Gavin Baker — GPUs, TPUs & The Economics of AI

Strategic Implications for the Tiger Portfolio Watchlist

Source: *Invest Like the Best*, EP.451 — "Nvidia v. Google, Scaling Laws, and the Economics of AI" **Published:** December 9, 2025 | **Host:** Patrick O'Shaughnessy **Research Date:** February 19, 2026 | **Portfolio:** Tiger ($100K AI Infrastructure)

Part I: The Interview — Key Thesis Points

Context & Backdrop

Gavin Baker is CIO of Atreides Management, a long/short tech fund regarded as one of the sharpest institutional voices on AI infrastructure. This episode was recorded in late November 2025 against the following backdrop:

- Gemini 3 just released by Google
- Blackwell (GB200/GB300) in early-scale deployment — only 3-4 months in at time of recording
- xAI's Colossus cluster operational
- DeepSeek open-source moment had occurred earlier in 2025
- OpenAI's Stargate announcement recent; Anthropic signed a $5B Nvidia deal

---

1. Scaling Laws: Intact and Multiplicative

> *"Gemini 3 was very important because it showed us that scaling laws for pre-training are intact. They stated that unequivocally."*

Baker frames our understanding of *why* scaling laws work like the ancient Egyptians measuring equinoxes: perfect empirical observation, zero theoretical grounding. Every confirmation matters enormously.

Three multiplicative scaling laws now exist — a step-function moment:

1. **Pre-training scaling** (compute × data × parameters)
2. **RLVR** — Reinforcement Learning with Verified Rewards
3. **Test-Time Compute** — reasoning at inference

ARC-AGI scores went from 8% to 95% in three months when the first reasoning model launched. *"Reasoning kind of saved AI"* — it bridged the 18-month Blackwell gap during which pre-training alone would have stalled progress.

---

2. Blackwell: The Most Complex Product Transition in Tech History

> *"The most complex product transition in the history of tech."*

Metric	Hopper (H100)	Blackwell (GB200/GB300)
Cooling	Air	Liquid (mandatory)
Rack weight	~1,000 lbs	~3,000 lbs
Power per rack	30 kW	130 kW
Max coherent cluster	~200,000 GPUs	1M+ GPUs

Baker's analogy: deploying Blackwell is like *"changing all outlets to 220V, putting in a Tesla Powerwall, generator, solar panels, and reinforcing the floor."* GB300 is drop-in compatible with GB200 liquid-cooled racks — meaning the liquid cooling infrastructure is **locked in permanently**.

Companies that master liquid-cooled Blackwell racks become the **low-cost token producers**.

---

3. NVIDIA Wins the GPU vs. ASIC War — With Two Exceptions

> *"I will be surprised if there are a lot of ASICs other than Trainium [Amazon] and TPU [Google] long-term. The economics make it absolutely inevitable."*

Baker's full-stack complexity argument — why ASICs fail:

> *"What's the NIC going to be? What's the CPU? What's the scale-up switch? Scale-up protocol? Scale-out switch? What kind of optics? What's the software? And then it's like, oh shit, I made this tiny little chip."*

Key data points on ASIC viability:

- TPU v1 was "an achievement just to exist" — not competitive until v3/v4 (~3 generations)
- Amazon Trainium same pattern
- Microsoft Maia, Meta MTIA — neither is commercially meaningful at GPU-replacement scale
- Baker: Nvidia's annual cadence (Blackwell → Rubin → next-gen) is itself a competitive weapon

The Broadcom tension around Google's TPU:

- Google pays Broadcom an estimated **$15B/year** (50-55% gross margin) for TPU back-end
- Broadcom's total semi opex is ~$5B — structural tension is unsustainable
- Google's MediaTek partnership = first warning shot; Apple-model full vertical integration is likely inevitable
- Baker's conclusion: Google's TPU cost advantage is **temporary** — once Blackwell/Rubin deploys at scale, the economics flip decisively

---

4. Token Economics: The New Competitive Battlefield

> *"AI is the first time in my career as a tech investor that being the low-cost producer has ever mattered."*

