

AI Infrastructure Valuations: Q2 2026

JUNE 2026

Valuation Bifurcation

Windsor Drake's working benchmark for public data-infrastructure software has settled near **11x EV/Revenue**.

- Capital-light platforms (Databricks ~25x, Datadog ~20x) earn software-plus multiples on consumption-led growth.
- Capital-intensive neoclouds re-rate on forward revenue (CoreWeave ~6.5x) as debt and depreciation weigh.
- The defining split is gross-margin quality, not the AI label.
- Buyer and builder appetite are converging on durable, contracted revenue.

Consumption Economics

Usage-based pricing ties revenue directly to AI workload growth.

- Snowflake net retention is **125%** and Databricks net retention exceeds **140%**.
- Metered pricing converts AI adoption into compounding revenue.
- Gross-margin spread (software ~75%+ vs neocloud ~50%) decides the multiple.
- Investors now price net revenue retention as the core quality signal.

The Capex Supercycle

Hyperscaler capital spending underwrites multi-year demand for the substrate beneath AI.

- The big four hyperscalers plan **\$725B** of 2026 capex, up **77%** year on year (Goldman Sachs).
- Amazon ~\$200B, Microsoft ~\$190B, Alphabet \$175-185B, Meta \$115-135B.
- About 75% of 2026 capex is AI-specific, roughly \$450B.
- Goldman models ~\$5.3T of hyperscaler AI-infra spend through 2030.

Macroeconomic Backdrop

Monetary policy is a cautious but constructive tailwind for long-duration assets.

- The Fed held the funds range at **3.50%-3.75%** in April 2026 and added an easing bias.
- The March dot plot points to no more than one further 2026 cut.
- Lower discount rates disproportionately lift long-duration compute and software.
- Abundant private capital is competing hard for scarce, scaled assets.

IPO Window Reopens

The public listing window reopened decisively for AI infrastructure.

- **Cerebras** priced at \$185 and surged 100%+ on its May 2026 Nasdaq debut.
- **CoreWeave** (Mar 2025) now trades near a \$62B market cap.
- Lambda and Crusoe are lining up listings; Databricks eyes an H2 2026 filing.
- The window rewards scale, contracted backlog and a credible margin path.

The Private AI Premium

Private AI-native leaders are running ahead of public software.

- **Databricks** raised at a **\$134B** valuation, about 25x annualised revenue.
- Inference platforms (Fireworks ~\$15B, Baseten ~\$11B) clear roughly 18-19x.
- Private marks are widening against public peers, not converging.
- A thin supply of scaled, independent assets sustains the premium.

Strategic Consolidation

Incumbents are buying the data substrate for enterprise AI.

- **IBM closed its \$11B acquisition of Confluent** in March 2026.
- Real-time data and pipelines are being bought as core AI infrastructure.
- Incumbents are buying, not building, the substrate for enterprise AI.
- Commoditising layers face a consolidate-or-be-consolidated choice.

Private Capital Pressure

Record private capital is intensifying competition for quality assets.

- About **\$3.7T** of global PE dry powder is seeking deployment.
- PE invested **~\$45.7B** in US data centers in 2025, a five-year high (S&P Global).
- The **\$40B** Aligned Data Centers sale set a record for the asset class.
- Sponsors are concentrating capital in larger, platform-scale infrastructure.

1. Consumption Durability Wins

Net revenue retention is the metric that gates premium multiples in AI infrastructure; spend that compounds with workloads is what investors pay for.

- Snowflake **125%** and Databricks **>140%** net retention set the bar.
- Instrument expansion revenue and report progress monthly.

2. Gross Margin Is Destiny

Software economics earn 2 to 3 times the multiple of resold compute; protecting margin is the clearest valuation lever.

- Capital-light platforms trade **10-25x**; compute resale **6-10x** forward.
- Climb from raw GPU resale toward higher-margin managed services.

3. Rule of 40 Still Decides

Growth plus margin remains the single best predictor of a premium multiple across data and infrastructure software.

- Clearing the threshold earns **2 to 3 times** the multiple of peers.
- Balance consumption growth with a credible path to margin.

4. Capital Intensity Is Penalised

Debt and depreciation compress forward multiples for asset-heavy compute, even at triple-digit growth.

- CoreWeave carries **>\$20B** of debt against a single-digit forward multiple.
- Favour asset-light or contracted structures to protect equity value.

5. Contracted Backlog De-Risks

Multi-year, take-or-pay backlog is what underwrites premium compute valuations and makes an asset financeable.

