Best Private Company Data Providers for Private Equity

Docket · Docket Research · 2026-09-18 · private company data providers · private equity database · pitchbook alternatives · deal sourcing · private market data · sourcescrub · grata · privco · docket

Executive Summary

The “best” private-company data provider for private equity is mandate-specific, not the one with the largest advertised universe. **PitchBook** says it tracks all private equity-backed companies, yet also says non-backed private companies are not comprehensively tracked in every region (pitchbook.com) (pitchbook.com). **SourceScrub** reports 17 million listings connected with 290,000 sources and lists, while **Grata** reports more than 22 million companies and 150 data points (sourcescrub.com) (grata.com). These are vendor-stated scale claims, not comparable measurements of usable lower middle market targets.

The shortlist should reflect the mandate. PitchBook supplies private-capital context and integrations. SourceScrub emphasizes source-led discovery. Grata supports keyword search. **PrivCo** focuses on established U.S. private companies and says material data points are corroborated across sources (privco.com) (privco.com). S&P Capital IQ Pro, Crunchbase, Tracxn, Dealroom, AlphaSense with Tegus, Gain, CB Insights, and Moody’s Orbis provide further combinations of breadth, financials, technology-market intelligence, documents, entity structure, and delivery. **Docket**, a direct provider with a different model, returns target-specific answers with sources, excerpts, and collection dates (docket.capital).

Select among them with a blinded field test. Freeze geography, entity type, ownership, size, sector, exclusions, evidence rules, and refresh tolerance. Test known positives, difficult positives, negatives, recent changes, aliases, acquired brands, duplicate traps, and deliberately unanswerable fields. Normalize legal entities, brands, facilities, contacts, and events before scoring. Measure **candidate precision**, **sample recall**, duplicate rate, ownership accuracy, evidence traceability, staleness, unknown handling, export usability, and analyst minutes. NIST defines precision and recall through true positives, false positives, and false negatives (nvlpubs.nist.gov) (nvlpubs.nist.gov). Report the latter as sample recall, never total-market coverage.

Use pass/fail disqualifiers before a preweighted score, then test whether reasonable weight changes alter the result. The economic denominator is accepted output. In the report's clearly fictional example, a **$42,700** combined license, review, and cleanup cost divided by **420 accepted companies** equals **$101.67 per accepted company**. At 300 accepted companies the figure is **$142.33**; at 550 it is **$77.64**. A complementary two-source stack can win if its incremental verified yield exceeds its added license and cleanup cost. Review current contracts before benchmarking or publishing results, and re-test volatile fields after purchase.

Introduction and Background

A private equity database is not intrinsically "best." It is best only for a defined acquisition mandate and a measurable workflow. A lower middle market buyout team seeking non-backed, founder-owned industrial service companies has a different coverage problem from a growth investor tracking venture rounds. That distinction matters because **PitchBook** says it tracks all private equity-backed companies, while also stating that it does not comprehensively track non-backed private companies in every region (pitchbook.com).

Private-company research is structurally fragmented. Unlike public issuers, private companies generally are not required to file with the U.S. Securities and Exchange Commission, and the Library of Congress says the work often requires "considerable creativity and patience" (guides.loc.gov) (guides.loc.gov). Profile totals therefore do not answer the buying question. The decisive unit is an in-scope, unique company with fields that an analyst can accept under the mandate's evidence rules.

This report compares documented capabilities, not unobserved accuracy. No controlled, like-for-like public benchmark found in this research establishes a universal winner. Vendor profile counts, update claims, and feature descriptions are treated as first-party claims. The recommendation is a **blinded, stratified pilot** that uses identical queries, cutoff dates, definitions, and adjudication rules. The output is a defensible decision for one mandate, not a league table.

Accordingly, this private company data providers comparison treats private equity deal sourcing databases, private market data platforms, PE deal origination software, private company intelligence tools, and private equity target research tools as overlapping categories. It answers which are the best databases for private equity firms and which PitchBook alternatives for private equity deserve a mandate-specific trial.

PitchBook

Capabilities

PitchBook is the broad private-capital reference option in this shortlist. Its documented scope includes all companies backed by private equity firms (pitchbook.com). It also offers an application programming interface, or **API**, that calls selected PitchBook data into internal systems (pitchbook.com). Documented customer relationship management, or **CRM**, options include Salesforce, Microsoft Dynamics, HubSpot, DealCloud, and Affinity (pitchbook.com).

