

AESOP ANSWERS

by Integrity Defence Solutions UK Ltd

Neurophos, Inc.

*Investment Analysis at a \$1.0B Hypothetical Entry — Photonic Prefill
Acceleration for AI Inference*

QUESTION ADDRESSED

Should we make a hypothetical \$1.0 billion primary investment in Neurophos, Inc. — and, if so, what range of outcomes should we expect?

Reference	AS/260914/194 (V1.1)
Date	16 September 2026
Prepared for	XXXXXX
Produced by	AESOP Answers, Integrity Defence Solutions UK Ltd

This report is provided under the AESOP Answers Terms and Conditions. It does not constitute investment, legal or financial advice, and no liability is accepted for decisions made in reliance on it. See the Terms and Conditions for permitted use, distribution and IP terms.

EXECUTIVE SUMMARY

Neurophos, Inc. is a pre-revenue, Austin-based photonic-computing startup (founded 2020, a Duke University metamaterials spinout) developing an optical processing unit (OPU). The company's own materials position it specifically as a prefill accelerator — hardware for the compute-bound phase of large-language-model inference, which Neurophos states represents roughly 93% of inference tokens — intended to run alongside, not replace, the GPU racks that handle the memory-bandwidth-bound decode phase. That is a narrower and more specific thesis than "photonic compute", and it is the frame this report uses to evaluate the company. Neurophos has raised ~\$118–120M through a January 2026 Series A; no post-money valuation has ever been disclosed. This report evaluates a hypothetical \$1.0B primary investment — roughly a 1.8x–3.2x step-up over the estimated Series A price, with no disclosed intervening milestone.

RECOMMENDATION: WAIT — PASS ON THIS PRICE, NOT ON THE COMPANY

A small toehold is defensible only if it secures contractual pro-rata into the Series B plus information rights. The full reasoning, and what would move this call in either direction, is set out in the Recommendation section (Section 6).

The two variables that dominate the outcome — usable numerical precision at deployable bit-width, and measured (not projected) system-level efficiency — will be substantially resolved by a specific, already-scheduled near-term event: an integrated demonstrator, followed by a 2027 pilot with Terakraft in Norway. Paying a rich step-up today for the option to wait for that evidence anyway is close to buying sentiment rather than de-risked progress.

Two things follow from that narrower frame. First, Neurophos's real competitive set is not primarily other photonics companies — it is a wave of well-funded, all-electronic transformer and prefill-specific accelerators (Etched's Sohu, Taalas, d-Matrix, Cerebras) attacking the identical compute-bound bottleneck without Neurophos's optical-specific unknowns, and several are further along: Etched has working silicon, over \$1B in signed contracts, and a reported \$21B valuation as of an August 2026 raise. Second, a body of published photonic-accelerator research (not specific to Neurophos) establishes that analogue-to-digital and digital-to-analogue (ADC/DAC) conversion and data-movement overhead routinely dominate system-level power in this device class — in some published designs, conversion circuitry alone consumes over half of chip area, and ADCs dominate energy even when the optical core itself is efficient. That is a specific, evidence-based reason to discount the headline efficiency claim before it reaches system-level reality.

At \$1.0B, Neurophos is priced as though it has cleared a bar its more mature electronic competitors have not yet definitively cleared for themselves, while carrying an additional, well-documented category of system-level power risk (electro-optic conversion overhead) that those competitors do not carry at all. At 7 years, the probability-weighted view puts failure/distress at roughly 35–48%, an indefinite stall at 22–33%, a defensible niche outcome at 10–18%, a strong outcome (comparable to Ayar Labs or Lightmatter, \$8–18B) at 6–13%, and a category-defining breakout at 1–3%. The illustrative expected MOIC is $\approx 1.1x$ excluding the breakout tail and $\approx 1.5x$ including it — but the median outcome is close to or below capital back, and at this entry price the single most probable good outcome for the company (a \$1–1.5B strategic acqui-hire) is actually a loss for this investor.

The core risk is that the headline efficiency claim — a projection at 4-bit precision, benchmarked against GPUs measured at 8-bit, with no published reconciliation, and one that has already drifted from "100x" at the January raise to the CEO's own "~10x Rubin-class GPUs" framing by July 2026 — has to survive the electro-optic conversion and data-movement overhead that typically dominates system power in this device class, and Neurophos has disclosed no measured breakdown. Against that, the strategic-acquirer environment for inference silicon is the strongest in a decade (Celestial AI → Marvell at \$3.25–5.5B; Groq's assets → Nvidia at ~\$20B; Etched's own trajectory to a reported \$21B), which is the real argument for retaining some optionality rather than passing outright.

This position also shares financing-environment risk with the existing Gravitics holding — both depend on continued access to large private rounds for unproven physical technology through 2027–2030.

A Note on Verification

This analysis draws solely on public information and was independently verified as of 15 September 2026. Internal company data may provide an alternate view.

TABLE OF CONTENTS

Executive Summary.....	2
1. The Opportunity at a Glance	5
2. The Evidence Base	7
3. Technology, Market and Regulatory Assessment.....	9
4. Outcome Scenarios and Returns.....	14
5. What Would Change This View.....	18
6. Recommendation.....	20
7. Stress-Testing This Conclusion	22
Appendices	23
Appendix A — Methodology Note	
Appendix B — Assumption Register	
Appendix C — Diligence Requirements	
Glossary	26
Annex A — Evidence Log.....	28

1. THE OPPORTUNITY AT A GLANCE

Initial Screening

Neurophos closed a \$110M Series A on 22 January 2026, led by Gates Frontier, with M12 (Microsoft), Aramco Ventures, Bosch Ventures, Carbon Direct Capital, Space Capital and others — total funding ~\$118–120M. No Series A post-money has ever been disclosed by the company, an SEC filing, or a COI-based data provider; the only public estimate remains PremierAlts' secondary-market figure of ~\$310M, which is treated here as low-confidence. Dilution arithmetic on a \$110M primary round implies a post-money band of roughly \$310M–\$550M. A \$1.0B entry is a ~1.8x–3.2x step-up over a seven-month-old round with no disclosed intervening milestone.

Sector context makes \$1.0B plausible rather than absurd: Ayar Labs sits near \$1B on ~\$155M raised (shipping optical I/O, Intel/AMD/NVIDIA ecosystem); Lightmatter was \$1.2B on ~\$420M raised in December 2023 and \$4.4B on \$850M raised by October 2024; Celestial AI was ~\$2.5B in March 2025 before its December 2025 acquisition by Marvell at \$3.25–5.5B. Neurophos at \$1.0B prices roughly at Ayar's level on materially less product maturity — a test chip, not a shipping product.

Enough is public to build a real assessment — including enough to identify Neurophos's actual product category (a prefill accelerator, not general photonic compute) and to name specific, more mature competitors within that category. Not enough to underwrite. Proceed — with the price, and a competitive-crowding flag, carried forward for scrutiny throughout.

