AI Valuations: Q3 2026

AUGUST 2026

Windsor Drake · Market Intelligence

Executive Summary: Most AI Valuations Are Not Multiples

WINDSOR DRAKE

The Priced Market

The 36-company public AI comparable set carries a median **15.0x** EV/Revenue on 21 August 2026.

- The interquartile range runs **9.3x to 23.8x**.
- The full observed range runs **0.6x to 68.7x**.
- Palantir clears 68.7x and Super Micro clears 0.6x on the same day.
- Six sub-segments, all constituents named.
- Computed from exchange closing prices and filed accounts.

The Denominator Test

A valuation becomes a multiple only when both halves are verifiable: a transactable price and an audited revenue figure.

- **Priced multiples** carry both. 36 companies qualify.
- **Private round marks** price a share class against a run rate the company asserts.
- **Headline valuations** carry neither: tenders and secondaries.
- Buyers underwrite the first. The other two set the mood.
- Every private figure in this deck carries its classification.

The Private Premium Is 81%

Seven AI primary rounds since November 2025 published a run rate. Their implied marks carry a median **27.1x**.

- OpenAI at **35.5x** on $2bn per month, 31 March 2026.
- Anthropic at **20.5x** on a $47bn run rate, 28 May 2026.
- Databricks at **27.1x** on a $7bn run rate, 13 August 2026.
- A run rate counts one month; trailing revenue counts twelve.
- Anysphere at **29.3x**, ElevenLabs at **33.3x**, Suno at **18.0x**.

Margin Structure Sets the Price

Gross margin explains most of the variation in multiple. The AI label explains almost none of it.

- Gross margin correlates with log EV/Revenue at **r = 0.82** across 16 constituents.
- Palantir: 84.7% gross margin, 68.7x.
- Super Micro: 10.8% fiscal 2026 gross margin, 0.6x.
- Only one of them keeps the margin.
- Arm at 97.2% gross margin clears 49.5x on royalty revenue.

Source: Windsor Drake computation from exchange closing prices of 21 August 2026 and SEC filings. Constituents named in the appendix.

Key Takeaways for AI Founders

WINDSOR DRAKE

1. Apply the Denominator Test Before You Anchor

The number a founder reads in the press is usually a private round mark or a tender price. A buyer underwrites a priced public multiple. Establish which category your benchmark comes from before you build a case on it.

- Priced public AI companies clear a **15.0x** median.
- Private primary rounds with a published run rate imply **27.1x** on that run rate.

2. Gross Margin Is the Lever, Not the Label

Gross margin correlates with the log of the multiple at r = 0.82 across the disclosing constituents. Compute and inspection cost sit directly on that line.

- A 10.8% gross margin clears 0.6x; an 84.7% gross margin clears 68.7x.
- Model costs are a cost of goods sold question before they are a technology question.

3. Contracted Revenue Changes the Ranking

In the AI cloud layer, trailing revenue and contracted revenue produce different orderings of the same companies. Choose the denominator deliberately, and disclose which one you used.

- Nebius clears 44.3x on trailing revenue and 1.50x on disclosed customer commitments.
- CoreWeave clears 10.3x on trailing revenue and 0.75x on its ASC 606 remaining performance obligation.

4. Your Buyer Is an Operator, Not a Sponsor

Of the 66 AI-capability transactions in the Windsor Drake Exit Index, 65 were acquired by a strategic operator and one by a financial sponsor. Build the buyer list accordingly.

- 45 of the 66 were announced in 2025 or 2026.
- Bloomberg entities bought Regology and Canoe Intelligence within eight weeks.

5. In an Opaque Market, Evidence Is Leverage

Fourteen of those 66 transactions disclose consideration, so the comparable set a seller builds is the benchmark that frames the negotiation.

- The buyer prices against private processes the seller cannot see.
- A reproducible public comp set with named constituents removes that asymmetry.

6. The Capital Backdrop Is Constructive and Cyclical

Enterprise AI spending is forecast at $2.60 trillion in 2026, up 47% (Gartner, May 2026), while official-sector bodies are flagging concentration and financing risk in the same market.

- AI companies are about half of S&P 500 market capitalisation, up from about a quarter in 2022 (Bank of England, July 2026).
- Conditions this favourable to a capability seller have historically compressed rather than persisted.

Source: Windsor Drake computation from exchange closing prices of 21 August 2026 and SEC filings. Constituents named in the appendix.

Founder FAQs

WINDSOR DRAKE

1 What multiple do AI companies trade at in 2026?

Windsor Drake's 36-company public AI comparable set carries a median 15.0x EV/Revenue as of 21 August 2026, with an interquartile range of 9.3x to 23.8x and a full range of 0.6x to 68.7x. The composite is wide enough that the layer of the stack a company occupies matters far more than the AI label attached to it.

2 Is a private AI valuation the same thing as a multiple?

No. A private round prices a share class; it rarely publishes an audited revenue figure. Where a company did publish a contemporaneous run rate, the implied mark can be computed: seven such rounds since November 2025 imply a median 27.1x on that run rate, against 15.0x for priced public companies. A tender or secondary price is a third category again, because no share class was priced by a primary financing.

3 Which AI sub-segment carries the highest multiple?

AI cloud and compute hosting, at a 33.2x median across six constituents, ahead of AI silicon and silicon IP at 23.9x across eight. Both sit well above AI application software at 11.1x. The ordering is the opposite of what gross margin alone would predict, because the cloud layer is priced against contracted revenue that trailing revenue has not yet caught.

4 Why does Palantir trade at 68.7x and C3.ai at 4.1x?

Margin structure and growth durability. Palantir reported 84.7% GAAP gross margin, 47.1% GAAP operating margin and 93% revenue growth in the quarter to 30 June 2026. C3.ai reported 30.9% gross margin for fiscal 2026 and revenue down 52.5% in its fourth quarter. Both are sold as enterprise AI.

5 What sets an AI company's gross margin?

Inference and inspection cost. The price of a given capability level has fallen sharply: Stanford HAI measured the cost of a GPT-3.5-equivalent query dropping from $20.00 to $0.07 per million tokens between November 2022 and October 2024. Frontier pricing has moved differently. Anthropic's top Opus tier fell from $15 to $5 per million input tokens over the year to August 2026 while a higher tier opened above it, and OpenAI's flagship input price rose from $1.25 to $2.00.

6 How many AI transactions disclose a valuation multiple?

Few. Fourteen of the 66 AI-capability transactions in the Windsor Drake Exit Index disclose consideration. The largest AI acquisition on record, SpaceX's $60 billion all-stock purchase of Anysphere, closed on 14 August 2026 with no target financials filed. Public comparables are the reproducible benchmark available to a seller.

7 Who is buying AI companies in 2026?

Operators. Sixty-five of the 66 AI-capability transactions in the Exit Index were acquired by a strategic buyer and one by a financial sponsor. Named acquirers in the trailing twelve months include Bloomberg, Stripe, Airwallex, SoFi, Coupa, Backbase, Duck Creek and OpenAI itself. IBM acquired Confluent for $11.3 billion in equity value, closing 17 March 2026, at roughly 9.1x fiscal 2025 revenue.

8 Is 2026 a good year to sell an AI company?

Enterprise AI spending is forecast at $2.60 trillion in 2026, up 47% (Gartner), and the listing window has reopened for scale issuers, with 104 US IPOs raising $145.8 billion year to date (Renaissance Capital). Official-sector bodies are simultaneously flagging concentration risk: the Bank of England puts AI companies at about half of S&P 500 market capitalisation. Whether it suits a specific asset depends on its margin structure and its readiness, not on the calendar.

Source: Windsor Drake computation from exchange closing prices of 21 August 2026 and SEC filings. Constituents named in the appendix.

The Denominator Test

WINDSOR DRAKE

A valuation headline becomes a multiple only when both halves are verifiable. Three categories, and only one of them is something a buyer will underwrite.

