WINDSOR DRAKE

AI Valuations: Q2 2026

MAY 2026

Windsor Drake · Market Intelligence

Executive Summary: Conviction Meets Scrutiny

WINDSOR DRAKE

A Bifurcated Premium

AI valuations sit far above the broader software market, but the premium is now conditional on demonstrated revenue, not narrative.

- AI-native software clears a Windsor Drake working benchmark near **11x EV/ Revenue**.
- Frontier labs are priced on strategic optionality at **15x to 30x** annualised revenue.
- Legacy software with no credible AI position has compressed toward **3x to 6x**.
- The distance between an AI label and an AI business is the central valuation question.

A Real Demand Supercycle

Underlying AI demand is compounding fast enough to underwrite premium multiples.

- Combined 2026 hyperscaler capital spending is tracking above **$650B**.
- Bain projects roughly **200 GW** of incremental AI compute demand by 2030.
- Inference at scale has become the centre of gravity for AI workloads.
- Enterprise adoption is broad, even where measured financial return still lags.

The Monetisation Gap

The quarter's defining tension is the distance between AI investment and proven return.

- McKinsey finds **94%** of enterprises not yet seeing significant value from AI spend.
- Revenue is concentrated in a short list of frontier labs and infrastructure names.
- Applied-AI monetisation is improving, but remains uneven across the cohort.
- The market is beginning to price proof of return, not the promise of it.

Bubble Risk in Plain View

Institutions are openly flagging concentration and repricing risk.

- The IMF has compared current pricing to the **dot-com** peak of 2000.
- The largest technology names now represent roughly **a fifth** of the MSCI World Index.
- A technology repricing would be felt across the whole market, not one sector.
- Windsor Drake reads the risk as a discipline on pricing, not a reason to wait.

Sources: McKinsey & Company; Bain & Company; International Monetary Fund. See appendix.

Executive Summary: Capital Floods In

WINDSOR DRAKE

Record Venture Funding

Venture capital is concentrating into AI at a scale without precedent.

- AI absorbed the large majority of all Q1 2026 global venture funding.
- Quarterly AI investment surpassed the **entire full-year 2025** total.
- OpenAI's **$122B** round set an **$852B** valuation in March 2026.
- Capital is concentrating sharply at the top of the quality curve.

The IPO Window Reopens

Public markets have reopened decisively for scaled AI infrastructure.

- **Cerebras** completed the largest US technology IPO since 2019.
- It raised **$5.55B** and reached a roughly **$95B** market value on debut.
- CoreWeave's March 2025 listing has compounded strongly since.
- OpenAI and Anthropic are both preparing potential public listings.

Strategic M&A Accelerates

AI has become the organising logic of the 2026 deal market.

- Goldman Sachs forecasts an AI-led **megadeal renaissance** for 2026.
- Morgan Stanley sees a full spectrum of AI M&A ahead across deal sizes.
- The contest for chips, power, data and talent is driving deal value.
- Hyperscalers are decisively buying capability rather than building it.

Private Capital Pressure

Record private capital is intensifying competition for AI-exposed assets.

- PE dry powder stood near **$3.7T** at the start of 2026.
- Nearly **half** of Bain's larger 2025 deals involved AI-native targets.
- Rate normalisation is gradually reopening leveraged deal math.
- Sponsors are converging on AI infrastructure and AI-native software.

Sources: PitchBook; CB Insights; Goldman Sachs; Bain & Company; SEC filings. See appendix.

Key Founder Takeaways: Q2 2026

WINDSOR DRAKE

1. Prove Revenue, Not Narrative

An AI label alone no longer earns a premium. The market now prices demonstrated, durable, recurring revenue and is hardening that test each quarter.

- AI-native software clears the **~11x** benchmark only on real revenue.
- Show paying customers, retention and gross margin, not pilots.

2. Rule of 40 Still Gates the Premium

For AI-native software with real revenue, the Rule of 40 remains the single cleanest predictor of a top-decile multiple.

- Top performers command **50% to 100%** valuation premiums.
- Only an estimated **15% to 20%** of AI-native software clears it.

3. AI Is Now the Acquirer's Mandate

Strategic buyers are acquiring AI capability across the full stack: chips, power, data, models and talent.

- Hyperscalers and software incumbents buy capability over building.
- Map your specific capability to a named buyer's declared gap.

4. Mind the Monetisation Gap

AI investment is running well ahead of measured enterprise return, and diligence is increasingly built around closing that gap.

- McKinsey finds **94%** of enterprises not yet seeing significant value.
- Bring hard ROI data, named customers and documented renewals.

5. Capital Intensity Cuts Both Ways

Compute-heavy models can command scarcity premiums, but they also carry real funding and obsolescence risk that buyers price in.

- AI infrastructure trades **10x to 22x** revenue on demand strength.
- Model a credible path through the next two funding cycles.

6. Prepare Before the Window Moves

A full process runs 12 to 18 months end to end, and today's alignment of demand, capital and pricing will not hold indefinitely.

- Clean financials, documented model and data rights, a tight data room.
- The prepared asset captures the competitive tension in a process.

Source: Windsor Drake advisory methodology. See appendix.

Founder FAQs: Valuations, Timing & Strategy

WINDSOR DRAKE

The questions AI founders ask most, answered against the Q2 2026 market.

Q1 Which valuation metric applies to my AI business?

Use **EV/Revenue**, often on a forward basis, for high-growth AI software and infrastructure, and **EV/EBITDA** only for the rare profitable, scaled names. Frontier labs are effectively priced on strategic optionality. Always reference the correct subsegment cohort, never a blended AI average.

Q2 What are the key subsegment ranges right now?

Foundation models lead at **15x to 30x** annualised revenue, with agentic AI platforms at **12x to 25x**. AI infrastructure and compute trade **10x to 22x**, AI-native application software **8x to 15x**, and legacy software without an AI position has compressed to **3x to 6x**.

Q3 Why do AI companies command such a premium?

Durable double- and triple-digit growth, software gross margins, and a credible claim on a multi-trillion-dollar productivity opportunity. McKinsey estimates **$2.6T to $4.4T** of annual value potential from generative AI. The premium is real, but increasingly conditional on proof.

Q4 How do public and private valuations compare?

A wide private premium persists. Private AI companies average near **24x** revenue against roughly **12x** for public AI-exposed software. Recent listings are re-rating the public set upward, but frontier-lab marks remain optionality-driven rather than revenue-driven.

Q5 Is there an AI bubble?

Institutions including the IMF have flagged concentration risk and drawn dot-com comparisons. Windsor Drake's reading: underlying demand is real, but the dispersion between proven and speculative assets is widening fast, and pricing discipline is returning to the market.

Q6 When is the optimal time to run a process?

After demonstrating **4 to 6 quarters** of predictable revenue and retention, while demand, capital availability and pricing remain aligned. A full process runs **12 to 18 months**, so preparation in the current cycle is what meets that window.

Q7 Who are the most active buyers today?

Hyperscalers and large software incumbents acquiring capability, **semiconductor and infrastructure consolidators** racing for compute and power, and **private equity** deploying record dry powder into AI-exposed infrastructure and software platforms.

Q8 How do we maximise our multiple, and expect earn-outs?

Clear the **Rule of 40** where revenue allows, evidence **NRR above 120%**, document model and data rights, and quantify customer ROI. Expect **earn-outs** to bridge valuation gaps on unproven capability, typically paid over 12 to 24 months.

Source: Windsor Drake advisory methodology. See appendix.

Q2 2026 Valuation Landscape Overview

WINDSOR DRAKE

Premium multiples concentrate in the model and infrastructure layers; AI-light software has re-rated down.

Median EV / Revenue Multiple by AI Subsegment (x)