This is Baker's most important reframe. The entire infrastructure investment thesis flows from it:

- Whoever produces tokens cheapest wins market share
- That prize goes to: best chips + lowest power costs + best cooling
- Google was the low-cost producer briefly (TPU + free power from owned infra)
- Blackwell + vertical integration transfers that advantage back to NVIDIA ecosystem operators

Implication for hyperscalers: If Google loses cost leadership, the rational strategy of running AI at negative 30% margins becomes untenable. Baker: *"It might start to impact their stock."*

---

5. ROI: Now Documented in Audited Fortune 500 Financials

> *"The ROI on AI has empirically, factually, unambiguously been positive."*

- ROIC of the biggest GPU buyers is **higher now** than before the spending ramped
- C.H. Robinson: went from quoting 60% of requests in 15-45 minutes → 100% in seconds, stock +20%
- AI is now showing up in audited financials — not just anecdotes
- The "prisoners dilemma" dynamic keeps spending elevated: stopping = permanent competitive disadvantage

The Blackwell ROI Air Gap (key near-term risk):

For ~3 quarters during Blackwell ramp, capex is massive but chips are mostly used for training, not inference. ROIC may dip. Baker: *"Meta printed a quarter where ROIC declined. And that was not good for the stock."* This is a **timing risk**, not a structural bear case.

---

6. Reasoning Changed the Industry Structure — Flywheels Are Spinning

Pre-reasoning AI had no data flywheel. Reasoning enables verifiable reward signals (users consistently like/dislike similar responses → feeds back into the model). Baker: *"That flywheel has started to spin."* This creates compounding moats — leaders widen the gap continuously.

Frontier model building is far harder than assumed:

- Meta: Spectacular failure despite massive spending. Zuckerberg's "best AI by 2025" prediction was *"as wrong as it was possible to be."*
- Microsoft: Failed post-Inflection AI acquisition; internal models never caught up
- Amazon: Nova models — *"I don't think they're in the top 20."*
- OpenAI: In "code red" on cost structure. Stargate = signal they know

they need to vertically integrate

Winners by Baker's framework:

- **Anthropic:** *"A good company"* — burning less cash than OpenAI, growing faster; pivoting to Blackwell via $5B Nvidia deal (Dario understands Rubin dynamics)
- **xAI:** Dominant share on OpenRouter; first mover on Blackwell inference at scale

---

7. Space Data Centers: The 10-20 Year Paradigm Shift

> *"The most important thing that's going to happen in the world in the next 3 to 4 years."*

First-principles case: 6× solar irradiance in LEO (no atmosphere), constant sun (no batteries), passive cooling toward absolute zero on dark side, laser transfer through vacuum faster than fiber. *"In every way, data centers in space from a first principles perspective are superior to data centers on earth."* Key implication for terrestrial plays: **power and cooling are the binding limits on AI scale** — reinforcing terrestrial thermal management as a persistent high-value market.

---

8. Semiconductor Supply Chain: Key Views

- **TSMC:** Making a mistake by not expanding capacity fast enough; met with Sam Altman and *"laughed and said he's a podcast bro"* — Baker sees this conservatism as a natural governor preventing overbuild
- **Intel:** Empty fabs; Lip-Bu Tan is a *"really good executive"*; Gelsinger's firing was *"shameful"* — Intel was on the only viable strategy. Eventually fabs fill due to TSMC supply constraints
- **DRAM:** First true DRAM cycle since the late '90s potentially brewing — historical cycles were 10× price increases; recent ones only 30-50%; *"if it starts to go up by X's instead of percentages, that's a whole different game"*
- **Semiconductor venture:** Baker's firm may have done more semi VC deals in 7 years than top 10 VCs combined; Nvidia's success *"singlehandedly ignited semiconductor venture"*; average founder is ~50 years old; Blackwell rack has thousands of parts, Nvidia makes only 200-300

Notable Quotes

Quote	Context
"Reasoning kind of saved AI."	RLVR + test-time compute bridged the Blackwell gap
"AI is the first time… being the low-cost producer has ever mattered."	Token economics thesis
"Oh shit, I made this tiny little chip."	ASIC builder's realization of full-stack complexity