Category	The numerator	The denominator	What it is worth in a negotiation
Priced multiple	**A transactable price** A listed share price struck on an exchange, adjusted to enterprise value using filed debt and cash.	**Audited revenue** Trailing twelve month revenue from filed accounts, covering a stated period an auditor has signed.	**The benchmark** 36 AI companies qualify, median 15.0x. This is the evidence an acquirer's committee will accept.
Private round mark	**A primary financing** A share class priced by new money entering the company, named investors, disclosed post-money valuation.	**An asserted run rate** Management annualises a recent month. Unaudited, and the period covered is rarely stated.	**Directional evidence** Seven rounds since November 2025 imply a median 27.1x on that run rate. Useful context, weak precedent.
Headline valuation	**A tender or a report** An employee tender, a secondary sale, or a figure attributed to people familiar with the matter.	**Usually absent** Many of the largest 2026 rounds published a usage metric instead: monthly active users, tokens per month, lawyer counts.	**Atmosphere** OpenAI's $852bn August tender and Clay's $5bn January tender priced no new shares. Groq's mark halved between two primary rounds.
The founder's test	**Can it be struck today?** If the price cannot be transacted at by a third party on the date quoted, it is not a numerator a committee will accept.	**Has an auditor signed it?** If the revenue figure annualises a month, or arrives as bookings, tokens or users, the multiple is not comparable to a listed one.	**Bring the priced set** The 36-company comparable set in this deck is reproducible from exchange data and filings, which is why it survives diligence.

Sources: SEC filings; company financing announcements; Bloomberg; Reuters; PitchBook. Windsor Drake framework.

AI Valuation Multiples by Sub-Segment

WINDSOR DRAKE

EV/Revenue, interquartile range with median marked. Windsor Drake computation from closing prices of 21 August 2026 and SEC filings.

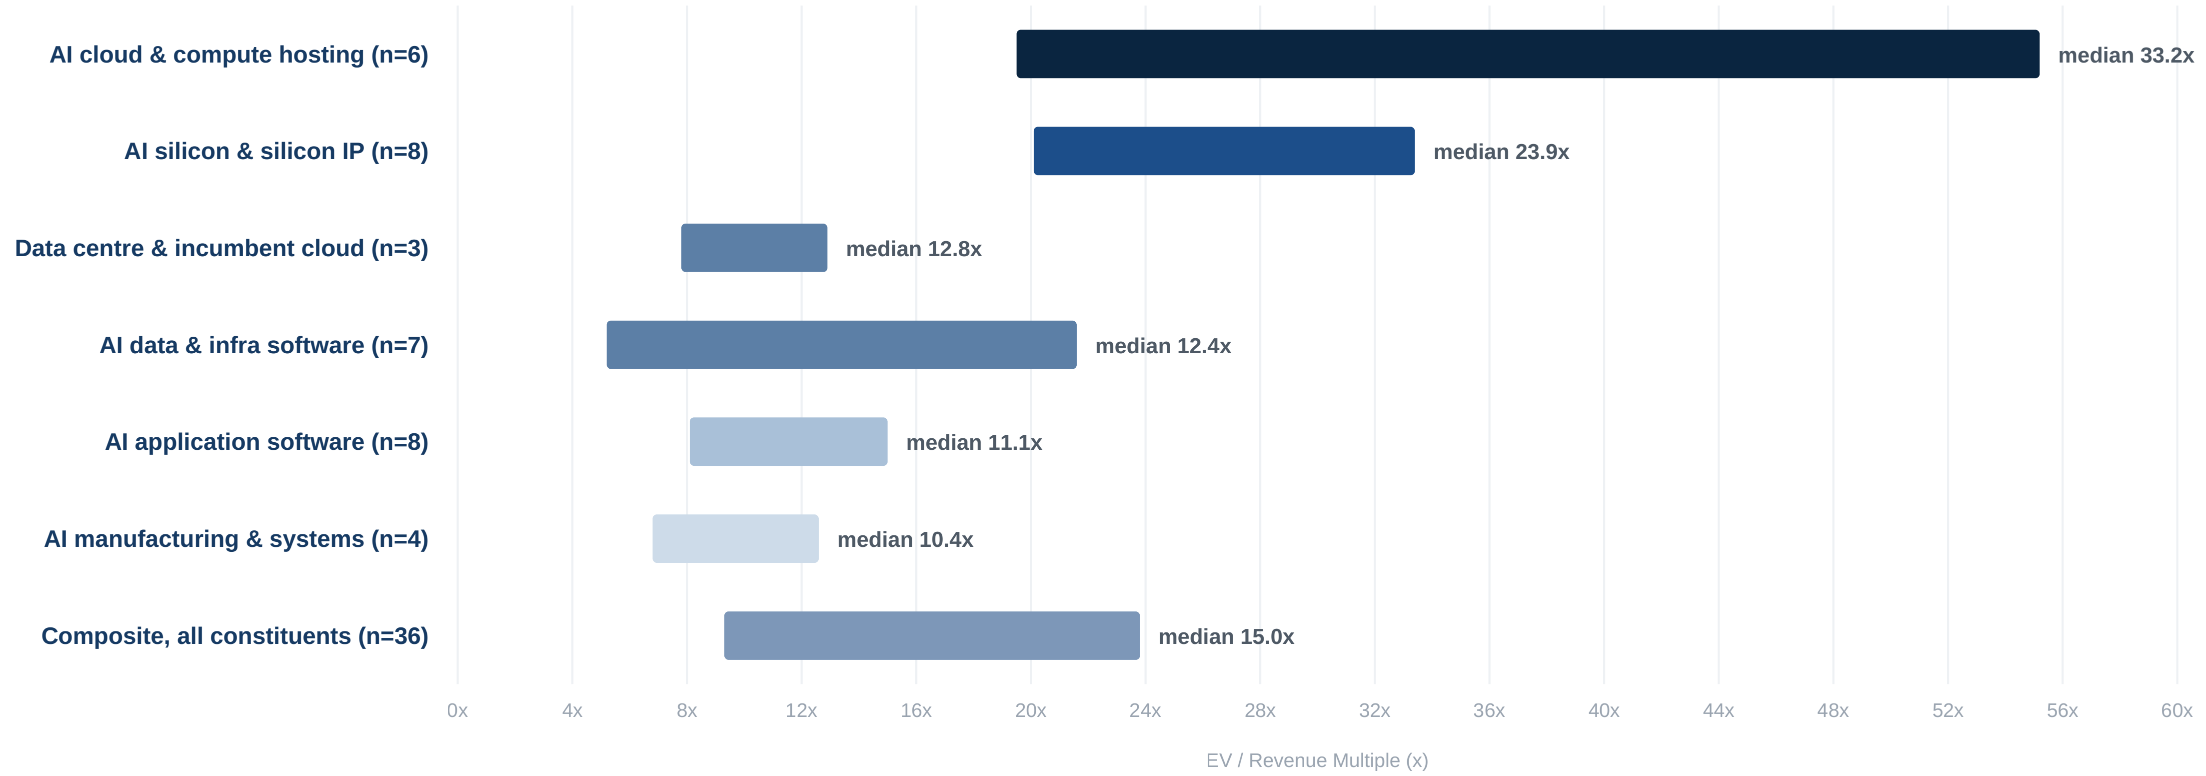

THE ORDERING IS NOT THE MARGIN ORDERING

Software is the highest-margin layer here and carries the lowest medians; the cloud layer carries the highest, because its multiple is computed against trailing revenue while its economics sit in contracted revenue not yet recognised. Applied Digital holds roughly **$36 billion** of contracted lease revenue against **$611 million** of fiscal 2026 revenue. Read the width as well as the level: AI application software carries the lowest median and the highest single constituent in the set.

Source: Windsor Drake computation from exchange closing prices of 21 August 2026 and SEC filings. Constituents named in the appendix.

The Spread Inside AI Is Larger Than Between Sectors

WINDSOR DRAKE

Individual constituents, EV/Revenue. One label, a hundred-fold range of outcomes.