- CoreWeave's **~\$100B** backlog anchors its equity re-rating.
- Convert pipeline into contracted, durable, financeable revenue.

6. Prepare for Competitive Tension

With about \$3.7T of dry powder and incumbents buying capability, the prepared asset captures the scarcity premium.

- Clean financials, documented architecture, and a defensible data room.
- Map specific capability gaps for each of your top five acquirers.

Founder FAQs: Valuations, Timing & Strategy

WINDSOR DRAKE

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

Q1 Which valuation metric applies to my business?

Use **EV/Revenue** for high-growth software and inference, **EV/EBITDA** for mature, cash-generative platforms, and **contracted-EBITDA or capacity** metrics for physical data centers. Always reference the correct subsegment cohort, never a broad average.

Q3 Why does consumption pricing matter so much?

It ties revenue directly to workload growth, so net retention of **125% to 140%+** compounds automatically as AI adoption scales. Investors pay up for that durability, but only where gross margins hold above the cost of serving.

Q5 When is the optimal time to run a process?

After demonstrating **4 to 6 quarters** of predictable consumption growth and contracted backlog, while still holding 12-18 months of runway. Negotiating from strength, not necessity, is what captures the scarcity premium.

Q7 Is the IPO window a viable alternative to M&A?

It reopened decisively (Cerebras, CoreWeave) but favours scale, contracted backlog and a credible margin path. For sub-scale assets, a strategic sale often delivers a superior risk-adjusted outcome via control premiums.

Q2 What are the key subsegment ranges right now?

Private AI-native data platforms trade **20-25x**, inference platforms **15-19x**, observability and data ops **9-20x** in public markets, and GPU cloud **6-10x** on forward revenue. The gap reflects gross margin and capital intensity, not the AI label.

Q4 How do public and private valuations compare?

Unlike mature software, the private AI premium is **widening**: Databricks trades near **25x** versus public data platforms at **9-10x**. Quality private rounds clear public marks, while non-AI laggards see flat marks and become M&A targets.

Q6 Who are the most active buyers today?

Hyperscalers and incumbents buying the data substrate (IBM/Confluent), **PE platforms** deploying about \$3.7T, and **infrastructure funds** acquiring contracted compute and powered capacity at record volumes.


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

Protect **gross margin**, sustain **NRR >120%**, contract the backlog, and clear the **Rule of 40**. Expect **earn-outs** to bridge gaps on unproven AI capability; they are standard structure for early-traction infrastructure assets, paid over 12-24 months.

Q2 2026 Valuation Landscape Overview

Premium multiples cluster in capital-light AI software; compute-intensive operators re-rate on forward revenue.

EV / Revenue Multiple by Representative Company (x)



Deep bifurcation: private AI-native platforms (Databricks ~25x) and inference startups (~18-19x) sit far above public data-infrastructure software (~9-10x), while neoclouds re-rate to single digits on forward revenue. (P) private last-round implied; (F) forward.

PUBLIC DATA-INFRA BENCHMARK

~11x

Windsor Drake's working forward EV/Revenue benchmark for public data-infrastructure software.

PRIVATE AI-NATIVE PREMIUM

~25x

Databricks' Series L implies roughly 25x annualised revenue, well above public peers.

HYPERSCALER 2026 CAPEX

\$725B

Planned 2026 capital spend, up 77% year on year, underwriting infrastructure demand (Goldman Sachs).


Key Driver

Capital-light, consumption-led revenue earns software-plus multiples; capital-intensive compute is priced on contracted backlog and margin, not topline alone.

Public vs Private: A Widening AI Premium

Unlike mature software, the AI-infrastructure private premium is widening as scaled, AI-native assets stay private.

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



PRIVATE / PUBLIC SPREAD

~10x

Up from about 4x in 2023, the inverse of mature software, where the gap compresses.

PUBLIC BENCHMARK ANCHOR

~11x

Public data-infrastructure software now sets the floor, not the ceiling, for the category.

PRIVATE CATALYST


Scarcity

A thin supply of scaled, independent AI-native assets sustains the premium.

A diverging market: public data-infrastructure software has re-rated steadily to about 11x, but private AI-native leaders have run to 20-25x, widening the spread from roughly 4x to about 10x. Windsor Drake analysis.

Exit Valuation Multiples by Subsegment

A sharp bifurcation persists between capital-light AI software (10-25x) and capital-intensive compute (6-10x forward).