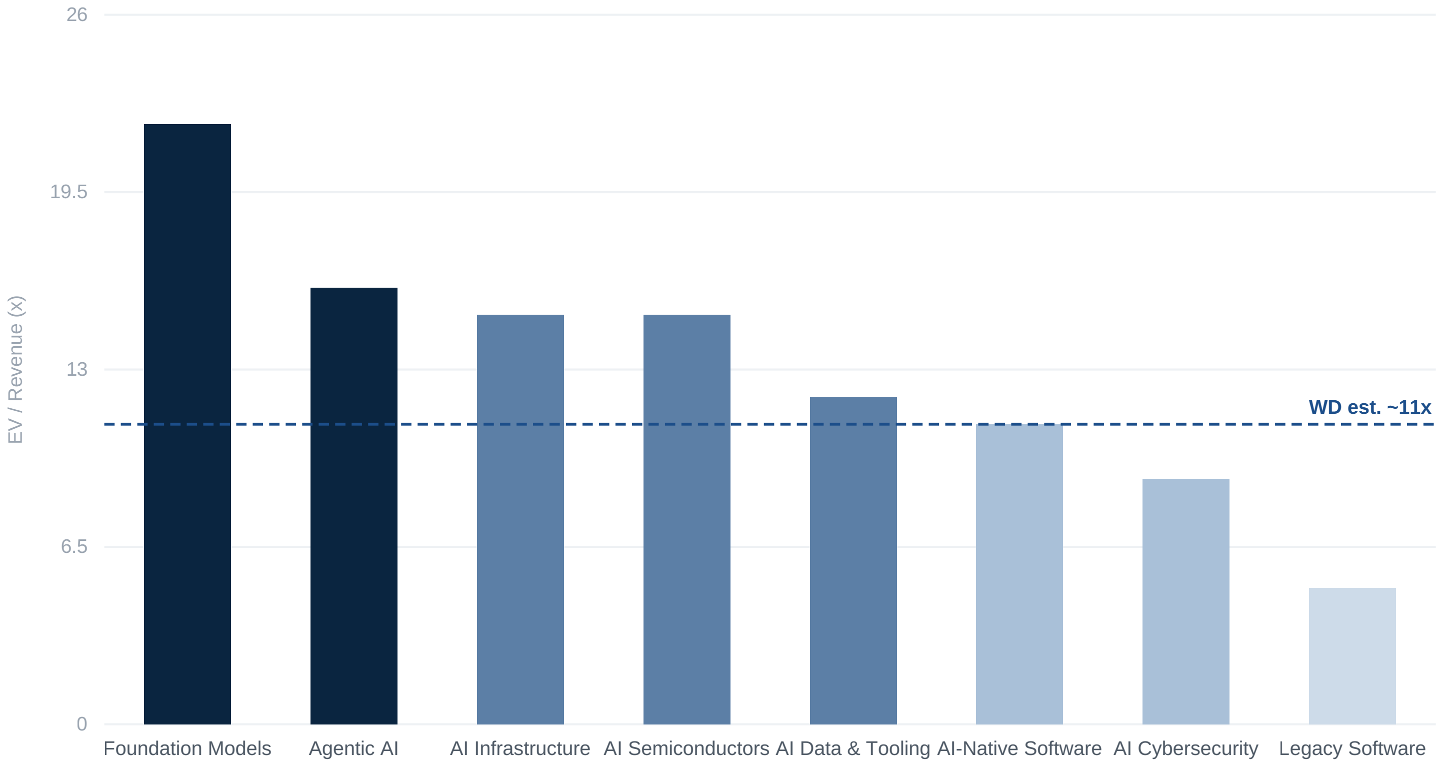

A layered premium: the model and compute layers carry double-digit-plus multiples on scarcity and demand, while application software is priced on proven revenue and legacy software without an AI position has compressed toward 3x to 6x.

AI-NATIVE BENCHMARK

~11x

Windsor Drake's working EV/Revenue benchmark for AI-native application software.

FRONTIER-LAB PREMIUM

15-30x

Foundation-model labs are priced on optionality, not current revenue.

GENERATIVE AI POTENTIAL

$2.6-4.4T

McKinsey's estimate of the annual value generative AI could create.

Key Driver

Conviction meets scrutiny: acquirers and investors still pay up for the AI stack, but the premium is now earned with demonstrated revenue and retention rather than narrative.

Sources: PitchBook; CB Insights; McKinsey & Company; Windsor Drake analysis. See appendix.

Public and Private Markets: A Persistent Premium

WINDSOR DRAKE

Public AI multiples are re-rating upward as listings validate the category, yet a wide private premium persists.

Average EV / Revenue Multiple, Public vs Private AI (x)

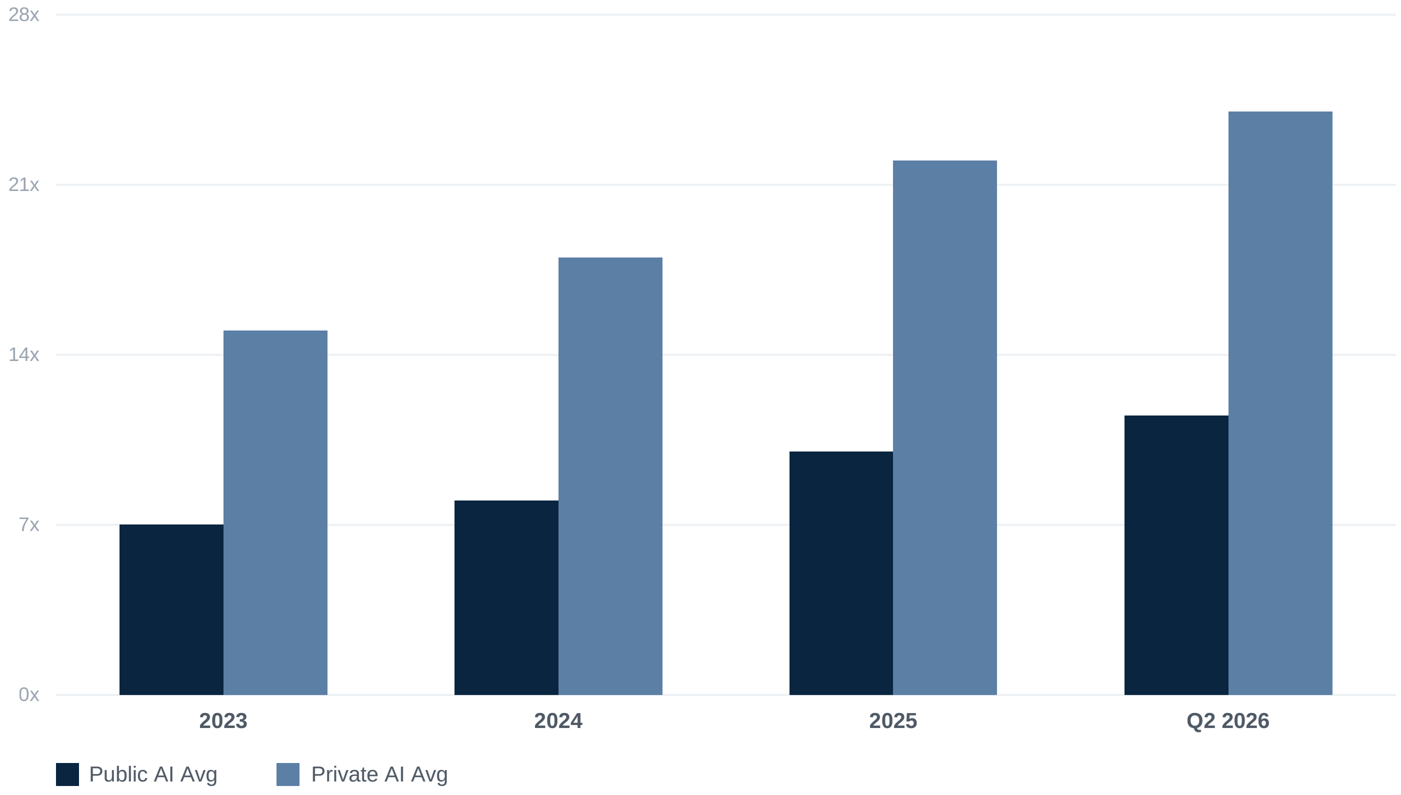

PUBLIC / PRIVATE SPREAD

~12x

Still wide; unlike prior cycles, the private premium has not yet closed.

PUBLIC RE-RATING

+4.5x

Public AI software has re-rated upward since 2023 as listings priced the category.

OPTIONALITY MARKS

Frontier

The widest private marks sit with frontier labs valued on strategic optionality.

An incomplete convergence: public AI-exposed software has re-rated from about 7x to 11.5x as the IPO class validated the category, but the private average near 24x still carries a wide premium, concentrated in frontier labs priced on optionality.

Sources: PitchBook; CB Insights; McKinsey & Company; Windsor Drake analysis. See appendix.

Exit Valuation Multiples by AI Subsegment

A steep gradient runs from frontier models (15-30x) to legacy software with no AI position (3-6x).

