

AI Valuations: Q3 2026

Most AI valuations quoted in 2026 are not multiples. Windsor Drake's 36-company public AI comparable set, computed on 21 August 2026, carries a median **15.0x EV/Revenue**, an interquartile range of **9.3x to 23.8x** and a full range of **0.6x to 68.7x**. Palantir clears 68.7x and Super Micro Computer clears 0.6x on the same day. Alongside that priced market sits a larger private one: PitchBook counts roughly 690 AI companies valued above \$1 billion, and nearly half of all unicorns have not raised in two years.

The mechanism separating them is this report's organising idea. A valuation becomes a multiple only when both halves are verifiable: a numerator that is a price at which shares actually change hands, and a denominator that is a revenue figure covering a stated period and signed by an auditor. Public equities carry both. A private financing carries a real numerator and, at best, an asserted run rate. A tender price carries neither.

The evidence is sharp. Seven AI primary financings since November 2025 published a contemporaneous run rate, making an implied mark computable: OpenAI at 35.5x, ElevenLabs at 33.3x, Anysphere at 29.3x, Databricks at 27.1x, Lovable at 26.6x, Anthropic at 20.5x and Suno at 18.0x. The median is **27.1x**, some **81% above** the priced public median, on a basis that flatters the private company because a run rate counts one month against twelve filed ones. Larger rounds published no revenue metric at all.

This report covers what listed AI companies trade at by layer, what predicts the multiple, how contracted revenue reorders the cloud layer, what private valuations are worth as precedent, and who is buying. The framework is **The Denominator Test**: classify the valuation before relying on it, because a buyer underwrites the priced category and treats the other two as atmosphere.

What multiples do AI companies trade at in 2026?

The table is computed rather than licensed. Enterprise value equals market capitalisation plus total debt less cash and short-term investments, using closing prices of 21 August 2026, the most recent SEC cover-page share count and trailing revenue from the latest reported quarter.

Table 1. AI Valuation Multiples by Sub-Segment, Q3 2026

Sub-segment	Constituents	Median	Interquartile range
AI cloud and compute hosting	6	33.2x	19.5x to 55.2x
AI silicon and silicon IP	8	23.9x	20.1x to 33.4x
Data centre and incumbent cloud	3	12.8x	7.8x to 12.9x
AI data and infrastructure software	7	12.4x	5.2x to 21.6x
AI application software	8	11.1x	8.1x to 15.0x
AI manufacturing and systems	4	10.4x	6.8x to 12.6x
Composite	36	15.0x	9.3x to 23.8x

Source: Windsor Drake computation of EV/Revenue from Nasdaq and NYSE closing prices of 21 August 2026 and SEC filings.

The ordering is not the one gross margin would produce. Software is the highest-margin layer here and carries the two lowest medians; the capital-intensive cloud layer carries the highest. That inversion is the denominator asserting itself. Applied Digital holds roughly \$36 billion of contracted lease revenue against \$611 million of fiscal 2026 revenue.

The spread inside AI exceeds the spread between most sectors

Palantir clears 68.7x, TeraWulf 62.9x, Arm Holdings 49.5x, Nebius 44.3x, Cloudflare 41.2x, Astera Labs 40.1x, Snowflake 22.8x, NVIDIA 20.3x, Taiwan Semiconductor 14.7x, MongoDB 12.4x, Elastic 4.7x, C3.ai 4.1x and Super Micro Computer 0.6x. All are priced as artificial intelligence companies.

Why does Palantir trade at 68.7x while C3.ai trades at 4.1x?

Margin structure, specifically gross margin after inference and inspection cost. Across the 16 constituents disclosing a GAAP gross margin, gross margin correlates with the natural logarithm of EV/Revenue at $r = 0.82$, far stronger than anything the AI label explains.

Table 2. Gross Margin Against Multiple, Selected Constituents

Company	GAAP gross margin	Latest quarterly growth	EV/Revenue
Arm Holdings	97.2%	+22.4%	49.5x
AppLovin	88.3%	+52.8%	15.0x
Palantir	84.7%	+92.8%	68.7x
NVIDIA	74.9%	+85.2%	20.3x
Astera Labs	73.3%	+104.5%	40.1x
C3.ai	30.9% FY2026	-52.5%	4.1x
Super Micro Computer	10.8% FY2026	+93.2%	0.6x

Source: Company earnings releases and SEC filings for the most recently reported fiscal quarter; Windsor Drake computation of EV/Revenue at 21 August 2026 closing prices.

Super Micro Computer is the clearest case. It grew revenue 93.2% year on year and generated \$39.1 billion across fiscal 2026, more than Palantir, Snowflake, Datadog, MongoDB, Cloudflare, Elastic, GitLab and Braze combined, and clears 0.6x on a 10.8% gross margin. The market prices what a company keeps.