Deal Snapshot

Field	Detail
Company	Neurophos, Inc. — Austin, TX. Founded 2020. Duke University (David Smith lab) / Metacept metamaterials spinout. Founders Patrick Bowen (CEO) and Tom Driscoll.
Product	Optical Processing Unit (OPU) for AI inference — metasurface optical modulators forming an in-memory optical tensor core. First part: "Tulkas T100." Company-stated as a prefill-specific accelerator (the compute-bound phase of LLM inference), designed to run alongside existing GPU racks that handle decode.
Stage	Pre-revenue, pre-product. Test chip fabricated and demonstrated; proof-of-concept chip in progress. No integrated system, no commercial tapeout. Full production not expected before mid-2028; initial volumes described by the CEO as "thousands, not tens of thousands" of chips.
Hypothetical entry	\$1.0B post-money, primary, Series-B-equivalent
Recommendation	WAIT — pass on the price, not the company (see Recommendation, Section 6)
Core thesis if it works	Prefill (input-token processing) is a compute-bound bottleneck distinct from memory-bound decode; optical matrix multiplication at low precision is a physically credible way to accelerate it, if system-level overhead does not erode the device-level gain. AI-silicon exits are currently clearing at \$3B–\$20B.
Biggest single risk	The headline efficiency figure — a projection at 4-bit precision, benchmarked against GPUs measured at 8-bit and never reconciled publicly — has to survive electro-optic conversion (ADC/DAC) and data-movement overhead that, per published literature on comparable device classes, typically dominates system power. Neurophos has disclosed no measured breakdown.
Biggest competitive risk	Well-funded, more mature electronic prefill/transformer ASICs (Etched, Taalas, d-Matrix, Cerebras) targeting the identical bottleneck without Neurophos's optical-specific unknowns.

Eclipse Space orbital-compute exploration was confirmed directly by CEO Bowen (July 2026 interview); the same interview has him citing "~10x Rubin at similar power" — down from "100x" at the January raise — while calling the numbers theoretical with no ship date before 2028.

2. THE EVIDENCE BASE

The facts below are graded by source reliability, from SEC filings down to third-party estimates that cannot be independently confirmed. Third-party valuation estimates in particular — including PremierAlts' ~\$310M figure below — are unverified and are used here only as directional context; no representation is made as to their accuracy.

Tier 1 — Primary source

- SEC Form D exempt-offering notice: offering amount only — no valuation, preference terms, burn, or cap table.

Tier 2 — Company-disclosed, multi-outlet corroborated

- Seed (\$7.2M, December 2023) and Series A (\$110M, 22 January 2026) amounts, dates, and investor syndicate.
- Duke/Metacept origin, founders, roadmap (evaluation hardware 2026, commercial ramp 2028).
- Terakraft (Norway) collaboration agreement, announced August 2025; pilot planned as part of a 2027 commercial early-access programme.
- Eclipse Space orbital-compute exploration — confirmed directly by CEO Bowen, July 2026.
- Prefill-specific positioning: the company states its focus is prefill ("the compute-bound bottleneck that represents 93% of inference tokens"), with decode described as memory-bandwidth-bound and left to existing GPUs. Stated on Neurophos's own site and corroborated independently by The Register, TechRadar and TechSpot (January 2026 coverage).
- T100 system specifications: dual-reticle package, 768GB HBM, projected 470 petaOPS FP4/INT4 compute at 1–2kW total package power under load. Reported by The Register, TechSpot and TechRadar (January 2026), consistent across outlets.
- Company engineers describe a "dense supporting architecture of vector units and SRAM" (internally nicknamed "the boondoggle") required to feed the 56GHz optical core fast enough to be useful. Reported by The Register.
- CEO Bowen states full production is not expected before mid-2028 and initial volumes will be "thousands, not tens of thousands" of chips (The Register, January 2026).

Tier 3 — Company claim (self-reported, unverified by a third party)

- Efficiency claims have moved materially: ~100x (January 2026) → ~300 TOPS/W (August 2025, Terakraft announcement) → "~10x Rubin at similar power" (Bowen, July 2026), with no published 4-bit/8-bit reconciliation.
- Terakraft engagement's commercial terms (paid vs. unpaid) undisclosed.
- The 93%-of-inference-tokens-are-prefill figure is Neurophos's own framing statistic; no independent, third-party-verified source for this precise figure was found. The true prefill/decode split varies by context length, batch size and model architecture, and is shifting as long-context and agentic workloads grow decode's share.

Tier 4 — External estimate or general field evidence, not Neurophos-specific

- Series A post-money ≈ \$310M (PremierAlts; not COI-based, unconfirmed).
- Employee counts vary by database — indicative only.
- Multiple 2025–2026 academic papers on photonic AI accelerators (general field literature, not about Neurophos) establishing that ADC/DAC conversion, modulator drive and data-movement overhead commonly dominate system-level power and area in this device class. Used here strictly as an outside-view stress test, not as a measurement of Neurophos's actual chip (see Section 3).

- Comparable-company disclosures for Etched (Sohu), Taalas, d-Matrix and Cerebras — funding, valuation, contract and shipping-timeline data as publicly reported (see Section 3).

Tier 5 — Unknown

- Series A pre/post-money, share price, preference structure. Burn rate and runway.
- Foundry/OSAT identity; packaging cost curve. Measured system-level efficiency or accuracy at any bit-width; any subsystem-level power breakdown (optical core / ADC-DAC / SRAM and vector units / HBM).
- Whether core metasurface IP is owned or licensed, and on what terms. Board composition and governance rights.

Tier 6 — Requires confidential diligence

- Cap table and preference waterfall; audited financials; test-chip characterization data; yield data; foundry/OSAT agreements; the Terakraft contract; hyperscaler NDA/eval agreements; IP/FTO opinion; the SiliconBee acquisition terms; pro-rata/investor-rights terms; export-control classification work.

3. TECHNOLOGY, MARKET AND REGULATORY ASSESSMENT

Technology and Product

What Neurophos needs to become

Neurophos's own public materials are explicit and specific. LLM inference splits into a compute-bound prefill phase (processing the input prompt) and a memory-bandwidth-bound decode phase (generating output tokens). Neurophos states that decode is "memory-bandwidth limited, not compute-limited" and that it is building specifically for prefill, which it states represents 93% of inference tokens. The company's stated go-to-market pairs T100 racks alongside existing GPU racks rather than replacing them — Neurophos is a compute co-processor for one phase of inference, not a GPU-replacement platform.

This framing changes three things that materially affect the analysis:

- The competitive set shifts from photonics companies to prefill and compute-bound accelerator companies generally, most of which are all-electronic (see Comparable Company Set below).
- The bar for "success" changes: Neurophos does not need to beat GPUs at everything. It needs to beat GPUs and its prefill-specific competitors, specifically at the compute-bound, low-precision matrix-multiply workload, on cost and power, at a scale a hyperscaler will actually deploy.
- The 93% prefill-token-share figure becomes a distinct, named and unverified assumption underlying the entire addressable-market case for the company — not something to accept at face value (see CU-9 below and A9 in Appendix B).

