

The Databricks partner ecosystem

Market Research · September 2026

A crowd-phase channel sitting on the highest-priced work in enterprise software. Data and analytics partners clear a 16.3x EBITDA median - and this ecosystem is almost entirely unconsolidated.

Now, for tomorrow

bakertilly

Executive Summary

Databricks is a crowd-phase ecosystem attached to the highest-priced category of partner work in the study. Data and analytics consultancies carry a 16.3x EBITDA median — above cybersecurity at 16.1x and roughly double the systems-integration band — and every AI workload an enterprise deploys needs a governed data foundation first. What the channel does not yet have is consolidation: Databricks carries no verified partner-firm deal count in the 1,364-deal dataset, entry remains easy and certifications discriminate weakly. The result is an unusual combination for an acquirer: premium end-work, commodity entry conditions, and no established buy-side to bid against.

Crowd

LIFECYCLE PHASE

Entry stays easy, partner count grows, certifications discriminate little. M&A has not yet begun.

16.3x

DATA & ANALYTICS EBITDA MEDIAN

The highest-valued subsector in the study — first quartile 10.7x, third quartile 27.1x.

None

VERIFIED PARTNER-FIRM DEALS

No entry in the per-ecosystem table. The channel is unmeasured, not consolidated.

2.2x+

REVENUE MULTIPLE, VALUE-ADDED POLE

Data and AI consultancies and IP-rich ISVs against 0.5x for licence resale.

Grows

FIRM-COUNT DIRECTION

Crowd ecosystems add firms at high single digits while mature channels shed them.

40% by 2026

ENTERPRISE APPS WITH AI AGENTS

Gartner projection, from under 5% in 2025 — and every agent needs governed data.

Time but act

SELL-TIMING READ

No scarcity clock, but early movers risk locking up the category.

Bought, not built

WHAT THE GSIS ARE DOING

Accenture's Keepler, RANGR Data and Faculty run is data and AI capability acquired outright.

Sources: Aventis Advisors subsector medians (data and analytics 16.3x EBITDA, first quartile 10.7x, third quartile 27.1x; value-added pole 2.2x-plus revenue; VAR 0.5x) and the 1,364-deal per-ecosystem table, which carries no Databricks entry; lifecycle-phase assignment (crowd); Gartner agentic-AI projection; Accenture capability-acquisition run; all as cited in Baker Tilly, *The Enterprise Software Partner Ecosystem* (Sept 2026).

Scope & Definition

The unit of analysis is the partner firm, not Databricks and not the end customer. Databricks consumption revenue is out of scope. This ecosystem carries no verifiable partner count in the reference table, so no density ratio is quoted — a deliberate omission rather than a gap, since inventing a denominator to produce a percentage is exactly the error this methodology exists to avoid. What can be stated with confidence is the phase, and the phase is what sets the acquisition logic.

Premium work, commodity entry conditions

The combination that makes this ecosystem unusual

What the work is worth

16.3x

EBITDA median, data & analytics

The highest-valued subsector in the study, ahead of cybersecurity at 16.1x and double the SI band.

THIRD QUARTILE: 27.1x

What the channel looks like

Crowd

lifecycle phase

Entry easy, count growing, certifications discriminate little, no verified partner M&A yet.

NO DENSITY QUOTED

THE PRACTICAL CONSEQUENCE

A buyer can acquire firms doing the study's best-paid work without paying a scarcity premium, because the program has not yet created scarcity. That gap between what the work earns and what the channel costs is the entire thesis here — and it closes the moment Databricks tightens its tiers.

In scope, by partner function

Resellers; global and regional systems integrators; data and AI implementation consultancies, **the prime buy-and-build target and the highest-priced partner type in the study**; managed service partners running platforms and pipelines; ISV and technology partners; training and enablement partners.

Counting rules applied

- → **No invented denominator.** Where the report has no verifiable partner count, it declines to compute density.
- → **Never sum across vendors.** The ten global system integrators hold top-tier Databricks status and top-tier status in nearly every other program simultaneously.
- → **Estimates are labelled.** Where a figure is a Baker Tilly estimate rather than a disclosure, this deck marks it INDICATIVE in place.

