




Global Artificial Intelligence (AI)

Market Report · July 2026

Now, for tomorrow



Executive Summary

The lower market, AI companies in the sub-\$100m enterprise-value band financed from pre-seed through small Series B rounds broadly under \$50m, is large in company count, thin in disclosed revenue, and capital-light. No provider sizes it as a category, so it must be triangulated from population, capital flow and round data. The decisive variable is not headline market size but bifurcation: defensible AI-native assets separate sharply from a disposable wrapper tail, and that split determines whether a target is a bolt-on at 4x to 6x revenue or an acqui-hire priced on headcount.

\$3–4.6m MEDIAN AI SEED ROUND About \$1m above the broader market; a 25–30% AI valuation premium sits on top.	~48% AI-NATIVE MEDIAN NRR Versus 32% for sub-\$50 plans and 85% above \$250. Retention is the valuation line.	+127% AI-TARGET DEAL VALUE, H1 2025 Up year on year even as deal volume fell about 20%. Capability, not volume.	~90% THIN WRAPPERS DEAD BY END-2026 One builder estimate; the wrapper tail is where most of the landmines sit.
68% SMB ADOPTION, 10–100 STAFF Up from 47% in a year; SMBs are the demand engine of the lower market.	~80% LLM API PRICE FALL IN A YEAR Every cut compresses wrapper margin while leaving vertical-agent economics intact.	\$220m MONGODB–VOYAGE AI The native-data bolt-on template: a 17-month-old company bought to own retrieval.	15–30x AI-NATIVE REVENUE MULTIPLE Against 3–5x for AI-threatened horizontal tools. The dispersion is the story.

Scope & Definition

This report covers the global artificial-intelligence sector as of mid-2026, written for the lower-market M&A practitioner hunting small AI companies in the sub-\$100m enterprise-value band. Its purpose is not to settle on a headline market figure but to locate the slice a small-cap buyer can actually own, and to make the AI-versus-wrapper distinction operational as the central deal-quality test. The substitution test frames every screen: could a technical user reproduce most of the output by pasting the prompt into a frontier model?

In scope

Vertical AI agents and AI-native vertical SaaS, horizontal applied-AI, MLOps, evaluation and vector tooling at the small-company edge, data labeling and curation, AI development services, and the governance and orchestration tools that attach to AI spend.

Adjacent / partial scope

Open-weight model ecosystems, neoclouds and the regulatory regime (EU AI Act) are analysed where they set small-company economics or create an acquisition rationale, not as target sets in themselves.

Out of scope

Frontier model labs, hyperscaler capex and the silicon and packaging tiers. These are economically decisive but unacquirable at sub-\$100m EV, and set the cost and margin floor below which small companies operate.

How the lower market is triangulated

Company population. Tracxn counts tens of thousands of AI companies, the great majority small and founder-led.

Capital flow. AI took roughly 80% of venture funding, but the dollars concentrate at the very top.

Disclosed revenue. Thin, because most targets are early and capital-light. The result is a high-population, low-disclosed-revenue cohort.

Analytical framework

Segmented by position in the stack, crossed with two decisive axes: model or infrastructure-native versus wrapper-heavy, and buying behavior.

The substitution test is the first screen on every candidate.

Nine analytical sections follow, each opening with a one-page executive summary so the report reads at two depths.


Size & Growth

The honest answer on the lower market's size is that it must be triangulated, because no single provider sizes sub-\$100m EV AI companies as a category. Three anchors frame it: a population of thousands of application-layer firms, small-round flow of roughly \$2.3bn across seed rounds in ten months of 2025, and a record \$256bn of total AI capital that tells you almost nothing about the small company because frontier labs absorbed the bulk. The 2025-26 multiple reset is the other half of the picture, because it converts founders into motivated sellers.

\$45–49m	6,678
MEDIAN AI SERIES A PRE-MONEY	AI/ML DEALS IN 2025
\$256bn	~80%
TOTAL AI CAPITAL, 2025 (~3 × 2024)	OF VENTURE FUNDING WENT TO AI

A founder-led 15-person AI startup lives almost entirely inside the sub-\$50m round world and, on exit, inside the sub-\$100m EV band. The reset that repriced public software is exactly what widens the pool of acquirable, motivated sellers below it.