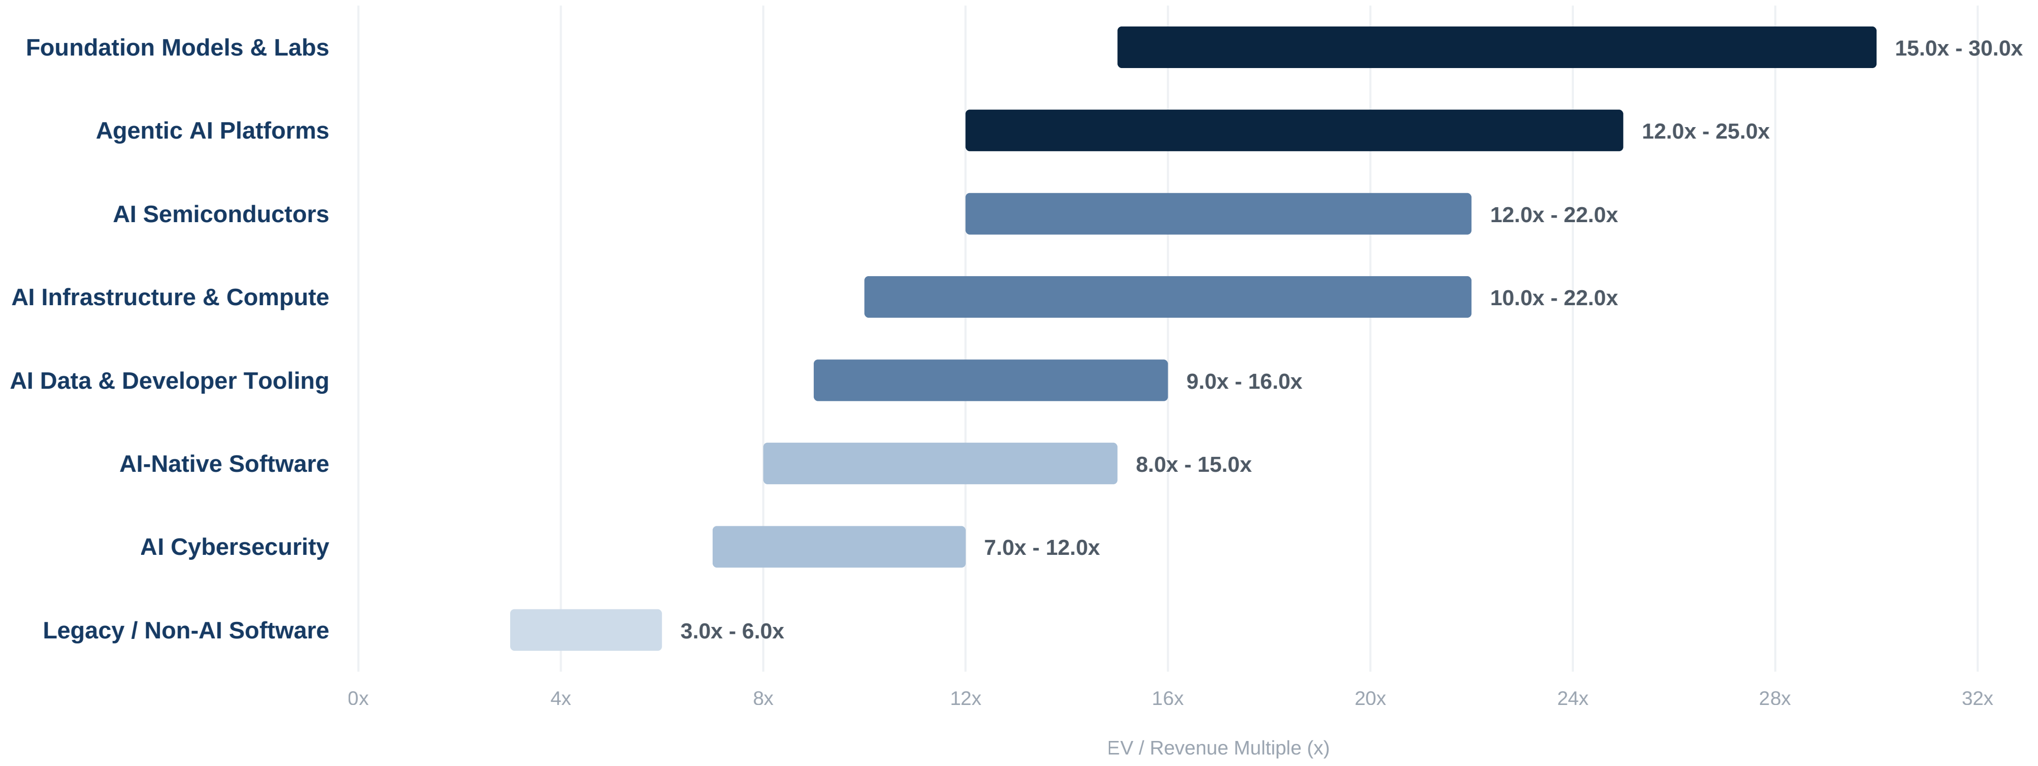

KEY OBSERVATION

The market is paying for position in the AI stack and for demonstrated demand, not for an AI label. The five-to-one gap between frontier models and AI-light software reflects scarcity of compute, models and proven monetisation.

Sources: PitchBook; CB Insights; McKinsey & Company; Windsor Drake analysis. See appendix.

Valuation Multiple Drivers: Expansion vs. Compression

WINDSOR DRAKE

Net expansion to roughly 11x for AI-native software is driven by demand and the IPO re-rating, partly offset by monetisation and bubble-risk concerns.

NET EXPANSION OF +2.0X

A genuine demand supercycle and the public re-rating from the 2026 IPO class outweigh a combined 1.5x drag from the enterprise monetisation gap and rising bubble-risk caution. The bridge reflects Windsor Drake analysis of the cited institutional data.

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

Capital Markets: IPO Window & Listing Benchmarks

WINDSOR DRAKE

The 2026 listing class reopened the window decisively for scaled AI infrastructure and compute.

Cerebras Systems (Nasdaq)

Listed May 2026, pricing above range at $185 for the largest US tech IPO since 2019.

- Raised roughly **$5.55B**, with shares closing the debut up about **68%**.
- Reached a market value near **$95B** on its first trading day.
- Confirmed deep public appetite for AI inference and compute infrastructure.
- Set a fresh public benchmark for the AI semiconductor and systems cohort.

CoreWeave (Nasdaq)

The March 2025 GPU-cloud listing has compounded strongly across the year.

- Raised about **$1.5B** at a roughly $23B valuation at IPO.
- Has since re-rated well above its listing price on AI compute demand.
- Validated the neocloud model as an investable public category.
- Anchored the case for the 2026 wave of AI infrastructure listings.

OpenAI & Anthropic (Pipeline)

The largest frontier labs are actively preparing potential public listings.

- OpenAI is preparing a possible listing later in 2026 after its **$852B** round.
- Anthropic is reported to be exploring a listing with major banks.
- A frontier-lab IPO would be the defining capital-markets event of the cycle.
- Both would test public appetite for optionality-priced model assets.

Private Mega-Marks

The strongest private names are setting record marks without listing.

- OpenAI's primary round priced the company at **$852B** in March 2026.
- Anthropic is reported to be raising near a **$900B** valuation.
- xAI was valued at **$250B** within the SpaceX combination.
- Tender and secondary markets give frontier labs liquidity without a listing.

Sources: SEC filings; financial press (Reuters, Bloomberg, CNBC). See appendix.

Geographic Valuation Variations

WINDSOR DRAKE

The United States commands a decisive AI premium; Europe and APAC trade at a capability and scale discount.

Share of Global AI Venture Investment by Region (%)

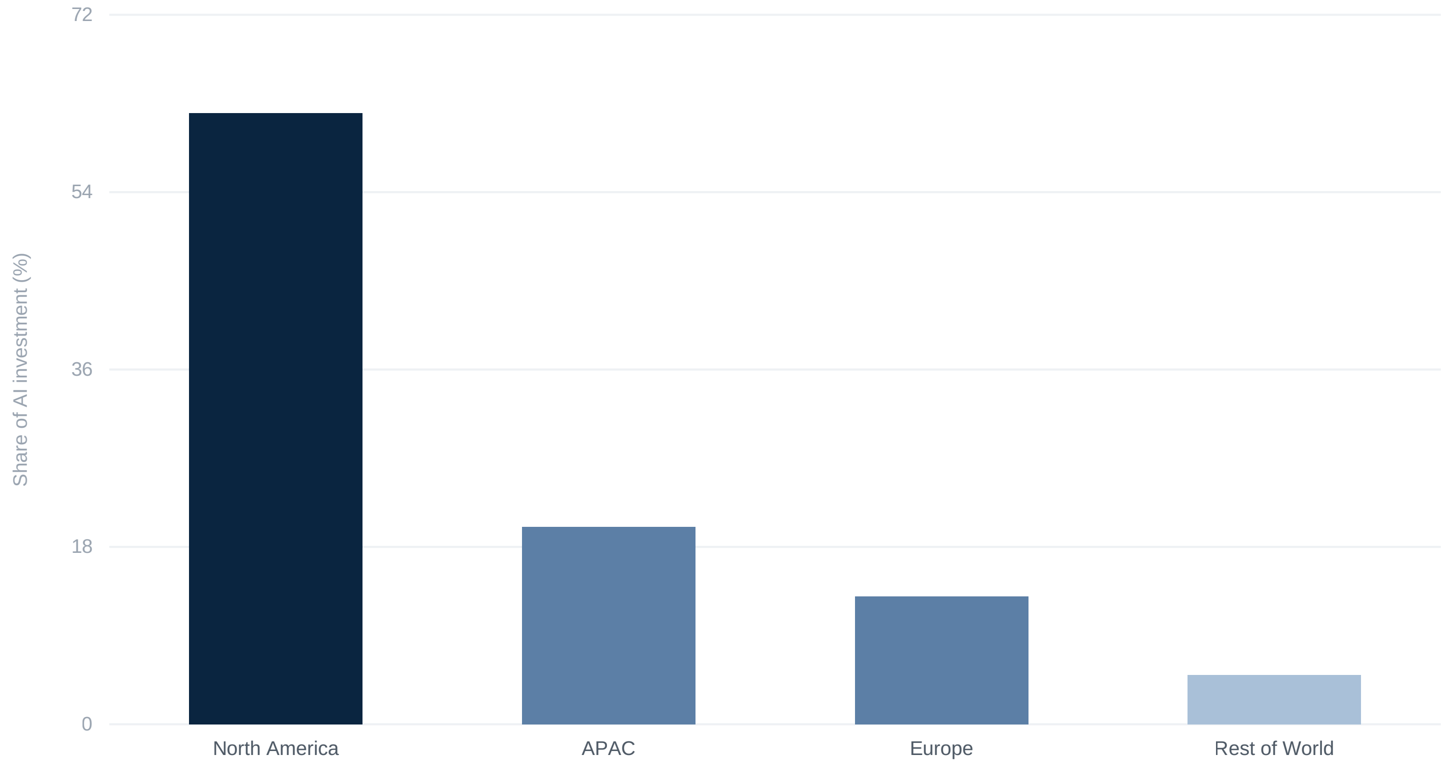

NORTH AMERICA

Premium

Roughly 62% of global AI venture investment, with the deepest capital and exit markets.

EUROPE

Discount

About 13% share; strong research depth, but a scale and capital gap versus US peers.

APAC

Scaling

About 20% share, led by national model programmes and large domestic compute build-outs.

A concentrated map: North America captured the clear majority of global AI venture investment in Q1 2026, on the strength of frontier labs, hyperscaler capital and the deepest exit markets. Europe and APAC trade at a discount that US acquirers are positioned to arbitrage.

Sources: PitchBook; CB Insights; Windsor Drake analysis. See appendix.

Foundation Models: The Optionality Premium

WINDSOR DRAKE

Frontier labs are priced on strategic optionality and compute access, not on current revenue multiples.