KEY OBSERVATION

The market pays for recurring, consumption-led revenue and software gross margins, not for resold compute. The multiple gap between AI-native data platforms and GPU cloud reflects margin quality and capital intensity, not demand.

Valuation Multiple Drivers: Expansion vs. Compression

Net expansion to roughly 11x is driven by the capex supercycle and consumption durability, partly offset by capital intensity and commoditisation.



NET EXPANSION OF +3.0X

The capex supercycle and consumption-led revenue durability outweigh a combined 1.5x drag from capital intensity, GPU commoditisation and power constraints. The bridge reflects Windsor Drake analysis of the cited institutional data.

Capital Markets: IPO Window & Listing Benchmarks

WINDSOR DRAKE

The 2025-26 listing class reopened the AI-infra window: scale, backlog and a margin path gate access.

Cerebras (Nasdaq: CBRS)

Priced at **\$185** in May 2026, raising about **\$5.55B**; shares surged 100%+ on debut.

- Kicked off the AI-infrastructure listing wave of 2026.
- Demand reflected scarcity of credible pure-play AI silicon.
- A premium debut, but one tied to concentration and capex risk.

CoreWeave (Nasdaq: CRWV)

IPO March 2025 at \$40 (~\$23B); now trades near a **\$62B** market cap.

- 2025 revenue \$5.1B (+170%); 2026 guidance \$12-13B.
- A ~\$100B contracted backlog anchors the equity story.
- >\$20B of debt keeps the forward multiple in single digits.

Lambda & Crusoe

The neocloud pipeline is filling behind the leaders.

- Lambda (Nvidia-backed, \$1.5B raised Nov 2025) eyes an H1 2026 listing.
- Crusoe targets a listing around a ~\$13B valuation.
- The window favours contracted, financeable capacity.

Databricks (Private)


The most-watched private candidate is staying private, for now.

- Run-rate **\$5.4B** (+65%), net retention >140%, free-cash-flow positive.
- Series L at a **\$134B** valuation, about 25x revenue.
- Analysts expect an H2 2026 S-1 filing.

Geographic Valuation Variations

AI-infrastructure value is concentrated in North America; sovereign-AI capital is the fastest-growing pool elsewhere.

Share of Global AI Infrastructure Investment (%)



NORTH AMERICA

Premium

The deepest capital markets and the bulk of hyperscaler capex and neocloud capacity.

EUROPE

Sovereign

Data-sovereignty and regulatory initiatives are pulling in new state-backed infrastructure capital.

APAC & GULF

Growth


Sovereign-AI buildouts (MGX and peers) are among the largest new sources of infrastructure capital.

Concentration premium: North America dominates AI-infrastructure investment on the strength of hyperscaler capex, deep capital markets and the leading neoclouds. Sovereign-AI initiatives in the Gulf, Europe and APAC are the fastest-growing pools of new capital. Windsor Drake analysis.

GPU & Accelerated Cloud: The Neoclounds

Neoclounds trade on contracted backlog and forward revenue; capital intensity caps the multiple despite hypergrowth.

EV / Revenue Multiple Range (x)



Valuation Drivers

Contracted Backlog

Multi-year, take-or-pay contracts, such as CoreWeave's roughly \$100B backlog, convert demand into financeable revenue and underwrite the equity story.

Capital Intensity

GPU capex, depreciation and debt loads above \$20B compress forward multiples even at triple-digit growth; equity value hinges on utilisation and the cost of capital.

Margin & Differentiation

Improving chip access erodes scarcity; orchestration software, networking and secured power separate durable platforms from pure GPU resellers.


KEY OBSERVATION

Neoclounds are priced like contracted infrastructure, not software: backlog quality and cost of capital, not topline growth, drive the multiple.

Model-Serving & Inference Platforms

Inference platforms command software-like multiples on explosive consumption growth, tempered by GPU-driven gross margins.

EV / Revenue Multiple Range (x)



Valuation Drivers

Consumption Explosion

Token and API volumes are compounding; Fireworks reached about \$800M and Baseten about \$600M annualised revenue within a few quarters, justifying premium last-round marks.

Gross-Margin Reality

GPU costs in cost of goods hold gross margins near 50%, below the 70%-plus of subscription software; durable winners must climb the software-margin curve.

Strategic Demand

Enterprises and labs are buying managed inference rather than building it; capability M&A in serving, routing and optimisation is accelerating.