PUBLIC SAAS VALUATION MULTIPLE
The reset that feeds motivated sellers
The 2025–26 de-rating widened the gap between AI-native and commoditizing software




Sources: Tracxn (agentic-AI population); CB Insights (Q2 2025 median AI deal size); Carta (Q3 2025 Series A valuations); PitchBook; SaaS Capital Index; QuantPillar; Jefferies. Round data reflects fundraising, not control sales.

Size & Growth

Rather than a single CAGR, scale is best treated as a sensitivity spectrum. The base case has the lower-market AI population growing mid-teens to low-twenties percent annually, tracking the broader vertical-SaaS rate of 11.5% to 23%, while native players compound and the wrapper tail washes out. That churn is not noise; it is the motivated-seller supply that feeds lower-market deal flow.

LOWER-MARKET AI REVENUE · INDEXED TO 2024 = 100
A sensitivity spectrum, not a single CAGR
Solid = tracked, dashed = base case, shaded band = bear to bull range to 2030



The three scenarios

SCENARIO	CAGR
Bear	~12%
Base	~18%
Bull	~26%

Bear: open-weight parity plus feature absorption destroys wrapper margin faster than vertical depth accrues. **Base:** native players compound while wrappers churn out, capital concentrating in defensible niches. **Bull:** vertical agents reach production reliability, marketplaces monetize small developers, and SMB adoption compounds.

Sources: Mordor Intelligence (vertical SaaS 11.5%); venture-tracked sub-segment estimates (16–23%); builder estimate on wrapper mortality. Forecasts are forward-looking projections presented as a range, not a single figure.


Size & Growth

Headline CAGRs invert the quality ranking. The steepest cited figure, agentic AI at roughly 46%, is a top-down projection that should be treated as perishable, and the same data shows agentic funding concentrating into fewer, larger rounds, which is a maturation signal rather than pure expansion. The durable small-company core, vertical AI and vertical SaaS, grows more slowly but compounds on proprietary workflow and data a foundation-model upgrade cannot easily erase.

PROJECTED REVENUE CAGR BY SEGMENT • 2026–2030

Growth is steepest where it is least durable

Lime marks the defensible small-company core; agentic growth is the most perishable



Why vertical eats horizontal

Bessemer finds AI-native vertical companies reach \$100m ARR faster than any prior SaaS generation, concentrated in healthcare, legal and housing.

"Shooting Stars" grow from \$3m to \$100m ARR over about four years at roughly 60% gross margin, against a pre-AI norm of seven years.

"Supernovas" hit \$100m ARR in about 1.5 years but often on fragile retention and roughly 25% gross margin. Speed is not durability.

Size & Growth

The most consequential structural trend is that open-weight models now rival closed frontier models on practical workloads, and self-hosting a 70B-class model runs roughly \$8k to \$12k per month against \$45k or more in API fees at 100k requests a day. Falling API prices sound pro-small-company but cut both ways: they lower input cost while letting the model vendor undercut any thin reseller. The moat therefore migrates toward proprietary data and workflow depth.

FOUNDATION-MODEL API PRICE • USD PER MILLION INPUT TOKENS

API prices fell roughly 80% in a year

Every cut compresses wrapper margin while open weights set the floor



Four trends reshaping the build

Open-weight parity converts a variable, vendor-controlled COGS line into a fixed, controllable one and enables data-sovereign deployment.

Agentic AI is the formation magnet but is concentrating; the seed window for generic agents is closing while niche vertical agents still raise.

No-code and low-code platforms lower the build barrier while commoditizing the build itself.

Regulatory divergence makes EU-compliant small targets more valuable to cross-border acquirers needing market access.

Size & Growth

Four driver families shape demand. Regulation creates non-discretionary demand for compliant, auditable, domain-specific AI. Economic forces push founders toward M&A exits and lift AI-native capital efficiency to revenue-per-employee an order of magnitude above traditional SaaS. Technology cuts the cost and time to build. Above all, SMB adoption is the customer-driven engine, and growing SMBs prioritize AI capability while struggling firms remain price-led, which segments the buyer base and the achievable price.

Regulatory · structural

Compliance-driven AI tooling is projected to grow 25% to 30% annually through 2030; a fine of up to 7% of global turnover under the EU AI Act dwarfs the subscription.

Economic · capital efficiency

Top lean-AI startups average \$3.48m revenue-per-employee against roughly \$610k for top traditional SaaS, widening the pool of bootstrapped, profitable targets.


Customer · accelerating

96% of small-business owners plan to adopt AI; the top use cases are marketing (53%), sales (49%) and support (46%).