EV / Revenue Multiple Range (x)

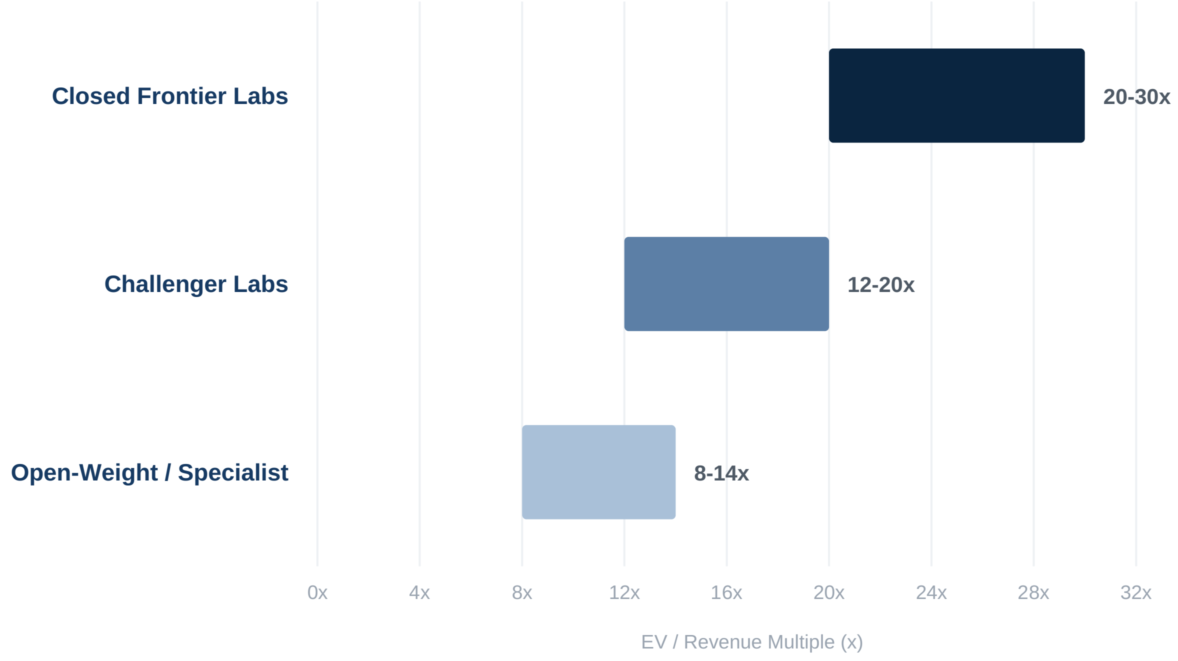

Valuation Drivers

Optionality, Not Revenue

Investors price frontier labs on the probability of dominating the model layer of a multi-trillion-dollar market, so headline revenue multiples understate the underwriting logic.

Compute and Capital Access

Valuation tracks secured access to compute, power and capital. Labs with hyperscaler backing and long-dated supply command the ceiling of the range.

Concentration Risk

Revenue and talent concentrate in two or three labs. Buyers and investors weigh a genuine winner-take-most outcome against a sharp repricing scenario.

KEY OBSERVATION

Foundation models are the highest-multiple subsegment, but the least revenue-anchored. The marks are a bet on strategic position, and they carry the cycle's widest dispersion of outcomes.

Sources: PitchBook; CB Insights; McKinsey & Company; Windsor Drake analysis. See appendix.

AI Infrastructure & Compute

Demand for training and inference capacity underwrites premium multiples for GPU cloud and data-centre platforms.

EV / Revenue Multiple Range (x)

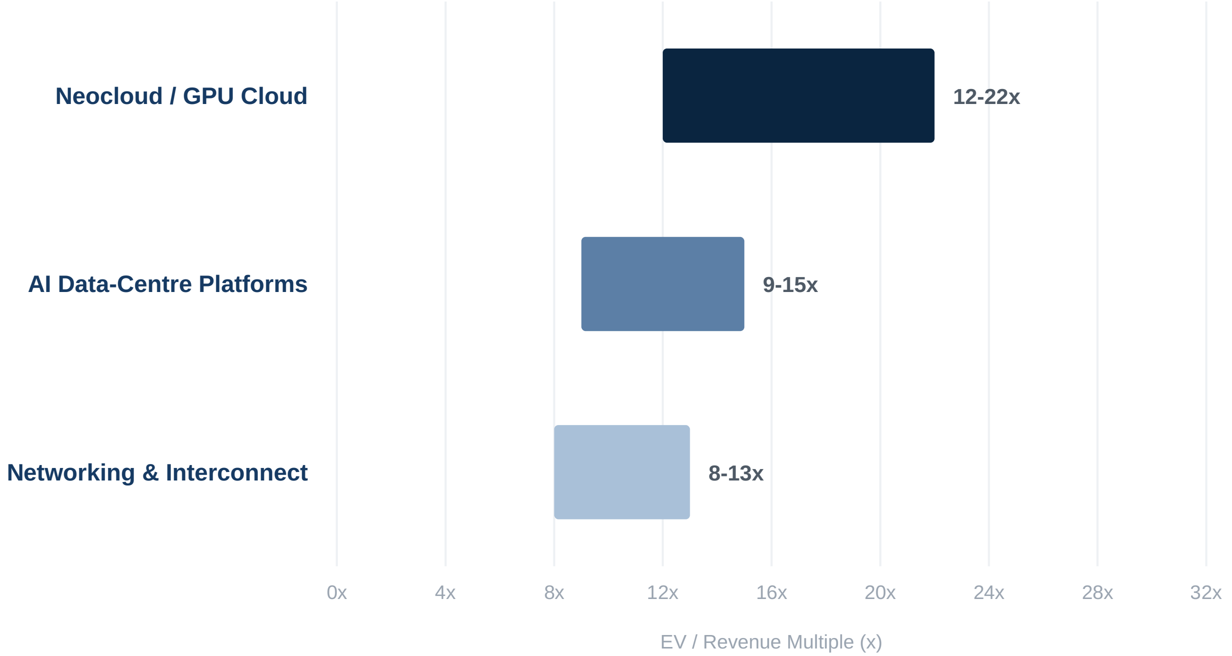

Valuation Drivers

Inference Is the Center of Gravity

Bain notes AI workloads are shifting toward inference at scale. Operators with low-latency, efficient inference capacity are re-rating ahead of pure training capacity.

Power Is the Binding Constraint

Access to firm power, not chips alone, now gates growth. Bain projects roughly 200 GW of incremental AI compute demand by 2030, with the US about half of it.

Capital-Intensity Discount

Heavy capital cycles, depreciation and obsolescence risk cap multiples below the model layer despite comparable demand growth and contracted revenue.

KEY OBSERVATION

Compute is priced as critical infrastructure, but the multiple carries a capital-intensity discount. The scarcest asset in the stack is contracted power, not silicon.

Sources: PitchBook; CB Insights; McKinsey & Company; Windsor Drake analysis. See appendix.

AI Semiconductors & Hardware

WINDSOR DRAKE

Pure-play AI accelerators command the premium; diversified and legacy silicon trades at a clear discount.

EV / Revenue Multiple Range (x)

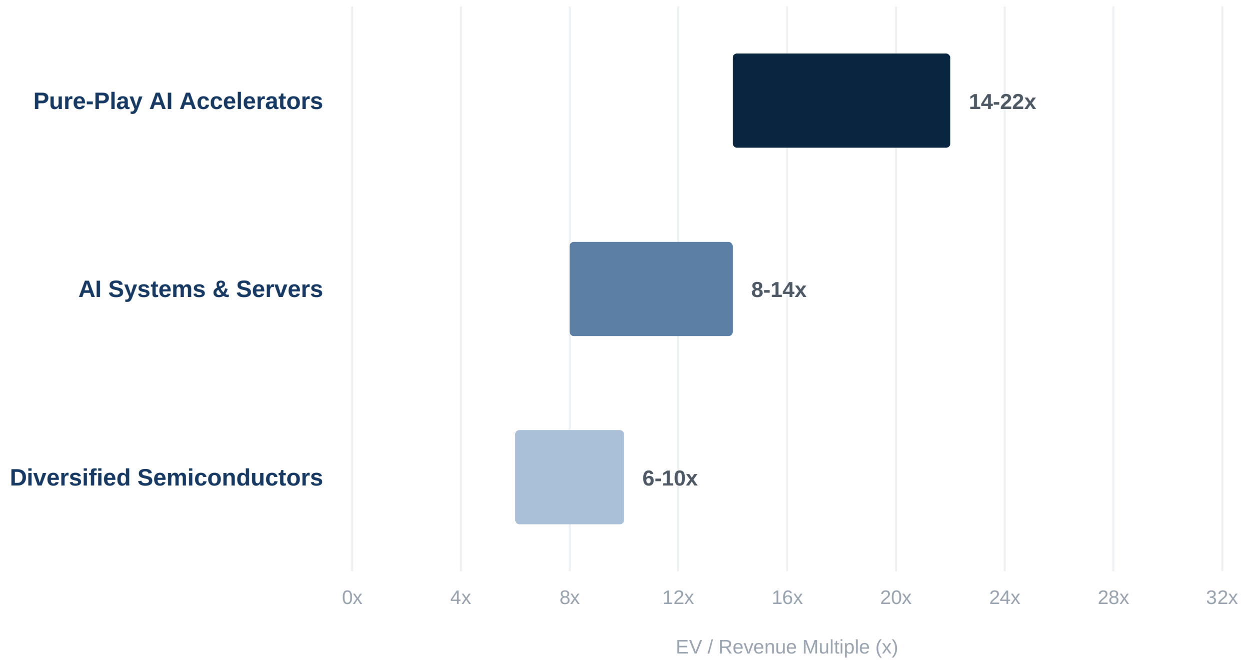

Valuation Drivers

Public Re-rating

The Cerebras listing re-rated the public AI-silicon cohort, demonstrating deep investor appetite for inference-optimised accelerator architectures.

Design Wins and Lock-In

Multiples track secured design wins with hyperscalers and labs. Long-dated supply agreements convert cyclical hardware into recurring, visible revenue.

Supply-Chain Concentration

Dependence on a narrow set of foundries and packaging capacity is a structural risk that buyers underwrite carefully into accelerator valuations.

KEY OBSERVATION

AI silicon is being valued more like infrastructure software than like cyclical hardware, provided the order book is contracted and the architecture is inference-efficient.

Sources: PitchBook; CB Insights; McKinsey & Company; Windsor Drake analysis. See appendix.

