AI Valuations: Q2 2026

Q2 2026 finds the AI sector in a state Windsor Drake characterises as conviction meeting scrutiny. Valuations remain far above the broader software market, but the premium is now conditional on demonstrated revenue rather than narrative. Windsor Drake's working benchmark for AI-native application software has settled near **11x EV/Revenue**, while foundation-model labs are priced on strategic optionality at **15x to 30x** annualised revenue and legacy software with no credible AI position has compressed toward **3x to 6x**.

The demand underwriting those multiples is real. Combined 2026 capital spending by the major hyperscalers is tracking above **$650B**, Bain & Company projects roughly **200 GW** of incremental AI compute demand by 2030, and McKinsey estimates generative AI could create **$2.6T to $4.4T** of annual value. The defining tension of the quarter is the distance between that investment and proven enterprise return: McKinsey finds **94%** of enterprises had not yet seen significant value from AI spend as of the end of 2025.

Capital markets have reopened in parallel. AI absorbed the large majority of Q1 2026 global venture funding, with quarterly AI investment surpassing the entire full-year 2025 total, and OpenAI's **$122B** round set an **$852B** valuation. Cerebras Systems completed the largest US technology IPO since 2019, and private equity entered the year with roughly **$3.7T** of dry powder. At the same time, institutions including the International Monetary Fund have flagged concentration risk and drawn comparisons to the dot-com peak, with the largest technology names now near a fifth of the MSCI World Index.

This report sets out institutional-grade analysis for navigating that market: one in which position in the AI stack and proof of monetisation, not an AI label, determine where an asset is priced.

What multiples are AI companies trading at?

The Q2 2026 valuation picture is best read as a gradient running down the AI stack. The model and compute layers carry the highest multiples on scarcity and demand; the application layer is priced on proven revenue; and software with no genuine AI position has re-rated sharply lower. The dispersion between the top and bottom of the table is the widest in the technology market, and it is widening.

Investors and acquirers are paying for position and for demonstrated demand. Foundation models sit at the top because their marks reflect a bet on owning the model layer of a multi-trillion-dollar market. Infrastructure and compute follow, underwritten by contracted demand but discounted for capital intensity. AI-native software is priced on revenue quality and retention, and the Windsor Drake benchmark of roughly 11x anchors that cohort.

Table 1. AI Valuation Multiples by Subsegment, Q2 2026

Subsegment	EV/Revenue Range	YoY Trend	Primary Driver
Foundation Models & Frontier Labs	15.0x - 30.0x	Rising	Strategic optionality, compute access
Agentic AI Platforms	12.0x - 25.0x	Rising	Workflow ownership, measurable ROI
AI Semiconductors & Hardware	12.0x - 22.0x	Rising	Design wins, contracted order books
AI Infrastructure & Compute	10.0x - 22.0x	Rising	Training and inference demand
AI Data & Developer Tooling	9.0x - 16.0x	Strengthening	Production-grade AI deployment
AI-Native Software	8.0x - 15.0x	Stable	Recurring revenue, retention
AI Cybersecurity	7.0x - 12.0x	Stable	Governance and trust demand
Legacy / Non-AI Software	3.0x - 6.0x	Compressing	Displacement by AI-native entrants

Source: Windsor Drake analysis of McKinsey & Company, PitchBook and CB Insights data.

Subsegment dynamics driving the dispersion

At the top of the table, foundation-model labs are not valued on revenue at all in any conventional sense. Their marks reflect the probability of dominating the model layer, and the headline multiple understates the underwriting logic. AI infrastructure and semiconductors are underwritten by contracted demand and, increasingly, by secured access to firm power, which Bain identifies as the binding constraint on AI growth. At the application layer, AI-native software is re-rating on proof: net revenue retention above 120% and demonstrated customer ROI separate the names that hold a premium from those that do not. Legacy software without an AI position is the clear loser, compressing as AI-native entrants threaten displacement.

Table 2. Subsegment Valuation Drivers and Principal Risks, Q2 2026

Subsegment	Premium Driver	Principal Risk
Foundation Models & Frontier Labs	Probability of winning the model layer	Sharp repricing, revenue concentration
Agentic AI Platforms	Capture of enterprise labour budgets	Dependency on third-party base models
AI Infrastructure & Compute	Contracted training and inference demand	Capital intensity, power constraints
AI Semiconductors & Hardware	Hyperscaler design wins and lock-in	Foundry and packaging concentration
AI-Native Software	Recurring revenue and high retention	Base-model commoditisation
Legacy / Non-AI Software	Installed base and switching costs	Displacement by AI-native software

Source: Windsor Drake analysis of McKinsey & Company, Bain & Company and CB Insights research.

How are AI companies valued in 2026?

