

Capacity Without Capex

Where the revenue actually comes from when AEP frees a megawatt — by business model

The honest claim

AEP does not create revenue from nothing — it creates **verified megawatts** from nothing: power that is already built, financed, cooled and contracted, freed from worst-case provisioning and enforced at the GPU driver, with evidence a lender, tenant or valuer can check independently. What that megawatt then earns — and who spends what to earn it — depends on the business model. The answer is better than it first looks in all three.

Model 1 — Colocation: no new capex at all

A colo operator's product is the contracted megawatt; **tenants bring the GPUs**, as they always have. Freed envelope converts to revenue two ways, neither of which touches the operator's capex line:

- **Verified facility-level oversubscription.** Today every tenant's block reserves bus capacity for nameplate worst case, so the facility sells fewer megawatts than it built. With driver-enforced envelopes provably holding aggregate draw inside the facility rating — sealed, continuously re-verified — the operator can contract more megawatts against the same build. More sellable product from assets already on the books.
- **Tenant densification.** Tenants deploy more IT inside their existing contracted blocks — tenant capex, tenant upside, faster absorption and stickier tenancy for the operator.

The valuation angle (owners and funders): when the forward yield curve is marking down DC valuations despite record contracting, capacity that raises near-term cash flow with **zero new facility debt** is the one growth lever the rate curve cannot discount. Illustratively, at colo lease rates of A\$3–5M per MW per year and infrastructure multiples, each verified freed megawatt represents tens of millions in enterprise value — arriving this forecast period, not after a construction programme.

Model 2 — GPU-as-a-Service: growth at GPU-only economics

A GPUaaS operator does need GPUs to monetise freed envelope — but look at **which half of the capex stack** that is. All-in AI-DC capacity costs US\$30–40M per MW, splitting roughly **56% servers / 30% facility / 13% network** (Epoch AI, 2026). Deploying into freed envelope means:

- **Skip the facility-and-grid tranche entirely** — US\$10–15M per MW of building, cooling and substation that is already sunk, plus the 5.1-year median US interconnection queue (LBNL 2026), plus markets like Singapore where new grid is unavailable at any price.
- **Spend only the financeable half.** GPU capex is short-duration, revenue-backed and routinely debt-financed against contracted demand — the least rate-sensitive way to add cash flow, and the lending improves further when the capacity underneath it is independently verified.
- In power-capped markets the comparison isn't "capex vs no capex" — it is "**grow at GPU-only economics vs don't grow at all.**"

Two paths that are genuinely zero-capex (both models)

- **Run installed GPUs harder.** Fleets constrained by facility power budgets throttle hardware they already own; managed envelopes let the existing estate deliver more throughput inside the same feed.
- **Defer the next build.** The coming expansion phase fits into freed envelope instead of new construction — capex pushed out, not spent.

We don't create revenue from nothing — we create verified megawatts from nothing. The revenue then costs only GPUs, not buildings, substations or a place in the grid queue. And for a colo, not even the GPUs.

Measured basis: 9.6% of provisioned envelope freed on a hard-run NVIDIA fleet at zero measured throughput cost (driver-exact, 2026-09-11); ~31% on a lighter fleet (2026-08-14). Verification: cryptographically sealed, tamper-evident chain, checkable without trusting the operator or Autropic.

Tom Pickett · Founder, Autropic · tom@autropic.com · Gold Coast, Australia