US SMB AI ADOPTION · 2024 — 2025

SMB adoption is the demand engine

Among firms with 10–100 employees, usage jumped from 47% to 68% in a year




Segmentation

The sector segments by position in the stack crossed with two decisive axes: model or infrastructure-native versus wrapper-heavy, which drives margin, retention and multiple, and buying behavior. The density of acquirable, founder-led, genuine-AI targets is highest in vertical AI agents and AI-native vertical SaaS, strong across MLOps, evaluation, vector tooling and premium data labeling, and effectively zero at the foundation-model and accelerator layers.

SEGMENT MAP · ACQUIRABLE DENSITY × DEFENSIBILITY

Where the sub-\$100m bolt-ons actually live

Bubble size ≈ deal activity; lime = the small-company hunting ground (top-right)



The screening rules

In the application family, the AI-versus-wrapper call decides whether a target is a bolt-on or an acqui-hire, and it correlates almost perfectly with NRR and gross margin.

Vertical agents take roughly 55% of agentic deals, the highest-quality acquirable agentic pool.

The wrapper tail is the largest population and the widest quality dispersion, and where most of the landmines sit.

Segmentation

SEGMENT	NATIVE / WRAPPER	REPRESENTATIVE	SMALL-COMPANY DENSITY
Vertical applied AI / agents	Native	EvenUp, Caseflood.ai, Corti, Tennr, EliseAI	Highest, moat resists model upgrades
AI-native vertical SaaS	Native	Jane Software, niche micro-SaaS	High, NRR 108–120%, bootstrapped scale
Vector / retrieval databases	Native	Qdrant, Zilliz/Milvus, Chroma	Strong, active bolt-on targets
MLOps / LLMOps / evaluation	Native	Arize, Weights & Biases, ZenML	Thin, consolidating fast into platforms
Data labeling / curation	Mixed	Surge AI, Argilla, Dappier	Premium tier defensible; low end commoditized
Horizontal applied AI	Wrapper-heavy	Midjourney, Copy.ai, HeyGen	Most populous, most dangerous, high churn
Agent infra / memory / orchestration	Native	Emerging, pre-revenue	Speculative, capital lags revenue-facing agents
GPU cloud · Accelerators · Foundation models	Native (frontier)	Neysa, FuriosaAI, Sarvam	Near-zero, capital-gated and unacquirable


End-Users

Demand runs through a different buyer than the enterprise mega-account: SMBs, prosumers, solo operators and single departments buying below the procurement threshold. Stated criteria diverge sharply from what wins, because workflow fit and vendor credibility beat RFP security checklists. A small company's user mix is therefore a leading indicator of its retention, its multiple, and its bolt-on-versus-acqui-hire classification.

NET REVENUE RETENTION BY TIER

Retention is the valuation line, not the AI label

The AI-versus-wrapper call correlates almost perfectly with NRR



Four buyer profiles

Regulated verticals (law, clinics, banks) rank certification, accuracy and data sovereignty above price, and retain best.

Growth SMBs are the demand engine, competitive rather than regulatory, product-led and ROI-elastic.

Prosumers and solo operators buy \$20–30/month tickets with near-zero switching cost and fragile ~32% NRR.

Single departments land below procurement thresholds, a fast bottom-up wedge that caps initial deal size.


Supplier Ecosystem

A fifteen-person AI startup in 2026 does not own its means of production: it rents intelligence from a handful of model providers, rents compute from a GPU oligopoly whose prices it cannot influence, and depends on a thin band of tooling and capital suppliers. The defining exposure is single-foundation-model lock-in, and the cruel twist is that falling API prices, the thing that should help, compress the wrapper's only margin. Switching cost, and therefore deal quality, is whatever proprietary data and workflow the company has manufactured.

\$2.99	600x +
MEDIAN H100 GPU-HOUR	SPREAD ACROSS THE MODEL MENU
>50m	3–6x
TOKENS/MO SELF-HOST BREAK-EVEN	HYPERSCALER PREMIUM VS NEOCLOUD

Where a small-cap acquirer can play: silicon is off-limits, but internalizing proprietary data, compliance tooling or inference optimization is achievable, and small targets there actually exist (Uber–Segments.ai, AMD–Brium, TrueFoundry–Seldon).

