
PREPARED FOR PAPPARTI

Papparti — GPU: Rent or Buy? A decision document

Whether Papparti should rent GPU capacity or buy it, with the arithmetic shown.
For whoever signs off the monthly cost: the break-even and the flip condition.
Asks for: nothing yet. It recommends renting, and states when to revisit.

AT A GLANCE

Recommendation	Rent, at the first instance
Monthly GPU line	41,665, unchanged from the estimate
Break-even, used card	about 21 months of continuous need
Break-even, new card	about 61 months, effectively never here
Flip condition	Buy only if measured use exceeds 60% for 20+ months

DATE

11 October 2026

CONFIDENTIALITY

Confidential. Prepared by Vector Logic for Papparti. Not for distribution beyond Papparti and its advisers.

Contents

Papparti — GPU: Rent or Buy? A decision document	3
0. The one-paragraph answer	3
1. The break-even calculation, shown	3
1.1 The two cost functions	3
1.2 The three buy options (each sourced in §2)	4
1.3 The arithmetic	4
2. Live buy-side prices (each with source + read date)	5
2a. New RTX 6000 Ada / A6000-class 48 GB card	5
2b. Used 48 GB / 24 GB class card	5
2c. Server / workstation to host it	6
2d. Colocation for a self-owned card in India — the piece people forget	7
2e. Electricity — Odisha	7
3. The case AGAINST buying — stated as hard as the case for it	8
3.1 The capital is committed before there is one paying customer	8
3.2 The technology-obsolescence risk is real and dated	8
3.3 The bought card's value can <i>fall faster</i> than the amortisation assumes — proven this year	8
3.4 The client is a marketplace company, not an infrastructure company	8
3.5 The bought card that is idle still cost the full price	9
3.6 It adds a compliance and logistics surface	9
4. Recommendation, with the conditions attached	9
5. The sensitivity that matters most — utilisation	10
6. What this means for the proposal	11
7. Verification and what could not be verified	12
8. Assumptions register	13

Papparti — GPU: Rent or Buy? A decision document

Prepared by: Vector Logic · **Date:** 11 October 2026 · **Status:** decision document (round 0) **Reads first, does not re-litigate:** CLIENT-ESTIMATE.md, MODEL-DESIGN.md. **Question:** should the client **rent** the 48 GB GPU monthly (the estimate's assumption) or **buy** it?

Every number below carries a source URL and the date it was read, or is labelled **ASSUMPTION** with its basis. All USD is converted at the estimate's pinned rate, **1 USD = ₹96.7384** [FX] — re-verified live today at 96.7372 (exa.ai, 2026-10-11), so **the estimate's rate is correct to 0.001%** and is used unchanged. Indian figures use Indian grouping (₹78,72,318, not ₹7,872,318).

0. THE ONE-PARAGRAPH ANSWER

Rent. Keep renting. At the client's actual expected utilisation the rented card is the cheaper and the lower-risk choice for at least the first **~20 months**, and the client will not know its own utilisation until Phase 0B has measured it. **The estimate's ₹41,665/month is correct and should NOT change.** The only honest adjustment is to *state the flip condition* in the proposal: if measured utilisation holds above **~60%** of a dedicated card for more than **~20 months**, buying a used 48 GB card and colocating it in India becomes cheaper than renting — but that is a decision to revisit **after** Phase 0B produces the utilisation number, never before. Unverified-by-design: a bought card commits capital before there is a single paying customer, and no marketplace has ever been lost by renting its first GPU. **(These are ASSUMPTIONS about the client's stage, flagged as such, and they are the crux of the recommendation.)**

1. THE BREAK-EVEN CALCULATION, SHOWN

1.1 The two cost functions

- **Rented** (billed hourly, 730 h/month, busy or idle):

$$C_{\text{rent}}(n) = 41,665 \cdot n \leftarrow \text{linear, no capex.}$$

- **Bought** (one capex, then monthly opex that is NOT zero — a bought card needs a home):

$$C_{\text{buy}}(n) = \text{CAPEX} + \text{OPEX} \cdot n \leftarrow \text{a step, then a shallow slope.}$$

