

Cerebras Systems Inc.

NASDAQ: CBRS

Wafer-Scale Inference at a Hyperscale Price

Last Price (27-May-26)	\$241.71	FY25 Revenue	\$510 m
Market Cap (basic)	~\$53.1 bn	FY25 Net Income (GAAP)	\$238 m
Market Cap (FD)	~\$60–66 bn	RPO / Backlog	\$24.6 bn
52-Week Range	\$185 – \$386	Trailing P/S	~104x
IPO Date / Price	14-May-2026 / \$185	Trailing P/E	~219x

ANALYTICAL VIEW

At \$241.71, CBRS is priced for execution on a \$20 bn OpenAI contract, sustained non-UAE customer diversification, and an inference moat that holds against CUDA + AMD MI400. Intrinsic value on an RPO-anchored framework with normalized margins is **\$100–150 per share**. Bull case (perfect execution + 2nd hyperscaler by 2027) lands at **\$200–275**. Current price sits at the top of the bull range. Cleanest re-entry: post-lockup unwind (Nov-2026) or after two clean quarters with disclosed third major customer.

Prepared 28 May 2026 • Verify live prices before sizing • Not investment advice

Executive Summary

Cerebras Systems IPO'd on 14-May-2026, pricing at \$185 (above a twice-raised range of \$150–160) and opening at \$385. The stock has since re-rated lower to \$241.71, down approximately 22% year-to-date and 37% from its debut high. The IPO debuted as the largest US technology offering of 2026, raising \$5.5 bn at a fully diluted valuation near \$66 bn on day one.

The investment debate is genuinely two-sided. On the positive side: Cerebras is the only credible non-GPU architecture at scale for frontier-model inference, with demonstrated 5–10x throughput advantage on large-language-model workloads; the OpenAI master relationship agreement (~\$20 bn over multi-year term) is the most significant architectural validation in AI infrastructure since Microsoft's first OpenAI investment; and the \$24.6 bn remaining performance obligation (RPO) provides forward revenue visibility exceeding 48x trailing sales. On the negative side: customer concentration is acute (UAE entities accounted for 86% of FY25 revenue); CUDA software lock-in remains a real moat for Nvidia; hyperscalers are increasingly building their own custom inference silicon (TPU, Trainium, MTIA, Maia); and the entry multiple of ~104x trailing P/S leaves no margin of safety for execution slips.

Our framework triangulates fair value across three lenses: trailing-multiple comparison, RPO-anchored DCF with normalized margins, and Street consensus. The base case lands at **\$100–150 per share**, implying the stock is currently trading 60–140% above intrinsic value under conventional assumptions. The bull case (\$200–275) requires perfect execution on the OpenAI ramp, a second non-UAE hyperscaler customer signed within 18 months, and the inference moat holding against CUDA-on-Blackwell economics.

Key takeaway

The technical and commercial story is real. The price is not. Re-entry windows: post-lockup supply unwind in November 2026 (preferred), or after two clean quarters with a disclosed third 10%+ customer. Long here requires you to underwrite all four of: OpenAI delivery schedule, second hyperscaler signing, CFIUS clearance on G42, and inference moat durability. Any one slipping compresses the multiple materially.

1. Company Overview & IPO Context

Cerebras Systems is a Sunnyvale, California-based AI infrastructure company founded in 2015 by Andrew Feldman, Gary Lauterbach, and Michael James. The company designs and manufactures the Wafer-Scale Engine (WSE), a wafer-sized AI accelerator, and sells it both as on-premise systems (CS-3) and as a cloud-based inference / training service. Cerebras is led by Andrew Feldman (CEO) and is funded in part by G42 of the United Arab Emirates, which holds a substantial pre-IPO stake.