EV / Revenue by company, 21 August 2026

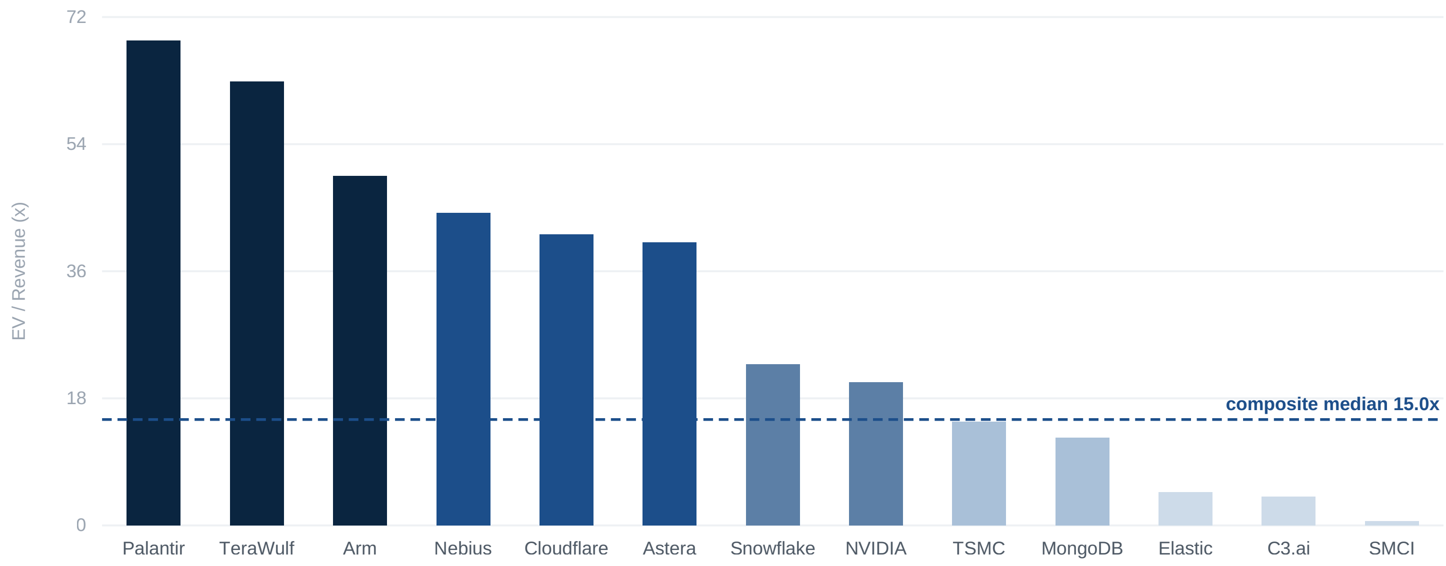

COMPOSITE MEDIAN

15.0x

Across 36 constituents with a filed revenue denominator.

INTERQUARTILE RANGE

9.3x to 23.8x

Half the set clears inside this band.

FULL RANGE

0.6x to 68.7x

Super Micro at the floor, Palantir at the ceiling.

GROSS MARGIN FIT

r = 0.82

Correlation of gross margin with log EV/Revenue, 16 disclosing constituents.

> Thirteen of the 36 constituents are plotted. An acquirer arrives with a comparable set chosen to support their number. So should a seller.

Source: Windsor Drake computation from exchange closing prices of 21 August 2026 and SEC filings. Constituents named in the appendix.

Priced Multiples Against Private Round Marks

WINDSOR DRAKE

Seven AI primary financings since November 2025 published a contemporaneous revenue run rate. Every implied mark sits above the priced public median.

Implied mark on company-asserted run rate, versus the priced public median

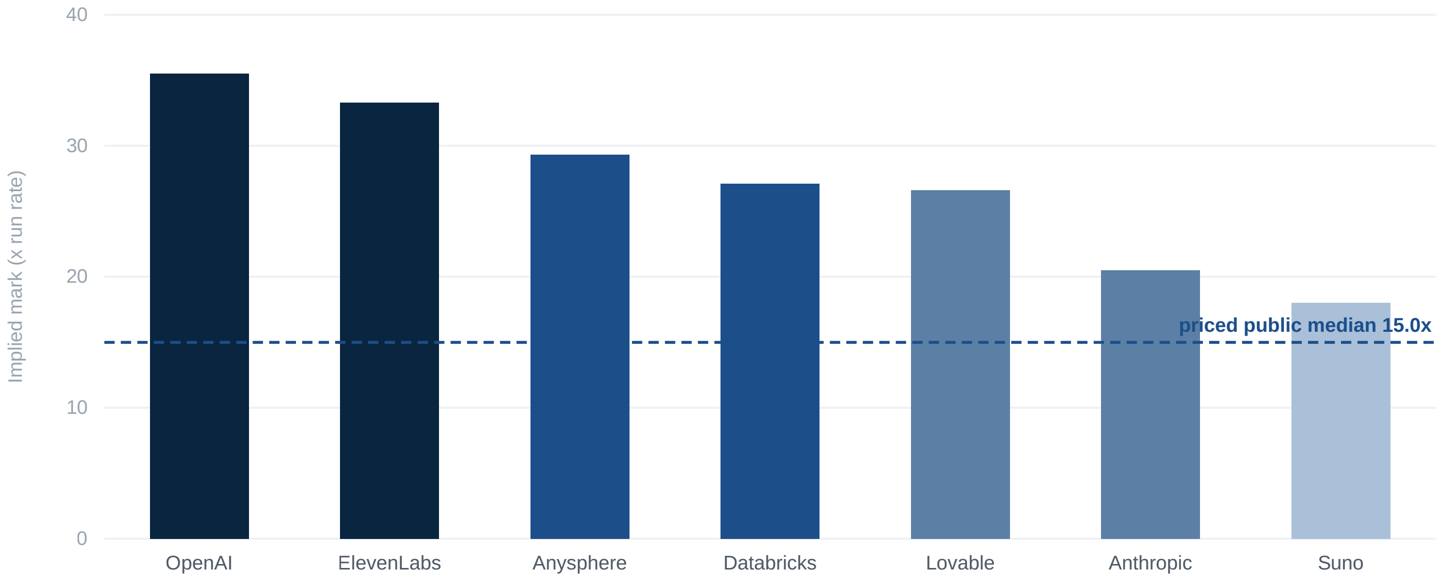

MEDIAN PRIVATE MARK

27.1x

On run rates the companies asserted in their own announcements.

PREMIUM TO PRICED

+81%

Against the 15.0x median of the public comparable set.

THE BASIS DIFFERENCE

1 month vs 12

A run rate annualises a recent month. Trailing revenue counts twelve filed months.

AI SHARE OF US VENTURE

86%

$355.9bn of $412.7bn deployed in H1 2026 (PitchBook-NVCA, July 2026).

> PitchBook counts roughly 690 AI companies valued above $1bn and notes that nearly half of all unicorns have not raised in more than two years. A mark no transaction has tested is the third category on the framework slide.

Sources: company financing announcements; PitchBook; CB Insights; Bloomberg; Reuters. See appendix.

AI Silicon: From Royalty to Assembly

WINDSOR DRAKE

The stack prices in a straight line from intellectual property down to physical integration. Same demand, eighty-fold spread.

EV / Revenue by role in the silicon stack (x)