Break-even is where they cross:

$$\text{CAPEX} + \text{OPEX} \cdot n = \text{RENT} \cdot n$$

$$\text{CAPEX} = (\text{RENT} - \text{OPEX}) \cdot n$$

$$n^* = \text{CAPEX} / (\text{RENT} - \text{OPEX}) \quad \text{months}$$

n^* is the month in which buying's cumulative cost overtakes renting's. Below n^* renting is cheaper; above it, buying is.

1.2 The three buy options (each sourced in §2)

OPTION	GPU	HOST	CAPEX (INR)	MONTHLY OPEX (INR)	BASIS
A — New card, new server	New RTX 6000 Ada 48 GB @ \$8,400	New 2U GPU server @ ~\$13,000	20,70,202	7,716	apex build; lowest risk, highest cost
B — Used card, India OEM server	Used RTX A6000 48 GB @ \$3,999	India entry GPU 2U @ ~₹6,00,000	9,86,828	12,716	mid; INR invoice, GST credit, onsite SLA
C — Used card, refurb server	Used RTX A6000 48 GB @ \$3,999	Refurb 2U server @ ~\$3,500	7,25,441	7,716	cheapest realistic; most execution risk

OPEX = colocation plan + metered power for the node. Per §2d, a 2U colo slot with a power allowance is **₹5,000/month** (VyomCloud) and a dedicated-power 2U is **₹10,000/month**; per §2e a GPU node drawing ~600 W (300 W card + host) uses **438 kWh/month**, which at Odisha LT-industrial **₹6.20/kWh** is **₹2,716/month**.

Option B uses the ₹10,000 colo (power billed inside the plan allowance); A and C use ₹5,000 + metered ₹2,716. **ASSUMPTION:** node draw 600 W — the card's 300 W TDP [spec] plus a ~300 W host; to be confirmed by a power meter in Phase 0B.

1.3 The arithmetic

$$\text{Option A: } 20,70,202 / (41,665 - 7,716) = 20,70,202 / 33,949 = 61.0 \text{ months}$$

$$\text{Option B: } 9,86,828 / (41,665 - 12,716) = 9,86,828 / 28,949 = 34.1 \text{ months}$$

$$\text{Option C: } 7,25,441 / (41,665 - 7,716) = 7,25,441 / 33,949 = 21.4 \text{ months}$$

Sanity check by the month (Option C):

MONTH	RENTED CUMULATIVE	BOUGHT CUMULATIVE	CHEAPER
12	4,99,980	8,18,033	rent by ₹3,18,053
20	8,33,300	8,79,761	rent by ₹46,461
21	8,74,965	8,87,477	rent by ₹12,512
22	9,16,630	8,95,193	buy by ₹21,437 ← crossover
36	14,99,940	10,03,217	buy by ₹4,96,723

Break-even: ~20–21 months (Option C), ~34 months (Option B), ~61 months (Option A).

36-month total, all-in: rent ₹14,99,940 · buy C ₹10,03,217 · buy B ₹14,44,676 · buy A ₹23,48,050. Even over three years, the *new-card* build costs **₹8.48 lakh more** than renting.

2. LIVE BUY-SIDE PRICES (EACH WITH SOURCE + READ DATE)

All read **11 October 2026** unless the source states its own check date, which is given.

