portfolio construction patterns that indicate natural fit with a specific fund type. That is not a report. That is an analyst working at machine speed.

03 Private Company Transactions

LP-ATTRIBUTED INVESTMENT RECORDS THAT CANNOT BE FOUND ANYWHERE ELSE

Dakota has compiled over 22,000 private company transactions and investments — each one attributed to the LP or investor that made it. This is, in many respects, the most powerful layer of all, because private market investment activity is precisely what cannot be found through public disclosure.

When a pension fund commits to a private equity fund, that is documented in the holdings data. But when an LP makes a direct co-investment, a secondary purchase, or a direct private company investment, that transaction often lives in no structured public record. Dakota's research team has spent years building this dataset from primary sources — and it now contains a picture of real private market investment behavior that exists nowhere else.

The implications for AI-powered targeting are direct: a fundraiser can now ask, in plain language, which LPs have made direct investments in healthcare technology companies at the growth equity stage — and get an answer backed by documented transaction history, not inference from profile descriptions.

04 Family Office Investment Intelligence

ACTUAL INVESTMENT ACTIVITY FROM THE MOST CONSEQUENTIAL PRIVATE CAPITAL POOL

Family offices represent one of the most consequential and least transparent pools of capital in private markets. They are significant allocators to private equity, venture, real assets, and credit strategies — yet because they have no public disclosure requirements, their investment activity has historically been almost impossible to track systematically. Most databases have a list of family office names and a few profile fields. That is not intelligence. That is a phone book.

Dakota has built what is, by any measure, one of the most comprehensive family office investment databases in the industry. This is not a collection of self-reported preferences or website-scraped descriptions. It is a structured record of actual family office investment activity — which GPs they have backed, which strategies they have allocated to, and which funds they have committed to over time. The dataset spans single-family offices, multi-family offices, and family office platforms, with investment records that reveal behavioral patterns no profile database can approximate.

For a GP raising capital from family offices — a channel that has grown dramatically as family offices have professionalized their investment programs — this data layer is the difference between cold outreach to a long list and targeted engagement with a short list of offices that have already demonstrated they invest in exactly what you're raising. When the AI Layer reasons over Dakota's family office investment records, it can answer questions that have never been answerable before: which family offices have backed first-time funds in your strategy, which have written checks at your target size, and which have co-invested alongside institutional LPs you already have in

your fund.

55,000+	13F	22,000+	★
PENSION FUND HOLDING RECORDS	RIA HOLDINGS DATA	PRIVATE TRANSACTIONS	FAMILY OFFICE INVESTMENTS

WHAT NATURAL LANGUAGE TARGETING ACTUALLY REQUIRES

The investment industry has enthusiastically embraced the concept of AI-powered search. Vendors across the landscape have rushed to put AI interfaces on top of their existing databases and describe the result as intelligent targeting. In most cases, it is not.

What those systems have done is give users a more comfortable way to query the same insufficient data they always had. Asking an AI to “find LPs who invest in private credit” against a database that contains only LP profiles and stated preferences is not AI-powered targeting. **It is autocomplete with branding.**

The reason is simple and worth stating plainly: **a language model is a reasoning engine, not a data source.** It is extraordinarily good at understanding what you mean, inferring what you need, and synthesizing patterns across large, complex datasets. But it cannot reason over data that does not exist. If the underlying database has no record of what an LP has actually invested in, the AI Layer cannot manufacture that record. It can only return what it has.

THE DEFINITIVE STATEMENT

Any fundraising database that does not include structured investment history, allocation data, and transaction records cannot leverage AI to deliver meaningful LP targeting. It can deliver a more conversational interface to the same incomplete results it always provided. The data is not the wrapper. The data is everything.

The Matching Problem: Why Exact Fit Requires Exact Evidence

Consider what a GP actually needs from an LP targeting system. They are not looking for a list of LPs who broadly invest in their asset class. Every LP/GP database can produce that. What they need is a ranked, specific list of LPs whose demonstrated investment behavior — strategy, stage, check size, GP relationships, consultant coverage — most closely matches what they are raising.

That match is only possible if the database contains evidence of actual investment behavior. Not stated preferences. Not survey responses. Not self-reported allocation targets. Actual documented investments, with names, amounts, and dates.

