# The AI capex stack: a primer in two cuts

## Who depends on whom, who loses first, and who is long the GPU collateral

Open-Sourced Rates Observatory · irates.today · a ruxton.ai research product

2 November 2026 · reference document for the AI capex page

### Why two stacks

Two structures are easy to confuse when people talk about “the stack” in AI infrastructure, and the readings on the Observatory’s AI capex page only make sense against both.

The first is the **industry stack**: who sells what to whom. Tokens are sold at the top and electrons are bought at the bottom; money flows down the stack as capex, rent and power bills, and capacity flows up. When the money stops, the risk flows back up — but not evenly, because some layers borrowed against an asset that only earns inside someone else’s building.

The second is each player’s **capital stack**: who gets paid first when that player’s money stops. A credit spread on the page is a price at one rung of one company’s capital stack; that company sits at one rung of the industry’s. The question the page keeps asking — who is actually long the GPU collateral — can only be answered by reading the two together.

The diagrams are schematic. The names are examples drawn from the page’s single-name board, not a census, and nothing here is a recommendation on any security.

## 1. The industry stack: who depends on whom

**The industry stack: who depends on whom, and who is exposed to the GPU collateral**




Blue: the anchor tenant whose commitments fund everything below. Amber: depends on GPU collateral — borrows against it, or is paid by someone who does. Red: gives the guarantees, so is exposed to the collateral it sold. Money and contracts flow down the left column; when they stop, the exposure flows back up through the amber and red rows first.

*The industry stack*

Reading it as a chain of dependencies rather than a list of sectors:

**Model owners and labs** sit at the top because their commitments fund everything below. They buy compute on multi-year take-or-pay contracts from hyperscalers and from capacity operators; those contracts are the hyperscalers' backlog and the operators' only revenue. The labs themselves depend on two things — powered capacity being delivered

on time, and their own fundraising to pay for it. A smaller raise at the anchor tenant is a smaller cushion under every contract below it.

**Hyperscalers** depend on the labs renewing, on chips being delivered, on landlords' long leases and on power bought under PPAs. They are hedged in the ways that matter: they own their fleets outright, they buy power forward, and they have investment-grade balance sheets. The exception is capacity built on a single tenant's commitment, which is where Oracle's protection has been pricing since September.

**Chips and systems** are the guarantor. They sell the asset and, increasingly, stand behind it downstream through capacity backstops, residual-value guarantees and equity stakes in their own customers. This is the 1990s aircraft structure — the manufacturer supporting the value of the thing it sells so that buyers who could not otherwise finance it can — and it means the party that sets the asset's obsolescence is also the party exposed to its residual. The exposure is contingent and unpriced until it is called.

**Capacity operators** — the neoclouds and the sponsor-financed vehicles — own the collateral. Their balance sheet is GPUs, financed on the GPUs themselves through delayed-draw secured loans and equipment leases, serving one or two tenants' take-or-pay contracts that are often terminable on notice. They lease the halls, and below about fifty megawatts they mostly buy power day-ahead at the node, unhedged. Everything that goes wrong in this industry goes wrong here first.

**Landlords and developers** depend on the tenant above paying rent for the term, on interconnection arriving, and — at the small end — on regional and community construction lenders rolling. Their exposure to the collateral is indirect: an operator-tenanted shell is worth its tenant. The leases are securitised and sold to insurers as investment grade, and reprice only when a lease is not renewed. This is where the de-hubbed-airport precedent applies, except that a data centre has no residual use agreement obliging the remaining tenants to cover the debt service.

**Power** depends on load actually arriving behind the interconnection queue. Its exposure runs the other way: the generators are long the option the unhedged operator is short when a freeze or a heat dome sends the node to its cap.

## 2. A capacity operator’s capital stack: who loses first

**A capacity operator's capital stack: who loses first, and who is long the collateral**




Lower rungs are paid first and seen least. The public tapes see the top three rungs; the statutory blanks see the bonds and the securitisations; the middle — where the GPU collateral is — is private. The direct rungs rise or fall with what a used GPU in someone else's building is worth; the indirect ones only when the operator stops paying.

*A capacity operator’s capital stack*

The capital stack of a capacity operator, read from first loss down, with the column that matters most on the right: which rungs are actually long the GPU collateral.

Common equity — founders, venture, sponsors, and the high-net-worth builders at the edge — loses first and is the only rung the public can see every day, through the single-stock swap tape. Preferred and convertibles, the neoclouds’ favoured instrument, are visible only in filings. Senior unsecured bonds are what a CDS contract references and delivers, so this is the rung the single-name board prices; its exposure to the collateral is indirect, because the secured lenders sit ahead of it.

