
EPISODE
436
THE FINANCIAL MESS BEHIND AI BOOM'S PROBLEM CHILDREN
CoreWeave
NEXTDC
HOSTED BY
DAMIEN KLASSEN
Nucleus Wealth

Disclaimer:

The information provided on this presentation is general in nature and does not constitute personal financial advice. The information has been prepared without taking into account your personal objectives, financial situation or needs. Because of this, you should consider the appropriateness of the information for your own objectives, financial situation and needs before acting on it.

Also, before you decide to invest in a financial product arranged by a representative of Nucleus Wealth Management Pty Ltd, ABN 54 614 386 266, corporate authorised representative of Nucleus Advice Pty Ltd AFSL 515796 (Nucleus Wealth or we or us), it is important that you read and consider the Product Disclosure Statement relating to the product before making any decision about whether to invest in it.

Where NEXTDC and Firmus sit in AI

The same AI workload pays several businesses along the chain

Layer	Example	What the customer buys
Powered facilities	NEXTDC	Space, power, cooling and connectivity
Equipped compute	Firmus	GPU clusters and cloud capacity
Frontier models	Anthropic / OpenAI	Model access and inference
Applications and users	IT and other businesses	Useful work and productivity

Revenue per GW differs at each layer. Customer payments and the compute bills they fund overlap.

NEXTDC sells powered space

Core exposure: facilities and contracted capacity. FY26 company summary.

175 MW

billing utilisation

740.1 MW

contracted utilisation

The operating question is conversion to billing

Buildings and power commitments precede cash receipts.

FY26 reported metrics include pro forma contracted capacity. The M3 delay is a company-specific execution event.

Firmus operates the compute factory

Company announcements, September 2026. Contracted capacity spans several sites.

900+ MW

contracted across all customers

2 operating sites and 5 in development

OpenAI: two Malaysian sites
Meta: contracted capacity plus expansion options

Contracted MW, capacity options and operating MW measure different stages. Firmus reports the figures above.

Different assets create different risks

A facility operator and a GPU operator capture different parts of the margin

NEXTDC

Longer-lived facilities. Returns depend on power delivery, occupancy, contracts and financing.

Firmus

Equipped compute adds GPU refresh, utilisation and compute-pricing exposure to the facility economics.

The later $50bn/GW installed-capital scenario includes equipment. It does not describe bare colocation capex.

Datacentre bans and the buildout

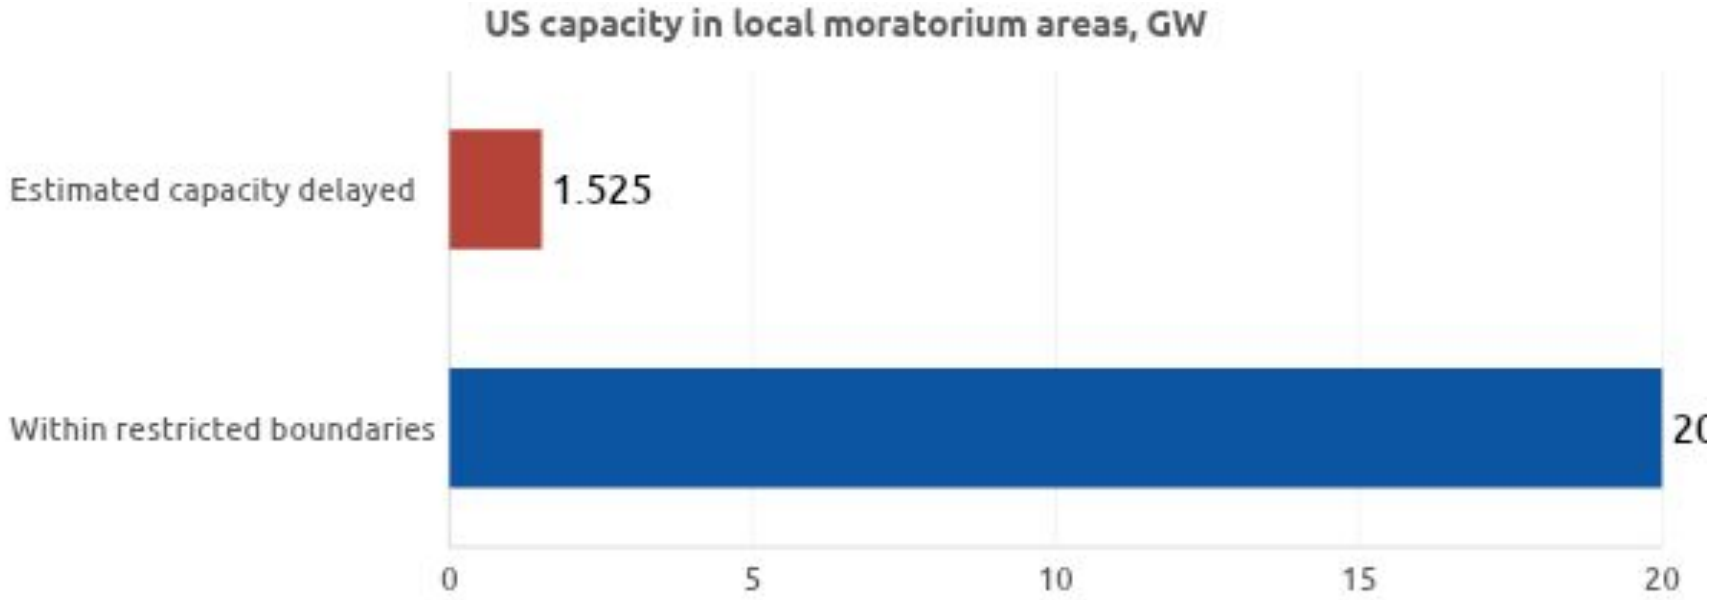
Damien’s view: much of the political response is performative

Policy action	Possible effect on supply
Ofgem: fees and readiness milestones	Clear speculative grid applications
Seattle: temporary siting freeze	Restrict new large projects locally
Existing approvals and relocation	Some capacity proceeds elsewhere

The relevant test is the change in credible MW delivered across regions, rather than the number of bans announced.

Exposure to a ban exceeds measured delay

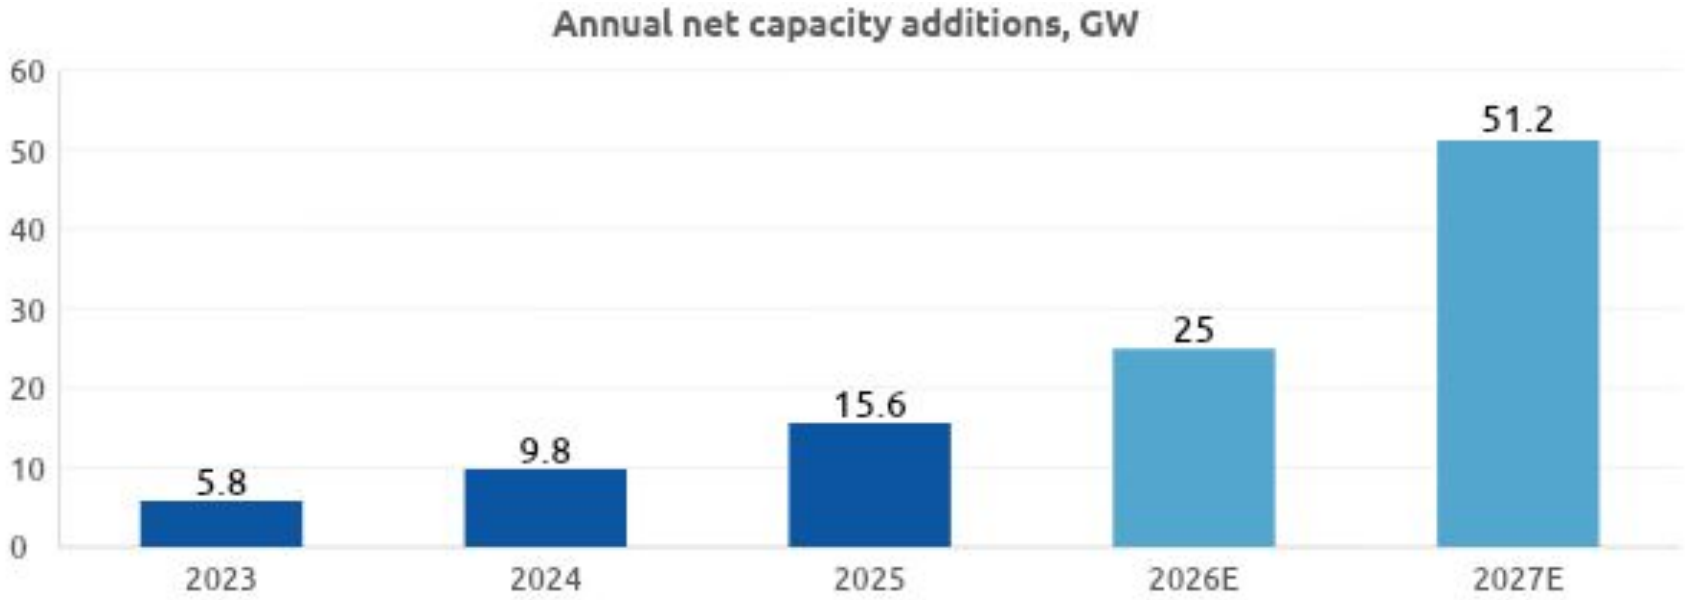
SemiAnalysis project-level audit, 15 September 2026. Local moratoriums.



