
Cerebras (CBRS) IPO Deep Dive — The NVIDIA Challenger Arrives

Analysis Memo · warband-2026-05

warband-2026-05

Ragnar, the HODLBarbarian

■■ Internal Analysis Memo ■■

Executive Summary

Cerebras Systems (CBRS) completed the largest IPO of 2026 on May 14, raising \$5.55B at \$185/share. The stock opened at \$350, hit \$386 intraday (circuit breaker triggered), and closed its first session at \$311 (+68%). As of Friday May 16, it trades at ~\$280 — down 28% from the intraday high but still +51% from IPO price. Market cap: ~\$60B.

The thesis in one sentence: Cerebras built a dinner-plate-sized chip (WSE-3) that runs AI inference 21x faster than NVIDIA's Blackwell GPU, and OpenAI bet \$20B on it. ARK Invest dumped \$51.5M of TSMC/AMD to load \$50M+ of CBRS across two days.

Bottom line: This is a WATCH, not a BUY. The technology is real, the demand is real, but the valuation is euphoric (130-190x revenue), 86% of revenue comes from two related UAE entities, and NVIDIA acquired Groq (\$20B) specifically to compete in this exact market. The stock will tell you when to buy — wait for the first earnings print and a pullback to \$200-220 range.

I. The IPO: What Happened

IPO price	\$185
First trade	\$350
Day 1 high	\$386.34 (circuit breaker)
Day 1 close	\$311.07 (+68%)
Friday close (May 16)	~\$280
Market cap	~\$60B
Shares outstanding	215.23M
Capital raised	\$5.55B (30M shares)
Largest IPO of 2026	Yes

Arm/SoftBank attempted a last-minute acquisition before listing. Cerebras rejected it. The company chose public markets — and the market responded with the most violent first-day pop since ARM's

own IPO.

II. The Technology: Why It Matters

WSE-3: The Dinner Plate Chip

Cerebras builds the Wafer-Scale Engine 3 — a processor that occupies an **entire silicon wafer** rather than the thumbnail-sized die used for traditional chips.

Area	46,225 mm ²	~814 mm ²
Transistors	4 trillion	208 billion
Compute cores	900,000	~18,000 CUDA
On-chip SRAM	44 GB	None (HBM)
Process node	5nm (TSMC)	4nm (TSMC)
Inference speed	2,100 tok/s	~100 tok/s

The key insight: Because all 44GB of model weights live on-chip (SRAM), inference never has to traverse the memory bus between chips. A single Cerebras system eliminates the inter-chip latency that bottlenecks GPU racks. This is why OpenAI's engineers found that no matter how they tuned NVIDIA parameters, Codex response speed was limited by GPU architecture — the bottleneck is structural, not optimization.

The Inference Moat

- Training = parallel compute (NVIDIA's GPU strength)
 - Inference = latency + memory bandwidth (Cerebras' wafer-scale strength)
 - The AI inference market is expected to grow **faster** than training over the long term
 - Agentic AI (1,000x more compute than chatbot inference) is the new demand driver
-

III. The OpenAI Deal: \$20B With Strings Attached

Structure

- **\$20B+ Master Relationship Agreement** — deploy 750MW of Cerebras inference capacity through 2028
- **\$1B working capital loan** from OpenAI to Cerebras at 6% interest
- **Warrants** for OpenAI to purchase 33.4M shares at near-zero exercise price
- At \$311 IPO close, those warrants = ~\$11.7B in value

The Clawback Risk

If the Master Relationship Agreement is terminated **for any reason other than OpenAI's own uncured breach**, OpenAI can:

- Seize control of the loan funds
- Demand immediate repayment
- Keep the warrants

Translation: OpenAI effectively has a kill switch on Cerebras. If OpenAI walks, Cerebras loses its largest customer, its \$1B loan, AND the warrants — simultaneously.

The Gross Profit Reality

Per Financial Times calculations: OpenAI's warrants are worth roughly **half the gross profit** Cerebras stands to make on the entire \$20B deal. OpenAI isn't just a customer — it's an equity partner extracting half the economics.

