AMP PBC

The independent AI grid: pooled compute and capital as a public benefit

Explainer · June 5, 2026

A public benefit corporation pooling AI compute into an "independent AI grid" and pairing it with capital to widen access beyond the largest tech companies.

COMPANY	AMP PBC	**OWNERSHIP**	Private
CEO	Anjney Midha		
HEADQUARTERS	Menlo Park / San Francisco, California, USA (as reported; not confirmed by primary filing)		
TOTAL RAISED	More than $1.3 billion in committed capital (reported funding window 2025–2026, as a first venture fund; not a Series A equity round)		

ABSTRACT

AMP PBC is a public benefit corporation with a stated mission to "maximize the world's frontier output" by building an "independent AI grid": pooled, automated compute infrastructure spanning clouds, models, and data centers, paired with flexible capital so labs, startups, and universities can reach frontier-scale resources without owning hyperscale data centers themselves. Founded by ex-Andreessen Horowitz general partner Anjney Midha, the company has more than $1.3 billion in committed capital and operates through two reinforcing units: AMP Infra PBC, the pooled grid, and AMP Foundry, a capital arm. Its public-benefit charter includes a commitment of up to $500 million in profits through 2030 to a fund for communities affected by the AI transition. The structure is novel and the mission is sincere on paper, but the central tension is whether the "grid" is genuinely neutral infrastructure or compute-bundled venture capital wearing a public-benefit badge. What is verifiable is the capital, the founder lineage, and the framing; the legal entity map and the precise mechanics are not.

Keywords: AI compute; public benefit corporation; AI infrastructure; Anjney Midha; Andreessen Horowitz; Anthropic; sovereign capital; compute grid

1. Snapshot

AMP PBC (amppublic.com) is a public benefit corporation (PBC: a for-profit company with a legally stated social mission) whose stated aim is to "maximize the world's frontier output" by pooling AI compute into an "independent AI grid" and pairing that infrastructure with capital to widen access beyond the largest technology companies. It was founded by **Anjney Midha**, previously a General Partner at Andreessen Horowitz (a16z), who earlier ran an internal a16z AI-infrastructure initiative called "Oxygen" that pooled Nvidia chips for portfolio companies, the conceptual precedent for AMP. The company has more than **$1.3 billion in committed capital**, raised primarily as a first venture fund drawing on institutions, pension funds, and sovereign funds rather than as conventional corporate equity or a Series A round. Third-party profiles place it in California, around Menlo Park and San Francisco. Founding year is unsettled, and as a private entity, its revenue, the allocation of capital between compute and venture deployment, customer count, and exact corporate-entity structure are not publicly disclosed.

2. Thesis: Why This Company, Why Now

The bet is that frontier AI is bottlenecked by compute, and that access to that compute is concentrated in a handful of hyperscalers, leaving labs, startups, and universities rationed or priced out. AMP's thesis is that an *independent* layer, pooling capacity across clouds, models, and data centers and bundling it with capital, can widen who gets to build at the frontier. The "independent AI grid" framing is the pitch: shared infrastructure plus flexible capital, governed by a public-benefit mandate rather than a single hyperscaler's strategic interest.

The timing leans on the 2024–2026 compute crunch, when GPU scarcity made allocation, not invention, the gating factor for many AI teams. The addressable opportunity AMP gestures at is enormous, effectively all non-hyperscaler frontier compute demand. The reachable market is narrower and turns on a hard question: whether neutral pooling is structurally real, or whether "independent" is mostly positioning around a capital-and-compute bundle.

3. The Core Idea in Plain English

AMP wants to be a power utility for AI compute. Just as an electrical grid aggregates generation from many sources and delivers it to anyone who plugs in, rather than forcing each user to build a power plant, AMP aggregates GPU capacity across providers and routes it to members who could never stand up a hyperscale data center alone. The capital arm is the financing layer that helps members actually afford to plug in.

Old world: each lab or startup negotiates its own scarce, expensive cloud contracts, or builds its own infrastructure. New world: a pooled, automated grid plus bundled capital lets smaller players reach frontier-scale resources, with a public-benefit charter governing how profits are shared.

4. The Technical Space

The category is AI compute aggregation and orchestration, the layer that sits between raw data-center capacity and the teams training and serving models. The underlying problem is matching: supply (GPUs across clouds, neoclouds, and dedicated facilities) is fragmented, lumpy, and unevenly priced, while demand (training runs, inference fleets) is bursty and capital-intensive. Standard approaches fall into a few buckets.

- **Hyperscaler reservations.** Long-term committed contracts with AWS, Azure, GCP, or Oracle, which give scale but lock buyers to one provider's pricing and availability.
- **Neoclouds and brokers.** Players like CoreWeave or Lambda that resell GPU capacity, optimizing for availability and price but without a capital layer or a neutrality mandate.
- **Capacity marketplaces.** Spot-style platforms like Vast.ai that match idle GPUs to buyers, strong on price discovery but weak on reliability and large reserved runs.