WITHOUT ALLOCATION DATA

A healthcare-focused growth equity fund can identify LPs who say they invest in healthcare growth equity. This is a list of approximately 800 LPs who have checked a box. It requires a skilled analyst to narrow down from there using judgment and relationship knowledge.

WITH DAKOTA'S DATA

The same fund can identify LPs who have actually committed capital to healthcare growth equity GPs in the last five years, at check sizes consistent with the fund's target, advised by consultants with track records in that strategy — pulled from 55,000 pension holdings, 22,000 transaction records, and 13F filings in a single natural language query.

THE DIFFERENCE

One list requires weeks of research to validate. The other is the validated research. The AI Layer does not help you narrow 800 names down. It delivers 23 high-confidence exact matches, with documented evidence behind each one.

The Data Depth Equation

The popular conversation about AI in the investment industry has focused almost entirely on the model: which model is most capable, which interface is most intuitive, which vendor has integrated which frontier AI. This is the wrong conversation.

The model is a commodity. The frontier AI models are accessible to any firm with a budget and an API key. What is not a commodity — what cannot be purchased, scraped, or replicated on a competitive timeline — is the underlying dataset. The depth, accuracy, and currency of the investment intelligence that the model reasons over is the only sustainable source of differentiation in AI-powered fundraising tools.

Dakota's dataset represents nearly three decades of practitioner-built intelligence, maintained today by over 60 dedicated data professionals whose sole function is enriching, updating, and expanding the coverage of every

record in the platform. More than one million lines of proprietary data span LP profiles, GP fund histories, public pension holdings, 13F filings, private transaction records, RIA intelligence, and financial advisor networks — all integrated into a single system that a fundraiser can query in plain English.

There is no comparable integrated dataset in the investment industry. The combination of LP profile intelligence, documented allocation history, 13F holdings, and private transaction records — all in one platform, all maintained by a dedicated research team, all queryable through natural language — does not exist anywhere else. Competitors may have pieces. None have the whole. And in AI-powered targeting, the whole is worth categorically more than the sum of the parts.

The Question Every GP Should Be Asking Their Data Provider

If you are evaluating any LP database, LP/GP intelligence platform, or AI-powered fundraising tool, there is a set of questions that will immediately reveal whether the product can deliver on its promises — or whether it is offering a better interface to the same old limitations.

→ *Do you have documented investment history for the LPs in your database — not stated preferences, but actual recorded allocations with GP names, amounts, and dates?*

→ *Do you have 13F holdings data that shows us what RIAs actually own, not just what they describe as their investment approach?*

→ *Do you have private company transaction records that capture LP investment behavior in private markets, including co-investments and direct deals?*

→ *How many data professionals are actively maintaining and enriching your dataset right now?*

The answers to those questions will tell you everything. A database without allocation data, transaction history, and active human maintenance is a database that cannot deliver meaningful AI-powered targeting regardless of which model it uses. **It is a directory wearing a lab coat.**

Dakota Marketplace: Built for This Moment

Dakota did not add investment and allocation data to an existing LP directory. We built the platform from the ground up as a practitioner intelligence system — because the founders knew from direct fundraising experience that profile data was never sufficient for serious targeting work. The allocation data, the transaction history, the holdings intelligence: these were not features added after the fact. They are the foundation on which everything else was built.

That foundation is what makes the Dakota AI-powered app capable of doing something that no other platform in private markets can do: returning exact, high-confidence LP matches based on documented investment behavior, delivered through a natural language interface that any team member can use without training.

THE CORE ARGUMENT

AI targeting is only as precise as the allocation data underneath it. A database with LP profiles but no investment history cannot tell an AI which LPs actually invest in your strategy. It can only tell it who might.

THE FOUR DATA LAYERS

01 Public pension fund holdings — LP, GP, consultant, amount
02 13F filings — RIA holdings, documented quarterly
03 Private company transactions — LP-attributed investment records
04 Family office investments — actual GP commitments and strategy allocations

DAKOTA DATA SCALE

55K+	22K+	1M+	60+	7 YRS	30 YRS
PENSION FUND HOLDING RECORDS	PRIVATE TRANSACTIONS	LINES OF PROPRIETARY DATA	DEDICATED DATA PROFESSIONALS	ALLOCATION HISTORY	PLATFORM INTELLIGENCE

THE VERDICT

> ***Without investment history, transaction data, and allocation records, no AI model can deliver exact LP targeting. It can only deliver a more polished version of the same insufficient results.***