IPO Mechanics

Cerebras priced its IPO at \$185 per share on 14 May 2026 — above a twice-raised marketing range of \$150–160. The offering of 30 million shares raised \$5.55 bn in gross proceeds, making it the largest US technology IPO of 2026. The stock opened at \$350 (89% above issue), traded as high as \$385.34 intraday, and closed day one at \$311.07 — implying a market cap of approximately \$66 bn on outstanding shares. Day-two trading saw a roughly 10% decline as profit-taking set in. Over the subsequent two weeks, the stock has continued to re-rate lower to \$241.71 on volume averaging 12.6 million shares per day.

Capital Structure

Item	Detail
Shares outstanding (basic)	~219.6 million
Shares outstanding (fully diluted)	~270–280 million (incl. options, RSUs, warrants)
Market cap (basic, 27-May)	~\$53.1 bn
Market cap (fully diluted)	~\$60–66 bn
Cash & equivalents (post-IPO)	~\$930 m
Total debt	\$262 m
Net cash position	~\$668 m
Lock-up expiry (180 days)	~10 November 2026

The lock-up unwind in November 2026 is a material technical overhang. Pre-IPO investors — Fidelity, Benchmark, Foundation Capital, Eclipse Ventures, and G42 — hold multi-billion-dollar stakes. Comparable high-multiple debuts (Snowflake, Coinbase, Rivian) saw 15–25% drawdowns into and through their lock-up unwinds. The cleanest re-entry window for new institutional positions sits on the other side of this event.

2. Technology & Architecture

Cerebras' core technical insight is architectural rather than process-driven. Rather than cutting a 12-inch silicon wafer into ~70 individual GPU dies (as Nvidia does with the H100/B200), Cerebras keeps the entire wafer intact and uses it as a single chip. This is the Wafer-Scale Engine (WSE).

Wafer-Scale Engine (WSE-3) Specifications

Specification	WSE-3 (Cerebras)	Nvidia H100	Multiple
Die size	46,225 mm ²	814 mm ²	57x larger
Transistor count	4 trillion	80 billion	50x more
AI cores	900,000	~16,000 (CUDA + Tensor)	56x more
On-chip SRAM	44 GB	60 MB	733x more
Memory bandwidth	21 PB/s (on-die)	3.35 TB/s (HBM3)	6,266x higher
Fabric bandwidth	214 Pb/s	900 GB/s (NVLink)	24,000x higher
Process node	TSMC 5nm	TSMC 4N	—

Why the Architecture Matters for Inference

In a conventional GPU cluster, large-model inference is gated by inter-chip communication. When a 70-billion-parameter model runs on 8 H100s, every attention layer requires all-to-all data exchange across the NVLink fabric. That fabric operates at ~900 GB/s — orders of magnitude slower than the on-die memory hierarchy. The interconnect is the bottleneck, not the compute.

By keeping the entire model on a single piece of silicon, Cerebras eliminates the slowest layer of the memory hierarchy entirely. The result on Llama 70B-class inference workloads is throughput in the range of 1,800–2,500 tokens per second, compared to 200–400 tokens per second on equivalent H100 clusters — a 5–10x speedup that is consistent, measured, and verified by third parties.

The strategic implication

Cerebras has architecturally optimized for inference at exactly the moment when inference is becoming the dominant AI compute workload. OpenAI's reported inference spend now exceeds its training spend; the same is broadly true across hyperscalers. Cerebras' advantage is not that it has a better GPU — it is that it does not have a GPU at all. Whether this is a five-year franchise or a transition vehicle depends on customer behaviour, not technology.

Performance-per-Watt Considerations

A second-order benefit, increasingly relevant as data center power becomes the binding constraint on AI buildouts, is performance-per-watt. Because Cerebras eliminates the inter-chip communication overhead, the same inference workload consumes 5–10x less energy than a comparable GPU cluster. In a world where every hyperscaler is power-constrained (see grid interconnect queue stretching to 4–7 years for new US facilities), this is a structural cost advantage that compounds with utilization.