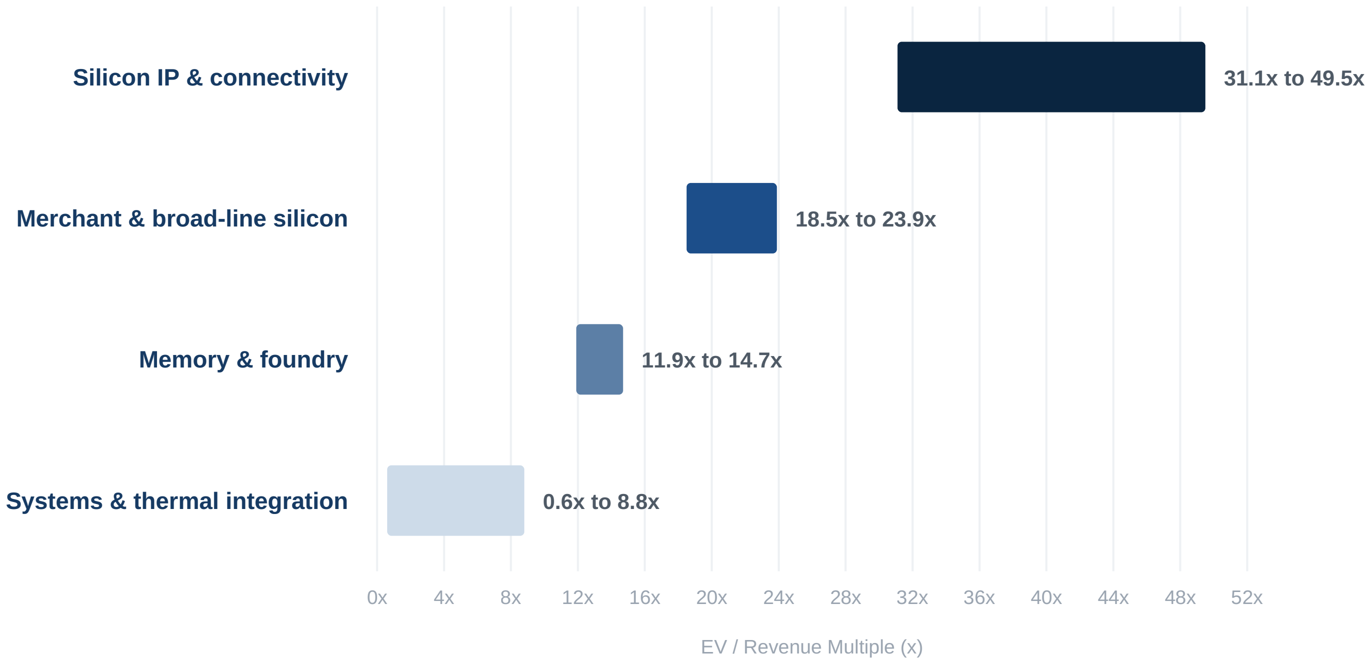

KEY OBSERVATION

The gradient from **49.5x** to **0.6x** runs almost exactly along gross margin. For a founder selling into this stack, the question a buyer will ask is which layer keeps the margin when compute prices fall, not which layer touches the most AI revenue.

Valuation Drivers

Royalty economics price highest

Arm reported 97.2% gross margin on $1,289m of quarterly revenue and clears 49.5x. Astera Labs at 73.3% gross margin and 104.5% growth clears 40.1x. The market pays for a revenue line that carries almost no cost of goods.

Scale does not lift the multiple

NVIDIA generated $253.5bn of trailing revenue at 74.9% gross margin and clears 20.3x, below Credo at 31.1x on $1.3bn. Being large in a merchant silicon position does not purchase a licensing multiple.

Assembly is the floor

Super Micro reported 10.8% fiscal 2026 gross margin on $39.1bn of revenue and clears 0.6x. Vertiv at 37.7% gross margin clears 8.8x. Proximity to the AI buildout does not confer AI economics.

Source: Windsor Drake computation from exchange closing prices of 21 August 2026 and SEC filings. Constituents named in the appendix.

AI Cloud: The Multiple Depends On Which Revenue You Count

WINDSOR DRAKE

Trailing revenue and contracted revenue produce different orderings of the same six companies. Both are computable, and each is labelled.

EV / trailing revenue by business model (x)

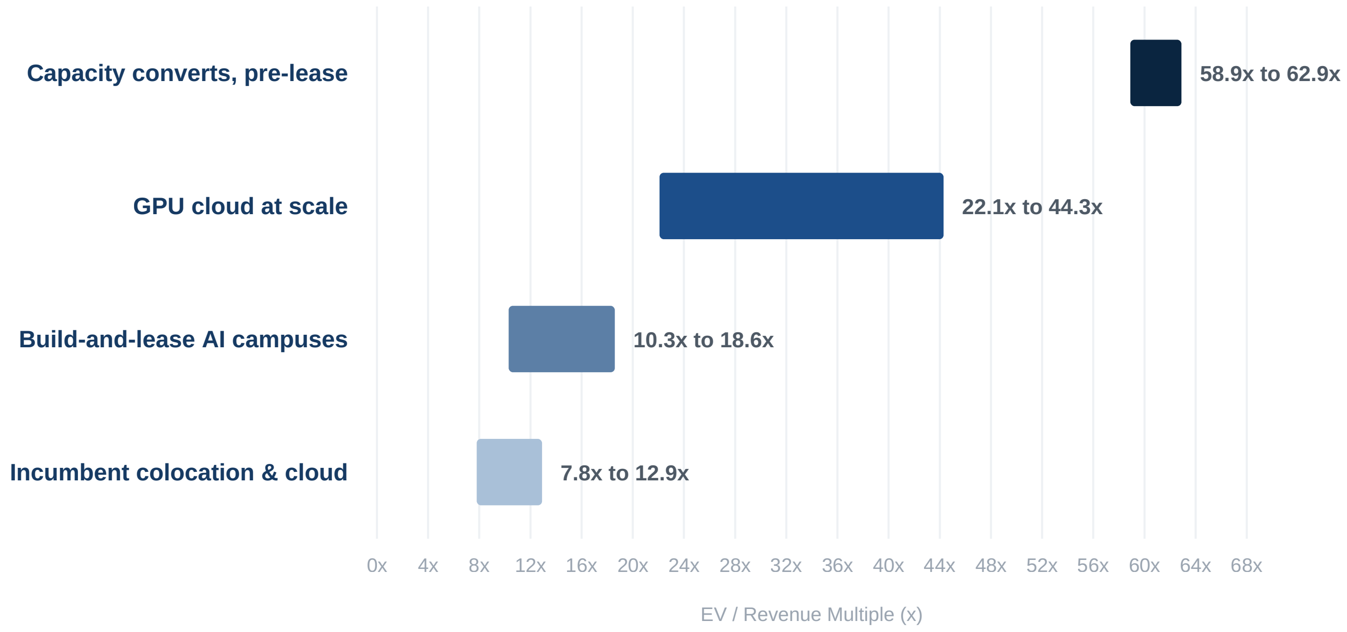

KEY OBSERVATION

Three different bases appear in this layer: ASC 606 remaining performance obligation, company-defined customer commitments and company-defined contracted lease revenue. Each is labelled in the report table. A seller in a contract-backed business should choose the basis deliberately and present it consistently, because the buyer will otherwise choose the one that suits them.

Valuation Drivers

Contracted revenue inverts the ranking

On trailing revenue Nebius clears 44.3x and CoreWeave 10.3x. On disclosed contracted revenue Nebius clears 1.50x and CoreWeave 0.75x. The gap narrows from four times to two, and Equinix moves from cheapest to dearest at 8.33x.

Coverage ratios are extraordinary

CoreWeave holds $103.7bn of ASC 606 remaining performance obligation against $7.59bn of trailing revenue, 13.7 times coverage. Applied Digital holds roughly $36bn of contracted lease revenue against $611m, 58.9 times. Oracle reports $638bn of RPO, up 363% year on year.

The leverage is the risk

CoreWeave carries $35.3bn of debt and $16.3bn of operating lease liabilities. The Bank of England notes that more than half of over $2 trillion of 2026 to 2028 data centre investment will be debt funded, against chip lives shorter than the debt.

Source: Windsor Drake computation from exchange closing prices of 21 August 2026 and SEC filings. Constituents named in the appendix.

AI Software: A Seventeen-Fold Range Under One Label

WINDSOR DRAKE

Fifteen listed companies, all marketed as AI software, clearing between 4.0x and 68.7x on the same date and the same metric.

EV / Revenue by software model (x)