KEY OBSERVATION

Inference is the fastest-repricing AI-infrastructure subsegment. The market is underwriting consumption growth ahead of margin; gross-margin convergence will decide which marks hold.

Vector Databases & Retrieval

Retrieval infrastructure is strategically central but commercially early; multiples run high on thin revenue bases.

EV / Revenue Multiple Range (x)



Valuation Drivers

RAG as Default

Retrieval-augmented generation has made vector search a default layer of the AI stack, expanding the addressable workload across nearly every enterprise deployment.

Feature vs Standalone

Incumbents such as MongoDB, Postgres and Elastic are embedding vector search, pressuring standalone players; differentiation now rests on scale, latency and cost.

Consolidation Pressure

Thin revenue and crowded competition invite M&A; well-capitalised data platforms are the natural acquirers of leading retrieval engines.


KEY OBSERVATION

Vector databases are strategically essential but face a feature-versus-product squeeze. Premium marks rest on category leadership, not yet on scaled revenue.

Data-Pipeline & Lakehouse Infrastructure

The lakehouse and streaming layer is the highest-value AI-data substrate, led by consumption platforms with strong retention.

EV / Revenue Multiple Range (x)



Valuation Drivers

Consumption Retention

Net retention of 125% at Snowflake to above 140% at Databricks compounds revenue as AI workloads scale, the single strongest driver of premium data-platform multiples.

AI Workload Gravity

Training and inference pull data gravity toward unified lakehouse platforms; owning the data layer increasingly means owning the AI control plane.

Strategic Consolidation

Incumbents are buying the substrate: IBM acquired Confluent for \$11B to make real-time data the engine of enterprise AI and agents.

VALUATION DRIVER


Data platforms monetise every downstream AI workload through consumption, the closest thing in AI infrastructure to a compounding software annuity.

Sources: PitchBook; CB Insights; S&P Global Market Intelligence; Windsor Drake analysis. See appendix.

AI Accelerators & Custom Silicon

Scarcity of credible alternatives to the GPU incumbent supports premium valuations for differentiated silicon.

EV / Revenue Multiple Range (x)



Valuation Drivers

Scarcity Premium

Cerebras' May 2026 IPO priced at \$185 and surged 100%-plus, evidence of deep investor appetite for credible alternatives to the dominant GPU supplier.

Concentration Risk

Heavy customer and end-market concentration, and dependence on the same capex cycle, inject volatility that tempers the through-cycle multiple.

Hyperscaler Custom Silicon

In-house accelerators such as TPUs and peers both validate and threaten merchant silicon, reshaping where value accrues across the stack.


KEY OBSERVATION

Silicon valuations price scarcity and optionality on the buildout. Concentration and capex sensitivity make the premium real but fragile.

Data Center & Power Infrastructure

Physical capacity is valued on contracted EBITDA and power access, not revenue multiples; private capital is setting records.

EV / EBITDA Multiple Range (x)



Valuation Drivers

Contracted Capacity

Long-term hyperscaler leases convert data centers into bond-like, contracted cash flows that infrastructure funds underwrite at premium EBITDA multiples.

Power Is the Constraint

Access to power and interconnection, not land or shells, is now the binding constraint and the primary source of scarcity value.

Record Private Capital

PE invested about \$45.7B in US data centers in 2025, a five-year high; the \$40B Aligned Data Centers sale reset the ceiling for the asset class.

KEY OBSERVATION

Data centers trade as contracted infrastructure: power access and lease quality, not technology, set the multiple. Capital is abundant and chasing scarce, powered capacity.

Comparable Transaction Analysis Framework

WINDSOR DRAKE

A quality-of-revenue and capital-intensity filter, not broad sector codes.

1. Select Peer Set

Identify genuinely comparable assets by model (software versus compute resale), revenue quality, gross margin and capital intensity, never by the AI label alone.

2. Normalise Metrics

Adjust to a pro-forma basis: separate committed from on-demand revenue, normalise for GPU-cost in cost of goods, and reconcile net retention across cohorts.

3. Adjust for Structure

Account for contracted backlog, debt load, earn-outs and the stock-versus-cash mix that pull headline valuation away from underlying economic value.

Proprietary Transaction Index

Calibration draws on Windsor Drake's proprietary index of **184 verified and reported transactions (2020 to 2026)**, a living comparable-set database refreshed each quarter.

Quality-of-Revenue Filter

Peer selection prioritises consumption durability (net retention), gross-margin profile (software ~75%+ versus compute ~50%) and customer concentration.