Growth counts, but only on top of margin

Palantir reported 93% growth, 84.7% gross margin and a 47.1% GAAP operating margin in the quarter to 30 June 2026. Across the 14 AI software constituents disclosing a GAAP operating margin, the two clearing the Rule of 40 carry a 41.9x median against 10.9x for the rest, a **3.8 times** premium on a cohort median score of 14.8.

The cost curve sits underneath all of it

Stanford HAI measured the cost of a GPT-3.5-capability query falling from \$20.00 to \$0.07 per million tokens between November 2022 and October 2024, and Goldman Sachs put inference cost declines at 60% to 70% per year in May 2026. The frontier moved differently: Anthropic's top Opus tier fell from \$15 to \$5 per million input tokens over the year to August 2026 while a new tier opened above it, OpenAI's flagship input price rose from \$1.25 to \$2.00, and Google discloses a doubling of Gemini Flash pricing on 1 January 2027.

How should a contract-backed AI business be valued?

The AI cloud layer is where The Denominator Test bites hardest, because two defensible denominators exist and they order the same companies differently.

Table 3. Trailing Revenue Against Contracted Revenue, AI Cloud and Incumbents

Company	EV / trailing revenue	Contracted revenue and basis	EV / contracted
Nebius Group	44.3x	over \$40.0bn commitments, company-defined	1.50x
Applied Digital	18.6x	about \$36.0bn lease revenue, company-defined	0.32x
Equinix	12.8x	\$15.0bn ASC 606 RPO	8.33x
CoreWeave	10.3x	\$103.7bn ASC 606 RPO	0.75x
Oracle	7.8x	\$638.0bn ASC 606 RPO	0.83x

Source: Company 10-Q, 10-K and 6-K filings and earnings releases for the most recently reported period; Windsor Drake computation at 21 August 2026 closing prices.

On trailing revenue Nebius looks more than four times as expensive as CoreWeave; on contracted revenue the gap narrows to roughly two times, and Equinix moves from mid-table to dearest at 8.33x. Coverage explains it: CoreWeave holds 13.7 times trailing revenue in remaining performance obligation, Nebius 29.5 times, Applied Digital 58.9 times and Equinix 1.5 times. Basis choice does not adjust the answer. It reorders it.

Three bases, each labelled

CoreWeave, Equinix and Oracle publish an ASC 606 remaining performance obligation. Nebius publishes customer commitments, a company-defined figure it does not reconcile to ASC 606. Applied Digital publishes contracted lease revenue over 15-year terms. Choose the basis deliberately and hold it, because a buyer given a choice will take the one that suits them.

Counterparty credit is now a valuation input

CoreWeave holds about \$22.4 billion cumulative with OpenAI and about \$21 billion with Meta through December 2032; Nebius \$17.4 billion with Microsoft and up to \$27 billion with Meta; IREN about \$9.7 billion with Microsoft including a 20% prepayment. Against that, CoreWeave carries \$35.3 billion of debt and \$16.3 billion of operating lease liabilities on \$7.6 billion of trailing revenue, and the Bank of England flags a maturity mismatch between short chip lives and the debt financing them.

What is a private AI valuation worth as a precedent?

Less than a founder hopes, more than nothing. The rounds below are the only large AI primary financings since November 2025 that published a contemporaneous run rate. Every revenue figure is a company assertion, and none is audited.

Table 4. Implied Marks on Company-Asserted Run Rates, AI Primary Financings

Company	Post-money	Date	Run rate disclosed with the round	Implied mark
OpenAI	\$852bn	31 Mar 2026	\$2bn per month	35.5x
ElevenLabs	\$11.0bn	4 Feb 2026	over \$330m ARR	33.3x
Anysphere (Cursor)	\$29.3bn	13 Nov 2025	over \$1bn annualised	29.3x
Databricks	\$190bn	13 Aug 2026	\$7bn run rate	27.1x
Lovable	\$13.3bn	12 Aug 2026	\$500m annualised	26.6x
Anthropic	\$965bn	28 May 2026	\$47bn run rate	20.5x
Suno	\$5.4bn	3 Jun 2026	about \$300m annual revenue	18.0x
Median implied mark			Company-asserted, unaudited	27.1x

Source: Company financing announcements and named executive statements to Bloomberg; Windsor Drake computation. Every run rate above is company-asserted and unaudited.

The 81% gap understates the difference, because the two are not computed on the same basis. A buyer's committee underwrites the 15.0x distribution and treats 27.1x as market colour.

Marks move on events that are not revenue events

Groq raised \$750 million at a \$6.9 billion post-money on 17 September 2025. In December 2025 NVIDIA licensed Groq's language processing unit technology non-exclusively and hired certain employees, recording \$17.0 billion of consideration of which \$14.4 billion was goodwill attributable primarily to the workforce, and stating that no customer contracts, products or equity interests were purchased. On 17 August 2026 Groq raised at \$3.5 billion, half its prior mark.