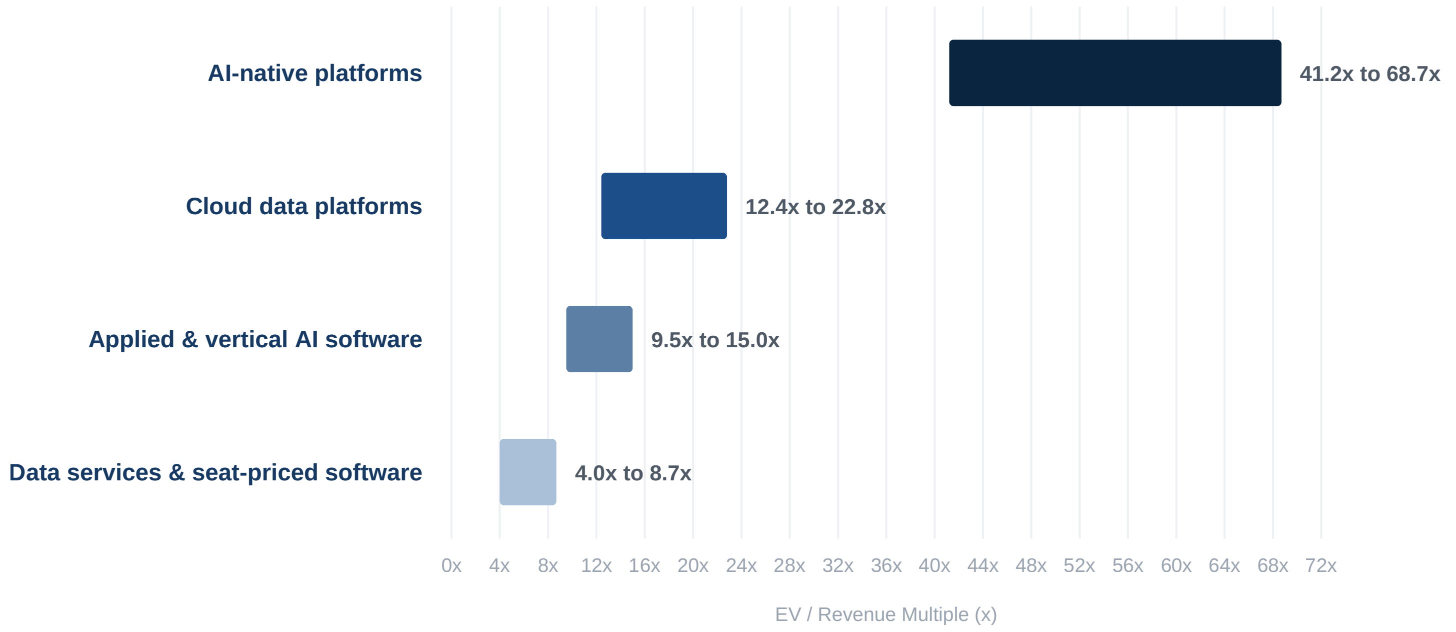

KEY OBSERVATION

The Rule of 40 deck covers the 2025 to 2026 repricing of software equities in detail, and the B2B SaaS Valuations Q3 2026 deck covers the committed-base mechanism. The finding relevant here is narrower: **the AI label moves no multiple on its own.** Every constituent on this chart carries it, and they span seventeen times.

Valuation Drivers

Growth plus margin, together

Palantir reported 93% revenue growth, 84.7% gross margin and 47.1% GAAP operating margin in the quarter to 30 June 2026, and clears 68.7x. AppLovin reported 53% growth and 77.7% GAAP operating margin and clears 15.0x. The market pays a further premium for growth that is still accelerating.

Seat pricing is being repriced

Braze at 4.0x, C3.ai at 4.1x, Elastic at 4.7x and GitLab at 5.7x sit at the floor. EY found in July 2026 that 82% of senior leaders expect per-seat software pricing to become less relevant within five years and 91% regard custom AI-built software as critical.

AI-native demand is measurable

Datadog disclosed that its AI-native customer cohort contributed a high single digit percentage point of total year on year revenue growth in the second quarter. Snowflake reported more than 13,600 accounts using its AI capabilities and 126% net revenue retention.

Source: Windsor Drake computation from exchange closing prices of 21 August 2026 and SEC filings. Constituents named in the appendix.

What Actually Sets an AI Multiple

WINDSOR DRAKE

Gross margin after inference

The single strongest statistical relationship in the set. Compute cost is a cost of goods sold line, and it moves with provider list prices a founder does not control.

- Gross margin correlates with log EV/Revenue at **r = 0.82** across 16 disclosing constituents.
- Stanford HAI measured the cost of a GPT-3.5-equivalent query falling from **$20.00 to $0.07** per million tokens.
- Goldman Sachs put semiconductor-driven inference cost declines at **60% to 70% per year** in May 2026.
- Anthropic's top Opus tier fell from $15 to $5 per million input tokens over the year to August 2026.
- OpenAI's flagship input price rose from $1.25 to $2.00 over the same period, and long context is now surcharged separately.

A denominator a buyer can audit

The highest multiples in the set belong to companies whose revenue line needs no explanation. The lowest belong to those whose revenue carries pass-through hardware or gross services cost inside it.

- Palantir: **84.7%** GAAP gross margin, revenue that is entirely software and services it performs.
- Arm: **97.2%** gross margin on a royalty and licensing line, clearing 49.5x.
- Super Micro: **10.8%** fiscal 2026 gross margin, revenue dominated by bought-in components.
- TeraWulf and Cipher report two defensible gross margin readings each, depending on treatment of operating cost.
- Together AI published annual bookings rather than revenue at its July 2026 financing.

Contract coverage and duration

Contracted revenue is now the primary asset in the AI cloud layer, and disclosure standards for it vary widely.

- CoreWeave: **$103.7bn** ASC 606 RPO, 41% recognisable within 24 months.
- Oracle: **$638bn** RPO at 31 May 2026, up 363% year on year.
- Nebius: over **$40bn** of customer commitments, a company-defined measure.
- Equinix: **$15.0bn** ASC 606 RPO, the tightest coverage in the layer at 1.5 times revenue.
- Digital Realty: **$1.9bn** of annualised backlog rent on signed but uncommenced leases.

Growth that is still accelerating

The market separates high growth from accelerating growth, and pays a further premium for the second.

- Palantir: revenue **+93%** year on year, US commercial **+149%**, closed TCV **$3.37bn**.
- Astera Labs: revenue **+104.5%** year on year and 27% sequentially.
- Nebius: revenue **+454%** year on year, ARR **$3.0bn** at 30 June 2026.
- C3.ai: revenue **down 52.5%** in its fourth fiscal quarter, clearing 4.1x.
- IREN: quarterly revenue flat year on year during the transition from mining to AI cloud.

Source: Windsor Drake computation from exchange closing prices of 21 August 2026 and SEC filings. Constituents named in the appendix.

The Largest AI Valuations, Categorised

WINDSOR DRAKE

Each entry labelled by what produced the number. Every revenue figure below is company-asserted and unaudited.

Company	The number	Category	The denominator behind it
Anthropic	**$965bn** Series H, 28 May 2026	**Private round mark** Primary financing led by Altimeter, Dragoneer, Greenoaks and Sequoia.	**20.5x** On a company-asserted $47bn run rate disclosed with the round.
OpenAI	**$852bn** Primary, 31 Mar 2026	**Private round mark** $122bn committed, anchored by Amazon, NVIDIA and SoftBank, co-led with Andreessen Horowitz.	**35.5x** On $2bn of monthly revenue stated in the company's own announcement.
Databricks	**$190bn** Strategic round, 13 Aug 2026	**Private round mark** Coatue lead. The most consistent self-published run-rate series in the private market.	**27.1x** On a $7bn run rate, growing over 80% year on year per the company.
xAI	**$230bn** Series E, 6 Jan 2026	**Headline valuation** The $20bn raise is company-disclosed; the valuation comes from press reporting rather than the company.	**Not computable** The company published monthly active users of about 600 million and no revenue figure.
Groq	**$3.5bn** Primary, 17 Aug 2026	**Private round mark** Roughly half the $6.9bn primary mark of 17 September 2025, after NVIDIA licensed the technology and hired the team.	**Not computable** No revenue disclosed at either round. The mark moved 50% on a personnel event.

Sources: Anthropic, OpenAI, Databricks and xAI announcements; Bloomberg; Reuters; NVIDIA Form 10-K, 25 February 2026.

Comparable Transaction Analysis

WINDSOR DRAKE

AI capability changes hands constantly and prices it rarely. That asymmetry runs in either direction.

What the transaction record shows

The Windsor Drake Exit Index held **522 records** on 24 August 2026 and **516 distinct transactions** after deduplication. **Sixty-six** name AI, machine learning or model capability, and **forty-five** of those were announced in 2025 or 2026.

The buyer is an operator

Sixty-five of the 66 went to a strategic operator and one to a financial sponsor. Across the whole index the split is 499 strategic, 14 sponsor and three sponsor-owned strategics.

What closes the information gap

Fourteen of the 66 disclose consideration, so the reproducible public comparable set a seller constructs is the benchmark that frames the negotiation.

516
DISTINCT INDEXED TRANSACTIONS

65 of 66
BOUGHT BY AN OPERATOR

36
PUBLIC CONSTITUENTS COMPUTED

Repeat AI buyers

Bloomberg entities bought Regology and Canoe Intelligence within eight weeks. Stripe bought Metronome and an AI gateway business.

