




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.

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

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.


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


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.

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.

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.

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.