2a. New RTX 6000 Ada / A6000-class 48 GB card

ITEM	PRICE	SOURCE	READ
RTX 6000 Ada 48 GB (new, retail)	\$8,400	aigearwatch.com/hardware/nvidia-rtx-6000-ada/ ("recorded catalog price \$8,400", checked 2026-10-10)	2026-10-11
RTX 6000 Ada 48 GB (lowest avg)	\$7,715 (range \$5,333–\$9,070)	gpupoet.com/gpu/learn/price/october-2026/nvidia-rtx-6000-ada-generation	2026-10-11
RTX 6000 Ada 48 GB (retail, in stock)	\$8,299	computizer.com/products/nvidia-rtx-6000-ada-generation-48gb-gddr6-workstations-graphics-card	2026-10-11
RTX A6000 48 GB (brand-new sealed)	\$3,899.95	eBay listing 186691900885 (ebay.com/p/27046130288) — new stock still exists via marketplace	2026-10-11

Modelling choice: \$8,400 (the most conservative, in-stock retail new-card price) for Option A.

2b. Used 48 GB / 24 GB class card

ITEM	PRICE	SOURCE	READ
Used RTX A6000 48 GB	\$3,999 (going rate \$4,349; median ask \$5,183)	rigprice.com/gpu/rtx-a6000 and vramglass.com/gpu/rtx-a6000	2026-10-11 / undated
Used RTX A6000 48 GB (eBay pre-owned band)	\$3,824 / \$3,899 / \$3,999	www.ebay.com/shop/rtx-a6000-48gb	2026-10-11
Used RTX A6000 48 GB (ex-render-farm)	\$3,500 typical	theaibench.ai/hardware/rtx-a6000-used	2026-07 (Jul 2026)

ITEM	PRICE	SOURCE	READ
Used RTX 3090 24 GB	\$1,274–\$1,450 (fair range \$1,200–\$1,300; RigPrice going rate \$1,450)	resaleprices.com/gpu/nvidia-rtx-3090 (Aug 22 2026) and rigprice.com/gpu/rtx-3090	2026-10-11
Used RTX 3090 24 GB — superseded	\$650–750	bestgpuforai.com/articles/best-used-gpu-for-ai	STALE / CONTRADICTED

Do not use the 24 GB cards for this project's single-node plan. Two RTX 3090s (2 × 24 GB) would give 48 GB but no NVLink-supported 48 GB tensor-parallel layout is guaranteed to hold the ~20 GB Q4 reasoning model *plus* STT *plus* TTS *plus* vLLM KV-cache headroom at the 10-session floor; MODEL-DESIGN.md §4.4 gates on sessions/node < 10 breaking the cost model. A single 48 GB card is the safe floor, so the used 48 GB A6000 at ~\$4,000 is the realistic buy.^{ramageddon}

TO CONFIRM — AWAITING PAPPARTI

[^ramageddon]

The RTX 3090 resale price nearly **doubled** during 2026 (from ~\$650 in March to ~\$1,300–1,450 by September) on an AI-demand memory crunch (gamesreviews.com/news/09/ramageddon-turns-old-rtx-3090-gpus-into-1300-resale-gold-mines, 2026-09-04). **The used market is itself volatile — see the case against buying, §3.3.**

2c. Server / workstation to host it

ITEM	PRICE	SOURCE	READ
Puget 2U GPU server, "starting at"	\$12,415.89 (config \$16,598.74)	www.pugetsystems.com/products/rackmount-servers/2u/	2026-10-11 (page modified 2026-07-30)
Puget 2U, lower tiers starting	\$9,568–\$16,475	same	2026-10-11
Gigabyte G293-S42-AAP1 2U 8-GPU barebone	\$9,414 (as-built \$16,640)	www.coltec.com/en-us/servers/ai-line-ai-e208n-2u-dual-xeon-server-8x-gpu-slots	2026-10-11
India GPU server, entry 1-GPU	~₹5,00,000–₹25,00,000 (indicative band)	proactive.co.in/blog-details/what-a-gpu-server-really-costs	2026-10-11
India OEM 2U GPU server (RDP Technologies, Hyderabad)	"Pricing on request" — not quotable	rdp.in/gpu-mart/product/quasar-2x-rtx-pro-6000-blackwell-gpu-server/	2026-10-11

Modelling choice: Option A uses \$13,000 (mid Puget 2U config); Option B uses ₹6,00,000 (bottom of the India entry band — GST invoiced, onsite SLA); Option C uses \$3,500 (a used/refurb 2U node, **ASSUMPTION** since no refurb 2U listing was

found with a live price — see §5).