3. Business Model & Customer Concentration

Cerebras monetizes its WSE-3 platform through two channels: direct hardware sales of CS-3 systems (purchase by the customer, deployed on-premise or in a partner data center), and inference cloud services (recurring revenue, customer pays per token or per reserved capacity). The hardware-to-cloud mix is shifting in favor of cloud services as Cerebras builds out its own data center footprint.

Revenue segment	FY25 (\$m)	% of total	FY24 (\$m)	YoY growth
Hardware (CS-3 systems)	358	70.2%	240 (est.)	+49%
Cloud services (inference + training)	152	29.8%	50 (est.)	+204%
Total revenue	510	100.0%	290	+76%

Customer Concentration — The Central Risk

This is the single most important quantitative issue in the Cerebras investment thesis. Per the S-1/A disclosure, UAE entities (principally G42 and its affiliates) represented 86% of FY25 revenue. G42 alone was 87% of H1-2024 revenue. The OpenAI master agreement begins to diversify this concentration from 2026 onward, but the legacy concentration remains acute and is the source of both CFIUS scrutiny and Street skepticism.

Customer	Status	Estimated FY25 contribution
G42 (UAE) and affiliates	Anchor customer; pre-IPO investor	~86% of revenue
OpenAI	\$20 bn multi-year MRA (deliveries 2026–2028)	Material from FY26
AWS (via Bedrock)	Integration March 2026	Early-stage
Other (sovereign, enterprise)	Mixed	<5%

The OpenAI Deal — The Validation

On 13 March 2026, Cerebras and OpenAI signed a Master Relationship Agreement covering 750 megawatts of low-latency AI compute, with deliveries commencing in 2026 and extending through 2028. The total contract value was disclosed initially as 'over \$10 bn'; subsequent disclosures took the figure to approximately \$20 bn across the multi-year term. OpenAI's first model running on Cerebras silicon launched in early 2026, providing public validation of architectural compatibility at production scale.

The strategic significance is that OpenAI — which has charter access to both Nvidia's Blackwell and Rubin roadmaps — chose to commit \$20 bn to a competing architecture. The decision is implicitly a statement that GPUs alone are insufficient for the inference scale OpenAI is building toward. This is the market's most important signal that Cerebras' architectural advantage is commercially significant, not just technically interesting.

AWS Bedrock — The Third Leg

In March 2026, Cerebras inference was integrated into Amazon Bedrock, AWS' managed foundation-model service. This is the first hyperscaler distribution channel for Cerebras outside the direct CS-3 sales model and represents the most credible path to non-UAE, non-OpenAI customer diversification. Early-stage revenue contribution; scale economics will not be visible before FY27.

4. Financial Performance

Cerebras posted reported revenue of \$510 million in FY25, up 76% from \$290 million in FY24, and a GAAP net income of \$238 million — the first full year of profitability in the company's history. Headline metrics are strong, but they require unpacking: the GAAP profit is meaningfully inflated by a non-cash gain on the conversion of pre-IPO forward contracts, and adjusted net income on a non-IFRS basis remains a loss.

Income Statement Item	FY24	FY25	YoY
Revenue	\$290 m	\$510 m	+76%
Hardware	\$240 m (est)	\$358 m	+49%
Cloud services	\$50 m (est)	\$152 m	+204%
Gross profit	\$120 m	\$215 m	+79%
Gross margin	41%	42%	+1 pp
Operating expenses	\$615 m	\$478 m	-22%
GAAP operating income (loss)	(\$495 m)	(\$263 m)	—
Non-cash items (forward contracts, SBC)	(\$5 m)	+\$501 m	—
GAAP net income (loss)	(\$500 m)	\$238 m	Swing
Non-IFRS adjusted net loss	(\$310 m)	(\$45 m)	—
Adjusted EBITDA	(\$265 m)	+\$15 m	Swing

Interpreting the Numbers

Two adjustments matter. **First**, the \$501 million non-cash gain in FY25 relates primarily to the revaluation and conversion of pre-IPO forward contracts on convertible preferred equity. This is a one-time accounting event tied to the IPO event and will not recur. Stripping this out, Cerebras is still operating at a loss on a normalized basis in FY25 — adjusted net loss of approximately \$45 million. **Second**, stock-based compensation continues to run heavy at approximately \$180 million in FY25, reflective of the typical pre-IPO equity loading; this will moderate but remain a material expense through 2027–28.