About 7.6% of local exposure. The separate local-plus-New York estimate is 2.3 GW, not a complete tally of every state restriction.

The credible construction pipeline remains large

Supplied SemiAnalysis forecast, October 2026. Global ex-China, all IT workloads.



2027: 39.4 GW operational / under construction and 11.9 GW planned. Source components differ from the headline by 0.1 GW.

Anthropic’s proposed float

Company announcement and preliminary filing, 26 October 2026

Item	Disclosed position
Filing	Public Form S-1
Proposed listing	Nasdaq Global Select Market, ANT
Price range and offering size	Undetermined in the announcement
Investor question	Can receipts fund the next compute ramp?

The filing describes a proposed offering. The announcement does not establish an IPO price or valuation.

The model lab’s operating economics

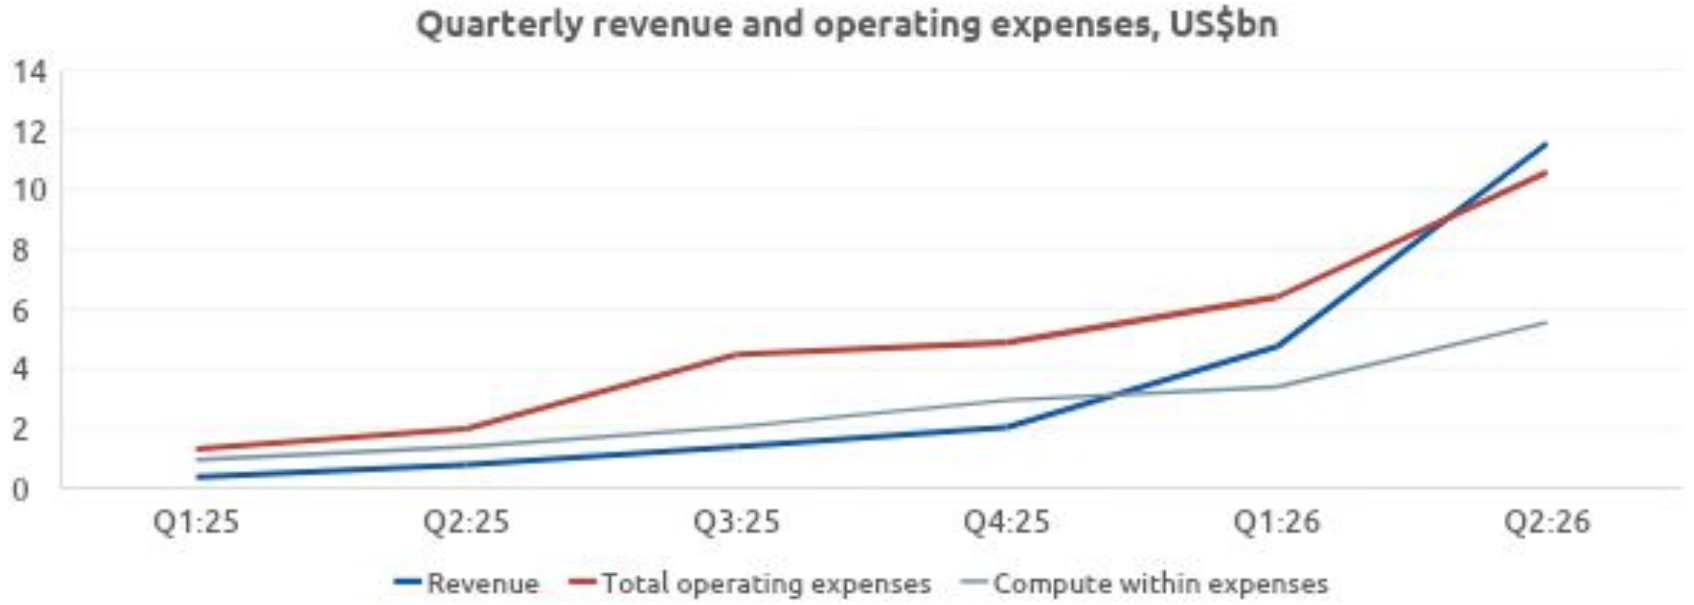
Research spending and customer serving respond to different drivers

Cost or receipt	Economic driver
Training and research	Model ambition and release cycle
Researchers and company costs	Talent, products and organisation
Paid inference	Usage and compute required per request
Customer receipts	Subscriptions and consumption

Research is recurring. Reserved compute makes part of serving cost fixed in the near term, even when usage changes.

Anthropic’s revenue has overtaken quarterly costs

Q2:26 operating profit $0.97bn. H1:26 operating loss $0.69bn.



Reported actuals. Compute includes research and serving. Operating profit differs from net income and cash flow.

Serving margins leave room for research

Anthropic H1:26. Cost per US$1bn of recognised revenue.

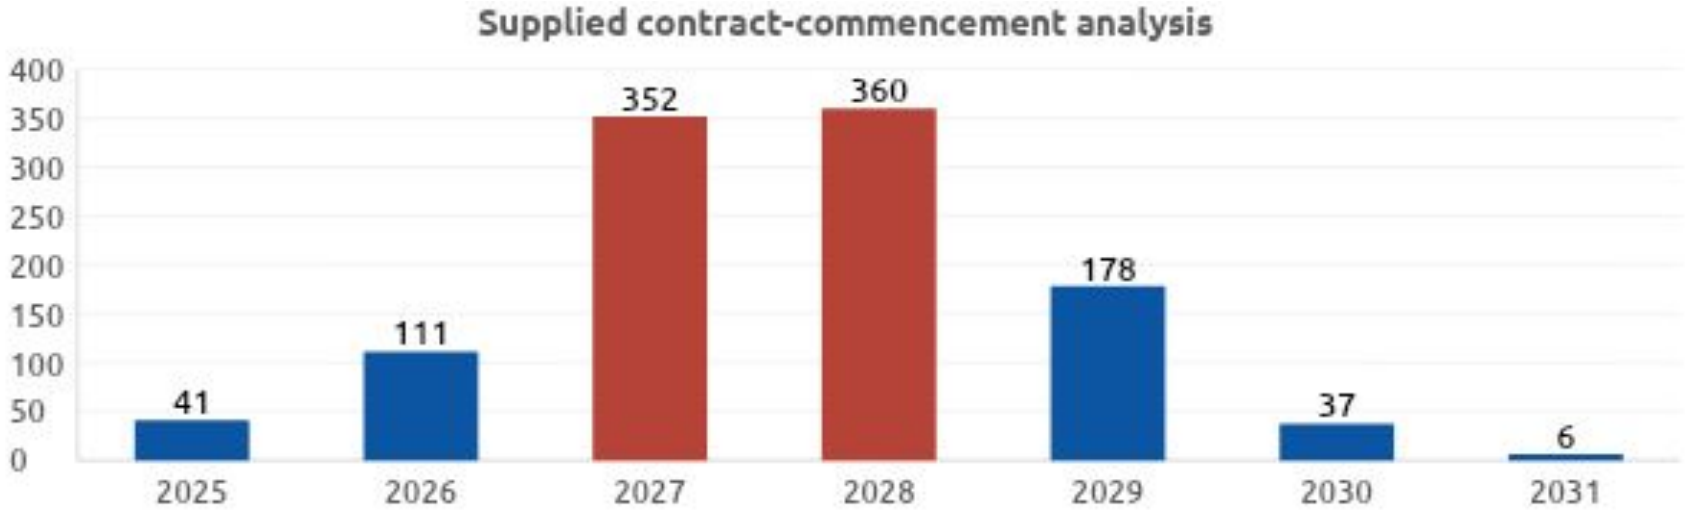
Expense scope	Cost, US$bn	Margin
Customer serving	0.146	85.4%
Serving plus cloud channel fees	0.312–0.319	68–69%
All operating expenses	1.043	(4.3%)