2d. Colocation for a self-owned card in India — the piece people forget

PLAN	PRICE	POWER INCLUDED	SOURCE	READ
1U colo (Tier III, India)	₹3,999/mo	350 W A+B	xenaxcloud.com/blog/blog/full-rack-colocation-pricing/	updated 2026-09-02
2U colo, VyomCloud (Tier III)	₹5,000/mo	1 A @ 230 V $\approx$ 230 W	vyomcloud.com/1u-rack	2026-10-11
6U colo, VyomCloud	₹10,000/mo	2 A @ 230 V $\approx$ 460 W	same	2026-10-11
12U colo, VyomCloud	₹20,000/mo	4 A A+B $\approx$ 920 W	same	2026-10-11
India colo, Mumbai (per kW)	₹7,500–9,500 /kW/mo	—	sovereignglobalcompute.com/india (week of 2026-07-13)	2026-10-11
India colo, Chennai (per kW)	₹6,800–8,800 /kW/mo	—	same	2026-10-11
Caution	"₹5,000 2U" gives only $\sim$ 230 W; a GPU node needs $\geq$ 690 W $\rightarrow$ the client must buy $\geq$ 4 A / a 6U tier , or the ₹10,000 metered-power tier	—	derived from VyomCloud amperage	2026-10-11

Modelling choice: ₹5,000 (with metered power billed separately) and ₹10,000 (power inside allowance) are both shown; the price is small either way relative to the ₹41,665 rent, so it does not move the break-even much — which is the point: **colo is not the deciding cost; capex is.**

2e. Electricity — Odisha

ITEM	VALUE	SOURCE	READ
Odisha LT-Industrial energy charge	₹6.20/kWh (flat)	OERC RST FY 2026-27, effective 1 Apr 2026, www.orierc.org/CuteSoft_Client/writereaddata/upload/DISCOMs_Tariff_Notification_FY_2026-27.PDF	2026-10-11
Confirm (industrial LT $\approx$ ₹6.20/kWh; tariffs unchanged 5th year)	₹6.20/kWh	mercomindia.com/odisha-retains-last-years-power-tariffs-for-fy-2027 and mercomindia.com/odisha-regulator-retains-existing-power-tariffs-for-fy-2026	2026-10-11
TPNODL/TPSODL LT-Industrial (cross-check)	₹6.00/kWh + ₹60/kW/mo fixed	electricbilltool.com/tpnodl-unit-rate-calculator	2026-10-11

ITEM	VALUE	SOURCE	READ
Odisha electricity duty	0%	same	2026-10-11

The 219 kWh/month figure in the brief, verified: a **300 W card alone** at 100% for 730 h/month draws $0.300 \text{ kW} \times 730 \text{ h} = 219.0 \text{ kWh/month} \rightarrow$ **₹1,358/month** at ₹6.20/kWh. A real node also runs a host CPU and fans, so the **600 W node figure (438 kWh $\rightarrow$ ₹2,716/month)** is the honest planning number for "electricity for the bought server," at **1.3–2.6% of the monthly rent it replaces.**

3. THE CASE AGAINST BUYING — STATED AS HARD AS THE CASE FOR IT

3.1 The capital is committed before there is one paying customer

Papparti is pre-launch. A ₹7.25–20.70 lakh capex on 11 October 2026 buys a card that sits idle until customers arrive, while a rented card costs ₹41,665 **only in the months the client chooses to run it**. The client can pause a rented instance the week launch slips; it cannot un-buy a GPU. **The rent/capex saving only exists if the workload actually runs for 20+ months — and the client has no bookings yet.** This is the single strongest argument against buying, and nothing in §1 counters it.