IV. The UAE Problem: 86% Revenue Concentration

The History

- First IPO attempt (late 2024) **collapsed** under CFIUS security review
- G42, Abu Dhabi AI conglomerate, was under investigation for Huawei ties
- G42 provided \$335M investment and 85% of revenue

The "Diversification" That Wasn't

The May 2026 S-1 shows:

- G42 revenue share fell to 24%
- **BUT** MBZUAI (Mohamed bin Zayed University of AI) — a related party of G42 — accounts for **62% of 2025 revenue**

- Combined: **86% of revenue from two related UAE entities**
- MBZUAI alone = 77.9% of accounts receivable

This is not diversification. It's reallocation between connected entities. The prospectus acknowledges G42 and MBZUAI "are considered related parties with respect to each other."

Geopolitical Risk

- Iran war / Middle East tensions directly impact UAE economic stability
 - U.S. export controls could tighten on advanced AI hardware to UAE
 - Any disruption to UAE payments = existential risk to Cerebras
-

V. The Profitability Illusion

Reported Numbers

- 2025 revenue: \$510M (+76% YoY)
- 2025 GAAP net income: \$237.8M
- Headline: "Cerebras is profitable!"

The Reality

- The \$237.8M profit was driven by a **one-time, non-cash gain of \$363.3M** from extinguishing a forward-contract liability tied to G42's original investment
- Strip that out + adjust for SBC: **non-GAAP net loss of \$75.7M** (247% worse YoY)
- GAAP operating loss: \$145.9M
- R&D spend: 48% of revenue

Valuation Context

At \$311 (Day 1 close): **130-190x trailing revenue**

- NVIDIA: ~26x revenue
 - AMD: ~21x revenue
 - Motley Fool: "borderline euphoric" — declined to recommend
-

VI. The Competition: NVIDIA's \$20B Counter-Move

NVIDIA Acquired Groq (December 2025)

- \$20B cash deal — NVIDIA's largest acquisition ever
- Groq's Language Processing Units use SRAM-based architecture **similar to Cerebras**
- NVIDIA announced Groq-based products at GTC March 2026
- Morningstar: Groq is "likely Cerebras' fiercest rival, since both are attacking AI inference using smaller but faster SRAM memory"

Other Competitors

- Google: custom TPU inference chips
- Microsoft: Maia AI accelerator
- Amazon: Trainium/Inferentia
- AMD: MI300X inference mode
- **Every hyperscaler is building custom inference silicon**

The 5nm Problem

WSE-3 is manufactured on TSMC's 5nm node. The most cutting-edge AI chips use 2nm. This is a process disadvantage competitors will exploit in next-gen products.

VII. ARK's Bet: \$50M+ Across Two Days***The Trades***

May 15 (IPO day)	Bought CBRS	\$4.85M (105,616 shares)
May 16	Bought CBRS	\$46.4M (149,176 shares)
May 15-16	Sold TSMC	\$40.6M
May 15-16	Sold AMD	\$10.9M

ARK's thesis: Cerebras is the next wave in AI compute. Selling "old guard" semis (TSM, AMD) to fund "new wave" inference chip.

The ARK Risk

- ARKK trails S&P 500 by 11% YTD in 2026

- ARK has a history of buying IPOs at peak euphoria and holding through drawdowns
 - CBRS at 130-190x revenue is exactly the kind of valuation ARK loves — and often gets burned on
-

VIII. The Backlog: \$24.6B But When?

- Remaining performance obligations: \$24.6B
- Management estimates: only ~15% recognized in 2026-2027 **combined**
- Bulk converts to revenue in **2028 and beyond**
- By 2028, NVIDIA (with Groq), AMD, Google, Amazon, Microsoft will all have competing inference products

The timing problem: The backlog is real but back-loaded. The stock is priced for 2028+ success today. Any delay, contract renegotiation, or competitive loss hits a stock priced for perfection.