“Meta was as wrong as it was possible to be.”	Frontier model failure analysis
“The flywheel has started to spin.”	Reasoning model compounding moat
“The most complex product transition in the history of tech.”	Blackwell infrastructure shift

---

Part II: GPU vs. ASIC Market Structure

Current Hardware Hierarchy

Generation	Baker Analogy	Coherent Cluster Scale	Status
Hopper (H100)	WWII P-51 Mustang	~200,000 GPUs	Deployed; being phased out
TPU Ironwood (v7)	F-4 Phantom	9,216-chip clusters	Deployed Nov 2025
Blackwell (GB200/300)	F-35	1M+ GPU potential	Deploying now
Rubin (coming)	Next-gen fighter	TBD	Baker: “gap will expand significantly”

Hyperscaler Custom Silicon Reality Check

Hyperscaler	Custom Chip	Baker’s Assessment	% of Capex
Google	TPU Ironwood (v7)	Temporary cost advantage; eroding	~20-25% of AI compute
Amazon	Trainium 2/3	“Best ASIC team outside Google”; viable long-term	~10-15% of AI compute
Microsoft	Maia/Athena	Supplementary; not GPU-replacement scale	<5% of AI compute
Meta	MTIA v3	Recommendation systems only; Llama runs on GPUs	<5% of AI compute
OpenAI	Broadcom ASIC (long-term)	Stargate = NVIDIA near-term; ASIC is a 5-year project	0% today

Key structural insight: NVIDIA commands ~80% of AI training chip market. Custom ASICs are a growing minority. Baker sees only TPU and Trainium as viable long-term GPU alternatives. Microsoft Maia, Meta MTIA = efficiency plays, not replacements.

AI Capex Cycle: Scale and Duration

Company	2025 Capex (est.)	2026 Guidance	YoY Growth
Amazon	~$130B	**$200B**	+54%
Alphabet	~$85B	**$175-185B**	+110%
Microsoft	~$80B	**$120B+**	+50%
Meta	~$70B	**$115-135B**	+65-85%
Oracle	~$20B	**$50B**	+150%
Big 5 Total	**~$385B**	**$660-690B**	**+75%**

Additional: Stargate Project (OpenAI/SoftBank/Oracle) = $500B 4-year commitment. All hyperscalers report **supply-constrained, not demand-constrained** markets.

Part III: Sector Implications

Interconnect Architecture: The Defining Tech Battle

Scale-up (within cluster): NVLink/NVSwitch — NVIDIA's proprietary domain. Custom ASIC clusters (TPU, Trainium) use proprietary equivalents. Third-party vendors serve scale-up mostly indirectly.

Scale-out (cluster to cluster): Ethernet is decisively winning vs. InfiniBand:

- Ultra Ethernet Consortium (UEC) 1.0 spec finalized June 2025; hyperscalers validating RoCE at scale
- Dell'Oro: 2025 was the turning point — Ethernet overtook InfiniBand in AI backend networks
- **This is a structural tailwind for ANET** — Arista is the dominant Ethernet switching vendor for AI clusters

Optics: Baker explicitly flags optics as the critical unsolved challenge:

> *"What kind of optics are you going to use?"*

- Pluggable transceivers: 400G → 800G → 1.6T progression
- Co-packaged optics (CPO): commercial deployment beginning 2025-2026
- NVIDIA announced silicon photonics switches at GTC 2025 (Spectrum-X, Quantum-X) — watershed moment
- Data center interconnect market: $15.99B (2024) → $32.73B by 2030 (CAGR: 12.7%)

Electrical interconnect:

- PCIe 6.0 adoption ramping with Blackwell
- CXL gaining traction for memory pooling across accelerators
- SerDes (224G PAM4) critical for chip-to-chip and chip-to-switch connections
- UALink 1.0 released; product samples 2026, revenue 2027

Liquid Cooling: Mandatory, Not Optional

Baker's explicit framing: liquid cooling is the defining infrastructure shift of Blackwell. 30kW/rack → 130kW/rack means air cooling is physically impossible at the required density.