Adoption and mandate fit

The documented scope makes PitchBook a natural control source when the mandate emphasizes sponsored companies, investors, funds, and transactions. It is less safe to assume that the same strength transfers to every non-backed regional company. The vendor's own limitation is unusually useful in procurement: "non-backed private company" recall should be a separate pilot stratum, not averaged into backed-company performance (pitchbook.com).

Strengths and limitations

The practical strength is workflow breadth around private capital. The practical limitation is that a large platform universe and mandate recall are different quantities. Buyers should also review contract language before publishing benchmark results. PitchBook's public terms say covered work product may not be published to more than 500 people without prior written consent (pitchbook.com). That clause does not decide what a negotiated order form permits, but it is enough to make legal review a precondition to any public comparison.

SourceScrub

Capabilities

SourceScrub represents a source-led origination model. It reports **17 million company listings** connected with **290,000 sources and lists**, figures that should be understood as vendor-stated scope rather than verified unique legal entities (sourcescrub.com). Its differentiating claim is that companies link back to their original sources (sourcescrub.com). SourceScrub also describes Data Connect as access to an always-growing, continuously refreshed database (sourcescrub.com).

Adoption and mandate fit

This model is relevant when a thesis can be expressed through discoverable source lists, such as conference exhibitors, association membership, awards, or trade directories. In a trial, the team should preserve the underlying source type and source date. A source-linked hit is not automatically an accepted company, but it lets an adjudicator evaluate why the candidate entered the funnel.

Strengths and limitations

Source traceability can reduce the gap between discovery and review. It does not prove ownership, revenue, or current operating status. Those fields require their own evidence. SourceScrub documents a two-way CRM integration workflow (live.sourcescrub.com), but the pilot still needs to test whether field names, identifiers, and update behavior survive export. Missing public trial terms should be resolved in the commercial documents, along with permitted uses of exported data and benchmark publication.

Grata

Capabilities

Grata is positioned around searchable private-company intelligence. Its API page states **150 data points across more than 22 million global companies** and describes keyword search against an investment thesis (grata.com) (grata.com). The same official page lists descriptions, revenue, headcount, ownership, funding, and industry classifications among its record fields (grata.com).

Adoption and mandate fit

Keyword search can be useful when standard industry codes are too coarse for a niche thesis. A buyer should nevertheless test query stability: run the same concept through a documented synonym set, record the query, and export the full result before adjudication. The trial should ask whether the result identifies a legal entity, an operating brand, or a website, because those objects may collapse into one profile.

Strengths and limitations

Grata documents pushing companies, notes, and tags into Salesforce, DealCloud, HubSpot, or Affinity (grata.com). That can support origination workflow, but integration availability does not establish data quality. The pilot should measure exported identifiers, ownership evidence, null semantics, and duplicates. Vendor-stated scale should be reported separately from in-mandate unique companies, sourced answers, and verified answers.

PrivCo

Capabilities

PrivCo is the focused private-financial option in this shortlist. Its official site describes coverage of established U.S. private companies (privco.com). Its methodology says every material data point is corroborated across sources before reaching a profile (privco.com). For downstream use, the company documents bulk extraction through JSON or CSV (privco.com).

Adoption and mandate fit

This scope is most relevant when the mandate is U.S.-centered and private-company financial fields carry substantial decision weight. A buyer should not treat corroboration as a substitute for field-level inspection. The pilot answer key should require the source, period, unit, currency, and whether a value is reported or estimated. An accepted headquarters address and an accepted revenue value are different answers and should earn separate scores.

Strengths and limitations

PrivCo publishes unusually specific List Builder trial limits: **five company profiles, 100 AI searches, and 125 contact views** (privco.com). Those limits may be too small for a statistically useful multi-stratum pilot. Procurement should request a controlled evaluation allowance large enough to include known positives, hard positives, negatives, aliases, and recent changes. If that is unavailable, the firm can still test record quality, but it should not infer broad recall from five profiles.

Docket

Capabilities

Docket is a direct provider of sourced target research rather than a general-purpose database. It says it runs Triage, Scout, and Audit agents across acquisition targets and returns a structured file on each company (docket.capital). The platform records supporting excerpts and collection dates, while unanswered questions retain the sources checked (docket.capital). Its public workflow also says unsupported findings are removed from accepted answers during a separate review pass (docket.capital).