GROSS MARGIN %
AI economics carry a variable inference layer
AI-product margins climb but stay below the classic SaaS band



Sources: published API pricing (DeepSeek, OpenAI, Anthropic); neocloud and hyperscaler GPU rate cards; ICONIQ (AI-product gross margin 41/45/52); Bessemer (Supernova / Shooting Star); The SaaS CFO.

Complementary Products

A supplier feeds the core, a substitute replaces it, and a complement extends it. Complements are the tools, marketplaces and services that sit beside a small AI company's core product, deepen lock-in, and convert one-off usage into recurring revenue. Several build data or workflow moats foundation-model providers cannot easily absorb, which is precisely why the complement layer is among the most actionable lower-market hunting grounds.

Data enrichment

Data moat

Clay raised a \$3.1bn Series C and waterfalls 130+ sources; proprietary signal a wrapper cannot replicate.

Agent marketplaces

Channel

Salesforce AgentExchange, Claude Marketplace and the GPT Store give a small company distribution it could never afford to build.

Observability / eval

Active M&A

Braintrust (\$80m Series B at \$800m), Arize and Langfuse prove the product works, a precondition for serious buyers.

Integration / MCP

10,000+

Active public MCP servers turn the N×M integration problem into N+M; durable value migrates to the governance layer above it.

The M&A thesis is to roll observability, evaluation and governance into one AI-operations platform whose moat is proprietary telemetry, and to acquire complements that own proprietary data or an embedded workflow while avoiding those the protocol standard or the marketplace owner is actively commoditizing.

Sources: TechCrunch (Clay Series C); Sacra; company disclosures (Braintrust, Langfuse); Anthropic MCP ecosystem update; Sales force Q4 FY26 (Agentforce ARR \$800m); FE International 2026 framing.

Substitute Products

A substitute does not extend a small AI company's product; it removes the reason to buy it, and for a wrapper or thin-application startup this is the single most lethal force in 2026 because the most credible substitute is often free. Feasibility depends entirely on what the adopter must supply: native feature absorption demands nothing, while open-weight self-hosting demands GPUs, MLOps competence and the willingness to own uptime. Only proprietary data, workflow depth and accreditation defend against it.

Where substitution is highest

Sherlocking. A native feature in a product the customer already uses, often free; OpenAI's cadence reportedly cannibalized 200+ funded wrapper startups in 2024 alone.

Open-weight erosion. Qwen, DeepSeek and GLM at frontier-class performance under permissive licenses erase any edge a wrapper claimed from model access.

Incumbent bundling. ServiceNow spent \$11.6bn on security acquisitions in 2025 to turn point solutions into features of a suite the customer already owns.

Where SaaS reasserts value

Proprietary data. The one thing a foundation lab cannot scrape or ship; buying a dataset converts a Sherlockable feature into a defensible product.

Workflow depth. Permissions, approvals and system-of-record integrations raise switching costs beyond what a native feature can match.

Compliance moat. A substitute that cannot produce SOC 2, HIPAA or CE-marked evidence cannot serve regulated buyers.

The deal rule follows directly: infrastructure or data-native businesses with proprietary data survive and earn ARR multiples, while pure wrappers are acqui-hires priced for the team.


Competitive Landscape

Small AI companies compete less on raw model quality than on proprietary data, vertical workflow depth and switching costs, because model access is commoditizing. The field splits into four archetypes: global incumbents who are the acquirers rather than the competitors, regional champions who own a national or linguistic niche, niche specialists who are the largest and richest hunting ground, and technology-enabled wrapper entrants who are acqui-hire fodder. Retention crossed with margin decides which is which.

RETENTION × MARGIN DECIDES THE DEAL

The bolt-on quadrant sits top-right

NRR versus gross margin; lime = bolt-on at a premium, charcoal = acqui-hire



Niche specialists. Voyage AI was 17 months old when MongoDB bought it for \$220m; the data asset, not the model, is the moat.

Regional champions. Legora consolidated Nordic legal AI; Doctrine built a 27,000-professional base before LexisNexis moved to acquire it.

Proprietary data. Noetica's 1.5-billion-term deal database is why Thomson Reuters acquired it.

Wrapper entrants. Roughly 80% projected to disappear within 18 months as platforms expose their lack of moats.

Competitive Landscape

Differentiation runs along four structural sources, none of which is a better prompt: proprietary data, vertical workflow depth, cost position through model portability, and niche accreditation. Competition is price-driven where the product is a thin wrapper and trust-driven where it is embedded in a regulated workflow, which is why vertical AI in healthcare, legal and finance commands the durable multiples.