- Liquid cooling market: $5-7B (2025) → $15-26B by 2030 (CAGR: ~25-30%)
- GB300 drop-in compatible with GB200 liquid-cooled racks → infrastructure locked in permanently
- Baker's space data center thesis reinforces: cooling is the **binding terrestrial constraint** on AI scale

---

Part IV: Ticker-by-Ticker Impact Analysis

Methodology

For each ticker: assess alignment with Baker's thesis across four dimensions: (1) Direct quote or explicit mention, (2) Sector alignment, (3) GPU-ASIC agnosticism, (4) Baker's actual portfolio action (Q4 2025 13F).

ALAB (Astera Labs) — PCIe/CXL Connectivity

Baker's Action: Q4 2025 13F shows increase from 62,050 shares → **1.6 million shares (+2,477%)**. This is his single most emphatic AI infrastructure bet outside NVDA calls.

Why It Aligns:

- ALAB makes PCIe Gen 5/6 retimers, CXL memory connectivity, smart cable modules — the "critical plumbing connecting GPUs in massive data center clusters"
- Baker's full-stack critique of ASIC builders highlights the NIC/connectivity layer as a key bottleneck — Astera solves this for GPU and ASIC clusters alike
- NVLink Fusion support: serves NVIDIA GPU clusters
- CXL/PCIe: serves ASIC clusters needing memory pooling and bandwidth expansion
- **Platform-agnostic by design** — open standards across PCIe, CXL, Ethernet, UALink
- Scorpio fabric switch and Neptun Smart Cable Modules ramping strongly
- Projected revenue: >$1B by 2026

GPU-ASIC Agnosticism Score: Very High — ALAB wins more if ASICs proliferate because ASICs need more third-party connectivity solutions than NVIDIA's vertically integrated GPU stacks

Risk: Post-Q4 2025 earnings margin concerns pushed stock down ~32%. Uncertain if Baker maintained position after earnings. Customer concentration.

Verdict: 🟢 STRONGLY SUPPORTED — Baker voted with his portfolio. Highest-conviction alignment.

CRDO (Credo Technology) — High-Speed SerDes

Baker's explicit mention: Flags optics and high-speed I/O as critical unresolved challenges for ASIC builders. Credo's SerDes and Active Electrical Cables (AECs) are direct answers.

Why It Aligns:

- SerDes is the universal physical layer interface — needed in every architecture
- Custom ASIC clusters actually **increase demand** for third-party SerDes: hyperscalers building their own accelerators source high-speed I/O externally
- Amazon (major CRDO customer) is Baker's #2 viable ASIC program — long-term Trainium → Credo SerDes
- CRDO stock: +2,050% from IPO; 63.8-65.8% gross margins; Q2 FY2026 earnings blowout
- 800G → 1.6T SerDes roadmap directly aligned with cluster bandwidth scaling
- Ethernet winning scale-out networking = structural tailwind for SerDes demand

GPU-ASIC Agnosticism Score: Very High — ASIC adoption accelerates CRDO's opportunity with hyperscaler custom silicon builders

Risk: Customer concentration (Amazon reportedly large share). If Amazon fully vertically integrates SerDes into Trainium, risk of customer loss (unlikely near-term given complexity).

Verdict: 🟢 STRONGLY SUPPORTED — Custom silicon is NET POSITIVE for CRDO.

FN (Fabrinet) — Optical/Electronic Contract Manufacturing

Baker's relevant thesis: Optical content per rack is expanding with Blackwell, and scale-out Ethernet networking requires 800G/1.6T transceivers at scale. Fabrinet manufactures the assemblies that every optical transceiver vendor needs.

Why It Aligns:

- 76% of FN revenue tied to data center networking
- As optical bandwidth demands increase (400G → 800G → 1.6T), precision manufacturing requirements increase → FN's competitive moat widens
- Data center interconnect spending: $15B (2025) → $26B (2030) flows through FN's facilities
- Manufactures for COHR, Lumentum, II-VI — rising optical volume flows through FN
- Google's unique optical circuit switch (OCS) for Ironwood clusters uses more optical infrastructure than standard NVLink — ASIC adoption is optically-intensive

GPU-ASIC Agnosticism Score: High — both GPU and ASIC clusters need optical transceivers; ASIC clusters may be more optical-intensive than GPU clusters

Risk: Contract manufacturer margin compression when customers insource. Thailand manufacturing risk. No direct Baker mention.