Other named acquirers

Airwallex, Backbase, Coupa, Duck Creek, Figure, MoonPay, SoFi, Willis Towers Watson and OpenAI itself.

Where the capability sits

Identity and fraud, capital markets technology, regtech and B2B payments account for 25 of the 66.

Public comparables

36 listed constituents computed on 21 August 2026, weighted to listed-company scale.

Source: Windsor Drake Exit Index, queried 24 August 2026; company announcements.

Rule of 40 Across the AI Software Cohort

WINDSOR DRAKE

Rule of 40 score against EV/Revenue, 14 AI software constituents that disclose a GAAP operating margin. Score is quarterly revenue growth plus GAAP operating margin.

Rule of 40 score versus EV / Revenue

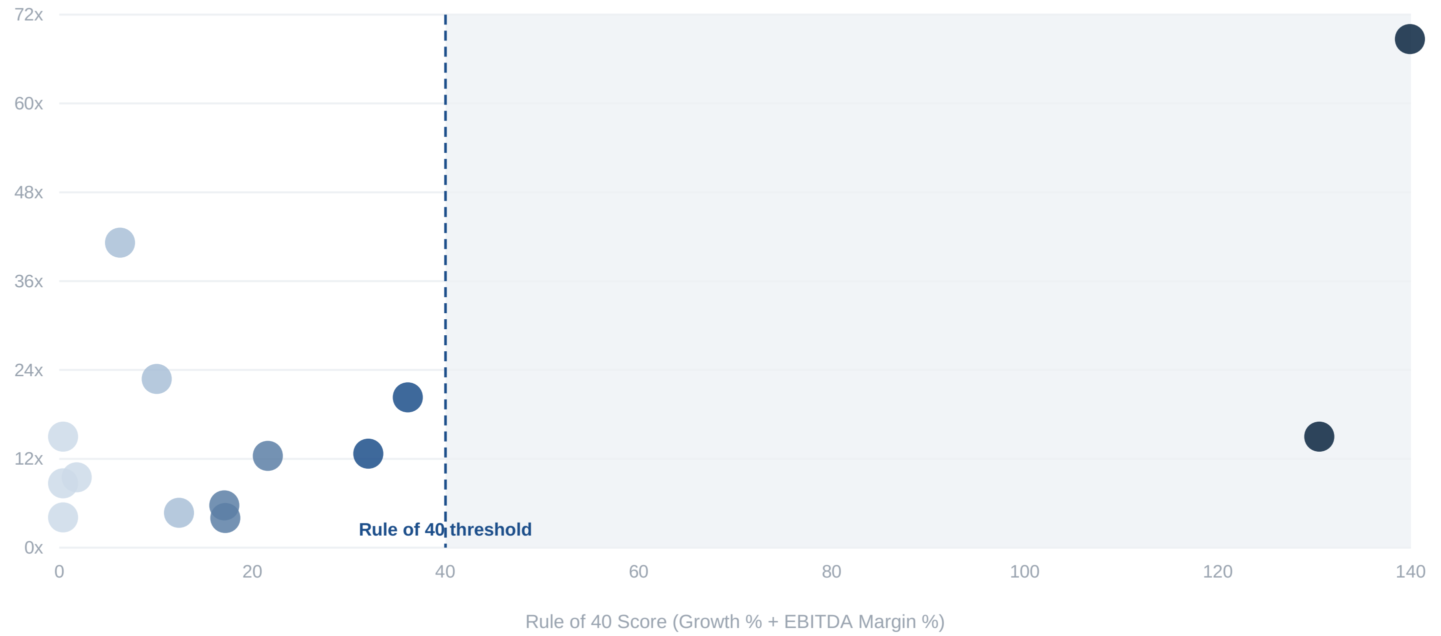

Three constituents sit below the axis floor and are shown at zero: SoundHound at minus 25.0, BigBear.ai at minus 64.5 and C3.ai at minus 287.3. The Windsor Drake Rule of 40 deck covers the broader software repricing this chart sits inside.

CLEARING 40

2 of 14

Palantir at 139.9 and AppLovin at 130.5, both far above the threshold.

MEDIAN SCORE

14.8

Against a median GAAP operating margin of minus 9.5%.

PREMIUM FOR CLEARING

3.8x

41.9x median for the two clearers against 10.9x for the rest.

MEDIAN GROWTH

30.4%

Growth is present across the cohort; operating leverage is not.

Source: Windsor Drake computation from exchange closing prices of 21 August 2026 and SEC filings. Constituents named in the appendix.

Capital Markets: Open for Scale, Selective Below It

WINDSOR DRAKE

Fewer Listings, Far More Money

104 US IPOs raised $145.8 billion year to date in 2026 against 202 pricings in all of 2025 (Renaissance Capital, 24 August 2026).

- Pricings are **down 24.6%** while proceeds are up **543.7%**.
- **12 US deals raised over $1 billion** in H1 2026 against four a year earlier (EY, July 2026).
- EY attributes the strength to AI-related demand in semiconductors, power and data centre infrastructure.
- SpaceX alone raised **$75.0 billion** on 12 June 2026, the largest IPO on record.

AI Listings Have Not Performed Uniformly

The window has rewarded infrastructure and treated application software harshly.

- **Cerebras** priced at $185 on 14 May 2026 at roughly **78x** trailing revenue and trades at $188.70.
- **CoreWeave** priced at $40 in March 2025 and trades at $86.81, up 117%.
- **Figma** priced at $33 in July 2025 and trades at $27.34, down 17%.
- Cerebras sits about 51% below its post-listing high of $386.34.

The Financing Mix Has Changed

AI capital expenditure has moved from operating cash flow onto balance sheets and into private credit.

- Private credit outstanding to AI firms went from near zero to over **$200 billion**, from under 1% to almost 8% of loan volumes (BIS, January 2026).
- AI loans average **$169m** against $90m elsewhere, at a 6.2 versus 6.1 percentage point spread (BIS).
- US technology net investment-grade issuance is **$192bn** year to date, 27% of all IG net issuance (J.P. Morgan Asset Management, July 2026).
- Meta financed Hyperion through a **$27bn** off-balance-sheet structure with Blue Owl.

Official-Sector Bodies Are Watching

Three central banks published concentration warnings within eight weeks of each other in 2026.

- AI companies are about **half of S&P 500 market capitalisation**, up from about a quarter in 2022 (Bank of England, 7 July 2026).
- **50%** of Federal Reserve survey contacts cited AI risk to financial stability in spring 2026, against 0% in autumn 2025.
- The BIS describes circular financing where chip makers take equity in labs that commit to multi-year purchases.
- The ECB flags energy price sensitivity as a distinct channel of AI valuation risk.

Sources: Renaissance Capital; EY; SEC filings; PitchBook-NVCA Venture Monitor. See appendix.

What Changes the Multiple Next

WINDSOR DRAKE

Four variables, each with a published measure attached.

Capital expenditure guidance

Alphabet guided 2026 capex to **$195 to $205 billion** on 22 July and expects a further significant increase in 2027. Amazon raised its 2026 figure to about **$220 billion** on memory prices. Meta guided **$130 to $145 billion**.

Power and permission

Texas 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%**. The IEA projects global data centre use rising from about 415 TWh in 2024 to about 945 TWh in 2030.

Enterprise return on deployment

McKinsey found 88% of organisations using AI, 39% reporting any EBIT impact and about **6%** above 5%. Bain found **40%** of companies measuring AI savings realised 10% or less.

$2.60T	~50%	86%
2026 AI SPENDING, UP 47%	AI SHARE OF S&P 500 MARKET CAP	AI SHARE OF US VENTURE, H1 2026

Model pricing direction

Google discloses a scheduled doubling of Gemini Flash pricing on 1 January 2027, and OpenAI's flagship rate is promotional.

The refinancing question

The Bank of England notes chip lives are shorter than the debt structures financing them.

Adoption broadens slowly

Eurostat measured 20.0% of EU enterprises using AI in 2025, and the US Census Bureau 17% to 20% of US firms.

Productivity is real and modest

A BIS study of 12,000 firms found AI raises labour productivity levels by 4% with no adverse employment effect.

Sources: Bank for International Settlements; Bank of England; Federal Reserve; Gartner; McKinsey & Company. See appendix.