The cleaner read is the adjusted EBITDA figure: a swing from negative \$265 million in FY24 to slightly positive \$15 million in FY25. The business is at break-even on a cash-operating basis, which is impressive given the capital-intensive ramp, but it is not yet generating the cash flow that would justify a 100x+ trailing sales multiple. That justification has to come from forward backlog conversion — which is the focus of Section 5.

Balance Sheet Position

Post-IPO, Cerebras held approximately \$930 million in cash and cash equivalents against \$262 million of total debt, providing a comfortable net cash position of roughly \$668 million. Combined with \$5.5 billion in IPO proceeds — most of which will be deployed into wafer inventory, data center buildout, and working capital — the company has sufficient runway to execute the OpenAI delivery schedule without near-term capital needs. Free cash flow generation is unlikely to be positive on a sustained basis before FY27 at the earliest, given the lumpy nature of CS-3 production cycles and the working capital needs of multi-year MRA delivery.

5. Remaining Performance Obligation (RPO) Analysis

The single number that drove the 20x IPO oversubscription is the \$24.6 billion remaining performance obligation — contractually committed but unrecognized revenue, equal to approximately 48x FY25 sales. Whether one views CBRS as a value trap or a generational opportunity depends largely on how one models the conversion of this backlog into recognized revenue and, ultimately, cash earnings.

Management's Disclosed Conversion Schedule

Recognition period	% of \$24.6 bn RPO	Implied \$	Annualized revenue
FY26–FY27 (combined)	15%	\$3.69 bn	~\$1.85 bn / year
FY28–FY29 (combined)	43%	\$10.58 bn	~\$5.29 bn / year
FY30 and beyond	42%	\$10.33 bn	Long-tail
Total	100%	\$24.60 bn	—

Implied Revenue Trajectory

FY	Revenue (\$bn)	Growth	Source
FY25A	\$0.51	+76%	Reported
FY26E	\$1.20	+135%	Management guidance + RPO start
FY27E	\$2.30	+92%	Mid-RPO conversion
FY28E	\$4.80	+109%	OpenAI deliveries peak
FY29E	\$5.50	+15%	Steady state
FY30E	\$5.80	+5%	Tail RPO + new customers

These figures assume management's conversion schedule holds without acceleration or delay. Two structural risks need to be flagged. **First**, RPO is not contracted revenue in the legal sense for all line items: portions are subject to delivery milestones, capacity commitments, and termination-for-convenience clauses standard in hyperscale infrastructure agreements. The \$24.6 bn figure is the company's most favorable disclosure; the legally enforceable floor is lower. **Second**, the bulk of recognition (43%) is concentrated in FY28–FY29, which means most of the value sits three to four years out, with substantial discount factor compression.

RPO sanity check

Even taking the disclosed schedule at face value, FY26 implied revenue of ~\$1.85 bn/year would put forward P/S at approximately 30x — high but not unreasonable for a hyper-growth AI infrastructure name. By FY28, however, P/S would compress to ~11x at the current market cap, which is closer to historical Nvidia at its 2023 inflection. The question is whether the multiple holds while waiting three years for the backlog to convert.

6. Valuation Framework — Three Lenses

We triangulate fair value across three distinct approaches: (i) trailing and forward multiples, (ii) RPO-anchored DCF with normalized margins, and (iii) Street consensus. Each approach has limitations; the convergence (or non-convergence) of the three is more informative than any single output.