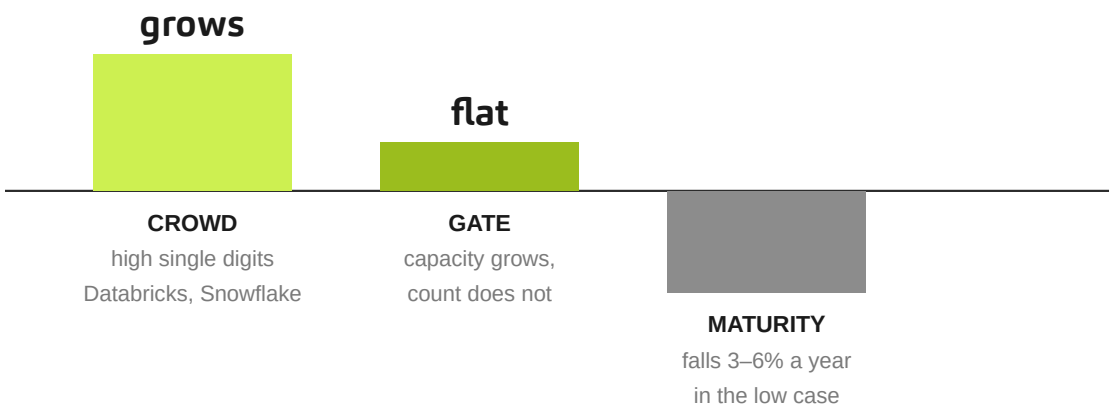
Sources: verified partner-count reference table and standing methodological rules in Baker Tilly, *The Enterprise Software Partner Ecosystem* (Sept 2026), which assigns Databricks to the crowd phase and carries no verifiable partner count for it; Aventis Advisors subsector medians (data and analytics 16.3x EBITDA, first quartile 10.7x, third quartile 27.1x; cybersecurity 16.1x).

Size, Growth, Trends & Drivers

Growth direction in a partner ecosystem is set by phase, not by the vendor's revenue. Open-door and crowd ecosystems add firms at high single digits; mature ones shed them at 3 to 6 percent a year in the low case. Databricks sits with Snowflake, Freshworks and Anthropic in the growing half, which makes it one of the few channels in the study where the addressable target population is getting larger. Underneath that, the demand driver is the strongest in the sector: agentic AI cannot be deployed on ungoverned data, and Gartner projects 40 percent of enterprise applications will carry task-specific agents by end-2026, from under 5 percent in 2025.

Phase sets the direction, not the vendor

Independent-firm count by lifecycle phase, base case



The base case across the whole partner economy is mid-to-high-single-digit growth in partner-attributed services revenue. A single blended CAGR for "the partner market" averages opposing motions into a meaningless middle.

Four drivers a partner firm can price

- → **Agentic AI needs governed data first.** The single strongest demand driver in the study. Forty percent of enterprise applications carrying task-specific agents by end-2026 means a data-foundation programme in front of each one.
- → **The highest-priced adjacency is the core work.** Unlike a cloud or ERP partner that has to acquire data capability to reach 16.3x, a Databricks partner starts there.
- → **GSIs are buying rather than building.** Accenture's Keepler, RANGR Data and Faculty acquisitions are data and AI capability bought outright — evidence of what the capability is worth to a strategic.
- → **Regulated AI.** The EU AI Act reaches full application in August 2026 and DORA carries fines up to 2 percent of worldwide turnover. Governance, lineage and auditability become billable requirements rather than best practice.

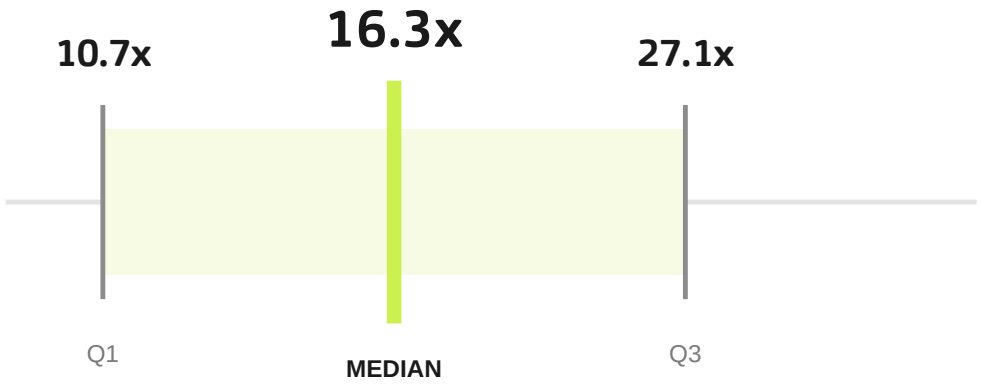
Sources: lifecycle growth-direction analysis and sensitivity spectrum in Baker Tilly, *The Enterprise Software Partner Ecosystem* (Sept 2026), assigning Databricks and Snowflake to the crowd phase; Gartner agentic-AI projection; Accenture capability-acquisition run (Keepler, RANGR Data, Faculty); DORA and EU AI Act application dates.