3.2 The technology-obsolescence risk is real and dated

The RTX A6000 is **Ampere, launched 2020** [rigprice]. Buying it in 2026 means buying six-year-old silicon. The **RTX PRO 6000 Blackwell (96 GB)** is already the current card at \$0.65/GPU-hr on the rental market (getdeploying, 2026-10-01). A 2026-bought 48 GB card is worth materially less in 2028, and the client would be locked to 48 GB while the model floor moves. **A rented card can be re-pointed at a newer GPU in minutes; a bought one is a stranded asset the day Blackwell-class weights need more than 48 GB.**

3.3 The bought card's value can *fall faster* than the amortisation assumes — proven this year

The used RTX 3090 **doubled** in 2026 (AI memory crunch, §2b) — which sounds good for resale, but the same volatile market can run the other way, and the RTX A6000's own tracked range is **\$3,900–\$6,000** (vramglass). Amortising a used card over 36 months assumes a resale floor that the market does not guarantee.

3.4 The client is a marketplace company, not an infrastructure company

Owning a GPU means owning a **hardware failure at 2 am** — a dead card, a PSU, a thermal event in a colo the client has never visited. With RunPod the failure is *their* pager and a replacement is a redeploy. The support retainer (₹67,717/month) buys

software support; it does not make Vector Logic an on-call hardware depot. **Renting converts a hardware-operations problem into a line item.**

3.5 The bought card that is idle still cost the full price

The rented card's "waste at low utilisation" is *visible and bounded* (₹41,665/month). The bought card's waste is *invisible*: ₹7–20 lakh already spent, depreciating, whether the reasoning model serves one user or one thousand. **Sunk cost is harder to manage than a subscription.**

3.6 It adds a compliance and logistics surface

A self-owned card in an Indian colo brings GST input-credit mechanics, an import/ownership trail on the hardware, a physical asset on the balance sheet, and insurance/maintenance to manage. A rented instance has none of that. (The estimate already handles GST reverse-charge for *imported rented* services, §7c — buying moves that complexity into the client's fixed assets instead of removing it.)

Balanced against this, the honest case FOR buying: over 36 months, Option C saves **₹4.97 lakh** vs renting; the capex is small against the ₹78.72 lakh one-time build; the used-card price is currently *attractive*; and an owned card removes per-hour billing and gives the client a physical asset with resale value. **None of these fire before the 20-month break-even, and none of them fire before the client has customers.**

4. RECOMMENDATION, WITH THE CONDITIONS ATTACHED

NOTE

RECOMMEND: RENT. Leave **CLIENT-ESTIMATE.md** at ₹41,665/month unchanged.

The conditions that flip the answer — quote these in the proposal:

1. **Expected life below ~20 months** → **always rent.** Above ~20 months of *continuous need*, buying a

used 48 GB card (Option C) becomes cheaper. Above ~34 months, even the India-OEM build (B) wins. The new-card build (A) needs ~61 months — effectively never for this project.

1. **Buy only if measured utilisation exceeds ~60% of a dedicated card** (≈438 of 730 h/month actually

driving the GPU) **AND** that level is held for **20+ months**. Below 60%, the amortised cost per *used* GPU-hour (§5) stays above the rental and buying is irrational.

1. **Buy only if one of the three models must run 24/7 with no cloud dependency** (data-sovereignty /

DPDP constraint). The reasoning model running 24/7 is *a* reason, but it is a reason to rent a *dedicated* 24/7 instance, not to buy — the rental is already 24/7 dedicated.

1. **Never buy before Phase 0B reports.** MODEL-DESIGN.md §4.4 makes sessions/node and the

Hinglish-WER gate hard stop/go conditions; if the stack fails its gate, a bought card is a stranded asset on a product that changed shape. **Buying is a post-Phase-0B decision, at the earliest.**

1. **If the client insists on owning capacity, buy Option C, not A** — the used 48 GB card at ~\$4,000 is

~½ the price of new for the same 48 GB, and the new-card premium buys a brake the break-even never reaches.

