

COMPUTE FUTURES AS A P&L TOOL FOR RELIANCE INTELLIGENCE

Hedging perishable GPU-hours at Jamnagar: use-cases, sensitivity analysis, and what the futures market does and does not solve

Companion to "The Token and GPU-Capacity Business: An End-to-End P&L Primer." Illustrative analysis prepared August 2026 for Ilan Gleiser. Public-source facts on Reliance Intelligence and the CME/ICE contracts are cited; all cluster and site economics are illustrative constructions, not Reliance's actual figures.

Summary

Reliance Intelligence is building one of the world's largest AI compute platforms: a ₹10 trillion (~\$110 billion), seven-year commitment announced in February 2026, with the first 120 MW of NVIDIA GB300 capacity at Jamnagar due online in the second half of 2026, powered by Reliance's Kutch solar assets and scalable toward a multi-gigawatt campus. The stated commercial ambition — to cut the cost of intelligence for India the way Jio cut the cost of mobile data — means the business will deliberately run at aggressive prices in a price-sensitive domestic market. That makes the core P&L problem identified in the primer even sharper: a GPU-hour is perishable inventory on a 4–6 year depreciating asset, so every unsold or under-priced hour is revenue lost forever, and every hour sold on a merchant basis carries full exposure to a rental market that reprices downward every hardware generation.

The compute futures market that is now forming — CME Group's Silicon Data H100 and B200 Rental Index Futures, listed on NYMEX and targeted for launch on October 5, 2026, and ICE's planned cash-settled contracts on the Ornn Compute Price Index covering H100, H200, B200 and other GPUs — gives an operator with a structurally long physical position a way to do what every other commodity producer does: sell forward the part of production it has not already contracted. This memo sets out four use-cases for Reliance (revenue hedging, financing, yield management on the uncommitted 30–40% of capacity, and deal-desk price governance), quantifies them on the primer's 1,024-GPU H100 cluster, scales the conclusions to a 120 MW GB300 site, and is candid about the limits — basis risk, liquidity, collateral, and the fact that futures hedge price, not utilization.

The headline numbers: on a 1,024-H100 cluster with 65% of capacity under take-or-pay at \$2.10 and the rest sold at spot, a fall in spot from \$2.10 to \$1.70 cuts pre-tax income from \$1.9M to \$0.8M per year (–56%), pushes equity payback from 2.5 to 3.0 years and DSCR from 1.70× to 1.57×. A short futures position covering 80% of the merchant hours at \$2.10 holds pre-tax income at \$1.8M across the whole \$2.50–\$1.50 range. At Jamnagar scale the same merchant sleeve is roughly 160 million GPU-hours a year, so each \$0.10 move in the Blackwell spot index is worth about \$16M of annual EBITDA — the hedge is not a rounding error, it is a line item.

1. Reliance's Position: Structurally Long Perishable Compute

Every cloud provider is long compute, but Reliance's exposure has three distinctive features. First, scale and speed: 120 MW of GB300 in phase one (a fleet Reliance describes as the equivalent of 75,000-plus H100s), scalable to 200,000-plus H100-equivalents, built on a seven-year capital plan — so at any given moment a large tranche of new capacity is coming online ahead of its contracted demand. Second, a strategic mandate to price low: the explicit goal is to make AI "dramatically more affordable for every Indian," which means the commercial team cannot simply defend margin with price; it must defend margin with utilization, cost, and risk management. Third, the revenue stack the primer describes — 60–70% of capacity sold as multi-year take-or-pay reservations ("sell certainty first"), with token serving and on-demand monetizing the remainder — leaves a merchant sleeve of 30–40% that is exposed to a spot rental market that has historically compressed 30–50% with each generation transition and is now globally benchmarked.

The economics of that merchant sleeve are the economics of a commodity producer without a forward market: the cost base (depreciation, interest, committed power from Kutch, colocation) is fixed in advance, while the price of the output is set hour by hour by a market Reliance does not control. Until now the only hedges were physical — discount a reservation, or chase spot volume. Exchange-traded compute futures add a financial hedge that does not require touching the physical book at all.