Verdict: 🟢 SUPPORTED — Best framed as picks-and-shovels for optical content growth Baker implies.

NVT (nVent Electric) — Thermal Management / Liquid Cooling

Baker's thesis: Liquid cooling is the defining infrastructure shift of Blackwell. Mandatory, not optional. The most direct expression of the physical constraint Baker highlights.

Why It Aligns:

- Baker centers his entire Blackwell infrastructure thesis on the air-cooling → liquid-cooling transition: 30kW → 130kW per rack
- GB300 drop-in compatible with GB200 racks = liquid cooling standard locked in permanently
- Baker's space data center thesis reinforces: **cooling is the binding terrestrial constraint** on AI scale for the next 10-20 years
- nVent: ~30% of revenue from data centers; rear-door heat exchangers, immersion cooling
- Liquid cooling market: $5-7B (2025) → $15-26B by 2030
- Platform-agnostic: liquid cooling serves GPU clusters and ASIC clusters equally

GPU-ASIC Agnosticism Score: Very High — thermal physics don't care what chip is being cooled

Risk: Baker's "ROI air gap" could cause pause in new data center builds. NVT has broader industrial exposure (non-AI). No direct Baker mention.

Verdict: 🟢🟢 STRONGLY SUPPORTED — Liquid cooling is Baker's single clearest infrastructure theme.

TSEM (Tower Semiconductor) — SiPho Foundry

Baker's relevant thesis: Optics are a critical unresolved challenge. Silicon photonics (SiPho) is the leading technology for co-packaged optics. Baker's NVIDIA-wins thesis supports Tower's role as NVIDIA's SiPho foundry partner.

Why It Aligns:

- TSEM is a primary SiPho foundry; NVIDIA partnership for optical I/O components directly aligns with Baker's NVIDIA-first worldview
- CPO commercial deployment 2025-2026 creates real near-term revenue opportunity for SiPho fabs
- Baker's skepticism about non-Google/Amazon ASICs focuses foundry demand on programs serving winners — Tower serves NVIDIA, which is Baker's winner

Why It's Complicated:

- Baker's ASIC skepticism means Tower's broader custom chip foundry

TAM may be smaller than bulls expect — many ASIC programs will not survive
- Tower is a specialty foundry without leading-edge capability — the real AI silicon action is at TSMC (3nm/2nm)
- Feb 11 earnings reversal on heavy volume = distribution signal independent of thesis fit
- No direct Baker mention of Tower/TSEM
- 43% above 200-day MA = extreme extension; analyst consensus limits stated upside

GPU-ASIC Agnosticism Score: Medium — SiPho opportunity is real but concentrated in NVIDIA partnership; ASIC proliferation is mixed (smaller ASIC programs = fewer SiPho contracts)

Verdict: 🟡 SUPPORTED WITH CAVEATS — SiPho/CPO is aligned but chart structure and extension risk are independent concerns requiring separate technical validation.

ANET (Arista Networks) — AI Ethernet Networking

Baker's most relevant thesis: Ethernet is winning vs. InfiniBand for scale-out networking. UEC 1.0 ratified. Hyperscalers deploying at scale. Arista is the dominant AI Ethernet vendor.