What Earns the Premium in AI

WINDSOR DRAKE

Distribution the model cannot bypass

The premium accrues to companies that own the workflow the model runs inside, not to the model itself.

- Palantir clears **68.7x** on 84.7% gross margin and 47.1% GAAP operating margin.
- Cloudflare clears **41.2x** with 120% dollar-based net retention and 4,698 customers above $100k.
- Snowflake reports **126%** net revenue retention and 779 customers above $1m of product revenue.
- Datadog reports net retention in the low 120s and names its AI-native cohort as a high single digit growth contributor.
- Owning the workflow survives a change of underlying model.

Contracted revenue with named counterparties

The AI cloud layer is priced on contract quality, and the quality is visible in who signed.

- CoreWeave: OpenAI at about **$22.4bn** cumulative, Meta at about **$21bn** through 2032, Anthropic undisclosed.
- Nebius: Microsoft at **$17.4bn**, Meta at up to **$27bn** over five years.
- IREN: Microsoft at about **$9.7bn** over five years with a 20% prepayment, NVIDIA at about **$3.4bn**.
- Cipher: Amazon Web Services at about **$5.5bn** over a fifteen-year base term with a parent guarantee.
- Counterparty credit is now a primary valuation input in this layer.

A revenue line that needs no explanation

The floor of the distribution is occupied by companies whose revenue includes pass-through hardware or gross services cost.

- Super Micro: **$39.1bn** of fiscal 2026 revenue at 10.8% gross margin, clearing 0.6x.
- Innodata: 42.7% gross margin, and 37% of second quarter revenue from a single customer.
- Together AI published annual **bookings** of $1.15bn rather than revenue at its July 2026 round.
- If an acquirer cannot compute the multiple, they will discount the ambiguity.
- Revenue quality is diligence work that starts long before a process.

Growth that survives the cost curve

Inference prices fall for a fixed capability level. A business whose pricing is anchored to tokens inherits that decline.

- Stanford HAI measured a **280-fold** fall in the cost of a GPT-3.5-equivalent query over roughly 18 months.
- Goldman Sachs put inference cost declines at **60% to 70% per year** in May 2026.
- Outcome and workflow pricing insulates the revenue line; per-token pricing does not.
- EY found **82%** of leaders expect per-seat pricing to lose relevance within five years.
- Pricing architecture is a valuation decision, not a commercial one.

Source: Windsor Drake computation from exchange closing prices of 21 August 2026 and SEC filings. Constituents named in the appendix.

Preparing an AI Exit: A 12 to 18 Month View

WINDSOR DRAKE

The work that moves an AI multiple is denominator work, and it takes longer than a process does. A decision to engage the market while current conditions hold is, in practice, a decision taken now.

Fix the denominator

Months 0-3

- Separate recurring contracted revenue from usage revenue and from services in management reporting.
- Isolate inference and inspection cost as a cost of goods sold line a buyer can trace.
- State the basis for any contracted revenue figure and hold it consistent across every document.
- Establish a trailing twelve month revenue series an auditor would recognise.

KEY MILESTONE
A revenue bridge a buyer can audit

Build the evidence

Months 3-9

- Assemble the public comparable set the seller will argue from, with named constituents.
- Document contract duration, renewal rates and counterparty credit quality.
- Quantify what happens to gross margin if provider list prices move 30% in either direction.
- Close the disclosure gaps diligence will surface anyway.

KEY MILESTONE
A defensible comp set and a clean data room

Test the buyer universe

Months 9-15

- Map operators first: 65 of 66 AI-capability transactions in the index went to a strategic buyer.
- Establish which acquirers have bought this capability in the trailing eighteen months.
- Understand the acquirer's own multiple, because it constrains what they can pay in stock.
- Prepare for structures other than a straight acquisition, including licence and hire.

KEY MILESTONE
A ranked, evidenced buyer list

Run the process

Months 15-18

- Enter with the benchmark established rather than discovered.
- Hold the priced-multiple evidence through diligence, not only at the letter of intent.
- Price against margin structure and contract quality, not against the AI composite.
- Maintain competitive tension across categories of buyer to close.

KEY MILESTONE
Competitive tension at close

Source: Windsor Drake computation from exchange closing prices of 21 August 2026 and SEC filings. Constituents named in the appendix.

AI Valuation Outlook

WINDSOR DRAKE

Composite EV/Revenue for the Windsor Drake AI comparable set. Scenarios are Windsor Drake's view; the current level is computed from market data.

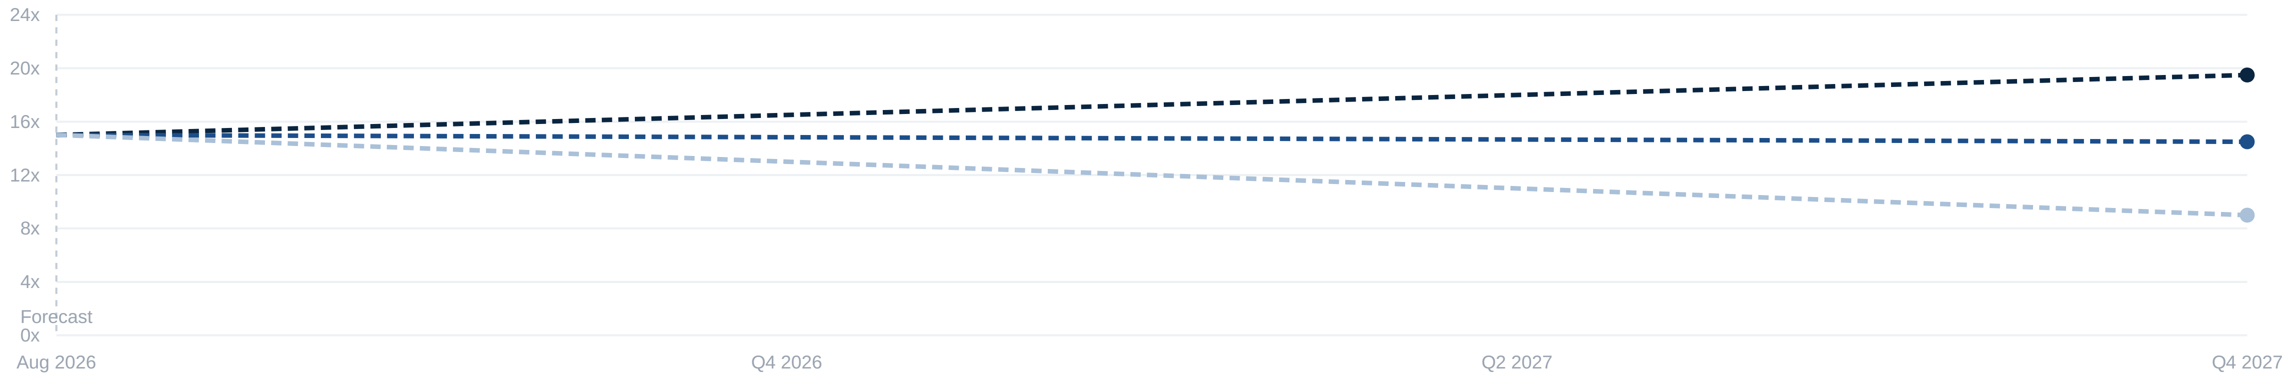

Bull 19.5x

Key Drivers

- Hyperscaler capital expenditure guidance holds at the Morgan Stanley $1.1 trillion 2027 path.
- Enterprise EBIT impact broadens beyond the 6% McKinsey measured at high performers.
- The OpenAI and Anthropic listings price well and pull the private complex onto an audited basis.

CAPABILITY SELLERS WITH CLEAN MARGIN STRUCTURE SHOULD MOVE; OPERATOR COMPETITION IS AT ITS MOST INTENSE.

Base 14.5x

Key Drivers

- Capital expenditure grows but power and permitting slow the buildout timetable.
- Dispersion by layer persists or widens further.
- Private marks stay above priced public multiples without a broad reset.

MARGIN STRUCTURE DETERMINES OUTCOME MORE THAN TIMING DOES. BUILD THE DENOMINATOR EVIDENCE AND RUN A DISCIPLINED PROCESS.

Bear 9.0x

Key Drivers

- A financing event resets the credit assumptions behind contracted AI capacity.
- The Bank of England's stated risk of a sharp equity correction materialises.
- Enterprise budgets re-anchor to the 10% or lower savings Bain measured rather than to projections.