On the dimensions that matter, "good" means reliable access to large contiguous blocks of accelerators, competitive effective cost per GPU-hour, low switching friction across providers, and, for AMP specifically, governance credibility, whether "independent" is enforceable rather than asserted. Capital availability is a fourth axis most pure-infrastructure plays don't address.

5. How Their Technology Works (and What's Proprietary)

What AMP actually offers, per its own materials, decomposes into two reinforcing units rather than a single product. The brief's caution applies here: this is a *functional* two-unit model, not a settled legal-entity map.

- **AMP Infra PBC, the AI Grid.** This is the pooled, automated infrastructure spanning clouds, models, and data centers. The company describes it as orchestration and pooling of compute that members can access, the technical substance being an aggregation and routing layer over heterogeneous capacity sources. How that compute is actually procured, whether reserved blocks, long-term leases, or opportunistic purchases, is not documented in available materials.
- **AMP Foundry, the capital arm.** Described on the site as wholly owned by AMP Infra PBC, it provides flexible capital to help grid members access compute. Functionally, this is the financing that converts pooled supply into something smaller members can afford.

The hard question is what is genuinely proprietary here. Compute aggregation and orchestration are not novel as engineering: brokers and neoclouds already do the matching, and a well-funded incumbent or a hyperscaler could replicate a pooling layer. The orchestration software, if it exists as differentiated IP, has not been demonstrated publicly, so it should be treated as a claim, not a verified asset. What is harder to copy is not the technology but the *structure*: the public-benefit charter and the bundling of capital with compute. Because the procurement mechanics are undocumented, the technical defensibility of the grid itself remains unproven. Read the differentiation as financial and institutional first, technical second.

6. Business and Go-to-Market

The commercial model is the distinctive part. AMP raised more than $1.3 billion primarily as a first venture fund, with LP capital from institutions, pension funds, and sovereign funds, rather than as startup equity. That structure suggests returns flow through compute-enabled deployment and equity positions in the teams it backs, not only through margin on resold compute. Allocation between compute acquisition and venture deployment is unspecified.

The go-to-market motion runs through membership. AMP names a set of "founding grid partners" spanning research labs, capital institutions, and compute providers, including Periodic Labs, ElevenLabs, Mistral, Sesame, Arena, Black Forest Labs, a16z, Y Combinator, NEA, and StepStone. Note carefully: the company's own site frames these as grid partners, not investors, so they should not be read as a cap table.

AMP has also taken a strategic investor position in Anthropic, a signal of how the grid-plus-capital model is meant to compound: backing a frontier lab while supplying it infrastructure. The size of that stake and AMP's exact role are not consistently reported, and AMP's own site does not mention Anthropic at all, so treat the relationship as confirmed in existence but unspecified in scale. The public-benefit charter adds a commitment of up to $500 million of profits through 2030 to an "AMP Public Wealth Fund" for communities affected by the AI transition. Unit economics, the gross margin on pooled compute versus venture upside, are not disclosed and are non-trivial given GPU cost structures.

7. Competitive Landscape and Moats

AMP sits at an unusual intersection of GPU broker, infrastructure fund, and public-benefit vehicle, which makes its comp set hybrid.

Closest direct rival: CoreWeave. CoreWeave is the sharpest head-on comparison as a dedicated AI-compute provider aggregating and reselling GPU capacity at scale. Where AMP wins: it pairs compute with flexible capital and a neutrality-and-public-benefit charter that a commercial neocloud cannot credibly claim, and it carries strategic equity exposure to the labs it serves. Where AMP loses: CoreWeave has operating data centers, proven reliability, large reserved capacity, and real revenue, while AMP's grid is so far more framing than demonstrated infrastructure.

The adjacent set rounds out the map: hyperscalers (AWS, Azure, GCP, Oracle) compete on scale and integrated reservations; Vast.ai competes at the capacity-marketplace layer on spot price; Together AI competes one layer up at model-services, narrowing AMP's wedge for buyers who want hosted models rather than raw compute; and a16z's own Oxygen-style pooling is the conceptual precedent AMP descends from.

Capital is the most real moat. More than $1.3 billion committed, with sovereign and pension LPs, is genuinely hard to assemble and buys both compute and strategic deals.

The public-benefit charter is a positioning moat, real as differentiation but only durable if governance is enforceable; otherwise it is marketing.

Network effects are asserted, not proven. A grid gets more valuable as partners and capacity join, but no evidence yet shows that flywheel turning. Platform risk is acute: a hyperscaler or a16z itself could replicate the pooling layer.