The underlying science is credible: metasurface modulators at $\sim 5\mu\text{m}$ pitch, over 1M elements on a $5\times 5\text{mm}$ die, fabricated in silicon and physically demonstrated — the enabling premise for a 1000×1000 optical tensor core.

Is the underlying science credible?

Yes — modulator miniaturization is demonstrated in silicon, the physics lineage is strong, and sophisticated technical investors (Microsoft's M12) underwrote the round.

Has a commercially viable product been demonstrated?

No, and not close. No commercially viable, integrated, prefill-benchmarked product has been demonstrated: no integrated system, no measured efficiency, no accuracy result on a real model, no commercial tapeout. The company is still validating a proof-of-concept chip ahead of a mid-2028 target for full production, with initial volumes the CEO himself describes as "thousands, not tens of thousands" of chips. "The boondoggle" — the company's own internal term for the vector-unit and SRAM architecture required to feed the 56GHz optical core fast enough to be useful — is precisely the kind of supporting electronic infrastructure where, per the published literature discussed next, efficiency gains are typically lost in this device class.

End-to-End System Stress Test — Device-Level Claim vs. System-Level Reality

The headline efficiency figure is a 4-bit projection benchmarked against GPUs measured at 8-bit, with no published reconciliation. Stated only that way, the risk is correctly identified but incomplete — it treats the gap as a precision-accounting technicality rather than naming the more fundamental system-level mechanism most likely to erode it.

A body of published photonic-accelerator research (multiple 2025–2026 papers; none specific to Neurophos) establishes a consistent pattern: analogue-to-digital and digital-to-analogue conversion (ADC/DAC), electro-optic modulator drive circuitry, and data movement and memory access routinely dominate the total energy and area budget of photonic AI accelerators. In some published designs, signal-conversion circuitry alone consumes over half of chip area; in others, ADCs specifically dominate

total energy consumption even when the optical core itself performs computation efficiently. This is a known, generalisable engineering reality of the device class Neurophos is building in — not a Neurophos-specific finding, since Neurophos has not published its own ADC/DAC count, resolution or measured power draw.

What public information does allow is cross-referencing this literature against what Neurophos has itself disclosed about the T100. The company's own description — a dual-reticle package, 768GB of HBM, and (per its own engineers, quoted in trade press) a "dense supporting architecture of vector units and SRAM" nicknamed "the boondoggle" needed to keep the 56GHz optical core fed — is itself an acknowledgment that a large share of the chip's area and power budget sits outside the optical tensor core. The company's own projected envelope — 470 petaOPS at 1–2kW total package power — is already a system-level number: the "100x" and "~10x Rubin" claims have to survive HBM power draw, SRAM access energy and conversion overhead to be real. That is a point in the company's favour on framing (the claim is not being compared unfairly to a system-level GPU number from a device-level-only photonic number), but it does not resolve whether the published conversion-overhead pattern applies at Neurophos's specific claimed magnitude — the company has disclosed no measured breakdown that would let an outside analyst check it.

Net assessment: this specific gap cannot be closed from public information alone. It requires Neurophos to publish, or a diligence process to obtain, a measured power breakdown by subsystem (optical core / ADC-DAC / SRAM and vector units / HBM / interconnect) benchmarked on a real transformer workload at deployable precision. What public information does provide is a sharper, evidence-based reason to weight this risk heavily (CU-1 and CU-2 in the Critical Uncertainty Register below), rather than treating it as a reconciliation footnote.

Comparable Company Set

Table A — Optical and interconnect comparables (useful for photonics-sector valuation levels, weaker as a product-level analogue)

Company	Founded	Raised	Mark	Status at mark
Ayar Labs (current)	2015	~\$964M (2026)	\$3.75B (Mar 2026)	Shipping optical I/O; Intel/AMD/NVIDIA ecosystem
Ayar Labs (earlier, cross-check)	2015	~\$155M	~\$1B	Cross-check reference
Lightmatter (current)	2017	~\$850M	\$4.4B (Oct 2024)	Passage (interconnect) shipping 2026
Lightmatter (earlier)	2017	~\$420M	\$1.2B (Dec 2023)	Two product lines in development
Celestial AI	2020 (same age)	~\$600M	Acquired by Marvell, Dec 2025, \$3.25–5.5B	Photonic Fabric, hyperscaler engagement
Neurophos (hypothetical)	2020	~\$120M	\$1.0B	Test chip; 2028 revenue target

Structural observation: every strong outcome in this set is fundamentally an interconnect play. Every full-compute attempt — including Lightmatter's own Envisé — stalled or was de-emphasized. Neurophos is the only well-funded company still pursuing full compute-in-memory, on roughly one-seventh of Lightmatter's capital.

Table B — Compute-bound and prefill-specific accelerator comparables (electronic, not photonic, but solving Neurophos's stated problem)

Company	Approach	Funding / status
Etched (Sohu)	Founded 2022. Transformer-only ASIC; hardwires the attention mechanism into silicon rather than running programmable kernels. Working A0 silicon on TSMC N4P; 8-chip rack-scale system unveiled.	Raised ~\$800M across four rounds. Reported \$500M round at ~\$5B (2025), then ~\$700M at a reported \$21B (August 2026). \$1B+ in signed customer contracts; first racks shipping mid-2026. Claims 500,000 tokens/sec on Llama 70B (~20x an 8-GPU H100 system) under prefill-favourable benchmark conditions (short output, large batch); no independent third-party benchmark yet.
Taalas	Finnish startup. Most radical approach: burns a specific model (Llama 3.1 8B) directly into silicon, trading all flexibility for throughput.	First hardware targeted 2027. Pre-commercial.
d-Matrix (Corsair)	Digital in-memory compute (DIMC) chiplets targeting low-latency LLM inference. 8 chiplets, 512MB SRAM per card, custom BFP12 (block floating point) format.	Commercialising; specific deployment scale not independently disclosed.
Cerebras (WSE-3)	Wafer-scale engine — the opposite architectural bet from Taalas (go huge rather than shrink); eliminates chip-to-chip communication overhead.	Established company, multiple generations shipped; deployed at scale with several customers.

Structural observation on Table B: Etched alone is a closer economic analogue to the \$1.0B Neurophos ask than any photonics comparable — both are single-purpose accelerators betting on a specific inference bottleneck, at a similar level of ambition. But Etched is already meaningfully ahead on de-risking milestones Neurophos has not reached: working silicon (vs. Neurophos's pre-integration test chip), disclosed customer contracts (vs. none disclosed for Neurophos), and a nearer shipping date (2026 racks vs. Neurophos's mid-2028 target) — all while carrying none of Neurophos's optical-specific conversion risk. If \$1.0B is roughly where the market prices a "credible, unproven, prefill-specific accelerator", Etched's own trajectory suggests the market rewards de-risking milestones very quickly and would likely punish falling behind on them just as quickly. Neurophos is being asked to clear that bar from further back, while carrying an additional risk class the electronic leaders do not have.