Lens 1 — Trailing & Forward Multiples

Metric	FY25A	FY26E (~\$1.2 bn rev)	FY27E (~\$2.3 bn rev)
Market cap (current)	\$53.1 bn	\$53.1 bn	\$53.1 bn
EV (net of cash)	\$52.4 bn	\$52.4 bn	\$52.4 bn
P/S (trailing/forward)	104x	44x	23x
EV/Revenue	103x	44x	23x
P/E (GAAP)	219x	n/a (loss)	n/a (loss)

For reference, Nvidia at its FY24 inflection traded at approximately 25–30x forward sales; Palantir today trades at 200x earnings on hyper-growth narrative; Cerebras' FY27E forward P/S of 23x is therefore not historically unprecedented, but it requires the forward revenue assumptions to be substantially met.

Lens 2 — RPO-Anchored DCF (Base Case)

FY	Revenue (\$bn)	Op margin	EBIT	Tax	NOPAT	Net income
FY26E	1.20	10%	0.12	10%	0.11	0.30 (incl. tax shield)
FY27E	2.30	18%	0.41	12%	0.37	0.58
FY28E	4.80	25%	1.20	15%	1.02	1.20
FY29E	5.50	27%	1.49	16%	1.25	1.38

Applying a terminal multiple of 25–30x to FY29E net income of \$1.38 bn yields a FY29 equity value of \$34.5–41.4 bn. Discounting back four years at a 12% cost of equity produces a present fair value of \$21.9–26.3 bn, or approximately **\$100–120 per share** on basic shares outstanding. Sensitivity: every 100 bp of normalized operating margin moves fair value by approximately \$8–10/share; every 5x of terminal multiple moves it by \$20–25/share.

Lens 3 — Street Consensus & Sell-Side Anchor

Street initiation coverage is still building post-IPO. The most rigorous published price target visible at the time of this note is Seeking Alpha's \$275 (analyst Daniel Jones), with a recommendation to wait for either post-lockup supply or two clean quarters before adding. Morningstar's DCF-based fair value is \$513.57 (Very High uncertainty rating, 1-star price \$484.43, 5-star price \$22.83) — a notably wide uncertainty band that essentially reflects the team's view that the outcome distribution is bimodal rather than normal. This is unusual and worth weighting accordingly: the model is telling us that point estimates are misleading for this name.

7. Fair-Value Range & Scenario Analysis

Scenario	Key Assumptions	Implied Value/Share
Bear	OpenAI renege on G42; CFIUS re-engagement on G42; Nvidia compresses inference moat with Blackwell software; backdoor converts at 60% of scheduled	\$60-80
Base	50% steady-state net margin; 25–30x terminal P/E; CUDA moat narrows but doesn't collapse; one additional 10%+ customer signed by Q2 FY26	\$100-150
Bull	Second hyperscaler MRA signed within 18 months; net margin holds at 30%; 30–35x terminal multiple; sovereign AI tailwind compounds; CFIUS clears G42 definitively	\$200-275

The current price of \$241.71 sits at the **top of the bull range**. To justify ownership at this level, the investor must underwrite essentially all four of: (i) OpenAI executing the \$20 bn contract on schedule, (ii) a second non-UAE hyperscaler materializing within 24 months, (iii) CFIUS finding no further cause for concern on G42 deepening involvement, and (iv) the inference moat holding against CUDA-on-Blackwell + AMD MI400 + hyperscaler custom silicon. All four conditions need to clear. The expected-value math is unfavourable here even with bullish individual probabilities, because conditional probabilities compound multiplicatively.

Sensitivity — Where Does the Stock Re-Rate?

Trigger	Direction	Magnitude
Q2 FY26 print misses revenue / margin guide	Down	-15 to -25%
Q2 FY26 discloses 3rd 10%+ customer (non-UAE, non-OpenAI)	Up	+10 to +20%
CFIUS announces G42 review	Down	-20 to -35%
OpenAI takes inference in-house (new ASIC announcement)	Down	-30 to -50%
Lock-up expiry (10-Nov-2026) — supply unwind	Down	-10 to -20% (technical)
Second hyperscaler MRA signed (any of Google/Meta/Oracle)	Up	+25 to +40%
AMD MI400 captures significant inference share	Down	-15 to -25%

8. Key Risks

Customer Concentration (G42 / UAE)

86% of FY25 revenue concentrated in a single sovereign-linked customer is the single largest quantitative risk. This is also the source of CFIUS scrutiny: G42's relationships with Chinese technology partners have been subject to US government review on multiple occasions. A material restriction on G42's ability to procure US-origin AI silicon would compress Cerebras' near-term revenue more than any other single event.