2. The Instruments

Venue	Contract	Underlying	Settlement	Status
CME Group (NYMEX)	Silicon Data H100 Rental Index Futures	Silicon Data H100 hourly rental index	Cash; one contract $\approx$ one month of rent per GPU	Launch targeted 5 Oct 2026, pending CFTC review
CME Group (NYMEX)	Silicon Data B200 Rental Index Futures	Silicon Data B200 hourly rental index	Cash; monthly	Same
ICE	GPU compute futures (with Ornn)	Ornn Compute Price Index (OCPI) — printed spot transactions; H100, H200, B200, RTX 5090, more to follow	Cash, USD	Announced; subject to regulatory approval

Sources: CME Group press release and SER-9785 listing notice (Aug 2026); ICE/Ornn announcement. Contract specifications (tick, months listed, position limits) were not yet public at the time of writing.

Two points matter for a GB300 operator. The contracts reference H100 and B200 (and at ICE, H200) indices, not GB300 — so a Jamnagar hedge is a cross-hedge, and its effectiveness depends on the correlation between Blackwell-Ultra rental prices and the B200 index. Blackwell generations are likely to move together under common drivers (next-generation chip supply, algorithmic efficiency, macro AI demand), but the basis will not be zero and should be measured once the contracts trade. And the indices are global/US spot benchmarks, while Reliance's physical sales are in India at prices it deliberately sets below global levels; the hedge protects against moves in the global benchmark, not against Reliance's own pricing decisions.

3. Four Use-Cases, Applied to Jamnagar

3.1 Hedging merchant revenue against price compression

Reliance sells B200 (or H100, as a proxy) rental index futures against the portion of GB300 capacity it expects to sell on-demand or as tokens over the next 1–12 months. If the index falls, the short futures gain offsets the lower realized rental price on the physical fleet; if the index rises, the futures lose and the physical book earns more. The position should be sized to the merchant sleeve only: take-or-pay revenue is already fixed-price and hedging it would create, not remove, exposure. A hedge ratio below 100% (the memo uses 80%) reflects basis risk and the fact that part of the merchant sleeve is token revenue, whose link to the rental index is real but looser.

3.2 Lowering the cost of capital

The ₹10 trillion programme will be financed substantially with debt, and lenders to GPU fleets underwrite contracted cash flows — this is why CoreWeave's \$24.9B debt stack sits against a \$99B backlog, and why its facilities are priced off delivery milestones. An unhedged merchant sleeve is the part of the revenue stack lenders haircut hardest. A liquid forward curve does two things: it gives lenders an observable mark for the uncommitted capacity (instead of a sponsor's spreadsheet), and it lets Reliance convert merchant revenue into quasi-contracted revenue by selling forward. In the worked example below, hedging the merchant sleeve holds DSCR at ~1.70× through a \$1.00 spot decline instead of letting it fall to 1.50× — the difference between comfortably inside and close to a typical 1.35–1.50× covenant. Better DSCR stability supports higher advance rates and lower spreads on the project debt; on a \$4B site, 50–100 bp of spread is \$14–28M a year.

3.3 Yield management on the uncommitted 30–40%

The "sell certainty first" strategy creates a recurring dilemma for the uncommitted capacity: discount reservations to fill it (locking in a low price for years), or run it merchant and hope. Futures add a third option: keep the physical capacity flexible — available for on-demand, token serving, sovereign or enterprise bursts — while locking the financial yield with a short futures position sized to expected merchant volume. When a good physical contract appears, the hedge can be unwound and replaced by the contract; in the meantime the P&L is protected. The futures market also prices the term structure, so the deal desk can see whether the market expects B200 rentals in 2027 to be above or below today's level before deciding how long a reservation discount to offer.

3.4 Price governance and deal structuring

A public, regulated benchmark lets Reliance's commercial team write contracts that reference the index: reservations priced at index-plus-or-minus a spread, on-demand rate cards that reset to the index monthly, and expansion options ("capacity calls") whose strike references the forward curve. It gives a clean margin-floor rule — never sell physical capacity below the forward price net of basis and cost of carry — and a transparent answer to large customers negotiating on price. It also aligns customer and provider: a customer that wants price certainty can buy futures itself, which lets Reliance avoid giving away fixed-price options for free inside physical contracts.

4. Worked Sensitivity: 1,024-GPU H100 Cluster

Assumptions carried over from the primer: 1,024 H100s, \$45.3M all-in capex (\$44.3k/GPU), 5-year straight-line depreciation, cash opex of \$0.485/GPU-hr (power, colocation, ops) incurred on all hours because power is committed, 70% debt at 8% amortizing over 5 years (level debt service \$7.94M/yr; year-1 interest \$2.54M), 65% of hours under take-or-pay at \$2.10, the remaining 35% sold merchant at 85% utilization at the spot price. Hedge case: short futures on 80% of merchant hours at \$2.10, settled against the spot price.