Research and development: $9.74bn in H1:26

Channel fees are an inferred range. The serving proxy includes some overhead. These are average accounting ratios.

The 2027–28 compute commencement wall

OpenAI + Anthropic. Multi-year contract notional at commencement, US$bn.



Anthropic: $161bn of five-year Broadcom equipment leases

The two vintages total $712bn of multi-year notional. Payment timing depends on delivery and contract terms.

Four sources of demand for compute

Funding sources and end customers overlap along the chain

Group	Main spending driver
OpenAI, Anthropic, SpaceX/xAI and smaller labs	Customer receipts plus new capital
Google, Meta, Microsoft and AWS	Defend and extend existing businesses
IT sector	Productivity and agent adoption
Other businesses	Adoption speed and demonstrated ROI

Consumer payments and advertising sit inside provider receipts. They do not create an additional amount to add again.

Damien’s view: the cost of losing the core business can dwarf an AI budget

The Kodak risk

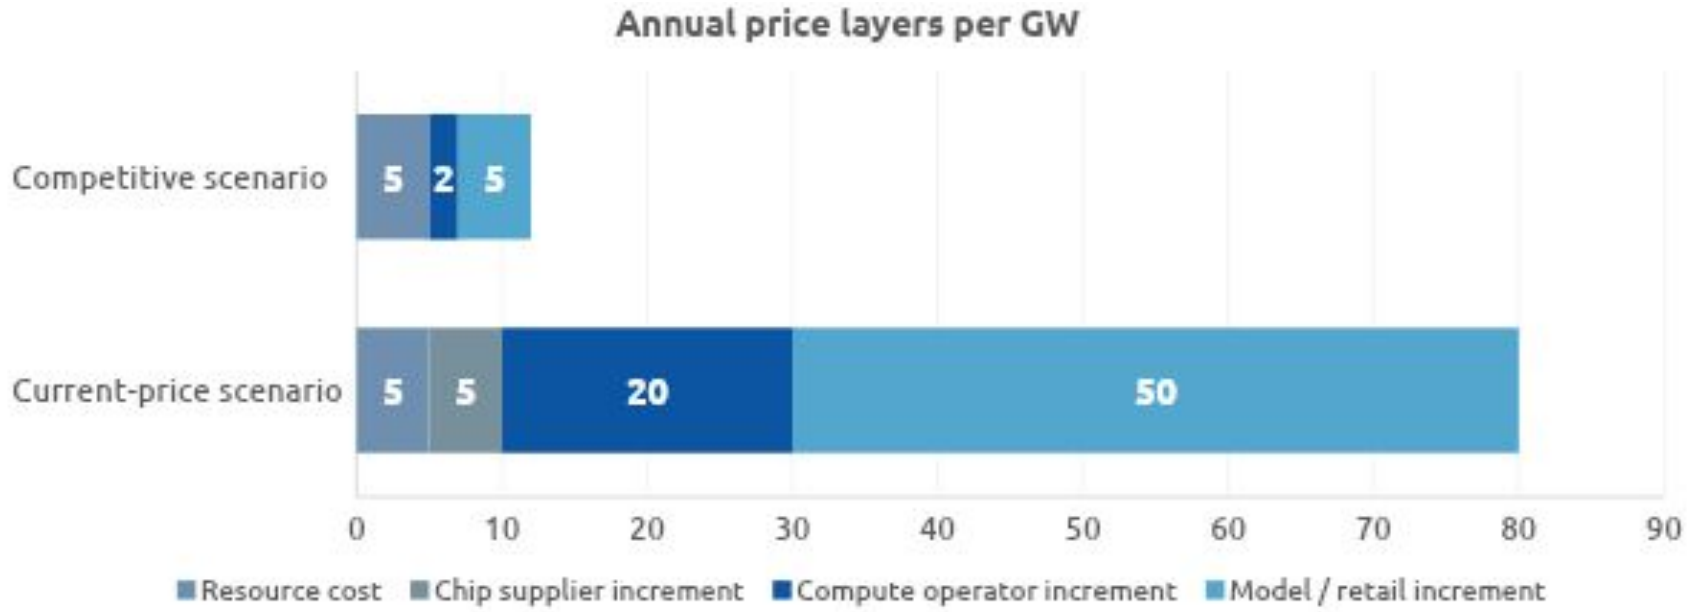
Search, advertising, software and cloud businesses have existing profit pools to protect.

Defensive spending can support compute demand before a standalone AI product earns its return.

The model’s matching-spend assumption is a scenario. Strategic pressure does not establish a disclosed spending floor.

Scarcity pricing across the chain

Hypothetical annual layers per GW. All amounts in US$bn.



User scenarios: $80bn versus $12bn of retail revenue/GW/year. The increments include services and normal returns.

Owned capacity, rented compute and inference

Each buyer faces a different annual cost or price per GW

Annual basis, US$bn/GW	Current	Competitive
Owner equipment use + operations	10	5
Wholesale compute rental	30	7
Retail inference receipts	80	12

The main case uses owner cost for incumbents, wholesale rental for labs, and retail receipts for IT customers.

The 2027 capacity hurdle

Increase in average operating AI capacity versus 2026, global ex-China

Contribution	Calculation	Average GW
2026 cohort: fuller availability	25 × 70% × 50%	8.75
2027 cohort: commissioning during year	51.2 × 70% × 50%	17.92
Additional capacity available in 2027		**26.67**

22.54 GW using the 39.4 GW construction component

70% AI share and even commissioning are assumptions. The supplied 51.2 GW forecast covers all workloads.

The annual compute budget scenario

Annual cost envelopes, US$bn. Anthropic and incumbent budgets are scenarios.

Compute envelope	2026	2027	Increase
OpenAI cloud services	31.0	112.0	81.0
Anthropic compute costs	30.0	90.0	60.0
Other frontier labs	12.2	40.4	28.2
Frontier total	73.2	242.4	169.2
Incumbent own use	**73.2**	**242.4**	**169.2**

OpenAI direct silicon excluded, roughly $6bn/$20bn. Anthropic $30bn/$90bn and matching incumbent budgets are scenarios.

Training, free inference and paid inference

Illustrative allocation of the annual budget increase, US$bn

Compute use	Lab share	Labs	Incumbents
Training and research	70%	118.4	118.4
Free inference	10%	16.9	n.a.
Paid customer inference	20%	33.8	n.a.
Internal / ad-funded inference	n.a.	n.a.	50.8

Incumbents: 70% research, 30% own inference. These are workload assumptions, not disclosed 2027 splits.

The IT customer spending estimate

Gross developer AI-tool spending. Scenario inputs, not a measured market total.

Driver	2026	2027
Professional developers, millions	30.0	31.5
Paid adoption	60%	80%
Monthly spend per paying developer	$100	$200
Annual customer spending, US$bn	**21.60**	**60.48**

The increase is $38.88bn. This is a customer view of revenue already flowing through AI providers.