Regulatory and Geopolitical

- Export controls: the AI Diffusion Rule was informally rescinded (May 2025) but not formally removed from the CFR; BIS is non-enforcing pending a replacement. The Affiliates Rule was published then paused. Net: real but currently fluid exposure.
- Deemed export: a rapidly-hiring photonics/AI team will include foreign nationals with access to controlled technology.
- Government/defense demand: confirmed optionality via Eclipse Space, unquantified — no contract, funding, or timeline disclosed.
- Foreign LP exposure: Aramco Ventures' Series A participation makes a CFIUS-at-exit question realistic.

None of this is disqualifying on its own. All of it belongs in Tier-1 diligence.

Portfolio Context

Checked against the existing Gravitics holding. Both are capital-intensive, long-dated, pre-revenue hardware exposures whose dominant shared risk is the availability of large private financing rounds for unproven physical technology through 2027–2030 — not diversified by holding both, doubled. A soft thematic thread: Space Capital is a Neurophos investor, and Neurophos separately explores orbital compute via Eclipse Space.

Critical Uncertainty Register

#	Variable	Impact	Confidence	Raises confidence	Lowens confidence
CU-1	Usable numerical precision (4-bit ceiling vs. production need) and accuracy parity	Very High	Very Low	Third-party benchmark at stated efficiency; published 4-bit/8-bit reconciliation	Continued reliance on 4-bit figures, no accuracy data
CU-2	System-level power: ADC/DAC conversion and data-movement overhead eroding the device-level efficiency claim (see System Stress Test above)	Very High	Very Low	Measured, subsystem-level power breakdown (optical core / ADC-DAC / SRAM and vector units / HBM) published or obtained in diligence	No breakdown published; continued reliance on device-level or fully projected system numbers
CU-3	Manufacturability at volume — uniformity, yield, packaging cost	Very High	Very Low	Named tier-1 foundry/OSAT; packaging within 2–3x electronic assembly	Foundry undisclosed; packaging cost 10x+
CU-4	Capital sufficiency — plausibly \$300–800M+ still required	Very High	Low	Series B at a clear step-up from a credible outside lead	Insider bridge; down round
CU-5	Commercial pull — engagement converting to a paid design win	High	Low	Paid NRE/production commitment; Terakraft converts to purchase	Terakraft slips; no named paying customer by 2028
CU-6	IP ownership / freedom to operate	High	Very Low	Clean FTO opinion; confirmed ownership	Royalty-bearing/field-limited license found
CU-7	Competitive erosion — Rubin-generation gains, hyperscaler internal silicon, CPO interconnect	High	Medium	Efficiency gap holds $\geq 10x$ at pilot time	Rubin gains narrow gap faster than roadmap slips
CU-8	Entry price integrity	High	Low	Verified Series A post-money + milestone	Post-money confirmed ~\$310M, no milestone
CU-9	Competitive crowding: mature, well-funded electronic prefill/transformer ASICs (Etched, Taalas, d-Matrix,	High	Medium	Neurophos demonstrates a decisive, independently verified advantage specifically over Etched-class electronic accelerators — not just over Nvidia GPUs	Etched/Taalas/d-Matrix ship commercially before Neurophos completes its proof-of-concept-to-production transition

#	Variable	Impact	Confidence	Raises confidence	Lowers confidence
	Cerebras) reaching commercial deployment faster, without Neurophos's optical-specific risk				

Highest-priority diligence: CU-1, CU-2, CU-3, CU-4, CU-9.

4. OUTCOME SCENARIOS AND RETURNS

Outcome Pathways

Entry: \$1.0B post-money. Retention factor sensitivity: ~0.55 (range 0.40–0.70) — 2–3 further rounds plus pool expansion are expected before a terminal event. Pathway conditions reference competitive crowding (CU-9) and system-level overhead (CU-2) directly.

#	Pathway	Conditions	Terminal state	Terminal value	MOIC
1	Failure / Distress	Precision/packaging unsolvable at cost; demonstrator fails or slips indefinitely; runway ends 2028–2030; or an electronic competitor commoditises the compute-bound prefill niche first	Shutdown or acqui-hire below preference stack	\$0–\$400M	0–0.3x
1b	Reverse-acquihire below entry	Team + IP absorbed (cf. Groq ~\$20B, Enfabrica ~\$900M)	Acqui-hire	\$300M–\$1.5B	0.2–0.9x — a loss at entry
2	Survives but Stalls	Real but 3–10x efficiency, not 100x; flat/down rounds; Nvidia closes most of the gap; or Etched-class ASICs capture the addressable prefill-offload market before Neurophos reaches production	Long private hold or modest sale	\$400M–\$1.5B	0.2–0.8x
3	Niche Commercial Success	Defensible subset (power-constrained, edge/orbital via Eclipse Space); several design wins; coexists with, rather than is displaced by, electronic prefill ASICs	Strategic acquisition or private hold	\$2B–\$5B	1.0–2.5x
4	Strong Commercial Success	Efficiency holds at production precision with a demonstrated advantage over Etched-class electronic competitors specifically, not just Nvidia; multiple hyperscaler deployments	Large acquisition or IPO path	\$8B–\$18B	4–8x
5	Category-Defining Breakout	Optical compute becomes a standard inference substrate, displacing both GPUs and electronic prefill ASICs	Very large acquisition or public listing	\$25B–\$60B+	12–30x

Probability Ranges

Explicitly subjective, analog-anchored ranges — no defensible base rate exists for this population. The ranges are deliberately wide: the competitive-crowding and conversion-overhead mechanisms sharpen what drives the outcome without producing a more precise number to attach to it. Treat these as brackets reflecting a considered judgment, not a calibrated statistical estimate — see Section 7 and Appendix A.

Outcome	3-year	5-year	7-year	Why broad
1. Failure/distress	6–14%	22–38%	35–48%	Two silicon generations, unsolved precision, 2–3 financings must all go right, and the company must avoid being commoditised by faster-moving electronic competitors
2. Survives but stalls	50–70%	30–45%	22–33%	Echodyne lineage; Lightmatter's own compute retreat; Etched/Taalas/d-Matrix capture much of the addressable prefill-offload demand
3. Niche success	3–8%	7–15%	10–18%	Real technical asset plus specialty and orbital markets support a floor even if electronic ASICs win the general prefill market
4. Strong success	1–3%	3–8%	6–13%	Celestial precedent; strategic-buyer appetite high; requires a demonstrated edge over electronic prefill ASICs specifically
5. Breakout	~0%	0.3–1.5%	1–3%	Groq shows it's possible; evidence supports only low single digits even before accounting for electronic competition

What moves these: CU-1/CU-2 resolving positively (a published, favourable subsystem power breakdown) shifts 10–15 points from Pathways 1–2 into 3–4. CU-9 resolving negatively (an Etched-class ASIC ships commercially, independently verified, before Neurophos exits proof-of-concept) shifts 10–15 points further into Pathways 1–2. CU-4 resolving negatively, or CU-7 showing Rubin outpacing the roadmap, shifts 10–20 points into Pathways 1–2.