Capital-Intensity Adjustment

A discount layer is applied for balance-sheet-heavy compute (depreciation and debt); a premium for asset-light, software-margin platforms.

Control Premium Calibration

Indications include a control-premium layer, typically **20-30%**, where cost and revenue synergies can be concretely underwritten.

184

INDEXED TRANSACTIONS

2020-26

INDEX COVERAGE

Strategic Acquirer Mapping by Subsegment

WINDSOR DRAKE

Hyperscalers and incumbents buy the data substrate; PE targets contracted cash flow; infrastructure funds buy capacity.

| Subsegment | Hyperscalers & Strategics | Private Equity | Infrastructure Funds |
|---------------------------|---|--|---|
| GPU & Cloud | MODERATE Capacity and capability tuck-ins; mostly build in-house. | HIGH Take-privates and roll-ups of contracted compute. | HIGH Contracted capacity held as a core infrastructure asset. |
| Inference | HIGH Buy managed serving to embed in platforms. | MODERATE Back scaled, margin-improving platforms. | LOW Outside the contracted-asset mandate. |
| Vector / Retrieval | HIGH Absorb retrieval into data and AI stacks. | MODERATE Buy-and-build within data tooling. | LOW Not an infrastructure-fund asset. |
| Data Platforms | HIGH Core substrate; IBM/Confluent is the template. | HIGH Cash-generative consumption platforms. | LOW Software, not physical capacity. |
| Custom Silicon | HIGH Strategic supply and vertical integration. | LOW Capex and concentration deter sponsors. | MODERATE Selective, via compute platforms. |
| Data Centers | MODERATE Self-build plus selective capacity buys. | HIGH Platform roll-ups of powered capacity. | HIGH Prime contracted-infrastructure target. |

High Activity / Priority Target

Moderate / Opportunistic

Low Activity / Niche Only

Hyperscalers & Strategics as Buyers

WINDSOR DRAKE

Incumbents are buying the data and AI substrate rather than building it.

| Buy vs. Build on the Substrate

The internal build cycle for real-time data, retrieval and serving is too slow against AI-native challengers. Incumbents are acquiring the substrate outright; IBM's \$11B purchase of Confluent is the clearest template.

| Acquisition Patterns

Strategics favour capability deals that plug directly into existing platforms, often after a commercial agreement or partnership de-risks the technology and confirms the customer overlap.

\$11B

IBM / CONFLUENT

Buy

OVER BUILD, ON DATA

Priority: Real-Time Data

Streaming and pipelines are valued as the engine of enterprise AI and agents, not as commodity middleware.

Priority: Retrieval & Context

Vector and retrieval layers are bought to ground enterprise AI in proprietary data with governance and low latency.

Priority: Managed Inference

Serving and optimisation capability is acquired to embed AI features without standing up GPU operations internally.

Priority: Governance & Lineage

Data-governance, lineage and observability assets are valued as the control plane for trustworthy enterprise AI.

Private Equity & Infrastructure Capital

WINDSOR DRAKE

Record dry powder is concentrating in contracted, cash-generative AI infrastructure.

Deployment Pressure

With roughly \$3.7T of global dry powder, sponsors face acute pressure to transact. PE invested about \$45.7B in US data centers in 2025, a five-year high and 72% of all sector investment (S&P Global).

Contracted-Cash-Flow Thesis

Long-term hyperscaler leases turn data centers into bond-like assets; the \$40B Aligned Data Centers sale, the largest ever, exemplifies infrastructure-fund appetite for powered capacity.

~\$3.7T

GLOBAL DRY POWDER

\$45.7B

PE INTO US DATA CENTERS, 2025

Ideal Target Profile

Sponsors prioritise contracted backlog, durable consumption revenue and a clear margin path, the profile that supports leverage capacity.

Value-Creation Playbook

Power procurement, utilisation optimisation, buy-and-build of fragmented capacity, and mix-shift toward higher-margin managed services.

Software-Platform Targets

Cash-generative data and tooling platforms with high net retention are prime buyout candidates as growth rates normalise.

Deal Structure Trends

Larger, debt-financed platform deals dominate; earn-outs bridge valuation gaps on unproven AI capability.

Competitive Moats Driving Premium Valuations

WINDSOR DRAKE

Multiples above 15x revenue are reserved for assets with structural defensibility in the AI stack.