Customer receipts and lab compute bills

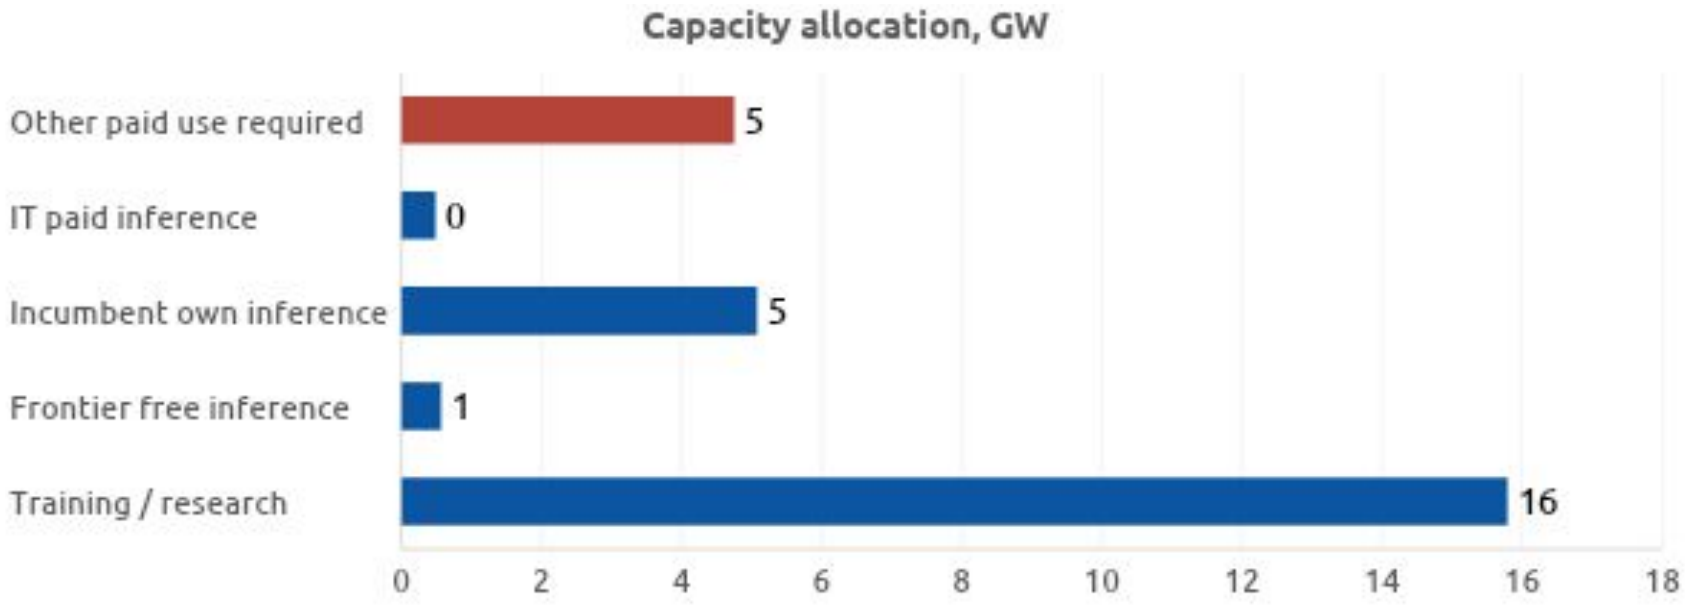
Incremental annual paid inference at the current price assumptions

View	US$bn
Lab paid-serving compute bill increase	33.84
Retail receipts implied at $30 / $80 per GW	90.24
IT receipts within that customer total	38.88
Other receipts implied within that total	51.36

The $90.24bn funds the $33.84bn compute bill. They represent the same serving workloads.

Demand at the $80bn retail endpoint

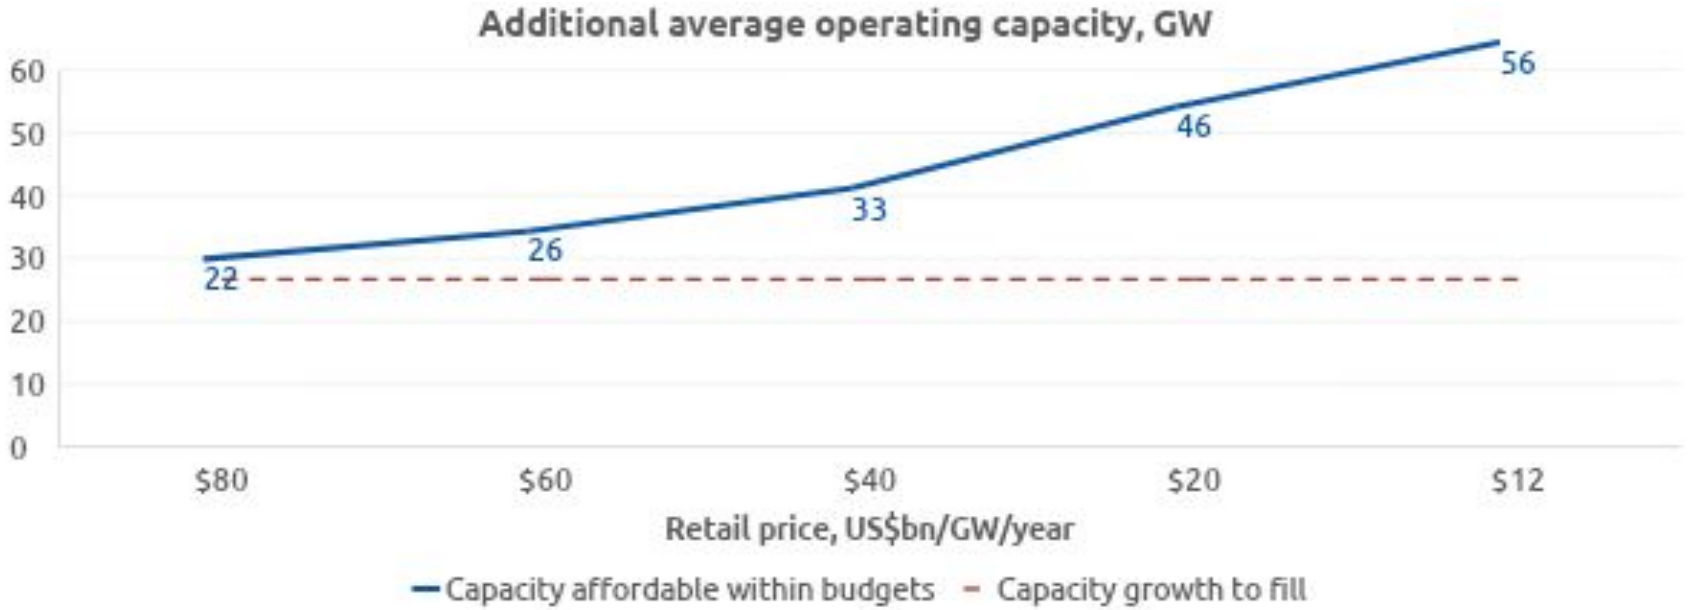
Additional annual-average GW versus 2026. Incumbents use the owner cost basis.



21.918 GW mapped + 4.752 GW other paid use = 26.67 GW. At $80bn/GW, the remaining customer bill is $380.16bn.

Lower prices increase budget purchasing power

Fixed incremental budgets. All savings fund extra workloads.



The curve shows potential purchasing power. Financing, actual workloads and contract repricing determine realised demand.