Return Ranges: MOIC and IRR

Horizon framing — why 3 years is treated differently

- **3 years (2029):** demonstrator result, tapeout status, Terakraft pilot data, a Series B/C outcome, and Etched/Taalas/d-Matrix's own commercial results by then are the realistic events by this point. Any "return" here is almost certainly a private mark from a follow-on financing, not a liquidity event — treat the 3-year column below as an illustrative interim mark conditional on which pathway is emerging, not a cash return.
- **5 years (2031):** first systems shipped or not; early production economics; whether a strategic buyer has moved. Liquidity is possible but not likely.
- **7 years (2033):** scaled revenue, acquisition, or terminal private value — the realistic liquidity window, and the horizon the pathway MOIC ranges above are built around.

Pathway	3-yr mark (illustrative, pre-liquidity)	5-yr MOIC (if resolved early)	Approx. 5-yr IRR	7-yr MOIC	Approx. 7-yr IRR
1. Failure	0–0.5x	0–0.3x	–100% to –22%	0–0.3x	–100% to –16%
1b. Reverse-acquire	0.2–0.9x (can occur this early)	0.2–0.9x	–27% to –2%	0.2–0.9x	–19% to –1.5%
2. Stall	0.5–1.2x	0.3–0.9x	–21% to –2%	0.2–0.8x	–19% to –3%
3. Niche	1.0–2.0x	1.0–2.0x	0% to +15%	1.0–2.5x	0% to +14%

Pathway	3-yr mark (illustrative, pre-liquidity)	5-yr MOIC (if resolved early)	Approx. 5-yr IRR	7-yr MOIC	Approx. 7-yr IRR
4. Strong	2.0–4.0x	3–6x	+25% to +43%	4–8x	+22% to +35%
5. Breakout	4.0–8.0x	8–15x	+51% to +72%	12–30x	+43% to +63%

The 3-year marks are wider and more speculative than the 5- or 7-year figures — at that point the CU-1/CU-2 demonstrator evidence is only just arriving, so the mark mostly reflects which narrative the market believes, not resolved fact.

MOIC ranges per pathway are unchanged by the competitive and system-level evidence above — that evidence affects which pathway is more or less likely, not the dollar outcome conditional on landing in a given pathway.

Feasibility test — feasibility is not probability

Target MOIC	Required terminal value	Precedent check
2x	~\$3.6B	Below Celestial's \$3.25B upfront — clearly feasible
10x	~\$18B	Near Groq's ~\$20B — feasible, once-a-decade outcome
20x	~\$36B	No photonics/AI-silicon startup has reached this
50x+	~\$90B+	Not credible on this evidence

Valuation cross-check

Company	Mark	Raised	Status at mark
Neurophos (hypothetical)	\$1.0B	\$120M	Test chip; 2028 revenue target
Ayar Labs (earlier)	~\$1B	~\$155M	Shipping optical I/O
Lightmatter (Dec 2023)	\$1.2B	~\$420M	Two product lines in development
Celestial AI (Mar 2025)	~\$2.5B	~\$600M	Hyperscaler engagement

Conclusion: \$1.0B prices Neurophos above Ayar Labs' own \$1B mark despite far less maturity.

Illustrative Expected MOIC

Using the probability midpoints and MOIC midpoints above (0.15x / 0.5x / 1.75x / 6x / 21x for Pathways 1–5), normalized because the ranges deliberately don't sum to exactly 100%:

Pathway	Probability (normalized)	MOIC midpoint	Contribution
1. Failure	43.9%	0.15x	0.066
2. Stall	29.1%	0.5x	0.146
3. Niche	14.8%	1.75x	0.259
4. Strong	10.1%	6x	0.606
5. Breakout	2.1%	21x	0.441

Sum $\approx 1.52x$ including the full tail. Excluding Pathway 5 — i.e. what a purely ex-tail, modal-outcome view looks like — renormalizing Pathways 1–4 gives $\approx 1.10x$.

Both figures are explicitly subordinated to the median, not treated as the headline number — the distribution is right-skewed, and the median outcome (Pathway 2, "stalls") is a mark near or below entry, not a return. This is arithmetic on the ranges above, not an independent estimate — it will move if the probability or MOIC ranges move. The competitive-crowding weight (CU-9) is what moves probability mass from Pathways 3–4 into Pathway 2 and holds both figures down.

5. WHAT WOULD CHANGE THIS VIEW

Sensitivity

Driver	Priority	Note
Ownership retention / dilution	Dominant	Driven by CU-4. ± 0.15 on the 0.55 retention factor swings Pathway-4 MOIC from $\sim 4.8x$ to $\sim 8.4x$.
Conversion-overhead severity (device-vs-system efficiency gap)	Dominant	Driven by CU-2. If published or diligenced data shows conversion overhead consumes the majority of system power (consistent with comparable published designs), the headline efficiency claim compresses toward the low end of every positive pathway simultaneously.
Competitive timing gap vs. Etched-class ASICs	Important	Driven by CU-9. Every year Neurophos remains pre-production while Etched, Taalas or d-Matrix ship commercially narrows the window in which Pathways 3–4 remain available.
Terminal exit multiple in Pathway 4	Important	$\$8B$ vs. $\$18B$ is $4x$ vs. $8x$ — sector-sentiment-dependent, largely uncontrollable.
Commercialization timing	Important	Every year of slip compounds dilution and competitive erosion (CU-7, CU-9) simultaneously. A 2-year slip likely converts Pathway 4 into Pathway 3.
Packaging cost per unit	Important	Determines whether Pathway 3 is profitable or merely alive.
Prefill token-share assumption (the 93% figure)	Secondary but structural	If the actual production prefill/decode split is meaningfully lower than 93% (e.g. long-context and agentic workloads shift weight toward decode), the addressable market underlying every positive pathway narrows — a market-sizing risk layered on top of the technical risks.
Breakout probability	Secondary — counterintuitively	Moving breakout from 1% to 3% changes expected MOIC by $\sim 0.4x$ but changes the median not at all. Shouldn't drive the decision.
TAM estimates	Immaterial	The market is obviously large enough. Nothing about the decision turns on TAM.
Regulatory/export outcomes	Secondary	Meaningful for exit mechanics and timing, not for whether the technology works.

The decision is dominated by price, system-level precision and overhead risk, and competitive timing — in that order.