Segmentation

The data and analytics subsector has the widest quartile spread of any category in the study — a first quartile of 10.7x EBITDA against a third quartile of 27.1x around a 16.3x median. That spread is the segmentation story: two firms doing nominally the same Databricks work can be valued two and a half times apart. What separates them is IP, recurring share and vertical depth, not certification, because in a crowd ecosystem the certification is easy to hold and therefore prices nothing.

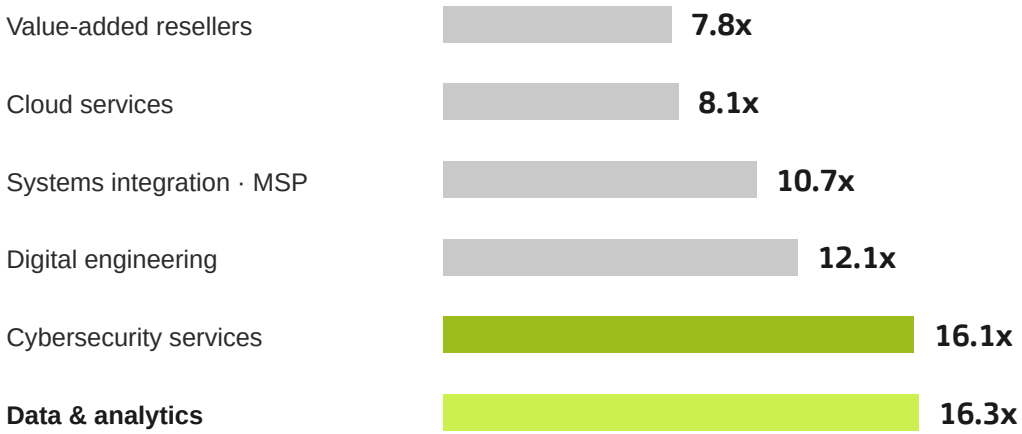
The widest spread of any subsector

Data & analytics EV / EBITDA, quartiles around the median



Where data sits against everything else

EV / EBITDA subsector medians



WHAT MOVES A FIRM ACROSS THE RANGE

Proprietary IP and accelerators, recurring share, and vertical depth. Not headcount, and in a crowd ecosystem not certification either — the badge is easy to hold.

SCALE STILL MATTERS

EBITDA medians rise from 6.7x below USD 5m enterprise value to 12.7x above USD 500m. Scale alone is worth roughly six turns, independent of subsector.

Sources: Aventis Advisors subsector EBITDA medians and quartiles (data and analytics 16.3x median, 10.7x first quartile, 27.1x third quartile; cybersecurity 16.1x, digital engineering 12.1x, systems integration and MSP 10.7x, cloud services 8.1x, VAR 7.8x) and deal-size bands (6.7x below USD 5m EV to 12.7x above USD 500m); crowd-phase segmentation analysis; all as cited in Baker Tilly, *The Enterprise Software Partner Ecosystem* (Sept 2026).

End-Users

Databricks buyers are the organisations with data problems large enough to need a platform, which skews the base toward large commercial and regulated enterprises with real engineering budgets. Stated criteria are cost, functional fit and references; revealed criteria are risk reduction, certified depth and speed to value. The distinguishing feature of this buyer set is that its need does not end at go-live: pipelines drift, schemas change and governance obligations accumulate, which makes data work structurally more annuity-shaped than implementation work in most other ecosystems.

Core

Large commercial

Retail, CPG, manufacturing, telecom, media. Data volumes and AI ambitions exceed internal capability. **Buys on vertical depth and accelerators.**

Highest

Regulated enterprise

BFSI, healthcare, pharma, insurance. Lineage, auditability and sovereignty demand certified partners and documented controls. **The EU AI Act sharpens this.**

Mandatory

Public sector

Sovereignty and certification requirements make onshore delivery compulsory, creating a distinct and less contested buyer pool.