The two amber rungs are the ones directly long the chips. Secured GPU-backed loans — delayed-draw, with the chips and the contract they serve as collateral — are marked to model, because there is no market mark for bare hardware; it clears in private requests for quote. Equipment leases and SPV debt go a step further: the lessor owns the chips, a vehicle

issues private bonds against them, and none of it is public until it is tested in court. The landlord's securitisations sit one tier down and are indirectly exposed through the operator's rent; they are where the insurers' real exposure sits, carried at NAIC 1–2.

Vendor support sits outside the stack rather than on it: capacity backstops, residual guarantees and equity stakes from the party that sold the asset. It is the rung most directly long the collateral's value and the least priced, visible only in the vendor's quarterly footnotes.

The pattern is the point. The public record is thickest at the top of the stack and thinnest where the collateral is. A GPU-secured lender is paid ahead of the bondholders the page can price, against an asset the page can only price by implication.

## 3. How the tiers fund themselves

| Tier | Typical funding | Where the risk ends up | Visible on the page |
|---|---|---|---|
| Model owners | Private equity rounds; compute bought on multi-year take-or-pay; little debt of their own | On their suppliers — the commitments are the suppliers' receivables and backlog | Indirectly, through Oracle's and the operators' credit; Meta and Alphabet directly |
| Hyperscalers | Operating cash flow; investment-grade bonds; increasingly off-balance-sheet joint ventures and leases for capacity | Bondholders and the statutory books (NAIC 1–2); the JV lenders | CDS, swap prices, Schedule D holdings |
| Chips & systems | Cash flow; vendor financing and backstops extended downstream | Contingent exposure to customers' capacity — the British Aerospace question | CDS (markets created this year), swap prices, footnotes |
| Capacity operators | Converts and preferred; high-yield bonds; GPU-backed delayed-draw loans; equipment leases through vehicles; vendor backstops | Secured lenders and lessors hold collateral with no mark; unsecured holders behind them; equity first | CDS on the 500 coupon, swap prices; the loans mostly not |
| Sponsor-financed capacity | Equipment-finance leases and private bonds out of a vehicle, funded by insurer-backed lenders, against the sponsor's broader credit | The vehicle's funders — insurance general accounts and private-credit funds | The sponsor's CDS; the vehicle not |
| Landlords & developers | REIT unsecured bonds; ABS and CMBS on lease cash flows; construction loans (regional and | Securitisation holders (insurers, NAIC 1–2); the construction lenders if the tenant never arrives | CDS and swap prices for the listed REITs; Schedule D for the ABS |

| | | | |
|---|---|---|---|
| | community banks at the small end); private credit | | |
| Power | Utility balance sheets; PPAs with the hyperscalers; merchant exposure at the node for the rest | The generators are long the option the unhedged operator is short | CDS and swap prices; the commodities tape for hedging volume (in preparation) |

## 4. Who is exposed to the GPU collateral, in order

Directly, in the order they would feel a fall in what a used GPU in someone else's building is worth: the operators' secured lenders and lessors, including the SPV bondholders behind sponsor-financed leases and the insurance general accounts and private-credit funds that fund them; the vendor, through the backstops and residual guarantees it has written; the operators' own equity and converts, which sit behind those lenders but in front of nothing; and the operators' unsecured bondholders, whose protection is what the single-name board prices.

Indirectly, through a tenant rather than the asset: the landlords and their securitisation holders, when an operator stops paying rent; the construction lenders behind the small developers, when the tenant never arrives; the hyperscaler whose capacity was built on one tenant's commitment; and, through the funding that pays the take-or-pay, the anchor tenant's own investors.

Not exposed, or exposed the other way: the hyperscalers that own their fleets and buy power forward; and the generators, who are long the power option the unhedged operator is short.

## 5. What the page shows of each

The single-name board (SEC credit file) prices the senior unsecured rung of every public name on the stack, and the coupon a name prints on is the market's own reclassification of it. The equity leg (SEC single-stock swap file) prices the first-loss rung and counts the synthetic financing written on it. The collateral panel (GPU rental rates, by generation) is the nearest public curve to the asset the amber rungs are long. The holders table (Schedule D) shows which rungs the insurers actually own — the bonds and the securitisations, not the loans. The watch list is the order in which these should move if the thesis holds: names first, equities second, the index last.

## Credits

The observations that bare GPU hardware clears bilaterally through private requests for quote and so has no public mark, that operators in power-capped halls retire working chips when the same megawatts earn more under denser racks, that the chip is better treated as a consumable option on power availability than as durable equipment, and that the sub-50 MW tier buys power unhedged at the node, are Dean Lee's (deanlee.info), made in correspondence on 26–28 October 2026 and used with his permission. The aircraft-finance

comparison and the capital-stack reading are the Observatory's. GPU rental data: GetDeploying, GPU rental price history, CC BY 4.0.

*Open-Sourced Rates Observatory. Derived aggregates from public data only; not investment advice. Research note 05, on the collateral, follows.*