IX. Signal Integration: Where CBRS Fits in the Thesis

Visser Framework

Cerebras sits at **Layer 2 (Chips & Memory)** and **Layer 4 (AI Infrastructure)** of Visser's five-layer cake. The IPO validates his thesis that AI inference demand will outstrip training demand. But Visser has been **rotating OUT of semis** (sold MU) into Layer 1 (power, commodities, chemicals). CBRS is a Layer 2/4 name in a market that's rotating toward Layer 1.

Weekly Catalyst Context

NVIDIA reports Q1 FY2027 earnings **tomorrow (May 20)**. Key implications:

- If NVIDIA beats on inference revenue → validates CBRS TAM → CBRS holds
- If NVIDIA guides down on inference → CBRS sympathy selloff
- If NVIDIA's Groq-based products show strong traction → CBRS competitive threat increases

Iran Peace Link

If Iran peace deal materializes → oil drops → rates drop → risk-on → growth/innovation stocks (including CBRS) benefit. But CBRS is too new and volatile to trade on macro — it trades on its own narrative.

X. Decision Framework

Conditions That Would Argue For Buying

1. **Pullback to \$200-220 range** (still 40-50% above IPO, but 100x revenue — more reasonable for 76% growth)
2. **First earnings print** showing revenue acceleration and UAE concentration declining
3. **NVIDIA earnings** confirming inference TAM expansion (not zero-sum)
4. **Iran peace deal** reducing geopolitical risk to UAE customer base

Conditions That Would Argue Against

1. **UAE concentration stays above 70%** through Q2 2026
2. **OpenAI clawback risk** materializes (any contract dispute)
3. **NVIDIA Groq products** show clear inference superiority at scale
4. **Stock stays above \$300** — valuation is simply too stretched

Exit Triggers (For Existing Holders)

1. UAE payment default or CFIUS re-investigation
 2. OpenAI terminates MRA
 3. Revenue growth decelerates below 50% YoY
-

XI. Key Quotes

"For our fast inference product, there's so much demand that our biggest challenge is actually trying to supply it. We are adding as much manufacturing and data center capacity as we possibly can, and we're still sold out into 2027." — Bob Komin, Cerebras CFO, May 14, 2026

"The OpenAI deal came as part of a \$20-billion arrangement that took the form of a commitment to buy billions of dollars of Cerebras' chips in exchange for a stake in the company and a \$1 billion loan." — Brian Colello, Morningstar Senior Equity Analyst

"Really, if you buy this stock up here, you're betting on the idea that Cerebras will have much better growth for many years in the future." — Jim Cramer, CNBC, May 14, 2026*

"There's some whales out there, there's some really big customers. That is one of the characteristics of this market." — Andrew Feldman, Cerebras CEO, on UAE concentration*

XII. Sources

- CNBC: "What you need to know about Nvidia competitor Cerebras after wild IPO" (May 15)
- TechTimes: "Cerebras Raises \$5.55 Billion in AI Chip IPO, But 86% Revenue Dependence on UAE Entities Unresolved" (May 15)
- Motley Fool: "Cerebras Systems Stock Soars 68% in Blockbuster IPO" (May 14)
- Rolling Out: "Cathie Wood's ARK increases Cerebras, cuts AMD stock" (May 17)
- Blockonomi: "ARK Invest Dumps \$40M in TSMC Stock, Pivots to Cerebras" (May 17)
- TipRanks: "Cathie Wood's Semiconductor Trades" (May 15)
- CNBC: "Nvidia buying AI chip startup Groq's assets for about \$20 billion" (Dec 24, 2025)
- Reuters: "Nvidia, joining Big Tech deal spree, to license Groq technology" (Dec 24, 2025)
- The Register: "Cerebras' wafer-scale AI bet delivers blockbuster IPO" (May 15)
- PANews: "Two \$20 billion deals: OpenAI and Nvidia are waging a war of inference" (April 2026)
- Motley Fool: "Nvidia's Aquil-Hire of Groq Eliminates a Potential Competitor" (Dec 28, 2025)

WAAS — *Viking intelligence delivered in real time. Not investment advice.*