THE CLOUD AND SYSTEMS LAYERS FACE THE MOST COMPRESSION. READINESS MATTERS MORE THAN PRICE DISCOVERY.

Source: Windsor Drake computation from exchange closing prices of 21 August 2026 and SEC filings. Constituents named in the appendix.

Case Study: Two Deals, One Benchmark

WINDSOR DRAKE

The largest AI acquisition on record produced no multiple. A deal a fifth its size produced one that any seller can use.

SpaceX and Anysphere

SpaceX agreed on 16 June 2026 to acquire Anysphere, developer of the Cursor coding agent, for **$60.0 billion** in stock, and closed on 14 August 2026. It is the largest AI acquisition on record.

- Consideration was **389,289,254** SpaceX Class A shares priced at the seven-day volume weighted average price.
- No target financial statements were filed, so no revenue multiple exists for the transaction.
- Anysphere's last published figure was a company-asserted run rate above $1bn at its November 2025 round.

IBM and Confluent

IBM acquired Confluent for **$31.00 per share in cash**, announced 7 December 2025 and closed 17 March 2026, at an equity value of about **$11.3 billion**.

- Total purchase price of **$11,590 million** against net cash acquired of $982 million.
- Confluent reported fiscal 2025 revenue of **$1,166.7 million**, up 21%.
- That gives an enterprise value of about **9.1x** trailing revenue, computable entirely from filings.

The read for a founder

Cash produces a benchmark

An all-cash deal for a listed target creates a public multiple. An all-stock deal for a private target creates a headline.

9.1x is the number that travels

The Confluent multiple will be quoted in AI data infrastructure processes for years. The $60 billion will be quoted as atmosphere.

Structure follows the buyer, not the seller

NVIDIA paid $17.0 billion for Groq's technology licence and team in December 2025, recording $14.4 billion of goodwill attributable primarily to the workforce. No acquisition was recorded.

The window for capability

Bain measured AI-related deal value at almost half of strategic technology deal value for transactions above $500 million in 2025. Operator appetite for capability is a function of the build-versus-buy calculus, and that calculus currently favours buying.

Sources: SpaceX Forms 8-K, 16 June and 14 August 2026; IBM Form 10-Q, 23 April 2026; Confluent Form 10-K, 11 February 2026; NVIDIA Form 10-K, 25 February 2026.

Emerging Buyer Trends

WINDSOR DRAKE

Three shifts a seller should account for when building an AI buyer universe.

Licence and hire is now standard

NVIDIA recorded **$17.0 billion** for a non-exclusive licence to Groq's technology plus certain employees, of which **$14.4 billion** was goodwill attributable primarily to the workforce, stating no contracts, products or equity were purchased. Meta and Scale AI and Google and Windsurf follow the pattern.

Incumbents buy capability in a workflow

Identity and fraud, capital markets technology, regtech and B2B payments account for 25 of the 66 AI transactions in the index. Buyers acquire a model embedded in a process they already sell.

Vendor financing reshapes the buyer list

NVIDIA committed up to $100 billion to OpenAI as each gigawatt deploys, and up to $10 billion to Anthropic alongside Microsoft's $5 billion. Suppliers now fund their own customers.

66	45	17%
AI TRANSACTIONS INDEXED	ANNOUNCED IN 2025 OR 2026	OF TOP-100 H1 2026 DEALS CITED AI

Deal language is shifting

PwC found AI cited in roughly a third of the 100 largest 2025 deals and 17% in H1 2026. The framing moved to workflow.

Concentration is the structure

PwC counts megadeals above $5 billion at 48% of 2026 deal value. Excluding them, deal value is down 4%.

The listing route is opening

OpenAI and Anthropic are reported to have filed confidentially in June 2026.

What this means for a buyer list

Rank by appetite for the specific capability, and include structures a buyer may prefer to an acquisition.

Sources: NVIDIA Form 10-K, 25 February 2026; PwC Global M&A Industry Trends 2026 Mid-Year Outlook; Bain & Company. See appendix.

Sources

Every third-party institution whose data appears in this deck. House computation is separated from third-party figures throughout.

Institution	Report / Source	Date
US Securities and Exchange Commission	*EDGAR filings: 10-K, 10-Q, 8-K, 6-K, 424B4 for all 36 constituents*	2026-08
Nasdaq and NYSE	*Official closing prices, 21 August 2026*	2026-08
Bank for International Settlements	*Bulletin 120, Financing the AI boom; Annual Economic Report 2026*	2026-06
Bank of England	*Financial Stability Report, July 2026*	2026-07
Federal Reserve	*Financial Stability Report, May 2026; FEDS Notes on AI adoption*	2026-05
European Central Bank	*Financial Stability Review, May 2026*	2026-05
International Energy Agency	*Energy and AI; Electricity 2026*	2026-02
US Energy Information Administration	*Short-Term Energy Outlook; Annual Energy Outlook 2026*	2026-08
Stanford HAI	*AI Index Report 2025 and 2026*	2026-04
Gartner	*Worldwide AI Spending Forecast; AI Platforms and Models Forecast*	2026-07
McKinsey & Company	*The State of AI; The cost of compute; MGI colocation data centres*	2026-06
Company investor relations	*Earnings releases and calls: Alphabet, Amazon, Meta, Microsoft, Oracle*	2026-07

Sources and Methodology

WINDSOR DRAKE

The remaining institutions, and how every figure was produced.

Institution	Report / Source	Date
Bain, PwC and EY	*Global M&A Report 2026; M&A Trends Mid-Year 2026; Global IPO Trends Q2 2026*	2026-07
PitchBook, NVCA and CB Insights	*Venture Monitor Q2 2026; Global Unicorn Tracker; State of AI Q1 2026*	2026-08
Goldman Sachs, Morgan Stanley and J.P. Morgan	*AI capital expenditure, inference cost and financing research*	2026-07
US Census Bureau and Eurostat	*Business Trends and Outlook Survey; EU enterprise AI use, 2025*	2026-05
BCG, Deloitte, KPMG and Renaissance Capital	*AI Radar 2026; State of AI in the Enterprise; Global AI Pulse; US IPO statistics*	2026-08

METHOD

How the multiples were computed

Enterprise value equals market capitalisation plus total debt less cash and short-term investments, using official Nasdaq and NYSE 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. TSMC is converted from New Taiwan dollars at the rate implied by its own quarterly revenue translation. No figure here is licensed from a data vendor.

Derived figures and the comparable set

36 listed constituents across six sub-segments, all named in the report. Four figures are derived: the gross margin correlation across the 16 constituents disclosing a GAAP gross margin, implied private marks dividing a post-money by a published run rate, Rule of 40 scores, and the IBM and Confluent multiple.

The Denominator Test

A Windsor Drake framework. Valuations are classified as a priced multiple where both a transactable price and a filed revenue figure exist, a private round mark where a primary financing priced a share class against a company-asserted run rate, and a headline valuation where the number derives from a tender, a secondary or press reporting. Every private figure in this deck carries its classification.

Windsor Drake

WINDSOR DRAKE

Sell-side M&A advisory

We advise technology founders on one side of the table only.

- Sell-side exclusively. No buy-side conflicts.
- Enterprise value **$5 million to $300 million**.
- AI, fintech, payments, SaaS and cybersecurity.
- Process control, not listings.
- Benchmarks computed, never licensed.

Proprietary data

The Windsor Drake Exit Index: 516 distinct sourced transactions, each with a source URL and a confidence grade.

- Built continuously, not purchased.
- 66 transactions carry AI or model capability.
- Used for comparable sets, buyer mapping and case work.
- Each record carries a source URL and a confidence grade.
- Updated every quarter.

How we work

A defined buyer universe, a controlled timeline, and a seller who knows what the evidence supports before the first conversation.

- Benchmarks established before the process, not during it.
- Buyer lists ranked by demonstrated appetite for the capability.
- Competitive tension held to close.
- Structures a buyer may prefer, anticipated in advance.
- No fabricated numbers, in any document we produce.

Contact

For a confidential valuation discussion on a specific business.

- **windsordrake.com**
- windsordrake.com/confidential-valuation
- Research library: windsordrake.com/research
- This deck: windsordrake.com/ai-valuation-multiples