Adoption and mandate fit

The model fits a team that already has a list, wants one built around a thesis, or needs evidence for fields that databases leave ambiguous. It does not make a public claim to a universal prebuilt company count, so it should be judged on accepted answers, candidate yield, latency, and review cost. Export usability should be tested in the same way as for database providers.

Strengths and limitations

The evidence model matches a field-level acceptance test, particularly for ownership, niche fit, and unresolved conflicts. The limitation is comparability: a researched answer is a different unit from a database profile. The pilot must allocate Docket’s research charge and review time to the same denominator used for competitors. It must also prevent a richer research process from receiving credit merely for writing more text. Only pre-specified accepted fields and accepted unique companies count.

Feature Comparison

No table can establish which provider has the highest accuracy without running the same cases through each product. It can, however, map documented operating models to mandates. *Table 1* distinguishes the primary shortlist and several credible alternatives using official pages available as of **September 18, 2026**.

Provider	Documented operating model	Potential mandate fit	Pilot question that matters
PitchBook	Private-capital platform with API and named CRM integrations (pitchbook.com)	Sponsored-company, investor, fund, and transaction context	What is sample recall for non-backed companies in the exact geography and size band?
SourceScrub	Source-led company discovery with links back to origin (sourcescrub.com)	Themes visible through conferences, directories, and industry lists	Does the origin source support current fit, and how many listings collapse into unique companies?
Grata	Keyword search and private-company fields exposed through an API (grata.com)	Niche theses not captured well by standard codes	Are concept-search results stable, precise, and exportable with identifiers?
PrivCo	U.S. private-company financial intelligence with JSON and CSV extraction (privco.com)	U.S. mandates where private financial fields are central	Which financial fields are reported, estimated, dated, sourced, and accepted?
Docket	Target-specific sourced research with separate evidence review (docket.capital)	Mandates requiring field-level evidence and explicit unknowns	What is cost per accepted answer after research and analyst review?
S&P Capital IQ Pro	Broad platform claiming more than 60 million private companies and 110,000 private equity and venture capital funds (spglobal.com)	Global teams combining company, market, and fund data	How much of the broad universe survives the mandate and evidence gates?
Crunchbase	Emerging and established organization coverage with platform and API access (about.crunchbase.com)	Growth, funding, and ecosystem discovery	Does coverage extend from venture-visible companies into the mandate's mature, non-backed segment?
Tracxn	Company database with funding, capitalization, financial, and competitor data (help.tracxn.com)	Startup, technology, and thematic landscape work	Are statutory-field updates timely enough for the target jurisdictions?
Dealroom	Startup, growth-company, and technology-ecosystem coverage (knowledge.dealroom.co)	European and global technology ecosystems	Does the mandate match the platform's startup and growth orientation?
AlphaSense with Tegus	Document and expert-insight workflow spanning origination through portfolio monitoring (alpha-sense.com)	Thesis research and document intelligence around targets	Do source-traceable documents answer the required company fields consistently?
Gain	Private-market graph with filings, disclosures, analyst review, CRM links, and bulk data (gain.ai)	European private-market mapping and workflow integration	Does claimed coverage translate into accepted unique companies in the precise region and size band?

Provider	Documented operating model	Potential mandate fit	Pilot question that matters
CB Insights	Private-market monitoring delivered through APIs, Snowflake, CRM, and AI connectors (cbinsights.com)	Technology-market and growth-company monitoring	Are signals reproducible and relevant to the investment committee's required fields?
Moody's Orbis	Private-company reference data with financials, deal histories, corporate structures, APIs, and connectors (moodys.com)	Legal-entity, ownership, and cross-border corporate research	Can the export preserve entity relationships and source dates at usable granularity?

The interpretation is deliberately conditional. **S&P Global** claims more than 200 datasets available through API, cloud, or Xpressfeed (spglobal.com). **Crunchbase** says it produces more than 30 million verified updates annually (about.crunchbase.com). **Moody's** says more than 99 percent of Orbis-covered companies are private (moodys.com). These facts describe scale and orientation. They do not supply mandate precision, recall, field accuracy, or cost per accepted company.