The remaining paid inference requirement

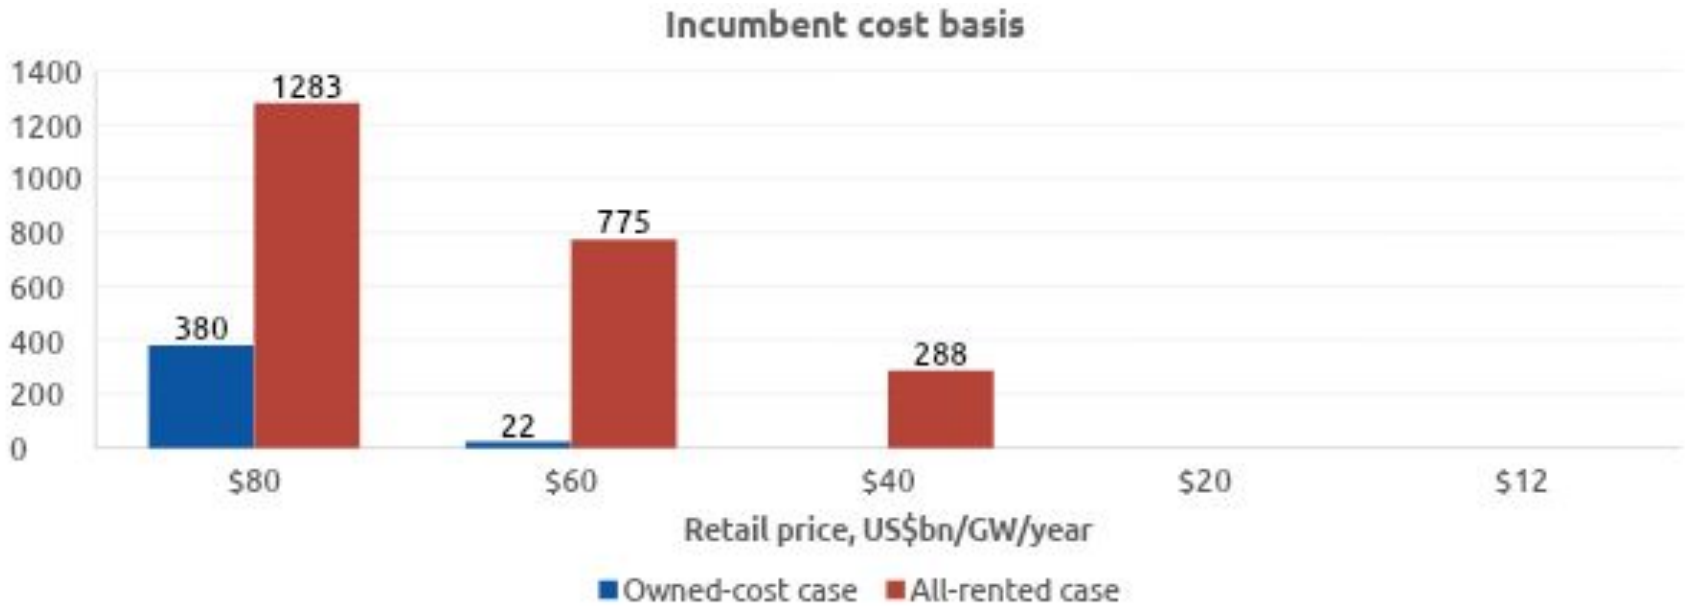
Incremental annual demand versus 2026. Fixed budgets and full reinvestment.

Retail $bn/GW	Research GW	Mapped inference GW	Other receipts US$bn
80	15.8	6.1	380
60	19.0	7.3	22
40	24.0	9.2	0
20	33.4	12.8	0
12	40.6	15.8	0

Zero means budget purchasing power reaches the 26.67 GW hurdle. It does not establish paying customer commitments.

Ownership changes the revenue hurdle

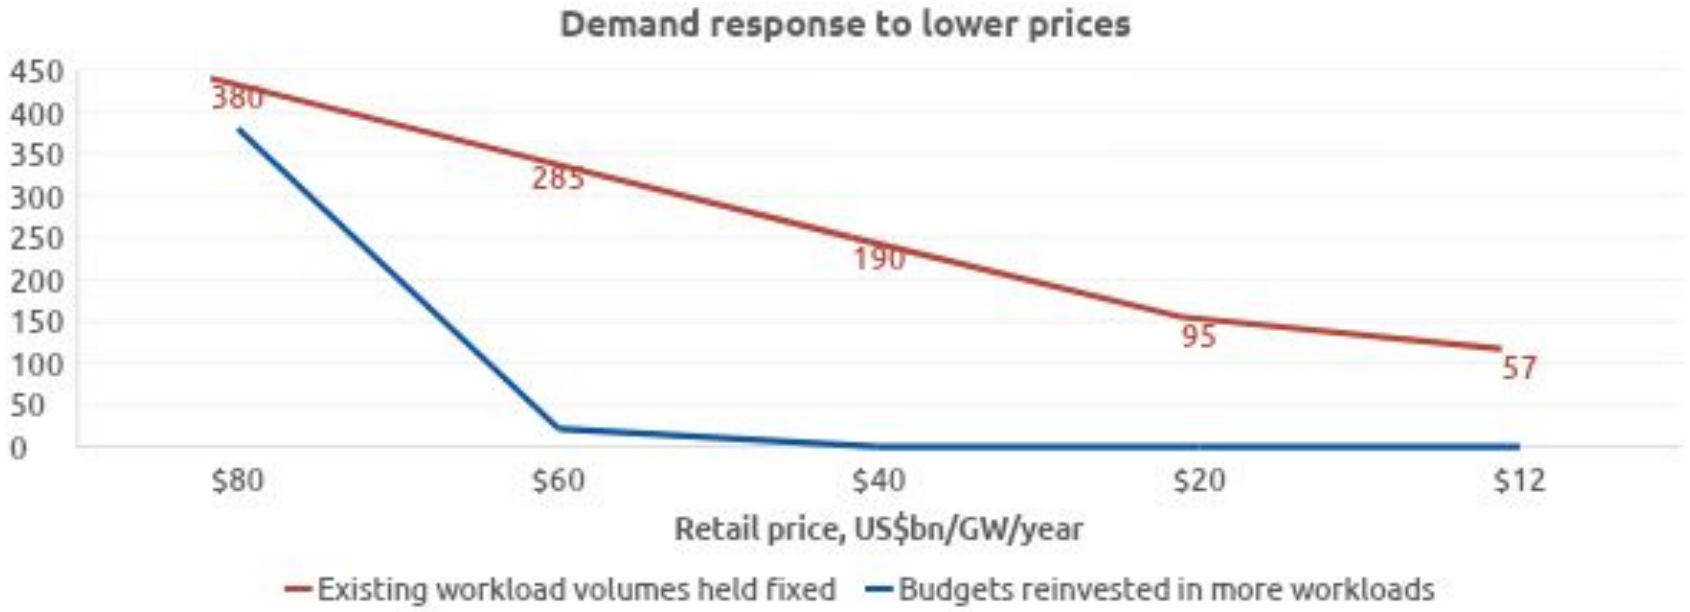
Other paid-inference receipts required, US$bn/year



At retail $80bn/GW, the hurdle is $380bn with owned compute versus $1,283bn if incumbents rent all compute.

Lower bills and more workloads have different effects

Other paid-inference revenue required, US$bn/year. Owned-incumbent case.



With fixed workloads, the gap stays at 4.752 GW. Its bill falls from $380bn to $57bn as retail pricing reaches $12bn/GW.

How much of the additional usage arrives in 2027?

At the $12bn retail endpoint, with partial use of the extra purchasing power

Extra workloads realised	0%	25%	100%
Capacity covered, average GW	21.9	30.5	56.4
Other receipts required, US$bn	57	0	0

Contracts and adoption determine the pace

The 25% illustration is a sensitivity. Anthropic procures much of its multi-year compute 12–18 months before availability.

More compute can coexist with lower investment returns

New cheap capacity can support more usage.
Existing assets retain debt and contracted costs.

The key variables are realised usage, annual cost per GW, and the returns retained at each layer.

Assumptions behind the 2027 calculation

The scenarios isolate scarcity pricing. Several inputs remain uncertain.

Assumption	Effect on interpretation
70% AI share and even commissioning	Changes capacity growth to fill
Anthropic budget and incumbent match	Sets annual compute costs
Training / inference mix and ownership	Changes GW supported per dollar
OpenAI owned silicon excluded	Owned capacity remains unestimated
Constant productivity and utilisation	Holds useful work per GW fixed
Developer receipts and geography	Customer location remains uncertain

The residual is a revenue hurdle for other paid use. It could also be filled by additional subsidised workloads.