Applied & Agentic AI

WINDSOR DRAKE

Agentic platforms that automate real workflows command the highest applied-layer multiples.

EV / Revenue Multiple Range (x)

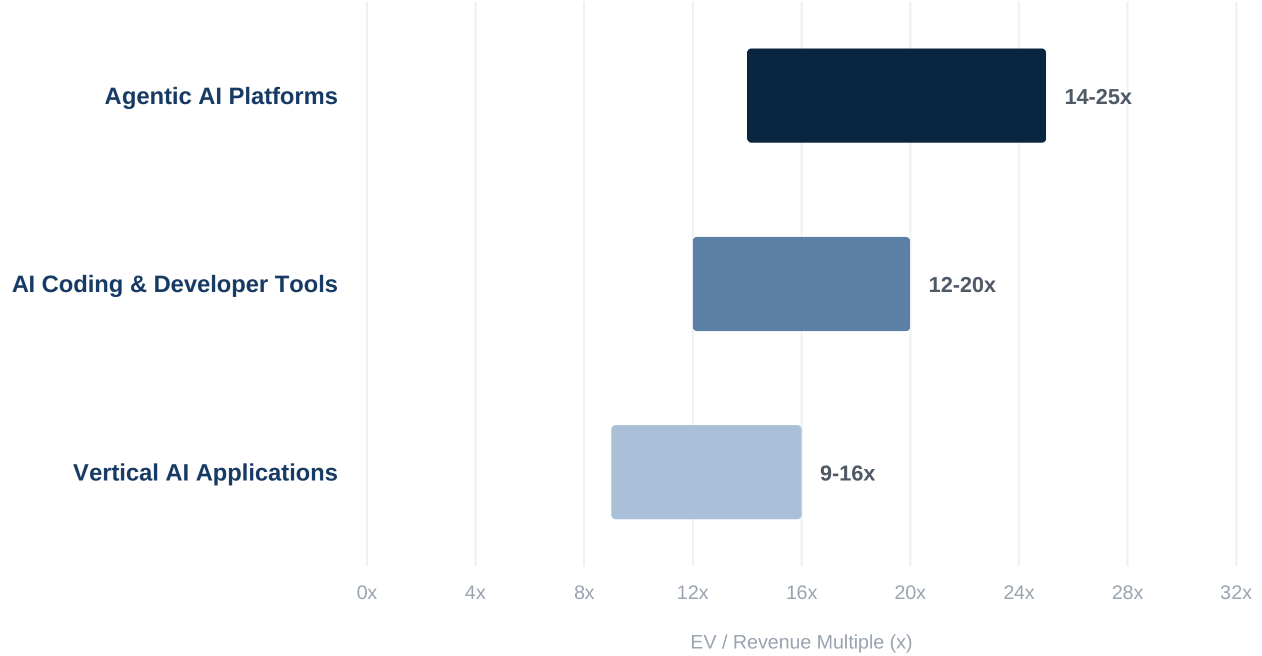

Valuation Drivers

Workflow Ownership

Agents that complete end-to-end tasks, rather than assist, command software-plus multiples by capturing labour budgets, not just software budgets.

Measurable ROI Wins

Applied-AI multiples increasingly track demonstrated customer ROI. Coding and developer tools lead because their productivity gains are the easiest to quantify.

Model-Dependency Risk

Applications built on third-party models face margin and moat questions. Proprietary data, distribution and workflow depth are what defend the premium.

KEY OBSERVATION

The applied layer is where the monetisation gap is closing fastest. Agentic platforms with proven ROI are re-rating toward the infrastructure layer.

Sources: PitchBook; CB Insights; McKinsey & Company; Windsor Drake analysis. See appendix.

AI-Native Software

WINDSOR DRAKE

AI-native applications trade well above legacy software, but the premium is now earned with proven revenue.

EV / Revenue Multiple Range (x)

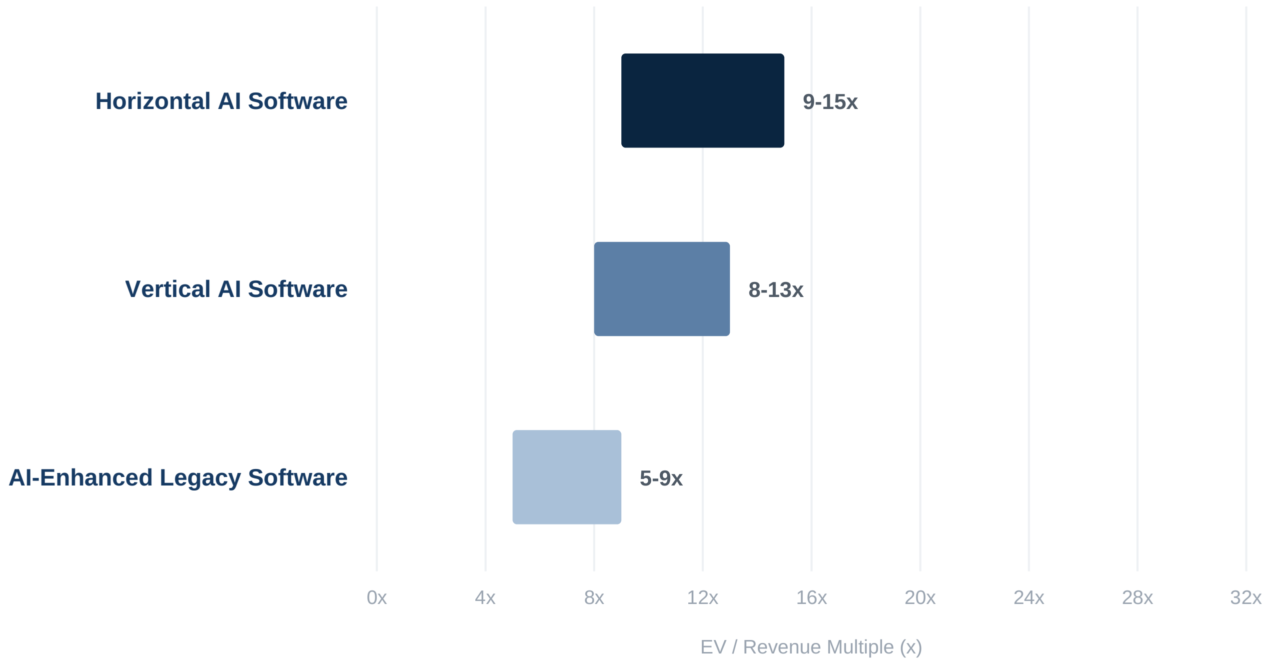

Valuation Drivers

Native vs. Enhanced

Software built around AI from inception trades well above legacy suites that bolt on AI features. Architecture, not marketing language, sets the multiple.

Retention Is the Proof

Net revenue retention above 120% is the clearest evidence that AI features are creating durable value rather than driving a one-time adoption spike.

Pricing-Model Transition

The shift from per-seat to outcome and consumption pricing is re-rating the names that capture usage-based upside as their AI features scale.

KEY OBSERVATION

AI-native software is the cohort the Windsor Drake ~11x benchmark anchors. The dispersion within it, from 5x to 15x, is the widest in the application layer.

Sources: PitchBook; CB Insights; McKinsey & Company; Windsor Drake analysis. See appendix.

AI Data, Tooling & Security

WINDSOR DRAKE

The connective layer of data, orchestration and governance is re-rating as AI deployment moves into production.

EV / Revenue Multiple Range (x)

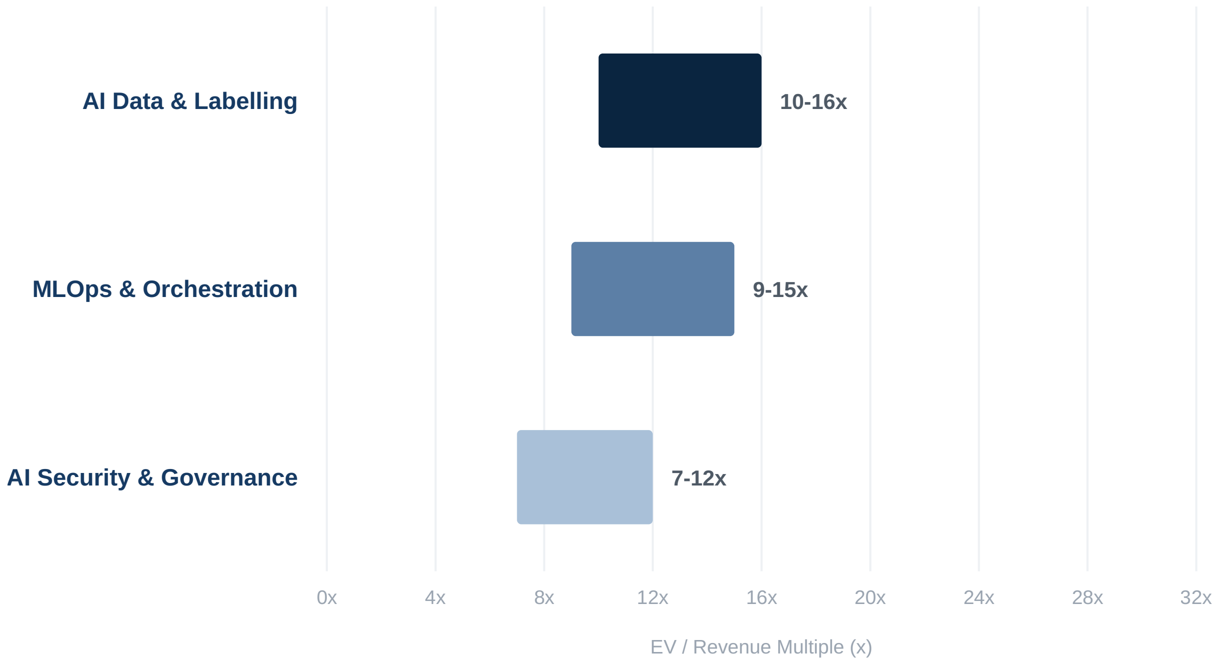

Valuation Drivers

Data as the Scarce Input

Proprietary, high-quality training and evaluation data is a durable moat. Owners of unique data pipelines command the top of the tooling range.