Performance and Benchmarks

Start with a written mandate

Before opening any trial, freeze a one-page specification. At minimum it should define:

- **Geography:** headquarters, incorporation, operating footprint, and acceptable border cases.
- **Company object:** legal entity, consolidated firm, operating brand, website, facility, or franchise location.
- **Company type:** independent, founder-owned, family-owned, sponsor-backed, corporate subsidiary, or carve-out.
- **Size:** revenue, earnings, employees, sites, or another measurable proxy, with period and currency.
- **Sector:** included activities, excluded activities, and evidence needed for a niche classification.
- **Ownership:** accepted ownership states, controlling-interest rule, and "unknown" treatment.
- **Evidence:** acceptable source types, source date, collection date, excerpt, and conflict policy.
- **Refresh tolerance:** maximum age by field, because ownership and address may need different limits.
- **Exclusions:** closed companies, acquired brands, distributors, nonprofits, and irrelevant geographies.
- **Deliverability:** export schema, API access, CRM mapping, rate limits, and record-level provenance.

The entity definition is not clerical. The U.S. Census Bureau defines an establishment as a fixed physical location where economic activity occurs, while it groups establishments under common operational control into a firm (census.gov) (census.gov). A provider returning ten branches and another returning one parent cannot be compared until those records are normalized.

Build a stratified test set and hidden answer key

Create a sample large enough to expose failure modes, not merely showcase obvious successes. Stratification means segmenting a population across levels of a factor (itl.nist.gov). A practical pilot can allocate cases across:

- **Known positives:** unambiguous in-scope companies with verified fields.
- **Difficult positives:** small, low-web-presence, non-backed, or multi-entity companies.

- **Known negatives:** similar-looking companies that fail one decisive criterion.
- **Recent changes:** acquisitions, name changes, new headquarters, or ownership transitions.
- **Aliases and DBAs:** “doing business as” names linked to the same legal entity.
- **Acquired brands:** brands that still have websites but are no longer independent targets.
- **Duplicate traps:** parent, subsidiary, brand, domain, and location variants.
- **Deliberately unanswerable fields:** cases where the correct output is unknown, not a guess.

Keep part of the answer key hidden from trial operators. Have at least two adjudicators independently review a subset, then report their agreement and reconciliation process. Reliability guidance specifically calls for reporting blinding and information availability, and for stating whether ratings were conducted independently (pmc.ncbi.nlm.nih.gov) (pmc.ncbi.nlm.nih.gov).

Normalize before scoring

Every export should be transformed to a common schema before anyone sees the provider name. Use a stable internal test-case identifier and retain the raw value separately from the normalized value. A Legal Entity Identifier, where one exists, is useful because each **LEI** is unique to one entity (gleif.org). It is not universal, so the crosswalk should also use legal name, jurisdiction, registration number, domain, address, and parent.

Record provenance as structured data, not a comment. The World Wide Web Consortium defines provenance as information about the entities, activities, and people involved in producing data (w3.org). Store both the source date and collection date. Do not convert “not available,” “not applicable,” “not requested,” and “not required” into one blank value; Eurostat guidance explicitly distinguishes those states (cros.ec.europa.eu).

Prevent a rigged trial

Use the same query concepts, filters, cutoff date, and export deadline for all providers. Do not accept a vendor-supplied showcase list as the test universe. Prohibit outreach to sample companies. Blind vendor identity during adjudication where presentation does not make that impossible. Count missing as missing, not automatically false. Census linkage standards support setting scoring weights and cutoffs and evaluating automated links through manual review (census.gov) (census.gov).

Finally, confirm in writing that the contract permits the intended test, exports, internal retention, and any publication. A product trial is access, not blanket permission to disclose comparative results.

Data Analysis and Evidence

The scorecard should keep funnel quantities separate. *Table 2* shows the required levels and prevents a large headline universe from being mistaken for verified coverage.

Quantity	Definition	Why it cannot substitute for the next level
Profiles in database	Vendor-stated records before mandate filters	May mix legal entities, brands, locations, public firms, and stale records.

Quantity	Definition	Why it cannot substitute for the next level
Profiles in mandate	Returned records satisfying machine-applied mandate filters	Filters may be missing, estimated, or based on different definitions.
Usable records	Unique normalized companies with required identifiers	A usable record can still contain unsupported field values.
Sourced answers	Field answers with an inspectable source and date	A source can be stale, ambiguous, or insufficient for the claimed value.
Verified answers	Sourced answers accepted by blinded adjudicators	Acceptance on sampled fields does not establish total-market completeness.