Why It Aligns:

- Baker: Ethernet winning scale-out is structural — ANET is the primary beneficiary
- AI cluster ethernet TAM: $15B by 2027 (vs. $3B in 2023); Arista growing AI revenue at 70%+ YoY
- Ultra Ethernet Consortium: ANET is a founding member
- NVIDIA Spectrum-X (Arista-compatible) vs. Quantum-X (InfiniBand) — both Blackwell-generation; Spectrum-X gaining share rapidly
- Both GPU clusters (NVIDIA Ethernet) and ASIC clusters (hyperscalers prefer Ethernet for scale-out) use Arista infrastructure
- Baker's concern about MSFT/Meta concentration in ANET is a valid risk worth monitoring

Why It's Complicated:

- Arista's MSFT/Meta concentration (~42% of revenue) means it is correlated to hyperscaler capex decisions — Baker's "ROI air gap" risk hits ANET directly if MSFT/Meta pause builds
- InfiniBand maintains dominance in HPC training clusters; Ethernet wins for scale-out but the battle is not over for dense training
- ANET has experienced a gap-and-fade pattern on prior earnings — distribution risk on extended chart

GPU-ASIC Agnosticism Score: High — both architectures converge on Ethernet for scale-out

Verdict: 🟡 NEUTRAL TO SUPPORTED — Ethernet-wins thesis is directly bullish, but customer concentration and chart history require technical confirmation before entry.

COHR (Coherent Corp) — Optical Transceivers

Baker's relevant thesis: Optics as a critical infrastructure challenge; CPO deployment beginning; 800G/1.6T transceiver upgrade cycle; Fabrinet's manufacturing tailwind implies optical component demand growth.

Why It Aligns:

- COHR supplies 800G transceivers and is developing 1.6T products — directly in Baker's capex beneficiary zone
- CPO is COHR's next major opportunity — silicon photonics integration for Blackwell-generation AI clusters
- Baker's view that both GPU clusters (NVIDIA) and ASIC clusters (Google OCS, Amazon optical) need more optics = COHR benefits from both
- Dell'Oro: data center transceiver market growing at 25%+ CAGR through 2027

Why It's Complicated:

- COHR has been integrating a complex merger (II-VI + JDSU + Finisar) — execution risk
- Thin chart structure and history of large drawdowns; Baker did not mention COHR specifically
- Faces pricing pressure from Chinese transceiver vendors at the commoditizing low end

GPU-ASIC Agnosticism Score: High — optics demand is driven by bandwidth requirements, not by the underlying accelerator type

Verdict: 🟢 SUPPORTED — Optical content growth is a direct Baker-aligned theme, but execution risk and chart structure require independent technical validation.

PWR (Quanta Services) — Data Center Construction / Power

Baker's relevant thesis: Blackwell requires full site infrastructure overhaul — power, cooling, floor reinforcement. Baker's analogy: *"changing all outlets to 220V, putting in a Tesla Powerwall, generator, solar panels, reinforcing the floor."* This is Quanta's core competency.

Why It Aligns:

- Quanta is the dominant electrical infrastructure contractor for data center buildouts
- Baker's $660-690B hyperscaler capex forecast (2026) = enormous civil and electrical work
- Power per rack 30kW → 130kW means every AI data center is a major electrical infrastructure project — not just a real estate transaction
- PWR backlog has surged; renewables + data center segment is fastest growing division
- Platform-agnostic: whether the data center runs GPUs or ASICs, the physical infrastructure is identical

Why It's Complicated:

- PWR is already near ATH — Baker's "buy at volume nodes, not thin air" principle applies

- Less direct leverage to Baker's specific thesis vs. ALAB or NVT
- Exposed to permitting, grid interconnection delays, and labor constraints

GPU-ASIC Agnosticism Score: Maximum — physical infrastructure is completely substrate-agnostic

Verdict: 🟢 SUPPORTED — Baker's capex thesis is a powerful structural tailwind, but near-ATH chart position means this is a "watch for pullback" name, not an initiating entry.

Part V: Portfolio Strategic Implications

Conviction Tiers (Baker-Aligned)

Tier 1 — Direct Baker Endorsement (buy on technical confirmation):

Ticker	Reason
ALAB	Baker bought 1.6M shares in Q4 2025. PCIe/CXL connectivity is the clearest infrastructure play. Platform-agnostic.
NVT	Liquid cooling is Baker's single most explicit infrastructure theme. Mandatory, not optional.