Spot \$/GPU-hr	Revenue \$M	EBITDA \$M	Pre-tax \$M	DSCR	Equity payback (yrs)	Full capex payback (yrs)
\$2.50	18.92	14.57	2.97	1.83×	2.05	3.77
\$2.10	17.85	13.50	1.90	1.70×	2.45	4.13
\$1.90	17.31	12.96	1.37	1.63×	2.71	4.34
\$1.70	16.78	12.43	0.83	1.57×	3.03	4.58
\$1.50	16.25	11.90	0.30	1.50×	3.44	4.84

Table 1 — Unhedged, 65% take-or-pay at \$2.10, merchant 35% at 85% utilization. Equity payback = equity (\$13.6M) ÷ (EBITDA – debt service). Full payback = capex ÷ (EBITDA – interest).

Spot \$/GPU-hr	Hedge P&L \$M	EBITDA \$M	Pre-tax \$M	DSCR	Equity payback (yrs)
\$2.50	–1.00	13.56	1.96	1.71×	2.42
\$2.10	0.00	13.50	1.90	1.70×	2.45
\$1.90	+0.50	13.47	1.87	1.70×	2.46
\$1.70	+1.00	13.44	1.84	1.69×	2.47
\$1.50	+1.51	13.40	1.81	1.69×	2.49

Table 2 — Same cluster with 80% of merchant hours sold forward at \$2.10. Residual drift reflects the 20% unhedged tail.




Figure 1 — Pre-tax income across spot prices for three postures. The take-or-pay book already absorbs most of the shock; the hedge removes what remains; a fully merchant fleet is underwater below ~\$2.10.

Reading the tables. The \$2.10 → \$1.70 move costs the unhedged cluster \$1.07M of EBITDA and pre-tax income (merchant hours × \$0.40), which is 56% of pre-tax profit because the cost base is fixed. DSCR slips from 1.70× to 1.57× and equity payback stretches from 2.5 to 3.0 years; at \$1.50 equity payback reaches 3.4 years — still inside the 5-year life, but only because 65% of the book is contracted. Strip the take-or-pay book out and the picture is stark: a fully merchant cluster at 85% utilization earns essentially zero pre-tax at \$2.10 and loses \$3.0M a year at \$1.70, with DSCR at 1.08× — a covenant breach on any normal facility. That comparison is the quantitative case for "sell certainty first"; the futures hedge is what finishes the job on the sleeve that certainty-selling leaves behind.

Two further sensitivities are worth the deal desk's attention. Utilization: at \$1.70, taking merchant utilization from 85% to 50% costs another \$1.9M of EBITDA and takes pre-tax to a \$1.0M loss — and no futures position helps, because futures hedge price, not volume. Depreciation life: at \$2.10, the same cluster shows pre-tax of –\$0.4M on a 4-year life, +\$1.9M on 5 years and +\$3.4M on 6; at \$1.70 those figures are –\$1.4M, +\$0.8M, +\$2.3M. The difference between a 4- and 6-year life is worth more to reported profit than the entire \$2.10 → \$1.70 price move — which is exactly why auditors, lenders and analysts will look at both together.

5. Scaling to Jamnagar: 120 MW of GB300

Treat 120 MW as IT load. At roughly 137 kW per GB300 NVL72 rack that is about 875 racks and 63,000 Blackwell-Ultra GPUs; at an illustrative \$4.5M per rack all-in (analyst rack estimates span \$3.7–6.5M, plus fabric, storage and integration) the compute capex is on the order of \$3.9B before the building, substation and cooling plant. Illustrative pricing: committed GB300 capacity at \$5.00/GPU-hr for 65% of hours (below the ~\$6.46 24-month reserved rate currently quoted by the one provider publishing GB300 prices, reflecting Reliance's stated intent to price aggressively), merchant hours at the spot index with \$4.31 (current published GB300 spot) as base case. Cash opex scales with power: ~1.9 kW per GPU against 1.275 kW for an H100 gives roughly \$0.72/GPU-hr for power, facilities and ops. Depreciation over 5 years and 70% debt at 8% complete the stack.