More from Nucleus Wealth

Find out what your personalised investment could look like in 10-20 minutes – commitment free & at no cost.

Find us on all major podcast platforms:

Spotify

Listen on Google Podcasts

Listen on Apple Podcasts

Nucleus Wealth on social media:

Facebook.com/nucleuswealth

Linkedin.com/company/nucleuswealth

Instagram.com/nucleus_wealth

Twitter.com/nucleuswealth

Read insightful blog articles:

nucleuswealth.com/blog

Have a guest or topic suggestion for the show? Leave a comment down below.

Credit markets are asking the question

~15%

secondary-market bond yield

<$80

bond price per $100 of face value

Distress pricing puts the funding model under scrutiny

Financing problem or industry problem?

The same credit sell-off can support two very different interpretations

1 The projects work; the balance sheet is stretched

Attractive operating economics can coexist with excessive borrowing and an unfunded expansion programme.

2 The projects do not earn their cost of capital

If realistic pricing, utilisation and replacement costs fail the test, the problem reaches beyond one operator.

The compute refinery

A useful economic analogy for a capital-intensive operator

Electricity → GPU clusters → Compute

The margin depends on the price of compute, operating efficiency and the cost of the equipment.

A capable operator still faces a capital cycle

Two asset lives, two capital commitments

Illustrative facility assumptions in the Economic Model

Commitment	Model horizon	Economic implication
Tenant infrastructure	15 years	Spread the fit-out over a longer life
GPU equipment	5 years	Replace the equipment repeatedly
Customer contracts	Often multiple years	Reduce revenue uncertainty during the term

A long lease secures capacity and creates a fixed obligation. Replacement life is a modelling assumption.

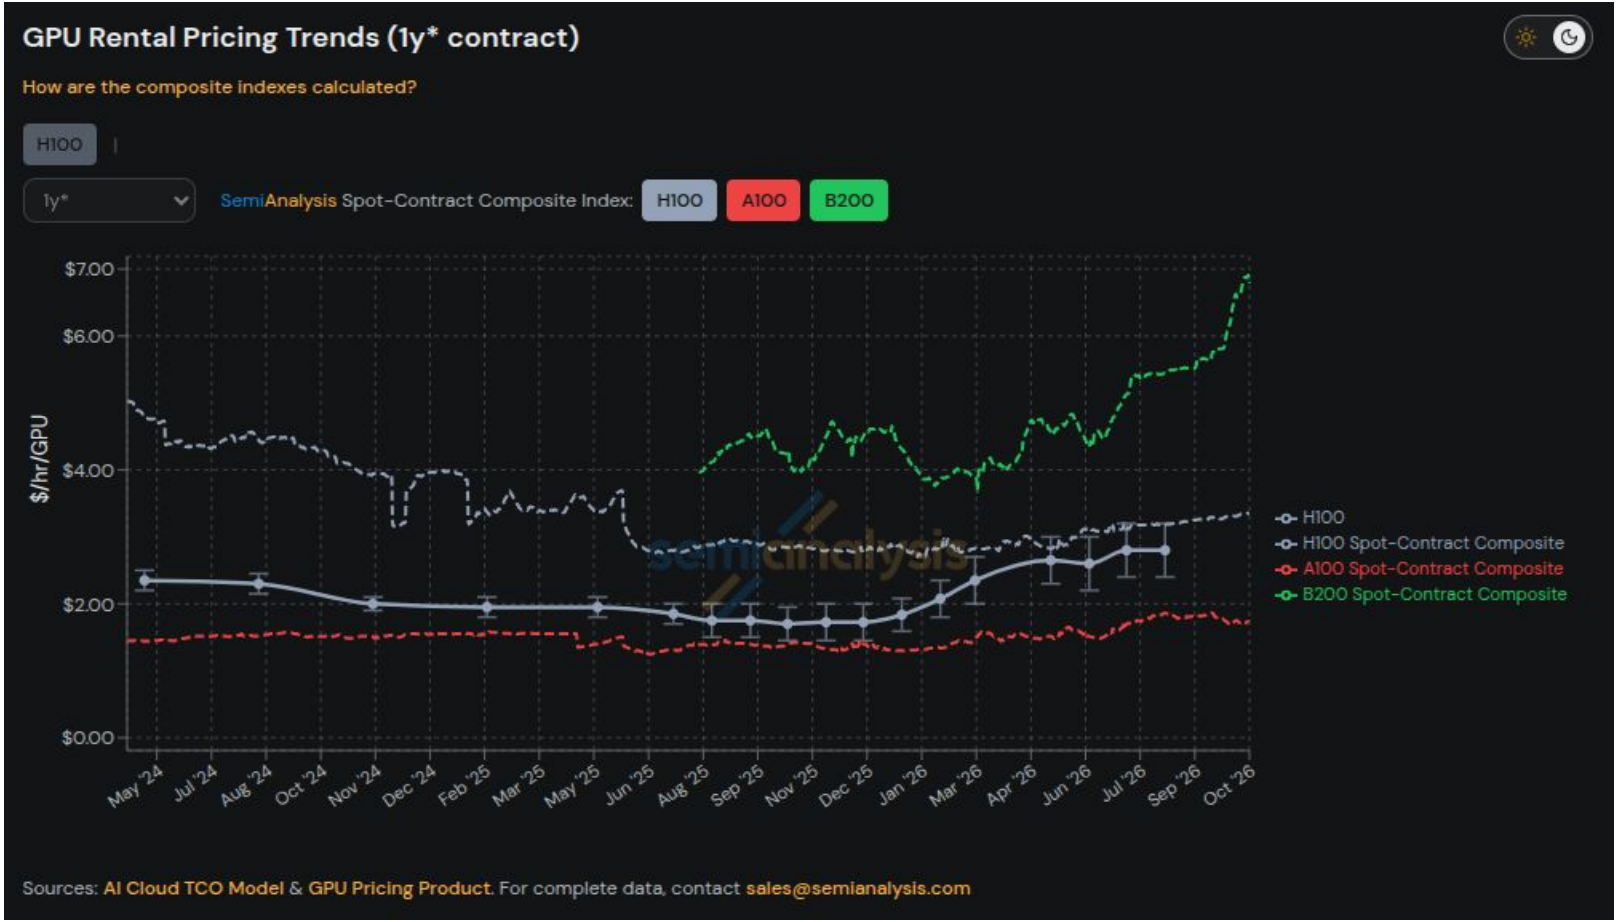
A100 and H100 unit economics

US$ per GPU-hour • model assumptions, not quoted market prices

GPU / item	Year 1	Year 2	Year 3	Year 4	Year 5
A100 revenue	**1.80**	**1.64**	**1.49**	**1.36**	**1.23**
A100 other opex	0.38	0.40	0.42	0.44	0.46
A100 power / cooling	0.07	0.07	0.07	0.08	0.08
H100 revenue	**3.25**	**2.96**	**2.69**	**2.45**	**2.23**
H100 other opex	0.70	0.74	0.77	0.81	0.85
H100 power / cooling	0.10	0.10	0.11	0.11	0.11

Annual assumptions: rental rates −9%; other operating cost +5%; power / cooling +3%.

A100 and H100 unit economics



Building the 100 MW example

H100 economics • initial investment and operating assumptions

Input	Model assumption
Installed GPU capacity	55,556 GPUs
Utilisation	90%
Equipment cost	US$40,000 per GPU → US$2.222bn
Tenant infrastructure	US$3m per MW → US$0.300bn
Total initial capex	**US$2.522bn**

What does 100 MW earn?