Who is buying AI companies, and what are they paying?

Operators, and rarely in a form that produces a public multiple. The Windsor Drake Exit Index held 522 records when queried on 24 August 2026 and 516 distinct transactions after deduplicating on acquirer and target and then on target and announcement date. Sixty-six name AI, machine learning or model capability, 45 of them announced in 2025 or 2026. **Sixty-five of the 66 went to a strategic operator and one to a financial sponsor.**

The capability sits close to a workflow rather than in the model layer: identity and fraud, capital markets technology, regulatory technology and B2B payments account for 25 of the 66. Named acquirers include Bloomberg, Stripe, Airwallex, Backbase, Coupa, Duck Creek, SoFi and OpenAI itself. Fourteen disclose consideration, so the comparable set a seller builds is the benchmark that frames the negotiation.

Two transactions, one usable benchmark

SpaceX agreed on 16 June 2026 to acquire Anysphere for \$60.0 billion in stock, closing 14 August. It is the largest AI acquisition on record and filed no target financials. IBM acquired Confluent for \$31.00 per share in cash, closing 17 March 2026 at a purchase price of \$11,590 million against \$982 million of net cash acquired, on fiscal 2025 revenue of \$1,166.7 million: **9.1x** trailing revenue, entirely from filings. The Confluent multiple will be quoted for years; the \$60 billion will be quoted as atmosphere.

Licence and hire is now a standard alternative to acquisition

Across NVIDIA and Groq at \$17.0 billion, Meta and Scale AI at \$14.3 billion for a 49% non-voting stake, and Google's Windsurf arrangement, more than \$40 billion moved through licences and minority stakes. None appears in an M&A league table.

What will move AI valuations over the next four quarters?

Table 5. The Variables and Where They Stand

Variable	Current measure	As of	Institution
Five-hyperscaler capex, 2026 to 2027	about \$805bn rising to \$1.1trn; Alphabet, Amazon and Meta raised 2026 guidance	3 May 2026	Morgan Stanley
Global data centre electricity use	about 415 TWh 2024 to about 945 TWh 2030	Apr 2025	International Energy Agency
Texas 2027 demand growth	cut from 14% to 6%	11 Aug 2026	US Energy Information Administration
Worldwide AI spending	\$2.60trn in 2026, up 47%	19 May 2026	Gartner
Enterprises with EBIT impact above 5%	about 6%, against 88% adopting	5 Nov 2025	McKinsey & Company
AI share of S&P 500 market cap	about half, from a quarter in 2022	7 Jul 2026	Bank of England
Private credit to AI firms	over \$200bn, almost 8% of loan volumes	7 Jan 2026	Bank for International Settlements

Source: As attributed per row.

Demand is committed; return is arriving slowly

BCG found companies planning to double AI spending in 2026 to about 1.7% of revenues, with 94% of chief executives investing regardless of payback inside a year. McKinsey found 88% of organisations using AI in a

function, 39% reporting any EBIT impact and about 6% above 5%. A BIS study of over 12,000 firms found adoption raises productivity levels by 4% with no adverse employment effect.

The constraint is shifting from capital to permission

The Texas governor paused new data centre interconnections on 3 August 2026, and the US Energy Information Administration then cut its Texas 2027 demand growth forecast from 14% to 6%. Capital for the buildout is abundant; sited, powered capacity is not. Permitted capacity is becoming the asset a buyer is acquiring.

The listing route is open for scale

Renaissance Capital counts 104 US IPOs raising \$145.8 billion year to date in 2026 against 202 pricings in all of 2025: pricings down 24.6%, proceeds up 543.7%. EY counts 12 US deals above \$1 billion in the first half against four a year earlier. Performance has been uneven: Cerebras priced at \$185 at roughly 78 times trailing revenue and trades at \$188.70, some 51% below its post-listing high, while Figma priced at \$33 and trades at \$27.34.

Table 6. The 36 Constituents by Sub-Segment

Sub-segment	Constituents
AI silicon and silicon IP	NVIDIA, Advanced Micro Devices, Broadcom, Marvell Technology, Astera Labs, Credo Technology, Arm Holdings, Monolithic Power Systems
AI manufacturing and systems	Taiwan Semiconductor, Micron Technology, Vertiv, Super Micro Computer
AI cloud and compute hosting	CoreWeave, Nebius Group, IREN, Applied Digital, Cipher Mining, TeraWulf
Data centre and incumbent cloud	Equinix, Digital Realty, Oracle
AI data and infrastructure software	Snowflake, Datadog, MongoDB, Elastic, Cloudflare, GitLab, Braze
AI application software	Palantir, AppLovin, Samsara, Tempus AI, SoundHound AI, Innodata, BigBear.ai, C3.ai

Source: Windsor Drake comparable set, computed 21 August 2026. Confluent appears only as a transaction, acquired by IBM on 17 March 2026.