Production-Grade Demand

As enterprises move AI from pilot to production, orchestration, evaluation and observability tools shift from optional to required infrastructure.

Governance Tailwind

Tightening AI regulation and enterprise risk policy are converting AI security and governance from a cost line into a budgeted, recurring spend category.

KEY OBSERVATION

The tooling and governance layer is an early-cycle opportunity: less crowded than the model layer, and directly leveraged to the shift from AI experiment to AI in production.

Sources: PitchBook; CB Insights; McKinsey & Company; Windsor Drake analysis. See appendix.

Comparable Transaction Analysis Framework

WINDSOR DRAKE

A rigorous capability-and-revenue-quality filter, not a broad AI label.

1. Select Peer Set

Identify genuinely comparable assets by position in the AI stack, revenue quality, model dependency and growth profile, never by a broad AI label alone.

2. Normalise Metrics

Adjust KPIs to a pro-forma basis: normalise for compute cost recognition, separate genuine recurring revenue from pilots, and reconcile gross versus net retention.

3. Adjust for Structure

Account for deal-specific terms (earn-outs, talent retention, compute commitments and stock-versus-cash mix) that pull headline valuation away from underlying value.

64	2019-26
TRACKED TRANSACTIONS	EXIT INDEX COVERAGE

Proprietary Transaction Index

Calibration draws on Windsor Drake's proprietary index of **64 verified and reported transactions (2019-2026)**, a living comparable-set database refreshed each quarter as AI-driven deals accelerate.

Capability-Acquisition Filter

Recent index entries such as the Coupa acquisition of Tonkean show buyers paying for AI orchestration and agentic capability rather than for installed customer bases.

Rule of 40 Premium Adjustment

A specific premium layer is applied for AI-native software in the top decile of efficiency; offsetting discounts apply for heavy earn-out or retention components.

Control Premium Calibration

Indications include a control-premium layer, typically **20% to 35%**, where capability, data and talent synergies can be concretely underwritten.

Source: Windsor Drake proprietary transaction database and methodology.

Strategic Acquirer Mapping by Subsegment

WINDSOR DRAKE

Hyperscalers buy capability and compute; PE targets cash-generative scale; software incumbents defend their platforms.

Subsegment	Hyperscalers	Private Equity	Software Incumbents
Foundation Models	**HIGH** Strategic stakes and compute partnerships to secure the model layer.	**LOW** Scale and burn rates sit outside the leveraged-buyout model.	**MODERATE** Minority investment and deep commercial partnerships.
AI Infrastructure	**HIGH** Capacity, power and neocloud assets to feed internal demand.	**HIGH** Contracted, cash-generative data-centre and power platforms.	**LOW** Outside core competency for most application vendors.
AI Semiconductors	**HIGH** Custom-silicon programmes to cut compute cost and dependency.	**MODERATE** Carve-outs and systems suppliers with contracted order books.	**LOW** Rarely acquired directly by software incumbents.
Applied & Agentic AI	**MODERATE** Selective buys of agent frameworks and developer ecosystems.	**HIGH** Buy-and-build of vertical agentic platforms with real revenue.	**HIGH** Capability M&A to embed agents into core product suites.
AI-Native Software	**LOW** Acquired only where strategically adjacent to platform goals.	**HIGH** Prime buy-and-build and multiple-arbitrage territory.	**HIGH** Defensive consolidation against AI-native disruptors.
AI Data & Security	**MODERATE** Data pipelines and evaluation tooling for internal model work.	**HIGH** Recurring-revenue tooling and governance platforms at scale.	**HIGH** Security and governance bolt-ons for enterprise trust.

High Activity / Priority Target

Moderate / Opportunistic

Low Activity / Niche Only

Hyperscalers & Strategic Buyers

WINDSOR DRAKE

The largest platforms are acquiring AI capability across the full stack rather than building it in time.

Strategic Motives: Buy vs. Build

The internal build cycle for frontier capability is too slow against a fast-moving field. That gap is compelling hyperscalers and software incumbents to acquire models, compute, data and talent outright, treating M&A as the primary route to capability.

Acquisition Patterns

Patterns span the full size range: strategic stakes and compute partnerships in frontier labs, mid-size capability buys in the applied layer, and talent-and-licence structures that secure teams without a full acquisition.

Full stack
BUYER FOCUS

Buy
OVER BUILD, ON AI

Priority: Compute and Power

Securing chips, data-centre capacity and firm power is the first-order acquisition driver, as access to compute now gates the pace of model and product progress.

Priority: Agentic Capability

Hyperscalers and software incumbents are buying agent frameworks and orchestration to embed autonomous workflows across existing enterprise product suites.

Priority: Data and Talent

Proprietary data assets and intact research teams are acquired as scarce inputs, often through structures designed specifically to retain key technical staff.

Integration Model

Acquirers increasingly run a federated model, preserving the target's research culture and pace while overlaying enterprise distribution, trust and compliance.

Sources: Goldman Sachs; Morgan Stanley; McKinsey & Company. See appendix.

Private Equity & Growth Capital Patterns

WINDSOR DRAKE

Record dry powder is meeting an AI-defined deal market, concentrating capital in cash-generative AI-exposed assets.

Deployment Pressure

With roughly $3.7T of global dry powder to deploy, sponsors face acute pressure to transact. Bain reports that nearly half of its larger 2025 deals involved AI-native companies or explicitly cited AI benefits in the thesis.

Infrastructure-First Thesis

Sponsors are concentrating on the intersection of AI and power: contracted data centres, energy assets and compute platforms, where cash flows are visible and the risk profile suits leverage better than venture-stage software.

~$3.7T
GLOBAL DRY POWDER

~50%
LARGER DEALS AI-LINKED

Ideal Target Profile

Sponsors prioritise contracted or recurring revenue, defensible margins and a credible AI position, the profile that supports leverage and a clean exit story.

Value-Creation Playbook

AI-enabled cost reduction inside portfolio companies, buy-and-build of fragmented vertical AI software, and operational professionalisation of fast-scaled assets.

Rate-Environment Catalyst

Gradual rate normalisation is reopening leveraged deal math, supporting larger AI infrastructure and software transactions through 2026.

Deal Structure Trends

All-cash transactions for deal certainty on contracted assets, with earn-outs and retention pools bridging gaps on unproven AI capability.

Sources: McKinsey & Company; Bain & Company; S&P Global Market Intelligence. See appendix.

Competitive Moats Driving Premium Valuations

WINDSOR DRAKE

Multiples above 15x revenue are reserved for AI companies that can demonstrate structural defensibility.

Proprietary Data

ASSET VALUE: HIGHEST

Unique training and evaluation data

- Creates a compounding quality advantage rivals cannot easily replicate.
- Powers model performance that survives the commoditisation of base models.
- Grows more valuable as the install base and usage history deepen.
- **Action:** secure exclusive data rights and first-party capture at every touchpoint.

Compute & Power Access

SCARCITY VALUE: HIGH

Secured capacity and firm power

- Long-dated compute and power agreements gate the pace of progress.
- Convert a cyclical cost into a contracted, defensible position.
- Are increasingly central to frontier and infrastructure valuations.
- **Action:** lock multi-year capacity and energy supply ahead of need.

Workflow Ownership

STICKINESS VALUE: HIGH

Agents embedded in core operations

- Owning the end-to-end workflow captures labour budgets, not just software spend.
- Deep integration raises switching costs far above feature-level AI.
- Generates the proprietary usage data that feeds the next product cycle.
- **Action:** move from assistant features to autonomous, outcome-priced workflows.

Distribution & Trust

BARRIER VALUE: MED-HIGH

Enterprise reach and governance

- Established enterprise distribution shortcuts the cost of AI go-to-market.
- Security, compliance and governance credentials gate large-enterprise budgets.
- Brand trust is decisive while buyers remain cautious on AI risk.
- **Action:** invest early in security posture and enterprise-grade governance.

Source: Windsor Drake analysis. See appendix.

Rule of 40 Performance Distribution

WINDSOR DRAKE

For AI-native software with real revenue, clearing the Rule of 40 still unlocks a 50 to 100% valuation premium.