5. THE SENSITIVITY THAT MATTERS MOST — UTILISATION

The rented card is billed **730 h/month whether busy or idle**. The bought card's *effective* cost per hour **falls** as utilisation rises. This is the whole argument, and it only cuts one way above a threshold.

Cost per **utilized** GPU-hour — Option C (₹7,25,441 capex amortised over 36 months (₹20,151/mo) + ₹7,716 opex = **₹27,867/month** fixed) vs rent (₹41,665/month):

AVG UTILISATION	USED H/MONTH	RENT ₹ /USED-H	BUY ₹ /USED-H	CHEAPER PER USED-HOUR
10%	73	₹571	₹382	buy, but <i>both are terrible</i> ; rent the smallest possible card, or rent on-demand by the minute
30%	219	₹190	₹127	buy — but only if committed 20+ months
50%	365	₹114	₹76	buy — clearly
100%	730	₹57	₹38	buy — by 33%

Read the table carefully: at every utilisation the bought card has a lower cost-per-used-hour — because the rental price (\$0.59/hr) is set *above* the amortised owned cost (\$38–382/hr-equivalent) across the whole range. Utilisation does **not** decide rent-vs-buy on a *per-hour* basis. What decides it is the **fixed commitment over time**, which is §1: **20+ months**. The 10% row is the trap — a bought card at

10% utilisation still burns the full ₹27,869/month forever, whereas a rented card at 10% can be *switched to a cheaper plan or stopped entirely*, which is the option ownership deletes.

Which case does Papparti fall into? Read MODEL-DESIGN.md:

- The **reasoning LLM runs 24/7** — that is a genuine **~100%-of-the-card baseline** for that one model,

and it is the strongest single fact for owning. (MODEL-DESIGN §0, §3.2: the LLM is "the heavy one".)

- But the **STT and TTS models are bursty** (§1.2–1.3: speech traffic peaks at booking time, idles

between), and MODEL-DESIGN.md §4.4 gates the whole cost model on **sessions/node ≥ 10** — a number that is **explicitly ESTIMATED and unmeasured** ("~14 sessions/node (ESTIMATED)"; "a 3× error here triples the bill"). The concurrency load test (§4.5) exists precisely because the real number is unknown.

- Therefore the client's **actual average utilisation of the card is UNKNOWN until Phase 0B measures it**,

and the estimate's own §10 says the concurrency capacity of a single 48 GB card under load is **UNVERIFIED. Planning a 20-month capex commitment on an unmeasured utilisation number is exactly the mistake this document exists to prevent.**

Plain statement: on the *known* fact (reasoning model 24/7) the card is heavily used and buying *would* pay off over 20+ months. On the *unknown* facts (real concurrency, whether the stack even passes Phase 0B) it is unsafe to commit capex now. **The uncertainty itself is the reason to rent until Phase 0B removes it.**

6. WHAT THIS MEANS FOR THE PROPOSAL

The recommendation equals the current assumption. Nothing in the money changes.

- **CLIENT-ESTIMATE.md** line that stays as-is: §1b / §5, the monthly row

| GPU serving – 1× 48 GB dedicated | 41,665 | RunPod RTX A6000 @ \$0.59/hr × 730 hr [1] |. **It does not change. ₹41,665/month is confirmed correct.**

- **The one-time total is NOT affected** — this is a monthly item; the ₹78,72,318 / ₹92,89,335 totals are

untouched. (The Phase 0 line's GPU sub-charge, ₹41,665, is likewise unchanged.)

- **Suggested addition (wording only, no arithmetic):** add one sentence under §5 or §9 noting that

renting is the deliberate choice, with a documented flip condition — if Phase 0B measures sustained utilisation above ~60% of a dedicated card for a planned life above ~20 months, a bought used 48 GB card in Indian colocation becomes cheaper and should be revisited at that point. This converts an unexamined assumption into a **stated, testable decision rule** — which is worth more to the client than the rupees.