No bench

Mid-market

Holds no data-engineering headcount. Buys fixed-fee predictability and single accountability — and is the segment a crowd-phase boutique can actually win.

Recurring

All embedded buyers

Pipelines drift and governance obligations accumulate every month. **The most natural annuity attachment of any partner category.**

Why data work retains

An implementation ends; a data estate does not. Schema changes, new sources, model retraining and audit requirements recur on a monthly cadence, which is why a data partner can convert project revenue into a retainer more easily than an ERP or ITSM implementer. Recurring MSP-style revenue prices at roughly 6 to 8x EBITDA against 3.5 to 5x for project revenue, and in this category the attachment is the natural next conversation rather than an upsell.

Pricing logic is moving

Value-based firms grew 8.7 percent against 2.1 percent for time-based, and buyers increasingly refuse to pay for hours AI has removed. Partner-firm retention benchmarks are not publicly disclosed with rigour; that gap is flagged, not filled.

Sources: end-user profiles and buying-criteria matrix in Baker Tilly, *The Enterprise Software Partner Ecosystem* (Sept 2026); Salt Creek Advisory recurring-vs-project pricing framework (6–8x versus 3.5–5x EBITDA); Deloitte 2025 Professional Services Benchmark; EU AI Act application date. Partner-firm retention benchmarks are not publicly disclosed with rigour; that gap is flagged, not filled.

Supplier Ecosystem

In a crowd ecosystem the vendor still supplies the non-substitutable inputs — product, certification, co-sell pipeline — but it is recruiting rather than gating, which gives a partner more negotiating room on economics than it will ever have again. The binding constraint sits in the labour market instead, and in this category it binds harder than anywhere: data and AI engineering skills are the specific skills driving above-inflation wage growth, at the same time as India's fresher intake collapsed from roughly 380,000 to roughly 120,000.

Databricks (tier 1)

Product, certification, co-sell pipeline and marketplace rails — non-substitutable. But the program is **expanding rather than gating**, so terms are more negotiable than in a mature ecosystem. That window closes as the channel matures.

Data & AI engineers

The hardest labour constraint in the study. **AI and ML skills drive above-inflation wage growth** while the graduate funnel narrowed sharply — which is why the GSIs are buying whole firms rather than recruiting.

Underlying cloud compute

For any managed service the compute sits with a hyperscaler, all of which were named among the 19 Critical ICT Third-Party Providers designated under DORA in November 2025. A second supplier layer beneath the platform.

Regulators & standards

SOC 2, ISO 27001, HIPAA, DORA and the EU AI Act supply the right to operate in regulated segments. The AI Act reaches full application August 2026 — **governance becomes a billable deliverable.**

The internalisation play

Because talent is the constraint, the highest-return supplier moves here are the talent moves. Acquiring offshore or nearshore delivery converts the largest variable cost into an owned asset and insulates the firm from wage inflation. Acquiring training internalises the capacity constraint and builds a funnel in the scarcest skill set in the sector. And converting repeatable pipeline and governance work into accelerators moves the firm from selling scarce hours to licensing product — which is what separates the 27.1x third quartile from the 10.7x first.

The vendor is also a competitor

As in every ecosystem, the vendor's own services motion takes the largest and most strategic accounts while the channel covers the long tail. In a crowd phase that matters less than it will later, but it is worth pricing — the direction of travel for a successful platform is always upmarket and always toward gating.

Sources: supplier-power and internalisation analysis in Baker Tilly, *The Enterprise Software Partner Ecosystem* (Sept 2026); NASSCOM and Indian IT-sector wage and intake data (fresher intake from ~380,000 to ~120,000; AI and ML skills driving above-inflation wage growth); European Supervisory Authorities CTPP designations (18 November 2025); DORA and EU AI Act application dates.

Complementary & Substitute Solutions

A Databricks partner already sits in the highest-priced services category, so the complement question is not how to reach a premium but how to hold it: productised accelerators and recurring governance retainers are what move a firm from the 10.7x first quartile toward the 27.1x third. The substitute picture is the one genuinely distinctive risk in this ecosystem — the same agentic AI that creates the demand also automates the pipeline work, and the platform itself keeps absorbing capabilities partners used to differentiate on.