Inference Moat Durability

Nvidia's CUDA software ecosystem remains the dominant moat in AI compute. Twenty years of developer tooling, optimised libraries, and trained engineers does not erode quickly. Cerebras' SDK, while improving, is not yet a credible substitute for CUDA at the average enterprise customer. As Nvidia continues to optimise its software stack specifically for inference (TensorRT-LLM, NIM microservices), the architectural advantage of WSE-3 may narrow in measured workloads even if it remains real on the largest models.

Hyperscaler Custom Silicon

Google (TPU v5/v6), Amazon (Trainium, Inferentia), Meta (MTIA), and Microsoft (Maia) are all building custom inference silicon optimised for their own workloads. If any one of these hyperscalers, or OpenAI itself, decides in 2027–28 that their model architecture has stabilised enough to commit to ASIC manufacture, the Cerebras MRA becomes a transition vehicle rather than a permanent franchise. This is the structural risk that matters most for the long-term thesis.

Wafer-Scale Manufacturing Economics

Cerebras' yield engineering compensates for defects through redundancy, but the unit economics of wafer-scale chips are structurally less favourable than die-cut GPUs at advanced nodes. As TSMC moves to 3nm and below for Cerebras' next-generation parts, the yield math becomes more stressed, and the cost-per-equivalent-compute advantage may erode.

Lock-Up & Technical Supply

The 180-day lock-up expires on approximately 10 November 2026. Pre-IPO holders — Benchmark, Foundation Capital, Eclipse, Fidelity, and G42 — collectively hold multi-billion-dollar positions. Comparable high-multiple IPOs (Snowflake, Coinbase, Rivian) saw 15–25% drawdowns in the 60 days surrounding their lock-up unwinds. This is a known technical overhang and is partially priced in, but supply at \$241 is not the same as supply at \$150–180; the level of the unwind matters.

Capital Intensity & FCF Timing

The OpenAI MRA requires Cerebras to procure substantial wafer and packaging capacity from TSMC and partners, build out dedicated data center capacity, and finance working capital across multi-year delivery cycles. Adjusted free cash flow positive on a sustained basis is unlikely before FY27. Investors holding the stock for cash returns are mispricing the horizon.

9. Catalysts & Watch Items

Catalyst	Expected Timing	Materiality
Q2 FY26 earnings — first datapoint on OpenAI delivery run-rate	Mid-August 2026	Very High
10-Q disclosure of 3rd 10%+ customer (or absence thereof)	Mid-August 2026	Very High
180-day lock-up expiry — supply unwind	~10 November 2026	High (technical)
CFIUS / G42 export-control review outcome	Ongoing / unscheduled	Very High (binary)
Q3 FY26 earnings — full-quarter visibility on OpenAI ramp	Early November 2026	High
FY26 full-year results	February 2027	High
AMD MI400 launch (competitive read)	Late 2026 / early 2027	Medium
Nvidia Rubin Ultra disclosure (competitive read)	GTC 2027	Medium
Second hyperscaler MRA signing (any)	Watch through 2026–27	Very High (bull case)

The single highest-value watch item is the Q2 FY26 print in mid-August. This will be the first quarter that includes any meaningful OpenAI delivery, and the revenue line will either confirm or contradict the implied \$1.85 bn/year run-rate from the disclosed RPO conversion schedule. A 10%+ beat on revenue with a disclosed third major customer would materially de-risk the bull thesis. A miss combined with margin pressure would likely compress the multiple by 20%+ in a single session.