EV / Revenue Multiple vs Rule of 40 Score

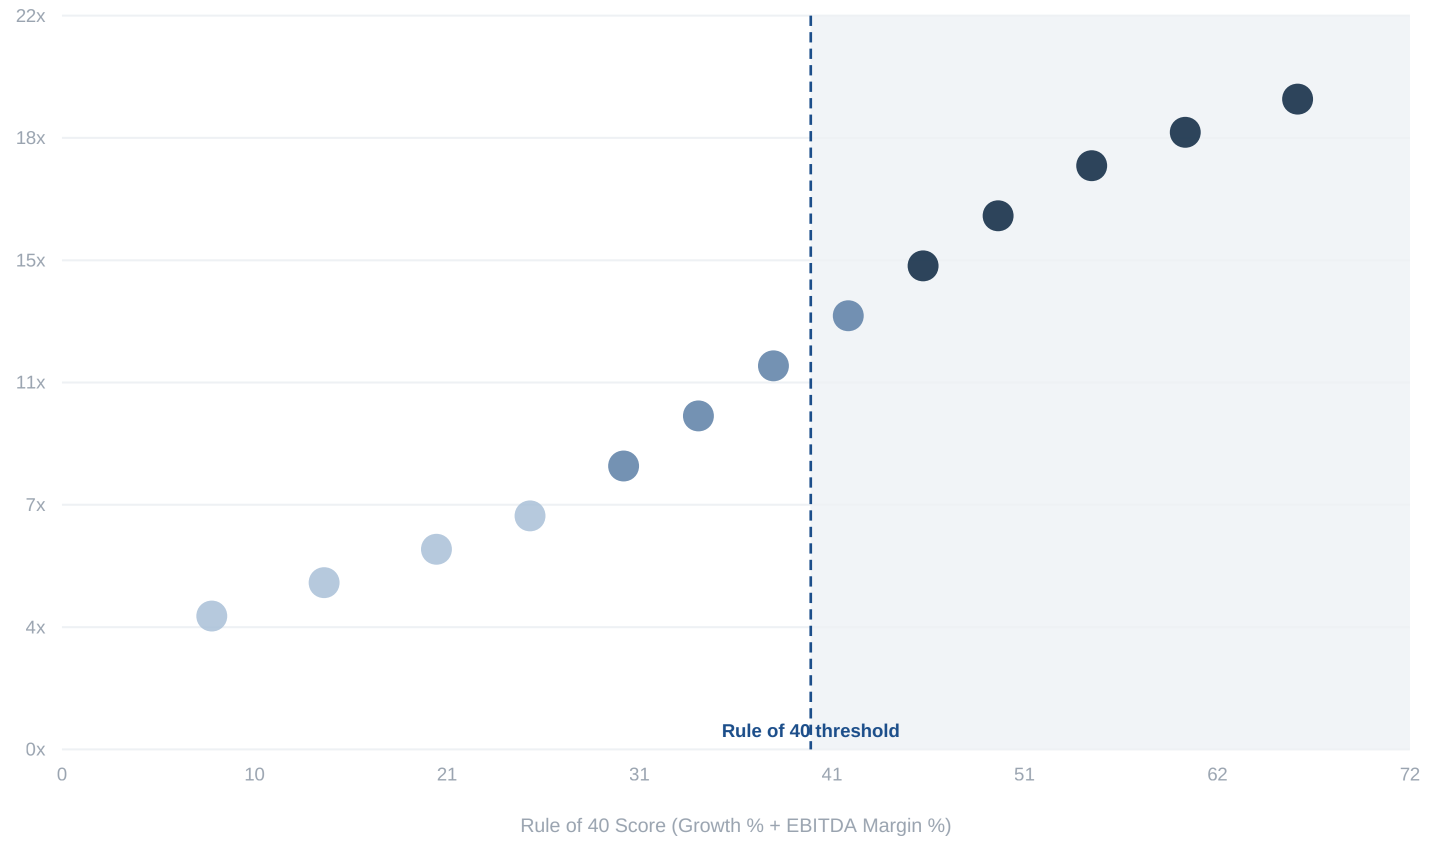

TOP QUARTILE (SCORE >50)

15x+

Scaled, efficient AI-native software; the primary targets for premium strategic M&A.

RULE OF 40 MET (40 TO 50)

11-15x

A healthy growth and efficiency balance; credible IPO-ready candidates.

BELOW THRESHOLD (<40)

4-10x

Growth-at-all-costs and unproven-monetisation zones; exposed to repricing.

Only an estimated 15 to 20% of AI-native software clears the Rule of 40, but it remains the cleanest predictor of a premium multiple. Frontier labs are exempt: they are pre-profit and priced on optionality, not on the Rule of 40.

Sources: McKinsey & Company; Bain & Company; Windsor Drake analysis. See appendix.

Cross-Border M&A Considerations

WINDSOR DRAKE

AI-specific regulation and export control are the primary deal risk, and the primary arbitrage, in 2026.

Regulatory & Security Review

Heightened scrutiny of AI models, compute and data infrastructure (CFIUS, FDI review and export controls on advanced chips), and divergence between the EU AI Act and US frameworks, materially complicate cross-border AI deals and lengthen timelines.

Capability Arbitrage

US acquirers are using premium domestic valuations and a strong capital position to acquire discounted European and APAC research talent and AI intellectual property at a clear relative discount.

12-18mo
CROSS-BORDER CYCLE

+30-50%
LONGER CLEARANCE

Extended Timelines

Regulatory clearance for cross-border AI deals now runs materially longer than domestic transactions; ensure runway to withstand delay without losing process leverage.

Export-Control Diligence

Confirm that compute supply, chip access and model weights can transfer cleanly across jurisdictions before signing; treat this as a gating diligence item.

Data Sovereignty & IP

Establish where training data, model weights and IP legally reside, and structure the transaction to keep data-residency and licensing obligations intact.

Talent Retention Structures

Cross-border AI value sits in research teams; retention pools, immigration planning and culture continuity must be designed well before the LOI.

Sources: EY; McKinsey & Company; Windsor Drake analysis. See appendix.

Exit Valuation Optimisation Strategies

WINDSOR DRAKE

Four levers that systematically de-risk an AI asset while amplifying its scarcity value.

1. Prove Monetisation

6 TO 12 MONTHS PRE-EXIT

Close the investment-to-return gap

- Convert pilots into paid, contracted, renewing revenue.
- Quantify customer ROI with auditable, named case studies.
- Move toward outcome and consumption pricing that captures usage upside.
- Evidence revenue durability across at least three to four quarters.

2. Revenue Quality

PREDICTABILITY

Engineer resilience into the model

- Increase genuine recurring revenue toward **80%+** of total.
- Sustain net revenue retention above **120%** through expansion.
- Reduce concentration so the top 10 accounts sit under 25%.
- Separate durable revenue clearly from one-time AI-adoption spikes.

3. Rule of 40 Efficiency

PREMIUM TIER

Prove scalable, efficient growth

- Use AI internally to decouple headcount from revenue growth.
- Reach **above 40%** on growth plus margin where revenue allows.
- Improve gross margin by optimising inference and compute cost.
- Track the score monthly with board-level visibility.

4. Defensibility Narrative

COMPETITIVE TENSION

Frame the moat, not the model

- Anchor the story on proprietary data, workflow and distribution.
- Show resilience to base-model commoditisation and price deflation.
- Map specific capability gaps for the top five acquirers.
- Run a structured process to manufacture competitive tension.

Source: Windsor Drake advisory methodology. See appendix.

Positioning for Strategic Acquisition

Strategic value is driven by capability fit, integration ease and defensible AI advantage.

Capability Fit

Demonstrate unique models, data or agentic IP that fill a specific, declared buyer gap, making the buy-versus-build decision self-evident for the acquirer.

Integration Ease

Acquirers pay premiums for assets that integrate cleanly. Document model architecture, data lineage and APIs, and present clean, audited financials.

Defensible AI Advantage

Show why the advantage survives base-model commoditisation: proprietary data, workflow ownership, distribution or regulatory and trust barriers.

Strategic Buyer Mapping

Run a structured gap analysis of potential acquirers and map your capabilities directly to each buyer's declared AI strategy and deficits.

Model & Data Rights Diligence

Pre-empt diligence with clean documentation of training-data provenance, licensing, model ownership and any open-source dependencies.

Talent Retention Planning

Much of an AI asset's value is its research team; design retention pools and continuity plans before engaging buyers.

Comprehensive VDR Readiness

Build a defensive data room addressing IP, data rights, compute commitments and customer-concentration risk before the first buyer engagement.

Source: Windsor Drake advisory methodology. See appendix.

Timing the Exit: 12-18 Month Roadmap

WINDSOR DRAKE

A full process runs 12 to 18 months end to end. Founders who prepare in the current cycle meet the market while today's alignment of AI demand, abundant capital and an open IPO window still holds.

Readiness & Hygiene
Q3 2026

- Audit completion to PCAOB standard
- Model and data-rights provenance review
- Inference and compute cost optimisation
- Clean up the cap table and option pool

KEY MILESTONE
Clean IP and data-rights scan

Strategic Positioning
Q4 2026

- Launch dual-track process preparation
- Build the strategic buyer-targeting list
- Draft the CIM and management presentation
- Lock research-team retention packages

KEY MILESTONE
Retention packages locked

Market Engagement
Q1 2027

- Fireside chats with priority strategics
- Solicit initial indications of interest
- Deliver management presentations
- Open the virtual data room

KEY MILESTONE
Competitive bid tension

Execution & Closing
Q2 2027

- Definitive agreement negotiation
- Regulatory filings (HSR and CFIUS)
- Confirmatory diligence support
- Closing and integration kickoff

KEY MILESTONE
No-MAC event verification

Source: Windsor Drake advisory methodology. See appendix.

2026 Valuation Forecast Scenarios

WINDSOR DRAKE

With AI-native software near 11x, forward trajectories diverge sharply on monetisation proof and the path of rates.