Valuation in 2026 splits cleanly into two regimes. At the frontier, labs are valued on optionality; everywhere else, AI companies are valued on a disciplined framework built on revenue quality, efficiency and a credible path to profitability. Understanding which regime applies is the first step in any defensible valuation.

Optionality at the frontier

Foundation-model labs are priced on strategic position, not current revenue multiples. OpenAI's March 2026 round set an $852B valuation, and xAI was valued at $250B within its combination with SpaceX. These marks reflect a bet on owning the model layer of a market McKinsey sizes at $2.6T to $4.4T of annual value potential. Valuation tracks secured access to compute, power, capital and talent. The approach carries the widest dispersion of outcomes in the market, spanning a genuine winner-take-most result and a sharp repricing.

The Rule of 40 for AI-native software

For AI-native software with real revenue, the Rule of 40, where revenue growth plus margin reaches at least 40%, remains the cleanest predictor of a premium multiple. Top performers command **50% to 100%** premiums over the median, yet only an estimated **15% to 20%** of AI-native software clears the bar. Frontier labs are exempt from this test: they are pre-profit and priced on optionality.

Table 3. Rule of 40 Performance Tiers for AI-Native Software, Q2 2026

Performance Tier	Rule of 40 Score	Avg EV/Revenue	Premium vs Median
Top Quartile	Above 50	15.0x and above	+50% to +100%
Rule of 40 Met	40 to 50	11.0x - 15.0x	Healthy premium
Near Miss	30 to 39	8.0x - 11.0x	Modest discount
Bottom Quartile	Below 30	4.0x - 8.0x	Deep discount

Source: Windsor Drake analysis of McKinsey and Bain & Company software value-creation research.

Proving monetisation

The defining diligence question of 2026 is whether AI investment is converting into return. McKinsey finds that 94% of enterprises had not yet seen significant value from AI spend as of the end of 2025. For any AI asset valued above 10x revenue, the market now expects auditable evidence that pilots are converting into paid, renewing revenue, that customer ROI is quantified, and that revenue is durable across at least three to four quarters rather than a one-time adoption spike.

What is driving AI valuations this quarter?

Valuations in Q2 2026 reflect an interplay of a genuine demand supercycle and the compressive realities of an unproven monetisation case and rising bubble-risk caution. Windsor Drake estimates the net effect has lifted the AI-native software benchmark from roughly 9x at the 2024 baseline to about 11x, a net expansion of 2.0x.

Table 4. Valuation Drivers, Expansion versus Compression, Q2 2026

Factor	Driver	Effect on Multiples	Notable Examples
Expansion	AI demand supercycle	+2.0x on AI-native software	Hyperscaler capex above $650B
Expansion	Public IPO re-rating	+1.0x as listings price the category	Cerebras, CoreWeave
Expansion	Rate normalisation	+0.5x as discount rates ease	Lower cost of long-duration capital
Compression	Monetisation gap	-0.8x on unproven enterprise ROI	94% not yet seeing value
Compression	Bubble-risk caution	-0.7x as concentration risk is priced	IMF dot-com comparison

Source: Windsor Drake analysis of McKinsey, Bain & Company and Federal Reserve data.

The macro backdrop is constructive but no longer a pure tailwind. The Federal Reserve funds range holds at **3.50% to 3.75%**, and the most recent Summary of Economic Projections points to roughly one further cut in 2026, easing the cost of capital for long-duration assets. Against that, the International Monetary Fund has compared current technology pricing to the dot-com peak, noting that the largest technology names now represent close to a fifth of the MSCI World Index, so a repricing would be felt market-wide.

Geographic variation

AI investment is concentrated to a degree unusual even by technology-sector standards. North America captured roughly **62%** of global AI venture investment in Q1 2026, on the strength of frontier labs, hyperscaler capital and the deepest exit markets. Europe trades at a capability and scale discount despite genuine research depth, and APAC is scaling through national model programmes and large domestic compute build-outs. The result is a clear arbitrage for well-capitalised US acquirers.

Table 5. Geographic Variation in AI Investment, Q2 2026

Region	Share of AI Investment	Posture	Key Drivers
North America	~62%	Premium	Frontier labs, hyperscaler capital, deep exit markets
APAC	~20%	Scaling	National model programmes, domestic compute build-outs
Europe	~13%	Discount	Strong research depth offset by a scale and capital gap
Rest of World	~5%	Selective	Sovereign compute initiatives and applied-AI niches

Source: Windsor Drake analysis of PitchBook and CB Insights data.

Public and private markets

A defining feature of the quarter is an incomplete convergence between public and private markets. Public AI-exposed software has re-rated from roughly 7x EV/Revenue in 2023 to about **11.5x** as the 2026 IPO class validated the category. Private AI companies, however, still average near **24x** revenue, a premium concentrated in frontier labs priced on optionality. Unlike prior technology cycles, the private premium has not yet closed, and public listings are pulling the public set upward toward it rather than the reverse.