8. Risks and Open Questions

The picture turns on whether "independent grid" is structure or slogan, and on facts that are not yet public. Five things would most change the read:

- **Neutrality risk.** Is the grid genuinely neutral infrastructure, or compute-bundled venture capital where access, capital, and strategic proximity arrive in one package, steering members toward AMP-aligned outcomes?
- **Technical defensibility.** What, if anything, is proprietary in the orchestration layer, versus a financing wrapper around capacity any broker can source?
- **Compute-sourcing mechanics.** How does AMP actually procure capacity, reserved blocks, long-term leases, or opportunistic buys, and what are the resulting unit economics?
- **AI-capex cyclicality.** The thesis rides the GPU scarcity wave; if compute loosens or AI capex contracts, both the urgency and the resale spread compress.
- **Governance of the charter.** What is the vehicle, board, and criteria for the $500 million AMP Public Wealth Fund, and what legally binds the public-benefit commitments?

I would also press on the corporate entity map: AMP PBC, AMP Infra PBC, and AMP Foundry are described functionally, but the precise legal structure is not established in available evidence. And I would want the size and role of the Anthropic stake reconciled, since AMP's own materials are silent on it.

9. Bottom Line

AMP is a genuinely novel attempt to fuse compute aggregation, institutional capital, and a public-benefit mandate into an "independent AI grid," led by a founder with directly relevant lineage from a16z's internal compute-pooling work. It works or fails on one question: whether the grid is structurally neutral infrastructure or simply venture capital with compute attached, which determines whether the moat is durable or just well-funded positioning. Watch next for hard evidence of operating capacity and the legal-and-governance substance behind the charter, the moment the framing has to become infrastructure.

10. For the Nerds

The interesting engineering question is whether AMP can deliver large *contiguous, reserved* accelerator blocks across heterogeneous providers, not just spot-match idle GPUs. Frontier training needs tightly interconnected fabric, high-bandwidth interconnect like InfiniBand or NVLink domains, and topology-aware scheduling. Pooling across multiple clouds and data centers fights physics: you cannot trivially stitch a single coherent training cluster across providers that don't share low-latency networking. The real risk is that the "grid" resolves in practice to a smaller number of homogeneous clusters that happen to be owned or leased by AMP rather than a hyperscaler, in which case the technical differentiation is thin and the moat lives entirely in the network and capital structure. That tension means AMP's grid is likely most useful for inference, fine-tuning, and elastic burst workloads, and least useful for the largest pretraining runs, where locality dominates.

A second open question is financial-structural: bundling capital with compute creates an implicit cost-of-capital advantage for members, but also a potential conflict, where the "independent" allocator is also an equity holder in the demand side. The neutrality claim is, ultimately, a mechanism-design problem as much as an engineering one.

WHERE SOURCES DISAGREE

Is the $1.3B a "Series A" or a "first fund"?

One funding-tracker source labels it a "$1.3B Series A," implying startup equity. Narrative reporting and the company's own positioning instead describe capital raised for a first venture fund from institutional, pension, and sovereign LPs

Unresolved: no primary filing is available and the only "Series A" label comes from a secondary aggregator that conflicts with higher-quality coverage

When was AMP founded?

Press coverage and profiles imply a 2026 founding tied to the funding announcement, while at least one account points to the AMP concept being announced earlier and the PBC launching before 2026

Unresolved: no incorporation document or primary disclosure states a founding year; "2026" is an inference from financing news

How large is the Anthropic investment, and what was AMP's role?

One report specifies a $300 million investment in Anthropic's roughly $30 billion round in February 2026; other reports describe AMP only as a "significant" or "co-leading" investor without a figure. The company's own site does not mention Anthropic at all

Unresolved: the specific size and role are not consistently sourced and there is no primary confirmation from AMP or Anthropic

Is the "independent grid" genuinely neutral, or a compute-bundled venture vehicle?

The company frames itself as a neutral, independent grid widening access to compute. Some reporting suggests it may function less as a neutral grid and more as bundled venture capital, where compute access, capital, and strategic proximity arrive in the same package

Unresolved: the characterization turns on interpretation of the same facts, and the sources actively disagree on whether "independent" is structural or marketing

Are "AMP PBC," "AMP Infra PBC," and "AMP Foundry" distinct legal entities?

The company's site describes a functional hierarchy in which AMP Foundry is wholly owned by AMP Infra PBC, but uses product/operational language. Other materials and databases treat them as separable units and reference a named fund

Unresolved: no charter, PBC filing, or fund/GP record in the available evidence establishes the precise legal entity map

How does AMP actually source compute?

Coverage describes AMP acquiring data-center capacity to redistribute to startups, universities, and labs, framed in some accounts as buying "excess" capacity. The company's own materials describe pooling and orchestration without specifying spot purchases, long-term leases, or reserved capacity

Unresolved: the financial structure of how compute is procured is not documented in any examined source