GLOBAL INCUMBENT

Salesforce · Thomson Reuters · MongoDB

The acquirers, not the competitors; they set the price of the APIs and cloud every small company rents

A feature shipped at a keynote can erase a product category overnight.

NICHE SPECIALIST

Voyage AI · Noetica · Seldon · Langfuse

Proprietary data or models; defensible margin; 15–30x revenue on sale

Voyage sold to MongoDB for \$220m at 17 months old; Langfuse to ClickHouse on 26m+ monthly SDK installs.

REGIONAL CHAMPION

Legora · Doctrine · Qura

Jurisdictional data moat; language and regulatory fit; sticky base

Curated, jurisdiction-specific data is a moat a US generalist model cannot scrape.

WRAPPER ENTRANT

Generic copilots and summarizers

Fast to market, low build cost; no moat, 52%-or-lower margin, 15–20% monthly churn post-Sherlocking

Acqui-hire soft landing or a pivot to proprietary data; not a platform.

Value Chain

Value in the AI stack is migrating away from raw model access, which is commoditizing fast, toward data interpretation and workflow ownership, which is defensible. For a small company the only durable positions are downstream of the model. Two structural dependencies define the position, the foundation-model API and the GPU, and value capture shifts toward whoever owns the data and the workflow, which is exactly the logic behind MongoDB buying Voyage and ClickHouse buying Langfuse.

VALUE-CHAIN STAGE	TYPICAL GROSS MARGIN	COMMODITIZED VS STRATEGIC	SMALL-COMPANY IMPLICATION
Compute / infrastructure (GPU, data center)	Thin	Commoditized (price-taker)	Avoid owning; rent. Not a sub-\$100m play
Data labeling / curation	30–50%	Mixed	Labor-heavy; acquired for data and clients
MLOps / orchestration / eval / vector DB	60–75%	Strategic if proprietary	Picks-and-shovels; prime bolt-on targets
Foundation models	50–70%	Boundary (not lower market)	Input cost only
Wrapper application	25–52%	Commoditized	Acqui-hire; fragile multiple
Native vertical application (proprietary data)	60–85%	Strategic (moat)	Highest durable multiple; switching costs

The most-acquired stages are the data and tooling layer, because it is fragmented and immediately accretive, and native vertical applications, because they carry proprietary data. Menlo Ventures frames the prize: vertical AI moves the ROI calculation from software budgets to far larger labor budgets.


M&A Landscape

The sub-\$100m AI M&A market is structurally opaque, with only about 21% of roughly 5,700 AI/ML acquisitions from 2020 to 2025 disclosing a value, but its shape is clear: capability and talent, not volume, drive activity. H1 2025 AI-target deal value rose 127% year on year even as volume fell about 20%, meaning acquirers paid premiums for capability rather than buying revenue. The native-versus-wrapper call is the central diligence question in the band.

ILLUSTRATIVE LOWER-MARKET AI ACQUISITIONS • USD M

Capability and talent, not volume, priced the band

H1 2025 AI-target deal value rose 127% YoY even as volume fell ~20%



Seven deal categories

Bolt-ons of niche specialists fold proprietary capability into a platform (MongoDB–Voyage the template).

Acqui-hires price wrapper and sub-scale teams at roughly \$1m to \$3m per engineer.

Small platform tuck-ins build out an agentic suite (Salesforce's eight 2025 AI deals).

PE buy-and-build add-ons consolidate vertical SaaS with AI inside it, funded by \$3.7tn of dry powder.

M&A Landscape

TRANSACTION	VALUE	TYPE	STRATEGIC SIGNAL
MongoDB – Voyage AI	\$220m	Native bolt-on	17-month-old company bought to make AI retrieval native to the database
Salesforce – Doti	~\$100m	Suite tuck-in	Raised only \$7m, an estimated 14x return; eighth Salesforce AI deal of 2025
Thomson Reuters – Noetica	~\$22m	Proprietary data	1.5-billion-term deal database folded into CoCounsel (Series A proxy)
Mistral – Koyeb	~\$15m	Vertical infra	13-person serverless-GPU team acquired to build Mistral Compute
ClickHouse – Langfuse	Undisc.	Tooling consolidation	Native open-source asset; low integration risk, architectures pre-aligned
TrueFoundry – Seldon	Undisc.	Adjacent capability	Production MLOps beneath an agentic gateway, both Kubernetes-native
Uber – Segments.ai	Undisc.	Team-plus-tech	Belgian lidar-annotation team internalized into Uber AI Solutions
Bow River – meQuilibrium + RippleWorx	\$10–49m	PE buy-and-build	Recapitalization merging two workforce-resilience platforms