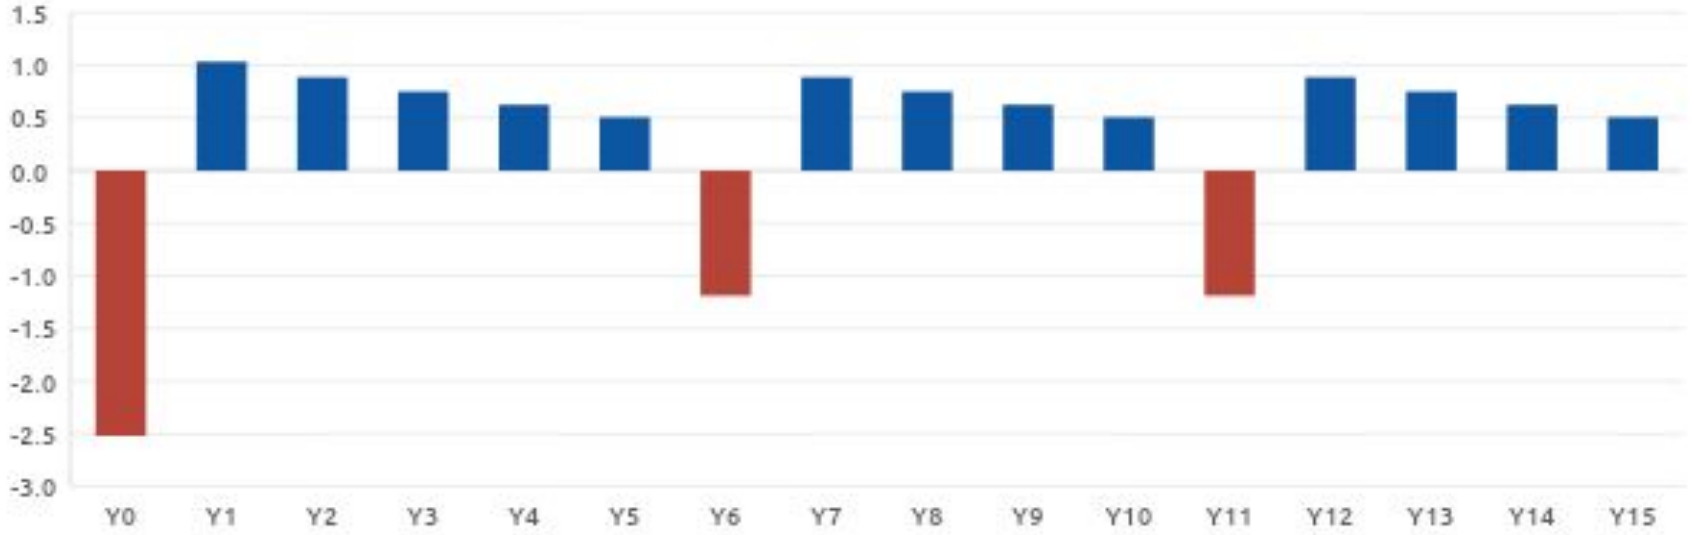
US$m per year • H100 worked example

Operating item	Year 1	Year 2	Year 3	Year 4	Year 5
Revenue	1,424	1,295	1,179	1,073	976
Other operating costs	-341	-358	-376	-394	-414
Power and cooling	-49	-50	-52	-53	-55
Cash / EBITDA	**1,034**	**888**	**752**	**625**	**507**

90% utilisation; 438m billed GPU-hours each year. EBITDA excludes interest, tax and equipment replacement.

A facility needs repeated reinvestment

US$bn • cash flow before interest

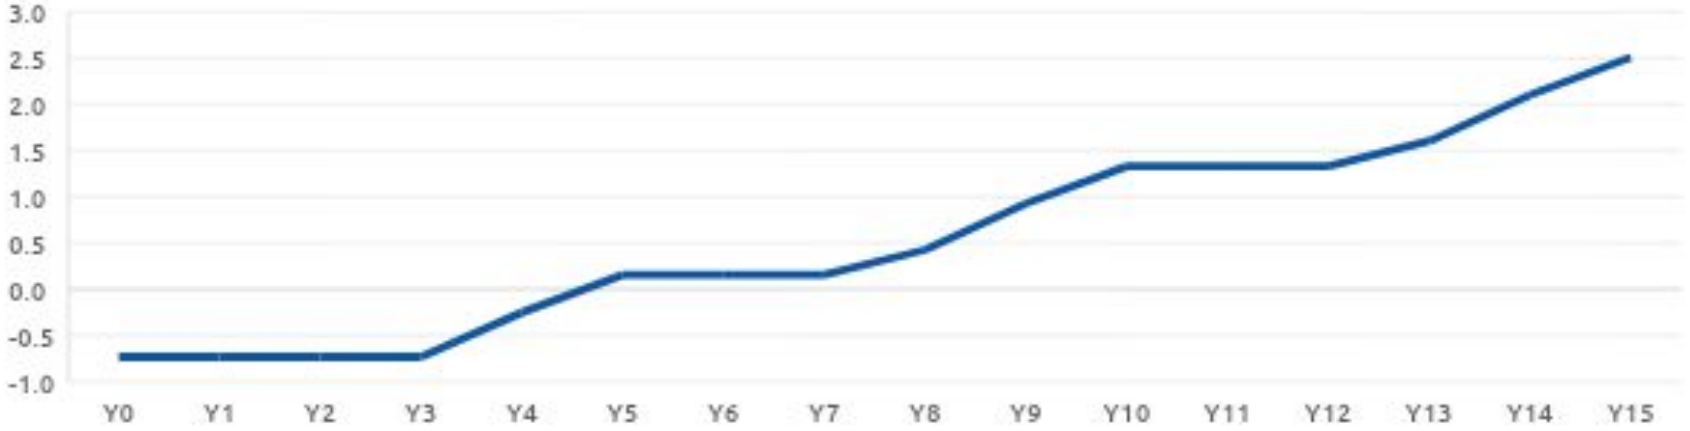


Initial capex in year 0; GPU refresh outflows in years 6 and 11. The facility continues operating.

20% modelled equity IRR

75% initial debt funding • 9% borrowing cost • simplified facility model

Capital is committed before equity cash returns



Debt is cleared during year 4; cumulative equity cash turns positive in year 5, within this model.