Complements, ranked by multiple effect

- → **Productised accelerators and IP.** Repeatable pipeline, lineage and governance work productises unusually well, and productised delivery carries a software multiple rather than a services one.
- → **Governance and compliance retainers.** The EU AI Act and DORA make auditability a standing obligation rather than a project. Cybersecurity services price at a 16.1x EBITDA median on the same logic.
- → **Managed platform and pipeline run.** Converts project revenue into annuity at roughly twice the multiple, and in a data estate the run work is genuinely continuous.
- → **Agent management and MCP integration.** Emerging billable categories: enterprises deploying agents need supervision, audit and integration on top of the data foundation.

Substitutes, priced honestly

- → **Agentic AI and code generation.** Compresses standardised pipeline construction toward zero marginal cost. It substitutes for billable hours, not for the firm — but hour-priced data shops carry the risk directly.
- → **Platform absorption.** Capability that ships natively is capability a partner stops billing for. The pattern is visible across the sector: ServiceNow bought Moveworks for 2.85bn dollars, IBM bought HashiCorp for 6.4bn.
- → **The captive GCC.** India hosts more than 1,700 global capability centres generating roughly 64.6bn dollars and employing 1.9 million people, heading toward 2,400 by 2030. Data engineering is exactly the recurring work captives insource.
- → **In-house data teams.** Enterprises with the budget to build a platform often have the budget to staff it. The partner's defence is specialisation and speed, not scale.

Sources: Aventis Advisors subsector multiples and quartiles (data and analytics 16.3x median, 10.7x Q1, 27.1x Q3; cybersecurity 16.1x); Salt Creek Advisory recurring-vs-project pricing framework; ServiceNow / Moveworks USD 2.85bn and IBM / HashiCorp USD 6.4bn (Mordor Intelligence); NASSCOM FY24 GCC data; Gartner agentic-AI projections; EU AI Act and DORA provisions; all as cited in Baker Tilly, *The Enterprise Software Partner Ecosystem* (Sept 2026).

Competitive Landscape

Crowd ecosystems are dominated by proliferating boutiques and mid-market specialists, with certification thin relative to headcount and no scarcity for a buyer to arbitrage. That describes the Databricks channel — but with one important difference from Freshworks, the other crowd ecosystem in this set: the work these boutiques do is the work strategics are paying premiums for elsewhere. The GSIs are already buying data and AI capability aggressively; they have simply not yet organised that buying around this ecosystem's badge.

Where Databricks sits in the lifecycle

Four-phase partner-ecosystem model

1 Open Door	2 Crowd	3 Gate	4 Maturity
Anthropic	Snowflake, Freshworks	AWS, ServiceNow	SAP, Salesforce, Oracle



Crowd

WHAT THE PHASE IMPLIES

Partner count grows, entry stays easy, certifications discriminate little and partner M&A has not begun. No scarcity premium to arbitrage — but also no competitive tension, on work that strategics are paying 16x-plus for in other ecosystems.

Who is buying this capability already

- → **Global system integrators.** Accenture leads the sector at 201 verified acquisitions with an FY26 acquisition-spend target near 9bn dollars, and its 2025–26 run — Keepler, RANGR Data, Faculty, NeuraFlash — is explicitly data and AI capability. Strategics pay a median 9.8x EBITDA.
- → **PE platforms.** Sponsors drove over 27 percent of all deals and pay a median 11.5x EBITDA, buying sub-5m-EBITDA firms at 5–8x into platforms valued 10–14x. No named Databricks platform exists yet.
- → **Frontier-AI services vehicles.** Ode with Anthropic at roughly 1.5bn dollars and the OpenAI Deployment Company at over 4bn dollars on a 14bn-dollar valuation recruit from the same talent pool.
- → **What is missing.** A platform organised around this ecosystem specifically. The capability is being bought; the channel is not being consolidated.

Sources: crowd-phase competitive-structure analysis and the five acquirer archetypes in Baker Tilly, *The Enterprise Software Partner Ecosystem* (Sept 2026); Aventis Advisors acquirer league table and sponsor-vs-strategic multiples; Accenture capability-acquisition run (Keepler, RANGR Data, Faculty, NeuraFlash); Ode with Anthropic and OpenAI Deployment Company formation data.