Data Gravity

ASSET VALUE: HIGHEST

Workloads migrate to where data lives

- Owning the data layer pulls training and inference onto the platform.
- Switching costs compound as pipelines and governance entrench.
- Each workload adds consumption revenue at high margin.
- **Action:** make the platform the system of record for AI data.

Consumption Retention

SCALE VALUE: HIGH

Revenue that compounds with usage

- Net retention of 125-140%+ turns adoption into automatic growth.
- Aligns vendor revenue with customer value realisation.
- Directly lifts the Rule of 40 score that gates premium multiples.
- **Action:** instrument and expand usage across every workload.

Contracted Backlog

BARRIER VALUE: HIGH

Financeable, multi-year revenue

- Take-or-pay contracts underwrite capex and equity value.
- A roughly \$100B backlog anchors the largest neocloud.
- De-risks the asset for lenders and acquirers alike.
- **Action:** convert pipeline into contracted, durable revenue.

Power & Supply Access

SCARCITY VALUE: HIGH


The binding constraint of the cycle

- Secured power and interconnection are the real scarcity.
- Priority accelerator allocation differentiates compute platforms.
- Hard to replicate on any near-term timeline.
- **Action:** lock long-dated power and supply agreements early.

Rule of 40 Performance Distribution

Clearing the Rule of 40 unlocks a 2 to 3 times multiple premium; capital-intensive laggards trade at deep discounts.

EV / Revenue Multiple vs Rule of 40 Score



TOP QUARTILE (SCORE >50)

18x+

AI-native data and inference leaders; the primary premium-M&A targets.

RULE OF 40 MET (40 TO 50)

11-18x

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

BELOW THRESHOLD (<40)

4-11x

Transition and capital-intensive zones; vulnerable to consolidation.

The Rule of 40 still decides: data and infrastructure software clearing the threshold trade at roughly 2 to 3 times the revenue multiple of those below it. AI-native platforms that pair high growth with consumption margins (top right) command the richest marks. Windsor Drake analysis.

Cross-Border M&A Considerations

WINDSOR DRAKE

Export controls and sovereign-AI policy are the primary deal risk, and the primary arbitrage, in 2026.

Export Controls & Security Review

Semiconductor export controls, national-security review (CFIUS and FDI regimes) and data-sovereignty rules materially complicate cross-border AI-infrastructure deals and lengthen approval timelines.

Sovereign-AI Capital

Gulf and European state-backed vehicles, MGX and peers among them, are deploying record capital into compute and data centers, creating both buyers and partners for cross-border assets.

12-18mo

CROSS-BORDER CYCLE

Sovereign

NEW CAPITAL POOLS

Extended Timelines

Clearance for cross-border infrastructure and silicon deals runs materially longer than domestic transactions; ensure runway to withstand delay without losing leverage.

Power & Siting Diligence

Cross-border data-center deals turn on power availability, grid access and permitting in the target jurisdiction, often the gating diligence item.

Chip-Access Conditionality

Allocation of constrained accelerators can be tied to jurisdiction and end-use; model the supply impact of export rules before signing.

Local Partners

Retain local management and sovereign partners to navigate permitting, energy procurement and post-close regulatory nuance.

Exit Valuation Optimisation Strategies

WINDSOR DRAKE

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

1. Protect Gross Margin

6 TO 12 MONTHS PRE-EXIT

Defend software-level economics

- Move from raw compute resale to higher-margin managed services.
- Optimise GPU utilisation and reserved-capacity costs.
- Target software-grade gross margin where the model allows.
- Evidence margin durability with clean cohort data.

2. Contract the Backlog

FINANCEABILITY

Turn pipeline into committed revenue

- Convert usage into multi-year, take-or-pay commitments.
- Lengthen contract duration to extend revenue visibility.
- Reduce concentration so no single customer dominates.
- Document backlog quality for lenders and buyers.

3. Prove Consumption Retention

PREMIUM TIER

Show revenue that compounds

- Sustain net revenue retention above **120%** via expansion.
- Separate committed from on-demand revenue clearly.
- Tie net retention directly to AI-workload growth.
- Track the metric monthly with board-level visibility.

4. Strategic Narrative

COMPETITIVE TENSION

Frame the asset as core substrate

- Position as the control plane, not a point tool.
- Present quantified synergy cases across revenue and cost.
- Map specific capability gaps for the top five acquirers.
- Run a structured process to manufacture competitive tension.

Positioning for Strategic Acquisition

WINDSOR DRAKE

Strategic value is driven by substrate fit, integration ease and synergy density.