Tier 2 — Strong Thesis Alignment:

Ticker	Reason
CRDO	Custom silicon proliferation is NET POSITIVE for SerDes vendors. Baker's ASIC skepticism benefits third-party I/O suppliers.
FN	Optical content per rack expanding; picks-and-shovels for Baker's optical theme.

Tier 3 — Supported, But Requires Chart Work:

Ticker	Reason
ANET	Ethernet-wins thesis is direct but MSFT/Meta concentration and chart history require Shannon-framework confirmation.
COHR	Optical content growth aligned; execution risk from merger integration is independent variable.
TSEM	SiPho/CPO real opportunity; NVIDIA partnership is the right horse; Feb 11 reversal and extension risk require fresh

	technical validation.
PWR	Baker's capex tailwind is powerful but near-ATH positioning requires pullback for entry.

Poorna's "Tops Are an Event, Bottoms Are a Process" Framework

Baker's thesis provides the **fundamental anchor** for applying this structure:

- **The thesis is intact:** Blackwell supercycle, optical interconnect buildout, liquid cooling transition — all confirmed by Baker's analysis with new supporting data points
- **The "bottoms are a process" candidates:** ALAB (down 32% post-earnings), CRDO (pulled back), COHR (below prior highs) — all have fundamental value confirmed by Baker's thesis
- **Ratio trade logic:** If you believe the thesis is intact but charts need to base, the ratio spread (long lower strikes, short higher strikes, net debit) lets you participate in a bottoming process without requiring a V-shaped recovery
- **Baker's own behavior confirms this:** He massively increased ALAB on the pullback (Q4 2025) — exactly the "frustrated seller, patient accumulator" dynamic Poorna described

The "catching falling knives" guard: Baker's framework helps here too. Companies that fail his full-stack thesis (Microsoft Maia, Meta MTIA, third-tier ASIC programs) have no fundamental anchor. The watchlist names that have fundamental anchors — confirmed by Baker's analysis — are the ones where bottoming processes are worth participating in.

Key Risk Flags from Baker's Thesis

Risk	Implication for Watchlist
Blackwell ROI Air Gap	~3 quarters of training capex without inference revenue → could pause new data center orders → temporary headwind for NVT, PWR, FN
DRAM shortage	*"First true DRAM cycle since the late '90s"* — 10× price increases possible → slows AI deployment → indirect headwind for all names
Google loses motivation	If Google loses cost leadership, could rationalize capex → ANET (MSFT/Meta concentrated), not directly watchlist-impactful
Custom silicon proliferation	If more hyperscalers succeed with ASICs (low Baker probability) → NET POSITIVE for ALAB/CRDO/FN/COHR, neutral for NVT/PWR, mixed for TSEM
Space data centers (10-20yr)	Long-term disruption to terrestrial data center growth — not a near/medium-term investment risk

---

Bottom Line: What This Interview Changed

Before this interview: Watchlist bull case rested on "AI capex supercycle" as a general thesis.

After this interview: Specific, attributable confirmation from one of the highest-conviction AI infrastructure investors that:

1. Liquid cooling is the clearest near-term infrastructure play (NVT)
2. PCIe/CXL/interconnect is where he put his own money (ALAB)
3. Ethernet wins scale-out networking (ANET structural tailwind)
4. Custom silicon is mostly dead ends — third-party I/O vendors win either way (CRDO, FN, COHR)
5. Optical content growth is the multi-year secular trend (FN, COHR, TSEM/SiPho)
6. The capex cycle is ROI-positive and driven by compounding competitive dynamics — not hype

The one thing Baker adds that wasn't in our thesis: The DRAM risk. If memory prices go up 10×, AI deployment could slow meaningfully. This is a macro variable worth monitoring — not a reason to exit, but a reason to keep cash available to add on weakness.

Sources: Invest Like the Best EP.451 (Dec 9, 2025), Atreides Management Q4 2025 13F, 247wallst.com, podchemy.com, podbrain.app, podpulse.ai, theneuron.ai, CNBC AI Chip Comparison (Nov 2025), Dell'Oro Group ethernet/InfiniBand data, Futurum AI Capex 2026 analysis, howaiworks.ai GPU/TPU/ASIC market analysis.