Value Chain

The Databricks chain is short and sits high. There is no meaningful licence-resale layer to speak of for most partners, and the stages that dominate — data engineering, governance, platform run — are the ones that in other ecosystems represent the premium adjacency rather than the core business. That is the structural advantage of this channel. The exposure is concentrated in one place: pipeline construction priced by the hour is the most automatable work in the chain.

STAGE	REVENUE MODEL	WHO CAPTURES IT	AI COMPRESSION	MULTIPLE SIGNAL
Demand generation & co-sell	Referral / co-sell	Databricks keeps most	Low	Neutral — but drives deal flow
Licence resale & provisioning	Pass-through margin	Reseller	High	**0.5x revenue · 7.8x EBITDA** — thin here
Advisory & architecture	T&M, fixed, outcome	IT consulting firm	Rising	**9.8x EBITDA**
Pipeline build & migration	T&M, fixed fee	Systems integrator	Highest	**10.7x EBITDA · 1.1x revenue**
Data engineering & analytics	Project plus recurring	Data & analytics firm	Medium	**16.3x EBITDA** — the ceiling, and the core
Platform run & managed pipelines	Recurring subscription	MSP	Low	**10.7x, to 11.2x above USD 500m**
Governance, lineage & compliance	Retainer	Cyber / governance specialist	Low	**16.1x EBITDA** (cybersecurity)

The consolidator's instruction

Buy firms already living in the bottom three rows, or firms that can be moved there. The median private IT-services transaction fell to 8.8x EBITDA in late 2025 from 13.6x in 2023 because buyers stress-tested billable-hours models; a data shop selling pipeline construction by the hour is exposed to exactly that test, while one selling governed, run and productised outcomes is not.

Geography sets the base

EV/EBITDA medians run 11.2x in Asia, 10.9x in North America, 10.0x in Europe and 8.3x elsewhere, with emerging-market firms discounted for scaling and political risk. North America carries the largest median deal size at 73 to 85m dollars.

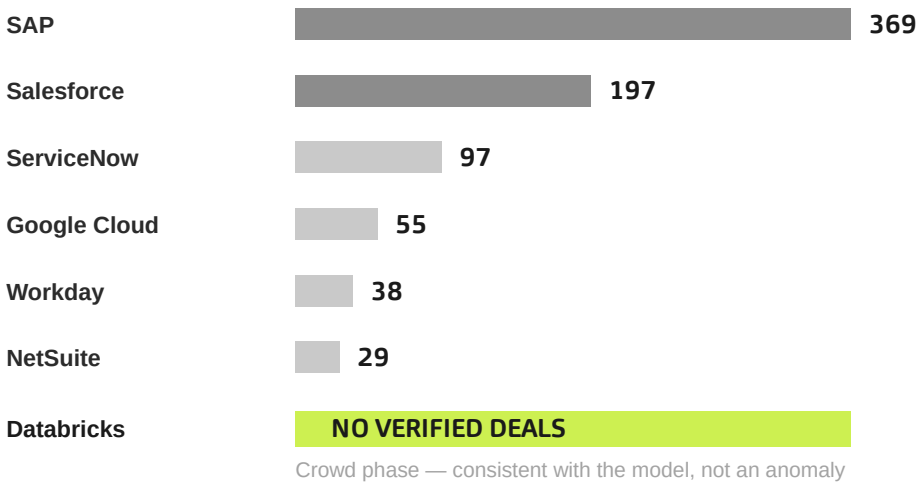
Sources: Aventis Advisors value-chain stage multiples, subsector medians, regional EV/EBITDA medians and the Q4 2025 median of 8.8x against a 2023 peak of 13.6x; all as cited in Baker Tilly, *The Enterprise Software Partner Ecosystem* (Sept 2026).

M&A Landscape

Twelve ecosystems carry verified partner-firm deal counts in the 1,364-deal dataset; Databricks carries none. That is consistent with the crowd phase rather than anomalous — partner M&A begins at the gate and intensifies through maturity, because that is when scarcity appears. But the absence is more striking here than in any other crowd ecosystem, because the capability these partners hold is precisely what the most acquisitive firms in the sector have been buying: Accenture's Keepler, RANGR Data and Faculty purchases are all data and AI capability bought rather than built.

M&A begins at the gate, not in the crowd

Verified partner-firm acquisitions 2015–2025



Crowd phase — consistent with the model, not an anomaly

Yet the capability itself is being bought hard: Accenture's Keepler, RANGR Data and Faculty acquisitions are all data and AI, at a 16.3x subsector median.