If the client nonetheless elects to buy, the line changes as follows (monthly): GPU serving ₹41,665 → **~₹7,716** (colo ₹5,000 + power ₹2,716, Option C), i.e. a monthly saving of **~₹33,949/month**; the monthly total ₹1,26,108 → **~₹92,161**; but the one-time total rises by the **₹7,25,441 capex** (→ ₹85.98 lakh excl. GST) — which is the trade the client is being asked to accept. **This is the flip the recommendation says NOT to take yet.**

7. VERIFICATION AND WHAT COULD NOT BE VERIFIED

Recomputed: every break-even quotient in §1.3, the month-by-month table, the 36-month totals, the per-used-hour table in §5, the 219 kWh and 438 kWh electricity figures, and the FX conversion. FX used: **96.7384**, matching the estimate exactly; live cross-check 96.7372 (exa.ai, 2026-10-11). All pass.

Could NOT verify / UNVERIFIED:

1. **No live price for a used/refurb 2U GPU server** — Option C's \$3,500 host is an **ASSUMPTION**; no

listing with a quotable price was found. Option B uses the bottom of the India *indicative band*, and RDP (the India OEM that sells exactly this) lists "**Pricing on request**" only.

1. **No quoted monthly colo price from a GPU-certified Indian facility.**

VyomCloud (₹5,000/₹10,000) and

Xenax (₹3,999) are Tier III **general** colo, not GPU-certified; Sify's DGX-ready GPU colo is **pay-per-use, price on request** (sifytechnologies.com, 2025-05-20). The GPU-colo real price is **UNVERIFIED** — and could be *higher* than the ₹5,000–10,000 assumed, which would *raise* the break-even and strengthen the rent recommendation.

1. **Node power draw (600 W)** is an **ASSUMPTION** (card TDP + host); a Phase 0B power meter would confirm it.
2. **The Odisha figure is the LT-Industrial energy charge (₹6.20/kWh).** A colo facility bills the client

inside its plan (power included), so the ₹6.20 applies only to the on-premises variant; treating it as the colo power price is an **ASSUMPTION**. Confirmed source: OERC RST FY 2026-27 + two Mercom reports.

1. **The client's actual future utilisation is unknowable now** — this is the honest core of the document,

not a gap to be papered over.

1. **Import duty / customs on a card bought from a US/China marketplace** was **not** modelled; landed cost

could be ~18–30% above the \$3,999–8,400 list once duty + freight + IGST are added. **Not priced — flag before any buy.**

8. ASSUMPTIONS REGISTER

#	ASSUMPTION	BASIS
A1	FX 1 USD = ₹96.7384	estimate's pinned rate [FX]; re-verified 96.7372 on 2026-10-11
A2	Rent ₹41,665/mo = \$0.59/hr × 730 × 96.7384	CLIENT-ESTIMATE.md §2 [1]
A3	Node draw 600 W → 438 kWh/mo	card 300 W TDP + ~300 W host — ASSUMPTION
A4	Odisha LT-industrial ₹6.20/kWh, duty 0%	OERC RST FY 2026-27; Mercom 2026
A5	Amortisation life 36 months	ASSUMPTION for per-used-hour illustration only; break-even uses no life assumption
A6	Used-card host \$3,500	ASSUMPTION — no live listing found
A7	India entry server ₹6,00,000	bottom of Proactive's indicative ₹5–25 lakh band
A8	Landed cost = list price (no duty modelled)	ASSUMPTION — understates buy cost; flagged in §7.6

[FX] [www.xe.com/currencyconverter/convert/?](https://www.xe.com/currencyconverter/convert/?Amount=1&From=USD&To=INR)

Amount=1&From=USD&To=INR — XE USD/INR mid-market, 11 Oct 2026 (estimate's source); re-verified at exa.ai/library/markets/forex/USDINR, 2026-10-11T08:16Z (96.737212).