Which valuation metric should apply?

Selecting the right metric is what separates a professional AI valuation from a careless one. Position in the AI stack, profitability and capital intensity each demand a different lens, and a generic EV/Revenue multiple applied across the sector will badly misprice both frontier labs and capital-heavy infrastructure.

EV/Revenue for AI-native software

EV/Revenue, often on a forward basis given the pace of growth, is the default lens for AI-native application software reinvesting ahead of profitability. The essential discipline is to anchor to the correct subsegment cohort and to let the Rule of 40 score dictate where in the 8x to 15x range an asset sits.

Forward revenue and optionality for frontier labs

Foundation-model labs cannot be valued on a conventional multiple. Their 15x to 30x marks are priced on strategic position: compute access, talent, capital and the probability of winning the model layer. Any valuation

should be framed explicitly as an optionality assessment rather than presented as a revenue multiple comparable to the application layer.

EV/EBITDA and capacity-based methods

EV/EBITDA applies only to the rare AI businesses with genuine positive EBITDA, including some diversified semiconductor and systems names. AI infrastructure is better assessed on contracted capacity, utilisation and order book, with capital intensity and depreciation capping the multiple below the software layer. Most AI companies remain pre-profit and are valued on revenue or forward revenue.

Table 6. AI Valuation Methodology Matrix, Q2 2026

Subsegment	Primary Metric	Typical 2026 Range	Key Adjustment
AI-Native Software	EV/Revenue	8x - 15x revenue	Rule of 40, retention
Foundation Models & Frontier Labs	Forward revenue / optionality	15x - 30x revenue	Compute access, strategic position
Agentic AI Platforms	EV/Revenue	12x - 25x revenue	Demonstrated ROI, model dependency
AI Infrastructure & Compute	Capacity and contract value	10x - 22x revenue	Utilisation, contracted demand
AI Semiconductors & Hardware	EV/Revenue or EV/EBITDA	12x - 22x revenue	Design wins, order book
AI Data, Tooling & Security	EV/Revenue	7x - 16x revenue	Recurring revenue, data moat

Source: Windsor Drake valuation methodology, calibrated to PitchBook and CB Insights comparables.

Key takeaways for founders

Translating the market picture into strategy means concentrating on six areas that consistently move AI valuations in the current environment.

1. Prove revenue, not narrative

An AI label alone no longer earns a premium. The market prices demonstrated, durable, recurring revenue, and AI-native software clears the roughly 11x benchmark only on that basis. Lead with paying customers, retention and gross margin rather than pilots and roadmaps.

2. Treat the Rule of 40 as the gate

For AI-native software with real revenue, the Rule of 40 remains the cleanest predictor of a top multiple. Only an estimated 15% to 20% of the cohort clears it, and those that do earn 50% to 100% premiums. Make the score a board-level priority with monthly tracking.

3. Close the monetisation gap

With 94% of enterprises not yet seeing significant value from AI spend, diligence is increasingly built around proof of return. Bring auditable ROI data, named reference customers and documented renewals, and be ready to show revenue durability across several quarters.

4. Understand your capital intensity

Compute-heavy models can command scarcity premiums, with AI infrastructure trading at 10x to 22x revenue, but they also carry funding and obsolescence risk that buyers price in. Model a credible path through the next two funding cycles before running a process.

5. Build a moat that survives commoditisation

Base models are commoditising and deflating in price. Premium valuations accrue to defensibility that survives that pressure: proprietary data, workflow ownership, secured compute and power, and enterprise distribution and trust. Frame the moat, not the model.

6. Prepare before the window moves

Public markets have reopened, but selectively, and demand, capital and pricing will not stay aligned indefinitely. A full process runs 12 to 18 months end to end, so a founder who intends to meet the market while today's alignment still holds is, in practice, preparing in the current cycle.

Sources

- McKinsey & Company, The Economic Potential of Generative AI
- McKinsey & Company, Where AI Will Create Value, and Where It Won't
- McKinsey & Company, Global Private Markets Report 2026
- Bain & Company, Global Technology Report
- Bain & Company, Global Private Equity Report 2026
- Goldman Sachs, 2026 Global M&A Outlook
- Morgan Stanley, Global M&A Activity Outlook 2026
- PitchBook, Q1 2026 AI VC Trends and Venture Monitor
- CB Insights, State of Venture and State of AI
- S&P Global Market Intelligence, private-capital and dry-powder analyses

- International Monetary Fund, World Economic Outlook Update
- Federal Reserve, FOMC statement and Summary of Economic Projections
- EY-Parthenon, M&A Outlook 2026
- Cerebras Systems Inc., Form S-1 and IPO prospectus