Forward Watchlist

Precise posteriors aren't defensible here — no credible prior/likelihood exists for "probability an independent benchmark confirms the claim." Directional updates instead:

- Integrated demonstrator measures within 2x of projection at production precision: CU-1/CU-2 confidence moves Very Low/Low to High; shift 10–15 points from Pathways 1–2 into 3–4; reconsider the recommendation.
- Demonstrator slips past mid-2027, or only component-level data circulates: CU-2 stays Low/Very Low; shift mass toward Pathways 1–2.
- Neurophos (or independent diligence) publishes a measured, subsystem-level power breakdown (optical core / ADC-DAC / SRAM and vector units / HBM) on a real transformer workload: the single most direct resolution of CU-1/CU-2.

- Etched's Sohu ships and independent (non-company) benchmarks confirm or refute its prefill throughput claims: recalibrates how hard the prefill-acceleration problem is and how fast it can be solved electronically — the baseline against which Neurophos should be judged (CU-9).
- Any hyperscaler or independent lab publishes data on the actual production prefill/decode token-share split: tests the 93% figure Neurophos self-reports (Section 3; A9).
- Taalas or d-Matrix hit or miss their own stated 2027 milestones: further competitive calibration for CU-9.
- Series B closes led by a credible new outside investor at a genuine step-up: CU-4/CU-8 improve; supports moving off WAIT.
- Nvidia's Rubin-generation efficiency data publishes and the claimed gap narrows materially: CU-7 confidence moves down; shift mass toward Pathways 1–2.
- Eclipse Space relationship converts to a disclosed contract or funded programme: CU-5 improves; Pathway 3 (niche/orbital) probability mass increases specifically.

6. RECOMMENDATION

Kill Criteria and Conviction Builders

Kill criteria — any one should end pursuit at this price:

- Integrated demonstrator fails to reach within ~3x of projected efficiency, or slips beyond 2027.
- Accuracy on production models proves unrecoverable at achievable precision (CU-1 resolves negative).
- Packaging cost per module lands more than 10x electronic assembly with no credible path down.
- Core metasurface IP proves licensed on royalty-bearing or field-limited terms.
- Series B is an insider bridge, a down round, or carries heavy structure.
- Terakraft pilot is cancelled or indefinitely deferred with no replacement.
- Rubin-generation gains close the claimed gap materially faster than Neurophos's roadmap slips.
- An Etched-class electronic prefill ASIC ships commercially at scale with independently verified performance before Neurophos completes integration — this would remove much of Neurophos's remaining option value.

Conviction builders — would justify moving from WAIT to INVEST:

- Measured, independently reviewed efficiency ≥ 100 TOPS/W at production-relevant precision.
- Accuracy parity on a standard transformer benchmark.
- Named tier-1 foundry and a disclosed, credible packaging path.
- A paid production commitment or funded NRE from a hyperscaler, neocloud, or a materialized Eclipse Space contract.
- Series B led by a credible new outside investor at a genuine step-up, on clean terms.
- Verified Series A post-money confirming the entry price is a defensible step-up.
- Neurophos publishes, or diligence obtains, a measured system-level power breakdown showing conversion overhead is a minority, not a majority, of total power (directly addresses the System Stress Test in Section 3).
- Neurophos demonstrates a decisive, independently verified efficiency or cost advantage specifically over Etched/Taalas/d-Matrix-class electronic accelerators on the actual compute-bound prefill workload — not just over Nvidia GPUs.

The first, second and fourth conviction builders above, together with the last two, are what would justify paying more than \$1.0B later.

The Recommendation: WAIT

Not Pass on the company. Pass on this price, today.

CU-1 and CU-2 will be substantially resolved by a specific, already-scheduled event: the integrated demonstrator, then the 2027 Terakraft pilot — and, independently, by whether Neurophos ever publishes the subsystem-level power breakdown that would let an outside party check the efficiency claim against the pattern documented in Section 3. Confidential diligence today sharpens the picture at the margin but cannot resolve either dominant uncertainty. Paying an estimated 1.8–3.2x step-up for the privilege of waiting for that data anyway is close to buying sentiment — and it is buying sentiment in a crowded field, since Etched, Taalas and d-Matrix are pursuing the same compute-bound opportunity on more mature timelines and without Neurophos's optical-specific conversion-overhead risk.

What drives the recommendation is valuation and system-level/competitive risk, in that order. At \$400–500M this would likely read INVESTIGATE FURTHER. At \$1.0B, the median outcome is capital back, the reverse-acquihire pathway flips from a win to a loss, and the tail that justifies venture risk is truncated (see the feasibility test above) — and is further truncated by the possibility that electronic competitors close the addressable prefill-offload opportunity before Neurophos reaches production.

WAIT, WITH A TOEHOLD PERMITTED IF AND ONLY IF IT SECURES CONTRACTUAL PRO-RATA INTO THE SERIES B PLUS INFORMATION RIGHTS.

The counter-argument deserves real weight (see Stress-Testing This Conclusion, Section 7): a zero-exposure strategy gets no option on the tail, and venture returns live in the tails. A small, option-like position with pro-rata captures most of that value while capping the price-error cost. That option's value erodes faster in a crowded, fast-moving field than it would in one with few players.

Position Sizing

Size band: none at \$1.0B on current evidence. If participation is required for relationship or information reasons, toehold only — the smallest position that secures information rights and pro-rata into the Series B. The pro-rata right is worth more than the position itself: it converts a bad entry price into an option on a well-informed entry at the next round — but that option is only valuable if Neurophos still has room to close the technical and competitive gap by the time it matures. A widening gap versus Etched-class competitors (CU-9) compresses the useful life of this option.

Sizing up to standard requires the first, second and fourth conviction builders, together with the two on system-level power and advantage over electronic accelerators. Any single kill criterion should trigger sizing down or exit — an insider-led Series B should trigger exit consideration rather than a defensive follow-on, since protecting a position through structure across 2–3 more financings is how a stall becomes a total loss.

Correlation flags — surfaced, not quantified

- Same-financing-environment dependency with the existing Gravitics holding: both depend on the availability of large private rounds for unproven physical technology through 2027–2030. Not diversified by holding both — doubled.
- Soft thematic overlap via Space Capital (a Neurophos investor) and the Eclipse Space relationship — worth tracking, not yet a hard correlation without more information on deal terms.
- AI-capex sentiment dependency: a contraction would compress every positive pathway simultaneously — the single biggest correlated macro risk across the whole table.
- Sector-sentiment dependency with the broader inference-ASIC financing wave (Etched's \$21B mark, Taalas, d-Matrix): a re-rating of that group, in either direction, likely moves Neurophos's achievable entry and exit multiples even if nothing about Neurophos itself changes.