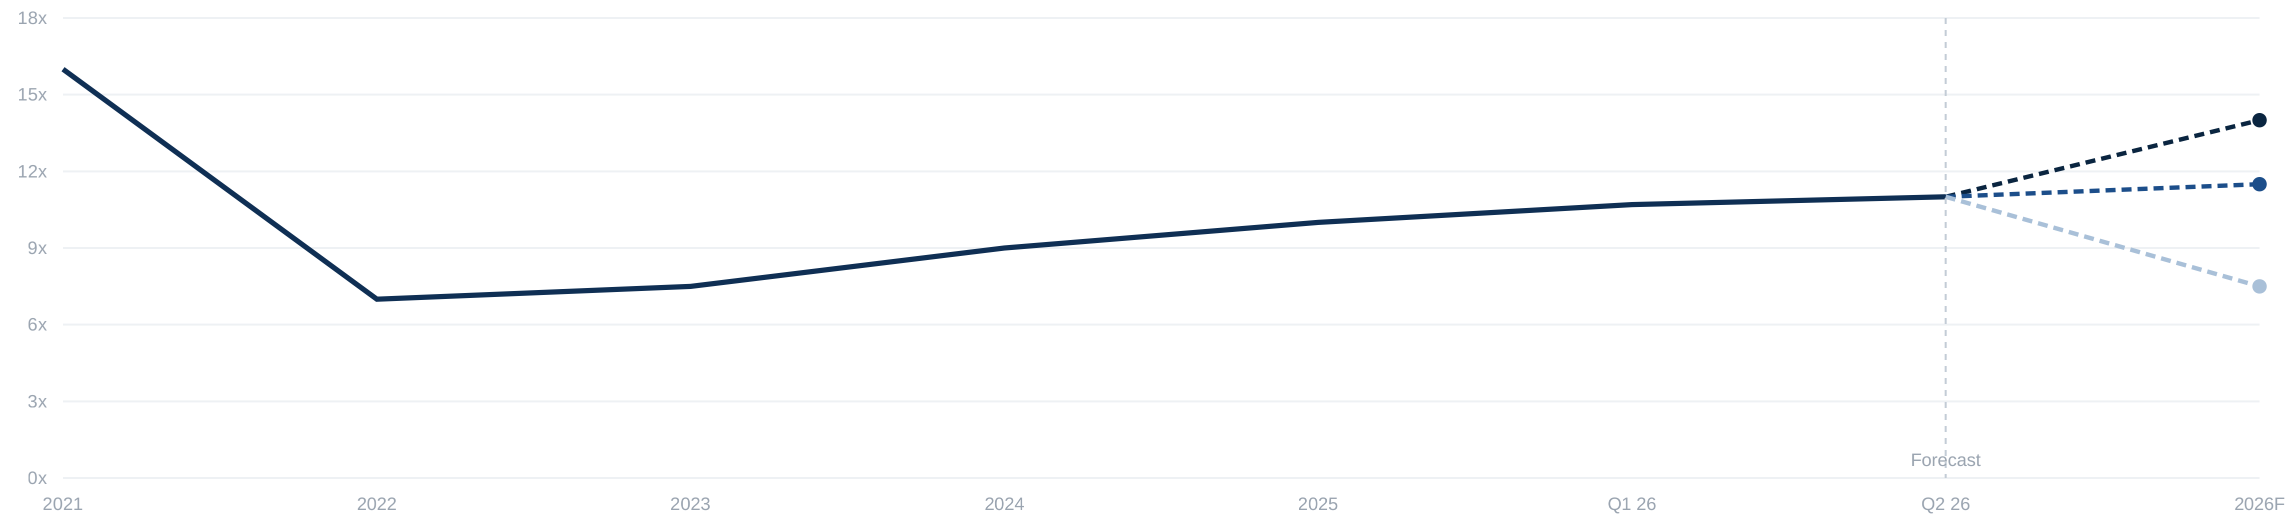

BULL CASE 14.0x

Key Drivers

- Enterprise AI ROI turns clearly positive
- Rate cuts and an open IPO window
- Frontier-lab listings re-rate the sector

STRATEGY: ACCELERATE GROWTH

BASE CASE 11.5x

Key Drivers

- Monetisation improves gradually but unevenly
- One further rate cut, steady demand
- Selective IPOs and active capability M&A

STRATEGY: BALANCE GROWTH & PROOF

BEAR CASE 7.5x

Key Drivers

- The monetisation gap fails to close
- A bubble-driven technology repricing
- Rate holds and a frozen IPO window

STRATEGY: CASH PRESERVATION

Source: Windsor Drake analysis; McKinsey; Federal Reserve projections. See appendix.

Emerging Opportunities & Buyer Trends

WINDSOR DRAKE

Capital is flowing into the connective infrastructure of an AI-first economy.

AI Power & Energy

Firm power has become the binding constraint on AI growth, drawing capital and M&A into generation, grid and data-centre energy assets at utility scale.

Agentic Enterprise Software

Autonomous agents that complete work, rather than assist with it, are extending AI from software budgets into far larger enterprise labour budgets.

AI Trust & Governance

Tightening regulation and enterprise risk policy are driving demand for AI security, evaluation, observability and governance as budgeted, recurring spend.

Capability Acquisitions

Hyperscalers and software incumbents are prioritising models, agents, data and talent over installed customer bases, buying to close capability gaps.

PE Platform Roll-Ups

Sponsors are consolidating fragmented vertical AI software and AI-exposed infrastructure to build scale and drive multiple expansion.

Sovereign & Regional Demand

National AI and compute programmes are creating a new class of large, well-funded buyers and partners outside the established hyperscalers.

Market Intelligence

Top-tier forecasts point to AI-led deal growth through 2026. Record dry powder and capability-driven demand are chasing a supply of proven, defensible AI assets that has not kept pace.

Sources: McKinsey & Company; Bain & Company; Goldman Sachs. See appendix.

Capital Markets Case Study: The Cerebras Listing

WINDSOR DRAKE

The defining capital-markets event of the quarter, and the playbook it sets for founders.

The Listing

AI chipmaker **Cerebras Systems** listed on Nasdaq in **May 2026**, pricing above its range at $185 per share and raising roughly **$5.55B**, the largest US technology IPO since 2019. Shares closed the debut up about **68%**, lifting the company to a market value near **$95B**.

Strategic Rationale

- **Inference focus:** Cerebras is built around AI inference, the workload Bain identifies as the new centre of gravity.
- **Public validation:** the listing proved deep public-market appetite for scaled AI compute and silicon.
- **Window opener:** the debut reset risk appetite ahead of expected listings from other AI infrastructure names.

Implications for Founders

The Window Is Open, and Selective

Public markets have reopened decisively, but for **scaled, demand-backed AI infrastructure**, not for every AI story. The window rewards proven capacity and contracted revenue, and windows of this kind do not stay open indefinitely.

Position in the Stack Sets the Multiple

Cerebras was priced on its **position in the AI stack**, inference compute, and on a contracted order book. Founders should be able to state precisely where they sit in the stack and why that position is defensible.

Demand Proof Beats Narrative

The listing cleared on **evidenced demand**, not on a story. The same test now governs private valuations: proof of revenue, contracts and retention is what converts AI conviction into a premium mark.

Sources: SEC filings; financial press (Reuters, Bloomberg, CNBC). See appendix.

Valuation Methodology: Choosing the Right Metric

WINDSOR DRAKE

The right metric depends on position in the AI stack, profitability and capital intensity.

EV / Revenue

8-15X+

High-growth AI software

- The default lens for AI-native software reinvesting ahead of profit.
- Often applied on a forward basis given the pace of revenue growth.
- The Rule of 40 score dictates where in the range an asset sits.
- Best suited to horizontal and vertical AI-native applications.

Forward Revenue / Optionality

15-30X

Foundation models & frontier labs

- Frontier labs are priced on strategic position, not current revenue.
- Valuation tracks compute access, talent and probability of winning.
- Headline multiples understate the underwriting logic of the round.
- Carries the widest dispersion of outcomes in the market.

EV / EBITDA

SELECTIVE

Profitable, scaled AI names

- Applies only to the rare AI businesses with real positive EBITDA.
- Relevant for some diversified semiconductor and systems names.
- Margin expansion and contracted demand are the key value drivers.
- Most AI companies remain pre-profit and are valued on revenue.

Capacity & Contract Value

INFRASTRUCTURE

Compute, data centres & power

- Infrastructure is valued on contracted capacity and order book.
- Firm power access and utilisation are central to the multiple.
- Capital intensity and depreciation cap the multiple below software.
- Best assessed alongside long-dated revenue and supply agreements.

Source: Windsor Drake valuation methodology. See appendix.

Appendix: Sources & Methodology (Part 1)

WINDSOR DRAKE

Institution	Report / Source	Date
McKinsey & Company	*The Economic Potential of Generative AI*	2025
McKinsey & Company	*Where AI Will Create Value, and Where It Won't*	2026
McKinsey & Company	*Global Private Markets Report 2026*	Mar 2026
Bain & Company	*Global Technology Report: AI Compute and Data-Centre Demand*	2025
Bain & Company	*Global Private Equity Report 2026*	Feb 2026
Goldman Sachs	*2026 Global M&A Outlook*	Jan 2026
Morgan Stanley	*Global M&A Activity Outlook 2026*	2026
PitchBook	*Q1 2026 AI VC Trends*	2026
PitchBook / NVCA	*Q1 2026 Venture Monitor*	Apr 2026
CB Insights	*State of Venture and State of AI, Q1 2026*	2026

Appendix: Sources & Methodology (Part 2)

WINDSOR DRAKE

Institution	Report / Source	Date
S&P Global Market Intelligence	*Private Capital and Dry Powder Analyses*	2026
International Monetary Fund	*World Economic Outlook Update and AI Market-Risk Commentary*	Jan 2026
Federal Reserve	*FOMC Statement (May 2026); Summary of Economic Projections (Mar 2026)*	2026
EY	*M&A Outlook 2026*	Jan 2026
Cerebras Systems Inc.	*Form S-1 and IPO prospectus*	2026
Financial press	*Reuters, Bloomberg and CNBC reporting on AI listings and transactions*	2026

VALUATION METHODOLOGY NOTES

Source Standard

Inputs are restricted to top-tier institutions: bulge-bracket banks, the major consultancies, elite data houses, and primary regulatory and filing sources. Reputable financial press is used only for factual reporting of listings and transactions.

Peer Set & Normalisation

Peers are filtered on position in the AI stack, revenue quality and model dependency. Financials are adjusted to a pro-forma basis, separating contracted recurring revenue from pilots and one-time adoption spikes.

Structural Adjustments

Valuations are adjusted for earn-outs, talent-retention pools, compute commitments and lack-of-marketability discounts, typically in the 20% to 35% range.

Synthesis & Attribution

Figures labelled as firm analysis or house estimate, including the roughly 11x AI-native software benchmark and the subsegment multiple ranges, are Windsor Drake's own synthesis of the cited institutional data and public comparables, presented as a house view rather than third-party consensus.