10. Conclusion & Investment View

Cerebras Systems is one of the most architecturally interesting public companies of this AI cycle. The Wafer-Scale Engine is a genuine technical achievement, the OpenAI MRA is a hard validation of commercial viability, and the inference-first positioning is well-timed against the broader shift in AI workload composition. None of these observations are controversial.

The investment debate is therefore not about *whether* Cerebras has a viable business — it does — but about *what that business is worth*. On the math, at \$241.71 the stock is priced for essentially perfect execution of an extraordinarily back-loaded contractual schedule, with no margin for any of the four primary risks (customer concentration, CUDA moat, hyperscaler ASICs, manufacturing economics) materialising adversely.

Our framework produces a base-case intrinsic value of **\$100–150 per share**, with a bull case at **\$200–275** requiring conditions that are individually plausible but jointly unlikely to all clear. The asymmetric setup favours patience over urgency. The cleanest re-entry windows are:

- (i) Post-lockup supply unwind in November 2026, where forced selling from pre-IPO holders is likely to take the stock into the \$150–180 range — a level at which the base case becomes investable on conventional risk-adjusted return assumptions; or
- (ii) Two consecutive quarters with on-schedule OpenAI revenue contribution and a disclosed third 10%+ customer, which would shift the customer-concentration probability set materially and validate the bull case.

For institutional positions sized against existing AI infrastructure exposure, the cleaner expression of the Cerebras thesis at this point in the cycle is a paired trade — long CBRS / short an over-stretched AI infrastructure peer with similar duration but worse customer concentration — rather than an absolute long position at the current entry. The technology story remains compelling; the price discipline matters.

Bottom line

Wait for the lockup. November 2026 is the cleanest re-entry. Holding cash and watching the Q2 FY26 print + customer disclosure is a higher-quality use of capital than chasing the post-IPO tape at 104x trailing sales. The thesis is real; the price isn't.

Important Disclosures

Nature of this document. This research note has been prepared for institutional review as an analytical framework summarising publicly available information on Cerebras Systems Inc. (NASDAQ: CBRS). It is not investment advice, an offer to sell, or a solicitation of an offer to buy any security. The analysis reflects the author's interpretation of publicly disclosed financial data, regulatory filings, broker research summaries, and management commentary as of 28 May 2026.

Price & data sources. Live price references reflect 27 May 2026 close prices as reported by Investing.com, The Motley Fool, Morningstar, and Robinhood. The user is reminded that securities prices change continuously; all upside, downside, and multiple computations in this document are point-in-time and should be re-verified against live market data before any investment decision is made.

Forward-looking statements. Any forward-looking statement (revenue projections, RPO conversion timing, margin trajectory, terminal multiples, fair value estimates) is an estimate based on the author's interpretation of disclosed company guidance, broker research, and sector comparables. Such estimates are subject to material change as new information becomes available. Cerebras has been a public company for fewer than 30 days as of this note's date; sell-side coverage is incomplete, and consensus estimates remain provisional.

Conflicts of interest. The author holds no position in CBRS or any peer security mentioned in this note and has received no compensation from any party referenced herein. This document has not been reviewed or endorsed by Cerebras Systems Inc., its underwriters, or any third party.

Use restriction. This document is prepared for internal institutional review. It is not to be distributed externally without the author's consent and is not intended for retail investor consumption.

Limitation of liability. The author and any associated party accept no liability for any loss arising from action taken on the basis of information contained in this document. Recipients should conduct independent due diligence and consult licensed financial, legal, and tax advisors before making investment decisions.

Risk warning. CBRS is a recently listed, high-volatility security (beta 18.82 per Motley Fool, reflecting <30-day trading history). The stock is subject to extreme short-term price movements and is unsuitable for investors without appropriate risk capacity. Past performance is not indicative of future returns.

Cerebras Systems Inc. (NASDAQ: CBRS) — Equity Research Note
Prepared 28 May 2026 • Last price reference: \$241.71 (27-May-2026 close)
Internal institutional review document.