7. STRESS-TESTING THIS CONCLUSION

- **Reference-class risk:** Echodyne (a small-TAM defense analog) informing a large-TAM AI-inference outcome is a judgment call, not a resolved fact — a reader weighting it differently would reasonably land on wider Pathway-3/4 ranges.
- **False precision:** the 0.55 retention factor and the pathway dollar ranges carry more apparent precision than the evidence supports — they are brackets, not point estimates.
- **Valuation-anchor risk:** the ~\$310M Series A post-money is a single low-tier source. At \$700M+, the step-up narrows to ~1.4x — the single highest-value cheap diligence item. The price objection survives even there: \$1.0B still exceeds Ayar Labs' mark on less maturity.
- **Capital-intensity risk:** Celestial needed ~\$600M, Lightmatter ~\$850M, on arguably easier problems. \$300–800M for Neurophos may be understated — this worsens CU-4 and every positive pathway, strengthening the case for WAIT.
- **Correlated-risk risk:** technical success, commercial pull, financing, and exit multiple are treated as loosely independent above — they are not; all sit downstream of AI-capex sentiment and the same demonstrator result. The true distribution is probably more barbelled: more failure and more breakout, less niche. This is the strongest argument against WAIT, since venture returns live in the tails — and it is why the recommendation includes a pro-rata toehold rather than a flat pass.
- **Wide-range risk:** the probability ranges in Section 4 are deliberately wide to avoid false precision. That is a defensible choice, but it has a failure mode of its own — wide ranges can be used to justify almost any prior conclusion, and a reader should treat the specific bounds chosen here as a considered judgment call, not a calibrated estimate. The more important analytical move is the naming of CU-9 (competitive crowding) as a distinct, first-order, trackable risk: that framing is falsifiable in a way the probability bands are not — it will be confirmed or refuted by whether Etched, Taalas or d-Matrix actually ship on their stated timelines.
- **Comparable-set risk:** Etched, Taalas and d-Matrix are themselves unproven at commercial scale. Etched has no independent (non-company) benchmark yet; Taalas does not ship until 2027; d-Matrix's actual deployed base is undisclosed. Treating them as a "more mature" comparable set is itself a judgment call — a reader could reasonably argue the electronic-ASIC space is just as speculative as the photonic one, simply further along a hype and financing cycle that has not yet been tested by real, independently verified deployment either. If that is right, CU-9 should carry somewhat less weight than this analysis gives it.

Taken together, these checks argue for retaining optionality rather than a flat pass — consistent with the toehold-plus-pro-rata recommendation above.

APPENDICES

Appendix A — Methodology Note

Scenario analysis was used rather than a single optimized valuation because Neurophos faces deep uncertainty, not statistical uncertainty. The 4-bit-precision question (CU-1) is not a parameter with a distribution — it is an unresolved physics-and-engineering question with a binary-ish character. Scenario pathways make the conditions for each outcome explicit and auditable.

A discounted cash flow model was not used: Neurophos is pre-revenue with first revenue no earlier than 2028, and no disclosed pricing, unit cost, or margin structure. A DCF would require inventing every input, and terminal value would dominate 80–95% of the result — producing false decimal precision rather than insight.

Precise probabilities were avoided for the same reason: no credible reference class exists. Roughly five companies have attempted this kind of full optical compute-in-memory approach, and roughly four more (Etched, Taalas, d-Matrix, Cerebras) are attempting the narrower compute-bound-inference problem electronically — none of either group at volume. Broad judgmental ranges anchored to named analogs, with the differences between Neurophos and those analogs stated explicitly, are the most defensible method available.

A Monte Carlo simulation was considered and set aside, on the same three-part test applied to the Gravitics analysis: is there at least one primary-source financial dataset to ground it (here, no — Neurophos has filed only a Form D, with no cap table, valuation, preferences, or burn disclosed anywhere); are there at least three input variables with empirically grounded distributions to draw from (no — there is no historical population of metasurface optical-compute companies to estimate a distribution from, and adding Etched, Taalas and d-Matrix as comparables does not change this — each is itself an $n=1$ unproven case, not a distribution to sample from); and would a simulation answer something the scenario tables cannot (no — the decision turns on price versus the median outcome and two binary-ish technical gates, not a path-dependent financing sequence). A simulation here would produce smooth confidence intervals describing sampling error in the analyst's own guesses, not real-world uncertainty — more rigorous-looking, but less honest.

For the same reasons, no Bayesian posterior probabilities are presented: no defensible prior or likelihood exists for a question like "probability an independent benchmark confirms the claim." The Forward Watchlist (Section 5) sets out how new evidence should move the assessment directionally instead, which is a more honest reflection of what the evidence actually supports.

It is useful to separate parameter uncertainty (where the probability bands should sit) from structural uncertainty (which comparison class applies at all, and whether the company's own framing of its addressable problem is correct). The product definition, system stress test and comparable-company material in Section 3, together with CU-9, address the structural layer. The probability bands in Section 4 remain judgmental, because no amount of public research manufactures a real reference class where none exists. Readers should weight Section 3 and CU-9 as the substantive analytical content, and treat Section 4's specific numbers with the scepticism the report itself applies to them.

Appendix B — Assumption Register

#	Assumption	Basis	If wrong
A1	Series A post-money \$310M–\$550M	One low-tier estimate plus dilution arithmetic on the disclosed \$110M round	If \$700M+, the step-up narrative weakens to ~1.4x; the price objection survives via the valuation cross-check
A2	Ownership retention ~0.55 (0.40–0.70)	2–3 further rounds plus pool expansion, typical for hardware reaching revenue in 2028	±0.15 swings Pathway-4 MOIC from ~4.8x to ~8.4x
A3	Additional capital required \$300–800M+	Celestial ~\$600M, Lightmatter ~\$850M, on easier problems, plus headcount trajectory	Likely understated per the stress-test; higher figures worsen A2 and all positive pathways
A4	Entry is a clean primary round, standard 1x non-participating preferred	Assumption in the absence of disclosed terms	A participating/ratchet/seniority structure would worsen Pathways 1–3
A5	No liquidity event before year 5	Roadmap (revenue 2028) plus sector exit timing	Early absorption at \$1–1.5B is a loss at this entry — Pathway 1b
A6	4-bit vs. 8-bit claim not reconciled publicly	Bowen's own July 2026 comments still call the numbers theoretical	If reconciled privately, CU-1 improves substantially — reopen the analysis
A7	Terakraft pilot unpaid or nominally paid	No disclosure either way	A materially paid pilot upgrades CU-5
A8	Strategic-buyer appetite persists through 2030–2033	Groq, Celestial, Enfabrica prints in 2025–2026	An AI-capex contraction compresses every positive pathway simultaneously
A9	Prefill token-share ~93%	Neurophos's own website only; no independent third-party source found	If actual production workloads split more evenly, or shift toward decode as context length and agentic patterns grow, the addressable-market thesis narrows, worsening all positive pathways
A10	Conversion and data-movement overhead severity applies to Neurophos at a comparable magnitude to published designs	Academic literature on other (non-Neurophos) photonic accelerators; applied here by analogy, not measurement	If Neurophos's specific architecture proves much more efficient at conversion than typical published designs, CU-1/CU-2 improve substantially — reopen the analysis
A11	Etched/Taalas/d-Matrix's maturity lead over Neurophos is real and durable	Public funding, contract and shipping-date disclosures as of September 2026	If electronic competitors stumble on their own scaling, accuracy or yield issues, the CU-9 concern eases, improving Neurophos's relative position without Neurophos doing anything differently