Lower operating risk can enable more leverage

Contracted revenue improves financeability; debt changes the equity risk

Multi-year contracts + customer prepayments

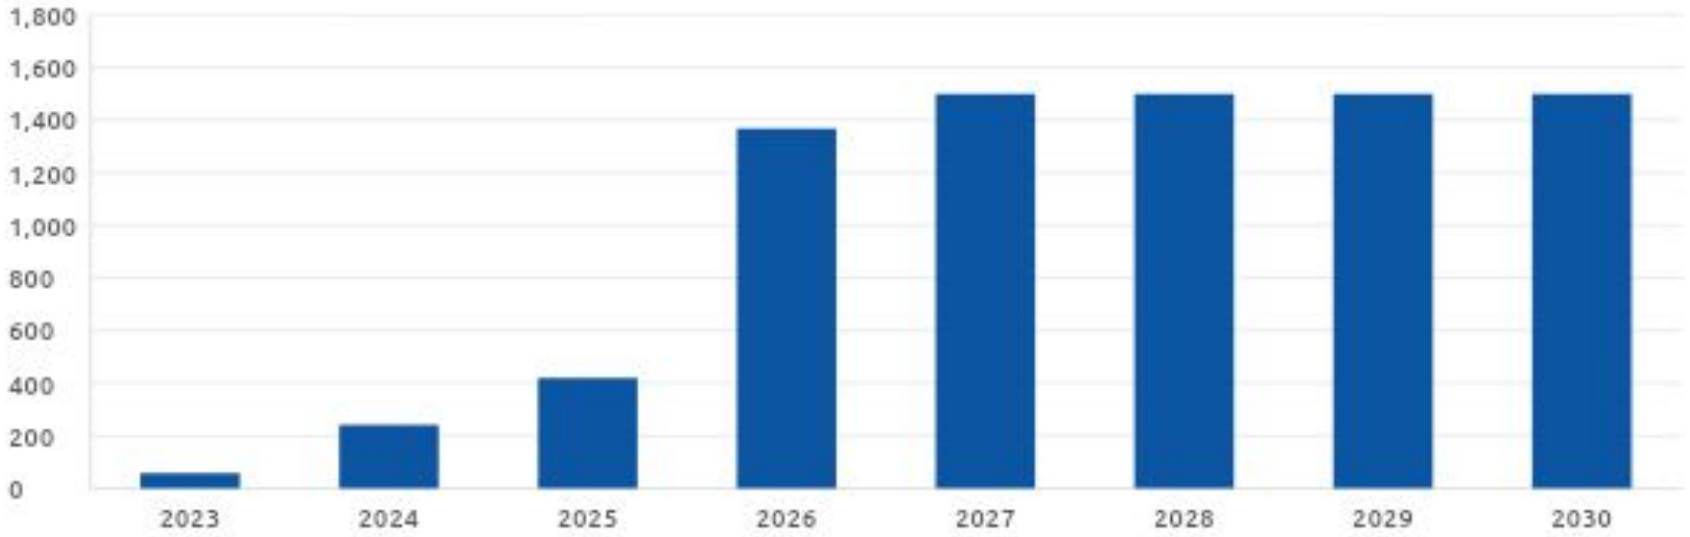
More visible cash receipts

More borrowing + faster construction

Less room for delays, overruns or refinancing stress

Construction grows faster than the cash harvest

Annual capacity additions in the Economic Model • MW



58.5 MW in 2023 → 241.5 MW in 2024 → 1,366 MW in 2026; then 1,500 MW annually to 2030.

Pre-tax earnings • US$bn

The rows show theoretical earnings from centres built each year

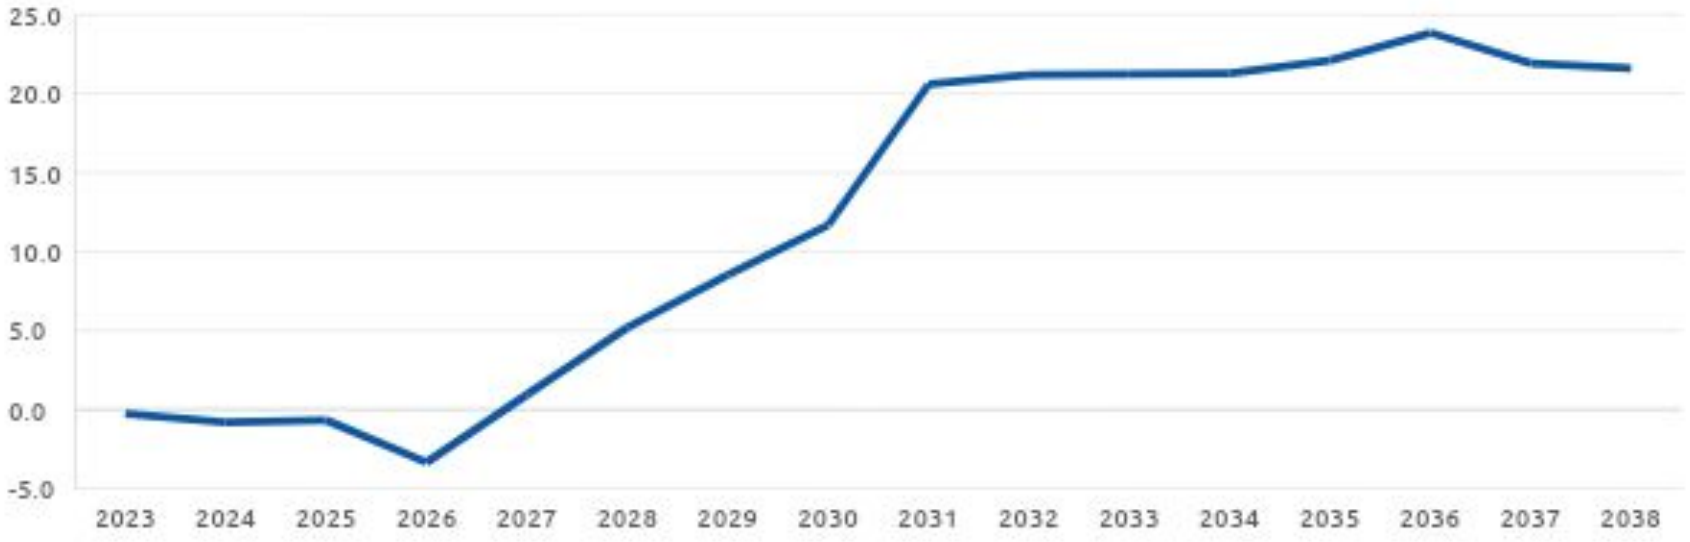
Vintage / MW	2023	2024	2025	2026	2027	2028	2029	2030
2023 / 59	-0.25	0.23	0.18	0.14	0.09	0.03	0.33	0.20
2024 / 242	—	-1.05	0.96	0.76	0.57	0.38	0.10	1.38
2025 / 420	—	—	-1.82	1.68	1.32	0.99	0.67	0.18
2026 / 1,366	—	—	—	-5.92	5.46	4.30	3.21	2.17
2027 / 1,500	—	—	—	—	-6.50	5.99	4.73	3.53
2028 / 1,500	—	—	—	—	—	-6.50	5.99	4.73
2029 / 1,500	—	—	—	—	—	—	-6.50	5.99
2030 / 1,500	—	—	—	—	—	—	—	-6.50
Total	**-0.25**	**-0.81**	**-0.67**	**-3.35**	**0.94**	**5.19**	**8.54**	**11.67**

The total doesn't include corporate overheads. Noting this is illustrative

First-year losses include depreciation and interest; they are not capex.

The earnings harvest arrives later

Aggregate pre-tax earnings • US$bn • simplified vintage model



New vintages end after 2030 in this exhibit. Replacement cycles continue; earnings are not free cash flow.

“Debt” needs a definition

Separate the obligations before adding up a headline number

Category	What the obligation represents
Financial borrowings	Principal, interest and refinancing obligations
Recognised lease liabilities	Present value of qualifying future lease payments
Future commitments	Contractual spending; may not yet be recognised debt
Customer advances	Cash received for services still to be delivered

Reconcile dates, gross versus net debt, lease treatment and potential double counting.

Customer prepayments are useful funding

Illustrative accounting example • US$100m paid before service delivery

When payment arrives	Effect
Cash asset	+US$100m
Deferred revenue liability	+US$100m
Revenue earned at receipt	US$0 until the service is provided

The liability is real. So is the funding benefit.

Illustrative only. Cash may subsequently fund construction; it need not remain available to offset the liability.

What if expansion stops?

Economic Model scenario outputs • US$bn

Last new vintage	Operating value[1]	Debt	Discounted equity[2]
2026	61.9	47.0	7.4
2027	106.4	82.8	11.7
2028	150.9	80.5	35.0
2029	195.4	89.1	52.9
2030	**239.9**	**90.8**	**74.2**

[1] 10× average 2031–35 EBIT. [2] (Operating value – debt) ÷ 1.15^5 for every scenario, as in the source.

The valuation depends on reaching the harvest

More capacity adds operating value—and a larger funding obligation

US$35.0bn

discounted equity: stop in 2028

US$74.2bn

discounted equity: stop in 2030

This is a scenario comparison, not a forecast.

Friendlier valuation multiple, shallower drop of in revenue and you can triple those valuations easily.

Both figures retain the workbook’s 10× EBIT valuation and common five-year discount convention.

What would break the financing explanation?

The conclusion depends on the operating assumptions surviving a stress test

Test	What would weaken the thesis
Compute pricing / utilisation	Receipts fall below the earning curve
GPU life / refresh cost	More frequent or more expensive replacement
Construction and delivery	Revenue starts later while funding costs continue
Financing and customer risk	Refinancing fails, cash is restricted, or contracts underperform

Stress secured borrowing costs separately from secondary unsecured bond yields.

A good operator can still be a fragile equity investment

The model supports the project economics.
Expansion and leverage determine whether equity captures them.

More from Nucleus Wealth

Find out what your personalised investment could look like in 10-20 minutes – commitment free & at no cost.

Find us on all major podcast platforms:

Spotify

Listen on Google Podcasts

Listen on Apple Podcasts

Nucleus Wealth on social media:

Facebook.com/nucleuswealth

Linkedin.com/company/nucleuswealth

Instagram.com/nucleus_wealth

Twitter.com/nucleuswealth

Read insightful blog articles:

nucleuswealth.com/blog

Have a guest or topic suggestion for the show? Leave a comment down below.