This hierarchy changes the procurement conversation. A claim such as Dealroom's **more than three million companies and more than 100 data points** belongs at the first level until the test shows what survives the mandate (knowledge.dealroom.co). Tracxn's stated **more than seven million companies** likewise describes its universe, not a lower middle market recall rate (help.tracxn.com).

Table 3 is a copy-ready field-level scorecard specification. Each row represents one vendor answer for one test case. A spreadsheet can add formulas and pivot tables without changing the raw observations.

Column	Required content	Scoring use
test_case_id	Internal blinded identifier	Joins output to the hidden answer key.
stratum	Known positive, hard positive, negative, recent change, alias, acquired brand, unanswerable	Exposes segment-specific performance.
expected_status	Positive, negative, unknown, or hidden	Prevents forced binary treatment.
vendor_value_raw	Exact exported value	Preserves what the provider returned.
normalized_value	Value after common entity and field rules	Supports like-for-like comparison.
source_url	Direct evidence URL, if supplied	Enables traceability scoring.
source_date	Date the underlying source describes	Supports staleness measurement.
collection_date	Date the provider or tester collected it	Separates collection from source age.
adjudication	Accepted, rejected, missing, ambiguous, or not applicable	Supplies field precision and yield.
freshness_days	Cutoff date minus applicable source or event date	Tests the field-specific refresh tolerance.
duplicate_cluster	Canonical company identifier	Measures duplicate burden.
analyst_minutes	Review and cleanup minutes	Feeds total-cost calculations.

Compute **candidate precision** as in-scope unique returned companies divided by all returned unique companies. Compute **sample recall** as known in-scope companies found divided by all known in-scope companies in the test set. NIST expresses the standard formulas as precision equal to true positives divided by true positives plus false positives, and recall equal to true positives divided by true positives plus false negatives (nvlpubs.nist.gov) (nvlpubs.nist.gov). Call the second measure sample recall, never total-market coverage, because the denominator is a constructed sample.

Use a weighted score only after pass/fail gates. Federal acquisition guidance says evaluation factors should be tailored to the acquisition and their relative importance disclosed in advance (acquisition.gov) (acquisition.gov). A suitable rule is:

Weighted score = sum of each metric score multiplied by its precommitted weight.

Apply separate disqualifiers first. For example, unsupported ownership on more than a chosen tolerance, inability to export source URLs, or impermissible license terms can fail a provider even if its total score is high. Government multi-criteria guidance similarly describes applying percentage weights and summing scores, then using sensitivity analysis to see whether assumptions change the ordering (gov.uk) (gov.uk).

Fictional cost example

The following example is fictional and does not represent any vendor result. Suppose a firm allocates **$24,000** of annual license cost to one mandate, spends **180 analyst hours** reviewing candidates at **$85 per hour**, and incurs **40 hours** of duplicate cleanup at the same rate. If **420 unique companies** are accepted, the calculation is:

Cost per accepted company = ($24,000 + 180 × $85 + 40 × $85) / 420 = $101.67.

At **300 accepted companies**, the same cost base yields **$142.33** per accepted company. At **550**, it yields **$77.64**. This sensitivity band is more decision-useful than a license quote alone because life-cycle costing includes costs incurred throughout a product or service's lifetime, not just purchase price (green-forum.ec.europa.eu). A parallel **cost per accepted answer** can expose providers that find many candidates but leave decisive fields unresolved.

Implications and Future Directions

The likely outcome is not always one database. Two complementary sources can beat one broad platform when they fail differently. A source-led discovery product might supply candidates, while an entity or research service verifies ownership and niche fit. The correct comparison is the combined cost and accepted yield, including deduplication between sources.

Refresh should also be field-specific. Tracxn says financials update according to the availability of statutory filings and disclosures (help.tracxn.com). Dealroom says its collection uses sources such as news, filings, domain and trade registries, and job boards (knowledge.dealroom.co). These inputs have different event and publication lags. "Updated yesterday" can mean the profile was processed yesterday, not that each underlying fact is current.