Appendix C — Diligence Requirements

Required before any commitment (all confidential)

- Full cap table, preference stack, participation, anti-dilution, actual Series A post-money.
- Burn rate, runway, board-approved plan to next financing.
- Test-chip characterization data: extinction ratio, insertion loss, array uniformity, thermal drift, ADC/DAC power share.
- Measured (not simulated) subsystem-level power breakdown: optical core / ADC-DAC / SRAM and vector units ("the boondoggle") / HBM / interconnect, on a real transformer workload at deployable precision.
- Accuracy modeling on real transformer inference at achievable bit-width, reconciling the 4-bit headline with 8-bit GPU comparisons.
- Independent, reproducible prefill throughput benchmark at deployable precision, directly comparable to published Etched/Taalas/d-Matrix figures where possible.
- The company's own internal data (if any) on production-workload prefill/decode token-share assumptions, and the investment thesis's sensitivity to that ratio.
- Named foundry/OSAT relationship and packaging cost curve.
- Core metasurface IP ownership documentation and freedom-to-operate opinion.
- Terms of the Eclipse Space relationship, now that its existence is confirmed — contract value, timeline, exclusivity.
- Export-control classification work and CFIUS analysis given the Aramco Ventures position.
- Board composition, governance rights, Gates Frontier's stated follow-on posture.

Reference calls

- Two independent photonics/analog-compute technologists, unaffiliated with the round, on precision and packaging specifically.
- Terakraft leadership, directly.
- A hyperscaler infrastructure buyer on actual deployment requirements for a photonic inference part.
- Someone with visibility into Etched's, Taalas's or d-Matrix's actual technical roadmap and timeline confidence, to calibrate "how far ahead is the electronic competition" independent of their own public messaging.

GLOSSARY

Term	Meaning
OPU	Optical Processing Unit — Neurophos's proposed chip category, intended to perform AI inference computation using light rather than electronic transistors.
Metasurface	An engineered surface with sub-wavelength structures that manipulate light; the core physical layer of Neurophos's optical modulators.
TOPS/W	Tera-operations per second per watt — a standard measure of computing efficiency; higher is better.
4-bit / 8-bit precision	The number of bits used to represent each number in a computation. Lower precision is faster and cheaper but can reduce accuracy; efficiency claims at different precisions are not directly comparable without reconciliation.
Prefill / decode	The two phases of large-language-model inference. Prefill processes the input prompt and is compute-bound (limited by arithmetic throughput); decode generates output tokens one at a time and is memory-bandwidth-bound. Neurophos targets prefill only.
HBM	High-Bandwidth Memory — stacked memory placed next to a processor to feed it data quickly; a significant contributor to a package's total power draw.
SRAM	Static Random-Access Memory — fast on-chip memory used to buffer data close to the compute units; the "vector units and SRAM" architecture Neurophos calls "the boondoggle" sits here.
ASIC	Application-Specific Integrated Circuit — a chip designed for one workload (here, transformer inference) rather than general-purpose computation; Etched, Taalas and d-Matrix build electronic ASICs for the same prefill bottleneck Neurophos targets.
MOIC	Multiple on invested capital — total value returned divided by capital invested.
IRR	Internal rate of return — the annualized return implied by a MOIC over a given holding period.
Post-money valuation	A company's valuation immediately after a funding round, including the new capital raised.
Step-up	The increase in valuation between one funding round and the next, expressed as a multiple.
Retention factor	The share of an investor's ownership that survives future dilution from later funding rounds and option pool expansions.
Non-participating preferred	A class of preferred stock that pays out its liquidation preference or converts to common (whichever is greater) but not both — more investor-friendly terms would include participation, ratchets, or seniority stacking.
Freedom to operate (FTO)	A legal assessment of whether a company can commercialize its technology without infringing others' patents.
Deemed export	Under U.S. export control law, the release of controlled technology or technical data to a foreign national, even within the U.S., which can itself require a license.
CFIUS	The Committee on Foreign Investment in the United States — a government body that reviews foreign investment for national security risk, relevant here given Aramco Ventures' participation.

Term	Meaning
BIS	The Bureau of Industry and Security — the U.S. agency administering export controls on sensitive technology.
Tapeout	The final step in chip design before manufacturing, when the design is sent to a foundry for fabrication.
NRE	Non-recurring engineering — a one-off payment for custom design or development work, as distinct from per-unit production revenue.
OSAT	Outsourced Semiconductor Assembly and Test — a third-party provider that packages and tests chips after fabrication.
Acqui-hire	An acquisition made primarily to obtain a company's team and technology, often at a price that does not represent a successful commercial outcome for earlier investors.
ADC / DAC	Analog-to-digital and digital-to-analog converters — components needed to interface an optical/analog computing core with conventional digital electronics, and a meaningful share of a system's total power draw.

ANNEX A — EVIDENCE LOG

Sources referenced in this analysis, grouped by category. SEC filings are the highest-reliability tier; private-market data providers are used only where flagged in the text and cross-checked against company disclosures.

Primary / regulatory sources

- SEC Form D exempt-offering notice (Neurophos, Inc.).
- U.S. Bureau of Industry and Security (BIS) guidance on the AI Diffusion Rule and the Affiliates Rule.

Company sources

- Neurophos funding announcements: seed round (December 2023) and Series A (22 January 2026).
- Terakraft (Norway) collaboration announcement, August 2025.
- CEO Patrick Bowen, interview, July 2026 (Eclipse Space confirmation; revised efficiency framing).
- Neurophos company website — prefill-specific positioning statements and T100 system specifications.

Industry / press and comparable-company sources

- Funding and valuation reporting on Ayar Labs, Lightmatter, and Celestial AI (including Marvell's December 2025 acquisition of Celestial AI).
- Reporting on Groq's asset sale to Nvidia and Enfabrica's acquisition, used as acqui-hire precedents.
- The Register, TechRadar and TechSpot coverage of the T100 announcement and CEO interview (January 2026) — prefill positioning, system specifications, "the boondoggle", production timing.
- Funding, valuation, customer-contract and shipping-timeline reporting on Etched (Sohu), Taalas, d-Matrix and Cerebras.

Academic and technical literature

- Published 2025–2026 academic papers on photonic AI accelerators covering ADC/DAC conversion, modulator-drive and data-movement overhead — general field literature, not Neurophos-specific, used as an outside-view stress test only.

Private-market data (lower evidence tier — used only where flagged and cross-checked)

- PremierAlts secondary-market valuation estimate (~\$310M) — not independently confirmed.
- Third-party employee-count databases — indicative only.