GB300 spot \$/GPU-hr	Revenue \$B	EBITDA \$B (unhedged)	Pre-tax \$B (unhedged)	Hedge P&L \$M (80% @ \$4.31)	Pre-tax \$B (hedged)
\$5.00	2.61	2.21	1.21	-107	1.10
\$4.31	2.50	2.10	1.09	0	1.09
\$3.50	2.37	1.97	0.96	+125	1.09
\$2.75	2.25	1.85	0.84	+241	1.08

Table 3 — Illustrative 120 MW GB300 site, 63,000 GPUs, \$3.9B compute capex, 65% committed at \$5.00, merchant at 85% utilization. Pre-tax after \$0.79B depreciation and \$0.22B interest. Hedge = 80% of merchant hours sold forward at \$4.31. Cross-hedge via B200 index futures assumed to track GB300 one-for-one; real basis will differ.

The arithmetic of scale is what changes the conversation. The merchant sleeve at Jamnagar is about 164 million GPU-hours a year; every \$0.10 in the Blackwell spot index is therefore ~\$16M of annual EBITDA, and the move from \$4.31 to \$2.75 — the kind of compression a generation transition has historically produced — is roughly \$256M a year, or about a quarter of pre-tax income. A futures position sized to 80% of that sleeve neutralizes it. At a ₹10 trillion programme scale, where phase one is a small fraction of the eventual fleet, the same logic applied across successive tranches is the difference between a P&L that looks like a regulated utility and one that looks like a merchant power producer.

The obvious caveat is pricing. If Reliance's domestic strategy means its committed price is materially below the illustrative \$5.00, or if its merchant price is set by Indian rather than global demand, the sensitivity shifts down the table and the cross-hedge basis widens. The framework does not change; the numbers do, and the deal desk should re-run them against its own rate card.

6. What Futures Do Not Solve — and How to Run the Programme

- **Basis risk** — the contracts reference H100/B200 (and H200 at ICE) global spot indices; Reliance is selling GB300 in India. Measure the correlation before sizing, hedge a fraction rather than the whole, and revisit the ratio as the contracts trade.
- **Volume risk** — futures hedge price, not utilization. A 35-point drop in merchant utilization costs more than a \$0.40 price move in the worked example, and only physical demand creation — token products, sovereign and enterprise contracts, Jio distribution — addresses it.
- **Liquidity and tenor** — these are new contracts; open interest will be thin at first and the curve may only extend a few months. Start with near-dated rolling hedges on a small fraction of the sleeve and scale as liquidity allows; do not assume a five-year hedge is available.
- **Collateral and liquidity** — a short futures position posts variation margin when spot rises — precisely when the physical fleet is earning more. The cash flows net over time but not daily; the treasury needs a margin facility and a cash-flow-at-risk view.
- **Accounting** — to qualify for cash-flow hedge accounting under Ind AS 109 / IFRS 9 the hedged forecast sales must be highly probable and the relationship documented and effectiveness-tested; a cross-hedge with meaningful basis may not qualify, leaving mark-to-market volatility in the P&L even when the economic hedge works.
- **Governance** — a hedge at \$2.10 gives up the upside at \$2.50 (the -\$1.0M in Table 2). That is the point, not a flaw — but the board should explicitly choose to be a price-taker on the merchant sleeve, and the hedge programme should have a written policy, limits, and a clear separation between hedging and proprietary trading.

A practical sequencing: establish the index-referenced price governance (use-case 4) immediately, because it costs nothing and disciplines the deal desk from day one; open a small rolling hedge on the merchant sleeve once the CME contracts trade in October, sized at perhaps 25–50% while basis and liquidity are measured; bring the forward curve and hedge policy into the financing discussions for the next tranche (use-case 2); and scale the hedge toward 70–80% of the merchant sleeve as the market matures. Throughout, the metering-to-settlement chain from the primer is what makes any of this credible: a hedge sized to forecast merchant hours is only as good as the meters that count them.

Sources

- Reliance Intelligence / Reliance AGM 2026 coverage: Business Standard, "Reliance Intelligence bets big on AI with Jamnagar compute platform" (19 Jun 2026); TechCrunch, "Reliance unveils \$110B AI investment plan" (19 Feb 2026).
- CME Group, "CME Group and Silicon Data to Launch Compute Futures on October 5" (11 Aug 2026); CME SER-9785, initial listing of Silicon Data H100 and B200 Rental Index Futures.
- ICE, "ICE and Ornn to Launch GPU Compute Futures Contracts" (2026).

- GetDeploying GB300 and H100 price indices; Spheron, "GB300 NVL72 vs GB200 NVL72: The Real Price Gap (2026)".
- Companion primer for cluster assumptions: "The Token and GPU-Capacity Business: An End-to-End P&L Primer" (Aug 2026).

All cluster and site figures are illustrative. This is general information, not investment, accounting, legal, or hedging advice; hedge programmes should be designed with qualified advisers.