After purchase, re-run a smaller holdout test quarterly for volatile fields and at renewal for full economics. Keep a change log containing:

- **Schema changes:** renamed, added, removed, or redefined fields.

- **Coverage changes:** new geographies, source classes, and company types.
- **Workflow changes:** export limits, API behavior, CRM mappings, and rate limits.
- **Evidence changes:** source URLs, source dates, retained excerpts, and conflict handling.
- **Quality changes:** precision, sample recall, duplicate rate, unknown rate, and freshness by stratum.
- **Cost changes:** license allocation, analyst minutes, cleanup time, and accepted output.

Do not label engagement, analyst time saved, or outreach volume as conversion or deal outcome. The data provider contributes to a pipeline, but acquisition outcomes depend on pricing, owner willingness, competition, diligence, financing, and investment judgment. Procurement claims should stop at what the trial can observe.

Frequently Asked Questions (FAQs)

What is the best private-company database for private equity?

There is no evidence-backed universal winner. **PitchBook** is a strong control for private-capital context; **SourceScrub** for source-led discovery; **Grata** for keyword-driven private-company search; **PrivCo** for established U.S. private-company financials; **Orbis** for entity and corporate-structure work; and **Docket** for sourced, mandate-specific target research. The best choice is the one that passes the firm's disqualifiers and produces the lowest acceptable cost per verified company or answer on a blinded sample.

What are the main PitchBook alternatives for PE deal sourcing?

Alternatives differ by job, not merely feature count. SourceScrub, Grata, Gain, Crunchbase, Tracxn, Dealroom, CB Insights, S&P Capital IQ Pro, AlphaSense with Tegus, Moody's Orbis, PrivCo, and Docket cover different combinations of discovery, financials, entity structure, documents, integrations, and research. For example, AlphaSense reports significant private-company coverage across **240,000 Tegus expert-insight transcripts**, which is a different asset from a company master database (alpha-sense.com).

How large should a vendor trial be?

There is no defensible universal number. Size follows the number of strata, the error rates the firm needs to detect, and available adjudication capacity. A five-profile trial may reveal schema and evidence problems but cannot support broad coverage conclusions. Include enough cases in every decision-critical stratum to prevent obvious, venture-visible companies from dominating the result.

How should unknown values be scored?

Unknown should be a distinct, often acceptable state. A provider should not be penalized as wrong when the answer is genuinely unavailable, but it should be penalized if the workflow requires that field and the source cannot supply it. Score missingness, correctness, and unsupported assertion separately. This preserves the difference between "no evidence found" and "evidence proves false."

Should a PE firm buy one platform or combine two?

Combine sources when their incremental accepted yield exceeds added license, integration, review, and cleanup cost. The test should run source A, source B, and the deduplicated union. Measure how many accepted companies and answers B adds after A, not B’s standalone profile total.

Can vendor-reported coverage counts be compared directly?

Usually not. Counts may refer to profiles, legal entities, brands, establishments, investors, or a mixture. S&P Global positions funding rounds, revenue growth, and employee headcount as sourcing fields (spglobal.com), while Crunchbase describes a source mix of participant input, engagement signals, ingestion systems, external sources, and expert validation (about.crunchbase.com). Normalize the returned objects before comparing coverage.

Conclusion

The best private-company data provider for a PE team is the provider that performs best on that team’s written mandate after entity normalization, blinded adjudication, evidence review, and full-cost accounting. Public pages can establish scope, features, integrations, and vendor-stated scale. They cannot establish paid-product accuracy for a particular geography, sector, ownership type, or size band.

A defensible buying process therefore begins with definitions, not demos. It uses a stratified set of known positives, difficult positives, negatives, recent changes, aliases, acquired brands, duplicates, and unanswerable fields. It measures candidate precision, sample recall, duplicate burden, ownership accuracy, traceability, freshness, unknown handling, analyst minutes, and export usability. It precommits weights, applies disqualifiers before totals, and tests whether rankings survive reasonable changes in assumptions.

The final commercial measure is cost per accepted unique company or accepted field answer, including allocated license cost, review labor, and cleanup. That measure permits fair comparison between broad databases, source-led discovery tools, document platforms, entity datasets, and researched-answer services. Re-test volatile fields during the term and repeat the complete protocol at renewal. The result will not be a universal vendor ranking, but it will be a reproducible procurement decision grounded in the evidence the investment team actually uses.