Substrate Fit

Demonstrate that your data, retrieval or serving layer fills a specific, declared gap in the acquirer's AI stack, making the buy-versus-build decision self-evident.

Integration Ease

Acquirers pay clear premiums for plug-and-play infrastructure: clean APIs, documented architecture, portable deployment and audited financials.

Synergy Density

Quantify the cross-sell into the acquirer's base, and model the cost synergies from shared compute, power and data infrastructure, to support a higher multiple.

Strategic Buyer Mapping

Run a structured gap analysis and map your capabilities directly to each buyer's declared AI-infrastructure deficits.

Proof-of-Integration

Provide reference architectures and benchmarks that demonstrate speed-to-value inside the acquirer's environment, pre-empting technical diligence.

Synergy Quantification

Model top-line and cost impact explicitly in the management presentation to anchor the valuation conversation on hard numbers.

Comprehensive VDR Readiness

Build a data room addressing security, data rights, power contracts and customer concentration before the first buyer engagement.

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 the capex supercycle, reopened listings and abundant capital still align.



Readiness & Hygiene

Q3 2026

- Audit completion to PCAOB standard
- Data-rights, security and model-governance review
- Gross-margin and utilisation optimisation
- Clean up the cap table and option pool

KEY MILESTONE

Clean security and data-rights review



Strategic Positioning

Q4 2026

- Launch dual-track process preparation
- Build the strategic buyer-targeting list
- Draft the CIM and management presentation
- Lock key-employee 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

2026 Valuation Forecast Scenarios

WINDSOR DRAKE

With the public data-infra benchmark near 11x, forward trajectories diverge on rates, capex durability and the pace of AI monetisation.



BULL CASE

15x

Key Drivers

- Aggressive Fed cuts beyond one
- Capex supercycle and AI monetisation broaden
- The listing window stays wide open

STRATEGY: ACCELERATE GROWTH

BASE CASE

12x

Key Drivers

- Steady rate normalisation, one cut
- Durable capex and consumption growth
- Selective IPOs and strategic consolidation

STRATEGY: BALANCE GROWTH & MARGIN

BEAR CASE

8x

Key Drivers

- Inflation resurgence or rate holds
- Capex digestion and an AI-spend pause
- A power or financing shock freezes deals

STRATEGY: PROTECT MARGIN & CASH

Emerging Opportunities & Buyer Trends

WINDSOR DRAKE

Capital is flowing into the connective tissue of an AI-first compute and data stack.

AI Data Substrate

Real-time pipelines, retrieval and governance, the layer that grounds enterprise AI in proprietary data, are the most strategically contested assets in the stack.

Inference Optimisation

Routing, caching, compilation and GPU-efficiency tooling are emerging as a high-value layer as enterprises seek to cut the cost of AI in production.

Power & Cooling

Energy procurement, on-site generation and liquid cooling are scarce capabilities; Eaton's \$9.5B purchase of Boyd Thermal signals value migrating to the physical layer.

Capability Acquisitions

Incumbents are prioritising data, retrieval and serving capability over raw distribution, buying technology to modernise their AI stacks.

Infrastructure-Fund Roll-Ups

Sponsors are consolidating powered data-center capacity and fragmented compute into platform-scale assets.

Sovereign Champions Go Global

Gulf, European and APAC state-backed vehicles are acquiring assets to build domestic AI capacity and capture scarcity value.

Market Intelligence

Top-tier forecasts point to a multi-year capex supercycle. Demand for scaled, independent AI-infrastructure assets is outpacing a supply that has not kept pace.

M&A Case Study: IBM & Confluent

WINDSOR DRAKE

The defining data-infrastructure acquisition of the cycle, and the playbook it sets for founders.

The Substrate Playbook

IBM's acquisition of **Confluent**, valued at **\$11B** and completed in **March 2026** at \$31 per share, roughly 10x revenue, is the defining data-infrastructure deal of the cycle. It confirms that real-time data and pipelines are now bought as core AI infrastructure, not middleware.

Strategic Rationale

- **Real-time data:** positioning streaming as the engine of enterprise AI and agents.
- **Buy over build:** acquiring a category leader rather than replicating Kafka at scale.
- **Consumption base:** layering a durable, expanding data-flow franchise onto an incumbent platform.

Implications for Founders

Own the Data Substrate

The layer that moves and grounds enterprise data is now strategic, and acquirable. Founders who own a durable slice of the AI data substrate sit in the path of incumbent demand, and the window to choose is narrowing.