Key takeaways for founders

1. Classify the benchmark before you build a case on it

A priced public multiple, a private round mark and a headline valuation are three different objects, and a committee underwrites the first. The priced AI median is 15.0x; the median private mark is 27.1x.

2. Gross margin is the lever, and it is a cost of goods sold question

Gross margin correlates with the log of the multiple at $r = 0.82$. A 10.8% gross margin clears 0.6x and an 84.7% gross margin clears 68.7x. Pricing architecture that decouples revenue from token consumption is a valuation decision.

3. Choose the denominator deliberately and hold it

Nebius clears 44.3x on trailing revenue and 1.50x on contracted revenue. Decide which basis represents the business, state it, and use it consistently. A buyer given a choice will select the lower number.

4. Build the buyer list around operators

Sixty-five of 66 AI-capability transactions went to a strategic buyer. Rank prospects by demonstrated appetite for the specific capability, understand the acquirer's own multiple because it constrains what they can pay in stock, and prepare for licence-and-hire structures.

5. Denominator work takes longer than a process does

Separating contracted from usage revenue, isolating inference cost as a traceable cost line and establishing a trailing revenue series an auditor would recognise takes twelve to eighteen months. A process runs a further six to nine.

6. The conditions favouring a capability seller are cyclical

Operator appetite is a function of a build-versus-buy calculus that currently favours buying. At the same time the Bank of England puts AI companies at about half of S&P 500 market capitalisation and states the risk of a sharp correction remains high, and 50% of Federal Reserve survey contacts cited AI risk in spring 2026 against none the previous autumn.

Sources

- [US Securities and Exchange Commission, EDGAR filings](#)
- [Nasdaq and NYSE, closing prices, 21 August 2026](#)
- [Bank for International Settlements, Bulletin 120, Annual Economic Report 2026 and Working Paper 1325](#)
- [Bank of England, Federal Reserve and European Central Bank financial stability reports, 2026](#)
- [International Energy Agency, Energy and AI; US Energy Information Administration, Short-Term Energy Outlook](#)
- [Stanford HAI, AI Index Report](#)
- [Gartner Worldwide AI Spending Forecast; McKinsey & Company, The State of AI](#)
- [Bain Global M&A Report 2026, PwC M&A Trends Mid-Year 2026 and EY Global IPO Trends Q2 2026](#)
- [PitchBook and NVCA Venture Monitor and Unicorn Tracker; CB Insights State of AI](#)

- [Goldman Sachs, Morgan Stanley and J.P. Morgan Asset Management AI research, 2026](#)
- [US Census Bureau Business Trends and Outlook Survey and Eurostat enterprise AI use](#)
- [Boston Consulting Group AI Radar 2026, Deloitte State of AI in the Enterprise 2026 and KPMG Global AI Pulse](#)
- [Renaissance Capital US IPO statistics; Bloomberg and Reuters reporting on AI financings](#)

Methodology note

Framework: The Denominator Test. Valuations are classified into three categories and every figure carries its classification. A **priced multiple** requires a transactable price and a revenue figure from filed accounts; 36 listed companies qualify. A **private round mark** requires a primary financing that priced a share class, and yields an implied mark only where a contemporaneous run rate was published. A **headline valuation** derives from a tender or press reporting.

Source standard. Every figure is either computed by Windsor Drake from primary market data or attributed to a named institution with a date. Primary sources are SEC filings, official exchange data, central banks and regulators, statistical agencies and the International Energy Agency. Institutional research comes from the major consultancies, bulge-bracket banks and elite data houses.

How the figures were computed. Enterprise value equals market capitalisation plus total debt less cash and short-term investments, using closing prices of 21 August 2026 and the most recent SEC cover-page share count. Balance-sheet and trailing revenue come from each company's latest reported quarter, so fiscal conventions differ across constituents and are stated in the deck. Taiwan Semiconductor's balance sheet is converted from New Taiwan dollars at the rate implied by its own second quarter revenue translation.

Derived estimates and attribution. Four figures are derived and stated as such: the gross margin correlation, across the 16 constituents disclosing a GAAP gross margin; implied private marks, dividing a disclosed post-money by a run rate published with the same round; Rule of 40 scores, adding quarterly growth to GAAP operating margin across 14 constituents; and the IBM and Confluent multiple. The Windsor Drake Exit Index, queried live on 24 August 2026, supplies acquirer identity, acquirer type mix and repeat buyers, counted after the deduplication stated above and on disclosed consideration only.