M&A Landscape

The spread between top-tier and bottom-tier deals has never been wider: two companies with identical ARR can close at prices differing three to four times on AI defensibility, NRR, profitability and process. A native target carries a defensible margin and a 15-to-30x revenue multiple, while a wrapper carries a 52%-or-lower gross margin, monthly churn risk after a platform ships the feature, and an acqui-hire price. Mispricing this is the most common and most expensive error in the band.

EV / REVENUE BY ASSET CLASS · 2026

The valuation barbell: native premium, wrapper discount

Proprietary-data natives clear 15–30x; AI-threatened tools sit at 3–5x



8–12x

QUALITY ASSETS, NRR 110%+

Versus 3–5x for 8%+ churn; retention translates directly into multiple.

\$3.7tn

GLOBAL PE DRY POWDER

Double 2019; underpins continued buy-and-build appetite at the small end.

3–5x

SDE · MICRO-SAAS WRAPPER

Deals under \$5m priced on earnings (~17.2x EBITDA), not revenue.

Sources: Aventis Advisors (~23.4x average, 6–50x dispersion); FE International 2026; SaaS Capital Index; ClearlyAcquired; Preqin (dry powder); QuantPillar. Private multiples reflect fundraising and small-deal comps, treated as directional.

M&A Landscape

Source native, structure for the team, and price the wrapper discount. Durable value accrues to proprietary data and infrastructure assets bought at a defensible multiple, while wrapper capabilities are better partnered with or licensed than acquired. The throughline is to own a control point the substitute cannot replicate, price substitution risk honestly, and let the valuation barbell time entry into mispriced assets.

#	PRIORITY	MOVE & STRUCTURE
1	Proprietary-data and infrastructure assets	Embeddings, eval, MLOps, vector and serverless inference (the Voyage / Langfuse / Seldon / Koyeb profile). Share deal, 60–70% cash at close, 30% founder roll-over, 2–3 year retention tied to capability milestones.
2	Compliance-native vertical applications	Healthcare, legal, finance and construction with proprietary datasets, where the EU AI Act and US sectoral rules create switching costs. Earnouts of 20–30% gated on retention and integration.
3	Regional champions as roll-up platforms	Jurisdictional data moats (the Legora / Doctrine profile) as add-ons or platforms; majority recap with management rollover and cross-border expansion.
4	Wrapper and sub-scale teams	Asset deal plus acqui-hire at roughly \$1m–3m per engineer, signing and RSU retention, and a founder carve-out negotiated ahead of the liquidation waterfall.
5	Defensive capability for portfolio companies	Bolt proprietary data, compliance tooling or inference optimization onto exposed applications to lift NRR and convert an acqui-hire into a bolt-on.

Sources & References

Figures are triangulated across multiple providers rather than taken from any single source, and discrepancies are flagged rather than smoothed. AI-native private valuations and ARR derive substantially from venture trackers and trade press rather than audited filings, and are treated as directional. Sub-\$100m deal values are frequently undisclosed and shown as funding-proxy estimates. Data reflects the most recent available figures as of mid-2026.

Population & capital flow

Tracxn · CB Insights · PitchBook · Crunchbase · Carta · Dealroom · S&P Capital IQ · Bessemer Venture Partners.

Benchmarks & multiples

Aventis Advisors · SaaS Capital · ICONIQ Growth · ChartMogul · Benchmarkit · Recurly · FE International · QuantPillar · ClearlyAcquired.

Deals, structural & regulatory

PwC · Software Equity Group · SaaSrise · Bain · Menlo Ventures · a16z · Stanford HAI · EU AI Act · company filings & press.

Contact us



Diego Gutiérrez Zarza

Corporate Finance Partner
dgutierrez@bakertilly.es



Cristina Salmón

Head of Tech Sellside
cristina.salmon@bakertilly.es

www.techma.bakertilly.es

© 2024 Baker Tilly (Spain) is an independent member of Baker Tilly International. Baker Tilly International Limited is an English company. Baker Tilly International does not provide professional services to its clients.