Consumption Proves Durability

Confluent's expanding, **consumption-led revenue** underwrote the multiple. Recurring, usage-based revenue with strong retention is what converts strategic interest into a premium price.

Platform vs. Point Solution

Assets framed as core **infrastructure** capable of absorbing into a larger platform trade at clear premiums. Clean architecture and integration readiness are themselves valuation levers.

Valuation Methodology: Choosing the Right Metric

WINDSOR DRAKE

The right metric depends on business model, margin profile and capital intensity.

EV / Revenue

10-25X+

High-growth software & inference

- Applied where consumption growth outruns near-term profit.
- Software revenue (~75%+ margin) valued far above resold compute.
- Net retention and the Rule of 40 set the place in the range.
- Best suited to data platforms, inference and retrieval.

EV / EBITDA

12-30X

Contracted & cash-generative

- The standard for data centers and mature, scaled platforms.
- Many compute assets shift from revenue to EBITDA as they scale.
- Contract quality and utilisation are the key value drivers.
- Captures the cash-flow reality of infrastructure.

EV / Contracted Backlog

COMPUTE

Forward, committed revenue

- Used where take-or-pay backlog dominates the story.
- Anchors neocloud valuations to financeable revenue.
- Weighs counterparty quality and contract duration.
- Pairs with cost of capital to value the equity.

Capacity & Power

PHYSICAL

Megawatts and interconnection

- Data centers valued on contracted megawatts and power access.
- Secured energy and interconnection drive scarcity value.
- Lease quality and tenant credit anchor the multiple.
- Increasingly the binding constraint on value.

Appendix: Sources & Methodology (Part 1)

WINDSOR DRAKE

| Institution | Report / Source | Date |
|--------------------------------|---|----------|
| Goldman Sachs | <i>AI Capex Outlook: Tracking the Build-Out (hyperscaler capex; ~\$5.3T through 2030)</i> | 2026 |
| Morgan Stanley | <i>Financing the AI Infrastructure Build-Out</i> | 2026 |
| McKinsey & Company | <i>Global Private Markets Report 2026</i> | Mar 2026 |
| McKinsey & Company | <i>The cost of compute: data-center capital needs to 2030</i> | 2026 |
| Bain & Company | <i>Global Private Equity Report 2026 (Outlook: Gaining Traction)</i> | Feb 2026 |
| S&P Global Market Intelligence | <i>US data-center deals hit five-year high on PE surge</i> | May 2026 |
| S&P Global Market Intelligence | <i>Private Equity Dry Powder and Deployment</i> | 2026 |
| PitchBook | <i>AI & Data Infrastructure Comp Sheet and Valuation Guide (Q1 2026)</i> | 2026 |
| CB Insights | <i>State of AI 2025 (infrastructure)</i> | 2026 |
| Bloomberg | <i>AI infrastructure funding and valuations (Fireworks, Baseten)</i> | May 2026 |

Appendix: Sources & Methodology (Part 2)

WINDSOR DRAKE

| Institution | Report / Source | Date |
|-----------------------|---|---------------------|
| Federal Reserve | FOMC Statement (Apr 2026); Summary of Economic Projections (Mar 2026) | 2026 |
| PwC | US Technology Deals 2026 Outlook | 2026 |
| EY | M&A Outlook 2026 (Deal Barometer) | Jan 2026 |
| Databricks Inc. | Series L financing and revenue run-rate disclosures | Dec 2025 - Feb 2026 |
| IBM Corporation | SEC filings and completion of the Confluent acquisition | Mar 2026 |
| Cerebras Systems Inc. | SEC Form S-1 and IPO pricing | May 2026 |
| Company filings | CoreWeave, Nebius, Snowflake, MongoDB and Datadog quarterly results | 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. Boutique and market-report vendors are excluded.

Structural Adjustments

Compute valuations are adjusted for contracted backlog, debt and depreciation; physical-capacity valuations rest on contracted EBITDA and power access; private marks are adjusted for lack of marketability, typically 20 to 30%.

Peer Set & Normalisation

Peers are filtered on business model, consumption durability (net retention), gross-margin profile and capital intensity. Financials are adjusted to a pro-forma basis excluding one-time items and stock-based compensation.

Synthesis & Attribution

Figures labelled as firm analysis or house estimate, including the roughly 11x public data-infrastructure benchmark, are the firm's own synthesis of the cited institutional data, presented as a house view rather than third-party consensus.