Five theses that would apply here

- → **Buy the premium category at commodity entry.** The core thesis: 16.3x end-work in a channel with no scarcity premium and no competing platform.
- → **Productisation arbitrage.** Move a firm from the 10.7x first quartile to the 27.1x third by converting repeatable pipeline and governance work into licensable accelerators.
- → **Multiple arbitrage.** Sub-3m-dollar-EBITDA specialists at 5–8x folded into platforms valued 10–14x.
- → **Governance as annuity.** The EU AI Act and DORA turn auditability into a standing retainer rather than a project deliverable.
- → **Talent acquisition, plainly.** Data and AI engineers cannot be recruited at scale in the current labour market. Buying the team is the only fast route.

Sell-timing read

Time it, but act. Crowd-phase owners face no scarcity clock and should expect to educate any buyer they meet. The risk is categorical rather than cyclical: if early movers plant flags and assemble the first platform, the independents that waited end up selling into a single-buyer market with no tension.

Sources: Aventis Advisors per-ecosystem verified deal counts (SAP 369, Dynamics 210, Salesforce 197, Oracle 115, Azure 106, ServiceNow 97, AWS 96, Google Cloud 55, Adobe 52, IBM 46 traced, Workday 38, NetSuite 29 — no Databricks entry) and lifecycle-timed sell-side strategy, including the crowd and open-door "time but act" read; Accenture capability-acquisition run; all as cited in Baker Tilly, *The Enterprise Software Partner Ecosystem* (Sept 2026).

bakertilly

Relevant Transactions & Takeaways

There are no Databricks-specific partner-firm transactions to report. The reference set below is what a buyer or seller in this channel will actually be priced against: the data and AI capability acquisitions happening right now in adjacent ecosystems, and the platforms assembled in crowd channels once they gated. Every one was done in a market that had no comparable transactions at the time either.

TRANSACTION	ACQUIRER	YEAR	VALUE / MULTIPLE	WHY IT IS THE ANALOGUE
Keepler Data Tech	Accenture	2026	**n.d.**	Cloud-native data engineering bought outright — the closest analogue to a Databricks partner asset in the dataset.
RANGR Data	Accenture	2026	**n.d.**	Same capability run, bought in sequence rather than assembled internally. Evidence of what data capability is worth to a strategic.
Faculty	Accenture	2026	**~$1bn**	Applied-AI services plus a proprietary decision-intelligence product; 400-plus staff. FT framing; terms undisclosed.
Artium	AlixPartners	2026	**n.d.**	78 employees bought for tiered relationships with both OpenAI and Anthropic — accreditation as the scarce asset.
Thirdera	Cognizant	2023	**~$430m**	What a platform built in a crowd channel becomes once the ecosystem gates — the exit an early consolidator underwrites.
Perficient	EQT / GIC	2024	**~18.5x EBITDA**	Sponsors pay far above the 11.5x median for scaled multi-platform assets. Estimate on disclosed reporting.

1 Premium work, no premium price

16.3x end-work in a channel with no scarcity clock and no competing platform. That gap is the thesis.

2 Underwrite IP, not certification

The badge prices nothing in a crowd phase. Accelerators, recurring share and vertical depth do.

3 Buy the team, plainly

Data and AI engineers cannot be recruited at scale. That is why the GSIs are buying whole firms.

4 Sell governance as annuity

The EU AI Act and DORA make auditability a standing obligation. Retainers, not deliverables.

5 Watch for tier tightening

The moment thresholds rise, scarcity appears and the buy-side arrives. That is the window being bought.

Sources: Aventis Advisors named-deals reference table (Accenture / Keepler and Accenture / RANGR Data 2026; Accenture / Faculty ~USD 1bn on FT framing with terms undisclosed; AlixPartners / Artium, announced 4 August 2026, 78 employees; Cognizant / Thirdera ~USD 430m 2023; EQT and GIC / Perficient ~18.5x EBITDA, estimate on disclosed reporting); all as cited in Baker Tilly, *The Enterprise Software Partner Ecosystem* (Sept 2026). No Databricks partner-firm transactions are recorded in that dataset. n.d. = value not disclosed.

Contact us

advisory@bakertilly.es

+34 946 42 41 42

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.

bakertilly